

Decision Procedures for the Temporal Verification of Concurrent Lists

Alejandro Sánchez¹ and César Sánchez^{1,2}

¹ The IMDEA Software Institute, Madrid, Spain

² Spanish Council for Scientific Research (CSIC), Spain
{alejandro.sanchez, cesar.sanchez}@imdea.org

Abstract. This paper studies the problem of formally verifying temporal properties of concurrent datatypes. Concurrent datatypes are implementations of classical data abstractions, specially designed to exploit the parallelism available in multiprocessor architectures. The correctness of concurrent datatypes is essential for the overall correctness of the client software. The main difficulty to reason about concurrent datatypes is due to the simultaneous use of unstructured concurrency and dynamic memory.

The first contribution of this paper is the use of deductive temporal verification methods, in particular verification diagrams, enriched with reasoning about dynamic memory. Proofs using verification diagrams are decomposed into a finite collection of verification conditions. Our second contribution is a decision procedure mixing memory regions, pointers and list-like lists with locks, that allows the automatic verification of the generated verification conditions. We illustrate our techniques proving safety and liveness properties of lock-coupling concurrent lists.

1 Introduction

Concurrent data structures [5] are an efficient approach to exploit the parallelism of multiprocessor architectures. In contrast with sequential implementations, concurrent datatypes allow the simultaneous access of many threads to the memory representing the data value of the concurrent datatype. Concurrent data structures are hard to design, difficult to implement correctly and even more difficult to formally prove correct.

The main difficulty in reasoning about concurrent datatypes comes from the interaction of concurrency and heap manipulation. The most popular technique to reason about the structures in the heap is separation logic [10]. Leveraging on this success, some researchers [6, 13] have extended this logic to deal with concurrent programs. However, in separation logic disjoint regions are implicitly declared (hidden in the separation conjunction), which makes the reasoning about unstructured concurrency more cumbersome.

In this paper, we propose a complementary approach. We start from temporal deductive verification in the style of Manna-Pnueli [7], in particular using general verification diagrams [4, 11] to deal with concurrency. Then, inspired by

regional logic [1], we enrich the state predicate language to reason about the different regions in the heap that a program manipulates. Finally, we build decision procedures capable of checking all generated verification conditions generated during our proofs, to aid in the automation of the verification process.

Explicit regions allow the use of a classical first-order assertion language to reason about heaps, including mutation and disjointness of memory regions. Regions correspond to finite sets of object references. Unlike separation logic, the theory of sets [14] can be easily combined with other classical theories to build more powerful decision procedures. Classical theories are also amenable of integration into SMT solvers [2]. Moreover, being a classical logic one can use classical Assume-Guarantee reasoning, for example McMillan proof rules [8], for reasoning compositionally about liveness properties. In practice, using explicit regions requires the annotation and manipulation of ghost variables of type *region*, but adding these annotations is usually straightforward.

Verification diagrams can be understood as an intuitive way to abstract the specific aspect of a program which illustrates why the program satisfies a given temporal property. We propose the following verification process to show that a datatype satisfies a property expressed in linear temporal logic. First, we build the *most general client* of the datatype, parametrized by the number of threads. Then, we annotate the client and datatype with ghost fields and ghost code to support the reasoning, if necessary. Second, we build a verification diagram that serves as a witness of the proof that all possible parallel executions of the program satisfy the given temporal formula.

The proof is checked in two phases. First, we check that all executions abstracted by the diagram satisfy the property, which can be solved through a fully-automatic finite state model checking method. Second, we must check that the diagram does in fact abstract the program, which reduces to verifying a collection of verification conditions, generated from the diagram. Each concurrent datatype maintains in memory a collection of nodes and pointers with a particular layout. Based on this fact, we propose to use an assertion language whose terms include predicates in specific theories for each layout. For instance, in the case of singly linked lists, we use a decision procedure capable of reasoning about ideal lists as well as pointers representing lists in memory. In this paper, we build a decision procedure extending the theory of linked lists [9] with locks. We illustrate the whole approach to prove thread termination on a simple implementation of concurrent lists.

Most previous approaches to verifying concurrent datatypes are restricted to safety properties. In comparison, the method we propose can be used to prove *all liveness* properties, relying on the completeness of verification diagrams.

The rest of the paper is structured as follows. Section 2 presents the running example: lock-coupling concurrent lists. Section 3 briefly introduces verification diagrams and explicit regions. Section 4 describes the proposed decision procedure for concurrent lists. Finally, Section 5 shows how to apply our approach to prove termination in one case of concurrent lists. Some proofs are missing due to space limitations.

2 Concurrent Lock-Coupling Lists

The running example in this paper is the verification of lock-coupling concurrent lists [5, 13]. Lock-coupling concurrent lists are ordered lists with non-repeating elements, in which each node is protected by a lock. A thread advances through the list acquiring the lock of the node it visits. This lock is only released after the lock of the next node has been acquired. The *List* and *Node* structures, shown in Fig. 1(a) are used to maintain the data of a concurrent list.

A *List* contains one field pointing to the *Node* representing the head of the list. A *Node* consists of a value, a pointer to the next *Node* in the list and a lock. We assume that the operating system provides the operations *lock* and *unlock* to acquire and release a lock. Every list has two sentinel nodes, *Head* and *Tail*, with phantom values representing the lowest and highest possible values. For simplicity, we assume such nodes cannot be removed or modified. Concurrent Lock-Coupling Lists are used to implement sets, so they offer three operations:

- *locate*, shown in Fig. 1(d), finds an element traversing the list. This operation returns the pair consisting of the desired node and the node that precedes it in the list. If the element is not found the *Tail* node is returned as the

<pre> class List { Node list; } class Node { Value val; Node next; Lock lock; } </pre>	<pre> 1: while true do 2: e := NondetPickElem 3: nondet [4: call search(e) or call add(e) or call remove(e)] 5: end while </pre>	<pre> 1: prev, curr := locate(e) 2: if curr.val = e then 3: result := true 4: else 5: result := false 6: end if 7: curr.unlock() 8: prev.unlock() 9: return result </pre>
(a) data structures	(b) <i>decide</i>	(c) <i>search</i>
<pre> 1: prev := Head 2: prev.lock() 3: curr := prev.next 4: curr.lock() 5: while curr.val < e do 6: prev.unlock() 7: prev := curr 8: curr := curr.next 9: curr.lock() 10: end while 11: return (prev, curr) </pre>	<pre> 1: prev, curr := locate(e) 2: if curr.val ≠ e then 3: aux := new Node(e) 4: aux.next := curr 5: prev.next := aux 6: result := true 7: else 8: result := false 9: end if 10: prev.unlock() 11: curr.unlock() 12: return result </pre>	<pre> 1: prev, curr := locate(e) 2: if curr.val = e then 3: aux := curr.next 4: prev.next := aux 5: result := true 6: else 7: result := false 8: end if 9: prev.unlock() 10: curr.unlock() 11: return result </pre>
(d) <i>locate</i>	(e) <i>add</i>	(f) <i>remove</i>

Fig. 1: Data structure and algorithms for concurrent lock-coupling list

current node. A search operation, shown in Fig. 1(c), that decides whether an element is in the list can be easily extended from *locate*.

