

POLING: SMT Aided Linearizability Proofs

He Zhu^(✉), Gustavo Petri, and Suresh Jagannathan

Purdue University, West Lafayette, USA

zhuhemail@gmail.com



Abstract. Proofs of linearizability of concurrent data structures generally rely on identifying linearization points to establish a simulation argument between the implementation and the specification. However, for many linearizable data structure operations, the linearization points may not correspond to their internal static code locations; for example, they might reside in the code of another concurrent operation. To overcome this limitation, we identify important program patterns that expose such instances, and describe a tool (POLING) that automatically verifies the linearizability of implementations that conform to these patterns.

1 Introduction

Linearizability [13] is the de facto correctness condition for the implementation of concurrent data structures. In a nutshell, linearizability establishes an observational equivalence between a multi-step fine-grained implementation, and an atomic coarse-grained specification of the data structure [7]. Thus, linearizability implies that each *operation* of a data structure implementation can be considered as executing atomically with respect to other concurrent operations.

While the definition of linearizability is intuitively simple, its proofs tend to be complex, and oftentimes depend on complicated simulation relations between the *abstraction* and the *implementation* (eg. [3, 18]). A common strategy to define linearizability simulations is to identify program points in the implementation of the data structure – known as *linearization points* (LP) – upon whose execution an operation can be considered to have happened atomically [21]. Technically, a linearization point indicates in a weak simulation argument a unique and atomic step in the execution of the implementation at which the specification and the implementation match their behaviors.

In this paper, we present a lightweight technique which reduces this complex task into a property checking problem. In our approach, the abstract specification of a data structure operation is defined using *recursive functions* on heaplets (assertions describing a portion of the heap). We model these assertions as a set of memory locations following [17], enabling the use of SMT solvers to discharge our proof obligations. For example, at a linearization point (LP) of a *stack push* method, the *set of locations conforming the stack* after the execution of LP, should be equal to the set of locations conforming the stack before the push, plus an additional location containing the new value pushed. The validity of the formulae relating these two states of this linearizability argument (over *sets*) is

decidable because they can be translated to quantifier-free formulas over integers and functions, and can be solved by using SMT solvers.

Tools like CAVE [22] are successful at automatically proving linearizability for data structures where the linearization points – static program locations – can be affixed to static program locations within each operation implementation. Unfortunately, a large class of lock-free linearizable data structures resist proofs by identifying such linearization points [3, 11, 16]. This should not be surprising since the definition of linearizability only requires the existence of a linearization for each invocation, which is intrinsically a *semantical* property of the implementation. Such a linearization needs not always correspond to the execution of predefined instructions within the code of the operation implementation. Our work extends previous efforts on the automation of linearizability proofs leveraging internal linearization points [22], by identifying data structure implementations whose linearization points may reside in code belonging to a concurrent operation.

We do so by identifying two common *patterns*, occurring in fine-grained concurrent algorithms which cannot be verified using linearization points. Our ideas extend the state of the art in automatic linearizability verification by extending CAVE [22] beyond linearization points. To the verification of linearizability by linearization points, POLING adds the following three notable features.

- Firstly, POLING converges faster than CAVE in all the benchmarks that both tools can handle. This can be attributed to the fact that POLING reasons about separation logic (SL) verification conditions by interpreting them as *sets* of memory locations, following [17], with efficient SMT solvers.
- Secondly, POLING can verify linearizability for implementations that use *helping* [14], a mechanism that allows an operation performed by one thread to be completed by a concurrent operation of a different thread.
- Finally, POLING is sensitive to *hindsight* [16], a specific pattern in which the linearization of operations that do not update the state (called *pure*) can only be established *a posteriori* of their execution, and it might depend on other threads operations.

2 Overview

In the tradition of CAVE, we focus on the linearizability of concurrent data structures implemented using linked lists. Some examples of such data structures are: stacks, queues and sets (see [12]). In our work we assume a garbage collected environment. Our approach is defined with respect to the *most general client* (MGC) program in which an unbounded number of threads interact through a single shared instance of the data structure [21], hence generalizing any possible client.

To introduce POLING, consider the implementation of a lock-free stack due to Treiber [19] shown in Fig. 1. The stack is organized as a linked list of nodes containing a payload (`val`), and a `next` pointer to the following node in the list. The head of the list is saved in the field `first` of a shared global variable

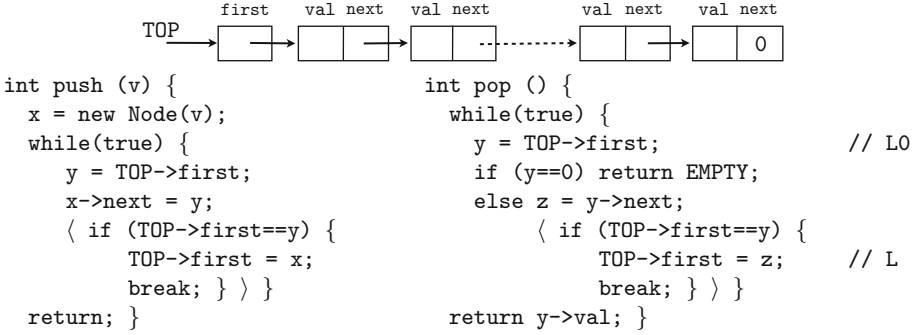


Fig. 1. Treiber stack.

TOP, which serves as the synchronization point for competing threads. Thus, TOP->first points to the last-in node, when there is one, and it contains 0 if the stack is empty. The operations `push` and `pop` are presented. As customary, we group blocks that should be executed atomically with “{” and “}” (in this case they represent an inline of a `cas` instruction). The operation `push` creates a new node - it reads the head of the linked list, and then tries to atomically update it. If some other thread modifies TOP->first between the read and the write, the update is aborted and the process is reiterated. A similar argument applies to `pop`.

Since linearizability relates the implementation to a specification, the programmer must provide the specification. In the case of CAVE this is done by means of a simple language with primitives for the interpretation of lists and sets. The programmer can use these primitives to operationally specify the data structure. On the other hand, in POLING the specification is done declaratively using *recursive functions*. The *specification* of `pop` is defined declaratively as:

$$\begin{aligned}
\text{vals}(\sigma, \text{TOP} \rightarrow \text{first}) = r :: \text{vals}(\sigma', \text{TOP} \rightarrow \text{first}) &\iff \text{pop}() = r \\
\text{vals}(\sigma, \text{TOP} \rightarrow \text{first}) = [] &\iff \text{pop}() = \text{Empty}
\end{aligned}$$

In these equations, TOP and first are program variables. The keyword `vals` is a primitive used to abstractly refer to the contents of a data structure, where the argument σ represents the current state, and the argument TOP->first is a program expression pointing to the first element of the list. The state resulting from executing `pop` is represented by the symbol σ' . Finally, the return value of `pop` is denoted by r . Then, the specification of `pop` establishes the relation between the abstract state of the list before and after its execution. In this example, POLING interprets `vals`(σ , TOP->first) as a *recursive function* [17], defining the values stored in the locations reachable from TOP->first in σ . Recursive functions will be formalized in Sect. 3.

