

Verifying Concurrent List–Manipulating Programs by LTL Model Checking

Joost–Pieter Katoen and Thomas Noll and Stefan Rieger

RWTH Aachen University
Software Modeling and Verification Group
52056 Aachen, Germany
`{katoen,noll,rieger}@cs.rwth-aachen.de`

Abstract

We present a novel approach to the verification of concurrent pointer–manipulating programs which operate on singly–linked lists. By abstracting from chains (i.e., non–interrupted sublists) in the heap, we obtain a finite–state representation of all possible executions of a given program. The combination of a simple pointer logic for expressing heap properties and of temporal operators then allows us to employ standard LTL model checking techniques. The usability of this approach is demonstrated by establishing correctness properties of a producer/consumer system and of a concurrent garbage collector.

Keywords: Software Model Checking, Abstraction, Heap Verification, Shape Analysis, LTL, Lists, Pointer Programs

1 Introduction

Techniques for the verification of elementary properties of concurrent pointer programs are indispensable. Programming with pointers is error–prone with potential pitfalls such as dereferencing null pointers and the creation of memory leaks. Pointer programming becomes even more vulnerable in a concurrent setting where data structures such as linked lists and trees are manipulated and inspected by several threads.

This paper presents a model–checking approach to the verification of concurrent programs that manipulate singly–linked lists. Existing approaches either make use of non–standard logics, advanced model–checking procedures or extended versions of Hoare logics with accompanying deduction techniques (see Sect. 6 about related work). In contrast, the approach advocated in this paper stays within the realm of traditional (linear–time) model checking. This facilitates the usage of standard (LTL) model checkers for validating temporal properties addressing absence of memory leaks, dereferencing of null pointers, dynamic creation of cells, and simple and position–dependent aliasing.

Our approach is illustrated by considering a simple concurrent programming language that besides the usual control structures offers primitives for pointer manipulation, cell creation and destruction, and (guarded) atomic regions that allow concurrency control constructs such as test–and–set primitives and monitors. An operational semantics is provided in terms of labeled transition systems in which states are equipped with a graph structure representing the current list configuration. List abstraction exploits a variant of summary nodes [45] that represent more than M chained list cells where constant M is directly obtained from the formula to be checked. Each configuration is shown to have a canonical representation (up to isomorphism). The abstract semantics of any concurrent program in our language is finite, obtained in a fully mechanized manner, and keeps the minimal “distance” between program variables and summary nodes invariant. Over–approximation oc-

curs in a very controlled manner; only assignments may yield nondeterminism as variables may get “too close” to summary nodes.

Properties are expressed in a first-order linear-time temporal logic (LTL) that is enriched with assertions on singly-linked lists such as reachability of cells, aliasing, and freshness of cells. Our logic is similar in spirit to NTL [19,20] and ETL [49]. Opposed to NTL, we avoid the use of temporal operators inside quantification. In this way, involved mechanisms to keep track of the identities of individual cells are not needed. As a result, standard LTL model checking algorithms can be employed. The differences with ETL are more of a technical nature. ETL has a three-valued interpretation, whereas our logical interpretation is a standard binary one. Moreover, ETL-formulas are translated in first-order logic with transitive closure for the evaluation on a trace, whereas in our case traces are generated by labeled transition systems and used in standard LTL model checking. The feasibility of our approach is shown by considering the verification of a simple concurrent garbage collection program. Furthermore a prototypical tool is currently under development for experimenting with real-life examples.

Please note that due to space constraints most of the proofs could not be included in this paper.

2 A List-Manipulating Programming Language

Given a universe PV of program variables, we define the set of *list-manipulating programs* (LM-programs) to be given by the following grammar (where $v_i, v \in PV$):

$$\begin{aligned} \text{LMP} &::= \mathbf{var} \ v_1, \dots, v_k (\text{Stmt}_1 \parallel \dots \parallel \text{Stmt}_l) \\ \text{Stmt} &::= \mathbf{skip} \mid \mathbf{signal} \mid v := \text{PExp} \mid *v := \text{PExp} \mid \text{Stmt}; \text{Stmt} \\ &\quad \mid \mathbf{if} \ \text{BExp} \ \mathbf{then} \ \text{Stmt} \ \mathbf{else} \ \text{Stmt} \ \mathbf{fi} \mid \mathbf{while} \ \text{BExp} \ \mathbf{do} \ \text{Stmt} \ \mathbf{od} \\ &\quad \mid \mathbf{new}(\text{PExp}) \mid \mathbf{del}(\text{PExp}) \mid \langle \text{BExp} : \text{Stmt} \rangle \\ \text{PExp} &::= \mathit{nil} \mid v \mid *v \mid \&v \\ \text{BExp} &::= \text{tt} \mid \text{ff} \mid \text{PExp} = \text{PExp} \mid \text{BExp} \wedge \text{BExp} \mid \neg \text{BExp} \end{aligned}$$

$V(\pi) := \{v_1, \dots, v_k\}$ denotes the set of variables for $\pi \in \text{LMP}$.

An LM-program thus consists of a declaration of global program variables and a series of statements to be executed in parallel. Each of these statements can either be a pointer assignment, a sequence of statements, a control structure, or a special statement such as **signal** which sets a global signal flag that can be tested in the logic, **new/del** for dynamic creation or deletion of objects at runtime (possibly leading to an unbounded number of allocated heap cells) and guarded atomic regions. If the Boolean guard g in $\langle g : s \rangle$ is true, s is executed atomically, i.e., with no interference by other processes. If g is evaluated to false, the process is blocked (until g becomes true).

```

var  $x, y, z$ (
  while  $\text{tt}$  do  $\langle \text{tt} :$ 
    if  $x = \mathit{nil}$ 
      then  $\mathbf{new}(y); x := y$ 
      else  $\mathbf{new}(*y); y := *y$ 
    fi
   $\rangle$  od
   $\parallel$  while  $\text{tt}$  do  $\langle x \neq \mathit{nil} :$ 
     $z := x; x := *x; \mathbf{del}(z)$ 
   $\rangle$  od
)
```

Fig. 1. Producer/Consumer

Pointer expressions comprise the special constant nil denoting an undefined

pointer value, a program variable, the dereferencing or referencing of a program variable. Note that for simplicity we do not allow arbitrary dereferencing depths; those can be emulated using a sequence of assignments within an atomic region.

Example 2.1 Figure 1 shows an LM-program implementing a producer inserting objects and a consumer deleting objects at the end (pointed to by y) and beginning (pointed to by x) of a queue, respectively. If the queue is empty the consumer cannot proceed due to the guard $x \neq \text{nil}$ until the producer has inserted at least one object. Insertion and deletion are executed atomically to prevent interferences.

Definition 2.2 A *heap configuration* of a program $\pi \in \text{LMP}$ is a tuple $\gamma = (N, A, \mu, F)$ with a set of nodes $N \supseteq V(\pi)$, a set of abstract nodes $A \subseteq N \setminus PV$, a successor function $\mu : N \rightarrow N_{\text{nil}}$ (where $N_{\text{nil}} := N \cup \{\text{nil}\}$), and a set of flags $F \subseteq \{\text{err}, \text{dl}, \text{leak}, \text{signal}, \text{new}\}$.

