

Ghost Signals

Verifying Termination of Busy-Waiting

Tobias Reinhard¹[0000-0003-1048-8735] and Bart Jacobs¹[0000-0002-3605-249X]

imec-DistriNet Research Group, KU Leuven, Belgium
 {tobias.reinhard,bart.jacobs}@kuleuven.be

Abstract. Programs for multiprocessor machines commonly perform busy-waiting for synchronization. We propose a separation logic using so-called *ghost signals* to modularly verify termination of such programs under fair scheduling. Intuitively spoken, ghost signals lift the runtime concept of *wait-notify* synchronization to the verification level and allow a thread to busy-wait for an event X while another thread promises to trigger X .

1 Introduction

Programs for multiprocessor machines commonly perform busy-waiting for synchronization [23,25]. In this paper, we propose a separation logic [26,33] to modularly verify termination of such programs under fair scheduling where some threads busy-wait for other threads to trigger a certain event, e.g., setting a memory flag. By modularly, we mean that we reason about each thread and each function in isolation. That is, we do not reason about thread scheduling or interleavings. We only consider these issues when proving the soundness of our logic. Assuming fair scheduling is necessary since busy-waiting for an event X only terminates if the thread responsible for the event is sufficiently often scheduled to trigger X .

The approach we present in this paper to verify termination of busy-waiting programs builds on two of our earlier works: (i) In [14] we proposed so-called *call permissions* to verify termination of programs involving recursion and loops. (ii) In [31,32] we proposed a separation logic to verify termination of busy-waiting for abrupt termination. We call said permissions *iteration permissions* as, for simplicity, in this paper we focus on loops rather than on recursion. An extension would, however, be straight-forward.

In order to prove that a busy-waiting loop terminates, we have to prove that it performs only finitely many iterations. To do this, we choose a well-founded set of degrees Δ , associate each iteration permission with a degree and start our verification with a finite number of these permissions. We let each loop iteration consume an iteration permission and allow to create new ones via weakening, i.e., replacing an iteration permission by finitely many of a lower degree. Then, we show that the finite stock of iteration permissions in the system is monotonically decreasing and strictly decreases with each loop iteration. The idea of using well-founded orders to verify termination is well-known in the literature [6,18,24].

However, the number of iterations that a busy-waiting loop performs typically depends on the runtime scheduling. In such a case, starting with a fixed number of iteration permissions and allowing weakening is not sufficient to prove termination. To overcome this limitation, we introduce *ghost signals* and *wait permissions*. This verification concept roughly corresponds to runtime *wait-notify* synchronization where one party waits for an event to happen and the party responsible for triggering the event notifies the waiting one when the event occurred. Ghost signals and wait permissions allow a thread to wait for an event.

We choose a well-founded set of levels \mathcal{Levs} and associate each signal with one as well as with a boolean to indicate whether it is *set* or *unset*. When a thread creates a signal, it is initially *unset* and the thread obtains a so-called *obligation* [9, 10, 19, 20] to set the signal. Generally, an obligation requires the thread holding the obligation to discharge it by performing a certain action. This way, when a thread *knows* that some signal s has not been set, yet, while it does not hold an obligation for s , it knows that another thread holds the obligation and will eventually set s .

Just as with iteration permissions, we associate each wait permission with a degree. Before a thread can wait for a signal s , it has to obtain a wait permission for s , which consumes an iteration permission of a higher degree. A thread holding a wait permission of degree δ for an *unset* signal s can generate an iteration permission of degree δ by waiting for s . We ensure that no thread (directly or indirectly) waits for itself by requiring the level of s to be lower than the level of each held obligation.

Ghost signals and permissions are *ghost resources* [16], i.e., resources that only exist on the verification level and hence do not affect the program’s runtime behaviour. In particular, setting a signal does not by definition correspond to any runtime event. So, in order to use a signal s effectively, we have to prove an invariant stating that s is set if and only if the event of interest has occurred.

We can prove that, assuming fair scheduling, in a verified program each signal is waited-for only finitely often. This allows us to view a wait permission as the finite collection of iteration permissions that will be generated by threads waiting for it. We extend the above soundness argument by (i) not just considering the stock of iteration permissions available at a certain point in time but also (ii) the stock of future iteration permissions that will be created by waiting for signals (via already available wait permissions). As above, this stock of iteration permissions decreases with each loop iteration. Hence, we can prove that verified programs terminate.

We start by gradually introducing the intuition behind our verification approach and the concepts we use. In § 2.1 and § 2.2, we introduce the aspects related to verifying data-race-freedom and deadlock-freedom, respectively, which are standard [5, 8, 14, 20]. In § 2.3, we introduce our approach for verifying termination of busy waiting. We continue in § 3 and § 4 by formally presenting the language and the logic, respectively. In the latter, we define the proof system and state the soundness theorem whose proof we outline in § 5. In § 6 we sketch the verification of a realistic example program to demonstrate our approach’s

usability and address fine grained concurrency in § 7. Further, we sketch an extension of our logic that allows permission transfer between threads in § 8, describe the available tool support in § 9, and discuss integrating higher-order features in § 10. We conclude by comparing our approach to related work and reflecting on it in § 11 and § 12.

2 Verifying Termination of Busy-Waiting in a Nutshell

When we try to verify termination of busy-waiting programs, multiple challenges arise. Throughout this section, we describe these challenges and our approach to overcome them. We will use the program presented in Fig. 1 as running example. It shows two threads communicating via a shared memory cell x that is initialized with 0 and protected by a mutex m . One thread sets x to 1 and the other thread busy-waits for x to change its value.

We introduce our program logic gradually. In § 2.1 and § 2.2, we introduce the aspects related to verifying data-race-freedom and deadlock-freedom, respectively, which are standard [5, 8, 14, 20]. In § 2.3, we introduce our approach for verifying termination of busy waiting.

```

let  $x := \text{cons}(0)$  in
let  $m := \text{new\_mutex}$  in
fork ( while(acquire  $m$ ;
             let  $y := [x]$  in
             release  $m$ ;
              $y = 0$ )
      do skip);
acquire  $m$ ;
 $[x] := 1$ ;
release  $m$ 

```

Fig. 1: Minimal Example Program.

2.1 Verifying Data Race Freedom

The two threads both access the shared heap cell x . Due to non-deterministic runtime scheduling, data races can occur if both try to access it simultaneously. To prevent this, we protect x by a mutex m and synchronize accesses. In this section, we introduce the aspects of our logic for verifying data-race-freedom; they are standard [5, 8].

Specifications We use Hoare triples $\{A\} c \{\lambda r. B(r)\}$ [12] to specify the behaviour of a program c . Such a triple expresses that if we execute c in a state that fulfils precondition A , then two properties hold: (i) c does not get stuck, i.e., it either terminates or diverges. (ii) If c terminates and r denotes the value returned by c , then postcondition $B(r)$ holds in its final state.

Proof System We define a proof relation \vdash which ensures that whenever we can prove $\vdash \{A\} c \{\lambda r. B(r)\}$, then no race condition occurs during the execution of c . Fig. 2b presents some of the proof rules we use to define \vdash . The full set of rules is presented in the appendix (cf. Fig. 9 and 10). Fig. 2a illustrates how we can apply them to verify data race freedom of our minimal example. For now, ignore the obligations chunks $\text{obs}(O)$ and levels marked in grey in the verification sketch. We are going to explain their use in the next subsection.

Accessing Heap Cells Our logic extends plain separation logic [26, 33]. When we allocate x by applying PR-CONS, we get a points-to chunk $\ell_x \mapsto 0$. It expresses that heap location ℓ_x , which we store in x , points to the value 0. We can split the chunk and give each half, i.e., $[\frac{1}{2}]\ell_x \mapsto 0$, to one thread [3, 4]. This allows both to read x , cf. proof rule PR-READHEAPLOC. However, in order to prevent data races, we require threads to own the full chunk $\ell_x \mapsto 0$ to write x , cf. PR-ASSIGNTOHEAP.

Lock Invariants Splitting x 's points-to chunk does not work for the example program since one thread requires write access. Hence, they have to share the chunk and we must synchronize their accesses of x . We use the mutex m for this. In order to specify which resources m protects, we associate it with a lock invariant P . The mutex protects x whose value varies over time. Hence, we choose the invariant $P := \exists v_x. \ell_x \mapsto v_x$ which abstracts over the concrete value v_x of x .

Ghost State & Ghost Proof Steps Directly after the creation of a mutex by proof rule PR-NEWMUTEX, it is uninitialized. That is, it has not been associated with any lock invariant, yet. We must choose this invariant explicitly by performing a proof step that consumes invariant P and binds it to m . This step does not affect the resources in any way that would be relevant to the program's runtime behaviour. The resources that P describes still exist but their ownership has been transferred to the mutex. Hence, we say that it changes the *ghost state* and call it a *ghost proof step*. We express such steps in the form of a *view shift* relation \Rightarrow [17]. A view shift $A \Rightarrow B$ expresses that we can reach postcondition B from precondition A by (i) drawing semantic conclusions or by (ii) manipulating the ghost state. View shift rule VS-MUTINIT' presented in Fig. 2b allows us to initialize m with invariant P . The full set of rules we use to define \Rightarrow are presented in the appendix (cf. Fig. 8).

Synchronization Since initializing m consumes lock invariant P , the shared heap cell x becomes unavailable to both threads. In order to read or write x , a thread must obtain the *exclusive* right to access x by acquiring m . The corresponding proof rule PR-ACQUIRE changes the mutex's state from *unlocked* to *locked* and produces the lock invariant P . When we release m , we must apply PR-RELEASE', which changes its state to *unlocked* and consumes P . That is, x becomes unavailable to the thread again, when it releases m . Mutex acquisition will not succeed during runtime if some thread already holds the mutex. Hence, our proof rules statically ensure that only one thread can access x at any time.

```

{obs(∅)}
let x := cons(0) in let m := new_mutex in
∀ℓx, ℓm.
{obs(∅) * ℓx ↦ 0 * uninit(ℓm)}
init_mutex; P := ∃vx. ℓx ↦ vx
{obs(∅) * mutex((ℓm, 0), P)}
fork ({obs(∅) * [ $\frac{1}{2}$ ]mutex((ℓm, 0), P)}
  while (acquire m
    {obs(⟦(m, 0)⟧) * locked((m, 0), P,  $\frac{1}{2}$ ) * ∃vx. ℓx ↦ vx}
    let y := [x] in release m; y = 0
    {obs(∅) * [ $\frac{1}{2}$ ]mutex((m, 0), P)}
  do skip);
{obs(∅) * [ $\frac{1}{2}$ ]mutex((ℓm, 0), P)}
acquire m; [x] := 1;
{obs(⟦(m, 0)⟧) * locked((ℓm, 0), P,  $\frac{1}{2}$ ) * ℓx ↦ 1}
release m
{obs(∅) * [ $\frac{1}{2}$ ]mutex((ℓm, 0), P)}

```

(a) Proof outline verifying data race and deadlock freedom. Hints on ghost proof steps highlighted in red. Auxiliary abbreviations marked in brown.

<p>PR-CONS $\vdash \{\mathbf{True}\} \mathbf{cons}(v) \{\lambda\ell. \ell \mapsto v\}$</p> <p>PR-ASSIGNTOHEAP $\vdash \{\ell \mapsto _ \} [\ell] := v \{\ell \mapsto v\}$</p>	<p>PR-READHEAPLOC $\vdash \{\llbracket f \rrbracket \ell \mapsto v\} [\ell] \{\lambda r. r = v * \llbracket f \rrbracket \ell \mapsto v\}$</p> <p>PR-NEWMUTEX $\vdash \{\mathbf{True}\} \mathbf{new_mutex} \{\lambda\ell. \mathbf{uninit}(\ell)\}$</p>
<p>PR-ACQUIRE</p> $\frac{\forall o \in O. \text{lev}(m) <_{\mathbb{L}} \text{lev}(o)}{\vdash \{\mathbf{obs}(O) * \llbracket f \rrbracket \mathbf{mutex}(m, P)\} \mathbf{acquire} m.\text{loc} \{\mathbf{obs}(O \uplus \{\llbracket m \rrbracket\}) * \mathbf{locked}(m, P, f) * P\}}$	
<p>PR-RELEASE'</p> $\vdash \{\mathbf{obs}(O \uplus \{\llbracket m \rrbracket\}) * \mathbf{locked}(m, P, f) * A\} \mathbf{release} m.\text{loc} \{\mathbf{obs}(O) * \llbracket f \rrbracket \mathbf{mutex}(m, P) * B\}$	<p>PR-VIEWSHIFT</p> $\frac{A \Rightarrow A' \quad B' \Rightarrow B \quad \vdash \{A'\} c \{B'\}}{\vdash \{A\} c \{B\}}$
<p>PR-FORK</p> $\frac{\vdash \{\mathbf{obs}(O_f) * A\} c \{\mathbf{obs}(\emptyset)\}}{\vdash \{\mathbf{obs}(O_m \uplus O_f) * A\} \mathbf{fork} c \{\mathbf{obs}(O_m)\}}$	
<p>VS-MUTINIT'</p> $\mathbf{uninit}(\ell) * P \Rightarrow \mathbf{mutex}((\ell, L), P)$	<p>VS-SEMIMP</p> $\frac{\forall H. \mathbf{consistent}_{\text{th}}(H) \wedge H \vDash_A A \Rightarrow H \vDash_A B}{A \Rightarrow B}$

(b) Proof rules and view shift rules for verification of data race and deadlock freedom.

Fig. 2: Verifying data race and deadlock freedom. Obligations chunks and levels marked in grey. Ignore them while reading § 2.1. Obligations and levels will be introduced in § 2.2.

2.2 Verifying Deadlock Freedom

Obligations During runtime, any attempt to acquire m will block if another thread already holds the mutex until this thread releases m again. If it does not release m , the thread trying to acquire m can never proceed with its execution. We use so-called *obligations* [9, 10, 19, 20] to prevent this. These are *ghost resources* [16], i.e., resources that do not exist during runtime and can hence not influence a program’s runtime behaviour. They carry, however, information relevant to the program’s verification. Generally, holding an obligation requires a thread to discharge it by performing a certain action. Whenever a thread acquires m we generate an obligation to release it and releasing m discharges this obligation, cf. proof rules PR-ACQUIRE and PR-RELEASE’ presented in Fig. 2b. We describe which obligations a thread holds by bundling them into obligations chunks $\text{obs}(O)$, cf. Fig. 2a (obs chunks marked in grey). Consequently, $\text{obs}(\emptyset)$ asserts that a thread does not hold any obligations. We can prove deadlock-freedom of a program c by proving that it discharges all its obligations, i.e., by proving $\vdash \{\text{obs}(\emptyset)\} c \{\text{obs}(\emptyset)\}$.

Whenever a thread forks a new thread, it can transfer ownership of resources to the newly forked thread, cf. PR-FORK. In particular, the forking thread can pass some of its obligations to the new thread. However, in order to prevent threads from dropping obligations via dummy forks, we only allow forking if the forked thread provably discharges all its obligations.

Levels It is not always straight-forward to see that a mutex will be released eventually. Consider a program with two mutexes m_1, m_2 and two threads. Let one thread execute **acquire** m_1 ; **acquire** m_2 ; **release** m_2 ; **release** m_1 ; and the other **acquire** m_2 ; **acquire** m_1 ; **release** m_1 ; **release** m_2 . Both threads’ execution can get stuck as both mutexes’ release depends on successful acquisition. To prevent such *wait cycles*, we choose a partially ordered set of levels $\mathcal{L}vs$ and associate every mutex with a level. In the example program, we choose 0 for m , cf. Fig. 2a (levels marked in grey). Additionally, we only allow a thread to acquire a mutex if its level is smaller than the level of each held obligation, cf. PR-ACQUIRE. This also prevents any thread from attempting to acquire mutexes twice, e.g., **acquire** m ; **acquire** m .