In Fig. 1, the program point marked with L0 corresponds to the linearization point of `pop` when the list is empty. The specification for the non-empty case is satisfied at the program point L, because the set of values of locations reachable

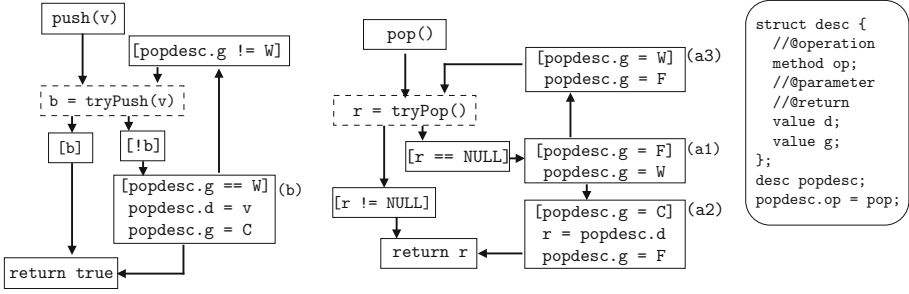


Fig. 2. Elimination based stack. The formulae in square brackets are assumes and each node denotes an atomic block. Struct `desc` is an annotated descriptor.

from `TOP`->`first` in a state σ (before the execution of the statement at `L`) equals to the set of values reachable in the resulting state (σ') plus the element pointed by `y`, which is `pop`'s return value. Importantly no other point in `pop` changes the values of the reachable locations. `POLING` using the set of memory locations at `L` calls an SMT solver to automatically check whether the update to these locations, before and after `L`, respects the specification of `pop`. We present verification details in Sect. 4.

The example shown in Fig. 2 (adapted from [6]) is a simplification of the HSY stack [10]. The stack uses an underlying Treiber implementation, but improves its performance for high-contention scenarios through an elimination layer, an occurrence of the *helping mechanism*. The example allows pairs of overlapping instances of `pop` and `push` to exchange their values without accessing the list. Each operation attempts an iteration of the Treiber implementation, and if it fails it applies the helping mechanism. The process is iterated until the operation succeeds. Helping is implemented by storing a descriptor of a pending `pop` into the shared state variable `popdesc`. In the descriptor, if the value of `g` is `W` (aiting) in location `a1`, there is an invocation of `pop` ready to be helped. If the value of `g` is `C` (ollided), then a pair of `push` and `pop` invocations is in the process of exchanging the value through `g` in location `b`. The helping completes or gives up when `pop` sets the value of `g` back to `F` (ree), in locations `a2` and `a3`, so that another instance of helping can happen. Importantly, the linearization of `pop` can happen in the code of a concurrent `push` in location `b` and, in location `a2`, the `pop` can witness that it has been helped because the transition through `a1` to `a2` can only be caused by an action that happened in location `b` of a `push` thread, where it was linearized.

We exploit this witnessing strategy in our linearization proofs. When an operation is linearized by another operation that helps it, we record the state at which the operation was linearized. Then, in the verification of the helped operation, we can use the recorded state as a witness to check that it was indeed helped by a concurrent operation. We observe that the witnessing is carried out through the descriptors (`popdesc`). To verify how a `push` can find and then help a pending `pop`, we allow programmers to use JML-style markers to annotate the

descriptor data structure, indicating for an operation’s name, the parameters and the return value (in Fig. 2 denoted via `@operation`, `@parameter` and `@return` respectively). We cover the details of the verification in Sect. 4.

3 Formal Model

We define a data structure D as a pair (D_Σ, D_M) consisting of D_Σ , the domain of the data structure; and D_M a set of method names, or *primitive operations* in the terminology of [13]. We call an *invocation* of the method m simply an *operation*, and use the metavariable op_m to range over the set \mathcal{Ops} of all possible invocations. In turn, operations can be decomposed into a tuple $op_m = (m, t, v, r)$ containing a thread identifier $t \in \mathcal{Tid}$, a vector with the arguments v , and if the operation is completed, a return value r .

Program Model. We omit a description of the program state and operational semantics of the programming language, assumed to be a standard first-order, shared-memory, concurrent, imperative language. We assume a set of states $\sigma \in D_\Sigma$, and an operational semantics with execution steps between states labeled by events $ev \in \mathcal{Evs}$ following the judgment: $\sigma \xrightarrow{t, ev} \sigma'$. Events capture data and control-flow actions (loads, stores, conditionals, etc.) with the addition of *invocation* and *response* events of operation op_m , denoted $inv(op_m)$ and $res(op_m)$ respectively. These latter two kinds of events serve to delimit the “duration” of operation invocations. We assume the obvious extension of this step judgment to traces of events ranged by $\mathbf{tr} \in \mathcal{Evs}^*$, denoted $\sigma \xrightarrow{\mathbf{tr}} \sigma'$, with the obvious inductive definition. Where unnecessary, we also omit the intermediate states.

Linearizability. Following [21], and without loss of generality, we assume a thread makes at most one invocation to an operation of the data structure. We then overload the notations for invocations and responses as $inv(t)$ and $res(t)$, for thread t . We define $history(\mathbf{tr})$ to be the projection of invocation and response events in trace \mathbf{tr} . Traces induce a partial order between operations. We say that a operation t precedes operation t' in \mathbf{tr} $t' \text{ intr } t$ ($t \prec_{\mathbf{tr}} t'$) as defined below¹:

$$t \prec_{\mathbf{tr}} t' \iff \exists \mathbf{tr}_0 \mathbf{tr}_1 \mathbf{tr}_2, \mathbf{tr} = \mathbf{tr}_0 \cdot res(t) \cdot \mathbf{tr}_1 \cdot inv(t') \cdot \mathbf{tr}_2$$

We say that a history is *sequential* if each invocation event $inv(t)$ is immediately followed by its corresponding response event $res(t)$.