Let $\mu^* : 2^N \rightarrow 2^N$ with $\mu^*(X) := \{n \in N \mid \exists k \in \mathbb{N}, \exists n' \in X : \mu^k(n') = n\}$ be the transitive closure of μ , i.e. all nodes reachable from a node in X (and X itself).

Thus the nodes represent both the dynamic objects created and deleted at runtime and the static program variables (which cannot be deleted). Edges, as formalized by the μ -function, encode the *points-to* information of a specific program state. The set A of abstract nodes will later be used for our abstraction technique and will be empty throughout the current section. Finally the flags give special information about a state, e.g., whether a runtime error or memory leak occurred, a new node was created, or the signal bit has been set using the **signal** command.

To ensure the finiteness of our abstraction we will automatically delete those heap nodes that are not reachable from the program variables. This is accomplished by the following *garbage collection* mapping. Whenever it removes an unreachable node, it sets the leak flag indicating a potential memory leak.

Definition 2.3 For $\gamma = (N, A, \mu, F)$ we define $\gamma \downarrow := (N', A \cap N', \mu \upharpoonright N', F \cup \{\text{leak} \mid (N \setminus N') \neq \emptyset\})$ where $N' = \mu^*(PV)$.

Γ denotes the set of all *garbage-free heap configurations*, i.e., $\forall \gamma \in \Gamma : \gamma \downarrow = \gamma$, and $\Gamma_c \subseteq \Gamma$ denotes the set of all concrete configurations, i.e., those with $A_\gamma = \emptyset$.

From now on we will always assume garbage freeness when mentioning heap configurations. This enforces a bound on the maximal number of incoming edges for a node (essentially the number of program variables).

Definition 2.4 Let $\gamma = (N, \emptyset, \mu, F) \in \Gamma_c$. Then we define the semantics of pointer expressions $\llbracket \cdot \rrbracket : \text{PExp} \rightarrow N_{\text{nil}}$ by¹:

$$\begin{aligned} \llbracket \text{nil} \rrbracket &:= \text{nil} & \llbracket *v \rrbracket &:= \mu(\llbracket v \rrbracket) \\ \llbracket v \rrbracket &:= \mu(v) & \llbracket \&v \rrbracket &:= v \end{aligned}$$

The semantics of Boolean expressions $\llbracket \cdot \rrbracket : \text{BExp} \rightarrow \mathbb{B}$ is standard and strict². Note that Def. 2.2 implies that $\mu(\text{nil}) = \perp$ and so $\llbracket \cdot \rrbracket$ can indeed yield undefined results for both pointer and Boolean expressions.

Definition 2.5 For $\pi = \text{var } v_1, \dots, v_k : (s_1 \parallel \dots \parallel s_l) \in \text{LMP}$ the *concrete operational semantics* is given by a transition system $T_\pi^c = (Q, q_0, \text{lab}, \rightarrow)$ with a set of states

¹ \rightarrow denotes a partial function and \perp the undefined value.

² One undefined operand yields an undefined expression.

$Q \subseteq \Gamma_c \times \text{Stmt}_\diamond(\{\|\}\text{Stmt}_\diamond)^*$ where $\text{Stmt}_\diamond = \text{Stmt} \cup \{\diamond\}\text{Stmt} \cup \{\varepsilon\}$, an initial state $q_0 = ((N_0, \emptyset, \mu_0, \emptyset), s_1 \| \dots \| s_l)$ where N_0 and μ_0 represent the “input heap”, a labeling $lab : Q \rightarrow \Gamma_c$ with $\forall (\gamma, s) \in Q : lab((\gamma, s)) = \gamma$, and a transition relation $\rightarrow \subseteq Q \times Q$.

In the following we will use the abbreviations \hat{F} for $F \setminus \{\text{signal}, \text{new}, \text{leak}\}$ and noerr for $\{\text{err}, \text{dl}\} \cap F = \emptyset$. γ_{err} and γ_{dl} will denote pointer error and deadlock states. Most transition rules are straightforward, thus here we will only consider some interesting examples.

$$\frac{\llbracket g \rrbracket = 1 \quad \gamma, s \rightarrow \gamma', s' \quad \text{noerr}}{\gamma, \langle g : s \rangle \rightarrow \gamma', \diamond s'} \quad (1)$$

$$\frac{\gamma, s \rightarrow \gamma', s' \quad s' \neq \varepsilon \quad \text{noerr}}{\gamma, \diamond s \rightarrow \gamma', \diamond s'} \quad \frac{\gamma, s \rightarrow \gamma', \varepsilon \quad \text{noerr}}{\gamma, \diamond s \rightarrow \gamma', \varepsilon} \quad (2)$$

$$\frac{\exists j \text{ s.t. } \gamma, s_j \rightarrow \gamma', s'_j \quad \forall i \neq j : \nexists s'_i \text{ s.t. } s_i = \diamond s'_i \quad \text{noerr}}{\gamma, s_1 \| \dots \| s_k \rightarrow \gamma', s_1 \| \dots \| s'_j \| \dots \| s_k} \quad (3)$$

$$\frac{\nexists j \text{ s.t. } \gamma, s_j \rightarrow \gamma', s'_j \quad \exists j : s_j \neq \varepsilon \quad \text{noerr}}{\gamma, s_1 \| \dots \| s_k \rightarrow \gamma_{\text{dl}}, \varepsilon} \quad (4)$$

$$\frac{}{\gamma, \varepsilon \| \dots \| \varepsilon \rightarrow \gamma, \varepsilon \| \dots \| \varepsilon} \quad (5)$$

$$\frac{\llbracket \alpha \rrbracket \neq \perp \quad \text{noerr}}{(N, A, \mu, F), v := \alpha \rightarrow (N, A, \mu[v/\llbracket \alpha \rrbracket], \hat{F})\downarrow, \varepsilon} \quad (6)$$

$$\frac{\text{noerr}}{(N, A, \mu, F), \mathbf{new}(v) \rightarrow (N \uplus \{n_{\text{new}}\}, A, \mu[v/n_{\text{new}}], \hat{F} \cup \{\text{new}\})\downarrow, \varepsilon} \quad (7)$$

$$\frac{\llbracket \alpha \rrbracket \in N \setminus PV \quad \text{noerr}}{(N, A, \mu, F), \mathbf{del}(\alpha) \rightarrow (N \setminus \{\llbracket \alpha \rrbracket\}, A, \mu[\llbracket \alpha \rrbracket / \perp, \mu^{-1}(\llbracket \alpha \rrbracket) / \text{nil}], \hat{F})\downarrow, \varepsilon} \quad (8)$$

Some remarks on the transition rules are in order. The leak, signal, and new flags are reset after each transition; they are only activated in the state directly following the corresponding “event”.

Regarding the concurrency rules we need to take care of the special semantics of the atomic regions. If a process is executing such a statement it must not be interrupted, and therefore the corresponding state is marked with \diamond (rule 1). The interleaving rule 3 excludes that any other than process j is in an atomic region. If no process can proceed (all are blocked) we reach the special deadlock state (rule 4). If all processes are terminated or an error or deadlock state is reached the program loops to ensure that all paths in the transition system are infinite (rule 5).

The treatment of assignments (rule 6) and the **new** statement (rule 7) is again straightforward, we though have to keep in mind in the first case that runtime errors might occur (dereferencing of *nil* pointers) and that garbage may be generated. Rule 8 handles the deletion of nodes. Please note that the next-pointers of the predecessors of the deleted node are set to *nil* (mainly to avoid case distinctions for undefined expressions in the semantics).

