

Proving Linearizability Using Forward Simulations

Ahmed Bouajjani¹, Michael Emmi², Constantin Enea^{1(✉)},
and Suha Orhun Mutluergil³

¹ IRIF, Univ. Paris Diderot, Paris, France
{`abou,cenea`}@irif.fr

² Nokia Bell Labs, Murray Hill, USA
`michael.emmi@nokia.com`

³ Koc University, Sariyer, Turkey
`smutluergil@ku.edu.tr`

Abstract. Linearizability is the standard correctness criterion for concurrent data structures such as stacks and queues. It allows to establish observational refinement between a concurrent implementation and an atomic reference implementation. Proving linearizability requires identifying linearization points for each method invocation along all possible computations, leading to valid sequential executions, or alternatively, establishing forward *and* backward simulations. In both cases, carrying out proofs is hard and complex in general. In particular, backward reasoning is difficult in the context of programs with data structures, and strategies for identifying statically linearization points cannot be defined for all existing implementations. In this paper, we show that, contrary to common belief, many such complex implementations, including, e.g., the Herlihy and Wing Queue and the Time-Stamped Stack, can be proved correct using only forward simulation arguments. This leads to simple and natural correctness proofs for these implementations that are amenable to automation.

1 Introduction

Programming efficient concurrent implementations of atomic collections, e.g., stacks and queues, is error prone. To minimize synchronization overhead between concurrent method invocations, implementors avoid blocking operations like lock acquisition, allowing methods to execute concurrently. However, concurrency risks unintended inter-operation interference, and risks conformance to atomic reference implementations. Conformance is formally captured by (*observational*) *refinement*, which assures that all behaviors of programs using these efficient implementations would also be possible were the atomic reference implementations used instead.

An extended version of this paper including the missing proofs can be found at [8].

Observational refinement can be formalized as a trace inclusion problem, and the latter can itself be reduced to an invariant checking problem, but this requires in general introducing history and prophecy variables [1]. Alternatively, verifying refinement requires in general establishing a forward simulation *and* a backward simulation [21]. While simulations are natural concepts, backward reasoning, corresponding to the use of prophecy variables, is in general hard and complex for programs manipulating data structures. Therefore, a crucial issue is to understand the limits of forward reasoning in establishing refinement. More precisely, an important question is to determine for which concurrent abstract data structures, and for which classes of implementations, it is possible to carry out a refinement proof using only forward simulations.

To get rid of backward simulations (or prophecy variables) while preserving completeness w.r.t. refinement, it is necessary to have reference implementations that are *deterministic*. Interestingly, determinism allows also to simplify the forward simulation checking problem. Indeed, in this case, this problem can be reduced to an invariant checking problem. Basically, the simulation relation can be seen as an invariant of the system composed of the two compared programs. Therefore, existing methods and tools for invariant checking can be leveraged in this context.

But, in order to determine precisely what is meant by determinism, an important point is to fix the alphabet of observable events along computations. Typically, to reason about refinement between two library implementations, the only observable events are the calls and returns corresponding to the method invocations along computations. This means that only the external interface of the library is considered to compare behaviors, and nothing else from the implementations is exposed. Unfortunately, it can be shown that in this case, it is impossible to have deterministic atomic reference implementations for common data structures such as stacks and queues (see, e.g., [24]). Then, an important question is what is the necessary amount of information that should be exposed by the implementations to overcome this problem?

One approach addressing this question is based on linearizability [18] and its correspondence with refinement [7, 12]. Linearizability of a computation (of some implementation) means that each of the method invocations can be seen as happening at some point, called *linearization point*, occurring somewhere between the call and return events of that invocation. The obtained sequence of linearization points along the computation should define a sequence of operations that is possible in the atomic reference implementation. Proving the existence of such sequences of linearization points, for all the computations of a concurrent library, is a complex problem [3, 5, 14]. However, proving linearizability becomes less complex when linearization points are fixed for each method, i.e., associated with the execution of a designated statement in its source code [5]. In this case, we can consider that libraries expose in addition to calls and returns, events signaling linearization points. By extending this way the alphabet of observable events, it becomes straightforward to define *deterministic* atomic reference implementations. Therefore, proving linearizability can be carried out using forward

simulations when linearization points are fixed, e.g., [2, 4, 27, 28]. Unfortunately, this approach is not applicable to efficient implementations such as the LCRQ queue [22] (based on the principle of the Herlihy and Wing queue [18]), and the Time-Stamped Stack [10]. The proofs of linearizability of these implementations are highly nontrivial, very involved, and hard to read, understand and automatize. Therefore, the crucial question we address is what is precisely the kind of information that is necessary to expose in order to obtain deterministic atomic reference implementations for such data structures, allowing to derive simple and natural linearizability proofs for such complex implementations, based on forward simulations, that are amenable to automation?

We observe that the main difficulty in reasoning about these implementations is that, linearization points of enqueue/push operations occurring along some given computation, depend in general on the linearization points of dequeue/pop operations that occur arbitrarily far in the future. Therefore, since linearization points for enqueue/push operations cannot be determined in advance, the information that could be fixed and exposed can concern only the dequeue/pop operations.

One first idea is to consider that linearization points are fixed for dequeue/pop methods and only for these methods. We show that under the assumption that implementations expose linearizations points for these methods, it is possible to define deterministic atomic reference implementations for both queues and stacks. We show that this is indeed useful by providing a simple proof of the Herlihy and Wing queue (based on establishing a forward simulation) that can be carried out as an invariant checking proof.

However, in the case of Time-Stamped Stack, fixing linearization points of pop operations is actually too restrictive. Nevertheless, we show that our approach can be generalized to handle this case. The key idea is to reason about what we call *commit points*, and that correspond roughly speaking to the last point a method accesses to the shared data structure during its execution. We prove that by exposing commit points (instead of linearization points) for pop methods, we can still provide deterministic reference implementations. We show that using this approach leads to a quite simple proof of the Time-Stamped Stack, based on forward simulations.

2 Preliminaries

We formalize several abstraction relations between libraries using a simple yet universal model of computation, namely labeled transition systems (LTS). This model captures shared-memory programs with an arbitrary number of threads, abstracting away the details of any particular programming language irrelevant to our development.

A *labeled transition system* (LTS) $A = (Q, \Sigma, s_0, \delta)$ over the possibly-infinite alphabet Σ is a possibly-infinite set Q of states with initial state $s_0 \in Q$, and a transition relation $\delta \subseteq Q \times \Sigma \times Q$. The i th symbol of a sequence $\tau \in \Sigma^*$ is denoted τ_i , and ε denotes the empty sequence. An *execution* of A

is an alternating sequence of states and transition labels (also called actions) $\rho = s_0, e_0, s_1 \dots e_{k-1}, s_k$ for some $k > 0$ such that $\delta(s_i, e_i, s_{i+1})$ for each i such that $0 \leq i < k$. We write $s_i \xrightarrow{e_i \dots e_{j-1}}_A s_j$ as shorthand for the subsequence $s_i, e_i, \dots, s_{j-1}, e_{j-1}, s_j$ of ρ , for any $0 \leq i \leq j < k$ (in particular $s_i \xrightarrow{\varepsilon} s_i$). The projection $\tau|\Gamma$ of a sequence τ is the maximum subsequence of τ over alphabet Γ . This notation is extended to sets of sequences as usual. A *trace* of A is the projection $\rho|\Sigma$ of an execution ρ of A . The set of executions, resp., traces, of an LTS A is denoted by $E(A)$, resp., $Tr(A)$. An LTS is *deterministic* if for any state s and any sequence $\tau \in \Sigma^*$, there is at most one state s' such that $s \xrightarrow{\tau} s'$. More generally, for an alphabet $\Gamma \subseteq \Sigma$, an LTS is Γ -*deterministic* if for any state s and any sequence $\tau \in \Gamma^*$, there is at most one state s' such that $s \xrightarrow{\tau'} s'$ and τ is a subsequence of τ' .

2.1 Libraries

Programs interact with libraries by calling named library *methods*, which receive *arguments* and yield *return values* upon completion. We fix arbitrary sets \mathbb{M} and \mathbb{V} of method names and argument/return values. We fix an arbitrary set \mathbb{O} of operation identifiers, and for given sets \mathbb{M} and \mathbb{V} of methods and values, we fix the sets

$$C = \{inv(m, d, k) : m \in \mathbb{M}, d \in \mathbb{V}, k \in \mathbb{O}\} \text{ and } R = \{ret(m, d, k) : m \in \mathbb{M}, d \in \mathbb{V}, k \in \mathbb{O}\}$$

of *call actions* and *return actions*; each call action $inv(m, d, k)$ combines a method $m \in \mathbb{M}$ and value $d \in \mathbb{V}$ with an *operation identifier* $k \in \mathbb{O}$. Operation identifiers are used to pair call and return actions. We may omit the second field from a call/return action for methods that have no arguments or return values. For notational convenience, we take $\mathbb{O} = \mathbb{N}$ for the rest of the paper.