Definition 1 (Linearizable History). *A history h is linearizable w.r.t. the specification of a data structure D if, and only if, there exists a sequential trace h_s of D such that $\prec_h \subseteq \prec_{H_s}$, and the set of operations of h and h_s coincide.*

This notion is trivially lifted to implementations by requiring every implementation trace to have a linearizable history.

¹ We eschew treating uncompleted invocations [13], which is a simple extension.

Set Specification. As hinted in Sect. 2, we abstract the data structure through an abstraction function: $\text{vals}(\sigma, \mathbf{x})$, which for a given state σ represents the set of values stored in the data structure pointed by the program variable \mathbf{x} . This allows us to express the abstract data structure specification as an equation relating the initial state σ and final state σ' . Then for *Set* data structures implemented as an ordered list, where we assume that the global variable $\text{head}(\text{Set})$ points to the beginning of the list, we write:

$$\begin{aligned} \text{vals}(\sigma', \text{head}(\text{Set})) &= \{v\} \cup \text{vals}(\sigma, \text{head}(\text{Set})) &\iff \text{add}(v) = \text{true} \\ \text{vals}(\sigma', \text{head}(\text{Set})) \cup \{v\} &= \text{vals}(\sigma, \text{head}(\text{Set})) &\iff \text{contains}(v) = \text{true} \\ v \notin \text{vals}(\sigma, \text{head}(\text{Set})) &&\iff \text{contains}(v) = \text{false} \end{aligned}$$

Order Preserving Specifications. Using the concatenation operator ($::$) instead of union (\cup) as above, we can capture the behavior of data structures whose temporal behavior imposes an ordering. For example a stack `pop` should always return the “last” value pushed. Assuming that $\text{head}(\text{D})$ is the global variable pointing to the first element of the list, we can specify `pop` and `dequeue` – where we omit all other methods – as:

$$\begin{aligned} \text{vals}(\sigma, \text{head}(\text{Stk})) &= r :: \text{vals}(\sigma', \text{head}(\text{Stk})) &\iff \text{pop}() = r \\ \text{vals}(\sigma, \text{head}(\text{Queue})) &= \text{vals}(\sigma', \text{head}(\text{Queue})) :: r &\iff \text{dequeue}() = r \end{aligned}$$

State Abstraction. Since we use RGSep [21], we abstract the program state σ using separation logic formulae, denoted by a metavariable ψ . The syntax of these formulae is given by the following grammars, where \approx ranges over binary relations of expression.

$$\begin{aligned} \psi &:: P * \boxed{P'} \mid \psi \vee \psi' & \Pi &:: \text{true} \mid \text{false} \mid E \approx E' \mid (\Pi \wedge \Pi') \\ P &:: \Pi \wedge \Gamma & \Gamma &:: \text{emp} \mid E_{tl} \mapsto \rho \mid \Gamma * \Gamma' \mid \text{lseg}_{tl, \rho}(E, E') \end{aligned}$$

We briefly describe these assertions (details can be found in [21]): 1. The formulae are given in disjunctive normal form, representing the different possible states reachable through different paths, 2. Each of the disjuncts has two parts, the local state, P predicating over the state local to the thread, and a shared state $\boxed{P'}$ – demarcated by a box [21] – which predicates over the state accessible to all threads, 3. Further, each of these states can be separated into a *pure* part Π , only concerned with stack allocated variables, and a *spatial* part Γ , which describes the heap, 4. Finally, heap assertions include the standard separation logic operators, where $E_{tl} \mapsto \rho$ denotes that the location E contains a record ρ (a mapping from field names to values), where the special field tl points to the next node in a linked list, and $\text{lseg}_{tl, \rho}(E, E')$ denotes a linked list segment starting at location E and ending at E' . The same convention applies to tl . All the nodes in this list segment share a same field-value mapping described in ρ .

Data Structure Abstraction. We use the method of [17] to discharge proof obligations about the state using an SMT solver. In [17] SL assertions are encoded as predicates on sets of memory locations. To that end, we define in Fig. 3 a data structure abstraction function that takes an RGSep assertion, and transforms it into a set of values. This is the function $(\text{vals})(\psi_\sigma, \mathbf{x})$, which is the symbolic

$$\begin{aligned}
\langle \text{vals} \rangle (P * \boxed{\Pi \wedge \Gamma}, E) &= \text{val}_{\text{aux}}(\Gamma, \text{Find}(\Gamma, E)) \\
\text{where} \\
\text{val}_{\text{aux}}(\Gamma, E_{tl} \mapsto \rho) &= \text{VAL}(E) \cup \langle \text{vals} \rangle (\boxed{\Gamma}, \rho(tl)) \\
\text{val}_{\text{aux}}(\Gamma, \text{lseg}_{tl, \rho}(E, E')) &= \text{VALS}(E, E') \cup \langle \text{vals} \rangle (\boxed{\Gamma}, E')
\end{aligned}$$

Fig. 3. Recursive abstraction definition

counterpart to the function $\text{vals}(\sigma, \mathbf{x})$ used for specifications. This recursive abstraction definition represents the set of elements that inhabit the data structure. The function $\text{Find}(\Gamma, E)$, simply finds the syntactic atomic occurrence of a node or a list starting from E in the spacial formula Γ . We omit its trivial recursive definition over RGSep formulae. Here, the capitalized expressions VALS are uninterpreted functions in the logic of the underlying theorem provers we use [17]. Finally, notice that if we substitute the concatenation operator ($::$) for all the occurrences of the union operator (\cup) in Fig. 3, we obtain a recursive definition for lists instead of sets. A refined definition of our data structure abstraction is given in [25] (that also considers reachable locations that are logically detached).

4 Verification

Our verification begins after a pass of the frontend of CAVE [22], which given the implementation of a data structure D , uses symbolic execution and shape analysis to produce D 's data structure invariant [4] and RGSep [21] rely-guarantee actions. To aid the verification of the helping mechanism, POLING requires the programmer to instrument descriptors (as exemplified in Fig. 2). A simple analysis could be implemented to instrument the descriptors automatically, or an interface could be implemented to indicate the thread descriptors, but we omit this unrelated step to simplify our development.

Central to our development are *LP (linearization point) functions*.