- *add*, shown in Fig. 1(e), inserts a new element in the list, using *locate* to determine the position at which the element must be inserted. The operation *add* returns *true* upon success, otherwise it returns *false*.
- *remove*, in Fig. 1(f), deletes a node from the list by redirecting the next pointer of the previous node appropriately.

Fig.1(b) shows the most general client of the concurrent-list datatype: the program *decide* that repeatedly chooses non-deterministically a method and its parameters. We construct a fair transition system $\mathcal{S}[N]$ parametrized by the total number of threads N , in which all threads run *decide*. Let ψ be the temporal formula that describes that the thread which holds the last lock in the list terminates. The verification problem is then casted as $\mathcal{S}[N] \models \psi$, for all N .

A sketch of a verification diagram is depicted in Fig. 2. We say that a thread is the rightmost owning a lock when there is no other thread owning a lock that protects a *Node* closer to the tail. Each diagram node is labeled with a predicate. This predicate captures the set of states of the transition system that the node abstracts. Edges represent transitions between states abstracted by the nodes.

Checking the proof represented by the verification diagram requires two activities. First, show that all traces of the diagram satisfy the temporal formula ψ , which can be performed by finite state model checking. Second, prove that all computations of $\mathcal{S}[N]$ are traces of the verification diagram. This process involves the verification of formulas built from several theories. For instance, considering the execution of line 5 of program *add* we should verify that the following condition holds:

$$at_add_5^{[k]} \wedge IsLast(k) \wedge \left(\begin{array}{l} r' = r \cup \langle aux^{[k]} \rangle \wedge \\ prev^{[k]}.next = aux^{[k]} \end{array} \right) \rightarrow at_add_6^{[k]} \wedge IsLast'(k)$$

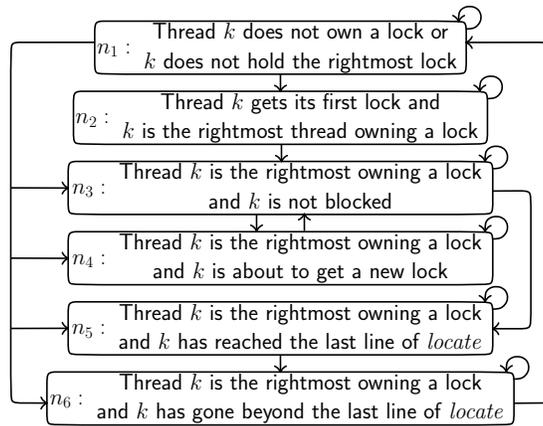


Fig. 2: Sketched verification diagram for $\mathcal{S}[N] \models \psi$

The predicate $prev'^{[k]}.next = curr^{[k]}$ is in the theory of pointers, while $r' = r \cup \langle curr^{[k]} \rangle$ is in the theory of regions. Moreover, some predicates belong to a combination of theories, like $IsLast(k)$, which among other things establishes that $List(h, x, r)$ holds. $List(h, x, r)$ expresses that in heap h , starting from pointer x , the pointers form a list of elements following the $next$ field, and that all nodes in this list form precisely the region r .

The construction of a verification diagram is a manual task, but it often follows the programmer’s intuitive explanation of why the property holds. The activity that we want to automate is checking that the diagram indeed proves the property. To accomplish this automation we must build a suitable decision procedure involving many theories, which we describe in the rest of the paper.

3 Preliminaries

We describe the temporal properties of interest in linear temporal logic, using operators such as \square (always), \diamond (eventually), \bigcirc (next) or \mathcal{U} (until) in conjunction with classical logic operations. The state predicates are built from the combination of theories that we present here.

Explicit Regions We use explicit regions to represent the manipulation of memory during the execution of the system. This reasoning is handled by extending the program code with ghost variables of type **rgn**, and ghost updates of these variables. Variables of type **rgn** represent finite sets of object references stored in the heap. Regional logic [1] provides a rich set of language constructs and assertions. However, it is enough for our purposes to use only a small fragment of regional logic. The term **emp** denotes the empty region and $\langle x \rangle$ represents the singleton region whose only object is the one referenced by x . Traditional set-like operators such as \cup , \cap and \setminus are also provided and can be applied to **rgn** variables. The assertion language allows reasoning involving mutation and separation. Given two **rgn** expressions r_1 and r_2 we can assert whether they are equal ($r_1 = r_2$), one is contained into the other ($r_1 \subseteq r_2$) or they are completely disjoint ($r_1 \# r_2$).

Verification Diagrams We sketch here the important notions from [4, 11]. Verification diagrams provide an intuitive way to abstract temporal proofs over fair transition systems (FTS). A FTS Φ is a tuple $\langle \mathcal{V}, \Theta, \mathcal{T}, \mathcal{J} \rangle$ where \mathcal{V} is a finite set of variables, Θ is an initial assertion, \mathcal{T} is a finite set of transitions and $\mathcal{J} \subseteq \mathcal{T}$ contains the fair transitions (in this paper we will not discuss strong fairness). A *state* is an interpretation of \mathcal{V} . We use \mathbf{S} to denote the set of all possible states. A transition $\tau \in \mathcal{T}$ is a function $\tau : \mathbf{S} \rightarrow 2^{\mathbf{S}}$, which is usually represented by a first-order logic formula $\rho_\tau(s, s')$ describing the relation between the values of the variables in a state s and in a successor state s' . Given a transition τ , the state predicate $En(\tau)$ denotes whether there exists a successor state s' such that $\rho_\tau(s, s')$.

A computation of Φ is an infinite sequence of states such that (a) the first state satisfies Θ ; (b) any two consecutive states satisfy ρ_τ for some $\tau \in \mathcal{T}$;

(c) for each $\tau \in \mathcal{J}$, if τ is continuously enabled after some point, then τ is taken infinitely many times. We use $\mathcal{L}(\Phi)$ to denote the set of computations of the FTS Φ . Given a formula φ , $\mathcal{L}(\varphi)$ denotes the set of sequences satisfying φ . A FTS Φ satisfies a temporal formula φ if all computations of Φ satisfy φ , i.e., $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\varphi)$.