A *library* is an LTS over alphabet Σ such that $C \cup R \subseteq \Sigma$. We assume that the traces of a library satisfy standard well-formedness properties, e.g., return actions correspond to previous call actions. Given a standard library description as a set of methods, the LTS represents the executions of its most general client (that calls an arbitrary set of methods with an arbitrary set of threads in an unbounded loop). The states of this LTS consist of the shared state of the library together with the local state of each thread. The transitions correspond to statements in the methods' bodies, or call and return actions. An operation k is called *completed* in a trace τ when $ret(m, d, k)$ occurs in τ , for some m and d . Otherwise, it is called *pending*.

The projection of a library trace over $C \cup R$ is called a *history*. The set of histories of a library L is denoted by $H(L)$. Since libraries only dictate methods executions between their respective calls and returns, for any history they admit, they must also admit histories with weaker inter-operation ordering, in which calls may happen earlier, and/or returns later. A history h_1 is *weaker* than a history h_2 , written $h_1 \sqsubseteq h_2$, iff there exists a history h'_1 obtained from h_1 by appending return actions, and deleting call actions, s.t.: h_2 is a permutation of

h'_1 that preserves the order between return and call actions, i.e., if a given return action occurs before a given call action in h'_1 , then the same holds in h_2 .

A library L is called *atomic* when there exists a set S of sequential histories such that $H(L)$ contains every weakening of a history in S . Atomic libraries are often considered as specifications for concurrent objects. Libraries can be made atomic by guarding their methods bodies with global lock acquisitions.

A library L is called a *queue implementation* when $\mathbb{M} = \{enq, deq\}$ (enq is the method that enqueues a value and deq is the method removing a value) and $\mathbb{V} = \mathbb{N} \cup \{\text{EMPTY}\}$ where **EMPTY** is the value returned by deq when the queue is empty. Similarly, a library L is called a *stack implementation* when $\mathbb{M} = \{push, pop\}$ and $\mathbb{V} = \mathbb{N} \cup \{\text{EMPTY}\}$. For queue and stack implementations, we assume that the same value is never added twice, i.e., for every trace τ of such a library and every two call actions $inv(m, d_1, k_1)$ and $inv(m, d_2, k_2)$ where $m \in \{enq, push\}$ we have that $d_1 \neq d_2$. As shown in several works [2, 6], this assumption is without loss of generality for libraries that are data independent, i.e., their behaviors are not influenced by the values added to the collection. All the queue and stack implementations that we are aware of are data independent. On a technical note, this assumption is used to define (Γ -)deterministic abstract implementations of stacks and queues in Sects. 4 and 5.

2.2 Refinement and Linearizability

Conformance of a library L_1 to a specification given as an “abstract” library L_2 is formally captured by (*observational*) *refinement*. Informally, we say L_1 refines L_2 iff every computation of every program using L_1 would also be possible were L_2 used instead. We assume that a program can interact with the library only through call and return actions, and thus refinement can be defined as history set inclusion. Refinement is equivalent to the *linearizability* criterion [18] when L_2 is an atomic library [7, 12].

Definition 1. *A library L_1 refines another library L_2 iff $H(L_1) \subseteq H(L_2)$.*

Linearizability [18] requires that every history of a concurrent library L_1 can be “linearized” to a sequential history admitted by a library L_2 used as a specification. Formally, a sequential history h_2 with only complete operations is called a *linearization* of a history h_1 when $h_1 \sqsubseteq h_2$. A history h_1 is *linearizable* w.r.t. a library L_2 iff there exists a linearization h_2 of h_1 such that $h_2 \in H(L_2)$. A library L_1 is *linearizable* w.r.t. L_2 , written $L_1 \sqsubseteq L_2$, iff each history $h_1 \in H(L_1)$ is linearizable w.r.t. L_2 .

Theorem 1 [7, 12]. *Let L_1 and L_2 be two libraries, such that L_2 is atomic. Then, $L_1 \sqsubseteq L_2$ iff L_1 refines L_2 .*

In the rest of the paper, we discuss methods for proving refinement (and thus, linearizability) focusing mainly on queue and stack implementations.

3 Refinement Proofs

Library refinement is an instance of a more general notion of refinement between LTSs, which for some alphabet Γ of *observable actions* is defined as the inclusion of sets of traces projected on Γ . Library refinement corresponds to the case $\Gamma = C \cup R$. Typically, Γ -refinement between two LTSs A_1 and A_2 is proved using *simulation relations* which roughly, require that A_2 can mimic every step of A_1 using a (possibly empty) sequence of steps. Mainly, there are two kinds of simulation relations, forward or backward, depending on whether the preservation of steps is proved starting from a similar state forward or backward. It has been shown that Γ -refinement is equivalent to the existence of *backward simulations*, modulo the addition of history variables that record events in the implementation, and to the existence of *forward simulations* provided that the right-hand side LTS, A_2 , is Γ -deterministic [1,21]. We focus on proofs based on forward simulations because they are easier to automatize.

In general, forward simulations are *not* a complete proof method for library refinement because libraries are not $C \cup R$ -deterministic (the same sequence of call/return actions can lead to different states depending on the interleaving of the internal actions). However, there are classes of atomic libraries, e.g., libraries with “fixed linearization points” (defined later in this section), for which it is possible to identify a larger alphabet Γ of observable actions (including call/return actions), and implementations that are Γ -deterministic. For queues and stacks, Sects. 4 and 5 define other such classes of implementations that cover all the implementations that we are aware of.

Let $A_1 = (Q_1, \Sigma, s_0^1, \delta_1)$ and $A_2 = (Q_2, \Sigma, s_0^2, \delta_2)$ be two LTSs over Σ_1 and Σ_2 , respectively, and Γ an alphabet, such that $\Gamma \subseteq \Sigma_1 \cap \Sigma_2$.

Definition 2. *The LTS A_1 Γ -refines A_2 iff $Tr(A_1)|\Gamma \subseteq Tr(A_2)|\Gamma$.*

The notion of Γ -refinement instantiated to libraries (i.e., to LTSs defining libraries) implies the notion of refinement in Definition 1 for every Γ such that $C \cup R \subseteq \Gamma$.

We define a notion of *forward simulation* that can be used to prove Γ -refinement

Definition 3. *A relation $F \subseteq Q_1 \times Q_2$ is called a Γ -forward simulation from A_1 to A_2 iff $F(s_0^1, s_0^2)$ and:*

- For all $s, s' \in Q_1$, $\gamma \in \Gamma$, and $u \in Q_2$, such that $(s, \gamma, s') \in \delta_1$ and $F(s, u)$, we have that there exists $u' \in Q_2$ such that $F(s', u')$ and $u \xrightarrow{\sigma} u'$ where $\sigma_i = \gamma$, for some i , and $\sigma_j \in \Sigma_2 \setminus \Gamma$, for all $j \neq i$.
- For all $s, s' \in Q_1$, $e \in \Sigma_1 \setminus \Gamma$, and $u \in Q_2$, such that $(s, e, s') \in \delta_1$ and $F(s, u)$, we have that there exists $u' \in Q_2$ such that $F(s', u')$ and $u \xrightarrow{\sigma} u'$ where $\sigma \in (\Sigma_2 \setminus \Gamma)^*$.

A Γ -forward simulation states that every step of A_1 is simulated by a sequence of steps of A_2 . To imply Γ -refinement, every step of A_1 labeled by an observable action $\gamma \in \Gamma$ should be simulated by a sequence of steps of A_2 where

exactly one transition is labeled by γ and all the other transitions are labeled by non-observable actions. The dual notion of *backward* simulation where steps are simulated backwards can be defined similarly.

The following shows the soundness and the completeness of Γ -forward simulations (when A_2 is Γ -deterministic). It is an instantiation of previous results [1, 21].

