

A Theory of Skiplists with Applications to the Verification of Concurrent Datatypes^{*}

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 presents a theory of skiplists with a decidable satisfiability problem, and shows its applications to the verification of concurrent skiplist implementations. A skiplist is a data structure used to implement sets by maintaining several ordered singly-linked lists in memory, with a performance comparable to balanced binary trees. We define a theory capable of expressing the memory layout of a skiplist and show a decision procedure for the satisfiability problem of this theory. We illustrate the application of our decision procedure to the temporal verification of an implementation of concurrent lock-coupling skiplists. Concurrent lock-coupling skiplists are a particular version of skiplists where every node contains a lock at each possible level, reducing granularity of mutual exclusion sections.

The first contribution of this paper is the theory TSL_K . TSL_K is a decidable theory capable of reasoning about list reachability, locks, ordered lists, and sublists of ordered lists. The second contribution is a proof that TSL_K enjoys a finite model property and thus it is decidable. Finally, we show how to reduce the satisfiability problem of quantifier-free TSL_K formulas to a combination of theories for which a many-sorted version of Nelson-Oppen can be applied.

1 Introduction

A skiplist [?] is a data structure that implements sets, maintaining several sorted singly-linked lists in memory. Skiplists are structured in multiple levels, where each level consists of a single linked list. The skiplist property establishes that the list at level $i + 1$ is a sublist of the list at level i . Each node in a skiplist stores a value and at least the pointer corresponding to the lowest level list. Some nodes also contain pointers at higher levels, pointing to the next element present at that level. The advantage of skiplists is that they are simpler and more efficient to implement than search trees, and search is still (probabilistically) logarithmic.

^{*} This work was funded in part by the EU project FET IST-231620 *HATS*, MICINN project TIN-2008-05624 *DOVES*, CAM project S2009TIC-1465 *PROMETIDOS*, and by the COST Action IC0901 *Rich ModelToolkit-An Infrastructure for Reliable Computer Systems*.




Fig. 1. A skiplist with 4 levels

Consider the skiplist shown in Fig. 1. Contrary to single-linked lists implementations, higher-level pointers allow to *skip* many elements during the search. A search is performed from left to right in a top down fashion, progressing as much as possible in a level before descending. For instance, in Fig. 1 a search for value 88 starts at level 3 of node *head*. From *head* the pointer at level 3 reaches *tail* with value $+\infty$, which is greater than 88. Hence the search algorithm moves down one level at *head* to level 2. The successor at level 2 contains value 22, which is smaller than 88, so the search continues at level 2 until a node containing a greater value is found. At that moment, the search moves down one further level again. The expected logarithmic search follows from the probability of any given node occurs at a certain level decreasing by $1/2$ as a level increases (see [?] for an analysis of the running time of skiplists).

We are interested in the formal verification of implementations of skiplists, in particular in temporal verification (liveness and safety properties) of sequential and concurrent implementations. This verification activity requires to deal with unbounded mutable data. One popular approach to verification of heap programs is Separation Logic [?]. Skiplists, however, are problematic for separation-like approaches due to the aliasing and memory sharing between nodes at different levels. Based on the success of separation logic some researchers have extended this logic to deal with concurrent programs [?,?], but concurrent datatypes follow a programming style in which the activities of concurrent threads are not structured according to critical regions with memory footprints. In these approaches based on Separation Logic memory regions are implicitly declared (hidden in the separation conjunction), which makes the reasoning about unstructured concurrency more cumbersome.

Most of the work in formal verification of pointer programs follows program logics in the Hoare tradition, either using separation logic or with specialized logics to deal with the heap and pointer structures [?,?,?]. However, extending these logics to deal with concurrent programs is hard, and though some success has been accomplished it is still an open area of research, particularly for liveness.

Continuing our previous work [?] we follow a complementary approach. We start from temporal deductive verification in the style of Manna-Pnueli [?], in particular using general verification diagrams [?,?] to deal with concurrency. This style of reasoning allows a clean separation in a proof between the temporal part (why the interleavings of actions that a set of threads can perform satisfy a certain property) with the underlying data being manipulated. A veri-

fication diagram decomposes a formal proof into a finite collection of verification conditions (VC), each of which corresponds to the effect that a small step in the program has in the data. To automatize the process of checking the proof represented by a verification diagram it is necessary to use decision procedures for the kind of data structures manipulated. This paper studies the automatic verification of VCs for the case of skiplists.

Logics like $[?,?,?]$ are very powerful to describe pointer structures, but they require the use of quantifiers to reach their expressive power. Hence, these logics preclude a combination a-la Nelson-Oppen [?] or BAPA [?] with other aspects of the program state. Instead, our solution starts from a quantifier-free theory of single-linked lists [?], and extends it in a non trivial way with order and sublists of ordered lists. The logic obtained can express skiplist-like properties without using quantifiers, allowing the combination with other theories. Proofs for an unbounded number of threads are achieved by parameterizing verification diagrams, splitting cases for interesting threads and producing a single verification condition to generalize the remaining cases. However, in this paper we mainly focus in the decision procedure. Since we want to verify concurrent lock-based implementations we extend the basic theory with locks, lock ownership, and sets of locks (and in general stores of locks). The decision procedure that we present here supports the manipulation of explicit regions, as in regional logic [?] equipped with *masked regions*, which enables reasoning about disjoint portions of the same memory cell. We use masked regions to “separate” different levels of the same skiplist node.

We call our theory TSL_K , that allows to reason about skiplists of height at most K . To illustrate the use of this theory, we sketch the proof of termination of every invocation of an implementation of a lock-coupling concurrent skiplist.

The rest of the paper is structured as follows. Section 2 presents lock-coupling concurrent skiplists. Section 3 introduces TSL_K . Section 4 shows that TSL_K is decidable by proving a finite model property theorem, and describes how to construct a more efficient decision procedure using the many-sorted Nelson-Oppen combination method. Finally, Section 5 concludes the paper. Some proofs are missing due to space limitation.

2 Fine-Grained Concurrent Lock-Coupling Skiplists

In this section we present a simple concurrent implementation of skiplists that uses lock-coupling [?] to acquire and release locks. This implementation can be seen as an extension of concurrent lock-coupling lists $[?,?]$ to multiple layers of pointers. This algorithm imposes a locking discipline, consisting of acquiring locks as the search progresses, and releasing a node’s lock only after the lock of the next node in the search process has been acquired. A naïve implementation of this solution would equip each node with a single lock, allowing multiple threads to access simultaneously different nodes in the list, but protecting concurrent accesses to two different fields of the same node. The performance can be improved by carefully allowing multiple threads to simultaneously access the

same node at different levels. We study here an implementation of this faster solution in which each node is equipped with a different lock at each level. At execution time a thread uses locks to protect the access to only some fields of a given node. A precise reasoning framework needs to capture those portions of the memory protected by a set of locks, which may include only *parts* of a node. Approaches based on strict separation (separation logic [?] or regional logic [?]) do not provide the fine grain needed to reason about individual fields of shared objects. Here, we introduce the concept of *masked regions* to describe regions and the fields within. A masked region consists of a set of pairs formed by a region ($Node$ cell) and a field (a skiplist level): $\mathbf{mrgn} \triangleq 2^{Node \times \mathbb{N}}$. We call the field a mask, since it identifies which part of the object is relevant. For example, in Fig. 2 the region within dots represents the area of the memory that thread j is protecting. This portion of the memory is described by the masked region $\{(n_2, 2), (n_5, 2), (n_2, 1), (n_4, 1), (n_3, 0), (n_4, 0)\}$. As with regional logic, an empty set intersection denotes separation. In masked regions two memory nodes at different levels do not overlap. This notion is similar to data-groups [?].

Fig. 3(a) contains the pseudo-code declaration of the $Node$ and $SkipList$ classes. Throughout the paper we use $//\textcircled{c}$ to denote ghost code added for verification purposes. Note that the structure is parametrized by a value K , which determines the maximum possible level of any node in the modeled skiplist. The fields val and key in the class $Node$ contains the value and the key of the element used to order them. Then, we can store key-value pairs, or use the skiplist as a set of arbitrary elements as long as the key can be used to compare. The $next$ array stores the pointers to the next nodes at each of the possible K different levels of the skiplist. Finally, the $lock$ array keeps the locks, one for each level, protecting the access to the corresponding $next$ field. The $SkipList$ class contains two pointer fields: $head$ and $tail$ plus a ghost variable field r . Field $head$ points to the first node of the skiplist, and $tail$ to the last one. Variable r , only used for verification purposes, keeps the (masked) region represented by all nodes in the skiplist with all their levels. In this implementation, $head$ and $tail$ are sentinel nodes, with $key = -\infty$ and $key = +\infty$, respectively. For simplicity, these nodes are not eliminated during the execution and their val field remains unchanged.

Fig. 3(b) shows the implementation of the insertion algorithm. The algorithms for searching and removing are similar, and omitted due to space limitations. The ghost variable m_r stores a masked region containing all the nodes




Fig. 2. A skiplist with the masked region given by the fields locked by thread j

<pre> class Node { Value val; Key key; Array(Node*)(K) next; Array(Node*)(K) lock; } </pre>	<pre> class SkipList { Node* head; Node* tail; //@ mrgrn r; } </pre>
---	--

(a) data structures

```

1: procedure INSERT(SkipList sl, Value newval)
2:   Vector(Node*) upd[0..K - 1]           //@ mrgrn mr := ∅
3:   lvl := randomLevel(K)
4:   Node* pred := sl.head
5:   pred.locks[K - 1].lock()               //@ mr := mr ∪ {(pred, K - 1)}
6:   Node* curr := pred.next[K - 1]
7:   curr.locks[K - 1].lock()              //@ mr := mr ∪ {(curr, K - 1)}
8:   for i := K - 1 downto 0 do
9:     if i < K - 1 then
10:      pred.locks[i].lock()               //@ mr := mr ∪ {(pred, i)}
11:      curr := pred.next[i]
12:      curr.locks[i].lock()              //@ mr := mr ∪ {(curr, i)}
13:      if i ≥ lvl then
14:        curr.locks[i + 1].unlock()       //@ mr := mr - {(curr, i + 1)}
15:        pred.locks[i + 1].unlock()       //@ mr := mr - {(pred, i + 1)}
16:      end if
17:    end if
18:    while curr.val < newval do
19:      pred.locks[i].unlock()             //@ mr := mr - {(pred, i)}
20:      pred := curr
21:      curr := pred.next[i]
22:      curr.locks[i].lock()              //@ mr := mr ∪ {(curr, i)}
23:    end while
24:    upd[i] := pred
25:  end for
26:  Bool valueWasIn := (curr.val = newval)
27:  if valueWasIn then
28:    for i := 0 to lvl do
29:      upd[i].next[i].locks[i].unlock()  //@ mr := mr - {(upd[i].next[i], i}
30:      upd[i].locks[i].unlock()          //@ mr := mr - {(upd[i], i}
31:    end for
32:  else
33:    x := CreateNode(lvl, newval)
34:    for i := 0 to lvl do
35:      x.next[i] := upd[i].next[i]
36:      upd[i].next[i] := x              //@ sl.r := sl.r ∪ {(x, i)}
37:      x.next[i].locks[i].unlock()       //@ mr := mr - {(x.next[i], i}
38:      upd[i].locks[i].unlock()          //@ mr := mr - {(upd[i], i}
39:    end for
40:  end if
41:  return ¬valueWasIn
42: end procedure

```

(b) insertion algorithm

Fig. 3. Data structure and insert algorithm for concurrent lock-coupling skiplist

and fields currently locked by the running thread. The set operations \cup and $-$ are used for the manipulation of the corresponding sets of pairs.

Let sl be a pointer to a skiplist (an instance of the class described in Fig. 3(a)). The following predicate captures whether sl points to a well-formed skiplist of height 4 or less:

$$SkipList_4(h, sl : SkipList) \hat{=} OList(h, sl, 0) \wedge \quad (1)$$

$$\left(\begin{array}{l} h[sl].tail.next[0] = null \wedge h[sl].tail.next[1] = null \\ h[sl].tail.next[2] = null \wedge h[sl].tail.next[3] = null \end{array} \right) \wedge \quad (2)$$

$$\left(\begin{array}{l} SubList(h, sl.head, sl.tail, 1, sl.head, sl.tail, 0) \wedge \\ SubList(h, sl.head, sl.tail, 2, sl.head, sl.tail, 1) \wedge \\ SubList(h, sl.head, sl.tail, 3, sl.head, sl.tail, 2) \end{array} \right) \quad (3)$$

The predicate $OList$ in (1) describes that in heap h , the pointer sl is an ordered linked-lists when repeatedly following the pointers at level 0 starting at $head$. The predicate (2) indicates all levels are *null* terminated, and (3) indicates that each level is in fact a sublist of its nearest lower level. Predicates of this kind also allow to express the effect of programs statements via first order transition relations. Consider the statement at line 36 in program *insert* shown in Fig. 3(b) on a skiplist of height 4, taken by thread with id t . This transition corresponds to a new node x at level i being connected to the skiplist. If the memory layout from pointer sl is that of a skiplist before the statement at line 36 is executed, then it is also a skiplist after the execution:

$$SkipList_4(h, sl) \wedge \varphi_{aux} \wedge \rho_{36}^{[t]}(V, V') \rightarrow SkipList_4(h', sl')$$

The effect of the statement at line 36 is represented by the first-order transition relation $\rho_{36}^{[t]}$. To ensure this property, i is required to be a valid level, and the key of the nodes that will be pointing to x must be lower than the key of node x . Moreover, the masked region of locked nodes remains unchanged. Predicate φ_{aux} contains support invariants. For simplicity, we use $prev$ for $upd^{[t]}[i]$. Then, the full verification condition is:

$$SkipList_4(h, sl) \wedge \left(\begin{array}{l} x.key = newval \quad \wedge \\ prev.key < newval \quad \wedge \\ x.next[i].key > newval \quad \wedge \\ prev.next[i] = x.next[i] \quad \wedge \\ (x, i) \notin sl.r \wedge 0 \leq i \leq 3 \end{array} \right) \wedge \left(\begin{array}{l} at_{36}[t] \quad \wedge \\ prev'.next[i] = x \quad \wedge \\ at'_{37}[t] \quad \wedge \\ h' = h \wedge sl = sl' \quad \wedge \\ x' = x \quad \dots \end{array} \right) \rightarrow SkipList_4(h', sl')$$

As usual, we use primed variables to describe the values of the variables after the transition is taken. Section 4 contains a full verification condition. This example illustrates that to be able to automatically prove VCs for the verification of skiplist manipulating algorithms, we require a theory that allows to reason about heaps, addresses, nodes, masked regions, ordered lists and sublists.

3 The Theory of Concurrent Skiplists of Height K : TSL_K

We build a decision procedure to reason about skiplist of height K combining different theories, aiming to represent pointer data structures with a skiplist layout, masked regions and locks. We extend the Theory of Concurrent Linked Lists (TLL3) [?], a decidable theory that includes reachability of concurrent list-like structures in the following way:

- each node is equipped with a *key* field, used to reason about element’s order.
- the reasoning about single level lists is extended to all the K levels.
- we extend the theory of regions with masked regions.
- lists are extended to ordered lists and sub-paths of ordered lists.

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 $\Sigma_2 = (S_2, F_2, P_2)$, we define $\Sigma_1 \cup \Sigma_2 = (S_1 \cup S_2, F_1 \cup F_2, P_1 \cup P_2)$. 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 from symbols in Σ to values. 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 a signature, \mathcal{A} an 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 skiplists of height K is defined as $\text{TSL}_K = (\Sigma_{\text{TSL}_K}, \mathbf{TSLK})$, where

$$\begin{aligned} \Sigma_{\text{TSL}_K} = & \Sigma_{\text{level}_K} \cup \Sigma_{\text{ord}} \cup \Sigma_{\text{thid}} \cup \Sigma_{\text{cell}} \cup \Sigma_{\text{mem}} \cup \Sigma_{\text{reach}} \cup \\ & \Sigma_{\text{set}} \cup \Sigma_{\text{setth}} \cup \Sigma_{\text{mrgn}} \cup \Sigma_{\text{bridge}} \end{aligned}$$

The signature of TSL_K is shown in Fig. 4. \mathbf{TSLK} is the class of Σ_{TSL_K} -structures satisfying the conditions depicted in Fig. 5. The symbols of Σ_{set} and Σ_{setth} follow their standard interpretation over sets of addresses and thread identifiers resp.

Informally, sort **addr** represents addresses; **elem** the universe of elements that can be stored in the skiplist; **ord** the ordered keys used to preserve a strict order in the skiplist; **thid** thread identifiers; **level_K** the levels of a skiplist; **cell** models *cells* representing a node in a skiplist; **mem** models the heap, mapping addresses to cells or to *null*; **path** describes finite sequences of non-repeating addresses to model non-cyclic list paths; **set** models sets of addresses – also known as regions –, while **setth** models sets of thread identifiers and **mrgn** masked regions.

Σ_{level_K} contains symbols for level identifiers $0, 1, \dots, K-1$ and their conventional order. Σ_{ord} contains two special elements $-\infty$ and ∞ for the lowest and highest values in the order \preceq . Σ_{thid} only contains, besides $=$ and \neq as for all the other theories, a special constant \emptyset to represent the absence of a thread identifier. Σ_{cell} contains the constructors and selectors for building and inspecting

cells, including *error* for incorrect dereferences. Σ_{mem} is the signature for heaps, with the usual memory access and single memory mutation functions. Σ_{set} and Σ_{setth} are theories of sets of addresses and thread ids resp. Σ_{mrgn} is the theory of masked regions. The signature Σ_{reach} contains predicates to check reachability of address using paths at different levels, while Σ_{bridge} contains auxiliary functions and predicates to manipulate and inspect paths and locks.