Definition 2 (Valid LP). Assume a trace $\text{tr} : \sigma_i \xrightarrow{\text{tr}} \sigma_f$ of the data structure D , and a function $\varrho : \mathcal{Ops} \rightarrow \mathcal{D}_\Sigma$, mapping each operation op of tr to the state in tr exactly prior to the linearization of op . We say that ϱ is a valid linearization point function for tr with respect to an abstract specification φ if:

1. every operation $op \in \text{tr}$, has an LP state (i.e. $\varrho(op)$) strictly between its invocation ($\text{inv}(op)$) and its response ($\text{res}(op)$). Formally:

$$\exists \text{tr}_{(1:4)}, \sigma_i \xrightarrow{\text{tr}} \sigma_f = \sigma_i \xrightarrow{\text{tr}_1} \cdot \xrightarrow{\text{inv}(op)} \cdot \xrightarrow{\text{tr}_2} \varrho(op) \xrightarrow{\text{tr}_3} \cdot \xrightarrow{\text{res}(op)} \cdot \xrightarrow{\text{tr}_4} \sigma_f$$

2. only the states that linearize operations can affect the abstract data structure:

$$\begin{aligned}
\sigma_1 \notin \{\varrho(op) \mid op \in \text{tr}\} \text{ and } \sigma_i \xrightarrow{\text{tr}} \sigma_f = \sigma_i \xrightarrow{\text{tr}_1} \sigma_1 \xrightarrow{ev} \sigma_2 \xrightarrow{\text{tr}_2} \sigma_f \\
\Rightarrow \text{vals}(\sigma_1, \text{head}(D)) = \text{vals}(\sigma_2, \text{head}(D))
\end{aligned}$$

3. for each operation $op \in \text{tr}$, we have that the LP state, and its subsequent state are related by the data structure specification φ . Formally, if the abstract specification of $op = (\mathbf{m}, t, v, r)$ is φ , for a trace of op : $\sigma_i \xrightarrow{\text{tr}} \sigma_f = \sigma_i \xrightarrow{\text{tr}_1} \varrho(op) \xrightarrow{ev} \hat{\sigma} \xrightarrow{\text{tr}_2} \cdot \xrightarrow{\text{res}(op)} \sigma_r \xrightarrow{\text{tr}_3} \sigma_f$ where $(\mathbf{m}(v) = \sigma_r(r))$:
- (1) if op is the only operation linearized in $\varrho(op)$ (i.e. there does not exist another op' , such that $op' \neq op \wedge \varrho(op') = \varrho(op)$) then

$$[\varrho(op)/\sigma][\hat{\sigma}/\sigma']\varphi$$

- (2) if there does exist one op' whose abstract specification is φ' , such that $op' \neq op \wedge \varrho(op') = \varrho(op)$ ² then, for any (ghost) state σ_g ,

$$([\varrho(op)/\sigma][\sigma_g/\sigma']\varphi \Rightarrow [\sigma_g/\sigma][\hat{\sigma}/\sigma']\varphi') \vee ([\varrho(op)/\sigma][\sigma_g/\sigma']\varphi' \Rightarrow [\sigma_g/\sigma][\hat{\sigma}/\sigma']\varphi)$$

In condition 3, to verify whether an operation op can be linearized when op is the single operation linearized in state $\varrho(op)$, we prove that the specification $[\varrho(op)/\sigma][\hat{\sigma}/\sigma']\varphi$ is respected by the execution step from the LP state $\varrho(op)$. In the substitution, the parametric pre state σ and post state σ' of the specification φ , are replaced with the LP state and its post state $\hat{\sigma}$. However, in the case of helping, many operations can be linearized in a single step (e.g. a `push` and a `pop` exemplified in Fig. 2). We handle this case by introducing *ghost states*. For example, if an event $\varrho(op) \xrightarrow{ev} \hat{\sigma}$ linearizes two operations; say $\varrho(op) = \varrho(op')$, we check that there exists an intermediate state σ_g such that $\varrho(op)$ and σ_g satisfy the specification of op and σ_g and $\hat{\sigma}$ satisfy the specification of op' , or viceversa. Intuitively this mediates and orders the linearization of the two operations.

Theorem 1. *A data structure implementation D is linearizable with respect to an abstract specification φ , if for every trace tr of the implementation, there exists a valid D -LP function with respect to the specification φ .³*

Witness States. Definition 2 paired with Theorem 1 provides us with a way of checking linearizability by constructing valid LP functions. However, since LP functions map operations to states, it could be the case that the LP state of a certain operation may precede an event from a thread other than the one whose operation is linearized. The same argument applies when multiple operations are linearized in a single step. This is true for the helping mechanism of Fig. 2. In this case, we define *witness states* as states from which one thread can make sure that it has been linearized, i.e. that a prior LP state ($\varrho(op)$) exists. For the simple case where the linearization point of an operation can be identified with its own program statement, the witness state is exactly the state before executing this statement. We prove linearizability by identifying *witness states*. In our approach, we distinguish effectful witness states where the abstract data structure is altered (like L in Fig. 1), from pure (or effect-less) witness state that leave the abstract data structure intact (like L0 in Fig. 1).

² Definition 2 is defined for at most two operations linearized in one step. It can be extended to handle the case when finitely many operations are helped in one step.

³ Theorems are proved in [25].

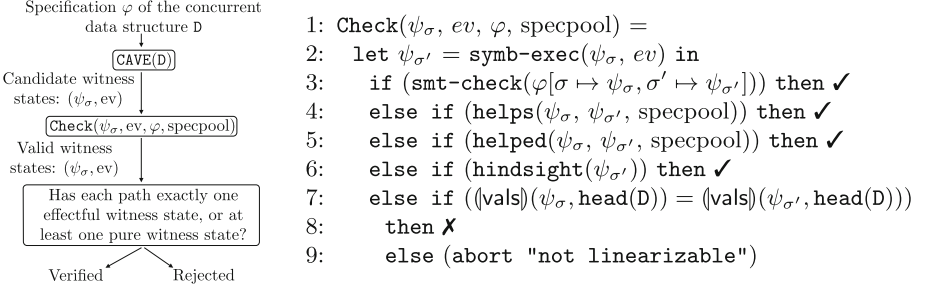


Fig. 4. General framework. (The full pseudo code is provided in [25]).

Algorithm. We present our overall verification strategy in Fig. 4. We first use heuristics derived from CAVE to identify a set of states as candidate witness states (ψ_σ), paired with the event (ev) to be executed next; such events include all the memory reads or writes. The abstract states (ψ_σ) are obtained through the symbolic execution of CAVE. The function **Check** verifies whether in a symbolic state ψ_σ the linearization of the operation can be witnessed w.r.t. its abstract specification φ . Programmers provide φ through the definition $\text{vals}(\sigma, \text{head}(\text{D}))$ which is translated into the symbolic version $(\text{vals})(\psi_\sigma, \text{head}(\text{D}))$.