Theorem 2. *If there is a Γ -forward simulation from A_1 to A_2 , then A_1 Γ -refines A_2 . Also, if A_1 Γ -refines A_2 and A_2 is Γ -deterministic, then there is a Γ -forward simulation from A_1 to A_2 .*

The linearization of a concurrent history can be also defined in terms of *linearization points*. Informally, a linearization point of an operation in an execution is a point in time where the operation is conceptually effectuated; given the linearization points of each operation, the linearization of a concurrent history is the sequential history which takes operations in the order of their linearization points. For some libraries, the linearization points of all the invocations of a method m correspond to the execution of a fixed statement in m 's body. For instance, when method bodies are guarded with a global-lock acquisition, the linearization point of every method invocation corresponds to the execution of the body. When the linearization points are fixed, we assume that the library is an LTS over an alphabet that includes actions $lin(m, d, k)$ with $m \in \mathbb{M}$, $d \in \mathbb{V}$ and $k \in \mathbb{O}$, representing the linearization point of the operation k returning value d . Let Lin denote the set of such actions. The projection of a library trace over $C \cup R \cup Lin$ is called an *extended history*. A trace or extended history is called *Lin-complete* when every completed operation has a linearization point, i.e., each return action $ret(m, d, k)$ is preceded by an action $lin(m, d, k)$. A library L over alphabet Σ is called *with fixed linearization points* iff $C \cup R \cup Lin \subseteq \Sigma$ and every trace $\tau \in Tr(L)$ is *Lin-complete*.

Proving the correctness of an implementation L_1 of a concurrent object such as a queue or a stack with fixed linearization points reduces to proving that L_1 is a $(C \cup R \cup Lin)$ -refinement of an abstract implementation L_2 of the same object where method bodies are guarded with a global-lock acquisition. As a direct consequence of Theorem 2, since the abstract implementation is $(C \cup R \cup Lin)$ -deterministic, proving $(C \cup R \cup Lin)$ -refinement is equivalent to finding a $(C \cup R \cup Lin)$ -forward simulation from L_1 to L_2 .

Sections 4 and 5 extend this result to queue and stack implementations where the linearization point of the methods *adding* values to the collection is *not* fixed.

4 Queues with Fixed Dequeue Linearization Points

The classical abstract queue implementation, denoted $AbsQ_0$, maintains a sequence of enqueued values; dequeues return the oldest non-dequeued value, at the time of their linearization points, or `EMPTY`. Some implementations, like the queue of Herlihy and Wing [18], denoted `HWQ` and listed in Fig. 1, are not forward-simulated by $AbsQ_0$, even though they refine $AbsQ_0$, since the order in

```

void enq(int x) {
    i = back++; items[i] = x;
}
int deq() {
    while (1) {
        range = back - 1;
        for (int i = 0; i <= range; i++) {
            x = swap(items[i], null);
            if (x != null) return x;
        }
    }
}
    
```

Fig. 1. The Herlihy and Wing Queue [18].

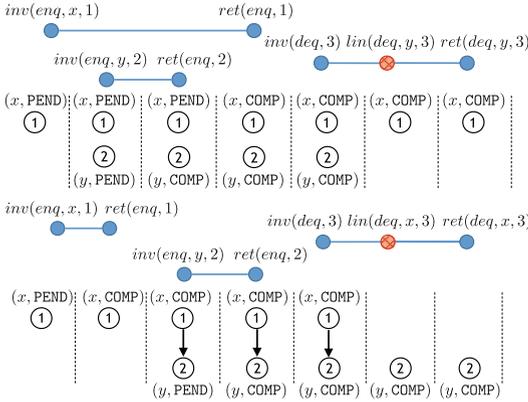


Fig. 2. Forward simulation with $AbsQ$. Lines depict operations, and circles depict call, return, and linearization point actions.

```

function loc:  $\mathbb{O} \rightarrow \{\text{inv}, \text{lin}, \text{ret}, \perp\}$ 
function arg, ret:  $\mathbb{O} \rightarrow \mathbb{V}$ 
function present, pending:  $\mathbb{O} \rightarrow \mathbb{B}$ 
function before:  $\mathbb{O} \times \mathbb{O} \rightarrow \mathbb{B}$ 
    
```

```

rule inv(enq, v, k):
    arg(k) := v
    present(k) := true
    pending(k) := true
    forall k1 with present(k1):
        if  $\neg$ pending(k1):
            before(k1, k) := true
    
```

```

rule ret(enq, k):
    pending(k) := false
    
```

```

rule inv(deq, k):
    pass
    
```

```

rule lin(deq, v, k):
    ret(k) := v
    if v = EMPTY:
        forall k' with present(k'):
            assert pending(k')
    else:
        let k1 =  $\text{arg}^{-1}(v)$ 
        assert present(k1)
        forall k2 with present(k2):
            assert  $\neg$ before(k2, k1)
        present(k1) := false
    
```

```

rule ret(deq, v, k):
    assert ret(k) = v
    
```

Fig. 3. The $AbsQ$ implementation; each rule $\alpha(_, k)$ implicitly begins with $\text{assert } \text{loc}(k) = \alpha$ and ends with the appropriate $\text{loc}(k) := \beta$.

which their enqueues are linearized to form $AbsQ_0$ s sequence is not determined until later, when their values are dequeued.

In this section we develop an abstract queue implementation, denoted $AbsQ$, which maintains a partial order of enqueues, rather than a linear sequence. Since $AbsQ$ does not force refining implementations to eagerly pick among linearizations of their enqueues, it forward-simulates many more queue implementations. In fact, $AbsQ$ forward-simulates all queue implementations of which we are aware that are not forward-simulated by $AbsQ_0$, including HWQ, The Baskets Queue [19], The Linked Concurrent Ring Queue (LCRQ) [22], and The Time-Stamped Queue [10].

4.1 Enqueue Methods with Non-Fixed Linearization Points

We describe HWQ where the linearization points of the enqueue methods are not fixed. The shared state consists of an array `items` storing the values in the queue and a counter `back` storing the index of the first unused position in `items`. Initially, all the positions in the array are `null` and `back` is 0. An enqueue method

starts by reserving a position in `items` (`i` stores the index of this position and `back` is incremented so the same position cannot be used by other enqueues) and then, stores the argument `x` at this position. The dequeue method traverses the array `items` starting from the beginning and atomically swaps `null` with the encountered value. If the value is not `null`, then the dequeue returns that value. If it reaches the end of the array, then it restarts.

The linearization points of the enqueues are not fixed, they depend on dequeues executing in the future. Consider the following trace with two concurrent enqueues ($i(k)$ represents the value of i in operation k): $inv(enq, x, 1)$, $inv(enq, y, 2)$, $i(1) = \text{bck}++$, $i(2) = \text{bck}++$, $items[i(2)] = y$. Assuming that the linearization point corresponds to the assignment of i , the history of this trace should be linearized to $inv(enq, x, 1)$, $ret(enq, 1)$, $inv(enq, y, 2)$, $ret(enq, 2)$. However, a dequeue executing until completion after this trace will return y (only position 1 is filled in the array `items`) which is not consistent with this linearization. On the other hand, assuming that enqueues should be linearized at the assignment of `items[i]` and extending the trace with $items[i(1)] = x$ and a completed dequeue that in this case returns x , leads to the incorrect linearization: $inv(enq, y, 2)$, $ret(enq, 2)$, $inv(enq, x, 1)$, $ret(enq, 1)$, $inv(deg, 3)$, $ret(deg, x, 3)$.

The dequeue method has a fixed linearization point which corresponds to an execution of `swap` returning a non-null value. This action alone contributes to the effect of that value being removed from the queue. Every concurrent history can be linearized to a sequential history where dequeues occur in the order of their linearization points in the concurrent history. This claim is formally proved in Sect. 4.3.

Since the linearization points of the enqueues are determined by future dequeue invocations, there exists no forward simulation from HWQ to $AbsQ_0$. In the following, we describe the abstract implementation $AbsQ$ for which such a forward simulation does exist.

4.2 Abstract Queue Implementation

Informally, $AbsQ$ records the set of enqueue operations, whose argument has not yet been removed by a matching dequeue operation. In addition, it records the happens-before order between those enqueue operations: this is a partial order ordering an enqueue k_1 before another enqueue k_2 iff k_1 returned before k_2 was invoked. The linearization point of a dequeue can either remove a minimal enqueue k (w.r.t. the happens-before stored in the state) and fix the return value to the value d added by k , or fix the return value to `EMPTY` provided that the current state stores only pending enqueues (intuitively, the dequeue overlaps with all the enqueue operations stored in the current state and it can be linearized before all of them).