Signt	Sort	Functions	Predicates
Σ_{level_K}	level_K	$0, 1, \dots, K - 1 : \text{level}_K$	$< : \text{level}_K \times \text{level}_K$
Σ_{ord}	ord	$-\infty, +\infty : \text{ord}$	$\preceq : \text{ord} \times \text{ord}$
Σ_{thid}	thid	$\emptyset : \text{thid}$	
Σ_{cell}	cell elem ord addr thid	<i>error</i> : cell <i>mkcell</i> : elem \times ord \times addr ^K \times thid ^K \rightarrow cell <i>..data</i> : cell \rightarrow elem <i>..key</i> : cell \rightarrow ord <i>..next</i> [_] : cell \times level _K \rightarrow addr <i>..lockid</i> [_] : cell \times level _K \rightarrow thid <i>..lock</i> [_] : cell \times level _K \rightarrow thid \rightarrow cell <i>..unlock</i> [_] : cell \times level _K \rightarrow cell	
Σ_{mem}	mem addr cell	<i>null</i> : addr <i>..</i> [_] : mem \times addr \rightarrow cell <i>upd</i> : mem \times addr \times cell \rightarrow mem	
Σ_{reach}	mem addr path	ϵ : path <i>..</i> [_] : addr \rightarrow path	<i>append</i> : path \times path \times path <i>reach</i> _K : mem \times addr \times addr \times level _K \times path
Σ_{set}	addr set	\emptyset : set $\{-\}$: addr \rightarrow set \cup, \cap, \setminus : set \times set \rightarrow set	\in : addr \times set \subseteq : set \times set
Σ_{setth}	thid setth	\emptyset_T : setth $\{-\}_T$: thid \rightarrow setth $\cup_T, \cap_T, \setminus_T$: setth \times setth \rightarrow setth	\in_T : thid \times setth \subseteq_T : setth \times setth
Σ_{mrgn}	mrgn addr level _K	emp _{mrgn} : mrgn $\langle \cdot, \cdot \rangle_{\text{mrgn}}$: addr \times level _K \rightarrow mrgn $\cup_{\text{mrgn}}, \cap_{\text{mrgn}}, \neg_{\text{mrgn}}$: mrgn \times mrgn \rightarrow mrgn	\in_{mrgn} : addr \times level _K \times mrgn \subseteq_{mrgn} : mrgn \times mrgn $\#_{\text{mrgn}}$: mrgn \times mrgn
Σ_{bridge}	mem addr set path	<i>path2set</i> : path \rightarrow set <i>addr2set</i> _K : mem \times addr \times level _K \rightarrow set <i>getp</i> _K : mem \times addr \times addr \times level _K \rightarrow path <i>fstlock</i> _K : mem \times path \times level _K \rightarrow addr	<i>ordList</i> : mem \times path

Fig. 4. The signature of the TSL_K theory

Interpret. of sorts: addr, elem, thid, level _K , ord, cell, mem, path, set, setth and mrgn	
Each sort σ in Σ_{TSL_K} is mapped to a non-empty set \mathcal{A}_σ such that:	
(a) $\mathcal{A}_{\text{addr}}$ and $\mathcal{A}_{\text{elem}}$ are discrete sets (b) $\mathcal{A}_{\text{thid}}$ is a discrete set containing \emptyset	
(c) $\mathcal{A}_{\text{level}_K}$ is the finite collection $0, \dots, K-1$ (d) \mathcal{A}_{ord} is a total ordered set	
(e) $\mathcal{A}_{\text{cell}} = \mathcal{A}_{\text{elem}} \times \mathcal{A}_{\text{ord}} \times \mathcal{A}_{\text{addr}}^K \times \mathcal{A}_{\text{thid}}^K$ (f) $\mathcal{A}_{\text{mem}} = \mathcal{A}_{\text{cell}}^{\mathcal{A}_{\text{addr}}}$	
(g) $\mathcal{A}_{\text{path}}$ is the set of all finite sequences of (pairwise) distinct elements of $\mathcal{A}_{\text{addr}}$ (h) \mathcal{A}_{set} is the power-set of $\mathcal{A}_{\text{addr}}$	
(i) $\mathcal{A}_{\text{setth}}$ is the power-set of $\mathcal{A}_{\text{thid}}$	
(j) $\mathcal{A}_{\text{mrgn}}$ is the power-set of $\mathcal{A}_{\text{addr}} \times \mathcal{A}_{\text{level}_K}$	
Signature	Interpretation
Σ_{ord}	$x \preceq^{\mathcal{A}} y \wedge y \preceq^{\mathcal{A}} x \rightarrow x = y$ $x \preceq^{\mathcal{A}} y \vee y \preceq^{\mathcal{A}} x$ for any $x, y, z \in \mathcal{A}_{\text{ord}}$ $x \preceq^{\mathcal{A}} y \wedge y \preceq^{\mathcal{A}} z \rightarrow x \preceq^{\mathcal{A}} z$ $-\infty^{\mathcal{A}} \preceq^{\mathcal{A}} x \wedge x \preceq^{\mathcal{A}} +\infty^{\mathcal{A}}$
Σ_{cell}	<ul style="list-style-type: none"> - $mkcell^{\mathcal{A}}(e, k, \vec{a}, \vec{t}) = \langle e, k, \vec{a}, \vec{t} \rangle$ - $error^{\mathcal{A}}.next^{\mathcal{A}} = null^{\mathcal{A}}$ - $\langle e, k, \vec{a}, \vec{t} \rangle.data^{\mathcal{A}} = e$ - $\langle e, k, \vec{a}, \vec{t} \rangle.key^{\mathcal{A}} = k$ - $\langle e, k, \vec{a}, \vec{t} \rangle.next^{\mathcal{A}}[j] = a_j$ - $\langle e, k, \vec{a}, \vec{t} \rangle.lockid^{\mathcal{A}}[j] = t_j$ - $\langle e, k, \vec{a}, \dots, t_{j-1}, t_j, t_{j+1}, \dots \rangle.lock^{\mathcal{A}}[j](t') = \langle e, k, \vec{a}, \dots, t_{j-1}, t', t_{j+1}, \dots \rangle$ - $\langle e, k, \vec{a}, \dots, t_{j-1}, t_j, t_{j+1}, \dots \rangle.unlock^{\mathcal{A}}[j] = \langle e, k, \vec{a}, \dots, t_{j-1}, \emptyset, t_{j+1}, \dots \rangle$ <p>for each $e \in \mathcal{A}_{\text{elem}}, k \in \mathcal{A}_{\text{ord}}, t_0, \dots, t_j, t_{j+1}, t_{j-1}, t' \in \mathcal{A}_{\text{thid}},$ $\vec{a} \in \mathcal{A}_{\text{addr}}^K, \vec{t} \in \mathcal{A}_{\text{thid}}^K$ and $j \in \mathcal{A}_{\text{level}_K}$</p>
Σ_{mem}	$m[a]^{\mathcal{A}} = m(a)$ $upd^{\mathcal{A}}(m, a, c) = m_{a \rightarrow c}$ $m^{\mathcal{A}}(null^{\mathcal{A}}) = error^{\mathcal{A}}$ for each $m \in \mathcal{A}_{\text{mem}}, a \in \mathcal{A}_{\text{addr}}$ and $c \in \mathcal{A}_{\text{cell}}$
Σ_{reach}	<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 $i_k \neq j_l$. - $(m, a_{\text{init}}, a_{\text{end}}, l, p) \in \text{reach}_K^{\mathcal{A}}$ iff $a_{\text{init}} = a_{\text{end}}$ and $p = \epsilon$, or there exist addresses $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ such that: <ul style="list-style-type: none"> (a) $p = [a_1 \dots a_n]$ (c) $m(a_r).next^{\mathcal{A}}[l] = a_{r+1}$, for $r < n$ (b) $a_1 = a_{\text{init}}$ (d) $m(a_n).next^{\mathcal{A}}[l] = a_{\text{end}}$
Σ_{mrgn}	<ul style="list-style-type: none"> - $\mathbf{emp}_{\text{mr}}^{\mathcal{A}} = \emptyset$ - $r \cup_{\text{mr}}^{\mathcal{A}} s = r \cup s$ - $(a, j) \in_{\text{mr}}^{\mathcal{A}} r \leftrightarrow (a, j) \in r$ - $\langle a, j \rangle_{\text{mr}}^{\mathcal{A}} = \{(a, j)\}$ - $r \cap_{\text{mr}}^{\mathcal{A}} s = r \cap s$ - $r \subseteq_{\text{mr}}^{\mathcal{A}} s \leftrightarrow r \subseteq s$ - $r -_{\text{mr}}^{\mathcal{A}} s = r \setminus s$ - $r \#_{\text{mr}}^{\mathcal{A}} s \leftrightarrow r \cap_{\text{mr}}^{\mathcal{A}} s = \mathbf{emp}_{\text{mr}}^{\mathcal{A}}$ for each $a \in \mathcal{A}_{\text{addr}}, j \in \mathcal{A}_{\text{level}_K}$ and $r, s \in \mathcal{A}_{\text{mrgn}}$
Σ_{bridge}	<ul style="list-style-type: none"> - $\text{path2set}^{\mathcal{A}}(p) = \{a_1, \dots, a_n\}$ for $p = [a_1, \dots, a_n] \in \mathcal{A}_{\text{path}}$ - $\text{addr2set}_K^{\mathcal{A}}(m, a, l) = \{a' \mid \exists p \in \mathcal{A}_{\text{path}} \cdot (m, a, a', l, p) \in \text{reach}_K^{\mathcal{A}}\}$ - $\text{getp}_K^{\mathcal{A}}(m, a_{\text{init}}, a_{\text{end}}, l) = \begin{cases} p & \text{if } (m, a_{\text{init}}, a_{\text{end}}, l, p) \in \text{reach}_K^{\mathcal{A}} \\ \epsilon & \text{otherwise} \end{cases}$ <p>for each $m \in \mathcal{A}_{\text{mem}}, p \in \mathcal{A}_{\text{path}}, l \in \mathcal{A}_{\text{level}_K}$ and $a_{\text{init}}, a_{\text{end}} \in \mathcal{A}_{\text{addr}}$</p> <ul style="list-style-type: none"> - $\text{fstlock}^{\mathcal{A}}(m, [a_1 \dots a_n], l) = \begin{cases} a_k & \text{if there is } k \leq n \text{ such that} \\ & \text{for all } j < k, m[a_j].lockid[l] = \emptyset \\ & \text{and } m[a_k].lockid[l] \neq \emptyset \\ null & \text{otherwise} \end{cases}$ - $\text{ordList}^{\mathcal{A}}(m, p)$ iff $p = \epsilon$ or $p = [a]$ or $p = [a_1 \dots a_n]$ with $n \geq 2$ and $m(a_i).key^{\mathcal{A}} \preceq m(a_{i+1}).key^{\mathcal{A}}$ for all $1 \leq i < n$, for any $m \in \mathcal{A}_{\text{mem}}$

Fig. 5. Characterization of a TSL_K -interpretation \mathcal{A}

4 Decidability of TSL_K

We show that TSL_K is decidable by proving that it enjoys the finite model property with respect to its sorts, and exhibiting upper bounds for the sizes of the domains of a small interpretation of a satisfiable formula.

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.*

The fact that TSL_K has the finite model property with respect to domains `elem`, `addr`, `ord`, `levelK` and `thid`, implies that TSL_K is decidable by enumerating all possible Σ_{TSL_K} -structures up to a certain cardinality. We now define the set of normalized TSL_K -literals.

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

$$\begin{array}{lll}
 e_1 \neq e_2 & a_1 \neq a_2 & l_1 \neq l_2 \\
 a = \text{null} & c = \text{error} & c = \text{rd}(m, a) \\
 k_1 \neq k_2 & k_1 \preceq k_2 & m_2 = \text{upd}(m_1, a, c) \\
 c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}) & & \\
 s = \{a\} & s_1 = s_2 \cup s_3 & s_1 = s_2 \setminus s_3 \\
 g = \{t\}_T & g_1 = g_2 \cup_T g_3 & g_1 = g_2 \setminus_T g_3 \\
 r = \langle a, l \rangle_{\text{mr}} & r_1 = r_2 \cup_{\text{mr}} r_3 & r_1 = r_2 -_{\text{mr}} r_3 \\
 p_1 \neq p_2 & p = [a] & p_1 = \text{rev}(p_2) \\
 s = \text{path2set}(p) & \text{append}(p_1, p_2, p_3) & \neg \text{append}(p_1, p_2, p_3) \\
 s = \text{addr2set}_K(m, a, l) & p = \text{getp}_K(m, a_1, a_2, l) & \\
 t_1 \neq t_2 & a = \text{fstlock}(m, p, l) & \text{ordList}(m, p)
 \end{array}$$

where e, e_1 and e_2 are `elem`-variables; $a, a_0, a_1, a_2, \dots, a_{K-1}$ 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; s, s_1, s_2 and s_3 are `set`-variables; g, g_1, g_2 and g_3 are `setth`-variables; r, r_1, r_2 and r_3 are `mr`-variables; k, k_1 and k_2 are `ord`-variables; l, l_1 and l_2 are `levelK`-variables and $t, t_0, t_1, t_2, \dots, t_{K-1}$ are `thid`-variables.

Lemma 1. *Every TSL_K -formula is equivalent to a collection of conjunctions of normalized TSL_K -literals.*

Proof (sketch). First, transform a formula in disjunctive normal form. Then each conjunct can be normalized introducing auxiliary fresh variables when necessary.

The phase of normalizing a formula is commonly known [?] as the “variable abstraction phase”. Note that normalized literals belong to just one theory.

Consider an arbitrary TSL_K -interpretation \mathcal{A} satisfying a conjunction of normalized TSL_K -literals Γ . We show that if \mathcal{A} consists of domains $\mathcal{A}_{\text{elem}}, \mathcal{A}_{\text{addr}}, \mathcal{A}_{\text{thid}}, \mathcal{A}_{\text{level}_K}$ and \mathcal{A}_{ord} then there are finite sets $\mathcal{B}_{\text{elem}}, \mathcal{B}_{\text{addr}}, \mathcal{B}_{\text{thid}}, \mathcal{B}_{\text{level}_K}$ and \mathcal{B}_{ord} with bounded cardinalities, where the finite bound on the sizes can be computed from Γ . Such sets can in turn be used to obtain a finite interpretation \mathcal{B} satisfying Γ , since all the other sorts are bounded by the sizes of these sets.

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

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

$$\begin{aligned} |\mathcal{B}_{\text{addr}}| &\leq \bar{a} + 1 + \bar{m} \bar{a} K + \bar{p}^2 + \bar{p}^3 + (K+2)\bar{m} \bar{p} & |\mathcal{B}_{\text{elem}}| &\leq \bar{e} + \bar{m} |\mathcal{B}_{\text{addr}}| \\ |\mathcal{B}_{\text{thid}}| &\leq \bar{t} + K \bar{m} |\mathcal{B}_{\text{addr}}| + 1 & |\mathcal{B}_{\text{ord}}| &\leq \bar{o} + \bar{m} |\mathcal{B}_{\text{addr}}| \\ |\mathcal{B}_{\text{level}_K}| &\leq K \end{aligned}$$

Proof. (2 \rightarrow 1) is immediate. (1 \rightarrow 2) is proved on a case analysis over the set of normalized literals of TSL_K . \square

4.1 A combination-based decision procedure for TSL_K

Lemma 2 enables a brute force method to automatically check whether a set of normalized TSL_K -literals is satisfiable. However, such a method is not efficient in practice. We describe now how to obtain a more efficient decision procedure for TSL_K applying a many-sorted variant [?] of the Nelson-Oppen combination method [?], by combining the decision procedures for the underlying theories. This combination method requires that the theories fulfill some conditions. First, each theory must have a decision procedure. Second, two theories can only share sorts (but not functions or predicates). Third, when two theories are combined, either both theories are stable infinite or one of them is polite with respect to the underlying sorts that it shares with the other. The stable infinite condition for a theory establishes that if a formula has a model then it has a model with infinite cardinality. In our case, some theories are not stable infinite. For example, T_{level_K} is not stably infinite, T_{ord} , and T_{thid} need not be stable infinite in same instances. The observation that the condition of stable infinity may be cumbersome in the combination of theories for data structures was already made in [?] where they suggest the condition of *politeness*:

Definition 3 (Politeness). *T is polite with respect to sorts $S : \{\sigma_1 \dots \sigma_n\}$ whenever:*

- (1) *Let φ be a satisfiable formula in theory T , \mathcal{A} be one model of φ and let $|\mathcal{A}_{\sigma_1}|, \dots, |\mathcal{A}_{\sigma_n}|$ be the cardinalities of the domains of \mathcal{A} for sorts in S . For every tuple of larger cardinalities $k_1 \geq |\mathcal{A}_{\sigma_1}|, \dots, k_n \geq |\mathcal{A}_{\sigma_n}|$, there is a model \mathcal{B} of φ with $|\mathcal{B}_{\sigma_i}| = k_i$.*
- (2) *There is a computable function that for every formula φ returns an equivalent formula $(\exists \bar{v})\psi$ (where $\bar{v} = V_\psi \setminus V_\varphi$) such that, if ψ is satisfiable, then there is an interpretation \mathcal{A} with $\mathcal{A}_\sigma = [V_\sigma(\psi)]^\mathcal{A}$ for each sort σ .*

Condition (1) is called *smoothness*, and guarantees that interpretations can be enlarged as needed. Condition (2) is called *finite witnessability*, and gives a procedure to produce a model in which every element is represented by a variable.