A verification diagram (VD) $\Psi : \langle N, N_0, E, \mu, \eta, \mathcal{F}, \Delta, f \rangle$ is a formula automaton with components:

- N is a finite set of nodes.
- $N_0 \subseteq N$ is the set of initial nodes.
- $E \subseteq N \times N$ is a set of edges.
- $\mu : N \rightarrow F(V)$ is a labeling function mapping nodes to assertions over V .
- $\eta : E \rightarrow 2^\tau$ is a labeling function assigning sets of transitions to edges.
- $\mathcal{F} \subseteq 2^{E \times E}$ is an edge acceptance set of the form $\{(P_1, R_1), \dots, (P_m, R_m)\}$.
- $\Delta \subseteq \{\delta | \delta : \mathbf{S} \rightarrow \mathcal{D}\}$ is a set of ranking functions from states to a well founded domain \mathcal{D} .
- f maps nodes into propositional formulas over atomic subformulas of φ .

If $n \in N$ then we use $next(n)$ to denote the set $\{\tilde{n} \in N | (n, \tilde{n}) \in E\}$ and $\tau(n)$ for $\{\tilde{n} \in next(n) | \tau \in \eta(n, \tilde{n})\}$. For each $(P_j, R_j) \in \mathcal{F}$ and for each $n \in N$, Δ contains a ranking function $\delta_{j,n}$. An infinite sequence of nodes $\pi = n_0, n_1, \dots$ is a path if $n_0 \in N_0$ and for each $i > 0$, $(n_i, n_{i+1}) \in E$. A path π is accepted if for each pair $(P_j, R_j) \in \mathcal{F}$ some edges of R_j occur infinitely often in π or all edges that occur infinitely often in π are also in P_j . An infinite path π is fair when, for any just transition τ , if τ is enabled on all nodes that appear infinitely often in π then τ is taken infinitely often.

Given a sequence of states $\sigma = s_0, s_1, \dots$ of Φ , a path $\pi = n_0, n_1, \dots$ is a trail of σ whenever $s_i \models \mu(n_i)$ for all $i \geq 0$. An infinite sequence of states σ is a computation of Ψ whenever there exists an accepting trail of σ such that is also fair. $\mathcal{L}(\Psi)$ is the set of computations of Ψ .

A verification diagram shows that $\Phi \models \varphi$ via the inclusions $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\Psi) \subseteq \mathcal{L}(\varphi)$. The map f is used to check $\mathcal{L}(\Psi) \subseteq \mathcal{L}(\varphi)$. To show $\mathcal{L}(\Phi) \subseteq \mathcal{L}(\Psi)$ it is enough to prove the following verification conditions:

- *Initiation*: at least one initial node from N_0 satisfies the initial condition of the fair transition system Φ .
- *Consecution*: for every node $n \in N$ and transition $\tau \in \mathcal{T}$,

$$\mu(n)(s) \wedge \rho_\tau(s, s') \rightarrow \mu(next(n))(s').$$

- *Acceptance*: for each $(P_j, R_j) \in \mathcal{F}$, if $(n_1, n_2) \in P_j \setminus R_j$ then

$$\rho_\tau(s, s') \wedge \mu(n_1)(s) \wedge \mu(n_2)(s') \rightarrow \delta_{j,n_1}(s) \succeq \delta_{j,n_2}(s')$$

and if $(n_1, n_2) \notin P_j \cup R_j$ then

$$\rho_\tau(s, s') \wedge \mu(n_1)(s) \wedge \mu(n_2)(s') \rightarrow \delta_{j,n_1}(s) \succ \delta_{j,n_2}(s')$$

- *Fairness*: For each $e = (n_1, n_2) \in E$ and $\tau \in \eta(e)$:
 1. τ is guaranteed to be enabled in every $\mu(n_1)(s)$.
 2. Any τ -successor of a state satisfying $\mu(n_1)$ satisfies the label of some node in $\tau(n)$.

4 Building a Suitable Decision Procedure

The automatic check of the proof represented by a verification diagram requires decision procedures to verify the generated verification conditions. These decision procedures must deal with formulas containing terms belonging to different theories. In particular, for concurrent lists the decision procedure must reason about pointer data structures with a list layout, regions and locks. To obtain a suitable decision procedure, we extend the Theory of Linked Lists (TLL) [9], a decidable theory including reachability of list-like structures. However, this theory lacks the expressivity to describe locked lists of cells, a fundamental component in our proofs.

We begin with a brief description of the basic notation and concepts. A signature Σ is a triple (S, F, P) where S is a set of sorts, F a set of functions and P a set of predicates. If $\Sigma_1 = (S_1, F_1, P_1)$ and $\Sigma_2 = (S_2, F_2, P_2)$ are two signatures, we define their union $\Sigma_1 \cup \Sigma_2 = (S_1 \cup S_2, F_1 \cup F_2, P_1 \cup P_2)$. Similarly we say that $\Sigma_1 \subseteq \Sigma_2$ when $S_1 \subseteq S_2$, $F_1 \subseteq F_2$ and $P_1 \subseteq P_2$. If $t(\varphi)$ is a term (resp. formula), then we denote with $V_\sigma(t)$ (resp. $V_\sigma(\varphi)$) the set of variables of sort σ occurring in t (resp. φ).

A Σ -interpretation is a map assigning a value to each symbol in Σ . A Σ -structure is a Σ -interpretation over an empty set of variables. A Σ -formula over a set X of variables is satisfiable whenever it is true in some Σ -interpretation over X . Let Ω be an interpretation, \mathcal{A} a Ω -interpretation over a set V of variables, $\Sigma \subseteq \Omega$ and $U \subseteq V$. $\mathcal{A}^{\Sigma, U}$ denotes the interpretation obtained from \mathcal{A} restricting it to interpret only the symbols in Σ and the variables in U . We use \mathcal{A}^Σ to denote $\mathcal{A}^{\Sigma, \emptyset}$. A Σ -theory is a pair (Σ, \mathbf{A}) where Σ is a signature and \mathbf{A} is a class of Σ -structures. Given a theory $T = (\Sigma, \mathbf{A})$, a T -interpretation is a Σ -interpretation \mathcal{A} such that $\mathcal{A}^\Sigma \in \mathbf{A}$. Given a Σ -theory T , a Σ -formula φ over a set of variables X is T -satisfiable if it is true on a T -interpretation over X .

Formally, the theory of linked lists is defined as $\mathbf{TLL} = (\Sigma_{\mathbf{TLL}}, \mathbf{TLL})$, where

$$\Sigma_{\mathbf{TLL}} := \Sigma_{\text{cell}} \cup \Sigma_{\text{mem}} \cup \Sigma_{\text{Reachability}} \cup \Sigma_{\text{set}} \cup \Sigma_{\text{Bridge}}$$

and \mathbf{TLL} is the class of $\Sigma_{\mathbf{TLL}}$ -structures satisfying the conditions shown in Fig. 4. The sorts, functions and predicates of \mathbf{TLL} correspond to the signatures shown in Fig. 3. (Note that Figs. 4 and 3 contain an extended signature and interpretation.) Informally, Σ_{cell} models *cells*, structures containing an element (data), an addresses (pointer) and a lock owner, which represent a node in a linked list. Σ_{mem} models the memory. $\Sigma_{\text{Reachability}}$ models finite sequences of non-repeating addresses, to represent paths. Σ_{set} models sets of addresses. Finally, Σ_{Bridge} is a *bridge theory* containing auxiliary functions. The sort **thid** contains thread identifiers. The sorts **addr**, **elem** and **thid** are uninterpreted, except that $\odot : \mathbf{thid}$ is different from all others thread ids. Otherwise, $\Sigma_{\text{addr}} = (\mathbf{addr}, \emptyset, \emptyset)$, $\Sigma_{\text{elem}} = (\mathbf{elem}, \emptyset, \emptyset)$ and $\Sigma_{\text{thid}} = (\mathbf{thid}, \emptyset, \emptyset)$.