Figure 2 pictures two executions of $AbsQ$ for two extended histories (that include dequeue linearization points). The state of $AbsQ$ after each action is pictured as a graph below the action. The nodes of this graph represent enqueue operations and the edges happens-before constraints. Each node is labeled by a value (the argument of the enqueue) and a flag `PEND` or `COMP` showing whether

the operation is pending or completed. For instance, in the case of the first history, the dequeue linearization point $lin(deq, y, 3)$ is enabled because the current happens-before contains a *minimal* enqueue operation with argument y . Note that a linearization point $lin(deq, x, 3)$ is also enabled at this state.

We define $AbsQ$ with the abstract state machine given in Fig. 3 which defines an LTS over the alphabet $C \cup R \cup Lin(deq)$. The state of $AbsQ$ consists of several updatable functions: `loc` indicates the abstract control point of a given operation; `arg` and `ret` indicate the argument or return value of a given operation, respectively; `present` indicates whether a given enqueue operation has yet to be removed, and `pending` indicates whether it has yet to complete; `before` indicates the happens-before order between operations. Initially, $loc(k) = inv$ for all k , and $present(k1) = pending(k1) = before(k1, k2) = false$ for all $k1, k2$. Each rule determines the next state of $AbsQ$ for the corresponding action. For instance, the $lin(deq, v, k)$ rule updates the state of $AbsQ$ for the linearization action of a dequeue operation with identifier k returning value v : when $v = EMPTY$ then $AbsQ$ insists via an assertion that any still-present enqueue must still be pending; otherwise, when $v \neq k$ then $AbsQ$ insists that a corresponding enqueue is present, and that it is minimal in the happens-before order, before marking that enqueue as not present. Updates to the `loc` function, implicit in Fig. 3, ensure that the invocation, linearization-point, and return actions of each operation occur in the correct order.

The following result states that the library $AbsQ$ has exactly the same set of histories as the standard abstract library $AbsQ_0$.

Theorem 3. *$AbsQ$ is a refinement of $AbsQ_0$ and vice-versa.*

A trace of a queue implementation is called *Lin(deq)-complete* when every completed dequeue has a linearization point, i.e., each return action $ret(deq, d, k)$ is preceded by an action $lin(deq, d, k)$. A queue implementation L over alphabet Σ , such that $C \cup R \cup Lin(deq) \subseteq \Sigma$, is called *with fixed dequeue linearization points* when every trace $\tau \in Tr(L)$ is *Lin(deq)-complete*.

The following result shows that $C \cup R \cup Lin(deq)$ -forward simulations are a sound and complete proof method for showing the correctness of a queue implementation with fixed dequeue linearization points (up to the correctness of the linearization points). It is obtained from Theorem 3 and Theorem 2 using the fact that the alphabet of $AbsQ$ is exactly $C \cup R \cup Lin(deq)$ and $AbsQ$ is deterministic. The determinism of $AbsQ$ relies on the assumption that every value is added at most once. Without this assumption, $AbsQ$ may reach a state with two enqueues adding the same value being both minimal in the happens-before. A transition corresponding to the linearization point of a dequeue from this state can remove any of these two enqueues leading to two different states. Therefore, $AbsQ$ becomes non-deterministic. Note that this is independent of the fact that $AbsQ$ manipulates operation identifiers.

Corollary 1. *A queue implementation L with fixed dequeue linearization points is a $C \cup R \cup Lin(deq)$ -refinement of $AbsQ_0$ iff there exists a $C \cup R \cup Lin(deq)$ -forward simulation from L to $AbsQ$.*

4.3 A Correctness Proof for Herlihy and Wing Queue

We describe a forward simulation F_1 from HWQ to $AbsQ$. The description of HWQ in Fig. 1 defines an LTS whose states contain the shared array `items` and the shared counter `back` together with a valuation for the local variables `i`, `x`, and `range`, and the control location of each operation. A transition is either a call or a return action, or a statement in one of the two methods `enq` or `deq`.

An HWQ state s is related by F_1 to $AbsQ$ states t where the predicate `present` is true for all the enqueues in s whose argument is stored in the array `items`, and all the pending enqueues that have not yet written to the array `items` (and only for these enqueues). We refer to such enqueues in s as `present` enqueues. Also, `pending(k)` is true in t whenever k is a pending enqueue in s , `arg(k) = d` in t whenever the argument of the enqueue k in s is d , and for every dequeue operation k such that `x(k) = d ≠ null`, we have that `y(k) = d` (recall that `y` is a local variable of the dequeue method in $AbsQ$). The order relation `before` in t satisfies the following constraints:

- (a) pending enqueues are maximal, i.e., for every two `present` enqueues k and k' such that k' is pending, we have that $\neg\text{before}(k', k)$,
- (b) `before` is consistent with the order in which positions of `items` have been reserved, i.e., for every two `present` enqueues k and k' such that `i(k) < i(k')`, we have that $\neg\text{before}(k', k)$,
- (c) if the position i reserved by an enqueue k has been “observed” by a non-linearized dequeue that in the current array traversal may “observe” a later position j reserved by another enqueue k' , then k can’t be ordered before k' , i.e., for every two `present` enqueues k and k' , and a dequeue k_d , such that

$$\text{canRemove}(k_d, k') \wedge (\text{i}(k) < \text{i}(k_d) \vee (\text{i}(k) = \text{i}(k_d) \wedge \text{afterSwapNull}(k_d))) \quad (1)$$

we have that $\neg\text{before}(k, k')$. The predicate `canRemove`(k_d, k') holds when k_d visited a `null` item in `items` and the position `i(k')` reserved by k' is in the range of (k_d) i.e., $(\text{x}(k_d) = \text{null} \wedge \text{i}(k_d) < \text{i}(k') \leq \text{range}(k_d)) \vee (\text{i}(k_d) = \text{i}(k') \wedge \text{beforeSwap}(k_d) \wedge \text{items}[\text{i}(k')]! = \text{null})$. The predicate `afterSwapNull`(k_d) (resp., `beforeSwap`(k_d)) holds when the dequeue k_d is at the control point after a `swap` returning `null` (resp., before a `swap`).

The constraints on `before` ensure that a `present` enqueue whose argument is about to be removed by a dequeue operation is minimal. Thus, let k' be a `present` enqueue that inserted its argument to `items`, and k_d a pending dequeue such that `canRemove`(k_d, k') holds and k_d is just before its swap action at the reserved position of k' i.e., `i(k_d) = i(k')`. Another pending enqueue k cannot be ordered before k' since pending enqueues are maximal by (a). Regarding the completed and `present` enqueues k , we consider two cases: $i(k) > i(k')$ and $i(k) < i(k')$. For the former case, the constraint (b) ensures $\neg\text{before}(k, k')$ and for the latter case the constraint (c) ensures $\neg\text{before}(k, k')$. Consequently, k' is a minimal element w.r.t. `before` just before k_d removes its argument.

Next, we show that F_1 is indeed a $C \cup R \cup \text{Lin}(\text{deq})$ -forward simulation. Let s and t be states of HWQ and $AbsQ$, respectively, such that $(s, t) \in F_1$. We omit

discussing the trivial case of transitions labeled by call and return actions which are simulated by similar transitions of *AbsQ*.

We show that each internal step of an enqueue or dequeue, except a `swap` returning a non-null value in dequeue (which represents its linearization point), is simulated by an *empty* sequence of *AbsQ* transitions, i.e., for every state s' obtained through one of these steps, if $(s, t) \in F_1$, then $(s', t) \in F_1$ for each *AbsQ* state t . Essentially, this consists in proving the following property, called *monotonicity*: the set of possible *before* relations associated by F_1 to s' doesn't exclude any order *before* associated to s .

Concerning enqueue rules, let s' be the state obtained from s when a pending enqueue k reserves an array position. This enqueue must be maximal in both t and any state t' related to s' (since it's pending). Moreover, there is no dequeue that can "observe" this position before restarting the array traversal. Therefore, item (c) in the definition of F_1 doesn't constrain the order between k and some other enqueue neither in s nor in s' . Since this transition doesn't affect the constraints on the order between enqueues different from k (their local variables remain unchanged), monotonicity holds. This property is trivially satisfied by the second step of enqueue which doesn't affect \mathbf{i} .