The Finite Model Property, Lemma 2 above, guarantees that every sub-theory of TSL_K is finite witnessable since one can add as many fresh variables as the bound for the corresponding sort in the lemma. The smoothness property can be shown for:

$$T_{\text{cell}} \oplus T_{\text{mem}} \oplus T_{\text{path}} \oplus T_{\text{set}} \oplus T_{\text{setth}} \oplus T_{\text{mrgn}}$$

with respect to sorts addr , level_K , elem , ord and thid . Moreover, these theories can be combined because all of them are stably infinite. The following can also be combined: $T_{\text{level}_K} \oplus T_{\text{ord}} \oplus T_{\text{thid}}$ because they do not share any sorts, so combination is trivial. The many-sorted Nelson-Oppen method allows to combine the first collection of theories with the second. Regarding the decision procedures for each individual theory, T_{level_K} is trivial since it is just a finite set of naturals with order. For T_{ord} we can adapt a decision procedure for dense orders as the reals [?], or other appropriate theory. For T_{cell} we can use a decision procedure for recursive data structures [?]. T_{mem} is the theory of arrays [?]. T_{set} , T_{setth} and T_{mrgn} are theories of (finite) sets for which there are many decision procedures [?,?]. The remaining theories are T_{reach} and T_{bridge} . Following the approaches in [?,?] we extend a decision procedure for the theory T_{path} of finite sequences of (non-repeated) addresses with the auxiliary functions and predicates shown in Fig. 6, and combine this theory to obtain:

$$T_{\text{SLKBase}} = T_{\text{addr}} \oplus T_{\text{ord}} \oplus T_{\text{thid}} \oplus T_{\text{level}_K} \oplus T_{\text{cell}} \oplus T_{\text{mem}} \oplus T_{\text{path}} \oplus T_{\text{set}} \oplus T_{\text{setth}} \oplus T_{\text{mrgn}}$$

Using T_{path} all symbols in T_{reach} can be easily defined. 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 term-generated structures that satisfy the axioms of distinctness, uniqueness and generation of sequences using constructors, as well as acyclicity (see, for example [?]). Let Σ_{path} be Σ_{fseq} extended with the symbols of Fig. 6 and let PATH be the set of axioms of T_{fseq} including the ones in Fig. 6. Then, we can formally define $T_{\text{path}} = (\Sigma_{\text{path}}, \text{ETGen})$ where ETGen is $\{\mathcal{A}^{\Sigma_{\text{path}}} | \mathcal{A}^{\Sigma_{\text{path}}} \models \text{PATH} \text{ and } \mathcal{A}^{\Sigma_{\text{fseq}}} \in \text{TGen}\}$. Next, we extend T_{SLKBase} with definitions for translating all missing functions and predicates from Σ_{reach} and Σ_{bridge} appearing in normalized TSL_K -literals by definitions from T_{SLKBase} . Let GAP be the set of axioms that define ϵ , $[-]$, append , reach_K , path2set , getp_K , fstlock and ordList . For instance: $\text{ispath}(p) \wedge \text{ordPath}(m, p) \leftrightarrow \text{ordList}(m, p)$. We now define $\widehat{\text{TSL}}_K = (\Sigma_{\widehat{\text{TSL}}_K}, \widehat{\text{ETGen}})$ where $\Sigma_{\widehat{\text{TSL}}_K}$ is $\Sigma_{T_{\text{SLKBase}}} \cup \{\text{append}, \text{reach}_K, \text{path2set}, \text{getp}_K, \text{fstlock}, \text{ordList}\}$ and $\widehat{\text{ETGen}} := \{\mathcal{A}^{\Sigma_{\widehat{\text{TSL}}_K}} | \mathcal{A}^{\Sigma_{\widehat{\text{TSL}}_K}} \models \text{GAP} \text{ and } \mathcal{A}^{\Sigma_{T_{\text{SLKBase}}}} \in \text{ETGen}\}$.

Using the definitions of GAP it is easy to prove that if Γ is a set of normalized TSL_K -literals, then Γ is TSL_K -satisfiable iff Γ is $\widehat{\text{TSL}}_K$ -satisfiable. Therefore, $\widehat{\text{TSL}}_K$ can be used in place of TSL_K for satisfiability checking. The reduction from $\widehat{\text{TSL}}_K$ into T_{SLKBase} is performed in two steps. First, by the finite model theorem (Lemma 2), it is always possible to calculate an upper bound in the number of elements of sort addr , elem , thid , ord and level in a model (if there is one model), based on the input formula. Therefore, one can introduce one variable

$app : \text{fseq} \times \text{fseq} \rightarrow \text{fseq}$	
$app(\text{nil}, l) = l$	$app(\text{cons}(a, l), l') = \text{cons}(a, app(l, l'))$
$fseq2set : \text{fseq} \rightarrow \text{set}$	
$fseq2set(\text{nil}) = \emptyset$	$fseq2set(\text{cons}(a, l)) = \{a\} \cup fseq2set(l)$
$ispath : \text{fseq}$	
$ispath(\text{nil})$	$ispath(\text{cons}(a, \text{nil}))$
$\{a\} \not\subseteq fseq2set(l) \wedge ispath(l) \rightarrow ispath(\text{cons}(a, l))$	
$last : \text{fseq} \rightarrow \text{addr}$	
$last(\text{cons}(a, \text{nil})) = a$	
$l \neq \text{nil} \rightarrow last(\text{cons}(a, l)) = last(l)$	
$isreach_{\kappa} : \text{mem} \times \text{addr} \times \text{addr} \times \text{level}_{\kappa}$	
$isreach_{\kappa}(m, a, a, l)$	
$m[a].next[l] = a' \wedge isreach_{\kappa}(m, a', b, l) \rightarrow isreach_{\kappa}(m, a, b, l)$	
$isreachp_{\kappa} : \text{mem} \times \text{addr} \times \text{addr} \times \text{level}_{\kappa} \times \text{fseq}$	
$isreachp_{\kappa}(m, a, a, l, \text{nil})$	
$m[a].next[l] = a' \wedge isreachp(m, a', b, l, p) \rightarrow isreachp(m, a, b, l, \text{cons}(a, p))$	
$fstmark : \text{mem} \times \text{fseq} \times \text{level}_{\kappa} \times \text{addr}$	
$fstmark(m, \text{nil}, l, \text{null})$	
$p \neq \text{nil} \wedge p = \text{cons}(a, q) \wedge m[a].lockid[l] \neq \emptyset \rightarrow fstmark(m, p, l, a)$	
$p \neq \text{nil} \wedge p = \text{cons}(a, q) \wedge m[a].lockid[l] = \emptyset \wedge fstmark(m, q, l, b) \rightarrow fstmark(m, p, l, b)$	
$ordPath : \text{mem} \times \text{fseq}$	
$ordPath(h, \text{nil})$	
$\left(\begin{array}{l} h[a].next[0] = a' \wedge h[a].key \preceq h[a'].key \wedge \\ p = \text{cons}(a, q) \wedge ordPath(h, q) \end{array} \right) \rightarrow ordPath(h, p)$	

Fig. 6. Functions, predicates and axioms of T_{path}

per element of each of these sorts and unfold all definitions in *PATH* and *GAP*, by symbolic expansion, leading to terms in Σ_{fseq} , and thus, in T_{SLKBase} . This way, it is always possible to reduce a $\widehat{\text{TSLK}}$ -satisfiability problem of normalized literals into a T_{SLKBase} -satisfiability problem. Hence, using a decision procedure for T_{SLKBase} we obtain a decision procedure for $\widehat{\text{TSLK}}$, and thus, for TSLK . Notice, for instance, that the predicate $subPath : \text{path} \times \text{path}$ for ordered lists can be defined using only $path2set$ as: $subPath(p_1, p_2) \doteq path2set(p_1) \subseteq path2set(p_2)$

For space reasons, we do not provide complete specification and proofs of the temporal properties. However, in [?] is detailed an example of a termination proof over concurrent lists, which easily carries over to skiplists. For illustration purposes, we now show the full verification condition for the verification of the safety property $\square(SkipList_4(h, sl))$ when executing transition 36 of program *insert* by a thread with id t , from Section 2. For clarity, we again use $prev$ as a short for $upd^{[t]}[i^{[t]}]$, and we use the auxiliary predicate $setnext(c, d, i, x)$ that makes the cell d identical to c except that $c.next[i] = x$.

$$setnext(c, d, i, x) \doteq \left(\begin{array}{l} d.data = c.data \wedge d.key = c.key \wedge d.lock[j] = c.lock[j] \wedge \\ (i \neq j) \rightarrow d.next[j] = c.next[j] \wedge d.next[i] = x \end{array} \right)$$

The VC is $(SkipList_4(h, sl) \wedge \varphi \rightarrow SkipList_4(h', sl'))$ where φ is:

$$\left(\begin{array}{l} x^{[t]}.key = newval \quad \wedge \\ prev.key < newval \quad \wedge \\ x^{[t]}.next[i^{[t]}].key > newval \quad \wedge \\ prev.next[i^{[t]}] = x^{[t]}.next[i^{[t]}] \wedge \\ (x^{[t]}, i^{[t]}) \notin sl.r \wedge 0 \leq i^{[t]} \leq 3 \end{array} \right) \wedge \left(\begin{array}{l} at_{36}[t] \wedge at'_{37}[t] \quad \wedge \\ prev'.next[i^{[t]}] = x^{[t]} \quad \wedge \\ setnext(h[prev], newcell, i^{[t]}, x^{[t]}) \wedge \\ h' = upd(h, prev, newcell) \quad \wedge \\ sl = sl' \wedge x'^{[t]} = x^{[t]} \quad \wedge \end{array} \right)$$

5 Conclusion and Future Work

In this paper we have presented TSL_K , a theory of skiplists of height at most K , useful for automatically prove the VCs generated during the verification of concurrent skiplist implementations. TSL_K is capable of reasoning about memory, cells, pointers, masked regions and reachability, enabling ordered lists and sublists, allowing the description of the skiplist property, and the representation of memory modifications introduced by the execution of program statements.

We showed that TSL_K is decidable by proving its finite model property, and exhibiting the minimal cardinality of a model if one such model exists. Moreover, we showed how to reduce the satisfiability problem of quantifier-free TSL_K formulas to a combination of theories using the many-sorted version of Nelson-Oppen, allowing the use of well studied decision procedures. The complexity of the decision problem for TSL_K is easily shown to be NP-complete since it properly extends TLL [?].

Current work includes the translation of formulas from T_{ord} , T_{level_K} , T_{set} , T_{seth} and T_{mrgn} into BAPA [?]. In BAPA, arithmetic, sets and cardinality aids in the definition of skiplists properties. Paths can be represented as finite sequences of addresses. We are studying how to replace the recursive functions from T_{reach} and Σ_{bridge} by canonical set and list abstractions [?], which would lead to a more efficient decision procedure, essentially encoding full TSL_K formulas into BAPA. The family of theories presented in the paper is limited to skiplists of a fixed maximum height. Typical skiplist implementations fix a maximum number of levels and this can be handled with TSL_K . Inserting more than 2^{levels} elements into a skiplist may slow-down the search of a skiplist implementation but this issue affects performance and not correctness, which is the goal pursued in this paper. We are studying techniques to describe skiplists of arbitrary many levels. A promising approach consists of equipping the theory with a primitive predicate denoting that the skiplist property holds above and below a given level. Then the reasoning is restricted to the single level being modified. This approach, however, is still work in progress.

Furthermore, we are working on a direct implementation of our decision procedure, as well as its integration into existing solvers. Future work also includes the temporal verification of sequential and concurrent skiplists implementations, including one at the `java.concurrent` standard library. This can be accomplished by the design of verification diagrams that use the decision procedure presented in this paper.

References

1. Armando, A., Ranise, S., Rusinowitch, M.: A rewriting approach to satisfiability procedures. *Information and Computation* 183(2), 140–164 (2003)
2. Banerjee, A., Naumann, D.A., Rosenberg, S.: Regional logic for local reasoning about global invariants. In: *Proc. of ECOOP'08*. pp. 387–411. Springer (2008)
3. Bouajjani, A., Dragoi, C., Enea, C., Sighireanu, M.: A logic-based framework for reasoning about composite data structures. In: *CONCUR'09*. pp. 178–195 (2009)
4. Bradley, A.R., Manna, Z.: *The Calculus of Computation*. Springer-Verlag (2007)
5. Browne, A., Manna, Z., Sipma, H.B.: Generalized verification diagrams. In: *Proc. of FSTTCS'95*. LNCS, vol. 1206, pp. 484–498. Springer (1995)
6. Herlihy, M., Shavit, N.: *The Art of Multiprocessor Programming*. Morgan-Kaufmann (2008)
7. Hobor, A., Appel, A.W., Nardelli, F.Z.: Oracle semantics for concurrent separation logic. In: *Proc. of ESOP'08*. LNCS, vol. 4960, pp. 353–367. Springer (2008)
8. Kuncak, V., Nguyen, H.H., Rinard, M.C.: An algorithm for deciding BAPA: Boolean Algebra with Presburger Arithmetic. In: *CADE'05*. pp. 260–277 (2005)
9. Lahiri, S.K., Qadeer, S.: Back to the future: revisiting precise program verification using smt solvers. In: *Proc. of POPL'08*. pp. 171–182. ACM (2008)
10. Leino, K.R.M.: Data groups: Specifying the modification of extended state. In: *OOPSLA'98*. pp. 144–153. ACM (1998)
11. Manna, Z., Pnueli, A.: *Temporal Verification of Reactive Systems*. Springer (1995)
12. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.* 1(2), 245–257 (1979)
13. Oppen, D.C.: Reasoning about recursively defined data structures. *J. ACM* 27(3), 403–411 (1980)
14. Pugh, W.: Skip lists: A probabilistic alternative to balanced trees. *Commun. ACM* 33(6), 668–676 (1990)
15. Ranise, S., Ringeissen, C., Zarba, C.G.: Combining data structures with nonstably infinite theories using many-sorted logic. In: *FRODOS'05*. pp. 48–64 (2005)
16. Ranise, S., Zarba, C.G.: A theory of singly-linked lists and its extensible decision procedure. In: *Proc. of SEFM 2006*. IEEE CS Press (2006)
17. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: *Proc. of LICS'02*. pp. 55–74. IEEE CS Press (2002)
18. Sánchez, A., Sánchez, C.: Decision procedures for the temporal verification of concurrent lists. In: *Proc. of ICFEM'10*. LNCS, vol. 6447, pp. 74–89. Springer (2010)
19. Sipma, H.B.: *Diagram-Based Verification of Discrete, Real-Time and Hybrid Systems*. Ph.D. thesis, Stanford University (1999)
20. Suter, P., Dotta, M., Kuncak, V.: Decision procedures for algebraic data types with abstractions. In: *Proc. of POPL'10*. pp. 199–210. ACM (2010)
21. Tarski, A.: *A decision method for elementary algebra and geometry*. University of California Press (1951)
22. Tinelli, C., Zarba, C.G.: Combining decision procedures for sorted theories. In: *JELIA'04*. LNCS, vol. 3229, pp. 641–653. Springer (2004)
23. Vafeiadis, V.: *Modular fine-grained concurrency verification*. Ph.D. thesis, University of Cambridge (2007)
24. Yorsh, G., Rabinovich, A.M., Sagiv, M., Meyer, A., Bouajjani, A.: A logic of reachable patterns in linked data-structures. In: *FOSSACS'06*. pp. 94–110 (2006)
25. Zarba, C.G.: Combining sets with elements. In: *Verification: Theory and Practice*. LNCS, vol. 2772, pp. 762–782. Springer (2003)

A Small Model Property

A.1 Normalized Literals

We show that TSL_K also has the finite model property with respect to domains `elem`, `addr`, `thid`, `ord` and `level`. Hence, TSL_K is decidable because one can enumerate Σ_{TSL_K} -structures up to a certain cardinality. To prove this result, we first define the set of normalized TSL_K -literals.

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

$$\begin{array}{lll}
 e_1 \neq e_2 & a_1 \neq a_2 & l_1 \neq l_2 \\
 a = \text{null} & c = \text{error} & \\
 k_1 \neq k_2 & k_1 \preceq k_2 & \\
 c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}) & & \\
 c = \text{rd}(m, a) & m_2 = \text{upd}(m_1, a, c) & \\
 s = \{a\} & s_1 = s_2 \cup s_3 & s_1 = s_2 \setminus s_3 \\
 g = \{t\}_T & g_1 = g_2 \cup_T g_3 & g_1 = g_2 \setminus_T g_3 \\
 r = \langle a, l \rangle_{\text{mr}} & r_1 = r_2 \cup_{\text{mr}} r_3 & r_1 = r_2 -_{\text{mr}} r_3 \\
 p_1 \neq p_2 & p = [a] & p_1 = \text{rev}(p_2) \\
 s = \text{path2set}(p) & \text{append}(p_1, p_2, p_3) & \neg \text{append}(p_1, p_2, p_3) \\
 s = \text{addr2set}_K(m, a, l) & p = \text{getp}_K(m, a_1, a_2, l) & \\
 t_1 \neq t_2 & a = \text{fstlock}(m, p, l) & \text{ordList}(m, p)
 \end{array}$$

where e, e_1 and e_2 are `elem`-variables; $a, a_1, a_2, \dots, a_{K-1}$ 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; s, s_1, s_2 and s_3 are `set`-variables; g, g_1, g_2 and g_3 are `setth`-variables; r, r_1, r_2 and r_3 are `mr`-variables; k, k_1 and k_2 are `ord`-variables; l, l_1 and l_2 are `levelK`-variables and $t, t_1, t_2, \dots, t_{K-1}$ are `thid`-variables.