We conclude that for the producer/consumer example (Fig. 1) the state space

becomes infinite when applying the operational semantics as defined above.

3 State–Space Abstraction

As we have seen in the previous section the state space of LM–programs can get infinite even for simple example programs making standard verification methods inapplicable. To tackle the problem we use abstraction techniques to generate an *abstract transition system* that incorporates the behavior of the concrete one, i.e., whose runs cover all concrete ones. This approach is correct but generally incomplete: although we can conclude from the satisfaction of a property in the abstract state space its validity in the concrete case, the inverse is impossible though. But since the abstraction is parameterized via a global constant $M \in \mathbb{N}$ we can refine the abstraction depending on our needs. For a given $M > 0$ we set $\mathbb{M} := \{0, 1, \dots, M, \star\}$, where \star represents all values greater than M .

Chain Abstraction

The main idea of our abstraction is to summarize subgraphs of a configuration into summary nodes [45], which will be exactly those contained in the A –component of a heap configuration. Summary nodes (also called abstract nodes) are not allowed to represent arbitrary structures but only so–called *chains*, i.e., non–interrupted lists. Our abstraction is based on [18,19] with the difference that nodes are either truly abstract or concrete, thus recording node multiplicities is not necessary.

Definition 3.1 Let $\gamma = (N, A, \mu, F) \in \Gamma$ be a configuration. A nonempty set of nodes $C \subseteq N$ is called a *chain* if either

- $|C| = 1$ and $C \subseteq PV$ or
- $C \cap PV = \emptyset$ and there exists a bijection $\pi : \{1, \dots, |C|\} \rightarrow C$ such that $\mu(\pi(i)) = \pi(i+1)$ for $i \in \{1, \dots, |C|\}$ and $\forall i \in \{2, \dots, |C|\} : |\mu^{-1}(\pi(i))| = 1$.

For a given chain C we will use the abbreviations $\overleftarrow{C} := \pi(1)$, and $\overrightarrow{C} := \pi(|C|)$. A chain is called *maximal* if no superset $C' \supset C$ is a chain.

Thus a chain is a sequence of pointer–connected nodes without interference of other incoming edges or a singleton set containing a program variable. It follows that the abstraction of chains preserves the graph structure. We will now introduce a type of functions, called *abstraction morphisms*, that is based on this concept.

Definition 3.2 Let $\gamma_i = (N_i, A_i, \mu_i, F_i) \in \Gamma$, $i \in \{1, 2\}$ be two heap configurations. An *abstraction morphism* $h : N_1 \rightarrow N_2$ satisfies for all $v \in PV \cap N_1$ and $n_i, n'_i \in N_i$:

1. $h(v) = v$
2. $h^{-1}(n_2)$ is a chain in N_1
3. $\mu_2(n_2) = n'_2 \Rightarrow \mu_1(\overrightarrow{h^{-1}(n_2)}) = \overleftarrow{h^{-1}(n'_2)}$
4. $\mu_1(n_1) = n'_1 \Rightarrow h(n_1) = h(n'_1) \vee \mu_2(h(n_1)) = h(n'_1)$
5. $n_2 \in A_2 \Leftrightarrow h^{-1}(n_2) \cap A_1 \neq \emptyset \vee |h^{-1}(n_2)| > M$
6. $F_1 = F_2$

We write $h : \gamma_1 \mapsto \gamma_2$ to denote that the abstraction morphism h abstracts γ_1 to γ_2 and $\gamma_2 \leq \gamma_1 \Leftrightarrow \exists h : \gamma_1 \mapsto \gamma_2$.

Abstraction morphisms abstract from concrete chains with minimal length $M+1$ (cond. 2 and 5). The preservation of the graph structure is enforced by conditions 3 and 4. Program variables, being special nodes, remain untouched (cond. 1).

Example 3.3 Figure 2 shows an abstraction morphism for $M = 1$. The dashed lines represent the mapping, and the black nodes denote the resulting abstract nodes. Note that for $M = 2$ the nodes 3 and 4 could not be projected onto the same abstract node (condition 5 of Def. 3.2). The chain $\{3, 4\}$ cannot be extended by node 5, since this node has two incoming edges which is only allowed for the first node of a chain. Although in this example the source configuration is concrete, this is of course not necessary by definition.

An important property of abstraction morphisms is their surjectivity. If, in addition a morphism is injective it becomes an isomorphism. Isomorphic configurations cannot be distinguished except for node naming, the graph structure is the same.




Fig. 2. An Abstraction Morphism

Canonical Configurations

Previously we have defined how configurations can be abstracted. It remains the problem that there can be different abstractions of a given source configuration. For this reason we need a normal form that implies uniqueness. In the following we define this normal form, assuming $\gamma = (N, A, \mu, F) \in \Gamma$.

Definition 3.4

- (i) Let $[N]_j := \{n \in N \mid \nexists v \in PV : \mu^k(v) = n, k < j\}$ be the set of nodes with a distance of at least j from the variable nodes. Analogously $[N]_j := N \setminus [N]_{j+1}$.
- (ii) A configuration γ is called *canonical* if $[N]_2 \cap A = \emptyset$ and for all maximal³ chains $C \subseteq [N]_3$ either $|C| = 1$ or $|C| \leq M \wedge C \cap A = \emptyset$. The set of all canonical configurations is denoted by Γ_{\natural} .

The notion of canonical configurations is quite intuitive: maximal chains are collapsed where possible but only up to a distance of three from variable nodes. The latter condition ensures that pointer expressions always evaluate to concrete nodes, which will simplify the definition of the abstract LMP semantics. The abstraction morphism in Fig. 2 yields a canonical configuration, as can be easily verified.

Theorem 3.5 (Existence) *For every $\gamma \in \Gamma$ with $[N]_2 \cap A = \emptyset$ there exists a $\gamma' \in \Gamma_{\natural}$ such that $\gamma' \leq \gamma$.*

It is easy to construct a morphism yielding a canonical configuration. It has to collapse maximal chains that are larger than M or contain abstract nodes, if they are sufficiently distant from the variable nodes. In the following we will call this morphism h_{\natural} . The precise definition does not matter as states the following theorem.

³ Here we refer to maximality in $[N]_3$.

Theorem 3.6 (Uniqueness) *Let $\gamma \in \Gamma$ and $\gamma_1, \gamma_2 \in \Gamma_{\natural}$ such that $h_1 : \gamma \mapsto \gamma_1$ and $h_2 : \gamma \mapsto \gamma_2$ are two abstraction morphisms. Then γ_1 and γ_2 are isomorphic.*

The proof of the uniqueness had to be omitted here. The consequence of these results is the appropriateness of canonical configurations as a normal form. The abstract semantics will operate on such configurations.

Abstract Semantics of List–Manipulating Programs

As already mentioned, our goal is to guarantee the correctness of our abstraction approach. This can be achieved by ensuring that every concrete execution of a given system can be “simulated” by an abstract computation, which necessarily introduces nondeterministic behavior on the abstract side.

Regarding the expression semantics nothing needs to be modified: in a canonical configuration, abstract nodes have a distance greater than two from the variable nodes such that every pointer expression refers to a concrete node. The expression semantics can therefore be chosen identical to the concrete case (Def. 2.4), now interpreted on canonical configurations.