To prove monotonicity in the case of dequeue internal steps different from its linearization point, it is important to track the non-trivial instantiations of item (c) in the definition of *before* over the two states s and s' , i.e., the triples (k, k', k_d) for which (1) holds. Instantiations that are enabled only in s' may in principle lead to a violation of monotonicity (since they restrict the orders *before* associated to s'). For the two steps that begin an array traversal, i.e., reading the index of the last used position and setting \mathbf{i} to 0, there exist no such new instantiations in s' because the value of \mathbf{i} is either not set or 0. The same is true for the increment of \mathbf{i} in a dequeue k_d since the predicate `afterSwapNull(k_d)` holds in state s . The execution of `swap` returning `null` in a dequeue k_d enables new instantiations (k, k', k_d) in s' , thus adding potentially new constraints $\text{--before}(k, k')$. We show that these instantiations are however vacuous because k must be pending in s and thus maximal in every order *before* associated by F_1 to s . Let k and k' be two enqueues such that together with the dequeue k_d they satisfy the property (1) in s' but not in s . We write $i_s(k)$ for the value of the variable \mathbf{i} of operation k in state s . We have that $i_{s'}(k) = i_{s'}(k_d) \leq i_{s'}(k')$ and `items` $[i_{s'}(k_d)] = \text{null}$. The latter implies that the enqueue k didn't execute the second statement (since the position it reserved is still `null`) and it is pending in s' . The step that swaps the null item does not modify anything except the control point of k_d that makes `afterSwapNull(k_d)` true in s' . Hence, $i_s(k) = i_s(k_d) \leq i_s(k')$ and `items` $[i_s(k_d)] = \text{null}$ is also true. Therefore, k is pending in s and maximal. Hence, `before` (k, k') is not true in both s and s' .

Finally, we show that the linearization point of a dequeue k of *HWQ*, i.e., an execution of `swap` returning a non-null value d , from state s and leading to a state s' is simulated by a transition labeled by `lin` (deq, d, k) of *AbsQ* from state t . By the definition of *HWQ*, there is a unique enqueue k_e which filled the position updated by k , i.e., $i_s(k_e) = i_s(k)$ and $\mathbf{x}_{s'}(k) = \mathbf{x}_s(k_e)$.

We show that k_e is minimal in the order before of t which implies that k_e could be chosen by $lin(deg, d, k)$ step applied on t . As explained previously, instantiating item (c) in the definition of before with $k' = k_e$ and $k_d = k$, and instantiating item (b) with $k = k_e$, we ensure the minimality of k_e . Moreover, the state t' obtained from t through a $lin(deg, d, k)$ transition is related to s' because the value added by k_e is not anymore present in `items` and `present(k_e)` doesn't hold in t' .

5 Stacks with Fixed Pop Commit Points

The abstract implementation in Sect. 4 can be adapted to stacks, the main modification being that the linearization point $lin(pop, d, k)$ with $d \neq \text{EMPTY}$ is enabled when k is added by a push which is maximal in the happens-before order stored in the state. However, there are stack implementations, e.g., Time-Stamped Stack [10] (*TSS*, for short), which cannot be proved correct using forward simulations to this abstract implementation because the linearization points of the pop operations are not fixed. Exploiting particular properties of the stack semantics, we refine the ideas used in *AbsQ* and define a new abstract implementation for stacks, denoted as *AbsS*, which is able to simulate such implementations. Forward simulations to *AbsS* are complete for proving the correctness of stack implementations provided that the point in time where the return value of a pop operation is determined, called *commit point*, corresponds to a fixed statement in the pop method.

5.1 Pop Methods with Fixed Commit Points

We explain the meaning of the commit points on a simplified version of the Time-Stamped Stack [10] (*TSS*, for short) given in Fig. 4. This implementation

```

struct Node{
    int data;
    int ts;
    Node* next;
    bool taken;
};

bool CAS(bool data, bool a, bool b);

Node* pools[maxThreads];
int TS = 0;

void push(int x) {
    Node* n =
        new Node(x, MAX_INT, null, false);
    n->next = pools[myTID];
    pools[myTID] = n;
    int i = TS++;
    n->ts = i;
}

int pop() {
    bool success = false;
    int maxTS = -1;
    Node* youngest = null;
    while ( !success ) {
        maxTS = -1; youngest = null;
        for(int i=0; i<maxThreads; i++) {
            Node* n = pools[i];
            while (n->taken && n->next != n)
                n = n->next;
            if(maxTS < n->ts) {
                maxTS = n->ts; youngest = n;
            }
        }
        if (youngest != null)
            success =
                CAS(youngest->taken, false, true);
    }
    return youngest->data;
}

```

Fig. 4. The time-stamped stack [10].

maintains an array of singly-linked lists, one for each thread, where list nodes contain a data value (field `data`), a timestamp (field `ts`), the next pointer (field `next`), and a Boolean flag indicating whether the node represents a value removed from the stack (field `taken`). Initially, each list contains a sentinel dummy node pointing to itself with timestamp -1 and the flag `taken` set to `false`.

Pushing a value to the stack proceeds in several steps: adding a node with maximal timestamp in the list associated to the thread executing the push (given by the special variable `myTID`), asking for a new timestamp (given by the shared variable `TS`), and updating the timestamp of the added node. Popping a value from the stack consists in traversing all the lists, finding the first element which doesn't represent a removed value (i.e., `taken` is `false`) in each list, and selecting the element with the maximal timestamp. A compare-and-swap (CAS) is used to set the `taken` flag of this element to `true`. The procedure restarts if the CAS fails.

The push operations don't have a fixed linearization point because adding a node to a list and updating its timestamp are not executed in a single atomic step. The nodes can be added in an order which is not consistent with the order between the timestamps assigned later in the execution. Also, the value added by a push that just added an element to a list can be popped before the value added by a completed push (since it has a maximal timestamp). The same holds for pop operations: The only reasonable choice for a linearization point is a successful CAS (that results in updating the field `taken`). Figure 5 pictures an execution showing that this action doesn't correspond to a linearization point, i.e., an execution for which the pop operations in every correct linearization are

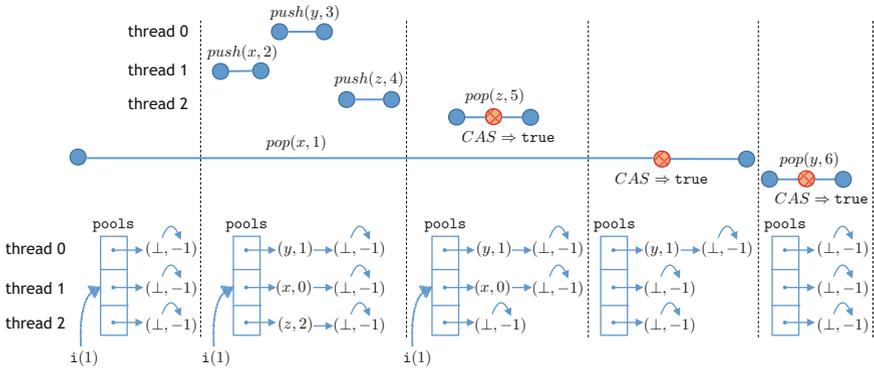


Fig. 5. An execution of *TSS*. An operation is pictured by a line delimited by two circles denoting the call and respectively, the return action. Pop operations with identifier k and removing value d are labeled $pop(d, k)$. Their representation includes another circle that stands for a successful CAS which is their commit point. The library state after an execution prefix delimited at the right by a dotted line is pictured in the bottom part (the picture immediately to the left of the dotted line). A pair (d, t) represents a list node with `data` = d and `ts` = t , and $i(1)$ denotes the value of `i` in the pop with identifier 1. We omit the nodes where the field `taken` is `true`.

not ordered according to the order between successful CASs. In every correct linearization of that execution, the pop operation removing x is ordered before the one removing z although they perform a successful CAS in the opposite order.

An interesting property of the successful CASs in pop operations is that they fix the return value, i.e., the return value is **youngest**->**data** where **youngest** is the node updated by the CAS. We call such actions *commit points*. More generally, commit points are actions that access shared variables, from which every control-flow path leads to the return control point and contains no more accesses to the shared memory (i.e., after a commit point, the return value is computed using only local variables).

When the commit points of pop operations are fixed to particular implementation actions (e.g., a successful CAS) we assume that the library is an LTS over an alphabet that contains actions $com(pop, d, k)$ with $d \in \mathbb{V}$ and $k \in \mathbb{O}$ (denoting the commit point of the pop with identifier k and returning d). Let $Com(pop)$ be the set of such actions.