The remaining literals can be rewritten from the normalized ones using the following equivalences:

$$\begin{array}{l}
 e = c.\text{data} \quad \leftrightarrow (\exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1}) \\
 \quad \quad \quad [c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})] \\
 k = c.\text{key} \quad \leftrightarrow (\exists_{\text{elem}} e \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1}) \\
 \quad \quad \quad [c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})] \\
 a = c.\text{next}[l] \quad \leftrightarrow (\exists_{\text{elem}} e \exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{l-1}, a_{l+1}, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{K-1}) \\
 \quad \quad \quad [c = \text{mkcell}(e, k, a_0, \dots, a_{l-1}, a, a_{l+1}, \dots, a_{K-1}, t_0, \dots, t_{K-1})] \\
 t = c.\text{lockid}[l] \quad \leftrightarrow (\exists_{\text{elem}} e \exists_{\text{ord}} k \exists_{\text{addr}} a_0, \dots, a_{K-1} \exists_{\text{thid}} t_0, \dots, t_{l-1}, t_{l+1}, \dots, t_{K-1}) \\
 \quad \quad \quad [c = \text{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1})]
 \end{array}$$

$$\begin{aligned}
c_1 = c_2.\text{lock}(l, t) &\leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{key} = c_1.\text{key} \wedge \\
& c_2.\text{next}[0] = c_1.\text{next}[0] \wedge \\
& \dots \\
& c_2.\text{next}[\mathbf{K} - 1] = c_1.\text{next}[\mathbf{K} - 1] \wedge \\
& c_2.\text{lockid}[0] = c_1.\text{lockid}[0] \wedge \\
& \dots \\
& c_2.\text{lockid}[l - 1] = c_1.\text{lockid}[l - 1] \wedge \\
& t = c_1.\text{lockid}[l] \wedge \\
& c_2.\text{lockid}[l + 1] = c_1.\text{lockid}[l + 1] \wedge \\
& \dots \\
& c_2.\text{lockid}[\mathbf{K} - 1] = c_1.\text{lockid}[\mathbf{K} - 1] \\
c_1 = c_2.\text{unlock}(l) &\leftrightarrow c_2.\text{data} = c_1.\text{data} \wedge c_2.\text{key} = c_1.\text{key} \wedge \\
& c_2.\text{next}[0] = c_1.\text{next}[0] \wedge \\
& \dots \\
& c_2.\text{next}[\mathbf{K} - 1] = c_1.\text{next}[\mathbf{K} - 1] \wedge \\
& c_2.\text{lockid}[0] = c_1.\text{lockid}[0] \wedge \\
& \dots \\
& c_2.\text{lockid}[l - 1] = c_1.\text{lockid}[l - 1] \wedge \\
& \emptyset = c_1.\text{lockid}[l] \wedge \\
& c_2.\text{lockid}[l + 1] = c_1.\text{lockid}[l + 1] \wedge \\
& \dots \\
& c_2.\text{lockid}[\mathbf{K} - 1] = c_1.\text{lockid}[\mathbf{K} - 1] \\
c_1 \neq_{\text{cell}} c_2 &\leftrightarrow c_1.\text{data} \neq c_2.\text{data} \vee c_1.\text{key} \neq c_2.\text{key} \vee \\
& c_1.\text{next}[0] \neq c_2.\text{next}[0] \vee \\
& \dots \\
& c_1.\text{next}[\mathbf{K} - 1] \neq c_2.\text{next}[\mathbf{K} - 1] \vee \\
& c_1.\text{lockid}[0] \neq c_2.\text{next}[0] \vee \\
& \dots \\
& c_1.\text{lockid}[\mathbf{K} - 1] \neq c_2.\text{next}[\mathbf{K} - 1] \\
m_1 \neq_{\text{mem}} m_2 &\leftrightarrow (\exists_{\text{addr } a} [rd(m_1, a) \neq rd(m_2, a)]) \\
g_1 \neq_{\text{setth}} g_2 &\leftrightarrow (\exists_{\text{thid } t} [t \in (g_1 \setminus_T g_2) \cup_T (g_2 \setminus_T g_1)]) \\
g = \emptyset_T &\leftrightarrow g = g \setminus_T g \\
g_3 = g_1 \cap_T g_2 &\leftrightarrow g_3 = (g_1 \cup_T g_2) \setminus_T ((g_1 \setminus_T g_2) \cup_T (g_2 \setminus_T g_1)) \\
t \in_T g &\leftrightarrow \{t\}_T \subseteq_T g \\
g_1 \subseteq_T g_2 &\leftrightarrow g_2 = g_1 \cup_T g_2
\end{aligned}$$

$r_1 \neq_{\text{seth}} r_2$	$\leftrightarrow (\exists_{\text{addr}} a \exists_{\text{level}_K} l) [(a, l) \in (r_1 -_{\text{mr}} r_2) \cup_{\text{mr}} (r_2 -_{\text{mr}} r_1)]$
$r = \mathbf{emp}_{\text{mr}}$	$\leftrightarrow r = r -_{\text{mr}} r$
$r_3 = r_1 \cap_{\text{mr}} r_2$	$\leftrightarrow r_3 = (r_1 \cup_{\text{mr}} r_2) -_{\text{mr}} ((r_1 -_{\text{mr}} r_2) \cup_{\text{mr}} (r_2 -_{\text{mr}} r_1))$
$(a, l) \in_{\text{mr}} r$	$\leftrightarrow \langle a, l \rangle_{\text{mr}} \subseteq_{\text{mr}} r$
$r_1 \subseteq_{\text{mr}} r_2$	$\leftrightarrow r_2 = r_1 \cup_{\text{mr}} r_2$
$r_1 \#_{\text{mr}} r_2$	$\leftrightarrow \mathbf{emp}_{\text{mr}} = (r_1 \cup_{\text{mr}} r_2) -_{\text{mr}} ((r_1 -_{\text{mr}} r_2) \cup_{\text{mr}} (r_2 -_{\text{mr}} r_1))$
$p = \epsilon$	$\leftrightarrow \mathit{append}(p, p)$
$\mathit{reach}_K(m, a_1, a_2, l, p)$	$\leftrightarrow a_2 \in \mathit{addr2set}_K(m, a_1, l) \wedge p = \mathit{getp}_K(m, a_1, a_2, l)$

this means that we can rewrite such literals using:

Flat:	$e = c.\mathit{data}$
Normalized:	$c = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$k = c.\mathit{key}$
Normalized:	$c = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$e, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$a = c.\mathit{next}[l]$
Normalized:	$c = \mathit{mkcell}(e, k, a_0, \dots, a_{l-1}, a, a_{l+1}, \dots, a_{K-1}, t_0, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{l-1}, a_{l+1}, a_{K-1}, t_0, \dots, t_{K-1}$ are fresh.
Flat:	$t = c.\mathit{lockid}[l]$
Normalized:	$c = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 = c_2.\mathit{lock}(l, t)$
Normalized:	$c_1 = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, t, t_{l+1}, \dots, t_{K-1}) \wedge$ $c_2 = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 = c_2.\mathit{unlock}(l)$
Normalized:	$c_1 = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \emptyset, t_{l+1}, \dots, t_{K-1}) \wedge$ $c_2 = \mathit{mkcell}(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1})$
Proviso:	$e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{l-1}, \tilde{t}, t_{l+1}, \dots, t_{K-1}$ are fresh.
Flat:	$c_1 \neq c_2$
Normalized:	$c_1.\mathit{data} \neq c_2.\mathit{data} \vee c_1.\mathit{key} \neq c_2.\mathit{key} \vee$ $c_1.\mathit{next}[0] \neq c_2.\mathit{next}[0] \vee \dots \vee c_1.\mathit{next}[K-1] \neq c_2.\mathit{next}[K-1] \vee$ $c_1.\mathit{lockid}[0] \neq c_2.\mathit{lockid}[0] \vee \dots \vee c_1.\mathit{lockid}[K-1] \neq c_2.\mathit{lockid}[K-1]$
Proviso:	-

Flat:	$m_1 \neq m_2$
Normalized:	$m[a] \neq m[b]$
Proviso:	a is fresh.
Flat:	$s_1 \neq s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_3 = s_{12} \cup s_{21} \wedge s = s_3 \cup \{a\} \wedge \{a\} \subseteq s$
Proviso:	s_{12}, s_{21}, s_3, s and a are fresh.
Flat:	$s = \emptyset$
Normalized:	$s = s \setminus s$
Proviso:	-
Flat:	$s_3 = s_1 \cap s_2$
Normalized:	$s_{12} = s_1 \setminus s_2 \wedge s_{21} = s_2 \setminus s_1 \wedge s_{u_1} = s_1 \cup s_2 \wedge s_{u_2} = s_{12} \cup s_{21} \wedge s_3 = s_{u_1} \setminus s_{u_2}$
Proviso:	s_{12}, s_{21}, s_{u_1} and s_{u_2} are fresh.
Flat:	$a \in s$
Normalized:	$s = \{a\} \cup s$
Proviso:	-
Flat:	$s_1 \subseteq s_2$
Normalized:	$s_2 = s_1 \cup s_2$
Proviso:	-
Flat:	$g_1 \neq g_2$
Normalized:	$g_{12} = g_1 \setminus_T g_2 \wedge g_{21} = g_2 \setminus_T g_1 \wedge g_3 = g_{12} \cup_T g_{21} \wedge g = g_3 \cup_T \{a\} \wedge \{a\} \subseteq_T g$
Proviso:	g_{12}, g_{21}, g_3, g and a are fresh.
Flat:	$g = \emptyset_T$
Normalized:	$g = g \setminus_T g$
Proviso:	-
Flat:	$g_3 = g_1 \cap_T g_2$
Normalized:	$g_{12} = g_1 \setminus_T g_2 \wedge g_{21} = g_2 \setminus_T g_1 \wedge g_{u_1} = g_1 \cup_T g_2 \wedge g_{u_2} = g_{12} \cup_T g_{21} \wedge g_3 = g_{u_1} \setminus_T g_{u_2}$
Proviso:	g_{12}, g_{21}, g_{u_1} and g_{u_2} are fresh.
Flat:	$t \in_T g$
Normalized:	$g = \{t\} \cup_T g$
Proviso:	-
Flat:	$g_1 \subseteq_T g_2$
Normalized:	$g_2 = g_1 \cup_T g_2$
Proviso:	-
Flat:	$r_1 \neq r_2$
Normalized:	$r_{12} = r_1 \text{--}_{\text{mr}} r_2 \wedge r_{21} = r_2 \text{--}_{\text{mr}} r_1 \wedge r_3 = r_{12} \cup_{\text{mr}} r_{21} \wedge r = r_3 \cup_{\text{mr}} \{(a, l)\} \wedge \{(a, l)\} \subseteq_{\text{mr}} r$
Proviso:	$r_{12}, r_{21}, r_3, r, a$ and l are fresh.
Flat:	$r = \mathbf{emp}_{\text{mr}}$
Normalized:	$r = r \text{--}_{\text{mr}} r$
Proviso:	-
Flat:	$r_3 = r_1 \cap_{\text{mr}} r_2$
Normalized:	$r_{12} = r_1 \text{--}_{\text{mr}} r_2 \wedge r_{21} = r_2 \text{--}_{\text{mr}} r_1 \wedge r_{u_1} = r_1 \cup_{\text{mr}} r_2 \wedge r_{u_2} = r_{12} \cup_{\text{mr}} r_{21} \wedge r_3 = r_{u_1} \text{--}_{\text{mr}} r_{u_2}$
Proviso:	r_{12}, r_{21}, r_{u_1} and r_{u_2} are fresh.

Flat:	$(a, l) \in_{\text{mr}} r$
Normalized:	$r = \{(a, l)\} \cup_{\text{mr}} r$
Proviso:	-
Flat:	$r_1 \subseteq_{\text{mr}} r_2$
Normalized:	$r_2 = r_1 \cup_{\text{mr}} r_2$
Proviso:	-
Flat:	$r_1 \#_{\text{mr}} r_2$
Normalized:	$r_{12} = r_1 -_{\text{mr}} r_2 \wedge r_{21} = r_2 -_{\text{mr}} r_1 \wedge r_{u_1} = r_1 \cup_{\text{mr}} r_2 \wedge$ $r_{u_2} = r_{12} \cup_{\text{mr}} r_{21} \wedge r_3 = r_{u_1} -_{\text{mr}} r_{u_2} \wedge r_3 = r_3 -_{\text{mr}} r_3$
Proviso:	$r_{12}, r_{21}, r_{u_1}, r_{u_2}$ and r_3 are fresh.
Flat:	$p = \epsilon$
Normalized:	$\text{append}(p, p, p)$
Proviso:	-
Flat:	$\text{reach}_{\mathbb{K}}(m, a_1, a_2, l, p)$
Normalized:	$a_2 \in \text{addr2set}_{\mathbb{K}}(m, a_1, l) \wedge p = \text{getp}_{\mathbb{K}}(m, a_1, a_2)$
Proviso:	-

A.2 The Small Model Property

Consider an arbitrary $\text{TSL}_{\mathbb{K}}$ -interpretation \mathcal{A} satisfying a conjunction of normalized $\text{TSL}_{\mathbb{K}}$ -literals Γ . We show that if there are sets $\mathcal{A}_{\text{elem}}, \mathcal{A}_{\text{addr}}, \mathcal{A}_{\text{thid}}, \mathcal{A}_{\text{level}_{\mathbb{K}}}$ and \mathcal{A}_{ord} then there are finite sets $\mathcal{A}'_{\text{elem}}, \mathcal{A}'_{\text{addr}}, \mathcal{A}'_{\text{thid}}, \mathcal{A}'_{\text{level}_{\mathbb{K}}}$ and $\mathcal{A}'_{\text{ord}}$ with bounded cardinalities (the bound depending on Γ). $\mathcal{A}'_{\text{elem}}, \mathcal{A}'_{\text{addr}}, \mathcal{A}'_{\text{thid}}, \mathcal{A}'_{\text{level}_{\mathbb{K}}}$ and $\mathcal{A}'_{\text{ord}}$ can in turn be used to obtain a finite interpretation \mathcal{A}' satisfying Γ .

Before proving that $\text{TSL}_{\mathbb{K}}$ enjoys of finite model property, we define some auxiliary functions. We start by defining the function $\text{first}_{\mathbb{K}}$. Let $\mathcal{B}_{\text{addr}} \subseteq \tilde{\mathcal{B}}_{\text{addr}}$, $m : \tilde{\mathcal{B}}_{\text{addr}} \rightarrow \mathcal{B}_{\text{elem}} \times \mathcal{B}_{\text{ord}} \times \tilde{\mathcal{B}}_{\text{addr}}^{\mathbb{K}} \times \mathcal{B}_{\text{thid}}^{\mathbb{K}}$, $a \in \mathcal{B}_{\text{addr}}$ and $l \in \mathcal{B}_{\text{level}_{\mathbb{K}}}$. The function $\text{first}_{\mathbb{K}}(m, a, l, \mathcal{B}_{\text{addr}})$ is defined by

$$\text{first}_{\mathbb{K}}(m, a, l, \mathcal{B}_{\text{addr}}) = \begin{cases} \text{null} & \text{if for all } r \geq 1 \ m^r(a).\text{next}(l) \notin \mathcal{B}_{\text{addr}} \\ m^s(a).\text{next}(l) & \text{if for some } s \geq 1 \ m^s(a).\text{next}(l) \in \mathcal{B}_{\text{addr}}, \\ & \text{and for all } r < s \ m^r(a).\text{next}(l) \notin \mathcal{B}_{\text{addr}} \end{cases}$$

where

- $m^1(a).\text{next}(l)$ stands for $m(a).\text{next}(l)$ and
- $m^{n+1}(a).\text{next}(l)$ stands for $m(m^n(a).\text{next}(l)).\text{next}(l)$ when $n > 0$.

Basically, given the original model \mathcal{A} and a subset of addresses $X \subseteq \mathcal{A}_{\text{addr}}$, function $\text{first}_{\mathbb{K}}$ chooses the next address in X that can be reached from a given address following repeatedly the $\text{next}(l)$ pointer. It is easy to see, for example, that if $m(a).\text{next}(l) \in X$ then $\text{first}_{\mathbb{K}}(m, a, l, X) = m(a).\text{next}(l)$. We will later filter out unnecessary intermediate nodes and use $\text{first}_{\mathbb{K}}$ to bypass properly the removed nodes, preserving the important connectivity properties.

Lemma 3. *Let $\mathcal{B}_{\text{addr}} \subseteq \tilde{\mathcal{B}}_{\text{addr}}$, $m : \tilde{\mathcal{B}}_{\text{addr}} \rightarrow \mathcal{B}_{\text{elem}} \times \mathcal{B}_{\text{ord}} \times \tilde{\mathcal{B}}_{\text{addr}}^{\mathbb{K}} \times \mathcal{B}_{\text{thid}}^{\mathbb{K}}$, $a \in \mathcal{B}_{\text{addr}}$ and $l \in \mathcal{B}_{\text{level}_{\mathbb{K}}}$. If $m(a).\text{next}(l) \in \mathcal{B}_{\text{addr}}$, then $\text{first}_{\mathbb{K}}(m, a, l, \mathcal{B}_{\text{addr}}) = m(a).\text{next}(l)$.*

Proof. Immediate from definition of $first_K$.