2.3 Verifying Termination

The approach we present in this paper to verify termination of busy-waiting programs builds on two of our earlier works: (i) In [14] we proposed so-called *call permissions* to verify termination of programs involving recursion and loops. (ii) In [31, 32] we proposed a separation logic to verify termination of busy-waiting for abrupt termination. We call said permissions *iteration permissions* as, for simplicity, in this paper we focus on loops rather than on recursion. An extension would, however, be straight-forward.

```

{obs(∅) * itperm(1)}
let x := cons(0) in let m := new_mutex in new_signal;
∀ℓx, ℓm, ids.
{obs({(ids, 1)}) * itperm(1) * ℓx ↦ 0 * uninit(ℓm) * signal((ids, 1), False)}
s := (ids, 1), mut := (ℓm, 0), P := ∃vx. ℓx ↦ vx * signal(s, vx ≠ 0)
{obs({s}) * itperm(1) * mutex(mut, P)}
new_waitPerm;
{obs({s}) * wperm(ids, 0) * mutex(mut, P)}
fork ({obs(∅) * wperm(ids, 0) * [ $\frac{1}{2}$ ]mutex(mut, P)}
  while (acquire m
    ∀vx.
    {obs({mut}) * wperm(ids, 0) * ℓx ↦ vx * signal(s, vx ≠ 0) * locked(mut, P,  $\frac{1}{2}$ )}
    let y := [x] in
    release m + release_view_shift
      {obs(∅) * wperm(ids, 0) * ℓx ↦ vx * signal(s, vx ≠ 0)}
      wait
      {vx = 0 ↔ itperm(0) * obs(∅) * wperm(ids, 0) * P};
    y = 0
    {obs(∅) * (vx = 0 → itperm(0) * wperm(ids, 0) * [ $\frac{1}{2}$ ]mutex(mut, P))}
    do skip);
    {obs(∅)}
    {obs({s}) * [ $\frac{1}{2}$ ]mutex(mut, P)}    lev(mut) < lev(s)
    acquire m; [x] := 1; set_signal;
    {obs({mut}) * locked(mut, P,  $\frac{1}{2}$ ) * ℓx ↦ 1 * signal(s, True)}
    release m
    {obs(∅) * [ $\frac{1}{2}$ ]mutex(mut, P)}
    {obs(∅)}

```

(a) Proof outline verifying termination. Hints on ghost proof steps highlighted in **red**. Auxiliary abbreviations marked in **brown**.

$$\begin{array}{c}
 \text{PR-WHILE} \\
 \frac{\vdash \{I\} c_b \left\{ \begin{array}{l} \lambda b. (b = \text{True} \vee b = \text{False}) \\ * (b \rightarrow \text{itperm}(\delta) * I) \\ * (\neg b \rightarrow B) \end{array} \right.}{\vdash \{I\} \text{ while } c_b \text{ do skip } \{B\}}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VS-WAITPERM} \\
 \frac{\delta' <_{\Delta} \delta}{\text{itperm}(\delta) \Rightarrow \text{wperm}(id, \delta')}
 \end{array}$$

$$\begin{array}{c}
 \text{PR-RELEASE} \\
 \frac{\text{obs}(O) * A \Rightarrow \text{obs}(O) * P * B}{\text{release } m.\text{loc} \quad \{ \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * A \} \Rightarrow \text{obs}(O) * [f]\text{mutex}(m, P) * B \}}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VS-WAIT} \\
 \frac{\forall o \in O. \text{lev}(s) <_{\mathbb{L}} \text{lev}(o)}{\text{obs}(O) * \text{wperm}(s.\text{id}, \delta) * \text{signal}(s, b) \Rightarrow \text{obs}(O) * \text{wperm}(s.\text{id}, \delta) * \text{signal}(s, b) * (\neg b \leftrightarrow \text{itperm}(\tau, \delta))}
 \end{array}$$

$$\begin{array}{c}
 \text{VS-WEAKPERM} \\
 \frac{\delta' <_{\Delta} \delta \quad N \in \mathbb{N}}{\text{itperm}(\delta) \Rightarrow \bigstar_{1, \dots, N} \text{itperm}(\delta')}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VS-NEWSIGNAL} \\
 \text{obs}(O) \Rightarrow \exists id. \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})
 \end{array}$$

$$\begin{array}{c}
 \text{VS-SET SIGNAL} \\
 \text{obs}(O \uplus \{s\}) * \text{signal}(s, _) \Rightarrow \text{obs}(O) * \text{signal}(s, \text{True})
 \end{array}
 \qquad
 \begin{array}{c}
 \text{VS-MUTINIT} \\
 \frac{\neg \text{mentionsPerms}(P)}{\text{uninit}(\ell) * P \Rightarrow \text{mutex}((\ell, L), P)}
 \end{array}$$

(b) Proof rules and view shift rules for termination verification.

Fig. 3: Verifying termination.

Iteration Permissions In order to prove termination we need to prove that loops perform only finitely many iterations. Fig. 3b shows the proof and view shift rules we use and Fig. 3a presents our termination verification of the minimal example. We will now explain their use. Consider the rule PR-WHILE. We only consider busy-waiting loops of the form **while** c_b **do skip**, where c_b is a command whose result controls the loop and **skip** is a no-op. We start our verification with a finite number of iteration permissions. Each loop iteration consumes one of these permissions. We have to prove that the condition c_b returns a boolean value to prevent the execution from getting stuck. As the loop busy-waits while c_b yields **True**, we must prove that, in this case, it preserves some invariant I and that we have an additional iteration permission for the iteration to consume.

To avoid the need of fixing the number of needed permissions at the beginning of a verification, we allow to generate finitely many permissions by *weakening*, cf. view shift rule VS-WEAKPERM. We choose a well-founded set of degrees Δ and associate each permission with one. Weakening allows us to replace one iteration permission of degree δ by finitely many of a lower degree δ' . The idea of using well-founded orders to verify termination is well-known in the literature [6,18,24].

Ghost Signals & Wait Permissions How many iterations a busy-waiting loop performs often depends on the concrete runtime scheduling as in the minimal example. In such a case, starting with a fixed number of iteration permissions and allowing weakening is not sufficient to prove termination. To overcome this limitation, we introduce *ghost signals* and *wait permissions*. This verification concept roughly corresponds to runtime *wait-notify* synchronization where one party waits for an event to happen and the party responsible for triggering the event notifies the waiting one when the event occurred. Ghost signals and wait permissions allow a thread to wait for an event, e.g., x to be set, in a similar manner.

Just like permissions, ghost signals are ghost resources. We associate each signal with a unique ID and a boolean $b \in \mathbb{B} = \{\mathbf{True}, \mathbf{False}\}$. We can create a signal via a view shift with VS-NEWSIGNAL. Initially, the signal is unset, i.e., associated with **False**. Creating a signal automatically creates an obligation to set the signal. In our termination verification proof, the main thread holds this obligation and will eventually discharge it by setting the signal when it sets x .

We associate each wait permission with a signal ID and a degree. A wait permission for some signal s enables a thread to wait until s is set. By applying VS-WAITPERM, we can replace an iteration permission by a wait permission of a lower degree. Rule VS-WAIT allows a thread holding a signal s and a wait permission for s of degree δ to generate an iteration permission of degree δ . The thread can use this iteration permission to justify another loop iteration.

In our termination verification, both threads share the created signal via the lock invariant P . The main thread eventually sets the signal and the busy-waiting thread uses it and the wait permission it holds to justify loop iterations until s is set. However, the latter thread actually loops until x is set. So, in order to use the ghost signal effectively, we have to prove an invariant stating that the signal is consistent with the event we wait for, i.e., we need to prove that s is

associated with `True` if and only if `x` has been set. We fix this as part of the lock invariant P .

Furthermore, we need to ensure that threads do not wait for themselves. Hence, we associate every signal with a level and only allow a thread to wait for s if its level is lower than the level of each held obligation, cf. `VS-WAIT`. In our example, the main thread acquires the mutex m while it holds the obligation for s and sets s while still holding m . The lock invariant stating consistency between s and x makes this necessary. Hence, we choose the level of m to be lower than the level of s as required by `PR-ACQUIRE`. This, however, means that the forked thread can only wait for s (i.e. apply `VS-WAIT`) after it discharged the obligation for m . We solve this by allowing a view shift when releasing m that takes place after the obligation for m has been discharged but before the mutex's lock invariant P is consumed. In the verification outline, this view shift is marked as *release_view_shift*.

So far we allowed threads to share arbitrary resources via lock invariants. However, this becomes problematic when permissions are involved. Consider a non-terminating program with two infinitely looping threads t_1 and t_2 sharing a signal s . t_1 holds the obligation for s and hence cannot wait for it. In each iteration t_2 can wait for s and thereby create two iteration permissions of which it transfers one through the lock invariant to t_1 . This way, both threads can justify their infinite loop and t_1 essentially busy-waits for itself to set s . We prevent such inconsistencies by prohibiting lock invariants from mentioning permissions and restrict `VS-MUTINIT'` accordingly. The new view shift rule `VS-MUTINIT` is presented in Fig. 3b.

Note that if we could prophesy ahead of time how many times we would need to busy-wait before the signal s was set, we could just use the weakening rule `VS-WEAKPERM`. Consider any execution of any program verified using our proof rules and the graph depicting its control flow at runtime, i.e., the graph where each path represents the execution of a particular thread including ghost steps corresponding to view shifts, e.g., creating a signal and waiting. We call this a *program order graph*. From the perspective of a path in this graph that only waits for signals that are waited for only finitely often, a wait permission is effectively an encapsulated prophecy: $\text{wperm}(id_s, \delta) = \bigstar_{1, \dots, N} \text{itperm}(\delta)$ where N is the number of times s will be waited for (or 0 if s is waited for infinitely often). By this interpretation of wait permissions, the stock of iteration permissions on such a path decreases at each loop iteration. It is not possible that a signal is waited for infinitely often. Indeed, suppose some signals S^∞ are. Take $s_{\min} \in S^\infty$ with minimal level. Then the path that carries the obligation for s_{\min} waits only for signals that are waited for finitely often. By the above argument, it is a finite path. This contradicts the fact that it carries an obligation that is never discharged. Notice that the above argument relies on the property that every nonempty set of levels has a minimal element. For this reason, for termination verification we require that $\mathcal{L}evs$ is not just partially ordered, but also well-founded.

3 Language

In this section and the next, we present our approach formally. In this section, we define the programming language; in § 4 we define the proof system. We state the soundness theorem and outline our soundness proof in § 5.

We consider a simple imperative programming language with support for multi-threading, shared memory and synchronization via mutexes. Fig. 4 presents its syntax and semantics. For its definition we assume (i) an infinite set of program variables $x \in \mathcal{X}$, (ii) an infinite set of heap locations $\ell \in \mathcal{Locs}$, (iii) a set of values $v \in \mathit{Values}$ which includes heap locations, booleans $\mathbb{B} = \{\mathit{True}, \mathit{False}\}$ and the unit value tt , (iv) a set of operations $op \in \mathit{Ops}$ and (v) an infinite, totally ordered and well-founded set of thread IDs $\theta \in \Theta$.

The language contains standard sets of pure expressions Exps and (potentially) side-effectful commands Cmds , cf. Fig. 4b. The latter includes commands for heap allocation and manipulation, forking and loops. We define *physical heaps* [14] (as opposed to *logical heaps* [14] presented in the next section) as a finite set of *physical resource chunks*, cf. Fig. 4c. A points-to chunk $\ell \mapsto v$ expresses that heap location ℓ points to value v [14, 33]. Moreover, we have chunks to represent unlocked and locked mutexes.

We represent a program state by a physical heap and a *thread pool*, which we define as a partial function mapping a finite number of thread IDs to threads, cf. Fig. 4d. Thread IDs are unique and never reused. Hence, we represent running threads by commands and terminated ones by **term** instead of removing threads from the pool. We define the operational semantics of our language in terms of two small-step reduction relations: $\rightsquigarrow_{\text{st}}$ for single threads and $\rightsquigarrow_{\text{tp}}$ for thread pools, cf. Fig. 4e and 4f. Since expressions are pure and their evaluation is deterministic (cf. Def. 13 in appendix) we identify closed expressions with their ascribed value. (i) $h, c \rightsquigarrow_{\text{st}} h', c', T$ expresses that heap h and command c are reduced in a single step to h' and c' and that this thread forks a set of threads T . This set is either empty or a singleton as no step forks more than one thread. (ii) $h, P \overset{\theta}{\rightsquigarrow}_{\text{tp}} h', P'$ expresses that heap h and thread pool P are reduced in a single step to h' and P' . ID θ identifies the thread reduced in this step.

As thread scheduling is non-deterministic, so is our thread pool reduction relation $\rightsquigarrow_{\text{tp}}$. Consider the minimal example presented in the previous section in Fig. 1. It does not terminate if the main thread is never scheduled after the new thread was forked. Hence, our verification approach relies on the assumption of fair scheduling. That is, we assume that every thread is always eventually scheduled while it remains running. Further, we represent program executions by sequences of reduction steps. As we primarily consider infinite sequences in this paper, we define reduction sequences to be infinite to simplify our terminology.

Definition 1 (Reduction Sequence). *Let $(h_i)_{i \in \mathbb{N}}$ and $(P_i)_{i \in \mathbb{N}}$ be infinite sequences of physical heaps and thread pools, respectively. We call $(h_i, P_i)_{i \in \mathbb{N}}$ a reduction sequence if there exists a sequence of thread IDs $(\theta_i)_{i \in \mathbb{N}}$ such that $h_i, P_i \overset{\theta_i}{\rightsquigarrow}_{\text{tp}} h_{i+1}, P_{i+1}$ holds for every $i \in \mathbb{N}$.*

Definition 2 (Fairness). We call a reduction sequence $(h_i, P_i)_{i \in \mathbb{N}}$ fair iff for all $i \in \mathbb{N}$ and $\theta \in \text{dom}(P_i)$ with $P_i(\theta) \neq \text{term}$ there exists some $k \geq i$ with

$$h_k, P_k \xrightarrow{\theta}_{\text{tp}} h_{k+1}, P_{k+1}.$$

4 Logic

In this section we formalize the logic we sketched in § 2. Fig. 5 defines the assertions we use to express which resources a thread holds. For the definition we assume (i) an infinite set of ghost signal IDs $id \in \mathcal{ID}$ and (ii) infinite, partially ordered and well-founded sets of levels $L \in \mathcal{Levs}$ and degrees $\delta \in \Delta$. We denote the latter sets' order relations by $<_{\mathcal{L}}$ and $<_{\Delta}$, respectively. We define ghost resources as described in § 2, cf. Fig. 5b, and the syntax of our assertion language to match the presented verification sketches, cf. Fig. 5c¹.

We define *logical heaps* and *logical resources* [14] in Fig. 5d. These concepts correspond to physical heaps and physical resources but additionally encompass ghost resources and ownership. For instance, logical resources include signal chunks and initialized mutex chunks are associated with a lock invariant. Rather than being a set of resources, logical heaps map logical resources to fractions. This allows us to express which portion of a resource a thread owns. We interpret assertions in terms of a model relation, cf. Fig. 5e. $H \models_{\mathcal{A}} a$ expresses that assertion a holds with respect to logical heap H . Further, we define various predicates to characterize a heap's contents.

Definition 3 (Logical Heap Predicates). We call a logical heap H complete and write $\text{complete}_{\text{h}}(H)$ if it contains exactly one obligations chunk, i.e., if there exist a bag of obligations O with $H(\text{obs}_{\text{Res}}(O)) = 1$ and if there does not exist any bag of obligations O' with $O \neq O'$ and $H(\text{obs}_{\text{Res}}(O')) > 0$.

We call a logical heap H finite and write $\text{finite}_{\text{h}}(H)$ if it contains only finitely many resources, i.e., if the set $\{r^l \in \mathcal{R}^{\text{log}} \mid H(r^l) > 0\}$ is finite.

We already presented the most important rules that we use to define our proof relation and view shift relation in Fig. 2b and 3b in § 2 (excluding auxiliary rules VS-MUTINIT' and PR-RELEASE'). We present the full collection of proof rules and view shift rules in the appendix.

Definition 4 (View Shift & Proof Relations). We define a view shift relation $\Rightarrow \subset \mathcal{A} \times \mathcal{A}$ and a proof relation $\vdash \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$ according to the rules presented in Fig. 8, 9 and 10 presented in the appendix. We state the provability of a Hoare triple in the form of $\vdash \{A\} c \{\lambda r. B(r)\}$ where r captures the value returned by c . To simplify the notation, we omit the result value if it is clear from the context or irrelevant.

¹ That is, we define the set of assertions \mathcal{A} as the least fixpoint of F where $F(A) = \{\text{True}, \text{False}\} \cup \{\neg a \mid a \in A\} \cup \{a_1 \wedge a_2 \mid a_1, a_2 \in A\} \cup \dots \cup \{\bigvee A' \mid A' \subseteq A\} \cup \dots$. Since F is a monotonic function over a complete lattice, it has a least fixpoint according to the Knaster-Tarski theorem [35].

$x \in \mathcal{X}$: Program variables $\ell \in \mathcal{Locs}$: Heap locations $op \in \mathcal{Ops}$: Operations
 $v \in \mathcal{Values} \supseteq \{\mathbf{tt}\} \cup \mathbb{B} \cup \mathcal{Locs}$: Values $\theta \in \Theta$: Thread IDs

(a) Assumed sets and variables. \mathcal{X} , \mathcal{Locs} , Θ infinite. Θ totally ordered and well-founded.

$e \in \mathcal{Exps} ::= x \mid v \mid e = e \mid \neg e \mid op(\bar{e})$
 $c \in \mathcal{Cmds} ::= e \mid \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \mid \mathbf{fork} \ c \mid \mathbf{let} \ x := c \ \mathbf{in} \ c \mid \mathbf{if} \ c \ \mathbf{then} \ c \mid$
 $\quad \mathbf{cons}(e) \mid [e] \mid [e] := e \mid \mathbf{new_mutex} \mid \mathbf{acquire} \ e \mid \mathbf{release} \ e$
 $E \in \mathcal{EvalCtxts} ::= \mathbf{if} \ \square \ \mathbf{then} \ c \mid \mathbf{let} \ x := \square \ \mathbf{in} \ c$

(b) Syntax.

$r^p \in \mathcal{R}^{\text{phys}} ::= \ell \mapsto v \mid \text{unlocked}_{\text{pRes}}(\ell) \mid \text{locked}_{\text{pRes}}(\ell)$
 $h \in \mathcal{Heaps}^{\text{phys}} := \mathcal{P}_{\text{fin}}(\mathcal{R}^{\text{phys}})$
 $\text{locs}_{\text{pRes}}(h) := \{\ell \in \mathcal{Locs} \mid \text{unlocked}_{\text{pRes}}(\ell) \in h \vee \text{locked}_{\text{pRes}}(\ell) \in h \vee \exists v. \ell \mapsto v \in h\}$

(c) Physical resources & heaps.

$P \in \mathcal{TP} := \Theta \rightarrow_{\text{fin}} (\mathcal{Cmds} \cup \{\text{term}\})$.
 $\emptyset_{\text{tp}} : \Theta \rightarrow_{\text{fin}} (\mathcal{Cmds} \cup \{\text{term}\})$, $\text{dom}(\emptyset_{\text{tp}}) = \emptyset$.
 $+_{\text{tp}} : \mathcal{TP} \times \{C \subset \mathcal{Cmds} \mid |C| \leq 1\} \rightarrow \mathcal{TP}$
 $P +_{\text{tp}} \emptyset := P$,
 $P +_{\text{tp}} \{c\} := P[\theta_{\text{new}} := c]$ for $\theta_{\text{new}} := \min(\Theta \setminus \text{dom}(P))$.

(d) Thread pools. Θ denotes a set of thread IDs and **term** a terminated thread.