5.2 Abstract Stack Implementation

We define an abstract stack $AbsS$ over alphabet $C \cup R \cup Com(pop)$ that essentially, similarly to $AbsQ$, maintains the happens-before order of the pushes whose value has not yet been removed by a matching pop. Pop operations are treated differently since the commit points are not necessarily linearization points. Intuitively, a pop can be linearized before its commit point. Each pop operation starts by taking a snapshot of the completed push operations which are maximal in the happens-before, more precisely, which don't happen before another completed push operation. Also, the library maintains the set of push operations overlapping with each pop operation. The commit point $com(pop, d, k)$ with $d \neq \text{EMPTY}$ is enabled if either d was added by one of the push operations in the initial snapshot, or by a push happening earlier when arguments of pushes from the initial snapshot have been removed, or by one of the push operations that overlaps with pop k . The commit point $com(pop, \text{EMPTY}, k)$ is enabled if all the values added by push operations happening before k have been removed. The effect of the commit points is explained below through examples.

Figure 6 pictures two executions of $AbsS$ for two extended histories (that include pop commit points). For readability, we give the state of $AbsS$ only after several execution prefixes delimited at the right by a dotted line. We focus on pop operations – the effect of push calls and returns is similar to enqueue calls and returns in $AbsQ$. Let us first consider the history on the top part. The first state we give is reached after the call of pop with identifier 3. This shows the effect of a pop invocation: the greatest completed pushes according to the current happens-before (here, the push with identifier 1) are marked as $\text{maxAtInvoc}(3)$, and the pending pushes are marked as $\text{overlap}(3)$. As a side remark, any other push operation that starts after pop 3 would be also marked as $\text{overlap}(3)$. The commit point $com(pop, x, 3)$ (pictured with a red circle) is enabled because x was added by a push marked as $\text{maxAtInvoc}(3)$. The effect of the commit point

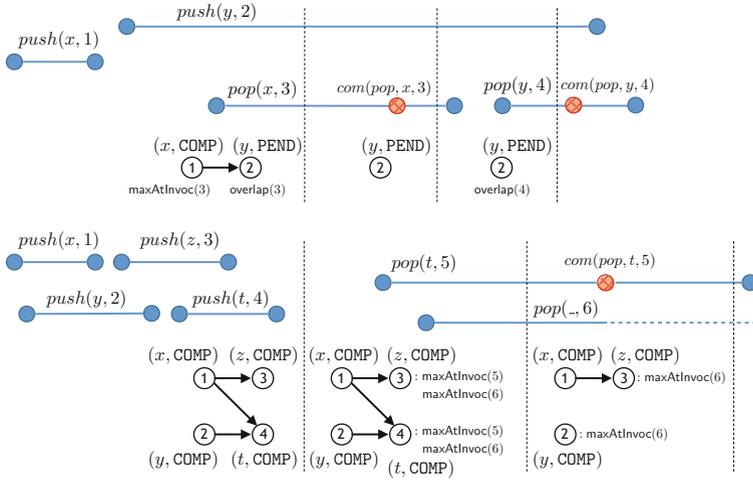


Fig. 6. Simulating stack histories with *AbsS*.

is that push 1 is removed from the state (the execution on the bottom shows a more complicated case). For the second pop, the commit point $com(pop, y, 4)$ is enabled because y was added by a push marked as $overlap(4)$. The execution on the bottom shows an example where the marking $maxAtInvc(k)$ for some pop k is updated at commit points. The pushes 3 and 4 are marked as $maxAtInvc(5)$ and $maxAtInvc(6)$ when the pops 5 and 6 start. Then, $com(pop, t, 5)$ is enabled since t was added by $push(t, 4)$ which is marked as $maxAtInvc(5)$. Besides removing $push(t, 4)$, the commit point produces a state where a pop committing later, e.g., pop 6, can remove y which was added by a predecessor of $push(t, 4)$ in the happens-before (y could become the top of the stack when t is removed). This history is valid because $push(y, 2)$ can be linearized after $push(x, 1)$ and $push(z, 3)$. Thus, push 2, a predecessor of the push which is removed, is marked as $maxAtInvc(6)$. Push 1 which is also a predecessor of the removed push is not marked as $maxAtInvc(6)$ because it happens before another push, i.e., push 3, which is already marked as $maxAtInvc(6)$ (the value added by push 3 should be removed before the value added by push 1 could become the top of the stack).

The description of *AbsS* as an abstract state machine is given in Fig. 7. Compared to *AbsQ*, the state contains two more updatable functions $maxAtInvc$ and $overlap$. For each pop operation k , $maxAtInvc(k)$ records the set of completed push operations which were maximal in the happens-before (defined by *before*) when pop k was invoked, or happening earlier provided that the values of all the pushes happening later than one of these maximal ones and before pop k have been removed. Also, $overlap(k)$ contains the push operations overlapping with a pop k . Initially, $loc(k) = inv$ for all k , $present(k1) = pending(k1) = before(k1, k2) = false$, and $maxAtInvc(k1) = overlap(k1) = \emptyset$, for all $k1, k2$. The rules for actions of push methods are similar to those for enqueues

```

function loc:  $\mathbb{O} \rightarrow \{\text{inv, com, ret, l}\}$    rule inv(pop, k):
function arg, ret:  $\mathbb{O} \rightarrow \mathbb{V}$            forall k1 with present(k1):
function present, pending:  $\mathbb{O} \rightarrow \mathbb{B}$      if pending(k1):
function before:  $\mathbb{O} \times \mathbb{O} \rightarrow \mathbb{B}$      overlap(k) := overlap(k)  $\cup$  {k1}
function maxAtInvoc:  $\mathbb{O} \rightarrow \wp(\mathbb{O})$     else:
function overlap:  $\mathbb{O} \rightarrow \wp(\mathbb{O})$         if  $\forall k2. \text{present}(k2) \wedge \text{before}(k1, k2) \Rightarrow \text{pending}(k2)$ :
                                           maxAtInvoc(k) = maxAtInvoc(k)  $\cup$  {k1}

rule inv(push, x, k):
  present(k) := true
  pending(k) := true
  arg(k) := x
  forall k1 with present(k1):
    if  $\neg$ pending(k1):
      before(k1, k) := true
  forall k1:
    overlap(k1) := overlap(k1)  $\cup$  {k}

rule com(push, k):
  pending(k) := false

rule ret(pop, y, k):
  assert ret(k) = y

rule com(pop, y, k):
  ret(k) := y
  if y = EMPTY:
    assert maxAtInvoc(k) =  $\emptyset$ 
  else:
    let k1 := arg-1(y)
    assert present(k1)
    assert k1  $\in$  maxAtInvoc(k)  $\cup$  overlap(k)
    present(k1) := false
    forall k2 with k1  $\in$  maxAtInvoc(k2):
      maxAtInvoc(k2) := maxAtInvoc(k2)  $\setminus$  {k1}
    forall k3 with before1(k3, k1):
      if  $\forall k4 \neq k1. \text{before}_1(k3, k4) \Rightarrow k4 \in \text{overlap}(k2)$ :
        maxAtInvoc(k2) := maxAtInvoc(k2)  $\cup$  {k3}

```

Fig. 7. The *AbsS* implementation; each rule $\alpha(-, k)$ implicitly begins with $\text{assert } \text{loc}(k) = \alpha$ and ends with the appropriate $\text{loc}(k) := \beta$; and before_1 denotes the transitive reduction of before .

in *AbsQ*, except that every newly invoked push operation k is added to the set $\text{overlap}(k1)$ for all pop operations $k1$ (since k overlaps with all the currently pending pops). The rule $\text{inv}(\text{pop}, k)$, marking the invocation of a pop, sets $\text{maxAtInvoc}(k)$ and $\text{overlap}(k)$ as explained above. The rule $\text{com}(\text{pop}, \text{EMPTY}, k)$ is enabled when the set $\text{maxAtInvoc}(k)$ is empty (otherwise, there would be push operations happening before pop k which makes the return value **EMPTY** incorrect). Also, $\text{com}(\text{pop}, y, k)$ with $y \neq \text{EMPTY}$ is enabled when y was added by a push $k1$ which belongs to $\text{maxAtInvoc}(k) \cup \text{overlap}(k)$. This rule may also update $\text{maxAtInvoc}(k2)$ for other pending pops $k2$. More precisely, whenever $\text{maxAtInvoc}(k2)$ contains the push $k1$, the latter is replaced by the immediate predecessors of $k1$ (according to before) that are followed exclusively by pushes overlapping with $k2$.