Secondly, we define the *compress* function which, given a path p and a set $\mathcal{B}_{\text{addr}}$ of addresses, returns the path obtained from p by removing all the addresses that do not belong to $\mathcal{B}_{\text{addr}}$.

$$\text{compress}([i_1, \dots, i_n], \mathcal{B}_{\text{addr}}) = \begin{cases} \epsilon & \text{if } n = 0 \\ [i_1] \circ \text{compress}([i_2, \dots, i_n], X) & \text{if } n > 0 \text{ and } i_1 \in X \\ \text{compress}([i_2, \dots, i_n], X) & \text{otherwise} \end{cases}$$

Third, the function *fstL* that, given a memory, a path and a level, chooses the first address in a path (at the given level), whose lock is not \emptyset , returning the address as a singleton set:

$$\text{fstL}(m, [i_1, \dots, i_n], l) = \begin{cases} \emptyset & \text{if } n = 0 \\ \{i_1\} & \text{if } m(i_1).lockid(l) \neq \emptyset \\ \text{fstL}(m, [i_2, \dots, i_n], l) & \text{if } m(i_1).lockid(l) = \emptyset \end{cases}$$

Fourth, the function *unordered* that given a memory m and a path p , returns a set containing two address that witness the failure to preserve the *key* order of elements in p :

$$\text{unordered}(m, [i_1, \dots, i_n]) = \begin{cases} \emptyset & \text{if } n = 0 \text{ or } n = 1 \\ \{i_1, i_2\} & \text{if } m(i_2).key \preceq m(i_1).key \text{ and} \\ & m(i_2).key \neq m(i_1).key \\ \text{unordered}(m, [i_2, \dots, i_n]) & \text{otherwise} \end{cases}$$

If two such addresses exist, *unordered* returns the first two consecutive addresses whose keys violate the order.

Lemma 4. *Let p be a path such that $p = [a_1, \dots, a_n]$ with $n \geq 2$ and let m be a memory. If exists a_i , with $1 \leq i < n$, such that $m(a_{i+1}).key \preceq m(a_i).key$ and $m(a_{i+1}).key \neq m(a_i).key$, then $\text{unordered}(m, p) \neq \emptyset$*

Proof. By induction. Let's consider $n = 2$ and let $p = [a_1, a_2]$ s.t., $m(a_2).key \preceq m(a_1).key$ and $m(a_2).key \neq m(a_1).key$. Then, by definition of *unordered*, we have that $\text{unordered}(m, p) = \{a_1, a_2\} \neq \emptyset$.

Now let's assume that $n > 2$ and let $p = [a_1, \dots, a_n]$. If $m(a_2).key \preceq m(a_1).key$ and $m(a_2).key \neq m(a_1).key$, then we have that $\text{unordered}(m, p) = \{a_1, a_2\} \neq \emptyset$. On the other hand, if $m(a_1).key \preceq m(a_2).key$, we still knows that

there is a a_i , with $2 \leq i < n$, s.t., $m(a_{i+1}).key \preceq m(a_i).key$ and $m(a_{i+1}).key \neq m(a_i).key$. Therefore, by induction we have that $unordered(m, [a_2, \dots, a_n]) \neq \emptyset$ and by definition of $unordered$, $unordered(m, p) = unordered(m, [a_2, \dots, a_n]) \neq \emptyset$. \square

Fifth, the function $diseq$ [?] that outputs a set of address accountable for the disequality of two given paths:

$$diseq([i_1, \dots, i_n], [j_1, \dots, j_m]) = \begin{cases} \emptyset & \text{if } n = m = 0 \\ \{i_1\} & \text{if } n > 0 \text{ and } m = 0 \\ \{j_1\} & \text{if } n = 0 \text{ and } m > 0 \\ \{i_1, j_1\} & \text{if } n, m > 0 \text{ and } i_1 \neq j_1 \\ diseq([i_2, \dots, i_n], [j_2, \dots, j_m]) & \text{otherwise} \end{cases}$$

Finally, the function $common$ [?] that outputs an element common to two paths (an element that witnesses that $path2set(p) \cap path2set(q) \neq \emptyset$):

$$common([i_1, \dots, i_n], p) = \begin{cases} \emptyset & \text{if } n = 0 \\ \{i_1\} & \text{if } n > 0 \text{ and } i_1 \in path2set(p) \\ common([i_2, \dots, i_n], p) & \text{otherwise} \end{cases}$$

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

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

$$\begin{aligned} |\mathcal{B}_{addr}| &\leq \bar{a} + 1 + \bar{m} \bar{a} K + \bar{p}^2 + \bar{p}^3 + (K + 2)\bar{m} \bar{p} \\ |\mathcal{B}_{elem}| &\leq \bar{e} + \bar{m} |\mathcal{B}_{addr}| \\ |\mathcal{B}_{thid}| &\leq \bar{t} + K \bar{m} |\mathcal{B}_{addr}| + 1 \\ |\mathcal{B}_{level_K}| &\leq K \\ |\mathcal{B}_{ord}| &\leq \bar{o} + \bar{m} |\mathcal{B}_{addr}| \end{aligned}$$

Proof. (2 \rightarrow 1) is immediate.

(1 \rightarrow 2). We prove this implication only for the new TSL_K -literals.

Bearing in mind the auxiliary functions we have defined, let now \mathcal{A} be a TSL_K -interpretation satisfying a set of normalized TSL_K -literals Γ . We use \mathcal{A} to construct a TSL_K -interpretation \mathcal{B} which satisfies Γ .

$$\begin{aligned}
\mathcal{B}_{\text{level}_K} &= \mathcal{A}_{\text{level}_K} = [0 \dots K - 1] \\
\mathcal{B}_{\text{addr}} &= V_{\text{addr}}^{\mathcal{A}} \cup \{ \text{null}^{\mathcal{A}} \} \cup \\
&\quad \{ m^{\mathcal{A}}(a^{\mathcal{A}}).\text{next}^{\mathcal{A}}(l) \mid m \in V_{\text{mem}}, a \in V_{\text{addr}} \text{ and } l \in \mathcal{B}_{\text{level}_K} \} \cup \\
&\quad \{ v \in \text{diseq}(p^{\mathcal{A}}, q^{\mathcal{A}}) \mid \text{the literal } p \neq q \text{ is in } \Gamma \} \cup \\
&\quad \{ v \in \text{common}(p_1^{\mathcal{A}}, p_2^{\mathcal{A}}) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
&\quad \quad \text{path2set}^{\mathcal{A}}(p_1^{\mathcal{A}}) \cap \text{path2set}^{\mathcal{A}}(p_2^{\mathcal{A}}) \neq \emptyset \} \cup \\
&\quad \{ v \in \text{common}(p_1^{\mathcal{A}} \circ p_2^{\mathcal{A}}, p_3^{\mathcal{A}}) \mid \text{the literal } \neg \text{append}(p_1, p_2, p_3) \text{ is in } \Gamma \text{ and} \\
&\quad \quad \text{path2set}^{\mathcal{A}}(p_1^{\mathcal{A}}) \cap \text{path2set}^{\mathcal{A}}(p_2^{\mathcal{A}}) = \emptyset \} \cup \\
&\quad \{ v \in \text{fstL}(m^{\mathcal{A}}, p^{\mathcal{A}}, l) \mid \text{fstlock}(m, p, l) \text{ is in } \Gamma \} \\
&\quad \{ v \in \text{unordered}(m^{\mathcal{A}}, p^{\mathcal{A}}) \mid \neg \text{ordList}(m, p) \text{ is in } \Gamma \} \\
\mathcal{B}_{\text{thid}} &= V_{\text{thid}}^{\mathcal{A}} \cup \{ \emptyset \} \cup \{ m^{\mathcal{A}}(v^{\mathcal{A}}).\text{lockid}^{\mathcal{A}}(l) \mid m \in V_{\text{mem}}, v \in \mathcal{B}_{\text{addr}} \text{ and } l \in \mathcal{B}_{\text{level}_K} \} \\
\mathcal{B}_{\text{elem}} &= V_{\text{elem}}^{\mathcal{A}} \cup \{ m^{\mathcal{A}}(v).\text{data}^{\mathcal{A}} \mid m \in V_{\text{mem}} \text{ and } v \in \mathcal{B}_{\text{addr}} \} \\
\mathcal{B}_{\text{ord}} &= V_{\text{ord}}^{\mathcal{A}} \cup \{ m^{\mathcal{A}}(v).\text{key}^{\mathcal{A}} \mid m \in V_{\text{mem}} \text{ and } v \in \mathcal{B}_{\text{addr}} \}
\end{aligned}$$

These domains satisfy the cardinality constraints expressed in the statement of the theorem. The interpretations of the symbols are:

$$\begin{aligned}
\text{error}^{\mathcal{B}} &= \text{error}^{\mathcal{A}} \\
\text{null}^{\mathcal{B}} &= \text{null}^{\mathcal{A}} \\
e^{\mathcal{B}} &= e^{\mathcal{A}} && \text{for each } e \in V_{\text{elem}} \\
a^{\mathcal{B}} &= a^{\mathcal{A}} && \text{for each } a \in V_{\text{addr}} \\
c^{\mathcal{B}} &= c^{\mathcal{A}} && \text{for each } c \in V_{\text{cell}} \\
t^{\mathcal{B}} &= t^{\mathcal{A}} && \text{for each } t \in V_{\text{thid}} \\
k^{\mathcal{B}} &= k^{\mathcal{A}} && \text{for each } k \in V_{\text{ord}} \\
l^{\mathcal{B}} &= l^{\mathcal{A}} && \text{for each } l \in V_{\text{level}_K} \\
m^{\mathcal{B}}(v) &= (m^{\mathcal{A}}(v).\text{data}^{\mathcal{A}}, m^{\mathcal{A}}(v).\text{key}^{\mathcal{A}}, && \text{for each } m \in V_{\text{mem}} \\
&\quad \text{first}_K(m^{\mathcal{A}}, v, 0, \mathcal{B}_{\text{addr}}), && \text{and } v \in \mathcal{B}_{\text{addr}} \\
&\quad \dots \\
&\quad \text{first}_K(m^{\mathcal{A}}, v, K - 1, \mathcal{B}_{\text{addr}}), \\
&\quad m^{\mathcal{A}}(v).\text{lockid}^{\mathcal{A}}[0], \dots, m^{\mathcal{A}}(v).\text{lockid}^{\mathcal{A}}[K - 1]) \\
s^{\mathcal{B}} &= s^{\mathcal{A}} \cap \mathcal{B}_{\text{addr}} && \text{for each } s \in V_{\text{set}} \\
g^{\mathcal{B}} &= g^{\mathcal{A}} \cap \mathcal{B}_{\text{thid}} && \text{for each } g \in V_{\text{setth}} \\
r^{\mathcal{B}} &= r^{\mathcal{A}} \cap (\mathcal{B}_{\text{addr}} \times \mathcal{B}_{\text{level}_K}) && \text{for each } r \in V_{\text{mrgn}} \\
p^{\mathcal{B}} &= \text{compress}(p^{\mathcal{A}}, \mathcal{B}_{\text{addr}}) && \text{for each } p \in V_{\text{path}}
\end{aligned}$$

Essentially, all variables and constants in \mathcal{B} are interpreted as in \mathcal{A} except that *next* pointers use *first*_K to point to the next reachable element that has been preserved in $\mathcal{B}_{\text{addr}}$, and paths filter out all elements except those in $\mathcal{B}_{\text{addr}}$. It can be routinely checked that \mathcal{B} is an interpretation of Γ . So it remains to be seen that \mathcal{B} satisfies all literals in Γ assuming that \mathcal{A} does, concluding that \mathcal{B} is

indeed a model of Γ . This check is performed by cases. The proof that \mathcal{B} satisfies all TSL_K -literals in Γ is not shown here. We just focus on the new functions and predicates that are not part of TLL . The proof for the missing literals can be found in [?]. For TSL_K -literals we must consider the following cases:

Literals of the form $l_1 \neq l_2, k_1 \neq k_2$ and $k_1 \preceq k_2$. Immediate

Literals of the form $c = mkcell(e, k, a_0, \dots, a_{K-1}, t_0, \dots, t_{K-1})$.

$$\begin{aligned} c^{\mathcal{B}} &= c^{\mathcal{A}} = (e^{\mathcal{A}}, k^{\mathcal{A}}, a_0^{\mathcal{A}}, \dots, a_{K-1}^{\mathcal{A}}, t_0^{\mathcal{A}}, \dots, t_{K-1}^{\mathcal{A}}) \\ &= (e^{\mathcal{B}}, k^{\mathcal{B}}, a_0^{\mathcal{B}}, \dots, a_{K-1}^{\mathcal{B}}, t_0^{\mathcal{B}}, \dots, t_{K-1}^{\mathcal{B}}) \end{aligned}$$

Literals of the form $c = rd(m, a)$. In this case we have that

$$\begin{aligned} [rd(m, a)]^{\mathcal{B}} &= m^{\mathcal{B}}(a^{\mathcal{B}}) \\ &= m^{\mathcal{B}}(a^{\mathcal{A}}) \\ &= \left(m^{\mathcal{A}}(a^{\mathcal{A}}).data^{\mathcal{A}}, m^{\mathcal{A}}(a^{\mathcal{A}}).key^{\mathcal{A}}, \right. \\ &\quad \left. first_K(m^{\mathcal{A}}, a^{\mathcal{A}}, 0, \mathcal{B}_{sAddr}), \dots, first_K(m^{\mathcal{A}}, a^{\mathcal{A}}, K-1, \mathcal{B}_{sAddr}), \right. \\ &\quad \left. m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[0], \dots, m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[K-1] \right) \\ &= \left(m^{\mathcal{A}}(a^{\mathcal{A}}).data^{\mathcal{A}}, m^{\mathcal{A}}(a^{\mathcal{A}}).key^{\mathcal{A}}, \right. \\ &\quad \left. m^{\mathcal{A}}(a^{\mathcal{A}}).next^{\mathcal{A}}(0), \dots, m^{\mathcal{A}}(a^{\mathcal{A}}).next^{\mathcal{A}}(K-1), \right. \\ &\quad \left. m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[0], \dots, m^{\mathcal{A}}(a^{\mathcal{A}}).lockid^{\mathcal{A}}[K-1] \right) \quad (\text{Lemma 3}) \\ &= m^{\mathcal{A}}(a^{\mathcal{A}}) \\ &= c^{\mathcal{A}} \\ &= c^{\mathcal{B}} \end{aligned}$$

Literals of the form $g = \{t\}_T$. We have that

$$g^{\mathcal{B}} = g^{\mathcal{A}} \cap \mathcal{B}_{sThId} = \{t^{\mathcal{A}}\}_T \cap \mathcal{B}_{sThId} = \{t^{\mathcal{B}}\}_T \cap \mathcal{B}_{sThId} = \{t^{\mathcal{B}}\}_T$$

Literals of the form $g_1 = g_2 \cup_T g_3$. In this case we have that

$$\begin{aligned} g_1^{\mathcal{B}} &= g_1^{\mathcal{A}} \cap \mathcal{B}_{sThId} = (g_2^{\mathcal{A}} \cup_T g_3^{\mathcal{A}}) \cap \mathcal{B}_{sThId} \\ &= (g_2^{\mathcal{A}} \cap \mathcal{B}_{sThId}) \cup_T (g_3^{\mathcal{A}} \cap \mathcal{B}_{sThId}) \\ &= g_2^{\mathcal{B}} \cup_T g_3^{\mathcal{B}} \end{aligned}$$

Literals of the form $g_1 = g_2 \setminus_T g_3$. We have that

$$\begin{aligned} g_1^{\mathcal{B}} &= g_1^{\mathcal{A}} \cap \mathcal{B}_{sThId} = (g_2^{\mathcal{A}} \setminus_T g_3^{\mathcal{A}}) \cap \mathcal{B}_{sThId} \\ &= (g_2^{\mathcal{A}} \cap \mathcal{B}_{sThId}) \setminus_T (g_3^{\mathcal{A}} \cap \mathcal{B}_{sThId}) \\ &= g_2^{\mathcal{B}} \setminus_T g_3^{\mathcal{B}} \end{aligned}$$

Literals of the form $r = \langle a, l \rangle_{\text{mr}}$. We have that

$$\begin{aligned}
r^{\mathcal{B}} &= r^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}}) \\
&= \langle a^{\mathcal{A}}, l^{\mathcal{A}} \rangle_{\text{mr}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}}) \\
&= \langle a^{\mathcal{B}}, l^{\mathcal{B}} \rangle_{\text{mr}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}}) \\
&= \langle a^{\mathcal{B}}, l^{\mathcal{B}} \rangle_{\text{mr}}
\end{aligned}$$

Literals of the form $r_1 = r_2 \cup_{\text{mr}} r_3$. In this case we have that

$$\begin{aligned}
r_1^{\mathcal{B}} &= r_1^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}}) \\
&= (r_2^{\mathcal{A}} \cup_{\text{mr}} r_3^{\mathcal{A}}) \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}}) \\
&= (r_2^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}})) \cup_{\text{mr}} (r_3^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{\text{level}_{\mathcal{K}}})) \\
&= r_2^{\mathcal{B}} \cup_{\text{mr}} r_3^{\mathcal{B}}
\end{aligned}$$