We extend \mathbf{TLL} into the theory of concurrent single linked lists $\mathbf{TLL3} := (\Sigma_{\mathbf{TLL3}}, \mathbf{TLL3})$, where $\Sigma_{\mathbf{TLL3}} = \Sigma_{\mathbf{TLL}} \cup \Sigma_{\text{setth}} \cup \{\text{lockid}, \text{lock}, \text{unlock}, \text{firstlocked}\}$. The sorts, functions and predicates of $\Sigma_{\mathbf{TLL3}}$ are described in Fig. 3. $\mathbf{TLL3}$ is the class of $\Sigma_{\mathbf{TLL3}}$ -structures satisfying the conditions listed in Fig. 4.

Signature	Sorts	Functions	Predicates
Σ_{cell}	cell elem addr thid	$error : \text{cell}$ $mkcell : \text{elem} \times \text{addr} \times \text{thid} \rightarrow \text{cell}$ $..data : \text{cell} \rightarrow \text{elem}$ $..next : \text{cell} \rightarrow \text{addr}$ $..lockid : \text{cell} \rightarrow \text{thid}$ $..lock : \text{cell} \rightarrow \text{thid} \rightarrow \text{cell}$ $..unlock : \text{cell} \rightarrow \text{cell}$	
Σ_{mem}	mem addr cell	$null : \text{addr}$ $..[-] : \text{mem} \times \text{addr} \rightarrow \text{cell}$ $upd : \text{mem} \times \text{addr} \times \text{cell} \rightarrow \text{mem}$	
$\Sigma_{\text{Reachability}}$	mem addr path	$\epsilon : \text{path}$ $[-] : \text{addr} \rightarrow \text{path}$	$append : \text{path} \times \text{path} \times \text{path}$ $reach : \text{mem} \times \text{addr} \times \text{addr} \times \text{path}$
Σ_{set}	addr set	$\emptyset : \text{set}$ $\{-\} : \text{addr} \rightarrow \text{set}$ $\cup, \cap, \setminus : \text{set} \times \text{set} \rightarrow \text{set}$	$\in : \text{addr} \times \text{set}$ $\subseteq : \text{set} \times \text{set}$
Σ_{setth}	thid setth	$\emptyset_T : \text{setth}$ $\{-\}_T : \text{thid} \rightarrow \text{setth}$ $\cup_T, \cap_T, \setminus_T : \text{setth} \times \text{setth} \rightarrow \text{setth}$	$\in_T : \text{thid} \times \text{setth}$ $\subseteq_T : \text{setth} \times \text{setth}$
Σ_{Bridge}	mem addr set path	$path2set : \text{path} \rightarrow \text{set}$ $addr2set : \text{mem} \times \text{addr} \rightarrow \text{set}$ $getp : \text{mem} \times \text{addr} \times \text{addr} \rightarrow \text{path}$ $firstlocked : \text{mem} \times \text{path} \rightarrow \text{addr}$	

Fig. 3: The signature of the TLL3 theory

Definition 1 (Finite Model Property). Let Σ be a signature, $S_0 \subseteq S$ be a set of sorts, and T be a Σ -theory. T has the finite model property with respect to S_0 if for every T -satisfiable quantifier-free Σ -formula φ there exists a T -interpretation \mathcal{A} satisfying φ such that for each sort $\sigma \in S_0$, \mathcal{A}_σ is finite.

TLL [9] enjoys the finite model property. We now show that TLL3 also has the finite model property with respect to domains `elem`, `addr` and `thid`. Hence, TLL3 is decidable because one can enumerate Σ_{TLL3} -structures up to a certain cardinality. To prove this result, we first extend the set of normalized TLL-literals.

Definition 2 (TLL3-normalized literals). A TLL3-literal is normalized if it is a flat literal of the form:

$$\begin{array}{lll}
e_1 \neq e_2 & a_1 \neq a_2 & \\
a = null & c = error & \\
c = mkcell(e, a) & c = rd(m, a) & m_2 = upd(m_1, a, c) \\
s = \{a\} & s_1 = s_2 \cup s_3 & s_1 = s_2 \setminus s_3 \\
p_1 \neq p_2 & p = [a] & p_1 = rev(p_2) \\
s = path2set(p) & append(p_1, p_2, p_3) & \neg append(p_1, p_2, p_3) \\
s = addr2set(m, a) & p = getp(m, a_1, a_2) & \\
k_1 \neq k_2 & c = mkcell(e, a, k) & a = firstlocked(m, p)
\end{array}$$

where e, e_1 and e_2 are `elem`-variables, a, a_1 and a_2 are `addr`-variables, c is a `cell`-variable, m, m_1 and m_2 are `mem`-variables, p, p_1, p_2 and p_3 are `path`-variables, and k, k_1 and k_2 are `thid`-variables.