The abstract state machine in Fig. 7 defines an LTS over the alphabet $C \cup R \cup \text{Com}(\text{pop})$. Let AbsS_0 be the standard abstract implementation of a stack where elements are stored in a sequence, push and pop operations adding and removing an element from the beginning of the sequence in one atomic step, respectively. For $\mathbb{M} = \{\text{push}, \text{pop}\}$, the alphabet of AbsS_0 is $C \cup R \cup \text{Lin}$. The following result states that the library *AbsS* has exactly the same set of histories as AbsS_0 .

Theorem 4. *AbsS is a refinement of AbsS_0 and vice-versa.*

A trace of a stack implementation is called *Com(pop)-complete* when every completed pop has a commit point, i.e., each return $\text{ret}(\text{pop}, d, k)$ is preceded by an action $\text{com}(\text{pop}, d, k)$. A stack implementation L over Σ , such that $C \cup R \cup \text{Com}(\text{pop}) \subseteq \Sigma$, is called *with fixed pop commit points* when every trace $\tau \in \text{Tr}(L)$ is *Com(pop)-complete*.

As a consequence of Theorem 2, $C \cup R \cup Com(pop)$ -forward simulations are a sound and complete proof method for showing the correctness of a stack implementation with fixed pop commit points (up to the correctness of the commit points).

Corollary 2. *A stack L with fixed pop commit points is a $C \cup R \cup Com(pop)$ -refinement of $AbsS$ iff there is a $C \cup R \cup Com(pop)$ -forward simulation from L to $AbsS$.*

Linearization points can also be seen as commit points and thus the following holds.

Corollary 3. *A stack implementation L with fixed pop linearization points where transition labels $lin(pop, d, k)$ are substituted with $com(pop, d, k)$ is a $C \cup R \cup Com(pop)$ -refinement of $AbsS_0$ iff there is a $C \cup R \cup Com(pop)$ -forward simulation from L to $AbsS$.*

5.3 A Correctness Proof for Time-Stamped Stack

We describe a forward simulation F_2 from TSS to $AbsS$, which is similar to the one from HWQ to $AbsQ$ for the components of an $AbsS$ state which exist also in $AbsQ$ (i.e., different from $maxAtInvo$ and $overlap$).

Thus, a TSS state s is related by F_2 to $AbsS$ states t where $present(k)$ is true for every push operation k in s such that k has not yet added a node to $pool$ s or its node is still present in $pool$ s (i.e., the node created by the push has $taken$ set to $false$). Also, $pending(k)$ is true in t iff k is pending in s .

To describe the constraints on the order relation $before$ and the sets $maxAtInvo$ and $overlap$ in t , we consider the following notations: $ts_s(k)$, resp., $TID_s(k)$, denotes the timestamp of the node created by the push k in state s (the ts field of this node), resp., the id of the thread executing k . By an abuse of terminology, we call $ts_s(k)$ the timestamp of k in state s . Also, $k \rightsquigarrow_s k'$ when intuitively, a traversal of $pool$ s would encounter the node created by k before the one created by k' . More precisely, $k \rightsquigarrow_s k'$ when $TID_s(k) < TID_s(k')$, or $TID_s(k) = TID_s(k')$ and the node created by k' is reachable from the one created by k in the list pointed to by $pool$ s[$TID_s(k)$].

The order relation $before$ satisfies the following: (1) pending pushes are maximal, (2) $before$ is consistent with the order between node timestamps, i.e., $ts_s(k) \leq ts_s(k')$ implies $\neg before(k', k)$, and (3) $before$ includes the order between pushes executed in the same thread, i.e., $TID_s(k) = TID_s(k')$ and $ts_s(k) < ts_s(k')$ implies $before(k, k')$.

The components $maxAtInvo$ and $overlap$ satisfy the following constraints (their domain is the set of identifiers of pending pops):

Frontiers: By the definition of TSS , a pending pop p in s could, in the future, remove the value added by a push k which is maximal (w.r.t. $before$) or a push k which is completed but followed only by pending pushes (in the order relation $before$). Therefore, for all pop operations p which are pending in s ,

we have that $k \in \text{overlap}(p) \cup \text{maxAtInvoc}(p)$, for every push k such that $\text{present}(k) \wedge (\text{pending}(k) \vee (\forall k'. \text{present}(k') \wedge \text{before}(k, k') \rightarrow \text{pending}(k')))$.

TraverseBefore: A pop p with $\text{youngest}(p) \neq \text{null}$ that reached the node n overlaps with every present push that created a node with a timestamp greater than $\text{youngest}(p) \rightarrow \text{ts}$ and which occurs in `pools` before the node n . Formally, if $\text{youngest}_s(p) = n_s(k) \neq \text{null}$, $n_s(p) = n_s(k_1)$, $k_2 \rightsquigarrow_s k_1$, $\text{present}(k_2)$, and $\text{ts}_s(k_2) \geq \text{ts}_s(k)$, then $k_2 \in \text{overlap}(p)$, for each p, k_1, k_2 .

TraverseBeforeNull: A pop p with $\text{youngest}(p) = \text{null}$ overlaps with every push that created a node which occurs in `pools` before the node reached by p , i.e., $\text{youngest}_s(p) = \text{null}$, $n_s(p) = n_s(k_1)$, $k_2 \rightsquigarrow_s k_1$, and $\text{present}(k_2)$ implies $k_2 \in \text{overlap}(p)$, for each p, k_1, k_2 .

TraverseAfter: If the variable `youngest` of a pop p points to a node which is not taken, then this node was created by a push in $\text{maxAtInvoc}(p) \cup \text{overlap}(p)$ or the node currently reached by p is followed in `pools` by another node which was created by a push in $\text{maxAtInvoc}(p) \cup \text{overlap}(p)$. Formally, for each p, k_1, k_2 , if $\text{youngest}_s(p) = n_s(k_1)$, $n_s(k_1) \rightarrow \text{taken} = \text{false}$, and $n_s(p) = n_s(k_2)$, then one of the following holds:

- $k_1 \in \text{maxAtInvoc}(p) \cup \text{overlap}(p)$, or
- there exists a push k_3 in s such that $\text{present}(k_3)$, $k_3 \in \text{maxAtInvoc}(p) \cup \text{overlap}(p)$, $\text{ts}_s(k_3) > \text{ts}_s(k_1)$, and either $k_2 \rightsquigarrow_s k_3$ or $n_s(k_2) = n_s(k_3)$ and p is at a control point before the assignment statement that changes the variable `youngest`.

The functions `maxAtInvoc` and `overlap` satisfy more constraints which can be seen as invariants of *AbsS*, e.g., $\text{maxAtInvoc}(p)$ and $\text{overlap}(p)$ do not contain predecessors of pushes from $\text{maxAtInvoc}(p)$ (for each p, k_1, k_2 , $\text{before}(k_1, k_2)$ and $k_2 \in \text{maxAtInvoc}(p)$ implies $k_1 \notin \text{maxAtInvoc}(p) \cup \text{overlap}(p)$). They can be found in [8].

Note that F_2 cannot satisfy the reverse of **Frontiers**, i.e., every push in $\text{overlap}(p) \cup \text{maxAtInvoc}(p)$, for some p , is maximal or followed only by pending pushes (w.r.t., `before`). This is because the linearization points of pop operations are not fixed and they can occur anywhere in between their invocation and commit points. Hence, any push operation which was maximal or followed only by pending pushes in the happens-before in between the invocation and the commit can be removed by a pop. And such a push may no longer satisfy the same properties in the state s .

Based on the values stored in $\text{youngest}_s(p)$ and $n_s(p)$, for some pop p , the other three constraints identify other push operations that overlap with p , or they were followed only by pending pushes when p was invoked. **TraverseBefore** and **TraverseBeforeNull** state that pushes which add new nodes to the pools seen by p in the past, are overlapping with p . **TraverseAfter** states that either the push adding the current youngest node $\text{youngest}_s(p)$ is in $\text{overlap}_s(p) \cup \text{maxAtInvoc}_s(p)$, or there is a node that p will visit in the future which is in $\text{overlap}_s(p) \cup \text{maxAtInvoc}_s(p)$.

The proof that F_2 is indeed a forward simulation from *TSS* to *AbsS* follows the same lines as the one given for the Herlihy and Wing Queue. It can be found in [8].