$\frac{h, c \rightsquigarrow_{\text{st}} h', c', T}{h, E[c] \rightsquigarrow_{\text{st}} h', E[c'], T} \quad h, \mathbf{fork} \ c \rightsquigarrow_{\text{st}} h, \text{tt}, \{c\} \quad h, \mathbf{if} \ \text{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, c$

$h, \mathbf{if} \ \text{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{st}} h, \text{tt} \quad h, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{st}} h, c[v/x]$

$h, \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rightsquigarrow_{\text{st}} h, \mathbf{if} \ c \ \mathbf{then} \ \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip}$

$\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{cons}(v) \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \ell} \quad \frac{\ell \mapsto v \in h}{h, [\ell] \rightsquigarrow_{\text{st}} h, v}$

$h \sqcup \{\ell \mapsto v'\}, [\ell] := v \rightsquigarrow_{\text{st}} h \sqcup \{\ell \mapsto v\}, \text{tt}$

$\frac{\ell \notin \text{locs}_{\text{pRes}}(h)}{h, \mathbf{new_mutex} \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \ell}$

$h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \mathbf{acquire} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \text{tt}$

$h \sqcup \{\text{locked}_{\text{pRes}}(\ell)\}, \mathbf{release} \ \ell \rightsquigarrow_{\text{st}} h \sqcup \{\text{unlocked}_{\text{pRes}}(\ell)\}, \text{tt}$

(e) Single thread reduction rules.

$\frac{P(\theta) = c \quad h, c \rightsquigarrow_{\text{st}} h', c', T}{h, P \overset{\theta}{\rightsquigarrow}_{\text{tp}} h', P[\theta := c'] +_{\text{tp}} T} \quad \frac{P(\theta) = v}{h, P \overset{\theta}{\rightsquigarrow}_{\text{tp}} h, P[\theta := \text{term}]}$

(f) Thread pool reduction rules

Fig. 4: Language definition.

In § 2 we verified the termination of the minimal example by proving that it discharges all its obligations (while starting with no obligations and a single iteration permission). The following theorem states that this verification approach is sound.

Theorem 1 (Soundness). *Let $\vdash \{\text{obs}(\emptyset) * \bigstar_{i=1,\dots,N} \text{itperm}(\delta_i)\} c \{\text{obs}(\emptyset)\}$ hold. There exists no fair, infinite reduction sequence $(h_i, P_i)_{i \in \mathbb{N}}$ with $h_0 = \emptyset$ and $P_0 = \{(\theta_0, c)\}$ for any choice of θ_0 .*

5 Soundness

We already sketched the high-level intuition behind our soundness argument in § 2.3. In this section, we go into more detail, provide the most important lemmas we need to prove our soundness theorem 1, sketch their proofs and finally prove the soundness theorem itself. A detailed soundness proof and all definitions can be found in the appendix and in the technical report [27].

Bridging the Gap During runtime, all threads share one physical heap where every thread is free to access every resource. This does not reflect the notions of ownership and lock invariants which we maintain on the verification level. It also does not allow us to restrict actions based on levels, e.g., only allowing the acquisition of a mutex m if its level is lower than the levels of all held obligations. Hence: (i) We annotate every thread by a logical heap to express which resources it owns (including ghost resources) and thereby obtain an *annotated thread pool*. (ii) We represent the program state by an *annotated heap* that keeps track of lock invariants and levels. In particular, we associate unlocked mutexes with logical heaps to represent the resources they protect. Since annotated heaps keep track of levels, they also keep track of signals. We denote annotated thread pools and heaps by P^a and h^a , respectively.

We define annotated versions $\rightsquigarrow_{\text{atp}}$ and $\rightsquigarrow_{\text{ast}}$ of the relations $\rightsquigarrow_{\text{tp}}$ and $\rightsquigarrow_{\text{st}}$, respectively. The annotated reduction semantics we thereby obtain needs to reflect ghost proof steps implemented by view shifts. Hence, we define $\rightsquigarrow_{\text{atp}}$ in terms of two relations: (i) $\rightsquigarrow_{\text{ghost}}$ for ghost steps and (ii) $\rightsquigarrow_{\text{real}}$ for actual program execution steps. The annotated semantics ensure that a reduction gets stuck if a thread violates any of the restrictions formulated by our proof rules. Fig. 6 presents three of the rules that we use to define these relations. The upper two rules show that a loop iteration consumes an iteration permission and that the reduction gets stuck if the looping thread does not own one. The lower rule shows that a thread can generate iteration permissions by waiting for an unset signal, but only if the signal's level is lower than the level of each held obligation. We see that these rules comply with PR-WHILE and VS-WAIT.

Interpreting Hoare Triples We interpret program specifications $\{A\} c \{B\}$ in terms of a model relation $\models_{\text{H}} \{A\} c \{B\}$ and an auxiliary safety relation

$id \in \mathcal{ID} : \text{Signal IDs}$ $L \in \mathcal{Levs} : \text{Levels}$ $\delta \in \Delta : \text{Degrees}$ $b \in \mathbb{B} : \text{Booleans}$

(a) Assumed sets and variables. \mathcal{ID} , \mathcal{Levs} , Δ infinite. \mathcal{Levs} , Δ partially ordered and well-founded.

$s \in \mathcal{S} := \mathcal{ID} \times \mathcal{Levs}$ $\Omega := \mathcal{ID} \times \Delta$ $\Lambda := \Delta$ $o \in \mathcal{O} := (\mathcal{Locs} \cup \mathcal{ID}) \times \mathcal{Levs}$
 $(id, L).id := id$ $(\ell, L).loc := \ell$ $lev(_, L) := L$

(b) Ghost resources: Ghost signals \mathcal{S} , wait/ iteration permissions Ω/ Λ , obligations \mathcal{O} .

$f \in \mathcal{F} := \{f \in \mathbb{Q} \mid 0 < f \leq 1\}$ $A \subseteq \mathcal{A}$ $O \in \text{Bags}(\mathcal{O})$ $b \in \mathbb{B}$ Index set I

$a \in \mathcal{A} ::= \text{True} \mid \text{False} \mid \neg a \mid a \wedge a \mid a \vee a \mid a * a \mid [f]\ell \mapsto v \mid \bigvee A$
 $\mid \text{uninit}(\ell) \mid [f]\text{mutex}((\ell, L), a) \mid \text{locked}((\ell, L), a, f)$
 $\mid \text{signal}((id, L), b) \mid \text{obs}(O) \mid \text{wperm}(id, \delta) \mid \text{itperm}(\delta)$
 $a_1 \rightarrow a_2 := \neg a_1 \vee a_2$ $a_1 \leftrightarrow a_2 := (a_1 \rightarrow a_2) \wedge (a_2 \rightarrow a_1)$
 $\exists i \in I. a(i) := \bigvee \{a(i) \mid i \in I\}$ $\forall i \in I. a(i) := \neg \exists i \in I. \neg a(i)$

(c) Assertion syntax. We omit quantification domain I when it is clear from the context.

$r^1 \in \mathcal{R}^{\text{log}} ::= \ell \mapsto v \mid \text{signal}_{\text{Res}}((id, L), b) \mid \text{uninit}_{\text{Res}}(\ell) \mid \text{mutex}_{\text{Res}}((\ell, L), a)$
 $\mid \text{locked}_{\text{Res}}((\ell, L), a, f) \mid \text{obs}_{\text{Res}}(O) \mid \text{wperm}_{\text{Res}}(id, \delta) \mid \text{itperm}_{\text{Res}}(\delta)$
 $H \in \text{Heaps}^{\text{log}} := \mathcal{R}^{\text{log}} \rightarrow \{q \in \mathbb{Q} \mid q \geq 0\}$
 $\emptyset_{\text{log}} \in \text{Heaps}^{\text{log}} : _ \mapsto 0$ $(H_1 + H_2)(r^1) := H_1(r^1) + H_2(r^1)$
 $\{r_1^1, \dots, r_n^1\} := \begin{cases} r_i^1 \mapsto 1 \\ x \mapsto 0 \text{ if } x \notin \{r_1^1, \dots, r_n^1\} \end{cases}$ $(q \cdot H)(r^1) := q \cdot (H(r^1))$

(d) Logical resources and heaps.

$H \vDash_{\mathbb{A}} \text{True}$
 $H \not\vDash_{\mathbb{A}} \text{False}$
 $H \vDash_{\mathbb{A}} \neg a$ if $H \not\vDash_{\mathbb{A}} a$
 $H \vDash_{\mathbb{A}} a_1 \wedge a_2$ if $H \vDash_{\mathbb{A}} a_1 \wedge H \vDash_{\mathbb{A}} a_2$
 $H \vDash_{\mathbb{A}} a_1 \vee a_2$ if $H \vDash_{\mathbb{A}} a_1 \vee H \vDash_{\mathbb{A}} a_2$
 $H \vDash_{\mathbb{A}} a_1 * a_2$ if $\exists H_1, H_2 \in \text{Heaps}^{\text{log}}. H = H_1 + H_2 \wedge H_1 \vDash_{\mathbb{A}} a_1 \wedge H_2 \vDash_{\mathbb{A}} a_2$
 $H \vDash_{\mathbb{A}} [f]\ell \mapsto v$ if $H(\ell \mapsto v) \geq f$
 $H \vDash_{\mathbb{A}} \bigvee A$ if $\exists a \in A. H \vDash_{\mathbb{A}} a$
 $H \vDash_{\mathbb{A}} [f]\text{uninit}(\ell)$ if $H(\text{uninit}_{\text{Res}}(\ell)) \geq f$
 $H \vDash_{\mathbb{A}} [f]\text{mutex}(m, P)$ if $H(\text{mutex}_{\text{Res}}(m, P)) \geq f$
 $H \vDash_{\mathbb{A}} [f]\text{locked}(m, P, f_u)$ if $H(\text{locked}_{\text{Res}}(m, P, f_u)) \geq f$
 $H \vDash_{\mathbb{A}} [f]\text{signal}(s, b)$ if $H(\text{signal}_{\text{Res}}(s, b)) \geq f$
 $H \vDash_{\mathbb{A}} \text{obs}(O)$ if $H(\text{obs}_{\text{Res}}(O)) \geq 1$
 $H \vDash_{\mathbb{A}} \text{wperm}(id, \delta)$ if $H(\text{wperm}_{\text{Res}}(id, \delta)) \geq 1$
 $H \vDash_{\mathbb{A}} \text{itperm}(\delta)$ if $H(\text{itperm}_{\text{Res}}(\delta)) \geq 1$

(e) Assertion model relation. $H \not\vDash_{\mathbb{A}} a$ expresses that $H \vDash_{\mathbb{A}} a$ does not hold.

Fig. 5: Assertions.

$h^a, H, \text{while } c \text{ do skip} \rightsquigarrow_{\text{ast}} h^a, H, \text{if } c \text{ then } (\text{consumeItPerm}; \text{while } c \text{ do skip})$

$$\frac{h^a, H + \{\text{itperm}_{\text{Res}}(\delta)\}, \text{consumeItPerm} \rightsquigarrow_{\text{ast}} h^a, H, \text{tt} \quad \begin{array}{l} \text{signal}_{\text{aRes}}(s, \text{False}) \in h^a \quad P^a(\theta) = (H, c) \\ H(\text{wperm}_{\text{Res}}(s.\text{id}, \delta)) \geq 1 \quad H(\text{obs}_{\text{Res}}(O)) \geq 1 \quad \forall o \in O. \text{lev}(s) <_{\text{L}} \text{lev}(O) \end{array}}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{itperm}_{\text{Res}}(\delta)\}, c)]}$$

Fig. 6: Example reduction rules for annotated reduction semantics. Auxiliary command **consumeItPerm** expresses consumption of an iteration permission.

$\text{safe}(H, c)$. Intuitively, a command c is *safe* under a logical heap H if H provides the necessary resources so that for every execution of c , there is a corresponding annotated execution of c that does not get stuck. We write: (i) $\text{annot}_{\text{tp}}(P^a, P)$ to express that P^a is an annotated version of P . (ii) $h^a \sim_{\text{ah} \sim_{\text{ph}}} h$ to express that h^a is *compatible* with h . (iii) And $\text{consistent}_{\text{conf}}(h^a, P^a)$ to express that the annotated machine configuration (h^a, P^a) is *consistent* in the sense that h^a is compatible with the logical heap H containing all resources owned by threads in P^a or protected by unlocked locks in h^a and that heap locations in H are unique.

Definition 5 (Safety). We define the safety predicate $\text{safe} \subseteq \text{Heaps}^{\text{log}} \times \text{Cmds}$ coinductively as the greatest solution (with respect to \subseteq) of the following equation:

$$\begin{aligned} & \text{safe}(H, c) \\ & \iff \\ & \text{complete}_{\text{th}}(H) \rightarrow \\ & \forall P, P'. \forall \theta \in \text{dom}(P). \forall h, h'. \forall P^a. \forall h^a. \\ & \text{consistent}_{\text{conf}}(h^a, P^a) \wedge h^a \sim_{\text{ah} \sim_{\text{ph}}} h \wedge \\ & P(\theta) = c \wedge P^a(\theta) = (H, c) \wedge \text{annot}_{\text{tp}}(P^a, P) \wedge h, P \xrightarrow{\theta}_{\text{tp}} h', P' \rightarrow \\ & \exists P^G, P^{a'}. \exists h^G, h^{a'}. \\ & h^a, P^a \xrightarrow{\theta}_{\text{ghost}}^* h^G, P^G \wedge h^G, P^G \xrightarrow{\theta}_{\text{real}} h^{a'}, P^{a'} \wedge \text{annot}_{\text{tp}}(P^{a'}, P') \wedge \\ & h^{a'} \sim_{\text{ah} \sim_{\text{ph}}} h' \wedge \\ & \forall (H_f, c_f) \in \text{range}(P^{a'}) \setminus \text{range}(P^a). \text{safe}(H_f, c_f). \end{aligned}$$

Definition 6 (Hoare Triple Model Relation). We define the model relation for Hoare triples $\models_{\text{H}} \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$ such that:

$$\begin{aligned} & \models_{\text{H}} \{A\} c \{ \lambda r. B(r) \} \\ & \iff \\ & \forall H_F. \forall E. (\forall v. \forall H_B. H_B \models_{\text{A}} B(v) \rightarrow \text{safe}(H_B + H_F, E[v])) \\ & \rightarrow \forall H_A. H_A \models_{\text{A}} A \rightarrow \text{safe}(H_A + H_F, E[c]) \end{aligned}$$

We can instantiate context E in above definition to $\text{let } x := \square \text{ in tt}$, which yields the consequent $\text{safe}(H_A + H_F, \text{let } x := c \text{ in tt})$. Note that this implies

$\text{safe}(H_A + H_F, c)$. Further, every specification we can derive in our proof system also holds in our model.

Lemma 1 (Hoare Triple Soundness). *Let $\vdash \{A\} c \{B\}$ hold; then also $\vDash_{\text{H}} \{A\} c \{B\}$ holds.*

Proof. Proof by induction on the derivation of $\vdash \{A\} c \{B\}$.

Constructing Annotated Executions Given a command c which provably discharges all obligations and a fair reduction sequence for c , we can construct a corresponding annotated reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$. This is a useful tool to analyse program executions as $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ carries much more information than the original sequence, e.g., which obligations and permissions a thread holds. Note that our definition of fairness forbids $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ to perform ghost steps forever and that we use $h^a \text{ah} \sim_{\text{lh}} H$ to express that h^a is compatible with H .

Definition 7 (Fairness of Annotated Reduction Sequences). *We call an annotated reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ fair iff for all $i \in \mathbb{N}$ and $\theta \in \text{dom}(P_i^a)$ with $P_i^a(\theta).\text{cmd} \neq \text{term}$ there exists some $k \geq i$ with*

$$h_k^a, P_k^a \overset{\theta}{\rightsquigarrow}_{\text{real}} h_{k+1}^a, P_{k+1}^a.$$

Lemma 2 (Construction of Annotated Reduction Sequences). *Suppose we can prove $\vDash_{\text{H}} \{A\} c \{\text{obs}(\emptyset)\}$. Let H_A be a logical heap with $H_A \vDash_A A$ and $\text{complete}_{\text{lh}}(H_A)$ and h_0^a an annotated heap with $h_0^a \text{ah} \sim_{\text{lh}} H_A$. Let $(h_i, P_i)_{i \in \mathbb{N}}$ be a fair plain reduction sequence with $h_0^a \text{ah} \sim_{\text{ph}} h_0$ and $P_0 = \{(\theta_0, c)\}$ for some thread ID θ_0 and command c .*

Then, there exists a fair annotated reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ with $P^a = \{(\theta_0, (H_A, c))\}$ and $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$ for all $i \in \mathbb{N}$.

Proof. We can construct the annotated reduction sequence inductively from the plain reduction sequence.

Program Order Graph In the remainder of this section, we prove that programs where every thread discharges all obligations terminate. For this, we need to analyse each thread's control flow, i.e., the subsequence of execution steps belonging to the thread. We do this by taking a sequence $(h_i, P_i)_{i \in \mathbb{N}}$ representing a program execution, constructing an annotation $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ carrying additional information and then analysing its *program order graph* $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ which represents the execution's control flow. The graph has the form $G((h_i^a, P_i^a)_{i \in \mathbb{N}}) = (\mathbb{N}, E)$ for $E \subset (\mathbb{N} \times \Theta \times N^r \times \mathbb{N})$ where N^r is the set of (annotated) reduction rule names.

Every node $i \in \mathbb{N}$ represents the i^{th} reduction step, i.e. $h_i^a, P_i^a \overset{\theta}{\rightsquigarrow}_{\text{atp}} h_{i+1}^a, P_{i+1}^a$. Consider an edge $(i, \theta, n, j) \in E$ between two nodes i, j . n describes the combination of reduction rules applied in step i . The edge expresses one of two things: (i) Either, steps i and j belong to the same thread θ and j denotes the first scheduling of θ after i . (ii) Or, step i forks the new thread θ and j is the first

time θ is scheduled. A path in this graph represents the control flow of a directly related line of threads, i.e., parent, child, etc. In case the reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ is clear from the context, we denote the graph by G to simplify the notation.