Definition 3.7 Given a program $\pi = \mathbf{var} \ v_1, \dots, v_k : (s_1 \parallel \dots \parallel s_l) \in \text{LMP}$, its *abstract operational semantics* is defined by the labeled transition system $T_{\pi}^a = (Q, [q_0]_{\cong}, \text{lab}, \rightarrow)$ with state set $Q \subseteq \Gamma_{\natural}/_{\cong} \times \text{Stnt}_{\diamond}(\{\|\}\text{Stnt}_{\diamond})^*$, initial state q_0 as in Def. 2.5, labeling function $\text{lab} : Q \rightarrow \Gamma_{\natural}$ where $\forall (K, s) \in Q : \text{lab}((K, s)) = K$, and transition relation \rightarrow as specified by the following rules (we focus on the assignments, since the other rules are analogous to the concrete case, but operating on isomorphism congruence classes).

$$\frac{\alpha \notin *V(\pi) \quad \text{noerr}}{[(N, A, \mu, F)]_{\cong}, v := \alpha \rightarrow [h_{\natural}((N, A, \mu[v/\llbracket \alpha \rrbracket], \hat{F})\downarrow)]_{\cong}, \varepsilon} \quad (1)$$

$$\frac{\gamma' \in \Gamma_{\natural} \text{ s.t. } h_{\natural}((N, A, \mu[v/\llbracket *w \rrbracket], \hat{F})\downarrow) \leq \gamma' \quad \llbracket w \rrbracket \neq \text{nil} \quad \text{noerr}}{[(N, A, \mu, F)]_{\cong}, v := *w \rightarrow [\gamma']_{\cong}, \varepsilon} \quad (2)$$

$$\frac{\llbracket v \rrbracket \neq \text{nil} \quad \llbracket \alpha \rrbracket \neq \perp \quad \text{noerr}}{[(N, A, \mu, F)]_{\cong}, *v := \alpha \rightarrow [h_{\natural}((N, A, \mu[\mu(v)/\llbracket \alpha \rrbracket], \hat{F})\downarrow)]_{\cong}, \varepsilon} \quad (3)$$

$$\frac{\llbracket \alpha \rrbracket = \perp \vee \llbracket \alpha' \rrbracket = \perp \quad \text{noerr}}{[\gamma]_{\cong}, \alpha := \alpha' \rightarrow [\gamma_{\text{err}}]_{\cong}, \varepsilon} \quad (4)$$

In Fig. 3 the semantic rules are visualized for an example configuration. In rule 2 there might be the necessity for both abstraction and concretion. The execution of the assignment and the following abstraction via h_{\natural} yields an intermediate configuration which is generally not canonical since the variable v could now be too close to an abstract node. Therefore we have to find a canonical configuration γ' that is at least as concrete as $\bar{\gamma}$ and related by an abstraction morphism to it. There might be more than one solution, thus this rule is nondeterministic (indicated by the dashed arrows), but remains the only source of nondeterminism.

In rules 1 and 3 the distance to an abstract node is not reduced, but the opposite case can occur: just imagine an assignment of the form $y := \text{nil}$. If y points into a list whose head is referred to by another variable, we possibly increase the distance from that variable to abstract nodes. The execution of the assignment





Fig. 3. Exemplary visualization of the abstract semantics ($M = 3$)


 Fig. 4. Producer/Consumer: Abstract State Space ($M = 1$)

therefore potentially yields a non-canonical configuration and we have to re-abstract to determine the corresponding canonical configuration. According to Thm. 3.6 the result is unique and thus these steps are deterministic.

Example 3.8 Figure 4 shows the finite abstract state space of the producer/consumer program from Fig. 1 for $M = 1$. The p - and c -transitions each summarize several producer/consumer steps. The dashed transitions are nondeterministic steps, since the abstract node, visualized in black color, represents at least two nodes in a chain. If now the consumer deletes one node from the beginning of the queue the distance of x to the abstract node becomes two and thus we need to *concretize* the graph to obtain a canonical configuration. For this we distinguish two cases: either the abstract node represents exactly two nodes, then we reach the graph to the right, or it represents more than two, in which case we stay in the same state since the abstract node still represents more than one concrete node.

Theorem 3.9 (Finiteness) *For every $\pi \in \text{LMP}$, T_π^a is finite.*

The idea of the proof is to establish a bound on the number of nodes of canonical configurations for a given number of program variables.

Theorem 3.10 (Correctness of the Abstraction) *Let $\pi \in \text{LMP}$. For every transition in T_π^c there exists a corresponding abstract transition in T_π^a such that the heaps are related by abstraction morphisms.*

The proof of the correctness theorem has been omitted due to space constraints.

4 A Logic for Concurrent List-Manipulating Programs

In the previous sections we have defined our programming language for concurrent pointer manipulation and both its concrete and abstract semantics. In this section we will present a logic which will allow us to reason about heap configurations and program behavior. In the following LV denotes a set of logical variables, where we always assume that $LV \cap PV = \emptyset$.

Pointer Logic

Pointer logic deals with single configurations and is employed to express graph properties as well as to inspect the special flags of heap configurations (see Def. 2.2).

Definition 4.1 The set PL of *Pointer Logic formulas* is given by the grammar

$$\begin{aligned} \text{NExp} & ::= \text{nil} \mid v \in PV \mid x \in LV \mid * \text{NExp} \\ \text{Atomic} & ::= \text{tt} \mid \text{ff} \mid \text{err} \mid \text{dl} \mid \text{leak} \mid \text{signal} \mid \text{new} \mid \text{NExp} = \text{NExp} \mid \text{NExp} \rightsquigarrow \text{NExp} \\ \text{PL} & ::= \text{Atomic} \mid \neg \text{PL} \mid \text{PL} \wedge \text{PL} \mid \exists x : \text{PL} \end{aligned}$$

Later on we will use the logical operations \vee , \rightarrow , \leftrightarrow , and \forall (defined as usual) as abbreviations. Note that in contrast to pointer expressions in LM–programs our logic supports dereferencing operations of arbitrary depth. The special operation $\alpha \rightsquigarrow \alpha'$ expresses the reachability of heap objects.

Definition 4.2 Let $\beta : LV \rightarrow N$ be a variable valuation and $\gamma \in \Gamma_c$ a concrete heap configuration. Then we define $\llbracket \cdot \rrbracket : \text{NExp} \rightarrow N_{\text{nil}}$ by:

$$\begin{aligned} \llbracket \text{nil} \rrbracket & := \text{nil} & \llbracket v \rrbracket & := v \text{ for } v \in PV \\ \llbracket x \rrbracket & := \beta(x) \text{ for } x \in LV & \llbracket * \alpha \rrbracket & := \mu_\gamma(\llbracket \alpha \rrbracket) \text{ for } \alpha \in \text{NExp} \end{aligned}$$

Note the semantic difference with respect to the programming language. In navigation expressions a variable v is interpreted by itself and not by the node it is referencing. This avoids the necessity of the referencing operator $\&$.