Consider the pseudo-code of **Check** (given in Fig. 4). In line 2 we symbolically execute the event ev from the state ψ_σ . To check if the abstract specification φ is fulfilled, in line 3, we replace the initial and final state with the ones obtained by symbolic execution and then unroll the definitions $(\text{vals})(\psi_\sigma, \text{head}(\text{D}))$ and $(\text{vals})(\psi_{\sigma'}, \text{head}(\text{D}))$ that are mentioned in the specification of the method, and encode them using first order logic (FOL) with set theories following [17]. We feed the unrolled formulae to an SMT solver (the arguments v and return values r in the specification are also replaced with proper program variables, not shown here). Notice that satisfiability of quantifier-free formulas over sets/multisets with set union (\cup) is decidable. Concatenation ($::$) is considered as uninterpreted. If the formula is provable we have identified a witness state.

This strategy only applies to the case when linearization can be syntactically associated to instructions of the operation’s own code, i.e. LPs. Lines 4-6 deal with the cases when the linearization point might reside in operations of concurrent threads, which will be covered in the subsequent sections. If we are not able to prove φ in ψ_σ after these checks, at line 7, before reporting this state is not a candidate witness, we check that the abstract data structure did not change. We recall that we assume that only linearization events can modify the abstract state of the data structure. Hence, if the state did change, we abort the process in line 9, and report that the implementation could not be proved linearizable. After all the witness states are validated, following the strategy in [22], we use a simple data-flow analysis to verify each program path has either exactly one witness state or at least one pure witness.

Example 1. Consider the `pop` method of the stack implementation of Fig. 1. With the method delineated above we obtain a symbolic state before the program point L (we only show the shared state):

$$\text{TOP} \mapsto (\mathbf{first} : \mathbf{y}) * \mathbf{y} \mapsto \mathbf{z} * \text{lseg}_{\text{next}}(\mathbf{z}, 0)$$

This state is rendered from the successful test in the `pop` ($\text{TOP} \mapsto \mathbf{first} = \mathbf{y}$). We will consider this state to be the witness state of `pop` (i.e. ψ_σ). The assignment of `z` to `TOP->first` would then be performed. To verify whether this implementation is faithful to the Stack specification, we first symbolically execute the instruction at L to render the post state after L (i.e. $\psi_{\sigma'}$):

$$\text{TOP} \mapsto (\mathbf{first} : \mathbf{z}) * \text{lseg}_{\text{next}}(\mathbf{z}, 0) * \mathbf{y} \mapsto \mathbf{z}$$

According to the abstract specification of `pop` (Definition 3), we have to prove:

$$\langle \text{vals} \rangle (\psi_\sigma, \text{TOP} \mapsto (\mathbf{first})) = r :: \langle \text{vals} \rangle (\psi_{\sigma'}, \text{TOP} \mapsto (\mathbf{first}))$$

After unfolding, and substituting the special return symbol r with the actual return value, we obtain:

$$\text{VAL}(\mathbf{y}) :: \text{VALS}(\mathbf{z}, 0) = \text{VAL}(\mathbf{y}) :: \text{VALS}(\mathbf{z}, 0)$$

which is clearly provable. Moreover,

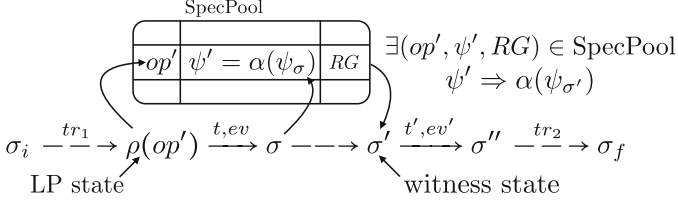
$$\langle \text{vals} \rangle (\psi_\sigma, \text{TOP} \mapsto (\mathbf{first})) = \langle \text{vals} \rangle (\psi_{\sigma'}, \text{TOP} \mapsto (\mathbf{first}))$$

holds for all the other states of `pop`. After all these verification steps, our method concludes that the program state before L is a valid effectful witness state.

Helping Verification. Consider a trace $\sigma_i \xrightarrow{\text{tr}_1} \varrho(\text{op}') \xrightarrow{t, \text{ev}} \sigma \xrightarrow{\text{tr}_2} \sigma_f$ corresponding to an execution of a data structure where op' is specified as (\mathbf{m}, t', v, r) . This trace is typical of algorithms implementing the helping mechanism. Here, the event (ev) that linearizes thread t' (the thread executing op') is taken by a concurrent thread t . A key ingredient of this pattern are the *descriptors*, which are used to keep information about ongoing invocations performed by different threads (c.f. `popdesc` in Fig. 2). A thread can acknowledge its concurrent threads through its descriptors which are used by the concurrent threads to complete helping. In our proof, a thread under verification can retrieve the specifications of the other concurrent operations through their descriptors.

To exploit such descriptors, we add to the symbolic state, a set that represents helped operations. We call this set the *Spec Pool*, and use it to keep track of the synchronization entailed through the descriptors. Operations that perform helping are assumed to affect the Spec Pool. In the Spec Pool, each helped operation is equipped with (1) the condition that must hold upon its linearization and (2) the rely actions (used by the helper thread) that linearize it.

We provide a pictorial description of the process in Fig. 5. Here we consider an event $\varrho(\text{op}') \xrightarrow{t, \text{ev}} \sigma$ from thread t which helps a concurrent operation op' . This event modifies the Spec Pool by inserting a tuple $(\text{op}', \alpha(\psi_\sigma), \text{RG})$ indicating


Fig. 5. Spec pool.

that op' has been helped at a state σ , and RG is the rely-guarantee actions extracted from this step [23], where α is a function that encodes a SL formula ψ_σ into a FOL formula. A verification step from op' (i.e. t') can observe the effects of t at state σ' (by checking a first order logical implication between $\alpha(\psi_\sigma)$ and $\alpha(\psi_{\sigma'})$). When verifying op' we also need to check that σ' can only be reached with the help of the rely action RG , absent of which σ' would be unreachable for t' . If the check is successful, σ' is considered as the witness state for op' . Figure 6 presents the definition of $\alpha(\psi)$. Simply stated, we keep the pure part of the SL formula and forget about the list segments. We encode the field-value mapping of a memory location ($E_{\text{tls}} \mapsto \rho$) into a conjunction of equations; each equation encodes the value $\rho(\text{pf})$ of a field pf on the location E . We encode E only if it is marked as a descriptor (local variables are implicitly existentially quantified).