Signal Capacity Threads can wait for signals to generate iteration permissions and thereby justify loop iterations. Consider some path p in a program order graph that represents the execution of some thread (and potentially some direct descendants) in a verified program. We prove that p is finite by (i) considering the stock of held iteration permissions and of future iteration permissions created via waiting and (ii) proving that this stock decreases. To do this, we need a notion that precisely captures these *future iteration permissions*. As a first step, we consider those signals that are waited for only finitely often along p . To each such signal s , we assign a *capacity* at every node i on p , i.e., $\text{sigCap}_p(i)$, which is the bag of iteration permissions that will be created by waiting for s after i along the remaining suffix of p .

For the following definition, note that we represent paths in a program order graph by subgraphs and that S_G^∞ denotes the set of signals waited for infinitely often within a subgraph G . Further, $\text{waitEdges}_G(s)$ denotes the set of edges (a, θ, n, b) in G where a represents a wait step. We call these edges *wait edges*. For any edge, $\text{itperms}_G((a, \theta, n, b))$ denotes the bag of iteration permissions generated by step a .

Definition 8 (Signal Capacity). *Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be a fair annotated reduction sequence and $G = (V, E)$ be a subgraph of the sequence's program order graph. We define the function $\text{sigCap}_G : (S \setminus S_G^\infty) \times \mathbb{N} \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$ mapping signals and indices to bags of iteration permissions as follows:*

$$\text{sigCap}_G(s, i) := \bigoplus_{\substack{(a, \theta, n, b) \in \text{waitEdges}_G(s) \\ a \geq i}} \text{itperms}_G((a, \theta, n, b)).$$

We call $\text{sigCap}_G(s, i)$ the capacity of signal s at index i .

Consider any path p representing some thread's execution on which no signal is waited for infinitely often. For every node i on p , we consider the stock of iteration permissions that are either held by the thread in step i or that will be created along the rest of the path. We can prove that this stock of permissions decreases at each loop iteration and that p must hence be finite.

Lemma 3. *Let $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ be a program order graph and let $p = (V, E)$ be a path in G with $S_p^\infty = \emptyset$. For every $\theta \in \text{dom}(P_0^a)$ let $P_0^a(\theta).\text{heap}$ be finite and complete. Then, p is finite.*

Proof (Sketch). Assume p is infinite. We prove a contradiction by assigning a finite capacity to every node along the path. Let θ_i be the ID of the thread reduced in step i .

Consider the function $\text{nodeCap} : V \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$ defined as

$$\text{nodeCap}(i) := \text{itPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap}) \uplus \biguplus_{\substack{id \in \{id' \mid (id', _)\in \text{waitPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})\} \\ L \in \mathcal{Levs}}} \text{sigCap}_p((id, L), i).$$

where $\text{itPerms}_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \text{Bags}(\Lambda)$ and $\text{waitPerms}_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \text{Bags}(\Omega)$ map any logical heap to the bag of contained iteration and wait permissions.

For every $i \in V$, the capacity of node i , i.e., $\text{nodeCap}(i)$, is the union of two finite iteration permission bags: (i) To the left $\text{itPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})$ captures all iteration permissions held by θ_i in step i . (ii) To the right $\biguplus \text{sigCap}_p((id, L), i)$ captures all iteration permissions that will be created along the suffix of p that starts at node i by waiting for signals for which thread θ_i already holds a wait permission (id, δ) in step i .

Note that for every $i \in V$, the bag of iteration permissions returned by $\text{nodeCap}(i)$ is indeed finite. All initial thread-local heaps are finite. Consequently, $\text{itPerms}_{\text{lh}}(P_0^{\text{a}}.\text{heap})$ and $\text{waitPerms}_{\text{lh}}(P_0^{\text{a}}.\text{heap})$ are finite. Our annotated semantics preserves this finiteness and hence $\text{itPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})$ and $\text{waitPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})$ are finite as well. Since signal IDs are unique, for every fixed choice of i and id , there is at most one level L , for which $\text{sigCap}_p((id, L), i) \neq \emptyset$. By assumption, along p all signals are waited for only finitely often, i.e., $S_p^\infty = \emptyset$. Hence, also the big union $\biguplus \text{sigCap}_p((id, L), i)$ is defined and finite.

The well-founded partial order $\prec_\Delta \subseteq \Lambda \times \Lambda$ induces a well-founded partial order $\prec_\Lambda \subseteq \text{Bags}_{\text{fin}}(\Lambda) \times \text{Bags}_{\text{fin}}(\Lambda)$ for finite iteration permission bags [6]. Consider the sequence $(\text{nodeCap}(i))_{i \in V}$. Since every element is a finite bag of iteration permissions, we can order it by \prec_Λ . We can prove that nodeCap is monotonically decreasing and that $\text{nodeCap}(j) \prec_\Lambda \text{nodeCap}(i)$ holds for edges (i, θ, n, j) corresponding to a loop iteration consuming an iteration permission. We can also prove that every infinite path, such as p , contains infinitely many such edges. Hence, $(\text{nodeCap}(i))_{i \in V}$ is an infinitely descending chain. This contradicts the well-foundedness of \prec_Λ .

We proceed to prove that no signals are waited for infinitely often.

Lemma 4. *Let $(h_i^{\text{a}}, P_i^{\text{a}})_{i \in \mathbb{N}}$ be a fair annotated reduction sequence with $P_0^{\text{a}} = \{(\theta_0, (H_0, c))\}$, $\text{finite}_{\text{ah}}(h_0^{\text{a}})$, $\text{complete}_{\text{lh}}(H_0)$, $\text{finite}_{\text{lh}}(H_0)$ and $\text{consistent}_{\text{conf}}(h_0^{\text{a}}, P_0^{\text{a}})$. Let H_0 contain no signal or wait permission chunks. Further, let h_0^{a} contain no chunks $\text{unlocked}_{\text{aRes}}(m, P, H_P)$ where H_P contains any signal chunks. Let G be the program order graph of $(h_i^{\text{a}}, P_i^{\text{a}})_{i \in \mathbb{N}}$. Then, $S_G^\infty = \emptyset$.*

Proof (Sketch). Suppose $S_G^\infty \neq \emptyset$. Since \mathcal{Levs} is well-founded, the same holds for the set $\{\text{lev}(s) \mid s \in S^\infty\}$. Hence, there is some $s_{\text{min}} \in S^\infty$ for which no $z \in S^\infty$ with $\text{lev}(z) <_{\text{L}} \text{lev}(s_{\text{min}})$ exists.

Since neither the initial logical heap H_0 nor any unlocked lock invariant stored in h_0^{a} does contain any signals, s_{min} must be created during the reduction sequence. The reduction step creating signal s_{min} simultaneously creates an obligation to set s_{min} . Analogous to our proof rules, the annotated semantics only allow

threads to wait for a signal s (i) if the signal's level is lower than the level of each held obligation and (ii) if s has not been set yet. Hence, we know for every wait edge (a, θ, n, b) referring to s_{\min} that (i) thread θ does not hold any obligation for s_{\min} (i.e., there is exactly one obligations chunk with $P_a^a(\theta).\text{heap}(\text{obs}_{\text{IRes}}(O)) = 1$ and for this chunk $s_{\min} \notin O$ holds) and (ii) s_{\min} has not been set, yet (i.e. $\text{signal}_{\text{aRes}}(s_{\min}, \text{False}) \in h_a^a$). Hence, in step a another thread $\theta_{\text{ob}} \neq \theta$ must hold the obligation for s_{\min} (i.e. $P_a^a(\theta_{\text{ob}}).\text{heap}(\text{obs}_{\text{IRes}}(O)) = 1$ for some bag of obligations O with $s_{\min} \in O$). Since there are infinitely many wait edges concerning s_{\min} in G , the signal is never set.

By fairness, for every wait edge as above, there must be a non-ghost reduction step $h_k^a, P_k^a \xrightarrow{\theta_{\text{ob}}}_{\text{atp}} h_{k+1}^a, P_{k+1}^a$ of the thread θ_{ob} holding the obligation for s_{\min} with $k \geq a$. Hence, there exists an infinite path p_{ob} in G where each edge $(e, \theta_{\text{ob}}, n, f) \in \text{edges}(p_{\text{ob}})$ concerns some thread θ_{ob} holding the obligation for s_{\min} . (Note that this thread ID does not have to be constant along the path, since the obligation can be passed on during fork steps.)

The path p_{ob} does not contain wait edges $(e, \theta_{\text{ob}}, n, f)$ for any $s^\infty \in S^\infty$, since this would require s^∞ to be of a lower level than all held obligations. This restriction implies $\text{lev}(s^\infty) <_{\text{L}} \text{lev}(s_{\min})$ and would hence contradict the minimality of s_{\min} . That is, $S_{p_{\text{ob}}}^\infty = \emptyset$.

The annotated reduction semantics preserve the finiteness of thread-local heaps. Since H_0 is finite, the same holds for every logical heap associated with the root of p_{ob} . This allows us to apply Lemma 3, by which we get that p_{ob} is finite. A contradiction.

Lemma 5. *Let $\models_{\text{H}} \{\text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\delta_i)\} c \{\text{obs}(\emptyset)\}$ hold. There exists no fair, infinite annotated reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ with $P_0^a = \{(\theta_0, (H_0, c))\}$, $h_0^a = \emptyset$ and $H_0 = \{\text{obs}_{\text{IRes}}(\emptyset), \text{itperm}_{\text{IRes}}(\delta_1), \dots, \text{itperm}_{\text{IRes}}(\delta_N)\}$.*

Proof. Suppose a reduction sequence as described above exists. We are going to prove a contradiction by considering its infinite program order graph G .

Since P_0^a contains only a single thread, G is a binary tree with an infinite set of vertices. By the Weak König's Lemma [34] G has an infinite branch, i.e. an infinite path p starting at root \emptyset .

The initial logical heap H_0 is complete and finite and the initial annotated machine configuration (h_0^a, P_0^a) is consistent. By Lemma 4 we know that $S_G^\infty = \emptyset$. Since $S_p^\infty \subseteq S_G^\infty$, we get $S_p^\infty = \emptyset$. This allows us to apply Lemma 3, by which we get that p is finite, which is a contradiction.

Theorem 1 (Soundness). *Let $\vdash \{\text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\delta_i)\} c \{\text{obs}(\emptyset)\}$ hold. There exists no fair, infinite reduction sequence $(h_i, P_i)_{i \in \mathbb{N}}$ with $h_0 = \emptyset$ and $P_0 = \{(\theta_0, c)\}$ for any choice of θ_0 .*

Proof. Assume that such a reduction sequence exists. By Hoare triple soundness, Lemma 1, we get $\models_{\text{H}} \{\text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$ from the assumption $\vdash \{\text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)\} c \{\text{obs}(\emptyset)\}$. Consider the logical heap $H_0 =$

$\{\text{obs}_{|\text{Res}}(\emptyset), \text{itperm}_{|\text{Res}}(\delta_1), \dots, \text{itperm}_{|\text{Res}}(\delta_N)\}$ and the annotated heap $h_0^a = \emptyset$. It holds $H_0 \vDash_A \text{obs}(\emptyset) * \bigstar_{i=1, \dots, N} \text{itperm}(\tau, \delta_i)$, $h_0^a \text{ah} \sim_{\text{lh}} H_0$ (since H_0 does not contain any logical resources with an annotated counterpart) and $h_0^a \text{ah} \sim_{\text{ph}} h_0$ (since both heaps are empty). This allows us to apply Lemma 2, by which we can construct a corresponding fair annotated reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ that starts with $h_0^a = \emptyset$ and $P_0^a = \{(\theta_0, (H_0, c))\}$. By Lemma 5 $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ does not exist. A contradiction.

6 A Realistic Example

To demonstrate the expressiveness of the presented verification approach, we verified the termination of the program presented in Fig. 7a. It involves two threads, a consumer and a producer, communicating via a shared bounded FIFO with a maximal capacity of 10. The producer enqueues numbers $100, \dots, 1$ into the FIFO and the consumer dequeues those. Whenever the queue is full, the producer busy-waits for the consumer to dequeue an element. Likewise, whenever the queue is empty, the consumer busy-waits for the producer to enqueue the next element. Each thread's termination depends on the other thread's productivity. This is, however, no cyclic dependency. For instance, in order to prove that the producer eventually pushes number i into the queue, we only need to rely on the consumer to pop $i + 10$. A similar property holds for the consumer.

Fine-Tuning Signal Creation To simplify complex proofs involving many signals we refine the process of creating a new ghost signal. For simplicity, we combined the allocation of a new signal ID and its association with a level and a boolean in one step. For some proofs, such as the one we outline in this section, it can be helpful to fix the IDs of all signals that will be created throughout the proof already at the beginning. To realize this, we replace view shift rule VS-NEWSIGNAL by the rules presented in Fig. 7b and adapt our logical resource chunks accordingly. With these more fine-grained view shifts, we start by allocating a signal ID, cf. VS-ALLOCSIGID. Thereby we obtain an *uninitialized* signal $\text{uninitSig}(\ell)$ that is not associated with any level or boolean, yet. Also, allocating a signal ID does not create any obligation because threads can only wait for *initialized* (and unset) signals. When we initialize a signal, we bind its already allocated ID to a level of our choice and associate the signal with **False**, cf. VS-SIGINIT. This creates an obligation to set the signal.

This change does not affect the soundness of our verification approach. However, it drastically simplifies proofs, such as the one discussed below, that would normally require the conditional creation of new signals within loops. Using the finer-grained view shifts, we can create all the signals we need centrally at the beginning of the proof. Since we thereby do not create any obligations, this step does not affect the acquisition of locks later on, as long as the signals remain uninitialized. Moreover, since we already know all signal IDs, we can centrally create all the wait permissions we need.

```

let fifo10 := cons(nil) in let m := new_mutex in
let cp := cons(100) in let cc := cons(100) in
fork (
  while (
    acquire m;
    let f := [fifo10] in
    if size(f) < 10 then (
      let c := [cp] in
      [fifo10] := f · (c :: nil);
      [cp] := c - 1
    );
    release m;
    let c := [cp] in
    c ≠ 0
  ) do skip;
);

while (
  acquire m;
  let f := [fifo10] in
  if size(f) > 0 then (
    let c := [cc] in
    [fifo10] := tail(f);
    [cc] := c - 1
  );
  release m;
  let c := [cc] in
  c ≠ 0
) do skip

```

(a) Realistic example program with a producer and a consumer thread communicating via a shared bounded FIFO.

<p>VS-ALLOC SIGID $\text{True} \Rightarrow \exists id. \text{uninitSig}(id)$</p>	<p>VS-SIGINIT $\text{obs}(O) * \text{uninitSig}(id)$ $\Rightarrow \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})$</p>
--	--

(b) Fine-grained view shift rules for signal creation.

Fig. 7: Realistic verification example.

Creating Permissions & Signals We choose $\Delta = \mathbb{N}$ and start our verification with a single iteration permission $\text{itperm}(2)$ and without obligations. Initially we create 400 iteration permissions of degree 1 by weakening the initial one. Additionally, we allocate 200 signal IDs $id_{\text{push}}^{100}, \dots, id_{\text{push}}^1, id_{\text{pop}}^{100}, \dots, id_{\text{pop}}^1$. We use 200 of our iteration permissions to create one wait permission of degree 0 for each signal. 100 of the remaining 200 iteration permissions will be used by each thread to justify the iterations in which it pushes or pops, respectively, i.e., the productive iterations in which it does not have to wait.

Using the Signals We are going to ensure that always at most one push signal and at most one pop signal are initialized and unset. The producer and consumer are going to hold the obligation for the push and pop signal, respectively. Eventually, the producer will hold the obligation for s_{push}^i while the signal is initialized and it will set s_{push}^i when it pushes the number i into the FIFO. Meanwhile, the consumer will use s_{push}^i and $\text{wperm}(id_{\text{push}}^i, 0)$ to wait for the number i to arrive in the queue when it is empty. Similarly, the consumer will hold the obligation for s_{pop}^i and will set s_{pop}^i when it pops the number i . The producer uses s_{pop}^i and $\text{wperm}(id_{\text{pop}}^i, 0)$ to wait for the consumer to pop i from the queue when it is full.

Choosing the Levels Note that we ignored the levels so far. The producer and the consumer both acquire the mutex while holding an obligation for a signal. Hence, we choose $\mathcal{L}evs = \mathbb{N}$, $\text{lev}(m) = 0$ and $\text{lev}(s) > 0$ for every signal s . Both threads will wait at the end of an iteration inside a release view shift, i.e., after they discharged the obligation to release the mutex. So, the mutex level does not interfere with the wait steps. The producer waits when the queue is full and the consumer when it is empty. That is, the producer waits for s_{pop}^{i+10} while holding an obligation for s_{push}^i and the consumer waits for s_{push}^i while holding an obligation for s_{pop}^i . So, we have to choose the signal levels such that $\text{lev}(s_{\text{pop}}^{i+10}) < \text{lev}(s_{\text{push}}^i)$ and $\text{lev}(s_{\text{push}}^i) < \text{lev}(s_{\text{pop}}^i)$ hold. We solve this by choosing $\text{lev}(s_{\text{pop}}^i) = 102 - i$ and $\text{lev}(s_{\text{push}}^i) = 101 - i$.

Distributing Resources Both threads count from 100 downwards. So, we use the allocated IDs $id_{\text{push}}^{100}, id_{\text{pop}}^{100}$ to initialize our first two signals $s_{\text{push}}^{100} = (id_{\text{push}}^{100}, 1)$, $s_{\text{pop}}^{100} = (id_{\text{pop}}^{100}, 2)$ and transfer their ownership to the mutex. We split the mutex chunk and pass one half to each thread such that both can access the initialized signals. Additionally, the producer thread receives the 100 iteration permissions, the obligation for s_{push}^{100} , the uninitialized signal chunks $\text{uninitSig}(id_{\text{push}}^{99}), \dots, \text{uninitSig}(id_{\text{push}}^1)$ and the wait permissions $\text{wperm}(id_{\text{pop}}^{100}, 0), \dots, \text{wperm}(id_{\text{pop}}^1, 0)$ upon its forking. The rest remains with the consumer, i.e., 100 iteration permissions, the obligation for s_{pop}^{100} , uninitialized signal chunks $\text{uninitSig}(id_{\text{pop}}^{99}), \dots, \text{uninitSig}(id_{\text{pop}}^1)$ and wait permissions $\text{wperm}(id_{\text{pop}}^{100}, 0), \dots, \text{wperm}(id_{\text{pop}}^1, 0)$.