Interpretation of sort symbols: cell, mem, path, set and setth	
Each sort σ in Σ_{TLL3} is mapped to a non-empty set \mathcal{A}_σ such that:	
(a) $\mathcal{A}_{\text{cell}} = \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{addr}} \times \mathcal{A}_{\text{thid}}$ (b) $\mathcal{A}_{\text{mem}} = \mathcal{A}_{\text{cell}}^{\mathcal{A}_{\text{addr}}}$	
(c) $\mathcal{A}_{\text{path}}$ is the set of all finite sequences of (pairwise) distinct elements of $\mathcal{A}_{\text{addr}}$ (d) \mathcal{A}_{set} is the power-set of $\mathcal{A}_{\text{addr}}$	
(e) $\mathcal{A}_{\text{setth}}$ is the power-set of $\mathcal{A}_{\text{thid}}$	
Signature	Interpretation
Σ_{cell}	<ul style="list-style-type: none"> – $mkcell(e, a, k) = \langle e, a, k \rangle$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $k \in \mathcal{A}_{\text{thid}}$ – $\langle e, a, t \rangle.data^{\mathcal{A}} = e$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $t \in \mathcal{A}_{\text{thid}}$ – $\langle e, a, t \rangle.next^{\mathcal{A}} = a$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $t \in \mathcal{A}_{\text{thid}}$ – $\langle e, a, t \rangle.lockid^{\mathcal{A}} = t$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $t \in \mathcal{A}_{\text{thid}}$ – $\langle e, a, t \rangle.lock^{\mathcal{A}}(t') = \langle e, a, t' \rangle$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $t, t' \in \mathcal{A}_{\text{thid}}$ – $\langle e, a, t \rangle.unlock^{\mathcal{A}} = \langle e, a, \emptyset \rangle$ for each $e \in \mathcal{A}_{\text{elem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $t \in \mathcal{A}_{\text{thid}}$ – $error^{\mathcal{A}}.next^{\mathcal{A}} = null^{\mathcal{A}}$
Σ_{mem}	<ul style="list-style-type: none"> – $m[a]^{\mathcal{A}} = m(a)$ for each $m \in \mathcal{A}_{\text{mem}}$ and $a \in \mathcal{A}_{\text{addr}}$ – $upd^{\mathcal{A}}(m, a, c) = m_{a \mapsto c}$ for each $m \in \mathcal{A}_{\text{mem}}$, $a \in \mathcal{A}_{\text{addr}}$ and $c \in \mathcal{A}_{\text{cell}}$ – $m^{\mathcal{A}}(null^{\mathcal{A}}) = error^{\mathcal{A}}$ for each $m \in \mathcal{A}_{\text{mem}}$
$\Sigma_{\text{Reachability}}$	<ul style="list-style-type: none"> – $\epsilon^{\mathcal{A}}$ is the empty sequence – $[i]^{\mathcal{A}}$ is the sequence containing $i \in \mathcal{A}_{\text{addr}}$ as the only element – $([i_1, \dots, i_n], [j_1, \dots, j_m], [i_1, \dots, i_n, j_1, \dots, j_m]) \in \text{append}^{\mathcal{A}}$ iff i_k and j_l are all distinct – $(m, i, j, p) \in \text{reach}^{\mathcal{A}}$ iff $i = j$ and $p = \epsilon$, or there exist addresses $i_1, \dots, i_n \in \mathcal{A}_{\text{addr}}$ such that: <ul style="list-style-type: none"> (a) $p = [i_1, \dots, i_n]$ (c) $m(i_r).next^{\mathcal{A}} = i_{r+1}$, for $1 \leq r < n$ (b) $i_1 = i$ (d) $m(i_n).next^{\mathcal{A}} = j$
Σ_{set}	The symbols \emptyset , $\{-\}$, \cup , \cap , \setminus , \in and \subseteq are interpreted according to their standard interpretation over sets of addresses.
Σ_{setth}	The symbols \emptyset_T , $\{-\}_T$, \cup_T , \cap_T , \setminus_T , \in_T and \subseteq_T are interpreted according to their standard interpretation over sets of thread identifiers.
Σ_{Bridge}	<ul style="list-style-type: none"> – $addr2set^{\mathcal{A}}(m, i) = \{j \in \mathcal{A}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t. } (m, i, j, p) \in \text{reach}\}$ – $path2set^{\mathcal{A}}(p) = \{i_1, \dots, i_n\}$ for $p = [i_1, \dots, i_n] \in \mathcal{A}_{\text{path}}$ – $getp^{\mathcal{A}}(m, i, j) = \begin{cases} p & \text{if } (m, i, j, p) \in \text{reach}^{\mathcal{A}} \\ \epsilon & \text{otherwise} \end{cases}$ – $firstlocked^{\mathcal{A}}(m, [a_1, \dots, a_n]) = \begin{cases} a_k & \text{if there is } 1 \leq k \leq n \text{ such that} \\ & \text{for all } 1 \leq j < k, m[a_j].lockid = \emptyset \\ & \text{and } m[a_k].lockid \neq \emptyset \\ null & \text{otherwise} \end{cases}$ <p style="text-align: center;">for each $m \in \mathcal{A}_{\text{mem}}$, $p \in \mathcal{A}_{\text{path}}$ and $i, j \in \mathcal{A}_{\text{addr}}$</p> <p style="text-align: center;">for each $m \in \mathcal{A}_{\text{mem}}$ and $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$</p>

Fig. 4: Characterization of a TLL3-interpretation \mathcal{A}

Lemma 1. *Deciding the TLL3-satisfiability of a quantifier-free TLL3-formula is equivalent to verifying the TLL3-satisfiability of the normalized TLL3-literals.*

Proof. By cases on the shape of all possible TLL3-literals. \square

Consider an arbitrary TLL3-interpretation \mathcal{A} satisfying a conjunction of normalized TLL3-literals Γ . We show that if there are sets $\mathcal{A}_{\text{elem}}$, $\mathcal{A}_{\text{addr}}$ and $\mathcal{A}_{\text{thid}}$

then there are finite sets $\mathcal{A}'_{\text{elem}}$, $\mathcal{A}'_{\text{addr}}$ and $\mathcal{A}'_{\text{thid}}$ with bounded cardinalities (the bound depending on Γ). $\mathcal{A}'_{\text{elem}}$, $\mathcal{A}'_{\text{addr}}$ and $\mathcal{A}'_{\text{thid}}$ can in turn be used to obtain a finite interpretation \mathcal{A}' satisfying Γ .

Lemma 2 (Finite Model Property). *Let Γ be a conjunction of normalized TLL3-literals. Let $\bar{e} = |V_{\text{elem}}(\Gamma)|$, $\bar{a} = |V_{\text{addr}}(\Gamma)|$, $\bar{m} = |V_{\text{mem}}(\Gamma)|$, $\bar{p} = |V_{\text{path}}(\Gamma)|$ and $\bar{k} = |V_{\text{thid}}(\Gamma)|$. Then the following are equivalent:*

1. Γ is TLL3-satisfiable;
2. Γ is true in a TLL3 interpretation \mathcal{A} such that

$$\begin{aligned} |\mathcal{A}_{\text{elem}}| &\leq \bar{e} + \bar{m} |\mathcal{A}_{\text{addr}}| \\ |\mathcal{A}_{\text{addr}}| &\leq \bar{a} + 1 + \bar{m} \bar{a} + \bar{p}^2 + \bar{p}^3 + \bar{m} \bar{p} \\ |\mathcal{A}_{\text{thid}}| &\leq \bar{k} + \bar{m} |\mathcal{A}_{\text{addr}}| + 1 \end{aligned}$$

Proof. (2 \rightarrow 1) is immediate. (1 \rightarrow 2), by case analysis on normalized TLL3 literals. \square

Lemma 2 justifies a brute force method to automatically check TLL3 satisfiability of normalized TLL3-literals. However, such a method is not efficient in practice. To find a more efficient decision procedure we decompose TLL3 into a combination of theories, and apply a many-sorted variant of the Nelson-Oppen combination method [12]. This method requires the theories to fulfill two conditions. First, each theory must have a decision procedure. Second, all involved theories must be stable infinite and share sorts only.