Definition 4.3 The (concrete) satisfaction relation \models for PL–formulas is given as follows⁴ (for $\gamma = (N, \emptyset, \mu, F)$):

$$\begin{aligned} \gamma, \beta \models f & \quad \text{iff } f \in F, \text{ where } f \in \{\text{err}, \text{dl}, \text{leak}, \text{signal}, \text{new}\} \\ \gamma, \beta \models \alpha_1 = \alpha_2 & \quad \text{iff } \llbracket \alpha_1 \rrbracket = \llbracket \alpha_2 \rrbracket \neq \perp \\ \gamma, \beta \models \alpha_1 \rightsquigarrow \alpha_2 & \quad \text{iff } \llbracket \alpha_1 \rrbracket \neq \perp \wedge \llbracket \alpha_2 \rrbracket \in \mu^*(\llbracket \alpha_1 \rrbracket) \\ \gamma, \beta \models \exists x : \varphi & \quad \text{iff } \exists n \in N : \gamma, \beta[x/n] \models \varphi \end{aligned}$$

Temporal Pointer Logic

Pointer Logic enables us to express properties of single configurations. However it cannot be used to specify (ongoing) computations, i.e., configuration sequences. To this aim we will now extend this logic by temporal operators.

Definition 4.4 The set TPL of *Temporal Pointer Logic formulas* is given as follows:

$$\text{TPL} ::= \text{PL} \mid \neg \text{TPL} \mid \text{TPL} \wedge \text{TPL} \mid \mathbf{X} \text{TPL} \mid \text{TPL} \mathbf{U} \text{TPL}$$

For $\varphi \in \text{TPL}$ we use the abbreviations $\mathbf{F}\varphi := \text{ttU}\varphi$ and $\mathbf{G}\varphi := \neg \mathbf{F}\neg\varphi$. Moreover $V(\varphi) \subseteq LV$ denotes the set of (bound or free) logical variables occurring in φ .

Note that it is *not* possible to nest quantifiers and temporal operators. To do so it would be necessary to keep track of the object identities between states, which is difficult in the presence of abstract nodes. In addition it would blow up the state space and exclude the use of standard model checking algorithms. To the best of our knowledge the only approach to support this idea is the one in [18,19,20]; other works in the area such as [46] consider only *shapes* of the heap. This results in a loss

⁴ For \wedge, \neg, tt and ff the semantics is standard and therefore omitted.

of expressivity, e.g., a property like $\forall x : \text{new}(x) \rightarrow \mathbf{F} \text{del}(x)$ which states that every produced object will eventually be consumed cannot be formulated. Nonetheless we can specify many interesting properties.

Example 4.5 For our producer/consumer system from Fig. 1 it holds true:

1. $\neg \mathbf{F}(\text{dl} \vee \text{err})$ (never deadlock or pointer errors)
2. $\mathbf{GF} \text{ new}$ (new objects are created infinitely often)
3. $\mathbf{G}((\ast x \neq \text{nil} \vee \ast y \neq \text{nil}) \rightarrow (x \rightsquigarrow \ast y \wedge \forall v : (v \neq y \rightarrow x \rightsquigarrow v)))$
(whenever the queue is not empty, the object y points to is reachable from x and between x and this object lies a chain)

More general correctness properties are:

4. $\mathbf{F} \ast x = \ast y$ (x and y will eventually become aliases)
5. $\mathbf{G}\neg(\exists z : (x \rightsquigarrow z \wedge y \rightsquigarrow z))$ (x and y always point to disjoint parts of the heap)
6. $\mathbf{G}(\forall y : (x \rightsquigarrow y \rightarrow (\neg \exists z : (y \rightsquigarrow z \wedge \ast z \rightsquigarrow y))))$
(x always points to a non-cyclic list)
7. $\mathbf{FG}(\neg \text{leak})$ (only finitely often a memory leak can occur)
8. $\mathbf{G}(\forall y : (x \rightsquigarrow y \rightarrow (\forall z : (z \rightsquigarrow y \rightarrow x \rightsquigarrow z))))$ (x always points to a chain)

As mentioned before, TPL specifies computation paths. The set of possible paths is represented by a transition system.

Definition 4.6 Let $T = (Q, q_0, \text{lab}, \rightarrow)$ be a (concrete) transition system with $\text{lab} : Q \rightarrow \Gamma_c$. A *path* in T is an infinite sequence of states $\rho = \rho_0 \rho_1 \rho_2 \dots \in Q^\omega$ such that $\rho_i \rightarrow \rho_{i+1}$ for all $i \in \mathbb{N}$. Then for $\varphi \in \text{PL}$ we have

$$\rho \models \varphi (\in \text{PL}) \text{ iff } \exists \beta : LV \rightarrow N_{\text{lab}(\rho_0)} \text{ s.t. } \text{lab}(\rho_0), \beta \models_{\text{PL}} \varphi$$

For the temporal operators the semantics is identical to the one of LTL. We write $T \models \varphi$ iff $\rho \models \varphi$ for all paths $\rho \in \{q_0\}Q^\omega$ in T .

Reasoning about Abstract Computations

As expected the concrete semantics is straightforward. When we switch to abstract configurations, however, we run into several complications since logical variables can be bound to both concrete and abstract nodes. In the latter case we have to record *which* concrete node, represented by the summary node, it is bound to. This could lead to undefinedness of Pointer Logic formulas. This problem occurs mainly in direct comparisons of the form $\alpha = \alpha'$. To tackle this problem we choose the global precision constant M in dependence of the formula as follows. If $\varphi \in \text{TPL}$ is the formula to check, then we assume from now on that

$$M \geq \sum_{x \in V(\varphi)} \{j + 1 \mid \ast^j x \text{ occurs in } \varphi\}.$$

Due to the presence of abstract nodes it is not sufficient anymore to evaluate logical variables by simple variable-to-node mappings. Additionally we must record the offset of a variable referring to an abstract node and the distance between variables pointing to the same abstract node. This leads to the concept of *abstract valuations*.

Given $\gamma \in \Gamma_{\natural}$ and $\varphi \in \text{TPL}$, an *abstract valuation* is of the form $\eta = (\beta, o, \delta)$, where $\beta : V(\varphi) \rightarrow N_\gamma$ maps logical variables to (abstract) nodes, $o : V(\varphi) \rightarrow \mathbb{M}$ denotes the offset for an abstract node, and $\delta : V(\varphi) \times V(\varphi) \rightarrow \mathbb{M}$ is a “distance matrix” for the logical variables with potentially undefined entries. δ is only defined if both arguments are mapped onto the same entity, and o is only different from 1 if the corresponding variable is mapped onto an abstract node. The set of all such valuations will be denoted by $\text{Val}_{\gamma, \varphi}$.

Using this concept one can define a function $d_{\gamma, \eta} : \text{NExp} \times \text{NExp} \rightarrow \{0, 1, \infty\}$ measuring the “distance” of pointer expressions, where distance here means either 0 if the expressions are mapped onto the same (concrete) entity, 1 if the the first case does not hold but the second argument is reachable from the first or ∞ if neither is the case.

The presence of abstract nodes plays a vital role in the abstract semantics. Without the global constraint for M we would not be able to resolve all possible cases of abstract valuations, a third truth value would thus become necessary. The distance function δ is required for the case that both variables are mapped onto an abstract node with offset \star . With the help of the distance function the abstract semantics of PL and TPL is straightforward.