Verifying Termination This setup suffices to verify the example program. Each thread holds (i) the iteration permissions it needs to justify its productive iter-

ations and (ii) the wait permissions that in combination with the acquired lock invariant allow to wait until the next productive iteration. Whenever the producer pushes a number i into the queue, it sets s_{push}^i which discharges the held obligation and decreases its counter. Afterwards, if $i > 1$, it uses the uninitialized signal chunk $\text{uninitSig}(id_{\text{push}}^{i-1})$ to initialize $s_{\text{push}}^{i-1} = (id_{\text{push}}^{i-1}, 101 - (i - 1))$ and replaces s_{push}^i in the lock invariant by s_{push}^{i-1} before it releases the lock. If $i = 1$, the counter reached 0 and the loop ends. In this case, the producer holds no obligation. The consumer behaves similarly. Since we proved that each thread discharged all its obligations, we proved that the program terminates. We present the verification sketch in the technical report [27]. Furthermore, we encoded [30] the proof in VeriFast [15].

7 Specifying Busy-Waiting Concurrent Objects

Our approach can be used to verify busy-waiting concurrent objects with respect to abstract specifications. For example, we have verified [29] the CLH lock [11] against a specification that is very similar to our proof rules for built-in mutexes shown in Fig. 2. The main difference is that it is slightly more abstract: when a lock is initialized, it is associated with a *bounded infinite set* of levels rather than with a single particular level. (To make this possible, an appropriate universe of levels should be used, such as the set of lists of natural numbers, ordered lexicographically.) To acquire a lock, the levels of the obligations held by the thread must be above the elements of the set; the new obligation’s level is an element of the set.

8 Allowing Permission Transfer

In the logic presented so far, transferring permissions between threads is not allowed, to prevent self-fueling busy-waiting loops. However, as we showed in earlier work [14], permission transfer is useful for verifying termination of non-blocking algorithms involving compare-and-swap loops. For this reason, we relax our logic as follows: we qualify each permission by the *thread phase* that produced it, and we allow a thread in phase p to consume a permission only if it was produced by an *ancestor thread phase* $p' \sqsubseteq p$. Initially, the main thread is in the *root phase* ϵ , and all permissions provided as part of the program’s precondition are considered to have been produced by the root phase. When a thread t running in phase p forks a new thread, the new thread’s initial phase is the new phase $p.\text{Forkee}$ and t itself moves into the new phase $p.\text{Forker}$. Weakening an iteration permission and converting an iteration permission into a wait permission do not affect the permission’s phase. As a result, in a program that does not involve busy-waiting, all permissions are qualified by the root phase and can be transferred between threads freely. Permissions produced by waiting in a phase p , however, can be transferred only between descendants of p . We formalize this relaxed logic in the technical report [27].

9 Tool Support

We have extended our VeriFast tool for separation logic-based modular verification of C and Java programs so that it supports verifying termination of busy-waiting C or Java programs. We had already added *call permissions* (similar to iteration permissions) to VeriFast as part of earlier work [14] on verifying termination of programs using primitive blocking constructs without assuming fairness of the thread scheduler. The only change we had to make to VeriFast’s symbolic execution engine was to enforce the thread phase rule mentioned in § 8. We encoded the other aspects of the logic simply as axioms in a *trusted header file*. We used this tool support to verify the bounded FIFO (§ 6) and the CLH lock (§ 7).

10 Integrating Higher-Order Features

The logic we presented in this paper does not support higher-order features such as assertions that quantify over assertions, or storing assertions in the (logical) heap as the values of ghost cells. While we did not need such features to carry out our example proofs, they are generally useful to verify higher-order program modules against abstract specifications. The typical way to support such features in a program logic is by applying *step indexing* [1,2], where the domain of logical heaps is indexed by the number of execution steps left in the (partial) program trace under consideration. Assertions stored in a logical heap at index $n + 1$ talk about logical heaps at index n ; i.e., they are meaningful only *later*, after at least one more execution step has been performed.

It follows that such logics apply directly only to *partial* correctness properties. Fortunately, we can reduce a termination property to a safety property by writing our program in a programming language *instrumented* with run-time checks that guarantee termination. Specifically, we can write our program in a programming language that tracks signals, obligations, iteration permissions, and wait permissions at run time, has constructs for iteration permission weakening, signal creation, wait permission creation, waiting, and setting a signal, and where the **fork** command takes as an extra operand the list of obligations to be transferred to the new thread (and the other constructs similarly take sufficient operands to eliminate any need for angelic choice), and that gets stuck when these constructs’ preconditions are not satisfied, such as when a thread waits for a signal while holding the obligation for that signal, or when it performs a loop iteration without holding an iteration permission. We can then use a step-indexing-based higher-order logic such as Iris [17] to verify that our program never gets stuck. Once we established this, we know none of the instrumentation has any effect and can be safely *erased* from the program.

11 Related & Future Work

Liang and Feng [21, 22] propose LiLi, a separation logic to verify liveness of blocking constructs implemented via busy-waiting. In contrast to our verification

approach, theirs is based on the idea of contextual refinement. In their approach, client code involving calls of blocking methods of the concurrent object is verified by first applying the contextual refinement result to replace these calls by code involving primitive blocking operations and then verifying the resulting client code using some other approach. In contrast, specifications in our approach are regular Hoare-style triples and proofs are regular Hoare-style proofs.

D’Osualdo et al. [7] propose TaDA Live, a separation logic for verifying termination of busy-waiting. This logic allows to modularly reason about fine-grained concurrent programs and blocking operations that are implemented in terms of busy-waiting and non-blocking primitives. It uses the concept of obligations to express thread-local liveness invariants, e.g., that a thread eventually releases an acquired lock. One difference with our work is that TaDA Live’s assertions are not syntactically stable; that is, certain proof rules require the user to prove stability of certain assertions with respect to environment actions as a side condition. This may make it more difficult to provide effective tool support: whereas we could easily implement our logic in our VeriFast tool, and we expect our logic to be easy to integrate into other separation logics such as Iris [17] and their tool support, we are not aware of tool support for TaDA Live at this time.

In recent work [13] we proposed a Hoare logic to verify liveness properties of the I/O behaviour of programs that do not perform busy waiting. By combining that approach with the present one, we expect to be able to verify I/O liveness of realistic concurrent programs involving both I/O and busy waiting, such as a server where one thread receives requests and enqueues them into a bounded FIFO, and another one dequeues them and responds. To support this claim, we encoded the combined logic in VeriFast and verified a simple server application where the receiver and responder thread communicate via a shared buffer [28].

12 Conclusion

In this paper we proposed a separation logic to verify termination of programs with busy-waiting. We proved our logic sound and demonstrated its usability by verifying a realistic example. Further, we encoded our logic and the realistic example in VeriFast [30] and used this encoding also to verify the CLH lock [29]. Moreover, we expect that our approach can be straightforwardly integrated into other existing concurrent separation logics such as Iris [17].

References

1. The category-theoretic solution of recursive metric-space equations. *Theoretical Computer Science* **411**(47), 4102 – 4122 (2010). <https://doi.org/10.1016/j.tcs.2010.07.010>
2. Appel, A.W., McAllester, D.: An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.* **23**(5), 657–683 (Sep 2001). <https://doi.org/10.1145/504709.504712>
3. Bornat, R., Calcagno, C., O’Hearn, P., Parkinson, M.J.: Permission accounting in separation logic. In: *POPL ’05* (2005). <https://doi.org/10.1145/1040305.1040327>

4. Boyland, J.: Checking interference with fractional permissions. In: SAS (2003). https://doi.org/10.1007/3-540-44898-5_4
5. Brookes, S., O’Hearn, P.: Concurrent separation logic. ACM SIGLOG News **3**, 47–65 (2016). <https://doi.org/10.1145/2984450.2984457>
6. Dershowitz, N., Manna, Z.: Proving termination with multiset orderings. In: ICALP (1979). https://doi.org/10.1007/3-540-09510-1_15
7. D’Ousaldo, E., Farzan, A., Gardner, P., Sutherland, J.: Tada live: Compositional reasoning for termination of fine-grained concurrent programs. CoRR **abs/1901.05750** (2019), <http://arxiv.org/abs/1901.05750>
8. Gotsman, A., Berdine, J., Cook, B., Rinetzky, N., Sagiv, S.: Local reasoning for storable locks and threads. In: APLAS (2007). https://doi.org/10.1007/978-3-540-76637-7_3
9. Hamin, J., Jacobs, B.: Deadlock-free monitors. In: ESOP (2018). https://doi.org/10.1007/978-3-319-89884-1_15
10. Hamin, J., Jacobs, B.: Transferring Obligations Through Synchronizations. In: Donaldson, A.F. (ed.) 33rd European Conference on Object-Oriented Programming (ECOOP 2019). Leibniz International Proceedings in Informatics (LIPIcs), vol. 134, pp. 19:1–19:58. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2019). <https://doi.org/10.4230/LIPIcs.ECOOP.2019.19>, <http://drops.dagstuhl.de/opus/volltexte/2019/10811>
11. Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming, Revised Reprint. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edn. (2012)
12. Hoare, C.A.R.: An axiomatic basis for computer programming. Commun. ACM **12**, 576–580 (1968). <https://doi.org/10.1145/363235.363259>
13. Jacobs, B.: Modular verification of liveness properties of the I/O behavior of imperative programs. Accepted at International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (2020), <https://people.cs.kuleuven.be/~bart.jacobs/isola2020.pdf>
14. Jacobs, B., Bosnacki, D., Kuiper, R.: Modular termination verification of single-threaded and multithreaded programs. ACM Trans. Program. Lang. Syst. **40**, 12:1–12:59 (2018). <https://doi.org/10.1145/3210258>
15. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: Verifast: A powerful, sound, predictable, fast verifier for c and java. In: Bobaru, M., Havelund, K., Holzmann, G., Joshi, R. (eds.) NASA Formal Methods (NFM 2011). vol. 6617, pp. 41–55. Springer (2011). https://doi.org/10.1007/978-3-642-20398-5_4, <https://lirias.kuleuven.be/95720>
16. Jung, R., Krebbers, R., Birkedal, L., Dreyer, D.: Higher-order ghost state. Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming (2016). <https://doi.org/10.1145/2951913.2951943>
17. Jung, R., Krebbers, R., Jourdan, J.H., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. J. Funct. Program. **28**, e20 (2018). <https://doi.org/10.1017/S0956796818000151>
18. Klop, J.: Term rewriting systems. In: LICS 1993 (1990)
19. Kobayashi, N.: A new type system for deadlock-free processes. In: CONCUR (2006). https://doi.org/10.1007/11817949_16
20. Leino, K.R.M., Müller, P., Smans, J.: Deadlock-free channels and locks. In: ESOP (2010). https://doi.org/10.1007/978-3-642-11957-6_22
21. Liang, H., Feng, X.: A program logic for concurrent objects under fair scheduling. In: POPL 2016 (2016). <https://doi.org/10.1145/2837614.2837635>

22. Liang, H., Feng, X.: Progress of concurrent objects with partial methods. *Proc. ACM Program. Lang.* **2**, 20:1–20:31 (2017). <https://doi.org/10.1145/3158108>
23. Mellor-Crummey, J.M., Scott, M.L.: Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Trans. Comput. Syst.* **9**, 21–65 (1991). <https://doi.org/10.1145/103727.103729>
24. Middeldorp, A., Zantema, H.: Simple termination revisited. In: CADE (1994). https://doi.org/10.1007/3-540-58156-1_33
25. Mühlemann, K.: Method for reducing memory conflicts caused by busy waiting in multiple processor synchronisation. *IEE Proceedings E - Computers and Digital Techniques* **127**(3), 85–87 (1980). <https://doi.org/10.1049/ip-e.1980.0017>
26. O’Hearn, P.W., Reynolds, J.C., Yang, H.: Local reasoning about programs that alter data structures. In: CSL (2001). https://doi.org/10.1007/3-540-44802-0_1
27. Reinhard, T., Jacobs, B.: Ghost signals: Verifying termination of busy-waiting (technical report) (2020), <https://people.cs.kuleuven.be/~tobias.reinhard/ghostSignals--TR.pdf>
28. Reinhard, T., Jacobs, B.: VeriFast proof of I/O liveness for a simple server with a receiver and a responder thread communicating via a shared buffer. (2020), https://github.com/verifast/verifast/blob/master/examples/busywaiting/ioliveness/echo_live_mt.c
29. Reinhard, T., Jacobs, B.: VeriFast proof of safety for CLH lock. (2020), <https://github.com/verifast/verifast/blob/master/examples/busywaiting/clhlock/clhlock.c>
30. Reinhard, T., Jacobs, B.: VeriFast proof of termination for consumer-producer problem with bounded FIFO. (2020), https://github.com/verifast/verifast/blob/master/examples/busywaiting/bounded_fifo.c
31. Reinhard, T., Timany, A., Jacobs, B.: A separation logic to verify termination of busy-waiting for abrupt program exit. Accepted at Formal Techniques for Java-like Programs (2020), <https://arxiv.org/abs/2010.07800>
32. Reinhard, T., Timany, A., Jacobs, B.: A separation logic to verify termination of busy-waiting for abrupt program exit: Technical report (2020), <https://arxiv.org/abs/2007.10215>
33. Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science* pp. 55–74 (2002). <https://doi.org/10.1109/LICS.2002.1029817>
34. Simpson, S.: Subsystems of second order arithmetic. In: *Perspectives in mathematical logic* (1999). <https://doi.org/10.1017/CBO9780511581007>
35. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* **5**, 285–309 (1955). <https://doi.org/10.2307/2963937>

Appendix

A General

Definition 9 (Projections). For any Cartesian product $C = \prod_{i \in I} A_i$ and any index $k \in I$, we denote the k^{th} projection by $\pi_k^C : \prod_{i \in I} A_i \rightarrow A_k$. We define

$$\pi_k^C((a_i)_{i \in I}) := a_k.$$

In case the domain C is clear from the context, we write π_k instead of π_k^C .

Definition 10 (Disjoint Union). Let A, B be sets. We define their disjoint union as

$$A \sqcup B := A \cup B$$

if $A \cap B = \emptyset$ and leave it undefined otherwise.

Definition 11 (Bags). For any set X we define the set of bags $\text{Bags}(X)$ and the set of finite bags $\text{Bags}_{\text{fin}}(X)$ over X as

$$\begin{aligned} \text{Bags}(X) &:= X \rightarrow \mathbb{N}, \\ \text{Bags}_{\text{fin}}(X) &:= \{B \in \text{Bags}(X) \mid \{x \in B \mid B(x) > 0\} \text{ finite}\}. \end{aligned}$$

We define union and subtraction of bags as

$$\begin{aligned} (B_1 \uplus B_2)(x) &:= B_1(x) + B_2(x), \\ (B_1 \setminus B_2)(x) &:= \max(0, B_1(x) - B_2(x)). \end{aligned}$$

For finite bags where the domain is clear from the context, we define the following set-like notation:

$$\begin{aligned} \emptyset &:= x \mapsto 0, \\ \{x\} &:= \begin{cases} x \mapsto 1 \\ y \mapsto 0 \text{ for } y \neq x, \end{cases} \\ \{x_1, \dots, x_n\} &:= \biguplus_{i=1}^n \{x_i\}. \end{aligned}$$

We define the following set-like notations for element and subset relationship:

$$\begin{aligned} x \in B &\Leftrightarrow B(x) > 0, \\ B_1 \subseteq B_2 &\Leftrightarrow \forall x \in B_1. B_1(x) \leq B_2(x), \\ B_1 \subset B_2 &\Leftrightarrow \exists C \subseteq B_1. C \neq \emptyset \wedge B_1 = B_2 \setminus C. \end{aligned}$$

For any bag $B \in \text{Bags}(X)$ and predicate $P \subseteq X$ we define the following refinement:

$$\{x \in B \mid P(x)\} := \begin{cases} x \mapsto B(x) & \text{if } P(x), \\ x \mapsto 0 & \text{otherwise.} \end{cases}$$

Definition 12 (Disjoint Union). Let A, B be sets. We define their disjoint union as

$$A \sqcup B := A \cup B$$

if $A \cap B = \emptyset$ and leave it undefined otherwise.

B Language

Definition 13 (Evaluation of Closed Expressions). We define a partial evaluation function $\llbracket \cdot \rrbracket : \text{Exps} \rightarrow \text{Values}$ on expressions by recursion on the structure of expressions as follows:

$$\begin{aligned}
\llbracket v \rrbracket &:= v && \text{if } v \in \text{Values} \\
\llbracket e = e' \rrbracket &:= \text{True} && \text{if } \llbracket e \rrbracket = \llbracket e' \rrbracket \neq \perp \\
\llbracket e = e' \rrbracket &:= \text{False} && \text{if } \llbracket e \rrbracket \neq \llbracket e' \rrbracket \wedge \llbracket e \rrbracket \neq \perp \wedge \llbracket e' \rrbracket \neq \perp \\
\llbracket \neg e \rrbracket &:= \text{False} && \text{if } \llbracket e \rrbracket = \text{True} \\
\llbracket \neg e \rrbracket &:= \text{True} && \text{if } \llbracket e \rrbracket = \text{False} \\
\llbracket e \rrbracket &:= \perp && \text{otherwise}
\end{aligned}$$

We identify closed expressions e with their ascribed value $\llbracket e \rrbracket$.