Our approach hence reduces the problem of verifying linearizability to the following proof obligations: (a) we must check *how* an operation can be helped at the valid LP state in Fig. 5 (i.e. linearized by another thread); this corresponds to the `helps` function in Fig. 4, (b) for the thread that is helped, we must check the code that detects *whether* the operation has been helped at the valid witness state in Fig. 5; this corresponds to the `helped` function in Fig. 4, and (c) we must check that the helped operation is linearized exactly once.

For (a) we prove whether a given execution step in thread t can linearize another thread t' (with $t' \neq t$), directly following Definition 2. Let us consider how this proof works for `push` and `pop` of Fig. 2. At the statement `b` of `push` we detect the descriptor `popdesc`, representing a concurrent `pop` thread. Assuming $\text{head}(Stk) = \text{TOP}$, we check:

$$\begin{aligned} \forall \psi_{gst}, \quad & (\langle \text{vals} \rangle(\psi_{gst}, \text{TOP}) = \mathbf{v} :: \langle \text{vals} \rangle(\psi_\sigma, \text{TOP}) \Rightarrow \\ & \langle \text{vals} \rangle(\psi_{gst}, \text{TOP}) = \text{popdesc.d} :: \langle \text{vals} \rangle(\psi_{\sigma'}, \text{TOP})) \\ \\ & \alpha(P * \boxed{II \wedge \Gamma}) = II \wedge \alpha_\Gamma(\Gamma) \\ \\ \alpha_\Gamma(\Gamma' * \Gamma'') = \alpha_\Gamma(\Gamma'') \wedge & \begin{cases} \bigwedge_{\text{pf} \in \text{dom}(\rho)} \text{pf}(E) = \rho(\text{pf}) & \text{if } \Gamma' = E_{\text{tls}} \mapsto \rho \ \& \ \text{Isdesc}(E) \\ \text{true} & \text{otherwise} \end{cases} \end{aligned}$$

Fig. 6. Data abstraction from SL formula to FOL formula.

According to Definition 2, ψ_{gst} is a necessary intermediate state in the abstract data structure between the push and pop operations (it does not exist in the actual execution). The precedent is obtained from the `push`'s operation specification, with the argument substituted with the formal parameter v . The consequent is the specification of `pop`'s operation, substituting the return value for `popdesc.d`, known from the descriptor. Since the stack is not updated by the instruction \mathbf{b} , after unrolling we can prove the above formula. Both operations are linearized in this step. After verifying that `pop` is helped, we create a Spec Pool item (op', ψ', RG) , representing the result of helping at statement \mathbf{b} :

$$(\text{pop}, \exists v. \text{popdesc.g} = C \wedge \text{popdesc.d} = v, \text{popdesc} \mapsto (\mathbf{g} : W) \rightsquigarrow \text{popdesc} \mapsto (\mathbf{g} : C))$$

As stated, ψ' is the data abstraction of state σ' while RG here only shows the key rely-guarantee action ([21]), i.e., \mathbf{g} is changed from $W(\text{aiting})$ to $C(\text{ollided})$.

We prove (b) by showing that if a thread t' is linearized by another thread t , this fact is manifest through the Spec Pool. To prove that t' has been helped in a state $\psi_{\sigma'}$, we need to find a pool element (op', ψ', RG) such that the operation of t' is with the same method name to op' , and prove with an SMT solver that:

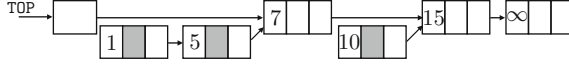
$$(\psi' \Rightarrow \alpha(\psi_{\sigma'})) \wedge \neg(\psi' \wedge \alpha(\psi_{\sigma' \setminus RG}))$$

The first conjunct implies that op' may have been linearized by another thread, and the second one ensures that this could only result from other threads' interference RG . To check that this linearization could not have been possible without the interference from another thread's helps, we compute the state $\psi_{\sigma' \setminus RG}$ by symbolically executing the method (using CAVE) to the code location of σ' dropping the rely action RG . Then the conditions recorded in the Spec Pool (the conditions hold upon helping) must contradict $\psi_{\sigma' \setminus RG}$. Consider the path reaching the statement `a2` in `pop` in Fig. 2. The conditions in the Spec Pool for `pop` (ψ') entail the data abstraction of state at `a2` (abstracted as `popdesc.g = C`). The only possible way to satisfy this assertion is by the rely RG , since originally we had `popdesc.g = W` at `a1`. We conclude that the `pop` was linearized by RG , the action made by a concurrent `push`. We also need to ensure a program path that witnesses helping must return the value of the return-field instrumented in the descriptor (e.g. `return popdesc.d` at `a2`).

We prove (c) by checking that an operation can only be helped once (e.g. helping for thread t' should be prohibited from state σ'' in the trace in Fig. 5). We leave the details in [25].

Our verification procedure maintains the Spec Pool as part of the abstract state, and calls function `Check` of Fig. 4 twice. In the first pass, we construct the Spec Pool by identifying helping scenarios; in the second pass, we exploit the Spec Pool to identify helped operations. Specifically, in Fig. 4, when a candidate state fails to fulfill the specification at line 3, we attempt to prove (a), calling function `helps` at line 4, which identifies a set of descriptors in the state that enable helping. If successful, the corresponding pool items are created (in the first pass). Otherwise, at line 5, by calling function `helped`, we attempt proof obligation \mathbf{b} to check if the operation has been helped (in the second pass).

Verification with Hindsight. This pattern is based on the Hindsight Lemma of [16]. In the interest of space we shall avoid presenting a full example like `Lazy_set` [9] (see [25]), which implements a set with an optimistic lock free *contains* operations using a linked list. As in the picture below, each node contains three fields: a value, a mark bit representing whether the item has been removed from the list (marked with grey), and a link (denoted as n) to the following node.



The fundamental invariants for this algorithm are: 1. the elements in the set are *ordered* for fast lookups through the lock-free *contains* method, 2. the elements *in* the list are all reachable from the TOP pointer, and are not marked, 3. removed elements are marked before being unlinked, and 4. the next pointer of a removed node never changes, hence it might still point to a node in the data structure, until this node is in turn removed. In the figure, the set contains the elements 7 and 15, but from the removed nodes we know that it contained the elements 1, 5 and 10 at some point in the past. A concurrent *contains* operation, which started before the elements were removed, may assume 1, 5 and 10 are still *contained*. Following [16], we shall call nodes that are reachable from TOP (including those that are marked) *backbone nodes* (e.g. 7 and 15). Conversely, nodes that cannot be reached from TOP are called *exterior nodes* (e.g. 1, 5 and 10).