Definition 4.7 Let $\gamma = (N, A, \mu, F) \in \Gamma_{\natural}$ and $\eta = (\beta, o, \delta) \in \text{Val}_{\gamma, \varphi}$. The satisfaction relation \models for PL-formulas on canonical configurations is then given as follows (omitting the trivial cases):

$$\begin{aligned} \gamma, \eta \models f & \quad \text{iff } f \in F, \text{ where } f \in \{\text{err}, \text{dl}, \text{leak}, \text{signal}, \text{new}\} \\ \gamma, \eta \models \alpha_1 = \alpha_2 & \quad \text{iff } d_{\gamma, \eta}(\alpha_1, \alpha_2) = 0 \\ \gamma, \eta \models \alpha_1 \rightsquigarrow \alpha_2 & \quad \text{iff } d_{\gamma, \eta}(\alpha_1, \alpha_2) \in \{0, 1\} \\ \gamma, \eta \models \exists x : \varphi & \quad \text{iff } \exists n \in N, \text{ off} \in \mathbb{M}, \text{ dist} : V(\varphi) \rightarrow \mathbb{M} \text{ s.t.} \\ & \quad \gamma, (\beta_\eta[x/n], o_\eta[x/\text{off}], \delta_\eta[x/\text{dist}]) \models \varphi \end{aligned}$$

Let $T = (Q, q_0, \text{lab}, \rightarrow)$ be an abstract transition system with $\text{lab} : Q \rightarrow \Gamma_{\natural}/\cong$ and $\rho \in Q^\omega$ a path in it. Then $\rho \models \varphi \in \text{PL}$ iff for $\gamma \in \text{lab}(\rho_0)$ there exists an $\eta \in \text{Val}_{\gamma, \varphi}$ s.t. $\gamma, \eta \models_{\text{PL}} \varphi$. Temporal operators and Boolean connectives are treated in the standard way. We write $T \models \varphi$ iff $\rho \models \varphi$ for all paths $\rho \in \{q_0\}Q^\omega$ in T .

The following theorem states that the abstract semantics of TPL and of the programming language is correct, i.e., that the validity of a formula under the abstract interpretation implies the validity under the concrete one. The converse though does not hold.

Theorem 4.8 *Let $\pi \in \text{LMP}$ and $\varphi \in \text{TPL}$. If $T_\pi^a \models \varphi$ then $T_\pi^c \models \varphi$.*

Proof. It suffices to show for all $\varphi \in \text{PL}$ and $\gamma \in \Gamma_c$ the proposition:

$$\exists \beta : LV \rightarrow N_\gamma \text{ s.t. } \gamma, \beta \models \varphi \Leftrightarrow \exists \eta \in \text{Val}_{\gamma, \varphi} \text{ s.t. } h_{\natural}(\gamma), \eta \models \varphi \quad (\star)$$

Note that the \Leftarrow -direction is sufficient for correctness, the \Rightarrow -direction though is trivial. In the proof the choice of the global constant M (depending on the formula) plays a central role. Imagine for example a property “the heap contains at least five objects different from program variables”. To formulate this property we need at least five different logical variables and the constraint on M implies that $M \geq 5$. For smaller M it can happen that a formula that is satisfied in the abstract case, does

not hold in all concrete configurations associated with the abstract one. E.g. for $M = 1$ and a graph with one abstract node our example property would be satisfied; in the corresponding concrete graph where the abstract node is represented by two concrete nodes not necessarily.

With (\star) we can infer from Thm. 3.10 the validity of the claim, since TPL does not allow path quantifiers. By construction of the abstract PL-semantics it is intuitively clear that (\star) holds. \square

Model Checking Temporal Pointer Logic

Because of the two-stage approach in defining the logic, we can reduce the TPL model checking problem to an LTL model checking problem, which can efficiently be verified by existing model checkers.

Algorithm 1 *Let $T = (Q, q_0, lab, \rightarrow)$ be the abstract transition system generated by a program $\pi \in \text{LMP}$ and $\varphi \in \text{TPL}$ the formula to verify. Let $\Psi := \{\psi \in \text{PL} \mid \psi \text{ maximal subformula of } \varphi\} = \{\psi_1, \dots, \psi_r\}$.*

Define a “traditional” transition system $T' = (Q, q_0, lab', \rightarrow)$ where $lab' : Q \rightarrow 2^{AP}$ with $AP = \{p_i \mid i \in \{1, \dots, r\}\}$ such that $p_i \in lab'(q) \Leftrightarrow lab(q) \models \psi_i$.

Now solve the LTL model checking problem $T' \models_{\text{LTL}}^? \varphi[\psi_1/p_1, \dots, \psi_r/p_r]$.

The idea is thus to replace all (maximal) PL-subformulas by atomic propositions to obtain an LTL-formula. To do so we first have to evaluate the PL-formulas on the transition system and to change its labeling from configurations to atomic propositions, where each atomic proposition represents the truth value of the corresponding PL-subformula on the given configuration. The correctness of this approach is clear.

Limitations

Due to the nondeterminism in the abstract semantics caused by the presence of abstract nodes we may obtain *false negatives*. This means that in the abstract transition system there may exist computations which do not correspond to concrete ones and on which the property to verify does not hold.

Consider a program creating a list (pointed to by v) with $M + 3$ elements and then deleting again $M + 3$ elements. The property to verify is $\mathbf{XF}(*v = nil)$, i.e. that the list becomes empty. It is obvious that due to the presence of an abstract node after the construction of the list in the abstract semantics there is a path that retains that abstract node and thus the list never becomes empty (see Def. 3.7, rule 2). In the concrete case however the formula is satisfied.

Due to the overapproximation and the LTL approach false *positives* though cannot occur. This means that the successful verification of a property in the abstract case implies the correctness in the concrete case. False negatives can only occur in cases where information on the precise number of objects is necessary.

5 Application: Concurrent Garbage Collection

In this section we will show we will employ our approach to find counterexamples of a concurrent garbage collection algorithm. More concretely we will consider a so-called *mark-and-sweep* collector, which maintains a bit for each object in the heap to record its reachability status. Here we model this information as an additional heap component, a (partial) function $r : N \rightarrow \mathbb{B}$ which indicates whether the

collector considers a node to be reachable (1) or not (0). This component is made accessible to the garbage collector program using the additional constructs

- **reset** \in Stmt, which resets the reachability value of every node to 0,
- **mark**(α) \in Stmt where $\alpha \in$ PExp, which sets the reachability information of the node $\llbracket \alpha \rrbracket$ to 1, and
- $r(\alpha) \in$ BExp where $\alpha \in$ PExp tests whether the reachability bit of $\llbracket \alpha \rrbracket$ is set.

We refrain from giving the formal details of the extended syntax and semantics of LM–programs; these are straightforward to formalize. The only modification we would like to mention explicitly is an adaptation of the automatic garbage collection procedure (cf. Def. 2.3), which is activated after the execution of every LM–statement which potentially causes nodes to become unreachable (we refer to the derivation rules in Def. 2.5). To ensure the finiteness of our abstraction, we still have to use it. However, we will adapt the handling of the leak flag such that it will be set only if the garbage collector considers an unreachable node n to be reachable, i.e., if $r(n) = 1$. Formally this means that for an extended configuration $\hat{\gamma} = (N, A, \mu, F, r)$ we define $\hat{\gamma} \downarrow := (N', A \cap N', \mu \upharpoonright N', F \cup \{\text{leak} \mid \exists n \in (N \setminus N') : r(n) = 1\}, r \upharpoonright N')$ with $N' = \mu^*(PV)$.