C Logic

Definition 14. We define the predicate $\text{mentionsPerms} \subset \mathcal{A}$ by recursion on the structure of assertions such that the following holds:

$$\begin{aligned}
&\text{mentionsPerms}(\text{itperm}(\delta)), \\
&\text{mentionsPerms}(\text{wperm}(\text{id}, \delta)), \\
&\text{mentionsPerms}(\neg a) && \text{if } \text{mentionsPerms}(a), \\
&\text{mentionsPerms}(a_1 \otimes a_2) && \text{if } \text{mentionsPerms}(a_1) \vee \text{mentionsPerms}(a_2), \\
&\text{mentionsPerms}(\bigvee A) && \text{if } \exists a \in A. \text{mentionsPerms}(a), \\
&\text{mentionsPerms}(\text{mutex}(m, a)) && \text{if } \text{mentionsPerms}(a), \\
&\text{mentionsPerms}(\text{locked}(m, a, f)) && \text{if } \text{mentionsPerms}(a),
\end{aligned}$$

where $\otimes \in \{\wedge, \vee, *\}$.

Definition 15. We define the function $\text{getHLocs}_{\text{Res}} : \mathcal{R}^{\text{log}} \rightarrow \mathcal{Locs}$ mapping logical resources to their respective (either empty or singleton) set of involved heap locations as

$$\begin{aligned}
\text{getHLocs}_{\text{Res}}(\ell \mapsto v) &:= \{\ell\}, \\
\text{getHLocs}_{\text{Res}}(\text{uninit}_{\text{Res}}(\ell)) &:= \{\ell\}, \\
\text{getHLocs}_{\text{Res}}(\text{mutex}_{\text{Res}}((\ell, L), a)) &:= \{\ell\}, \\
\text{getHLocs}_{\text{Res}}(\text{locked}_{\text{Res}}((\ell, L), a, f)) &:= \{\ell\}, \\
\text{getHLocs}_{\text{Res}}(_) &:= \emptyset \text{ otherwise.}
\end{aligned}$$

Definition 16 (Logical Heap Consistency). We call a logical heap H consistent and write $\text{consistent}_{\text{h}}(H)$ if (i) it contains only full obligations, wait and iteration permission chunks, i.e., if

$$\begin{aligned}
H(\text{obs}_{\text{Res}}(O)) &\in \mathbb{N}, \\
H(\text{wperm}_{\text{Res}}(\text{id}, \delta)) &\in \mathbb{N}, \\
H(\text{itperm}_{\text{Res}}(\delta)) &\in \mathbb{N}
\end{aligned}$$

holds for all $O \in \text{Bags}(\mathcal{O})$, $id \in \mathcal{ID}$ and $\delta \in \Delta$ and if (ii) heap locations are unique in H , i.e., if there are no $r_1^l, r_2^l \in \mathcal{R}^{\text{log}}$ with $r_1^l \neq r_2^l$, $H(r_1^l) > 0$, $H(r_2^l) > 0$ and with $\text{getHLocs}_{\text{Res}}(r_1^l) \cap \text{getHLocs}_{\text{Res}}(r_2^l) \neq \emptyset$.

Definition 17 (View Shift). We define a view shift relation $\Rightarrow \subset \mathcal{A} \times \mathcal{A}$ according to the rules presented in Fig 8.

$$\begin{array}{c}
\text{VS-SEMIMP} \\
\frac{\forall H. \text{consistent}_{\text{th}}(H) \wedge H \models_{\mathcal{A}} A \Rightarrow H \models_{\mathcal{A}} B}{A \Rightarrow B}
\end{array}
\qquad
\begin{array}{c}
\text{VS-TRANS} \\
\frac{A \Rightarrow C \quad C \Rightarrow B}{A \Rightarrow B}
\end{array}$$

$$\begin{array}{c}
\text{VS-OR} \\
\frac{A_1 \Rightarrow B \quad A_2 \Rightarrow B}{A_1 \vee A_2 \Rightarrow B}
\end{array}
\qquad
\begin{array}{c}
\text{VS-NEWSIGNAL} \\
\text{obs}(O) \Rightarrow \exists id. \text{obs}(O \uplus \{(id, L)\}) * \text{signal}((id, L), \text{False})
\end{array}$$

$$\begin{array}{c}
\text{VS-SET SIGNAL} \\
\text{obs}(O \uplus \{s\}) * \text{signal}(s, _) \Rightarrow \text{obs}(O) * \text{signal}(s, \text{True})
\end{array}
\qquad
\begin{array}{c}
\text{VS-WAITPERM} \\
\frac{\delta' <_{\Delta} \delta}{\text{itperm}(\delta) \Rightarrow \text{wperm}(id, \delta')}
\end{array}$$

$$\begin{array}{c}
\text{VS-WAIT} \\
\frac{\forall o \in O. \text{lev}(s) <_{\text{L}} \text{lev}(o)}{\text{obs}(O) * \text{wperm}(s.\text{id}, \delta) * \text{signal}(s, b) \Rightarrow \text{obs}(O) * \text{wperm}(s.\text{id}, \delta) * \text{signal}(s, b) * (\neg b \leftrightarrow \text{itperm}(\tau, \delta))}
\end{array}$$

$$\begin{array}{c}
\text{VS-WEAKPERM} \\
\frac{\delta' <_{\Delta} \delta \quad N \in \mathbb{N}}{\text{itperm}(\delta) \Rightarrow \bigstar_{1, \dots, N} \text{itperm}(\delta')}
\end{array}
\qquad
\begin{array}{c}
\text{VS-MUTINIT} \\
\frac{\neg \text{mentionsPerms}(P)}{\text{unit}(\ell) * P \Rightarrow \text{mutex}((\ell, L), P)}
\end{array}$$

Fig. 8: View shift rules.

Definition 18 (Proof Relation). We define a proof relation $\vdash \subset \mathcal{A} \times \text{Cmds} \times (\text{Values} \rightarrow \mathcal{A})$ according to the rules presented in Fig. 9 and 10. We state the provability of a Hoare triple in the form of $\vdash \{A\} c \{\lambda r. B(r)\}$ where r captures the value returned by c . To simplify the notation, we omit the result value if it is clear from the context or irrelevant.

D Soundness

D.1 Annotated Semantics

Definition 19 (Intermediate Representation). We define an extended set of commands Cmds^+ according to the syntax presented in Fig. 11.

$$\begin{array}{c}
\text{PR-FRAME} \\
\frac{\vdash \{A\} c \{B\}}{\vdash \{A * F\} c \{B * F\}} \\
\\
\text{PR-EXP} \\
\frac{\llbracket e \rrbracket \in \text{Values}}{\vdash \{\text{True}\} e \{\lambda r. r = [e]\}} \\
\\
\text{PR-FORK} \\
\frac{\vdash \{\text{obs}(O_f) * A\} c \{\text{obs}(\emptyset)\}}{\vdash \{\text{obs}(O_m \uplus O_f) * A\} \text{fork } c \{\lambda r. \text{obs}(O_m) * r = \text{tt}\}}
\end{array}
\qquad
\begin{array}{c}
\text{PR-VIEWSHIFT} \\
\frac{A \Rightarrow A' \quad \vdash \{A'\} c \{B'\} \quad B' \Rightarrow B}{\vdash \{A\} c \{B\}} \\
\\
\text{PR-EXISTS} \\
\frac{\forall a \in A. \vdash \{a\} c \{B\}}{\vdash \{\bigvee A\} c \{B\}}
\end{array}$$

(a) Basic Proof Rules.

$$\begin{array}{c}
\text{PR-IF} \\
\frac{\vdash \{A\} c_b \{\lambda b. C(b) \wedge (b = \text{True} \vee b = \text{False})\} \quad \vdash \{C(\text{True})\} c_t \{B\} \quad C(\text{False}) \Rightarrow B}{\vdash \{A\} \text{if } c_b \text{ then } c_t \{B\}} \\
\\
\text{PR-WHILE} \\
\frac{\vdash \{I\} c_b \left\{ \begin{array}{l} \lambda b. (b = \text{True} \vee b = \text{False}) \\ * (b \rightarrow \text{itperm}(\delta) * I) \\ * (\neg b \rightarrow B) \end{array} \right\}}{\vdash \{I\} \text{while } c_b \text{ do skip } \{B\}} \\
\\
\text{PR-LET} \\
\frac{\vdash \{A\} c \{\lambda r. C(r)\} \quad \forall v. \vdash \{C(v)\} c'[v/x] \{B\}}{\vdash \{A\} \text{let } x := c \text{ in } c' \{B\}}
\end{array}$$

(b) Control Structures.

Fig. 9: Proof rules (part 1).

$$\begin{array}{c}
\text{PR-ACQUIRE} \\
\frac{\forall o \in O. \text{lev}(m) <_{\perp} \text{lev}(o)}{\{\text{obs}(O) * [f]\text{mutex}(m, P)\}} \\
\vdash \text{acquire } m.\text{loc} \\
\{\lambda r. r = \text{tt} * \text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * P\} \\
\\
\text{PR-RELEASE} \\
\frac{\text{obs}(O) * A \Rightarrow \text{obs}(O) * P * B}{\{\text{obs}(O \uplus \{m\}) * \text{locked}(m, P, f) * A\}} \\
\vdash \text{release } m.\text{loc} \\
\{\lambda r. r = \text{tt} * \text{obs}(O) * [f]\text{mutex}(m, P) * B\} \\
\\
\text{PR-NEWMUTEX} \\
\vdash \{\text{True}\} \text{new_mutex } \{\lambda \ell. \text{uninit}(\ell)\}
\end{array}$$

(a) Mutexes.

$$\begin{array}{c}
\text{PR-CONS} \qquad \qquad \qquad \text{PR-READHEAPLOC} \\
\vdash \{\text{True}\} \text{cons}(v) \{\lambda \ell. \ell \mapsto v\} \qquad \vdash \{[f]\ell \mapsto v\} [\ell] \{\lambda r. r = v * [f]\ell \mapsto v\} \\
\\
\text{PR-ASSIGNTOHEAP} \\
\vdash \{\ell \mapsto _ \} [\ell] := v \{\lambda r. r = \text{tt} * \ell \mapsto v\}
\end{array}$$

(b) Heap Access.

Fig. 10: Proof rules (part 2).

$$c \in \text{Cmds}^+ ::= e \mid \text{while } c \text{ do skip} \mid \text{fork } c \mid \text{let } x := c \text{ in } c \mid \text{if } c \text{ then } c \\
\mid \text{cons}(e) \mid [e] \mid [e] := e \mid \text{new_mutex} \mid \text{acquire } e \\
\mid \text{release } e \\
\mid \text{consumeItPerm}$$

Fig. 11: Extended set of commands for intermediate representation.

For the rest of the appendix, commands c refer to the extended set of commands, i.e., $c \in \text{Cmds}^+$.

Definition 20 (Annotated Resources). We define the set of annotated resources AnnoRes as

$$r^a \in \text{AnnoRes} ::= \ell \mapsto v \mid \text{uinit}_{\text{aRes}}(\ell) \mid \\ \text{unlocked}_{\text{aRes}}((\ell, L), a, H) \mid \text{locked}_{\text{aRes}}((\ell, L), a, f) \mid \\ \text{signal}_{\text{aRes}}(\text{id}, L), b)$$

where H does not contain any obligations chunks.

Definition 21 (Annotated Heaps). We define the set of annotated heaps as

$$\text{Heaps}^{\text{annot}} := \mathcal{P}_{\text{fin}}(\text{AnnoRes}),$$

the function $\text{locs}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{Locs})$ mapping annotated heaps to the sets of allocated heap locations as

$$\text{locs}_{\text{ah}}(h^a) := \{ \ell \in \mathcal{Locs} \mid \exists v \in \text{Values}. \exists L \in \mathcal{Levs}. \exists a \in \mathcal{A}. \\ \exists H \in \text{Heaps}^{\text{log}}. \exists f \in \mathcal{F}. \\ \ell \mapsto v \in h^a \vee \text{uinit}_{\text{aRes}}(\ell) \in h^a \vee \\ \text{unlocked}_{\text{aRes}}((\ell, L), a, H) \in h^a \vee \\ \text{locked}_{\text{aRes}}((\ell, L), a, f) \in h^a \}$$

and the function $\text{ids}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{ID})$ mapping annotated heaps to sets of allocated signal IDs as

$$\text{ids}_{\text{ah}}(h^a) := \{ \text{id} \in \mathcal{ID} \mid \exists L \in \mathcal{Levs}. \exists b \in \mathbb{B}. \text{signal}_{\text{aRes}}(\text{id}, L), b \in h^a \}.$$

We denote annotated heaps by h^a .

We call an annotated heap h^a *finite* and write $\text{finite}_{\text{ah}}(h^a)$ if there exists no chunk $\text{unlocked}_{\text{aRes}}((\ell, L), a, H) \in h^a$ for which $\text{finite}_{\text{lh}}(H)$ does not hold.

Definition 22 (Compatibility of Annotated and Physical Heaps). We inductively define a relation $\text{ah} \sim_{\text{ph}} \subset \text{Heaps}^{\text{annot}} \times \mathcal{R}^{\text{phys}}$ between annotated and physical heaps such that the following holds:

$$\begin{array}{ll} \emptyset & \text{ah} \sim_{\text{ph}} \emptyset, \\ \ell \mapsto v \cup h^a & \text{ah} \sim_{\text{ph}} \ell \mapsto v \cup h, \\ \text{uinit}_{\text{aRes}}(\ell) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{unlocked}_{\text{aRes}}((\ell, L), P, H_P) \cup h^a & \text{ah} \sim_{\text{ph}} \text{unlocked}_{\text{pRes}}(\ell) \cup h, \\ \text{locked}_{\text{aRes}}((\ell, L), P, f) \cup h^a & \text{ah} \sim_{\text{ph}} \text{locked}_{\text{pRes}}(\ell) \cup h, \\ \text{signal}_{\text{aRes}}(s, b) \cup h^a & \text{ah} \sim_{\text{ph}} h, \end{array}$$

where $h^a \in \text{Heaps}^{\text{annot}}$ and $h \in \text{Heaps}^{\text{phys}}$ are annotated and physical heaps with $h^a \sim_{\text{ph}} h$.

Definition 23 (Compatibility of Annotated and Logical Heaps). We inductively define a relation $\text{ah} \sim_{\text{lh}} \subset \text{Heaps}^{\text{annot}} \times \text{Heaps}^{\text{log}}$ between annotated and logical heaps such that the following holds:

$$\begin{array}{ll}
\emptyset & \text{ah} \sim_{\text{lh}} \emptyset_{\text{log}}, \\
h^a \cup \{\ell \mapsto v\} & \text{ah} \sim_{\text{lh}} H + \{\ell \mapsto v\}, \\
h^a \cup \{\text{uninit}_{\text{aRes}}(\ell)\} & \text{ah} \sim_{\text{lh}} H + \{\text{uninit}_{\text{lRes}}(\ell)\}, \\
h^a \cup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\} & \text{ah} \sim_{\text{lh}} H + \{\text{mutex}_{\text{lRes}}(m, P)\} + H_P, \\
h^a \cup \{\text{locked}_{\text{aRes}}(m, P, f)\} & \text{ah} \sim_{\text{lh}} H + \{\text{locked}_{\text{lRes}}(m, P, f)\} \\
& \quad + (1 - f) \cdot \{\text{mutex}_{\text{lRes}}(m, P)\}, \\
h^a \cup \{\text{signal}_{\text{aRes}}(s, b)\} & \text{ah} \sim_{\text{lh}} H + \{\text{signal}_{\text{lRes}}(s, b)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{obs}_{\text{lRes}}(O)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{wperm}_{\text{lRes}}(\text{id}, \delta)\}, \\
h^a & \text{ah} \sim_{\text{lh}} H + \{\text{itperm}_{\text{lRes}}(\delta)\},
\end{array}$$

where $h^a \in \text{Heaps}^{\text{annot}}$ and $H \in \text{Heaps}^{\text{log}}$ are annotated and logical heaps with $\ell, m.\text{loc} \notin \text{locs}_{\text{ah}}(h^a)$, $s.\text{id} \notin \text{ids}_{\text{ah}}(h^a)$ and $h^a \sim_{\text{ah}} H$.

Definition 24 (Annotated Single Thread Reduction Relation). We define a reduction relation $\rightsquigarrow_{\text{ast}}$ for annotated threads according to the rules presented in Fig. 12. A reduction step has the form

$$h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T^a$$

for a set of annotated forked threads $T^a \subset \text{Heaps}^{\text{log}} \times \text{Cmds}$ with $|T^a| \leq 1$.

It indicates that given annotated heap h^a and a logical heap H , command c can be reduced to annotated heap $h^{a'}$, logical heap H' and command c' . The either empty or singleton set T^a represents whether a new thread is forked in this step.

For simplicity of notation we omit T^a if it is clear from the context that no thread is forked and $T^a = \emptyset$.

Definition 25 (Annotated Thread Pools). We define the set of annotated thread pools \mathcal{TP}^a as the set of finite partial functions mapping thread IDs to annotated threads:

$$\mathcal{TP}^a := \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds} \cup \{\text{term}\}).$$

We denote annotated thread pools by P^a and the empty thread pool by \emptyset_{atp} , i.e.,

$$\begin{aligned}
\emptyset_{\text{atp}} &: \Theta \rightarrow_{\text{fin}} \text{Heaps}^{\text{log}} \times (\text{Cmds} \cup \{\text{term}\}), \\
\text{dom}(\emptyset_{\text{atp}}) &= \emptyset.
\end{aligned}$$

We define the extension operation $+_{\text{atp}}$ analogously to $+_{\text{tp}}$, cf. Definition 4d.