Lemma 1 (Hindsight [16]). *Let tr be an execution of the set data structure presented above satisfying:*

1. *An exterior node never becomes a backbone node.*
2. *The successor of an exterior node will never change.*
3. *If the successor of a backbone node changes, the node remains a backbone node in the immediate following state.*

Then, for any states $\sigma_i = \text{tr}(i)$, $\sigma_k = \text{tr}(k)$ such that $0 \leq i \leq k < |\text{tr}|$ and for any nodes u, v, w such that $u.n \mapsto v$ is a backbone link in σ_i , and $v.n \mapsto w$ is a link (not necessarily in the backbone) in σ_k , there exists a state $\sigma_j = \text{tr}(j)$ such that $i \leq j \leq k$ and $v.n \mapsto w$ is a backbone link in σ_j .

Lemma 1 allows us to use exterior nodes and links in the current state to infer that there existed a past state in which the exterior nodes were in the backbone. Using this information we attempt to linearize the *contains* method, even if the found node is, in the current state, an exterior node. However, Lemma 1 cannot be used directly because, although an exterior link $v.n \mapsto w$ might be found in the current state, its premise, that a link $u.n \mapsto v$ was present in the backbone in a previous state, cannot be immediately established by looking at the current state in the symbolic execution. To resolve this problem we propose Theorem 2 which we exploit to automate the application of Lemma 1 in POLING.

Theorem 2. *If there is an exterior link $v.n \mapsto w$ in σ , a past state of σ in which the link is a backbone link exists provided the following conditions:*

1. *The premises of Lemma 1 hold, and*
2. *$\text{Reach}(\text{head}(\text{D}), v)$ can be proved in the sequential state σ^{seq} .*⁴

The sequential counterpart σ^{seq} of a state σ in a trace of an operation op , $\sigma_i \xrightarrow{\text{inv}(op)} \cdot \xrightarrow{\text{tr}_2} \sigma \xrightarrow{\text{tr}_3} \cdot \xrightarrow{\text{res}(op)} \sigma_f$, is obtained by execution from σ_i to σ dropping all the steps from op 's concurrent operations (executing op sequentially).

Note that the second condition ensures a temporal traversal to v (see [16]) and hence guarantees that $v.n \mapsto w$ was once a backbone link. The verification of this pattern (e.g. *contains* operation) is implemented in the `hindsight` function in line 6 in the `Check` function (Fig. 4). In this function, exploiting Theorem 2, if an exterior link $v.n \mapsto w$ is found in a candidate state σ and $\text{Reach}(\text{head}(\text{D}), v)$ holds in σ^{seq} (we compute $\psi_{\sigma^{\text{seq}}}$ by utilizing symbolic execution with an empty set of rely-guarantee actions in the implementation), we construct a past state σ_p and substitute it for σ when verifying the method's specification. If the verification succeeds, σ is a pure witness state for the verifying thread's linearization, that is, we can deduce the existence of LP state (σ_p) from witness state σ . We also customize the symbolic execution engine to verify all the three premises in Lemma 1: for each execution step $\sigma \xrightarrow{t, ev} \sigma'$, we collect exterior nodes (symbolically) in σ (ψ_σ), and verify that the step ev does not change their successors and they do not become reachable from $\text{head}(\text{D})$; we also collect backbone nodes and check, if their successors are changed by ev , then they remain reachable from $\text{head}(\text{D})$ in σ' ($\psi'_{\sigma'}$). If any of these checks fails, the `hindsight` function (Fig. 4) returns false.

$$\begin{array}{l} \psi_{Inv}^{\text{D}} : \bigvee_i \Pi_i \wedge \Gamma_i \quad v.n \mapsto w \in \sigma \quad \begin{array}{l} \text{exp}(v, w, \text{true}) = v.n \mapsto w * \text{true} \\ \text{exp}(v, w, z \mapsto z') = v = z \wedge v \mapsto z' \end{array} \\ \hline \sigma_p \equiv \bigvee_i \left(\bigvee_{S \in \Gamma_i} \Pi_i \wedge \text{exp}(v, w, S) * (\Gamma_i \setminus S) \right) \quad \begin{array}{l} \text{exp}(v, w, \text{lseg}(z, z')) = \text{lseg}(z, v) * v.n \mapsto w \\ \quad \quad \quad \quad \quad \quad \quad \quad * \text{lseg}(w, z') \end{array} \end{array}$$

Fig. 7. Hindsight application rule.

To reconstruct a past state σ_p as above, we introduce the Hindsight Lemma application rule in Fig. 7. The rule is an adaptation of May-subtraction [23]. Intuitively, `May-Subtract`(P, Q) considers the ways in which an `RGSep` assertion Q can be removed from another assertion P . Our application rule works as `May-Subtract`($\psi_{Inv}^{\text{D}}, v.n \mapsto w$) to subtract an exterior link out of data structure invariant, and return the remaining state with the link added back. The auxiliary function `exp` (expose) considers all the ways in which v can be matched to a `node` or `linked list` assertion. Notice that since the only thing that is assumed in the rule (the hypotheses) is the data structure shape invariant ψ_{Inv}^{D} (derived from `CAVE`), the resulting symbolic state is an abstraction of an actual past state. The correctness of this rule is guaranteed by the proof of Theorem 2.

⁴ `Reach` is the obvious reachability predicate over SL formulas.

Limitations. Although POLING can automatically handle concurrent data structures with non-internal linearization points, we acknowledge that it cannot verify a class of concurrent data structure whose linearization points depend on future behaviors [11]. We expect to extend POLING to support this class of programs in the future.

5 Experimental Results

We evaluated POLING⁵ on 11 examples, divided into 3 categories shown in the tables of Fig. 8. In the first table we present algorithms provable using internal linearization points. We compare the times that CAVE (version 2.1) and POLING take to verify the algorithms and notice that for all these programs POLING outperforms CAVE. This can be attributed to our usage of SMT solvers following [17] to efficiently discharge linearizability proof obligations.

The second table presents algorithms falling under the hindsight pattern. We considered set implementation algorithms that perform an optimistic `contains` (or `lookup`) operation. `Optimistic_set` [16] traverses the list optimistically (without acquiring any locks, or synchronizing with other threads) to find a node. In contrast, `Lazy_set` [9], and its variant `Vechev_CAS_set` [24] use a bit for marking nodes before deletion.

The last table includes programs that implement the helping mechanism. Conditional compare-and-swap (CCAS) [20] is a simplified version of the well known RDCSS algorithm [8]. If a CCAS method finds a (thread) descriptor in its targeting shared memory location, they attempt to help complete the operation in that descriptor before performing its own. Finally, `HSY_stack` is the full HSY stack implementation [10]. Our running time in this complex example is comparable to a rewriting technique illustrated in [6]. As expected, CAVE cannot prove all the programs in the second and third categories.