Definition 3 (stable-infiniteness). *A Σ -theory T is stably infinite if for every T -satisfiable quantifier-free Σ -formula φ there exists a T -interpretation \mathcal{A} satisfying φ whose domain is infinite.*

All theories involved in TLL [9] are stably-infinite, so the only missing theory is the one defining *firstlocked*. We define the theory T_{Base3} as follows:

$$T_{\text{Base3}} = T_{\text{addr}} \oplus T_{\text{elem}} \oplus T_{\text{cell}} \oplus T_{\text{mem}} \oplus T_{\text{path}} \oplus T_{\text{set}} \oplus T_{\text{setth}} \oplus T_{\text{thid}}$$

where T_{path} extends the theory of finite sequences of addresses with the auxiliary functions and predicates shown in Fig. 5.

The theory of finite sequences of addresses is defined by $T_{\text{fseq}} = (\Sigma_{\text{fseq}}, \text{TGen})$, where $\Sigma_{\text{fseq}} = (\{\text{addr}, \text{fseq}\}, \{\text{nil} : \text{fseq}, \text{cons} : \text{addr} \times \text{fseq} \rightarrow \text{fseq}, \text{hd} : \text{fseq} \rightarrow \text{addr}, \text{tl} : \text{fseq} \rightarrow \text{fseq}\}, \emptyset)$ and TGen as the class of multi-sorted term-generated structures that satisfy the axioms of T_{fseq} . These axioms are the standard for a theory of lists, such as distinctness, uniqueness and generation of sequences using the constructors *cons* and *nil*, as well as acyclicity of sequences (see, for example [3]). Let *PATH* be the set of axioms of T_{fseq} including all in Fig. 5. Then, we can formally define $T_{\text{path}} = (\Sigma_{\text{path}}, \text{ETGen})$ where ETGen is $\{\mathcal{A}^{\Sigma_{\text{path}}} \mid \mathcal{A}^{\Sigma_{\text{path}}} \models \text{PATH} \text{ and } \mathcal{A}^{\Sigma_{\text{fseq}}} \in \text{TGen}\}$. Next, we extend T_{Base3} defining the missing functions and predicates from $T_{\text{Reachability}}$ and Σ_{Bridge} . For example:

$$\text{ispath}(p) \wedge \text{firstmarked}(m, p, i) \leftrightarrow \text{firstlocked}(m, p) = i$$

$app : fseq \times fseq \rightarrow fseq$
$app(nil, l) = l$ $app(cons(a, l), l') = cons(a, app(l, l'))$
$fseq2set : fseq \rightarrow set$
$fseq2set(nil) = \emptyset$ $fseq2set(cons(a, l)) = \{a\} \cup fseq2set(l)$
$ispath : fseq$
$ispath(nil)$ $ispath(cons(a, nil))$ $\{a\} \not\subseteq fseq2set(l) \wedge ispath(l) \rightarrow ispath(cons(a, l))$
$last : fseq \rightarrow addr$
$last(cons(a, nil)) = a$ $l \neq nil \rightarrow last(cons(a, l)) = last(l)$
$isreachable : mem \times addr \times addr$
$isreachable(m, a, a)$ $m[a].next = a' \wedge isreachable(m, a', b) \rightarrow isreachable(m, a, b)$
$isreachablep : mem \times addr \times addr \times fseq$
$isreachablep(m, a, a, nil)$ $m[a].next = a' \wedge isreachablep(m, a', b, p) \rightarrow isreachablep(m, a, b, cons(a, p))$
$firstmarked : mem \times fseq \times addr$
$firstmarked(m, nil, null)$ $p \neq nil \wedge p = cons(j, q) \wedge m[j].lockid \neq \circlearrowleft \rightarrow firstmarked(m, p, j)$ $p \neq nil \wedge p = cons(j, q) \wedge m[j].lockid = \circlearrowleft \wedge firstmarked(m, q, i) \rightarrow firstmarked(m, p, i)$

Fig. 5: Functions, predicates and axioms of T_{path}

Let GAP be the set of axioms that define ϵ , $[_]$, $append$, $reach$, $path2set$, $addr2set$ and $getp$. We define $\widehat{TLL3} = (\Sigma_{\widehat{TLL3}}, \widehat{ETGen})$ where $\Sigma_{\widehat{TLL3}}$ is $\Sigma_{TLL} \cup \{getp, append, path2set, firstlocked\}$ and $\widehat{ETGen} := \{\mathcal{A}^{\Sigma_{\widehat{TLL3}}} | \mathcal{A}^{\Sigma_{\widehat{TLL3}}} \models GAP \text{ and } \mathcal{A}^{\Sigma_{path}} \in ETGen\}$.

Using the definitions of GAP it is easy to prove that if Γ is a set of normalized TLL3-literals, then Γ is TLL3-satisfiable iff Γ is $\widehat{TLL3}$ -satisfiable. Therefore, $\widehat{TLL3}$ can be used in place of TLL3 for satisfiability checking. We reduce $\widehat{TLL3}$ into T_{Base3} in two steps. First we do the unfolding of the definition of auxiliary functions defined in $PATH$ and GAP , getting rid of the extra functions, and obtaining a formula in $\widehat{TLL3}$ and T_{Base} . Then, we use the known reduction [9] from \widehat{TLL} into T_{Base} . All theories involved in T_{Base3} share only sorts symbols, are stably-infinite and for all of them there is a decision procedure. Hence, the multi-sorted Nelson-Oppen combination method can be applied, obtaining a decision procedure for TLL3.

We now define some auxiliary functions and predicates using TLL3, that aid in the reasoning about concurrent linked-lists (see Fig. 6). For example, predicate $List(h, a, r)$ expresses that in heap h , starting from address a there is sequence of cells all of which form region r . Function $LastMarked(h, p)$, on the other hand,

$List : \text{mem} \times \text{addr} \times \text{set}$
$List(h, a, r) \leftrightarrow null \in \text{addr2set}(h, a) \wedge r = \text{path2set}(\text{getp}(h, a, null))$
$f_a : \text{mem} \times \text{addr} \rightarrow \text{path}$
$f_a(h, n) = \begin{cases} \epsilon & \text{if } n = \text{null} \\ \text{getp}(h, h[n].\text{next}, \text{null}) & \text{if } n \neq \text{null} \end{cases}$
$LastMarked : \text{mem} \times \text{path} \rightarrow \text{addr}$
$LastMarked(m, p) = \text{firstlocked}(m, \text{rev}(p))$
$NoMarks : \text{mem} \times \text{path}$
$NoMarks(m, p) \leftrightarrow \text{firstlocked}(m, p) = \text{null}$
$SomeMark : \text{mem} \times \text{path}$
$SomeMark(m, p) \leftrightarrow \text{firstlocked}(m, p) \neq \text{null}$

Fig. 6: Auxiliary functions to reason about concurrent lists

returns the address of the last locked node in path p on memory h . All these functions can be used in verification conditions. Then, using the equivalences in Fig. 6 the predicates are removed, generating a pure $\overline{\text{TLL3}}$ formula whose satisfiability can be checked with the procedure described above.