For convenience of notation we define selector functions for annotated threads as

$$\begin{aligned}
(H, c).\text{heap} &:= H, \\
(H, c).\text{cmd} &:= c.
\end{aligned}$$

$$\frac{\text{AST-RED-EVALCTXT} \quad h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T}{h^a, H, E[c] \rightsquigarrow_{\text{ast}} h^{a'}, H', E[c'], T} \quad \frac{\text{AST-RED-FORK} \quad h^a, H_m + \{\text{obs}_{\text{IRes}}(O_m \uplus O_f)\} + H_f, \mathbf{fork} \ c \rightsquigarrow_{\text{ast}}}{h^a, H_m + \{\text{obs}_{\text{IRes}}(O_m)\}, \mathbf{tt}, \{(\{\text{obs}_{\text{IRes}}(O_f)\} + H_f), c\}}$$

(a) Basic constructs.

$$\text{AST-RED-WHILE} \quad h^a, H, \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip} \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{if} \ c \ \mathbf{then} \ (\mathbf{consumeItPerm}; \mathbf{while} \ c \ \mathbf{do} \ \mathbf{skip})$$

$$\frac{\text{AST-RED-IFTRUE} \quad h^a, H, \mathbf{if} \ \mathbf{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, c}{h^a, H, \mathbf{if} \ \mathbf{True} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, c} \quad \frac{\text{AST-RED-IFFALSE} \quad h^a, H, \mathbf{if} \ \mathbf{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt}}{h^a, H, \mathbf{if} \ \mathbf{False} \ \mathbf{then} \ c \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt}}$$

$$\frac{\text{AST-RED-CONSUMEITPERM} \quad h^a, H + \{\text{itperm}_{\text{IRes}}(\delta)\}, \mathbf{consumeItPerm} \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt}}{h^a, H + \{\text{itperm}_{\text{IRes}}(\delta)\}, \mathbf{consumeItPerm} \rightsquigarrow_{\text{ast}} h^a, H, \mathbf{tt}}$$

$$\frac{\text{AST-RED-LET} \quad h^a, H, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{ast}} h^a, H, c[v/x]}{h^a, H, \mathbf{let} \ x := v \ \mathbf{in} \ c \rightsquigarrow_{\text{ast}} h^a, H, c[v/x]}$$

(b) Control structures.

$$\frac{\text{AST-RED-CONS} \quad \ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \mathbf{cons}(v) \rightsquigarrow_{\text{ast}} h^a \cup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, \ell} \quad \frac{\text{AST-RED-READHEAPLOC} \quad \ell \mapsto v \in h^a}{h^a, H, [\ell] \rightsquigarrow_{\text{ast}} h^a, H, v}$$

$$\frac{\text{AST-RED-ASSIGN} \quad h \sqcup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, [\ell] := v \rightsquigarrow_{\text{ast}} h \sqcup \{\ell \mapsto v'\}, H + \{\ell \mapsto v'\}, \mathbf{tt}}{h \sqcup \{\ell \mapsto v\}, H + \{\ell \mapsto v\}, [\ell] := v \rightsquigarrow_{\text{ast}} h \sqcup \{\ell \mapsto v'\}, H + \{\ell \mapsto v'\}, \mathbf{tt}}$$

(c) Heap access.

$$\frac{\text{AST-RED-NEWMUTEX} \quad \ell \notin \text{locs}_{\text{ah}}(h^a)}{h^a, H, \mathbf{new_mutex} \rightsquigarrow_{\text{ast}} h^a \cup \{\text{uninit}_{\text{aRes}}(\ell)\}, H + \{\text{uninit}_{\text{IRes}}(\ell)\}, \ell}$$

$$\frac{\text{AST-RED-ACQUIRE} \quad f \in \mathcal{F} \quad \forall o \in O. \text{lev}(m) <_{\text{L}} \text{lev}(o)}{h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, a, H_P)\}, H + \{\text{obs}_{\text{IRes}}(O)\} + f \cdot \{\text{mutex}_{\text{IRes}}(m, P)\}, \mathbf{acquire} \ m.\text{loc} \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, H + \{\text{obs}_{\text{IRes}}(O \uplus \{m\})\}, \text{locked}_{\text{IRes}}(m, P, f) + H_P, \mathbf{tt}}$$

$$\frac{\text{AST-RED-RELEASE} \quad H_P \models_{\text{A}} P \quad \text{consistent}_{\text{lh}}(H_P) \quad \exists O. H(\text{obs}_{\text{IRes}}(O)) \geq 1}{\neg \exists \delta. \text{id.}(H_P(\text{itperm}_{\text{IRes}}(\delta)) > 0 \vee H_P(\text{wperm}_{\text{IRes}}(\text{id}, \delta)) > 0)}{h^a \sqcup \{\text{locked}_{\text{aRes}}(m, P, f)\}, H + \{\text{obs}_{\text{IRes}}(O \uplus \{m\})\}, \text{locked}_{\text{IRes}}(m, P, f) + H_P, \mathbf{release} \ m.\text{loc} \rightsquigarrow_{\text{ast}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}(m, P, H_P)\}, H + \{\text{obs}_{\text{IRes}}(O)\} + f \cdot \{\text{mutex}_{\text{IRes}}(m, P)\}, \mathbf{tt}}$$

(d) Mutexes.

Fig. 12: Annotated single thread reduction rules.

Definition 26 (Ghost Reduction Relation). We define a thread pool reduction relation $\rightsquigarrow_{\text{ghost}}$ according to the rules presented in Fig. 13 to express ghost steps. A ghost reduction step has the form

$$h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^{a'}, P^{a'}.$$

We denote its reflexive transitive closure by $\rightsquigarrow_{\text{ghost}}^*$.

$$\begin{array}{c} \text{GTP-RED-NEWSIGNAL} \\ \frac{P^a(\theta) = (H + \{\text{obs}_{\text{IRes}}(O)\}, c) \quad id \notin \text{idS}_{\text{ah}}(h^a) \quad H' = H + \{\text{signal}_{\text{IRes}}((id, L), \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{id, L\}\})\}}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}((id, L), \text{False})\}, P^a[\theta := (H', c)]} \\ \\ \text{GTP-RED-SET SIGNAL} \\ \frac{P^a(\theta) = (H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O \uplus \{\{s\}\})\}, c) \quad H' = H + \{\text{signal}_{\text{IRes}}(s, \text{False}), \text{obs}_{\text{IRes}}(O)\}}{h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{False})\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{signal}_{\text{aRes}}(s, \text{True})\}, P^a[\theta := (H', c)]} \\ \\ \text{GTP-RED-WAIT PERM} \\ \frac{\delta' <_{\Delta} \delta \quad P^a(\theta) = (H + \{\text{itperm}_{\text{IRes}}(\delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{wperm}_{\text{IRes}}(id, \delta')\}, c)]} \\ \\ \text{GTP-RED-WAIT} \\ \frac{\text{signal}_{\text{aRes}}(s, \text{False}) \in h^a \quad P^a(\theta) = (H, c) \quad H(\text{wperm}_{\text{IRes}}(s.\text{id}, \delta)) \geq 1 \quad H(\text{obs}_{\text{IRes}}(O)) \geq 1 \quad \forall o \in O. \text{lev}(s) <_{\text{L}} \text{lev}(o)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + \{\text{itperm}_{\text{IRes}}(\delta)\}, c)]} \\ \\ \text{GTP-RED-WEAKITPERM} \\ \frac{\delta' <_{\Delta} \delta \quad N \in \mathbb{N} \quad P^a(\theta) = (H + \{\text{itperm}_{\text{IRes}}(\delta)\}, c)}{h^a, P^a \xrightarrow{\theta}_{\text{ghost}} h^a, P^a[\theta := (H + N \cdot \{\text{itperm}_{\text{IRes}}(\delta')\}, c)]} \\ \\ \text{GTP-RED-MUTINIT} \\ \frac{P^a(\theta) = (H + \{\text{uninit}_{\text{IRes}}(\ell)\} + H_P, c) \quad H' = H + \{\text{mutex}_{\text{IRes}}((\ell, L), H_P)\} \quad H_P \models_{\text{A}} P \quad \text{consistent}_{\text{lh}}(H_P) \quad \exists O. H(\text{obs}_{\text{IRes}}(O)) \geq 1 \quad \neg \exists \delta, id. (H_P(\text{itperm}_{\text{IRes}}(\delta)) > 0 \vee H_P(\text{wperm}_{\text{IRes}}(id, \delta)) > 0)}{h^a \sqcup \{\text{uninit}_{\text{aRes}}(\ell)\}, P^a \xrightarrow{\theta}_{\text{ghost}} h^a \sqcup \{\text{unlocked}_{\text{aRes}}((\ell, L), a, H_P)\}, P^a[\theta := (H', c)]} \end{array}$$

Fig. 13: Ghost thread pool reduction rules.

Definition 27 (Non-ghost Thread Pool Reduction Relation). We define a thread pool reduction relation $\rightsquigarrow_{\text{real}}$ according to the rules presented in Fig. 14 to express real (i.e. non-ghost) reduction steps. A reduction step has the form

$$h^a, P^a \xrightarrow{\theta}_{\text{real}} h^{a'}, P^{a'}.$$

$$\begin{array}{c}
\text{RTP-RED-LIFT} \\
\frac{\theta_f = \min(\Theta \setminus \text{dom}(P^a)) \quad P^a(\theta) = (H, c) \quad h^a, H, c \rightsquigarrow_{\text{ast}} h^{a'}, H', c', T^a}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^{a'}, P^a[\theta := (H', c')] +_{\text{atp}} T^a} \\
\\
\text{RTP-RED-TERM} \\
\frac{P^a(\theta) = (H, v) \quad H.\text{obs} = \emptyset}{h^a, P^a \xrightarrow{\theta}_{\text{real}} h^a, P^a -_{\text{atp}} \theta}
\end{array}$$

Fig. 14: Non-ghost thread pool reduction rules.

Definition 28 (Annotated Thread Pool Reduction Relation). We define the annotated thread pool reduction relation $\rightsquigarrow_{\text{atp}}$ as

$$\rightsquigarrow_{\text{atp}} := \rightsquigarrow_{\text{ghost}} \cup \rightsquigarrow_{\text{real}}.$$

Definition 29 (Annotated Reduction Sequence). Let $(h^a_i)_{i \in \mathbb{N}}$ and $(P^a_i)_{i \in \mathbb{N}}$ be infinite sequences of annotated heaps and annotated thread pools, respectively. Let $\text{sig} : \mathbb{N} \rightarrow \mathcal{S}$ be a partial function mapping indices to signals.

We call $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$ an annotated reduction sequence if there exists a sequence of thread IDs $(\theta_i)_{i \in \mathbb{N}}$ such that the following holds for every $i \in \mathbb{N}$:

- $h^a_i, P^a_i \xrightarrow{\theta_i}_{\text{atp}} h^a_{i+1}, P^a_{i+1}$
- If this reduction step results from an application of GTP-RED-WAIT to some signal s , then $\text{sig}(i) = s$ holds and otherwise $\text{sig}(i) = \perp$.

In case the signal annotation sig is clear from the context or not relevant, we omit it and write $(h^a_i, P^a_i)_{i \in \mathbb{N}}$ instead of $((h^a_i, P^a_i)_{i \in \mathbb{N}}, \text{sig})$.

We call (h^a_i, P^a_i) an annotated machine configuration.

Lemma 6 (Preservation of Finiteness). Let $(h^a_i, P^a_i)_{i \in \mathbb{N}}$ be an annotated reduction sequence with $\text{finite}_{\text{ah}}(h^a_0)$ and $\text{finite}_{\text{lh}}(P^a_0(\theta).\text{heap})$ for all $\theta \in \text{dom}(P^a_0)$.

Then, $\text{finite}_{\text{lh}}(P^a_i(\theta).\text{heap})$ holds for all $i \in \mathbb{N}$ and all $\theta \in \text{dom}(P^a_i)$.

Proof. Proof by induction on i .

Lemma 7 (Preservation of Completeness). Let $(h^a_i, P^a_i)_{i \in \mathbb{N}}$ be an annotated reduction sequence with $\text{complete}_{\text{lh}}(P^a_0(\theta).\text{heap})$ for all $\theta \in \text{dom}(P^a_0)$. Then, $\text{complete}_{\text{lh}}(P^a_i(\theta).\text{heap})$ holds for every $i \in \mathbb{N}$ and every $\theta \in \text{dom}(P^a_i)$.

Proof. Proof by induction on i .

Every thread of an annotated thread pool is annotated by a thread-local logical heap that expresses which resources are owned by this thread. In the following we define a function to extract the logical heap expressing which resources are owned by threads of a thread pool (i.e. the sum of all thread-local logical heaps).

Definition 30. We define the function $\text{ownedResHeap}_{\text{atp}} : \mathcal{TP}^a \rightarrow \text{Heaps}^{\text{log}}$ mapping annotated thread pools to logical heaps as follows:

$$P^a \mapsto \sum_{\theta \in \text{dom}(P^a)} P^a(\theta).\text{heap}$$

Annotated resources representing unlocked locks, i.e., $\text{unlocked}_{\text{aRes}}(m, a, H_a)$, contain a logical heap H_a that expresses which resources are protected by this lock. In the following, we define a function that extracts a logical heap from an annotated heap h^a expressing which resources are protected by unlocked locks in h^a .

Definition 31. We define the function $\text{protectedResHeap}_{\text{ah}} : \text{Heaps}^{\text{annot}} \rightarrow \text{Heaps}^{\text{log}}$ mapping annotated heaps to logical heaps as follows:

For any annotated heap h^a let

$$\text{LockInvs}(h^a) := \{ \{ H_P \in \text{Heaps}^{\text{log}} \mid \exists m \in \mathcal{Locs} \times \mathcal{Levs}. \exists P \in \mathcal{A}. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \} \}$$

be the auxiliary set aggregating all logical heaps corresponding to lock invariants of unlocked locks stored in h^a . We define $\text{protectedResHeap}_{\text{ah}}$ as

$$h^a \mapsto \sum_{H_P \in \text{LockInvs}(h^a)} H_P.$$

We consider a machine configuration (h^a, P^a) to be *consistent* if it fulfils three criteria: (i) Every thread-local logical heap is consistent, i.e., for all used thread IDs θ , $P^a(\theta).\text{heap}$ only stores full obligations, wait permission and iteration permission chunks. (ii) Every logical heap protected by an unlocked lock in h^a is consistent. (iii) h^a is compatible with the logical heap that represents (a) the resources owned by threads in P^a and (b) the resources protected by unlocked locks stored in h^a .

Definition 32 (Consistency of Annotated Machine Configurations). We call an annotated machine configuration (h^a, P^a) consistent and write $\text{consistent}_{\text{conf}}(h^a, P^a)$ if all of the following hold:

- $\text{consistent}_{\text{lh}}(P^a(\theta).\text{heap})$ for all $\theta \in \text{dom}(P^a)$,
- $\forall m. \forall P. \forall H_P. \text{unlocked}_{\text{aRes}}(m, P, H_P) \in h^a \rightarrow \text{consistent}_{\text{lh}}(H_P)$,
- $h^a \text{ ah} \sim_{\text{lh}} \text{ownedResHeap}_{\text{atp}}(P^a) + \text{protectedResHeap}_{\text{ah}}(h^a)$.

Lemma 8 (Preservation of Consistency). Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be an annotated reduction sequence with $\text{consistent}_{\text{conf}}(h_0^a, P_0^a)$. Then, $\text{consistent}_{\text{conf}}(h_i^a, P_i^a)$ holds for every $i \in \mathbb{N}$.

Proof. Proof by induction on i .

Definition 33 (Command Annotation). We define the predicate $\text{annot}_{\text{cmd}} \subset \text{Cmds} \times \text{Cmds}$ such that $\text{annot}_{\text{cmd}}(c', c)$ holds iff c' results from c by removing all occurrences of **consumeItPerm**.

Definition 34 (Thread Pool Annotation). We define a predicate $\text{annot}_{\text{tp}} \subset \mathcal{TP}^a \times \mathcal{TP}$ such that:

$$\begin{aligned} & \text{annot}_{\text{tp}}(P^a, P) \\ & \iff \\ & \text{dom}(P^a) = \text{dom}(P) \wedge \forall \theta \in \text{dom}(P). \text{annot}_{\text{cmd}}(P^a(\theta).\text{cmd}, P(\theta)) \end{aligned}$$

D.2 Soundness Proof

In this subsection we present detailed versions of proof sketches presented in § 5.

Definition 35 (Program Order Graph).

Let $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$ be an annotated reduction sequence. Let N^r be the set of names referring to reduction rules defining the relations $\rightsquigarrow_{\text{real}}$, $\rightsquigarrow_{\text{ghost}}$ and $\rightsquigarrow_{\text{ast}}$. We define the set of annotated reduction rule names N^a where GTP-RED-WAIT is annotated by signals as

$$\begin{aligned} N^a & := (N^r \setminus \{\text{GTP-RED-WAIT}\}) \\ & \cup (\{\text{GTP-RED-WAIT}\} \times \mathcal{S}). \end{aligned}$$

We define the program order graph $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})) = (\mathbb{N}, E)$ with root 0 where $E \subset \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$.

A node $a \in \mathbb{N}$ corresponds to the sequence's a^{th} reduction step, i.e., to the step $h_a^a, P_a^a \rightsquigarrow_{\text{atp}}^{\theta} h_{a+1}^a, P_{a+1}^a$ for some $\theta \in \text{dom}(P_a^a)$. An edge from node a to node b expresses that the b^{th} reduction step continues the control flow of step a . For any $\ell \in \mathbb{N}$, let θ_ℓ denote the ID of the thread reduced in step ℓ . Furthermore, let n_ℓ^a denote the name of the reduction rule applied in the ℓ^{th} step, in the following sense:

- If $h_\ell^a, P_\ell^a \rightsquigarrow_{\text{atp}}^{\theta} h_{\ell+1}^a, P_{\ell+1}^a$ results from an application of RTP-RED-LIFT in combination with single-thread reduction rule n^{st} , then $n_\ell^a = n^{\text{st}}$.
- If $h_\ell^a, P_\ell^a \rightsquigarrow_{\text{atp}}^{\theta} h_{\ell+1}^a, P_{\ell+1}^a$ results from an application of GTP-RED-WAIT, then $n_\ell^a = (\text{GTP-RED-WAIT}, \text{sig}(\ell))$.
- Otherwise, n^a denotes the applied (real or ghost) thread pool reduction rule.