Literals of the form $r_1 = r_2 -_{\text{mr}} r_3$. We have that

$$\begin{aligned}
r_1^{\mathcal{B}} &= r_1^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{s\text{Level}_{\mathcal{K}}}) \\
&= (r_2^{\mathcal{A}} -_{\text{mr}} r_3^{\mathcal{A}}) \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{s\text{Level}_{\mathcal{K}}}) \\
&= (r_2^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{s\text{Level}_{\mathcal{K}}})) -_{\text{mr}} (r_3^{\mathcal{A}} \cap (\mathcal{B}_{s\text{Addr}} \times \mathcal{B}_{s\text{Level}_{\mathcal{K}}})) \\
&= r_2^{\mathcal{B}} -_{\text{mr}} r_3^{\mathcal{B}}
\end{aligned}$$

Literals of the form $s = \text{addr2set}_{\mathcal{K}}(m, a, l)$. Let $x = a^{\mathcal{B}} = a^{\mathcal{A}}$. Then, we have that

$$\begin{aligned}
s^{\mathcal{B}} &= s^{\mathcal{A}} \cap \mathcal{B}_{s\text{Addr}} \\
&= \left\{ y \in \mathcal{A}_{\text{addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^{\mathcal{A}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{A}} \right\} \cap \mathcal{B}_{s\text{Addr}} \\
&= \left\{ y \in \mathcal{B}_{s\text{Addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^{\mathcal{A}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{A}} \right\} \\
&= \left\{ y \in \mathcal{B}_{s\text{Addr}} \mid \exists p \in \mathcal{B}_{\text{path}} \text{ s.t.}, (m^{\mathcal{B}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{B}} \right\}
\end{aligned}$$

It just remains to see that the last equality holds. Let

$$\begin{aligned}
- S_{\mathcal{B}} &= \left\{ y \in \mathcal{B}_{s\text{Addr}} \mid \exists p \in \mathcal{B}_{\text{path}} \text{ s.t.}, (m^{\mathcal{B}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{B}} \right\}, \text{ and} \\
- S_{\mathcal{A}} &= \left\{ y \in \mathcal{B}_{s\text{Addr}} \mid \exists p \in \mathcal{A}_{\text{path}} \text{ s.t.}, (m^{\mathcal{A}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{A}} \right\}
\end{aligned}$$

We first show that $S_{\mathcal{A}} \subseteq S_{\mathcal{B}}$. Let $y \in S_{\mathcal{A}}$. Then exists $p \in \mathcal{A}_{\text{path}}$ such that $(m^{\mathcal{A}}, x, y, l, p) \in \text{reach}_{\mathcal{K}}^{\mathcal{A}}$. Then, by definition of $\text{reach}_{\mathcal{K}}$ there are two possible cases.

- If $p = \epsilon$ and $x = y$, then $(m^{\mathcal{B}}, x, y, l, \epsilon^{\mathcal{B}}) \in reach_{\mathcal{K}}^{\mathcal{B}}$ and therefore $y \in S_{\mathcal{B}}$.
- Otherwise, there exists $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ s.t.,

$$\begin{array}{ll} i) p = [a_1, \dots, a_n] & iii) m^{\mathcal{A}}(a_r).next^{\mathcal{A}}(l) = a_{r+1}, \text{ for } 1 \leq r < n \\ ii) x = a_1 & iv) m^{\mathcal{A}}(a_n).next^{\mathcal{A}}(l) = y \end{array}$$

Then, we only need to find $\tilde{a}_1, \dots, \tilde{a}_m \in \mathcal{B}_{\text{addr}}$ s.t.,

$$\begin{array}{ll} i) q = [\tilde{a}_1, \dots, \tilde{a}_m] & iii) m^{\mathcal{B}}(\tilde{a}_r).next^{\mathcal{B}}(l) = \tilde{a}_{r+1}, \text{ for } 1 \leq r < m \\ ii) x = \tilde{a}_1 & iv) m^{\mathcal{B}}(\tilde{a}_m).next^{\mathcal{B}}(l) = y \end{array}$$

We define $\tilde{a}_1 = a_1 = x$ and $\tilde{a}_2 = first_{\mathcal{K}}(m^{\mathcal{A}}, \tilde{a}_1, l, \mathcal{B}_{sAddr})$. Then we know that $\tilde{a}_2 = m^{\mathcal{B}}(\tilde{a}_1).next^{\mathcal{B}}(l)$ and that $\tilde{a}_2 \in \mathcal{B}_{sAddr}$ and thus $\tilde{a}_2 \in \mathcal{B}_{\text{addr}}$. Then, if $\tilde{a}_2 = y$ there is nothing else to prove. On the other hand, if $\tilde{a}_2 \neq y$ then we proceed in the same way to define \tilde{a}_3 and so on until $\tilde{a}_{m+1} = y$. Notice that this way, y is guaranteed to be found in at most n steps.

To show that $S_{\mathcal{B}} \subseteq S_{\mathcal{A}}$ we proceed in a similar way. Let $y \in S_{\mathcal{B}}$. Then $x = y$ and $p = \epsilon$ and thus $(m^{\mathcal{A}}, x, y, l, \epsilon^{\mathcal{A}}) \in reach_{\mathcal{K}}^{\mathcal{A}}$, or exists $a_1, \dots, a_n \in \mathcal{B}_{sAddr}$ such that

$$\begin{array}{ll} i) p = [a_1, \dots, a_n] & iii) m^{\mathcal{B}}(a_r).next^{\mathcal{B}}(l) = a_{r+1}, \text{ for } 1 \leq r < n \\ ii) x = a_1 & iv) m^{\mathcal{B}}(a_n).next^{\mathcal{B}}(l) = y \end{array}$$

As we know that $a_1, \dots, a_n, y \in \mathcal{B}_{sAddr}$, by definition of $first_{\mathcal{K}}$ we know that exists $s \geq 1$ s.t.,

$$m^{\mathcal{A}}\left(\dots \underbrace{(m^{\mathcal{A}}(a_1).next^{\mathcal{A}}(l)) \dots}_s\right).next^{\mathcal{A}}(l) = a_2$$

Let then $a_1^1, \dots, a_1^{s-1} \in \mathcal{A}_{\text{addr}}$ such that

$$\begin{aligned} m^{\mathcal{A}}(a_1).next^{\mathcal{A}}(l) &= a_1^1 \\ m^{\mathcal{A}}(a_1^1).next^{\mathcal{A}}(l) &= a_1^2 \\ &\vdots \\ m^{\mathcal{A}}(a_1^{s-1}).next^{\mathcal{A}}(l) &= a_2 \end{aligned}$$

We then use a_1^1, \dots, a_1^{s-1} to construct the section of a path q that goes from a_1 up to a_2 . Finally we use the same approach to finish with the construction of such path in \mathcal{A} . Then we have that $(m^{\mathcal{A}}, x, y, l, q^{\mathcal{A}}) \in reach_{\mathcal{K}}^{\mathcal{A}}$. Then, $y \in S_{\mathcal{A}}$.

Literals of the form $p = getp_{\mathcal{K}}(m, a, b, l)$. We consider two possible cases.

- Case $b^A \in \text{addr2set}_{\mathbb{K}}(m^A, a^A, l)$.
Since $(m^A, a^A, b^A, l, p^A) \in \text{reach}_{\mathbb{K}}^A$, it is enough to prove:

$$(m^A, x, y, l, q) \in \text{reach}_{\mathbb{K}}^A \quad \rightarrow \quad (m^B, x, y, l, \text{compress}(q, \mathcal{B}_{sAddr})) \in \text{reach}_{\mathbb{K}}^B$$

for each $x, y \in \mathcal{B}_{sAddr}$ and $q \in \mathcal{A}_{\text{path}}$. Assume that $(m^A, x, y, l, q) \in \text{reach}_{\mathbb{K}}^A$. If $x = y$ and $q = \epsilon$, then $(m^B, x, y, l, \text{compress}(q, \mathcal{B}_{sAddr})) \in \text{reach}_{\mathbb{K}}^B$. Otherwise, there exists $a_1, \dots, a_n \in \mathcal{A}_{\text{addr}}$ such that:

$$\begin{array}{ll} i) q = [a_1, \dots, a_n] & iii) m^A(a_r).next^A(l) = a_{r+1}, \text{ for } 1 \leq r < n \\ ii) x = a_1 & iv) m^A(a_n).next^A(l) = y \end{array}$$

Then, we proceed by induction on n .

- If $n = 1$, then $q = [a_1]$ and therefore $\text{compress}(q, \mathcal{B}_{sAddr}) = [a_1]$, since $x = a_1 \in \mathcal{B}_{sAddr}$. Besides $m^A(a_1).next^A(l) = y$ which implies that $m^B(a_1).next^B(l) = y$. Then $(m^B, x, y, l, \text{compress}(q, \mathcal{B}_{sAddr})) \in \text{reach}_{\mathbb{K}}^B$.
- If $n > 1$, then let $a_i = \text{first}_{\mathbb{K}}(m^A, x, l, \mathcal{B}_{sAddr})$. As

$$q = [x = a_1, a_2, \dots, a_i, a_{i+1}, \dots, a_n]$$

we have that

$$\text{compress}(q, \mathcal{B}_{sAddr}) = [x = a_1] \circ \text{compress}([a_i, a_{i+1}, \dots, a_n], \mathcal{B}_{sAddr})$$

Besides, as $(m^A, a_i, y, l, [a_i, a_{i+1}, \dots, a_n]) \in \text{reach}_{\mathbb{K}}^A$, by induction we have that

$$(m^B, a_i, y, l, \text{compress}([a_i, a_{i+1}, \dots, a_n], \mathcal{B}_{sAddr})) \in \text{reach}_{\mathbb{K}}^B$$

Moreover $m^B(x).next^B(l) = a_i$ and therefore

$$(m^B, x, y, l, \text{compress}(q, \mathcal{B}_{sAddr})) \in \text{reach}_{\mathbb{K}}^B$$

- Case $b^A \notin \text{addr2set}_{\mathbb{K}}(m^A, a^A, l)$.
In such case we have that $p^A = \epsilon$, which implies that $p^B = \epsilon$. Then using a reasoning similar to the previous case we can deduce that $b^B \notin \text{addr2set}_{\mathbb{K}}(m^B, a^B, l)$.

Literals of the form $a = \text{fstlock}_{\mathbb{K}}(m, p, l)$. If we consider the case $p = \epsilon$, then we know that $\text{fstlock}_{\mathbb{K}}^A(m^A, \epsilon^A, l^A) = \text{null}^A$. At the same time, we know that $\epsilon^B = \text{compress}(\epsilon^A, \mathcal{B}_{sAddr})$ and so $\text{fstlock}_{\mathbb{K}}^B(m^B, \epsilon^B, l^B) = \text{null}^B$. Let's now consider the case at which $p = [a_1, \dots, a_n]$. There are two scenarios to consider.

- If for all $1 \leq k \leq n$, $m^A(a_k^A).lockid(l) = \emptyset$, then we have that

$$fstlock_{\kappa}^A(m^A, p^A, l^A) = null^A$$

Notice that function *compress* returns a subset of the path it receives with the property that all addresses in the returned path belong to the received set. Then, if $[\tilde{a}_1, \dots, \tilde{a}_m] = p^B = compress(p^A, \mathcal{B}_{sAddr})$, we know that $\{\tilde{a}_1, \dots, \tilde{a}_m\} \subseteq \mathcal{B}_{sAddr}$ and therefore for all $1 \leq j \leq m$, $m^B(\tilde{a}_j).lockid^B(l^B) = \emptyset$. Then, we can finally conclude that in fact $fstlock_{\kappa}^B(m^B, p^B, l^B) = null^B$.

- If exists a $1 \leq k \leq n$ such that for all $1 \leq j < k$, $m^A(a_j^A).lockid(l) = \emptyset$ and $m^A(a_k^A).lockid(l) \neq \emptyset$ then since $a^B = a^A$, we can say that $a^B = a^A = x \in \mathcal{B}_{sAddr}$. It then remains to verify whether

$$x = fstlock_{\kappa}^A(m^A, p^A, l^A) \rightarrow x = fstlock_{\kappa}^B(m^B, compress(p^A, \mathcal{B}_{sAddr}), l^B)$$

By definition of *fstlock_κ* we have that $x = a_k^A$ and by κ we know that $a_k^A \in \mathcal{B}_{sAddr}$. Let $[\tilde{a}_1, \dots, \tilde{a}_i, \dots, \tilde{a}_m] = compress(p^A, \mathcal{B}_{sAddr})$ such that $\tilde{a}_i = a_k^A$. We also know that $\tilde{a}_j \in \mathcal{B}_{sAddr}$ for all $1 \leq j \leq m$. Then, as *compress* preserves the order and for all $1 \leq j < k$, $m^A(a_j^A).lockid^A(l^A) = \emptyset$, we have that for all $1 \leq j < i$, $m^B(\tilde{a}_j).lockid^B(l^B) = \emptyset$. Besides $m^B(\tilde{a}_i).lockid^B(l^B) \neq \emptyset$. Then:

$$\begin{aligned} fstlock_{\kappa}^B(m^B, compress(p^A, l^B)) &= fstlock_{\kappa}^B(m^B, [\tilde{a}_1, \dots, \tilde{a}_m], l^B) \\ &= \tilde{a}_i \\ &= a^A \\ &= x \end{aligned}$$

Literals of the form *ordList*(*m*, *p*). Assume that $(m^A, p^A) \in ordList^A$. We want to see that $(m^B, p^B) \in ordList^B$ i.e., $(m^B, compress(p^A, \mathcal{B}_{sAddr})) \in ordList^B$. We proceed by induction on *p*.

- If $p = \epsilon$, by definition of *compress* and *ordList*, we have that $(m^B, \epsilon^B) \in ordList^B$.
- If $p = [a_1]$, we know that $(m^A, [a_1]^A) \in ordList^A$ and that $p^B = compress(p^A, \mathcal{B}_{sAddr})$. Then, if $a_1^A \in \mathcal{B}_{sAddr}$, we have that $p^B = [a_1]^B$ and then clearly $(m^B, p^B) \in ordList^B$ holds. On the other hand, if $a_1^A \notin \mathcal{B}_{sAddr}$, then $p^B = \epsilon^B$ and once more $(m^B, p^B) \in ordList^B$ holds.
- If $p = [a_1, \dots, a_{n+1}]$ with $n \geq 1$, then we have two possible cases to bear in mind. If we consider the case at which $a_1^A \notin \mathcal{B}_{sAddr}$ then $compress(p^A, \mathcal{B}_{sAddr}) = compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{sAddr})$ and as by induction we have that $(m^B, compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{sAddr})) \in ordList^B$ we conclude that $(m^B, compress([a_1, a_2, \dots, a_{n+1}]^A, \mathcal{B}_{sAddr})) \in ordList^B$. On the other hand, if $a_1^A \in \mathcal{B}_{sAddr}$ then once more, by induction, $(m^B, compress([a_2, \dots, a_{n+1}]^A, \mathcal{B}_{sAddr})) \in ordList^B$. Besides, as we have

that $m^A(a_1^A).key^A \preceq m^A(a_2^A).key^A$ we can deduce that $m^B(a_1^A).key^B \preceq m^B(a_2^A).key^B$. And so, $(m^B, compress([a_1, a_2, \dots, a_{n+1}]^A, \mathcal{B}_{sAddr})) \in ordList^B$.

Literals of the form $\neg ordList(m, p)$. Let's assume that $(m^A, p^A) \notin ordList^A$.

We want to see that $(m^B, p^B) \notin ordList^B$. If $(m^A, p^A) \notin ordList^A$, then it means that $p = [a_1, \dots, a_n]$ with $n \geq 2$ and $m^A(a_{i+1}).key^A \preceq m^A(a_i).key^A$ and $m^A(a_{i+1}).key^A \neq m^A(a_i).key^A$ for some $i \in 1, \dots, n-1$. Let that i be the one such that for all $j < i$, $m^A(a_j).key^A \preceq m^A(a_{j+1}).key^A$. Then, by Lemma 4 we know that $unordered(m^A, [a_1, \dots, a_n]^A) \neq \emptyset$ and besides $\{a_i^A, a_{i+1}^A\} \subseteq unordered(m^A, [a_1, \dots, a_n]^A) \subseteq \mathcal{B}_{sAddr}$. This means that $compress([a_1, \dots, a_n]^A, \mathcal{B}_{sAddr}) = [\tilde{a}_1, \dots, a_i, a_{i+1}, \dots, \tilde{a}_m]^B$. Therefore, since $m^B(a_{i+1}).key^B \preceq m^B(a_i).key^B$ and $m^B(a_{i+1}).key^B \neq m^B(a_i).key^B$, we have that $(m^B, compress([a_1, \dots, a_n]^A, \mathcal{B}_{sAddr})) \notin ordList^B$.