Using these concepts we can now proceed by describing how a concurrent garbage collector can be added to a given LM–program, called a *mutator*. For $\pi = \mathbf{var} v_1, \dots, v_k : (s_1 \parallel \dots \parallel s_l) \in$ LMP, we define $\pi' := \mathbf{var} v_1, \dots, v_k, t : (s_1 \parallel \dots \parallel s_l \parallel c)$ with garbage collector c as in Fig. 5.

Thus the garbage collector is running concurrently with the mutator. It executes an infinite loop, starting by resetting the reachability bit of every node in the heap. Using the auxiliary variable t , it then marks every reachable node, beginning with the roots of the heap which are accessible by the program variables. Here the statement **with** $v \in PV$ **do** s **od** is a meta construct which is expanded to $s[v/v_1]; s[v/v_2]; \dots; s[v/v_k]$ for $PV = \{v_1, \dots, v_k\}$. Whenever it encounters a node which has already been marked (**if** statement), it continues with the next program variable to avoid redundant assignments. Finally it employs the signaling mechanism of our programming language to indicate that now the actual collection phase would start, i.e., that all nodes whose reachability bit is 0 would be removed.

Note, however, that we are still using our automatic garbage collection procedure such that we can guarantee that in every configuration of the system, all nodes are reachable. In other words, whenever the signal occurs there should not exist any unmarked node in the heap. This observation is the key idea for specifying the *soundness* of the garbage collector c as a safety property in TPL. Here we assume that the underlying Pointer Logic (cf. Def. 4.1) is extended by atomic propositions of the form $r(\alpha)$ which allow us test the reachability information of the node to which the navigation expression α refers:

```

while tt do
  reset;
  with  $v \in PV$  do
     $t := v$ ;
    while  $t \neq nil$  do
      if  $r(t)$  then  $t := nil$ 
      else mark( $t$ );
       $t := *t$ 
    fi
  od
od;
signal
od
    
```

Fig. 5. A naive garbage collector

$$\mathbf{G}(\text{signal} \rightarrow \forall x : r(x))$$

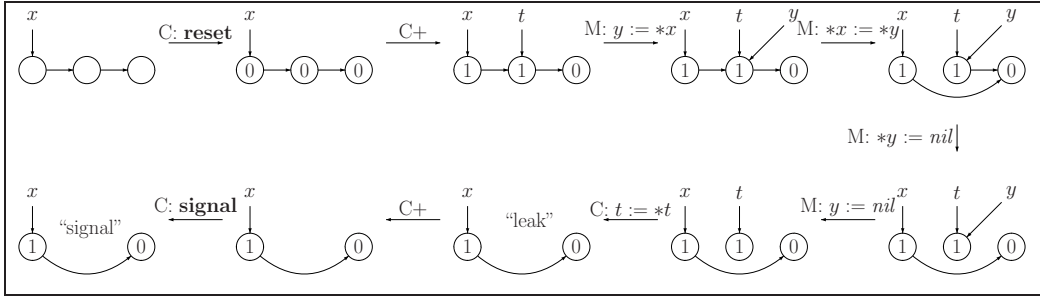


Fig. 6. Possible erroneous run of garbage collector and mutator

Another important issue is the *completeness* of the garbage collector, which means that every node which has become unreachable in the course of the computation, will eventually be removed. This, however, cannot be directly expressed for two reasons. First, verifying this property would require to keep track of the identity of objects between different configurations, which in turn involves the nesting of quantifiers and temporal operators. This is not supported by our logic. Second, our automatic garbage collection procedure immediately removes nodes that have become unreachable.

What we can formulate instead, however, is a safety property which comes very close to the actual completeness. It expresses that a node which has become unreachable will never be marked by the garbage collector. Employing the modified handling of the leak flag, this property can simply be formulated as

$$\mathbf{G} \neg \text{leak}$$

Note that this formalization is only justified since the garbage collector is monotonic in the following sense: once a node has been marked, its reachability information will not be reset before the collection signal occurs. Moreover completeness can only be expected (just as the above soundness property) if it is guaranteed that the mutator does not modify the reachability bits.

The example computation in Fig. 6 shows that the above garbage collector violates both of these requirements. Here the mutator program is assumed to be of the form $y := *x; *x := *y; *y := \text{nil}; y := \text{nil}$; it simply discards the second node of the list whose head is referenced by x (assuming that this node exists). Here C and M stand for operations of the collector and the mutator, respectively, which are either concretely given or summarized by a “+” sign. The bits labeling the nodes indicate the reachability information as set by the collector.

The computation shows that the collector is neither correct nor complete. In the final step involving the signal flag, the reachability value 0 of the list’s tail node means that it would be removed by the collector although it is reachable. Two steps earlier, the leak flag indicates that garbage has automatically been deleted which has been marked as reachable by the collector. Both of these problems are caused by the uncontrolled interaction between the mutator and the collector; they can be avoided by placing the body of the collector loop in an atomic region.

6 Related Work

Related work on the topic of analyzing pointer-manipulating programs can be classified into the following (often overlapping) categories.

Predicate abstraction abstracts the state space of the program by evaluating it under a number of given predicates. This yields a Boolean program which conservatively simulates all potential executions [25]. Successful software model checkers such as BLAST [28] and SLAM [3] are based on this approach. There are several papers that use classical predicate abstraction for pointer analysis [2,14]. In particular, [15,16] study concurrent garbage collection using predicate abstraction.

Shape analysis is a static analysis framework that represents recursive data structures of unbounded size by finite structures, called “shape graphs”. The idea is to apply to the heap the same abstraction that is applied to the program’s states in predicate abstraction: it is defined in terms of equivalence classes of heap objects that are induced by a finite set of predicates on those objects. The usual approach is to formalize shape graphs by three-valued logical structures [46]. This approach has been implemented TVLA [34] and in BLAST [5] which makes use of TVLA.

Recent developments comprise the development of adaptive methods which automatically adjust to the data structures that occur in the given program [31,35,48], demand-driven techniques [5,27], efficiency improvements [33], and interprocedural shape analysis [26,30,43,44].

It is often argued that the application of predicate abstraction to pointer structures does not work well because it is difficult to find predicates which abstract heap structures in an appropriate and compact way [5]. This claim is substantiated by the results in [36] which investigates the application of both predicate abstraction and shape analysis to programs operating on singly-linked lists, employing a similar abstraction as ours: elements on unshared list segments are summarized. It is shown that standard predicate abstraction requires an exponential number of predicates in comparison to the number of predicates in shape analysis. Also [41] considers both techniques, but in a very restricted programming-language setting which only supports single assignments.

Regular model checking is a framework for unified verification of infinite-state systems based on automata theory. It represents states using words (trees) over a finite alphabet and sets of states using finite (tree) automata [10]. Like in our approach, singly-linked lists are also considered in [8,9], but only safety and termination properties are verified.

Dataflow analysis is a technique for gathering information about certain aspects of a program using its control flow graph. This approach is generally efficient but restricted to rather shallow properties of programs such as aliasing relations [17,39], points-to information [47,51], or pointer range analysis [50].

Hoare-style approaches: first-order reasoning typically breaks down when it comes to prove properties of pointer-manipulating programs. The main reason is that it is impossible to express an invariant of all members of a data structure in first-order logic. The latter has to be extended therefore to support the definition of a reachability predicate [1,12,22,32,37,38]. However such deductive techniques usually involve user interaction, or otherwise only restricted properties such as dereferencing of nil pointers or aliasing effects can be analyzed.