6 Related Work

Many techniques for linearizability verification, e.g., [2, 4, 27, 28], are based on forward simulation arguments, and typically only work for libraries where the linearization point of every invocation of a method m is fixed to a particular statement in the code of m . The works in [9, 11, 25, 29] deal with *external* linearization points where the action of an operation k can be the linearization point of a concurrently executing operation k' . We say that the linearization point of k' is external. This situation arises in read-only methods like the `contains` method of an optimistic set [23], libraries based on the elimination back-off scheme, e.g., [15], or flat combining [13, 16]. In these implementations, an operation can do an update on the shared state that becomes the linearization point of a concurrent read-only method (e.g., a `contains` returning `true`) or an operation may update the data structure on behalf of other concurrently executing operations (whose updates are published in the shared state). In all these cases, every linearization point can still be associated syntactically to a statement in the code of a method and doesn't depend on operations executed in the future (unlike *HWQ* and *TSS*). However, identifying the set of operations for which such a statement is a linearization point can only be done by looking at the whole program state (the local states of all the active operations). This poses a problem in the context of compositional reasoning (where auxiliary variables are required), but still admits a forward simulation argument. For manual proofs, such implementations with external linearization points can still be defined as LTSs that produce *Lin*-complete traces and thus still fall in the class of implementations for which forward simulations are enough for proving refinement. These proof methods are not complete and they are not able to deal with implementations like *HWQ* or *TSS*.

There also exist linearizability proof techniques based on backward simulations or alternatively, prophecy variables, e.g., [20, 24, 26]. These works can deal with implementations where the linearization points are not fixed, but the proofs are conceptually more complex and less amenable to automation.

The works in [6, 17] propose reductions of linearizability to assertion checking where the idea is to define finite-state automata that recognize violations of concurrent queues and stacks. These automata are simple enough in the case of queues and there is a proof of *HWQ* based on this reduction [17]. However, in the case of stacks, the automata become much more complicated and we are not aware of a proof for an implementation such as *TSS* which is based on this reduction.

Acknowledgements. This work is supported in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 678177).

References

1. Abadi, M., Lamport, L.: The existence of refinement mappings. *Theoret. Comput. Sci.* **82**(2), 253–284 (1991)
2. Abdulla, P.A., Haziza, F., Holík, L., Jonsson, B., Rezine, A.: An integrated specification and verification technique for highly concurrent data structures. In: Piterman, N., Smolka, S.A. (eds.) *TACAS 2013*. LNCS, vol. 7795, pp. 324–338. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36742-7_23](https://doi.org/10.1007/978-3-642-36742-7_23)
3. Alur, R., McMillan, K.L., Peled, D.: Model-checking of correctness conditions for concurrent objects. *Inf. Comput.* **160**(1–2), 167–188 (2000)
4. Amit, D., Rinetzky, N., Reps, T.W., Sagiv, M., Yahav, E.: Comparison under abstraction for verifying linearizability. In: Damm, W., Hermanns, H. (eds.) *CAV 2007*. LNCS, vol. 4590, pp. 477–490. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-73368-3_49](https://doi.org/10.1007/978-3-540-73368-3_49)
5. Bouajjani, A., Emmi, M., Enea, C., Hamza, J.: Verifying concurrent programs against sequential specifications. In: Felleisen, M., Gardner, P. (eds.) *ESOP 2013*. LNCS, vol. 7792, pp. 290–309. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37036-6_17](https://doi.org/10.1007/978-3-642-37036-6_17)
6. Bouajjani, A., Emmi, M., Enea, C., Hamza, J.: On reducing linearizability to state reachability. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) *ICALP 2015*. LNCS, vol. 9135, pp. 95–107. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47666-6_8](https://doi.org/10.1007/978-3-662-47666-6_8)
7. Bouajjani, A., Emmi, M., Enea, C., Hamza, J.: Tractable refinement checking for concurrent objects. In: *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, 15–17 January 2015*, pp. 651–662 (2015)
8. Bouajjani, A., Emmi, M., Enea, C., Mutluergil, S.O.: Proving linearizability using forward simulations. *CoRR*, abs/1702.02705 (2017). <http://arxiv.org/abs/1702.02705>
9. Derrick, J., Schellhorn, G., Wehrheim, H.: Verifying linearisability with potential linearisation points. In: Butler, M., Schulte, W. (eds.) *FM 2011*. LNCS, vol. 6664, pp. 323–337. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-21437-0_25](https://doi.org/10.1007/978-3-642-21437-0_25). ISBN 978-3-642-21437-0
10. Dodds, M., Haas, A., Kirsch, C.M.: A scalable, correct time-stamped stack. In: *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, 15–17 January 2015*, pp. 233–246 (2015)
11. Drăgoi, C., Gupta, A., Henzinger, T.A.: Automatic linearizability proofs of concurrent objects with cooperating updates. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 174–190. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-39799-8_11](https://doi.org/10.1007/978-3-642-39799-8_11)
12. Filipovic, I., O’Hearn, P.W., Rinetzky, N., Yang, H.: Abstraction for concurrent objects. *Theoret. Comput. Sci.* **411**(51–52), 4379–4398 (2010)
13. Gorelik, M., Hendler, D.: Brief announcement: an asymmetric flat-combining based queue algorithm. In: *ACM Symposium on Principles of Distributed Computing, PODC 2013, Montreal, QC, Canada, 22–24 July 2013*, pp. 319–321 (2013)
14. Hamza, J.: On the complexity of linearizability. In: Bouajjani, A., Fauconnier, H. (eds.) *NETYS 2015*. LNCS, vol. 9466, pp. 308–321. Springer, Cham (2015). doi:[10.1007/978-3-319-26850-7_21](https://doi.org/10.1007/978-3-319-26850-7_21)

15. Hendler, D., Shavit, N., Yerushalmi, L.: A scalable lock-free stack algorithm. In: SPAA 2004, pp. 206–215. ACM (2004)
16. Hendler, D., Incze, I., Shavit, N., Tzafrir, M.: Flat combining and the synchronization-parallelism tradeoff. In: SPAA 2010: Proceedings of the 22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures, Thira, Santorini, Greece, 13–15 June 2010, pp. 355–364 (2010)
17. Henzinger, T.A., Sezgin, A., Vafeiadis, V.: Aspect-oriented linearizability proofs. In: CONCUR, pp. 242–256 (2013)
18. Herlihy, M., Wing, J.M.: Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.* **12**(3), 463–492 (1990)
19. Hoffman, M., Shalev, O., Shavit, N.: The baskets queue. In: Tovar, E., Tsigas, P., Fouchal, H. (eds.) OPODIS 2007. LNCS, vol. 4878, pp. 401–414. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-77096-1_29](https://doi.org/10.1007/978-3-540-77096-1_29). ISBN 978-3-540-77095-4
20. Liang, H., Feng, X.: Modular verification of linearizability with non-fixed linearization points. In: ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2013, Seattle, WA, USA, 16–19 June 2013, pp. 459–470 (2013)
21. Lynch, N.A., Vaandrager, F.W.: Forward and backward simulations: I. untimed systems. *Inf. Comput.* **121**(2), 214–233 (1995)
22. Morrison, A., Afek, Y.: Fast concurrent queues for x86 processors. In: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2013, Shenzhen, China, 23–27 February 2013, pp. 103–112 (2013)
23. O’Hearn, P.W., Rinetsyky, N., Vechev, M.T., Yahav, E., Yorsh, G.: Verifying linearizability with hindsight. In: PODC 2010, pp. 85–94. ACM (2010)
24. Schellhorn, G., Wehrheim, H., Derrick, J.: How to prove algorithms linearisable. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 243–259. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31424-7_21](https://doi.org/10.1007/978-3-642-31424-7_21)
25. Vafeiadis, V.: Automatically proving linearizability. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 450–464. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-14295-6_40](https://doi.org/10.1007/978-3-642-14295-6_40)
26. Vafeiadis, V.: Modular fine-grained concurrency verification. Ph.D. thesis, University of Cambridge (2008)
27. Vafeiadis, V.: Shape-value abstraction for verifying linearizability. In: Jones, N.D., Müller-Olm, M. (eds.) VMCAI 2009. LNCS, vol. 5403, pp. 335–348. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-93900-9_27](https://doi.org/10.1007/978-3-540-93900-9_27)
28. Vafeiadis, V., Herlihy, M., Hoare, T., Shapiro, M.: Proving correctness of highly-concurrent linearisable objects. In: PPOPP 2006, pp. 129–136. ACM (2006)
29. Zhu, H., Petri, G., Jagannathan, S.: POLING: SMT aided linearizability proofs. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9207, pp. 3–19. Springer, Cham (2015). doi:[10.1007/978-3-319-21668-3_1](https://doi.org/10.1007/978-3-319-21668-3_1)