5 Termination of Concurrent Lock-Coupling Lists

In this section we show the proof of a simple liveness property of concurrent lock-coupling lists: termination of the leading thread.

To aid in the verification of this property we annotate the code in Fig. 1 with ghost fields and ghost updates, as shown in Fig. 7, where the boxes represent the annotations introduced. The predicate $c.\text{lockid} = \emptyset$ denotes that the lock of list node c is not taken. The predicate $c.\text{lockid} = k$ establishes that the lock at list node c is owned by thread k . We enrich $List$ objects with a ghost field r of type region that keeps track of all the nodes in the list. The code for add and $remove$ is extended with ghost updates to maintain r .

T_k denotes thread k . We want to prove that if a thread has acquired a lock at node n and no other thread holds a lock ahead of n , then thread k eventually terminates. The predicate $at_add_n^{[k]}$ means that thread k is executing line n of program add . Similarly, $at_add_{n_1, \dots, n_m}^{[k]}$ is a short for thread k is running some of the lines n_1, \dots, n_m of program add . To reduce notation, $\tau_{a_n}^{[k]}$, $\tau_{r_n}^{[k]}$ and $\tau_{l_n}^{[k]}$ denote $\tau_{add_n}^{[k]}$, $\tau_{remove_n}^{[k]}$ and $\tau_{locate_n}^{[k]}$ respectively. The instance of a local variable v in thread k is represented by $v^{[k]}$. We define $DisjList$ as an extension of $List$ enriching it with the property that new nodes created during insertion are all disjoint one from each other, including all nodes that are already part of the list:

$$\begin{aligned}
DisjList(h, a, r) \hat{=} & List(h, a, r) \wedge \forall j : T_{ID}.at_a_{4,5}^{[j]} \rightarrow \langle aux^{[j]} \rangle \# r \wedge \\
& \forall i, j : T_{ID}.i \neq j \wedge at_a_{4,5}^{[i]} \wedge at_a_{4,5}^{[j]} \rightarrow \langle aux^{[i]} \rangle \# \langle aux^{[j]} \rangle \# r
\end{aligned}$$

We now define the following auxiliary predicate:

$$\begin{aligned}
IsLast(k) \hat{=} & DisjList(h, l.list, l.r) \wedge SomeMark(h, getp(h, l.list, null)) \\
& \wedge LastMarked(h, getp(h, l.list, null)) = a \quad \wedge \quad h[a].lockid = k
\end{aligned}$$

The formula $IsLast(k)$ identifies whether T_k is the thread owning the last lock in the list (i.e., the closest node towards the end of the list). Using these predicates we define the parametrized temporal formula we want to verify as:

$$\psi(k) \hat{=} \square \left(at_locate_{3..10}^{[k]} \wedge IsLast(k) \rightarrow IsLast(k) \mathcal{U} at_locate_{11}^{[k]} \right)$$

This temporal formula states that if thread k is running *locate* and it owns the last locked node in the list, then thread T_k will still own the last locked node until T_k reaches the last line of *locate*. Reachability of the last line of *locate* implies termination of the invocation to the concurrent datatype because *locate* is the only program containing potentially blocking operations.

We proceed with the construction of a verification diagram that proves the parallel execution of all threads guarantee the satisfaction of formula $\psi(k)$. Given N , we build the transitions system $\mathcal{S}[N]$, in which threads T_1, \dots, T_N run in parallel the program *decide* and show that $\mathcal{S}[N] \models \psi(k)$. The verification diagram

```

class List {
  Node list;
  rgn r;
}

class Node {
  Value val;
  Node next;
  Lock lock;
}

```

(a) data structure

```

1: prev := Head
2: prev.lock()
3: curr := prev.next
4: curr.lock()
5: while curr.val < e do
6:   prev.unlock()
7:   prev := curr
8:   curr := curr.next
9:   curr.lock()
10: end while
11: return (prev, curr)

```

(b) locate

```

1: prev, curr := locate(e)
2: if curr.val ≠ e then
3:   aux := new Node(e)
4:   aux.next := curr
5:   prev.next := aux
6:   l.r := l.r ∪ {aux}
7: else
8:   result := false
9: end if
10: prev.unlock()
11: curr.unlock()
12: return result

```

(c) add

```

1: prev, curr := locate(e)
2: if curr.val = e then
3:   aux := curr.next
4:   prev.next := aux
5:   l.r := l.r - {curr}
6: else
7:   result := false
8: end if
9: prev.unlock()
10: curr.unlock()
11: return result

```

(d) remove

Fig. 7: Concurrent lock-coupling list extended with ghost fields

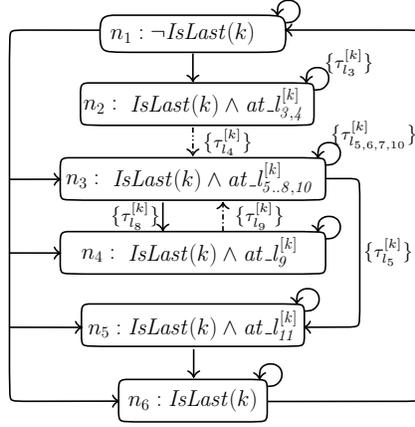


Fig. 8: Verification diagram Ψ for $\parallel_{j < N} T_j \models \psi(k)$

is depicted in Fig. 8. Dashed arrows in the diagram denote transitions that strictly decrement the ranking function δ . Formally, the verification diagram is:

- $N_0 = \{n_1\}$
- $\mathcal{F} = \{(P, R)\}$ where
 - $P = \{(n_3, n_4), (n_3, n_5), (n_5, n_6), (n_6, n_1)\} \cup \{(n_1, n_j) | j \in 2..6\} \cup \{(n_j, n_j) | j \in 1..6\}$
 - $R = \emptyset$
- $\delta(n, s) = \begin{cases} \{a \mid a \in \text{dom}(h)\} & n = n_1, n_2 \\ \text{path2set}(f_a(h, \text{LastMarked}(h, \text{getp}(h, \text{prev}^{[k]}, \text{null})))) & \text{otherwise} \end{cases}$
- $f(n) = \begin{cases} \emptyset & \text{if } n = n_1, n_6 \\ \text{at_locate}_{5..8,10}^{[k]} & \text{if } n = n_3 \\ \text{at_locate}_{11}^{[k]} & \text{if } n = n_5 \\ \text{at_locate}_{3,4}^{[k]} & \text{if } n = n_2 \\ \text{at_locate}_9^{[k]} & \text{if } n = n_4 \end{cases}$