□

B Missing Implementations

```

1: procedure SEARCH(SkipList sl, Value v)
2:   int i := K - 1 //@ mr mr := ∅
3:   Node* pred := sl.head
4:   pred.locks[i].lock() //@ mr := mr ∪ {(pred, i)}
5:   Node* curr := pred.next[i]
6:   curr.locks[i].lock() //@ mr := mr ∪ {(curr, i)}
7:   while 0 ≤ i ∧ curr.val ≠ v do
8:     if i < K - 1 then
9:       pred.locks[i].lock() //@ mr := mr ∪ {(pred, i)}
10:      curr := pred.next[i]
11:      curr.locks[i].lock() //@ mr := mr ∪ {(curr, i)}
12:      pred.next[i + 1].locks[i + 1].unlock() //@ mr := mr - {(pred.next[i + 1], i + 1)}
13:      pred.locks[i + 1].unlock() //@ mr := mr - {(pred, i + 1)}
14:     end if
15:     while curr.val < v do
16:       pred.locks[i].unlock() //@ mr := mr - {(pred, i)}
17:       pred := curr
18:       curr := pred.next[i]
19:       curr.locks[i].lock() //@ mr := mr ∪ {(curr, i)}
20:     end while
21:     i := i - 1
22:   end while
23:   Bool valueIsIn := (curr.val = v)
24:   if i = K - 1 then
25:     curr.locks[i].unlock() //@ mr := mr - {(curr, i)}
26:     pred.locks[i].unlock() //@ mr := mr - {(pred, i)}
27:   else
28:     curr.locks[i + 1].unlock() //@ mr := mr - {(curr, i + 1)}
29:     pred.locks[i + 1].unlock() //@ mr := mr - {(pred, i + 1)}
30:   end if
31:   return valueIsIn
32: end procedure

```

Fig. 7. Algorithm for searching on a concurrent lock-coupling skiplist

```

1: procedure REMOVE(SkipList sl, Value v)
2:   Vector < Node* > upd[0..K - 1]           //@ mrngn mr := ∅
3:   Node* pred := sl.head
4:   pred.locks[K - 1].lock()                 //@ mr := mr ∪ {(pred, K)}
5:   Node* curr := pred.next[K - 1]
6:   curr.locks[K - 1].lock()                 //@ mr := mr ∪ {(curr, K)}
7:   for i := K - 1 downto 0 do
8:     if i < K - 1 then
9:       pred.locks[i].lock()                 //@ mr := mr ∪ {(pred, i)}
10:      curr := pred.next[i]
11:      curr.locks[i].lock()                 //@ mr := mr ∪ {(curr, i)}
12:    end if
13:    while curr.val < v do
14:      pred.locks[i].unlock()                 //@ mr := mr - {(pred, i)}
15:      pred := curr
16:      curr := pred.next[i]
17:      curr.locks[i].lock()                 //@ mr := mr ∪ {(curr, i)}
18:    end while
19:    upd[i] := pred
20:  end for
21:  for i := K - 1 downto 0 do
22:    if upd[i].next[i] = curr ∧ curr.val = v then
23:      upd[i].next[i] := curr.next[i]       //@ sl.r := sl.r - {(curr, i)}
24:      curr.locks[i].unlock()                 //@ mr := mr - {(curr, i)}
25:    else
26:      upd[i].next[i].locks[i].unlock()     //@ mr := mr - {upd[i].next[i], i}
27:    end if
28:    upd[i].locks[i].unlock()                 //@ mr := mr - {(upd[i], i)}
29:  end for
30:  Bool valueWasIn := (curr.val = v)
31:  if valueWasIn then
32:    free (curr)
33:  end if
34:  return valueWasIn
35: end procedure

```

Fig. 8. Algorithm for deletion on a concurrent lock-coupling skiplist

C No Thread Overtakes

In this section we prove that, in fact, no thread can overtake another thread, considering the region of the skiplist that can potentially be modified by the latter one. This does not mean that no thread can overtake another one using a higher level pointer. Instead, we want to verify that no thread can go through the region to be modified by another thread.

Let's consider the skiplist shown in Fig. 9(a). Imagine a thread j trying to insert a one level node with value 11. Then, after reaching the position where the node must be inserted the skiplist may look as the one depicted in Fig. 9(b). If then another thread, lets say i , wants to insert a node with value 19, it will undoubtedly jump over the nodes locked by thread j . This is not the situation we are trying to prevent, since this is a correct behavior for a concurrent skiplist. Moreover, the modifications introduced by thread i will not interfere with thread i at all.

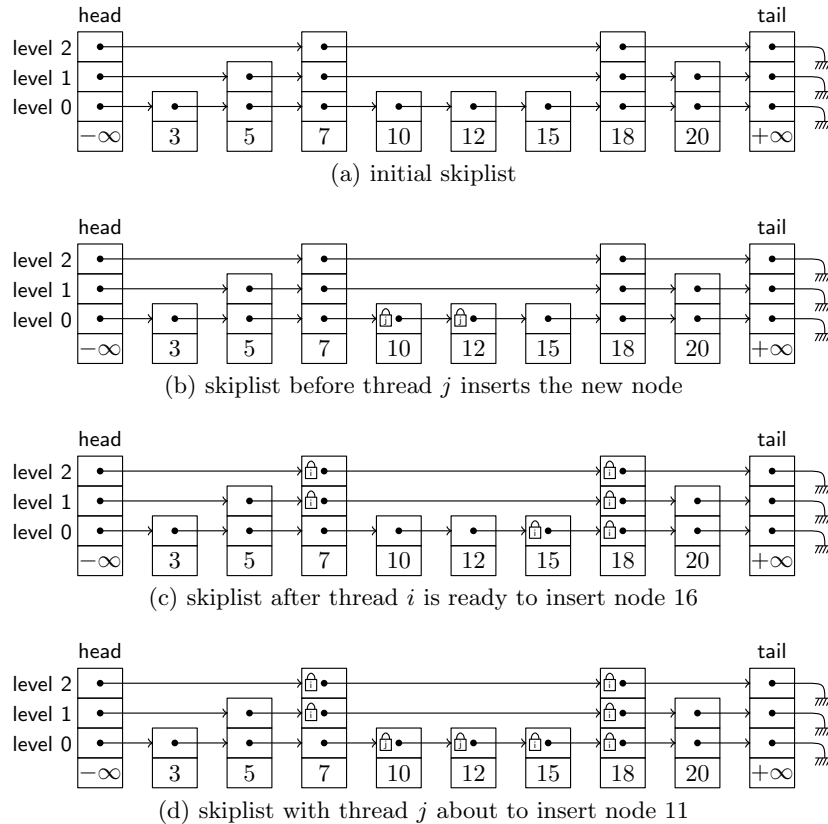


Fig. 9. Example of skiplist

However, let's now consider the same skiplist depicted in Fig. 9(a). Imagine now that thread i wants to insert a node of three levels with value 16. In such case, before the insertion is accomplished, the skiplist will have the aspect depicted in Fig. 9(c). Is in this case we want to show that thread j will not be able to progress up to the position where node 11 must be inserted, ending up in a scenario as the one shown in Fig. 9(d). An informal reasoning let us deduce that thread j cannot reach such position since node 7 should be reached in a top-down fashion, something that cannot happen since thread i has locked the upper levels of such node. We now proceed to formalize this reasoning.

We begin by extending the actual code with ghost code to aid the verification. We add the ghost variables L, U, H denoting the limits of the minimum region we are sure a thread can potentially modified (PM). Such region is then formally defined as a masked region by:

$$PM \doteq \bigcup_{i=0}^H \{(n, i) \mid n \in \text{getp}_\kappa(h, L, U, i)\}$$

assuming that the skiplist resides in the heap h . Considering once more the situation of thread i trying to insert a node with value 16 depicted in Fig. 9(c), we can represent the PM region of such thread as shown with dashed lines in Fig. 10.




Fig. 10. PM region for thread i when inserting node 16

We first extend the algorithm for *search*, *insert* and *remove* with the new ghost variables. Then, we define a function $\text{skipRgTh} : \text{addr} \times \text{addr} \times \text{level}_\kappa \rightarrow \text{setth}$ and a predicate $\text{lastTh} : \text{addr} \times \text{addr} \times \text{level}_\kappa$. Given a lower address L , and upper address U and a level L , the function $\text{skipRgTh}(L, U, H)$ returns the set of threads identifiers which has a locked node within the PM region described by L , U and H . Meanwhile, $\text{lastTh}(L, U, H)$ holds whenever $\text{skipRgTh}(L, U, H) = \emptyset_T$.

The new algorithms extended with ghost variables are depicted in Fig. 11, 12 and 13. Notice that when setting H to -1 we are just saying that PM becomes empty.

The main idea is that ever moment, $PM^{[t]}$ represents the minimum region of the skiplist we are sure thread t can potentially modify. After every transition is taken, we end up with a subregion (possibly the same one) as before. We would like to ensure that every transition, taken by thread t or any other thread of the system, does not increment the number of threads within its PM region.

```

1: procedure SEARCH(SkipList sl, Value v)
2:   int i := K - 1                                //@ mrng  $m_r := \emptyset$ 
                                                    //@  $L := sl.head$ 
                                                    //@  $U := sl.tail$ 
                                                    //@  $H := K - 1$ 

3:   Node* pred := sl.head
4:   pred.locks[i].lock()                            //@  $m_r := m_r \cup \{(pred, i)\}$ 
5:   Node* curr := pred.next[i]
6:   curr.locks[i].lock()                            //@  $m_r := m_r \cup \{(curr, i)\}$ 
7:   while  $0 \leq i \wedge curr.val \neq v$  do
8:     if  $i < K - 1$  then
9:       pred.locks[i].lock()                        //@  $m_r := m_r \cup \{(pred, i)\}$ 
                                                    //@  $U := curr$ 
10:      curr := pred.next[i]
11:      curr.locks[i].lock()                          //@  $m_r := m_r \cup \{(curr, i)\}$ 
12:      pred.next[i + 1].locks[i + 1].unlock()
                                                    //@  $m_r := m_r - \{(pred.next[i + 1], i + 1)\}$ 
13:      pred.locks[i + 1].unlock()                    //@  $m_r := m_r - \{(pred, i + 1)\}$ 
                                                    //@  $H := i$ 
14:     end if
15:     while  $curr.val < v$  do
16:       pred.locks[i].unlock()                      //@  $m_r := m_r - \{(pred, i)\}$ 
                                                    //@  $L := curr$ 
17:       pred := curr
18:       curr := pred.next[i]
19:       curr.locks[i].lock()                          //@  $m_r := m_r \cup \{(curr, i)\}$ 
20:     end while
21:     i := i - 1
22:   end while
23:   Bool valueIsIn := (curr.val = v)
24:   if  $i = K - 1$  then
25:     curr.locks[i].unlock()                          //@  $m_r := m_r - \{(curr, i)\}$ 
                                                    //@  $U := L$ 
26:     pred.locks[i].unlock()                          //@  $m_r := m_r - \{(pred, i)\}$ 
                                                    //@  $H := -1$ 
27:   else
28:     curr.locks[i + 1].unlock()                      //@  $m_r := m_r - \{(curr, i + 1)\}$ 
                                                    //@  $U := L$ 
29:     pred.locks[i + 1].unlock()                      //@  $m_r := m_r - \{(pred, i + 1)\}$ 
                                                    //@  $H := -1$ 
30:   end if
31:   return valueIsIn
32: end procedure

```

Fig. 11. Algorithm for searching on a concurrent lock-coupling skiplist

```

1: procedure INSERT(SkipList sl, Value newval)
2:   Vector(Node*) upd[0..K - 1]           //@ mrng  $m_r := \emptyset$ 
                                           //@  $L := sl.head$ 
                                           //@  $U := sl.tail$ 
                                           //@  $H := K - 1$ 

3:   lvl := randomLevel(K)
4:   Node* pred := sl.head
5:   pred.locks[K - 1].lock()             //@  $m_r := m_r \cup \{(pred, K - 1)\}$ 
6:   Node* curr := pred.next[K - 1]
7:   curr.locks[K - 1].lock()             //@  $m_r := m_r \cup \{(curr, K - 1)\}$ 
8:   for  $i := K - 1$  downto 0 do
9:     if  $i < K - 1$  then
10:      pred.locks[i].lock()             //@  $m_r := m_r \cup \{(pred, i)\}$ 
                                           //@  $U := curr$ 
11:      curr := pred.next[i]
12:      curr.locks[i].lock()             //@  $m_r := m_r \cup \{(curr, i)\}$ 
13:      if  $i \geq lvl$  then
14:        pred.next[i + 1].locks[i + 1].unlock()
                                           //@  $m_r := m_r - \{(pred.next[i + 1], i + 1)\}$ 
15:        pred.locks[i + 1].unlock()     //@  $m_r := m_r - \{(pred, i + 1)\}$ 
                                           //@  $H := i$ 
16:      end if
17:    end if
18:    while curr.val < newval do
19:      pred.locks[i].unlock()           //@  $m_r := m_r - \{(pred, i)\}$ 
                                           //@ if ( $i = lvl$ ) { $L := curr$ }
20:      pred := curr
21:      curr := pred.next[i]
22:      curr.locks[i].lock()             //@  $m_r := m_r \cup \{(curr, i)\}$ 
23:    end while
24:    upd[i] := pred
25:  end for
26:  Bool valueWasIn := (curr.val = newval)
27:  if valueWasIn then
28:    for  $i := 0$  to lvl do
29:      upd[i].next[i].locks[i].unlock() //@  $m_r := m_r - \{(upd[i].next[i], i)\}$ 
30:      upd[i].locks[i].unlock()       //@  $m_r := m_r - \{(upd[i], i)\}$ 
31:    end for
32:  else
33:    x := CreateNode(lvl, newval)
34:    for  $i := 0$  to lvl do
35:      x.next[i] := upd[i].next[i]
36:      upd[i].next[i] := x             //@  $sl.r := sl.r \cup \{(x, i)\}$ 
37:      x.next[i].locks[i].unlock()     //@  $m_r := m_r - \{(x.next[i], i)\}$ 
38:      upd[i].locks[i].unlock()       //@  $m_r := m_r - \{(upd[i], i)\}$ 
39:    end for
40:  end if
                                           //@  $H = -1$ 

41:  return  $\neg valueWasIn$ 
42: end procedure

```

Fig. 12. Algorithm for insertion on a concurrent lock-coupling skiplist

```

1: procedure REMOVE(SkipList sl, Value v)
2:   Vector < Node* > upd[0..K - 1]           //@ mrngn  $m_r := \emptyset$ 
                                                //@  $L := sl.head$ 
                                                //@  $U := sl.tail$ 
                                                //@  $H := K - 1$ 

3:   Node* pred := sl.head
4:   pred.locks[K - 1].lock()                   //@  $m_r := m_r \cup \{(pred, K)\}$ 
5:   Node* curr := pred.next[K - 1]
6:   curr.locks[K - 1].lock()                   //@  $m_r := m_r \cup \{(curr, K)\}$ 
7:   Node* aux
8:   for  $i := K - 1$  downto 0 do
9:     if  $i < K - 1$  then
10:      pred.locks[i].lock()                   //@  $m_r := m_r \cup \{(pred, i)\}$ 
                                                //@  $U := curr$ 
11:      aux := curr
12:      curr := pred.next[i]
13:      curr.locks[i].lock()                   //@  $m_r := m_r \cup \{(curr, i)\}$ 
14:      if aux.val  $\neq e$  then
15:        aux.locks[i + 1].unlock()
16:        pred.locks[i + 1].unlock()           //@  $H := i$ 
17:      end if
18:    end if
19:    while curr.val < v do
20:      pred.locks[i].unlock()                 //@  $m_r := m_r - \{(pred, i)\}$ 
                                                //@ if (aux.val  $\neq e$ ) { $L := curr$ }
21:      pred := curr
22:      curr := pred.next[i]
23:      curr.locks[i].lock()                   //@  $m_r := m_r \cup \{(curr, i)\}$ 
24:    end while
25:    upd[i] := pred
26:  end for
27:  for  $i := K - 1$  downto 0 do
28:    if upd[i].next[i] = curr  $\wedge$  curr.val = v then
29:      upd[i].next[i] := curr.next[i]       //@  $sl.r := sl.r - \{(curr, i)\}$ 
30:      curr.locks[i].unlock()                 //@  $m_r := m_r - \{(curr, i)\}$ 
31:    else
32:      upd[i].next[i].locks[i].unlock()     //@  $m_r := m_r - \{upd[i].next[i], i\}$ 
33:    end if
34:    upd[i].locks[i].unlock()                 //@  $m_r := m_r - \{upd[i], i\}$ 
                                                //@  $H := i - 1$ 
35:  end for
36:  Bool value WasIn := (curr.val = v)
37:  if value WasIn then
38:    free (curr)
39:  end if
40:  return value WasIn
41: end procedure

```

Fig. 13. Algorithm for deletion on a concurrent lock-coupling skiplist

D Non Termination Under Weak Fairness

Here we prove that the implementation given in Fig. 3 does not ensure termination of all threads under the assumption of weak-fairness. For such purpose, consider the skiplist depicted in Fig. 14.




Fig. 14. An example of skiplist

We use values $L^{[t]}$, $U^{[t]}$ and $H^{[t]}$ to denote the section of the skiplist that can be potentially modified by thread t . $L^{[t]}$ describes the lowest address bound while $U^{[t]}$ denotes the upper address bound. Meanwhile, $H^{[t]}$ represents the higher level to be modified. Considering the skiplist at Fig. 14 we consider two thread running concurrently. Thread 1 (called T_1) will insert value 14 with height 1, while thread 2 (denoted T_2) inserts value 16 with height 2.




Fig. 15. Progress of T_1 towards insertion of value 14

We start executing T_1 . This thread grabs the lock at level 2 on node $-\infty$ and 18, as shown in Fig. 15(a). As it detects that it has gone beyond the position

where 16 should be inserted, it decides to go down a level. The algorithm proceeds as depicted in Fig. 15(b) and 15(c).