An edge $(a, \theta, n^a, b) \in \mathbb{N} \times \Theta \times N^a \times \mathbb{N}$ is contained in E if $n^a = n_a^a$ and one of the following conditions applies:

- $\theta = \theta_a = \theta_b$ and $b = \min(\{k > a \mid h_k^a, P_k^a \rightsquigarrow_{\text{atp}}^{\theta_a} h_{k+1}^a, P_{k+1}^a\})$.
In this case, the edge expresses that step b marks the first time that thread θ_a is rescheduled for reduction (after step a).
- $\text{dom}(P_{a+1}^a) \setminus \text{dom}(P_a^a) = \{\theta\}$ and
 $b = \min\{k \in \mathbb{N} \mid h_k^a, P_k^a \rightsquigarrow_{\text{atp}}^{\theta} h_{k+1}^a, P_{k+1}^a\}$.
In this case, θ identifies the thread forked in step a . The edge expresses that step b marks the first reduction of the forked thread.

In case the choice of reduction sequence $((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig})$ is clear from the context, we write G instead of $G(((h_i^a, P_i^a)_{i \in \mathbb{N}}, \text{sig}))$.

Observation 1. *Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be an annotated reduction sequence with $|\text{dom}(P_0^a)| = 1$. The sequence's program order graph $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ is a binary tree.*

For any reduction sequence $(h_i^a, P_i^a)_{i \in \mathbb{N}}$, the paths in its program order graph $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ represent the sequence's control flow paths. Hence, we are going to use program order graphs to analyse reduction sequences' control flows.

We refer to a program order graph's edges by the kind of reduction step they represent. For instance, we call edges of the form $(a, \theta, \text{ST-RED-WHILE}, b)$ *loop edges* because they represent a loop backjump and we call edges of the form $(a, \theta, (\text{GTP-RED-WAIT}, s), b)$ *wait edges*. Any wait edge of this form represents an application of GTP-RED-WAIT to signal s .

In the following, we prove that any path in a program order graph that does not involve a loop edge is finite. This follows from the fact that the size of the command reduced along this path decreases with each non-ghost non-loop step.

Lemma 9. *Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be a fair annotated reduction sequence. Let $p = (V, E)$ be a path in $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$. Let $L = \{e \in E \mid \pi_3(e) = \text{AST-RED-WHILE}\}$ be the set of loop edges contained in p . Then, p is infinite if and only if L is infinite.*

Proof. If L is infinite, p is obviously infinite as well. So, suppose L is finite.

For any command, we consider its size to be the number of nodes contained in its abstract syntax tree. By structural induction over the set of commands, it follows that the size of a command $c = P^a(\theta).\text{cmd}$ decreases in every non-ghost reduction step $h^a, P^a \xrightarrow{\theta}_{\text{atp}} h^{a'}, P^{a'}$ that is not an application of RTP-RED-LIFT in combination with AST-RED-WHILE.

Since L is finite, there exists a node x such that the suffix $p_{\geq x}$ starting at node x does not contain any loop edges. By fairness of $(h_i^a, P_i^a)_{i \in \mathbb{N}}$, every non-empty suffix of $p_{\geq x}$ contains an edge corresponding to a non-ghost reduction step. For any edge $e = (i, \theta, n, j)$ consider the command $c_e = P_i^a(\theta).\text{cmd}$ reduced in this edge. The size of these commands decreases along $p_{\geq x}$. So, $p_{\geq x}$ must be finite and thus p must be finite as well.

Corollary 1. *Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be a fair annotated reduction sequence. Let $p = (V, E)$ be a path in $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$. Let*

$$C = \{e \in E \mid \pi_3(e) = \text{AST-RED-CONSUMEITPERM}\}$$

be the set of consume edges contained in p . Then, p is infinite if and only if C is infinite.

Proof. Follows from Lemma 9 by the fact that the set $\{e \in E \mid \pi_3(e) = \text{AST-RED-WHILE}\}$ is infinite if and only if C is infinite.

Definition 36. *Let $G = (V, E)$ be a subgraph of some program order graph. We define the function $\text{waitEdges}_G : \mathcal{S} \rightarrow \mathcal{P}(E)$ mapping any signal s to the set of wait edges in G concerning s as:*

$$\text{waitEdges}_G(s) := \{(a, \theta, (\text{GTP-RED-WAIT}, s'), b) \in E \mid s' = s\}.$$

Furthermore, we define the set $\mathcal{S}_G \subset \mathcal{S}$ of signals being waited for in G and its subset $\mathcal{S}_G^\infty \subseteq \mathcal{S}_G$ of signals waited-for infinitely often in G as follows:

$$\begin{aligned}\mathcal{S}_G &:= \{s \in \mathcal{S} \mid \text{waitEdges}_G(s) \neq \emptyset\}, \\ \mathcal{S}_G^\infty &:= \{s^\infty \in \mathcal{S}_G \mid \text{waitEdges}_G(s^\infty) \text{ infinite}\}.\end{aligned}$$

Definition 37. Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be a fair annotated reduction sequence and let $G = (V, E)$ be a subgraph of the sequence's program order graph. We define the function $\text{itperms}_G : E \rightarrow \text{Bags}_{\text{fin}}(A)$ mapping any edge e to the (potentially empty) finite bag of iteration permissions derived in the reduction step corresponding to e as follows:

Let $(i, \theta, n, j) \in E$ be an edge.

- If $n = (\text{GTP-RED-WAIT}, s)$ for some signal $s \in \mathcal{S}$, then the i^{th} reduction step spawns a single iteration permission of degree δ , i.e., $P_{i+1}^a = P_i^a[\theta := (P_i^a(\theta).\text{heap} + \{\text{itperm}_{\text{IRes}}(\delta)\}, P_i^a(\theta).\text{cmd})]$. In this case, we define

$$\text{itperms}_G((i, \theta, (\text{GTP-RED-WAIT}, s), j)) := \{\{\delta\}\}.$$

- If $n = \text{GTP-RED-WEAKITPERM}$, then the i^{th} reduction step consumes an iteration permission of degree δ and produces N permissions of a lower degree δ' , i.e., $P_i^a(\theta).\text{heap} = H + \{\text{itperm}(\delta)\}$ for some heap H and $P_{i+1}^a = P_i^a[\theta := (H', P_i^a(\theta).\text{cmd})]$ for

$$H' = H + N \cdot \{\text{itperm}_{\text{IRes}}(\delta')\}.$$

In this case, we define

$$\text{itperms}_G((i, \theta, \text{GTP-RED-WEAKITPERM}, j)) := \underbrace{\{\{\delta', \dots, \delta'\}\}}_{N \text{ times}}.$$

- Otherwise, we define

$$\text{itperms}_G((i, \theta, n, j)) := \emptyset.$$

Definition 38 (Partial Order on Finite Bags).

Let X be a set and $<_X \subset X \times X$ a partial order on X . We define the partial order $\prec_X \subset \text{Bags}_{\text{fin}}(X) \times \text{Bags}_{\text{fin}}(X)$ on finite bags over X as the Dershowitz-Manna ordering [6] induced by $<_X$:

$$\begin{aligned}A \prec_X B &\iff \exists C, D \in \text{Bags}_{\text{fin}}(X). \emptyset \neq C \subseteq B \\ &\quad \wedge A = (B \setminus C) \uplus D \\ &\quad \wedge \forall d \in D. \exists c \in C. d <_X c.\end{aligned}$$

We define $\preceq_X \subset \text{Bags}_{\text{fin}}(X) \times \text{Bags}_{\text{fin}}(X)$ such that

$$A \preceq_X B \iff A = B \vee A \prec_X B$$

holds.

Corollary 2. *The partial order $\prec_\Lambda \subset \text{Bags}_{\text{fin}}(\Lambda) \times \text{Bags}_{\text{fin}}(\Lambda)$ is well-founded.*

Proof. Follows from [6] and well-foundedness of \prec_Δ .

We view paths in a program order graph as single-branched subgraphs. This allows us to apply above definitions on graphs to paths. In particular, this allows us to refer to the capacity of a signal s on a path p by referring to sigCap_p .

For the following definition, remember that a bag $B \in \text{Bags}(X)$ is a function $B : X \rightarrow \mathbb{N}$ while a logical heap $H \in \text{Heaps}^{\text{log}}$ is a function $H : \mathcal{R}^{\text{log}} \rightarrow \mathbb{Q}_{\geq 0}$.

Definition 39. *We define the functions $\text{itPerms}_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \text{Bags}(\Lambda)$ and $\text{waitPerms}_{\text{lh}} : \text{Heaps}^{\text{log}} \rightarrow \text{Bags}(\Omega)$ mapping logical heaps to bags of iteration and wait permissions, respectively, as follows:*

$$\begin{aligned} \text{itPerms}_{\text{lh}}(H)(\delta) &:= \lfloor H(\text{itperm}_{\text{lRes}}(\delta)) \rfloor \\ \text{waitPerms}_{\text{lh}}(H)(id, \delta) &:= \lfloor H(\text{wperm}_{\text{lRes}}(id, \delta)) \rfloor \end{aligned}$$

Note that for consistent logical heaps (h^a, P^a) the above flooring is without any affect.

Lemma 10. *Let $(h_i^a, P_i^a)_{i \in \mathbb{N}}$ be an annotated reduction sequence such that every initial thread-local heap contains only finitely many permissions, i.e., such that for every $\theta \in \text{dom}(P_0^a)$ the set*

$$\begin{aligned} &\{\delta \in \Lambda \mid P_0^a(\theta).\text{heap}(\text{itperm}_{\text{lRes}}(\delta)) > 0\} \\ &\cup \\ &\{(id, \delta) \in \Omega \mid P_0^a(\theta).\text{heap}(\text{wperm}_{\text{lRes}}(id, \delta)) > 0\} \end{aligned}$$

is finite. Then, $\text{itPerms}_{\text{lh}}(P_i^a(\theta).\text{heap})$ and $\text{waitPerms}_{\text{lh}}(P_i^a(\theta).\text{heap})$ are finite for every choice of $i \in \mathbb{N}$ and $\theta \in \text{dom}(P_i^a)$.

Proof. Proof by induction on i .

Lemma 3. *Let $G((h_i^a, P_i^a)_{i \in \mathbb{N}})$ be a program order graph and let $p = (V, E)$ be a path in G with $S_p^\infty = \emptyset$. For every $\theta \in \text{dom}(P_0^a)$ let $P_0^a(\theta).\text{heap}$ be finite and complete. Then, p is finite.*

Proof. Assume p is infinite. We prove a contradiction by assigning a finite capacity to every node along the path. Consider the function $\text{nodeCap} : V \rightarrow \text{Bags}_{\text{fin}}(\Lambda)$ defined as

$$\text{nodeCap}(i) := \text{itPerms}_{\text{lh}}(P_i^a.\text{heap}) \uplus \bigoplus_{\substack{id \in \{id' \mid (id', _)\in \text{waitPerms}_{\text{lh}}(P_i^a.\text{heap})\} \\ L \in \text{Levs}}} \text{sigCap}_p((id, L), i).$$

For every $i \in V$, the capacity of node i , i.e., $\text{nodeCap}(i)$, is the union of two finite iteration permission bags: (i) To the left $\text{itPerms}_{\text{lh}}(P_i^a.\text{heap})$ captures all iteration permissions held by θ_i in step i . (ii) To the right $\bigoplus \text{sigCap}_p((id, L), i)$ captures all iteration permissions that will be created along the suffix of p that

starts at node i by waiting for signals for which thread θ_i already holds a wait permission (id, δ) in step i .

Note that for every $i \in V$, the bag of iteration permissions returned by $\text{nodeCap}(i)$ is indeed finite. All initial thread-local heaps contain only finitely many permissions. By Lemma 10 $\text{itPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})$ and $\text{waitPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap})$ are finite. Since signal IDs are unique, for every fixed choice of i and id , there is at most one level L , for which $\text{sigCap}_p((id, L), i) \neq \emptyset$. By assumption, along p all signals are waited for only finitely often, i.e., $S_p^\infty = \emptyset$. Hence, also the big union $\biguplus \text{sigCap}_p((id, L), i)$ is defined and finite.

Consider the sequence $(\text{nodeCap}(i))_{i \in V}$. Since every element is a finite bag of iteration permissions, we can order it by \prec_A . We are going to prove a contradiction by proving that the sequence is an infinitely descending chain.

Consider any edge $(i, \theta, n, j) \in E$. There are only three cases in which $\text{nodeCap}(i) \neq \text{nodeCap}(j)$ holds.

– $n = \text{GTP-RED-WAITPERM}$:

In this case, there are degrees δ, δ' with $\delta' \prec_\Delta \delta$, a signal s and $N \in \mathbb{N}$ for which we get

$$\text{nodeCap}(j) = (\text{nodeCap}(i) \setminus \{\delta\}) \uplus \underbrace{\{\delta', \dots, \delta'\}}_{N \text{ times}}.$$

That is, $\text{nodeCap}(j) \prec_A \text{nodeCap}(i)$.

– $n = \text{GTP-RED-WEAKITPERM}$: Same as above.

– $n = \text{AST-RED-CONSUMEITPERM}$:

In this case, we know that $\text{nodeCap}(j) = \text{nodeCap}(i) \setminus \{\delta\} \prec_A \text{nodeCap}(i)$ holds for some δ .

(Note that in case of $n = \text{GTP-RED-WAIT}$, we have $\text{nodeCap}(i) = \text{nodeCap}(j)$ since

$$\begin{aligned} \text{itPerms}_{\text{lh}}(P_j^{\text{a}}.\text{heap}) &= \text{itPerms}_{\text{lh}}(P_i^{\text{a}}.\text{heap}) \uplus \{\delta\}, \\ \biguplus \text{sigCap}_p((id, L), j) &= \left(\biguplus \text{sigCap}_p((id, L), i) \right) \setminus \{\delta\} \end{aligned}$$

holds for some δ .) So, nodeCap is monotonically decreasing.

By assumption p is infinite. According to Corollary 1 this implies that the path contains infinitely many consume edges, i.e., edges with a labelling $n = \text{AST-RED-CONSUMEITPERM}$. Hence, the sequence $(\text{nodeCap}(i))_{i \in V}$ forms an infinitely descending chain. However, according to Corollary 2, \prec_A is well-founded. A contradiction.

Lemma 4. *Let $(h_i^{\text{a}}, P_i^{\text{a}})_{i \in \mathbb{N}}$ be a fair annotated reduction sequence with $P_0^{\text{a}} = \{(\theta_0, (H_0, c))\}$, $\text{finite}_{\text{ah}}(h_0^{\text{a}})$, $\text{complete}_{\text{lh}}(H_0)$, $\text{finite}_{\text{lh}}(H_0)$ and $\text{consistent}_{\text{conf}}(h_0^{\text{a}}, P_0^{\text{a}})$. Let H_0 contain no signal or wait permission chunks. Further, let h_0^{a} contain no chunks $\text{unlocked}_{\text{aRes}}(m, P, H_P)$ where H_P contains any signal chunks. Let G be the program order graph of $(h_i^{\text{a}}, P_i^{\text{a}})_{i \in \mathbb{N}}$. Then, $S_G^\infty = \emptyset$.*

Proof. Suppose $S_G^\infty \neq \emptyset$. Since $\mathcal{L}evs$ is well-founded, the same holds for the set $\{\text{lev}(s) \mid s \in S^\infty\}$. Hence, there is some $s_{\min} \in S^\infty$ for which no $z \in S^\infty$ with $\text{lev}(z) <_{\mathbb{L}} \text{lev}(s_{\min})$ exists.

Since neither the initial logical heap H_0 nor any unlocked lock invariant stored in h_0^a does contain any signals, s_{\min} must be created during the reduction sequence. The reduction step creating signal s_{\min} is an application of GTP-RED-NEWSIGNAL, which simultaneously creates an obligation to set s_{\min} . By preservation of completeness, Lemma 7, every thread-local logical heap $P_i^a(\theta).\text{heap}$ annotating some thread θ in some step i is complete. According to reduction rule GTP-RED-WAIT, every wait edge $(a, \theta, (\text{GTP-RED-WAIT}, s_{\min}), b)$ implies together with completeness that in step a (i) thread θ does not hold any obligation for s_{\min} (i.e. $P_a^a(\theta).\text{heap}(\text{obs}_{\text{Res}}(O)) = 1$ for some bag of obligations O with $s_{\min} \notin O$) and (ii) s_{\min} has not been set, yet (i.e. $\text{signal}_{\text{aRes}}(s_{\min}, \text{False}) \in h_a^a$). Hence, in step a another thread $\theta_{\text{ob}} \neq \theta$ must hold the obligation for s_{\min} (i.e. $P_a^a(\theta_{\text{ob}}).\text{heap}(\text{obs}_{\text{Res}}(O)) = 1$ for some bag of obligations O with $s_{\min} \in O$). Since there are infinitely many wait edges concerning s_{\min} in G , the signal is never set.

By fairness, for every wait edge as above, there must be a non-ghost reduction step $h_k^a, P_k^a \xrightarrow{\theta_{\text{ob}} \text{atp}} h_{k+1}^a, P_{k+1}^a$ of the thread θ_{ob} holding the obligation for s_{\min} with $k \geq a$. Hence, there exists an infinite path p_{ob} in G where each edge $(e, \theta_{\text{ob}}, n, f) \in \text{edges}(p_{\text{ob}})$ concerns some thread θ_{ob} holding the obligation for s_{\min} . (Note that this thread ID does not have to be constant along the path, since the obligation can be passed on during fork steps.)

The path p_{ob} does not contain wait edges $(e, \theta_{\text{ob}}, (\text{GTP-RED-WAIT}, s^\infty), f)$ for any $s^\infty \in S^\infty$, since reduction rule GTP-RED-WAIT would (together with completeness of $P_e^a(\theta_{\text{ob}}).\text{heap}$) require s^∞ to be of a lower level than all held obligations. This restriction implies $\text{lev}(s^\infty) <_{\mathbb{L}} \text{lev}(s_{\min})$ and would hence contradict the minimality of s_{\min} . That is, $S_{p_{\text{ob}}}^\infty = \emptyset$.

By preservation of finiteness, Lemma 6, we get that every logical heap associated with the root of p_{ob} is finite. This allows us to apply Lemma 3, by which we get that p_{ob} is finite. A contradiction.