We can now describe the verification conditions:

initialization Trivial, since in the initial state $l.\text{list}$ forms an empty list, and consequently $\neg \text{IsLast}(k)$.

consecution We will show, for illustration purposes, transition $\tau_{l_9}^{[j]}$ on node n_2 with $j \neq k$. The verification condition is:

$$\left(\begin{array}{c} \overbrace{\text{IsLast}(k)}^{\text{TLL3}} \wedge \overbrace{j \neq k}^{T_{\text{thid}}} \wedge \\ \text{at_l}_{3,4}^{[k]} \wedge \text{at_l}_9^{[j]} \wedge \\ \underbrace{\text{curr}^{[j]}.lockid = \emptyset}_{\text{TLL3}} \end{array} \right) \wedge \overbrace{\text{curr}^{[j]}.lock(j)}^{\text{TLL3}} \rightarrow \left(\begin{array}{c} \overbrace{\text{IsLast}(k') \wedge \text{at}'_{l_{3,4}}^{[k']}}^{\text{TLL3}} \wedge \\ \text{at}'_{l_{10}^{[j']}} \wedge \text{pres}(\mathcal{V} - \text{curr}^{[j]}) \wedge \\ \underbrace{\text{curr}'^{[j']}.lockid = j'}_{\text{TLL3}} \end{array} \right)$$

where pres is the predicate denoting variable preservation. Note that all fragments of such verification condition belong to theories for which we have

already defined a decision procedure, including propositional logic for the (finite) locations of the program counters.

acceptance The ranking function δ maps, at a given state, the set of list nodes accessible from the last node with an owned lock. This set remains identical for all transitions except $\tau_{l_A}^{[k]}$ and $\tau_{l_g}^{[k]}$, for which the set decrements (in the inclusion order on sets). The decision procedure presented in Section 4 proves this automatically (using \subset operation and equality over sets of addresses).

fairness Only two conditions must be verified. First, all transitions labeling an edge are enabled since the only potentially blocking operation is $\tau_{l_g}^{[k]}$ and $IsLast(k)$ implies that $\tau_{l_g}^{[k]}$ is enabled. Second, for all nodes and labelled edges, starting from a state that satisfies the predicate of the incoming node satisfies the predicate of the outgoing node via taking the transition. Sequential progress of thread k is guaranteed by fairness, since all idling transitions for thread k are in fact a diagram idiom to represent the expansion of such nodes to a sequence of nodes with a single program position on each node.

satisfaction $\mathcal{L}(\Psi) \subseteq \mathcal{L}(\psi(k))$ is automatically checkable via a finite LTL model-checking problem.

6 Conclusion

We have presented a method for the verification of temporal properties (safety and liveness) of an imperative implementation of concurrent lists. The verification is performed using verification diagrams – a complete method to prove temporal properties of reactive systems – and explicit reasoning of memory regions. The verification process usually requires the aid of ghost variables. Checking a proof is reduced to proving a finite number of verification conditions, which requires decision procedures in the appropriate theories, including regions, pointers, locks and specific theories for memory layouts, in this case single linked-lists. This paper also presents a decision procedure built as a combination of theories.

There are some key differences with other approaches in the literature. Building on the success of separation logic in proving sequential programs, the most popular approach has been extending separation logic to concurrent programs. These extensions require adapting techniques like rely-guarantee that cannot be directly used with separation logic. Our decision to use explicit regions (finite sets of addresses) allows the direct use of classical techniques like assume-guarantee and the combination of decision procedures. Furthermore, in concurrent separation logic, it is critical to describe memory footprints of sections of code. This description becomes very cumbersome when the code is not organized in mutual exclusion regions, as in fine-grain synchronization algorithms. Moreover, the integration into SMT solvers is quite straightforward with classical logics, but it is still an open question with separation logic.

The technique we propose can be seen as a method to separate the reasoning about concurrency (with verification diagrams) from the reasoning about the memory (with decision procedures). The former is independent of the data structure under consideration. We are currently extending our approach to the

verification of other pointer-based concurrent data structures like skip-lists or concurrent hash maps. Again, the sharing of these data structures makes it very hard to reason using separation logic. For our approach, these extensions will require the design of suitable decision procedures. Future work also includes building a generic VCgen for verification diagrams, implementing an ad-hoc version of the decision procedure described here, and later integrating this decision procedure into state-of-the-art SMT solvers.

Acknowledgment

We are grateful to the anonymous reviewers for their detailed comments and suggestions.

References

1. Banerjee, A., Naumann, D.A., Rosenberg, S.: Regional logic for local reasoning about global invariants. In: Proc. of Europ. Conf. Object-Oriented Programming (ECOOP'08). LNCS, vol. 5142, pp. 387–411. Springer (2008)
2. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Handbook of Satisfiability, chap. Satisfiability Modulo Theories. IOS Press (2008)
3. Bradley, A.R., Manna, Z.: The Calculus of Computation. Springer (2007)
4. Browne, A., Manna, Z., Sipma, H.B.: Generalized verification diagrams. In: Proc. of 15th Conf. on the Foundations of Software Theory and Theoretical Computer Science (FSTTCS'95). LNCS, vol. 1206, pp. 484–498. Springer (1995)
5. Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming. Morgan-Kaufmann (2008)
6. Hobor, A., Appel, A.W., Nardelli, F.Z.: Oracle semantics for concurrent separation logic. In: Proc. of European Symposium on Programming (ESOP'08). LNCS, vol. 4960, pp. 353–367. Springer (2008)
7. Manna, Z., Pnueli, A.: Temporal Verification of Reactive Systems. Springer (1995)
8. McMillan, K.L.: Circular compositional reasoning about liveness. In: Proc. of CHARME'99. LNCS, vol. 1703, pp. 342–345. Springer (1999)
9. Ranise, S., Zarba, C.G.: A theory of singly-linked lists and its extensible decision procedure. In: Proc. of Software Engineering and Formal Methods (SEFM'06). IEEE Computer Society Press (2006)
10. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Proc. Logic in Computer Science (LICS'02). pp. 55–74. IEEE Computer Society Press (2002)
11. Sipma, H.B.: Diagram-Based Verification of Discrete, Real-Time and Hybrid Systems. Ph.D. thesis, Stanford University (1999)
12. Tinelli, C., Zarba, C.G.: Combining decision procedures for sorted theories. In: Proc. Logic in Artificial Intelligence (JELIA'04). LNCS, vol. 3229, pp. 641–653. Springer (2004)
13. Vafeiadis, V., Herlihy, M., Hoare, T., Shapiro, M.: Proving correctness of highly-concurrent linearisable objects. In: Principles & Practice of Parallel Programming (PPOPP'06). pp. 129–136. ACM (2006)
14. Wies, T., Piskac, R., Kuncak, V.: Combining theories with shared set operations. In: Proc. of Frontiers of Combining Systems (FroCoS'09). LNCS, vol. 5749, pp. 366–382. Springer (2009)