At this moment, T_2 starts its execution. Fig. 16 shown the progress made by T_2 toward the insertion of a level 2 node with value 16.

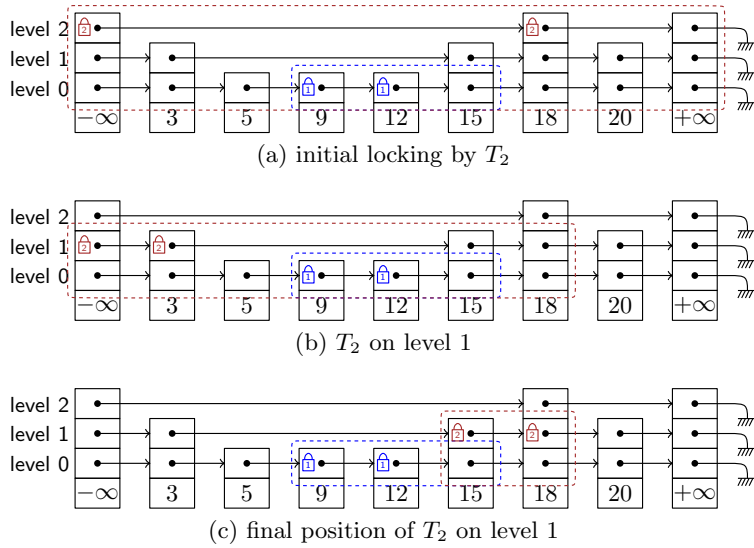


Fig. 16. Progress of T_2 towards insertion of value 16

Notice that the potentially modifiable regions by thread T_1 and T_2 intersects, as shown in Fig. 16(c). In this case, it is quite easy to see that under the assumption of weak-fairness, if T_2 continuously perform the same insertion, it prevents T_1 from progressing. However, no problem exists under the assumption of strong fairness, since T_1 is not continuously enabled. Notice that T_1 becomes disable every time T_2 gets the lock at level 0 on node 15.

E Optimistic Lock-Coupling Skiplist

```

1: procedure INSERT(SkipList sl, Value newval)
2:   Vector(Node*) upd[0..K - 1]           //@ mrng  $m_r := \emptyset$ 
                                           //@  $L := sl.head$ 
                                           //@  $U := sl.tail$ 
                                           //@  $H := K - 1$ 

3:   lvl := randomLevel(K)
4:   Node* curr := sl.head
5:   curr.locks[K - 1].lock()             //@  $m_r := m_r \cup \{(curr, K - 1)\}$ 
6:   Node* pred
7:   for  $i := K - 1$  downto 0 do
8:     if  $i < K - 1$  then
9:       curr.locks[i].lock()             //@  $m_r := m_r \cup \{(curr, i)\}$ 
                                           //@  $U := curr.next[i + 1]$ 
10:      if  $i \geq lvl$  then
11:        curr.locks[i + 1].unlock()       //@  $m_r := m_r - \{(curr, i + 1)\}$ 
                                           //@  $H := i$ 
12:      end if
13:    end if
14:    while curr.next[i].val < newval do
15:      curr.next[i].locks[i].lock()       //@  $m_r := m_r \cup \{(curr.next, i)\}$ 
16:      pred := curr
17:      curr := curr.next[i]
18:      pred.locks[i].unlock()             //@  $m_r := m_r - \{(pred, i)\}$ 
                                           //@ if ( $i = lvl$ )  $\{L := curr\}$ 
19:    end while
20:    upd[i] := curr
21:  end for
22:  Bool valueWasIn := (curr.next[i].val = newval)
23:  if valueWasIn then
24:    for  $i := 0$  to lvl do
25:      upd[i].locks[i].unlock()           //@  $m_r := m_r - \{(upd[i], i)\}$ 
26:    end for
27:  else
28:    x := CreateNode(lvl, newval)
29:    for  $i := 0$  to lvl do
30:      x.next[i] := upd[i].next[i]
31:      upd[i].next[i] := x                 //@  $sl.r := sl.r \cup \{(x, i)\}$ 
32:      upd[i].locks[i].unlock()           //@  $m_r := m_r - \{(upd[i], i)\}$ 
33:    end for
34:  end if                                     //@  $H = -1$ 

35:  return  $\neg valueWasIn$ 
36: end procedure

```

Fig. 17. Optimistic algorithm for insertion on a concurrent lock-coupling skiplist

F Pessimistic Lock-Coupling Skiplist

```

1: procedure SEARCH(SkipList sl, Value v)
2:   int i := K - 1                                //@ mr := ∅
                                                    //@ L := sl.head
                                                    //@ U := sl.tail
                                                    //@ H := K - 1

3:   Node* pred := sl.head
4:   pred.locks[i].lock()                          //@ mr := mr ∪ {(pred, i)}
5:   Node* curr := pred.next[i]
6:   curr.locks[i].lock()                          //@ mr := mr ∪ {(curr, i)}
7:   while 0 ≤ i ∧ curr.val ≠ v do
8:     if i < K - 1 then
9:       pred.locks[i].lock()                      //@ mr := mr ∪ {(pred, i)}
                                                    //@ U := curr
10:      curr := pred.next[i]
11:      curr.locks[i].lock()                       //@ mr := mr ∪ {(curr, i)}
12:      pred.next[i + 1].locks[i + 1].unlock()
                                                    //@ mr := mr - {(pred.next[i + 1], i + 1)}
13:      pred.locks[i + 1].unlock()                 //@ mr := mr - {(pred, i + 1)}
                                                    //@ H := i
14:     end if
15:     while curr.val < v do
16:       pred.locks[i].unlock()                   //@ mr := mr - {(pred, i)}
                                                    //@ L := curr
17:       pred := curr
18:       curr := pred.next[i]
19:       curr.locks[i].lock()                      //@ mr := mr ∪ {(curr, i)}
20:     end while
21:     i := i - 1
22:   end while
23:   Bool valueIsIn := (curr.val = v)
24:   if i = K - 1 then
25:     curr.locks[i].unlock()                      //@ mr := mr - {(curr, i)}
                                                    //@ U := L
26:     pred.locks[i].unlock()                      //@ mr := mr - {(pred, i)}
                                                    //@ H := -1
27:   else
28:     curr.locks[i + 1].unlock()                  //@ mr := mr - {(curr, i + 1)}
                                                    //@ U := L
29:     pred.locks[i + 1].unlock()                  //@ mr := mr - {(pred, i + 1)}
                                                    //@ H := -1
30:   end if
31:   return valueIsIn
32: end procedure

```

Fig. 18. Pessimistic algorithm for searching on a concurrent lock-coupling skiplist

G Proving Termination of All Threads

In this section we describe how to extend our decision procedure in order to reason about the termination of all threads. We use the (L, U, H) regions introduced in Section C to represent the minimum portion of the skiplist we are sure a thread can potentially modify. We say that two threads conflict when the (L, U, H) of one is included into the (L, U, H) of the other one. The idea of the proof is that if a thread does not conflict with other thread, then it terminates. For such purpose we consider the optimistic and pessimistic version of the algorithms presented in Section E and F respectively.

To apply this idea, we need to define a $min : \text{mem} \times \text{addr} \times \text{addr} \times \text{level}_K \rightarrow \text{setth}$ function which, given a memory representation and a skiplist regions denoted by a (L, U, H) triple, returns the set of thread identifiers whose (L, U, H) is contained in such region. To aid the definition of min , we extend TSL_K with a new function $skipRgTh : \text{mem} \times \text{addr} \times \text{addr} \times \text{level}_K \times \text{addr} \times \text{addr} \times \text{level}_K \rightarrow \text{setth}$, defined into Σ_{bridge} . The function $skipRgTh$ returns the set of thread identifiers that contain a lock on any node in the list that goes from a to b considering level l of the skiplist. Besides, we ask that the (L, U, H) of all thread identifiers in the set returned by $skipRgTh$ are contained into the region denoted by the L , U and H given as parameter.

Before we proceed with the definition of such function, there are some assumptions we need to make. In particular, we require that (L, U, H) values are not kept locally, as depicted in previous algorithms. We need them to be shared among all threads. Therefore, we assume that the *SkipList* class is extended with ghost mappings m_L , m_U and m_H which goes from *thid* to *addr*. These mappings are updated by the algorithms such that at every moment, for every thread identifier t , $(m_L(t), m_U(t), m_H(t))$ matches with the (L, U, H) of thread t .

We begin extending our decision procedure by adding a *cont* predicate to *PATH*. This predicate holds when a region (L', U', H') is contained into a region (L, U, H) :

$$\begin{aligned} cont : \text{mem} \times \text{addr} \times \text{addr} \times \text{level}_K \times \text{addr} \times \text{addr} \times \text{level}_K \\ cont(h, L', U', H', L, U, H) \hat{=} & reach_K(h, L, L', 0) \quad \wedge \\ & reach_K(h, U', U, 0) \quad \wedge \\ & (K' < K \vee K' = K) \end{aligned}$$

Basically, predicate *cont* says that considering level 0 of the skiplist, from L it is possible to reach L' , from U' it can be reached U and K' should be lower or equal to K . Of course, we can ensure that (L', U', H') represents a valid rectangular region by adding the constraint $reach_K(h, L', U', 0)$ to the predicate. Then, we can extend *PATH* further by adding a recursive *contTh* predicate which takes as argument:

- a memory layout, h
- a lower bound address, L
- an upper bound address, U

- a bound on the skiplist’s level, H
- an initial address, a
- a final address, b
- a level, l
- a set of thread identifiers, s

Then, the predicate $contTh(h, L, U, H, a, b, l, s)$ holds whether s is the set of threads identifiers owing a lock in the list that goes from address a to b , through level l of the skiplist. Moreover, we require that the (L, U, H) of each thread identifiers in s must be a subset of the region denoted by the L, U and K given as parameter. Formally, we define the predicate $contTh$ as:

$contTh : mem \times addr \times addr \times level_K \times addr \times addr \times level_K \times setth$

$$\left(\begin{array}{l} t = h[a].lockid[l] \wedge \\ t \neq \emptyset \quad \wedge \\ contained \end{array} \right) \rightarrow contTh(h, L, U, H, a, a, l, \{t\}_T)$$

$$\left(\begin{array}{l} t = h[a].lockid[l] \quad \wedge \\ (t = \emptyset) \vee (t \neq \emptyset \wedge \neg contained) \end{array} \right) \rightarrow contTh(h, L, U, H, a, a, l, \emptyset_T)$$

$$\left(\begin{array}{l} t = h[a].lockid[l] \quad \wedge \\ h[a].next[l] = a' \quad \wedge \\ contTh(h, L, U, H, a', b, l, s) \wedge \\ t \neq \emptyset \quad \wedge \\ contained \end{array} \right) \rightarrow contTh(h, L, U, H, a, b, l, s \cup_T \{t\}_T)$$

$$\left(\begin{array}{l} t = h[a].lockid[l] \quad \wedge \\ h[a].next[l] = a' \quad \wedge \\ contTh(h, L, U, H, a', b, l, s) \wedge \\ (t == null \vee \neg contained) \end{array} \right) \rightarrow contTh(h, L, U, H, a, b, l, s)$$

where $contained \hat{=} cont(h, m_L(t), m_U(t), m_H(t), L, U, H)$

Notice that the definition of $contTh$ is similar to the one of $reach_K$. Finally, we need to add to GAP the following equivalence:

$$contTh(h, L, U, H, a, b, l, s) \leftrightarrow skipRgTh(h, L, U, H, a, b, l) = s$$

We can now use function $skipRgTh$ to define the function min we required. We can do so through the following equivalence:

$$\begin{aligned}
t = \min(h, L, U, H) &\leftrightarrow s = \text{skipRgTh}(h, L, U, H) && \wedge \\
&t \in s && \wedge \\
&\emptyset = \text{skipRgTh}(h, m_L(t), m_U(t), m_H(t))
\end{aligned}$$

As usual, every time we find a literal of the form $t = \min(h, L, U, H)$ we proceed to replace it by the equivalent definition we have given above. Then, we replace the invocations of *skipRgTh* by invocations to *contTh*. We can finally unroll the occurrences of *contTh* according to its recursive definition up to the bound given by the small model property.

TODO: Remains to verify whether the SMP still holds.

```

1: procedure INSERT(SkipList sl, Value newval)
2:   Vector(Node*) upd[0..K - 1]           //@ mrgn  $m_r := \emptyset$ 
                                           //@  $L := sl.head$ 
                                           //@  $U := sl.tail$ 
                                           //@  $H := K - 1$ 

3:   lvl := randomLevel(K)
4:   Node* pred := sl.head
5:   pred.locks[K - 1].lock()             //@  $m_r := m_r \cup \{(pred, K - 1)\}$ 
6:   Node* curr := pred.next[K - 1]
7:   curr.locks[K - 1].lock()             //@  $m_r := m_r \cup \{(curr, K - 1)\}$ 
8:   Node* cover
9:   for  $i := K - 1$  downto 0 do
10:    if  $i < K - 1$  then
11:      pred.locks[i].lock()             //@  $m_r := m_r \cup \{(pred, i)\}$ 
                                           //@  $U := pred.next[i + 1]$ 
12:      if  $i \geq lvl$  then
13:        pred.locks[i + 1].unlock()     //@  $m_r := m_r - \{(pred, i + 1)\}$ 
                                           //@  $H := i$ 
14:      end if
15:      if  $i < k - 2 \wedge i > lvl - 2$  then
16:        cover.locks[i + 2].unlock()
17:      end if
18:      cover := curr
19:      curr := pred.next[i]
20:      curr.locks[i].lock()             //@  $m_r := m_r \cup \{(curr, i)\}$ 
21:    end if
22:    while curr.val < newval do
23:      pred.locks[i].unlock()           //@  $m_r := m_r - \{(pred, i)\}$ 
                                           //@ if ( $i = lvl$ ) { $L := curr$ }
24:      pred := curr
25:      curr := pred.next[i]
26:      curr.locks[i].lock()             //@  $m_r := m_r \cup \{(curr, i)\}$ 
27:    end while
28:    upd[i] := pred
29:  end for
30:  Bool valueWasIn := (curr.val = newval)
31:  if valueWasIn then
32:    for  $i := 0$  to lvl do
33:      upd[i].next[i].locks[i].unlock()  //@  $m_r := m_r - \{(upd[i].next[i], i)\}$ 
34:      upd[i].locks[i].unlock()         //@  $m_r := m_r - \{(upd[i], i)\}$ 
35:    end for
36:  else
37:    x := CreateNode(lvl, newval)
38:    for  $i := 0$  to lvl do
39:      x.next[i] := upd[i].next[i]
40:      upd[i].next[i] := x               //@  $sl.r := sl.r \cup \{(x, i)\}$ 
41:      x.next[i].locks[i].unlock()       //@  $m_r := m_r - \{(x.next[i], i)\}$ 
42:      upd[i].locks[i].unlock()       //@  $m_r := m_r - \{(upd[i], i)\}$ 
43:    end for
44:  end if
                                           //@  $H = -1$ 

45:  return  $\neg valueWasIn$ 
46: end procedure

```

Fig. 19. Pessimistic algorithm for insertion on a concurrent lock-coupling skiplist

```

1: procedure REMOVE(SkipList sl, Value v)
2:   Vector < Node* > upd[0..K - 1]           //@ mrng mr := ∅
                                                //@ L := sl.head
                                                //@ U := sl.tail
                                                //@ H := K - 1

3:   Node* pred := sl.head
4:   pred.locks[K - 1].lock()                 //@ mr := mr ∪ {(pred, K)}
5:   Node* curr := pred.next[K - 1]
6:   curr.locks[K - 1].lock()                 //@ mr := mr ∪ {(curr, K)}
7:   for i := K - 1 downto 0 do
8:     if i < K - 1 then
9:       pred.locks[i].lock()                 //@ mr := mr ∪ {(pred, i)}
                                                //@ U := curr
10:      curr := pred.next[i]
11:      curr.locks[i].lock()                 //@ mr := mr ∪ {(curr, i)}
12:      if pred.next[i + 1].val ≠ e then
13:        pred.next[i + 1].locks[i + 1].unlock()
14:        pred.locks[i + 1].unlock()         //@ H := i
15:      end if
16:    end if
17:    while curr.val < v do
18:      pred.locks[i].unlock()                 //@ mr := mr - {(pred, i)}
                                                //@ if (aux.val ≠ e) {L := curr}
19:      pred := curr
20:      curr := pred.next[i]
21:      curr.locks[i].lock()                 //@ mr := mr ∪ {(curr, i)}
22:    end while
23:    upd[i] := pred
24:  end for
25:  for i := K - 1 downto 0 do
26:    if upd[i].next[i] = curr ∧ curr.val = v then
27:      upd[i].next[i] := curr.next[i]       //@ sl.r := sl.r - {(curr, i)}
28:      curr.locks[i].unlock()                 //@ mr := mr - {(curr, i)}
29:    else
30:      upd[i].next[i].locks[i].unlock()     //@ mr := mr - {upd[i].next[i], i}
31:    end if
32:    upd[i].locks[i].unlock()                 //@ mr := mr - {(upd[i], i)}
                                                //@ H := i - 1
33:  end for
34:  Bool valueWasIn := (curr.val = v)
35:  if valueWasIn then
36:    free (curr)
37:  end if
38:  return valueWasIn
39: end procedure

```

Fig. 20. Pessimistic algorithm for deletion on a concurrent lock-coupling skiplist