Linearization Points			Hindsight	
Program	CAVE	POLING	Program	POLING
LockCoupling_set [12]	13.28s	4.01s	Vechev_CAS_set [24]	868.44s
Vechev_DCAS_set [24]	73.90s	3.15s	Optimistic_set [16]	27.51s
2lock_queue [15]	2.91s	2.51s	Lazy_set [9]	321.78s
Treiber [19]	0.28s	0.06s		
MSqueue [15]	7.66s	1.12s	Helping	
DGLMqueue [5]	9.40s	1.47s	Program	POLING
			CCAS [20]	0.82s
			HSY_stack [10]	5.98s

Fig. 8. Experimental results.

6 Related Work and Conclusion

Related Work. Most techniques on linearizability verification (e.g., [1, 2]) are based on forward simulation arguments, and typically only work for methods

⁵ Project page: <https://www.cs.purdue.edu/homes/zhu103/poling/index.html>.

with internal linearization points local to their own code locations. To deal with external linearization points, [3] proposed a technique limited to the case where only read-only operations may have external linearization points.

Complete backward simulation strategies have been proposed in [18]. However, they are often difficult to automate. Other methods combine both forward and backward simulations, using history and/or prophecy variables [21], instrumenting the program with auxiliary state [14], or using logical relations to construct relational proofs [20]. A general and modular proof strategy is proposed in [14], that, along with lightweight instrumentation, leverages rely-guarantee reasoning to manually verify algorithms with external linearization points. In contrast, our method exploits witness states to infer a proof automatically. The helping mechanism is also considered in [6] (which does not deal with the hindsight pattern) by rewriting the implementation so that all operations have their linearization points within their rewritten code. Our technique does not rely on rewritings because the relevant witness is found within the Spec Pool.

Our technique can be considered an adaptation of [17] which verifies sequential data structures using recursive definitions on heaplets. Similar to POLING, the automata based approach [1] is also a property checking algorithm which formalizes linearizability specifications as automata, and checks the cross-product of a symbolic encoding of the program with the specification automata for safety. The main difference between POLING and [1] resides in the verification of implementations with external linearization points.

Conclusion. We describe a procedure and a tool POLING that automatically checks the linearizability of fine-grained concurrent data structures. POLING abstracts concurrent data structure into sets of locations following [17] and considers linearizability verification as a property checking technique, which are efficiently solved with an SMT solver. POLING extends prior art by incorporating important concurrent programming patterns: algorithms using *helping*, and algorithms that can be proved using the *hindsight* lemma [16]. Our experimental results provide evidence of the effectiveness of our tool.

References

1. 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 (ETAPS 2013). LNCS, vol. 7795, pp. 324–338. Springer, Heidelberg (2013)
2. Amit, D., Rinetzy, N., Reps, T., 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)
3. 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)
4. Distefano, D., O’Hearn, P.W., Yang, H.: A local shape analysis based on separation logic. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 287–302. Springer, Heidelberg (2006)

5. Doherty, S., Groves, L., Luchangco, V., Moir, M.: Formal verification of a practical lock-free queue algorithm. In: de Frutos-Escrig, D., Núñez, M. (eds.) FORTE 2004. LNCS, vol. 3235, pp. 97–114. Springer, Heidelberg (2004)
6. 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)
7. Filipovic, I., O’Hearn, P.W., Rinetzký, N., Yang, H.: Abstraction for concurrent objects. *Theor. Comput. Sci.* **411**, 252–266 (2010)
8. Harris, T.L., Fraser, K., Pratt, I.A.: A practical multi-word compare-and-swap operation. In: Malkhi, D. (ed.) DISC 2002. LNCS, vol. 2508, pp. 265–279. Springer, Heidelberg (2002)
9. Heller, S., Herlihy, M.P., Luchangco, V., Moir, M., Scherer III, W.N., Shavit, N.N.: A lazy concurrent list-based set algorithm. In: Anderson, J.H., Prencipe, G., Wattenhofer, R. (eds.) OPODIS 2005. LNCS, vol. 3974, pp. 3–16. Springer, Heidelberg (2006)
10. Hendler, D., Shavit, N., Yerushalmi, L.: A scalable lock-free stack algorithm. In: SPAA (2004)
11. Henzinger, T.A., Sezgin, A., Vafeiadis, V.: Aspect-oriented linearizability proofs. In: D’Argenio, P.R., Melgratti, H. (eds.) CONCUR 2013 – Concurrency Theory. LNCS, vol. 8052, pp. 242–256. Springer, Heidelberg (2013)
12. Herlihy, M., Shavit, N.: *The Art of Multiprocessor Programming*. MorganKaufmann, San Francisco (2008)
13. Herlihy, M., Wing, J.: Linearizability: a correctness condition for concurrent objects. *ACM TOPLAS* **12**(3), 463–492 (1990)
14. Liang, H., Feng, X.: Modular verification of linearizability with non-fixed linearization points. In: PLDI (2013)
15. Michael, M., Scott, M.: Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In: PODC (1996)
16. O’Hearn, P., Rinetzký, N., Vechev, M., Yahav, E., Yorsh, G.: Verifying linearizability with hindsight. In: PODC (2010)
17. Qiu, X., Garg, P., Stefanescu, A., Madhusudan, P.: Natural proofs for structure, data, and separation. In: PLDI (2013)
18. 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)
19. Treiber, P.: *System programming: coping with parallelism*. In: Technique report RJ 5118, IBM Almaden Research Center (1986)
20. Turon, A., Thamsborg, J., Ahmed, A., Birkedal, L., Dreyer, D.: Logical relations for fine-grained concurrency. In: POPL (2013)
21. Vafeiadis, V.: *Modular fine-grained concurrency verification*. Ph.D. thesis, University of Cambridge (2008)
22. Vafeiadis, V.: Automatically proving linearizability. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 450–464. Springer, Heidelberg (2010)
23. Vafeiadis, V.: RGSep action inference. In: Barthe, G., Hermenegildo, M. (eds.) VMCAI 2010. LNCS, vol. 5944, pp. 345–361. Springer, Heidelberg (2010)
24. Vechev, M., Yahav, E.: Deriving linerizable fine-grained concurrent objects. In: PLDI (2008)
25. Zhu, H., Petri, G., Jagannathan, S.: Poling: Smt aided linearizability proofs. Technical report, Purdue University (2015). <https://www.cs.purdue.edu/homes/zhu103/poling/tech.pdf>