Separation logic has been proposed as an extension to Hoare logic that permits local reasoning about linked structures, supporting features to support modular correctness proofs for pointer-manipulating programs [40,42]. It has been employed for termination proofs of heap-manipulating programs [4], for interprocedural shape

analysis [24], for handling abstract data types [7], and for verifying garbage collection algorithms [6]. However most of the work on separation logic focuses on verifying programs manually.

In summary, many of the characterizing features of our approach are already present in earlier papers: the restriction to singly-linked lists without data fields, the introduction of abstract entities which represent a potentially unbounded number of heap cells (called “summary nodes” in [13]); see e.g. [2,9,36], and the observation that, in this setting, the number of sharing points in heap structures is bounded by the number of program variables [9,36].

However none of these combines the strengths of our approach which supports concurrent programs with dynamic memory allocation and destructive updates such that arbitrary (cyclic) linked lists can be constructed, integrates both abstraction and model checking in a fully automated way, supports a linear-time logic in which both safety and liveness properties can be expressed, and which allows to use standard LTL model checkers.

In comparison, many of the existing approaches suffer from the poor programming environment, the exclusion of cyclic data structures, the requirement of user interaction, or the restriction to safety properties. Notable exceptions are [2], which also offers liveness properties but requires user-defined ranking functions, [20], which employs extended tableau-based techniques for model checking, and [49], which has a non-standard interpretation.

7 Conclusions and Future Work

We have presented a framework for the verification of concurrent pointer-manipulating programs with unbounded heap size and destructive updates. The correctness properties are specified using temporal pointer logic which is essentially pointer logic for expressing heap properties enriched with temporal operators. Instead of requiring dedicated algorithms, the TPL model checking problem is reduced to an LTL model checking problem that can be verified effectively with a broad variety of existing model checkers. The tradeoff is the restriction to list-like data structures as well as the limitation in expressiveness of the logic because object identities are not tracked between configurations.

Currently we are implementing our method to verify more realistic examples in the future. In particular we will extend the analysis of concurrent garbage collectors by defining a “hardest mutator”, i.e., a general mutator program which is capable of simulating the behavior of any other mutator. This will enable us to establish the correctness of garbage collectors independent of the concrete mutator.

Furthermore due to the extensive use of concurrency, state space reduction and optimization techniques such as partial order reduction [21,23] will have to be employed and integrated in the implementation. We also plan to extend our framework with dynamic (unbounded) creation of threads. Another interesting aspect could be the combination of existing finite-state modeling languages like Promela [29] and pointer manipulation. Finally in the long run we have plans to increase the expressivity of the logic as well as to generalize our approach to richer data structures, for which new abstractions will be necessary. Moreover an automata-theoretic approach to defining a storeless semantics, as it is studied in [11] for a (concrete) semantics for pointer programs seems promising.

References

- [1] T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In *POPL '06*, pages 91–102. ACM Press, 2006.
- [2] I. Balaban, A. Pnueli, and L. D. Zuck. Shape analysis by predicate abstraction. In *VMCAI '05*, volume 3385 of *LNCS*, pages 164–180. Springer-Verlag, 2005.
- [3] T. Ball and S. K. Rajamani. The SLAM project: debugging system software via static analysis. In *POPL '02*, pages 1–3. ACM Press, 2002.
- [4] J. Berdine, B. Cook, D. Distefano, and P. W. O'Hearn. Automatic termination proofs for programs with shape-shifting heaps. In *CAV '06*, volume 4144 of *LNCS*, pages 386–400. Springer-Verlag, 2006.
- [5] D. Beyer, T. A. Henzinger, and G. Théoduloz. Lazy shape analysis. In *CAV '06*, volume 4144 of *LNCS*, pages 532–546. Springer-Verlag, 2006.
- [6] L. Birkedal, N. Torp-Smith, and J. C. Reynolds. Local reasoning about a copying garbage collector. In *POPL '04*, pages 220–231. ACM Press, 2004.
- [7] R. Bornat, C. Calcagno, P. O'Hearn, and M. Parkinson. Permission accounting in separation logic. In *POPL '05*, pages 259–270. ACM Press, 2005.
- [8] A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, and T. Vojnar. Programs with lists are counter automata. In *CAV '06*, volume 4144 of *LNCS*, pages 517–531. Springer-Verlag, 2006.
- [9] A. Bouajjani, P. Habermehl, P. Moro, and T. Vojnar. Verifying programs with dynamic 1-selector-linked list structures in regular model checking. In *TACAS '05*, volume 3440 of *LNCS 3440*, pages 13–29. Springer-Verlag, 2005.
- [10] A. Bouajjani, P. Habermehl, A. Rogalewicz, and T. Vojnar. Abstract regular tree model checking of complex dynamic data structures. In *SAS '06*, volume 4134 of *LNCS*, pages 52–70. Springer-Verlag, 2006.
- [11] M. Bozga, R. Iosif, and Y. Lakhnech. Storeless semantics and alias logic. *ACM SIGPLAN Not.*, 38(10):55–65, 2003.
- [12] M. Bozga, R. Iosif, and Y. Lakhnech. On logics of aliasing. In *SAS '04*, volume 3148 of *LNCS*, pages 344–360. Springer-Verlag, 2004.
- [13] D. R. Chase, M. Wegman, and F. K. Zadeck. Analysis of pointers and structures. In *PLDI '90*, pages 296–310. ACM Press, 1990.
- [14] D. Dams and K. S. Namjoshi. Shape analysis through predicate abstraction and model checking. In *VMCAI '03*, volume 2575 of *LNCS*, pages 310–323. Springer-Verlag, 2003.
- [15] S. Das and D. L. Dill. Successive approximation of abstract transition relations. In *LICS '01*, pages 51–58. IEEE, 2001.
- [16] S. Das, D. L. Dill, and S. Park. Experience with predicate abstraction. In N. Halbwachs and D. Peled, editors, *CAV '99*, volume 1633 of *LNCS*, pages 160–171. Springer-Verlag, 1999.
- [17] A. Deutsch. Interprocedural may-alias analysis for pointers: beyond k-limiting. In *PLDI '94*, pages 230–241. ACM Press, 1994.
- [18] D. Distefano. *On Model Checking the Dynamics of Object-Based Software: a Foundational Approach*. PhD thesis, Univ. of Twente, 2003.
- [19] D. Distefano, J.-P. Katoen, and A. Rensink. Who is pointing when to whom? – on the automated verification of linked list structures. In *FSTTCS '04*, volume 3328 of *LNCS*, pages 250–262. Springer-Verlag, 2004.
- [20] D. Distefano, J.-P. Katoen, and A. Rensink. Safety and liveness in concurrent pointer programs. In *FMCO '06*, volume 4111 of *LNCS*, pages 280–312. Springer-Verlag, 2006.
- [21] C. Flanagan and P. Godefroid. Dynamic partial-order reduction for model checking software. In *POPL '05*, pages 110–121. ACM Press, 2005.
- [22] P. Fradet, R. Gaugne, and D. L. Métayer. Static detection of pointer errors: an axiomatisation and a checking algorithm. In *ESOP '96*, volume 1058 of *LNCS*, pages 125–140. Springer-Verlag, 1996.
- [23] P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem*, volume 1032 of *LNCS*. Springer-Verlag, 1996.
- [24] A. Gotsman, J. Berdine, and B. Cook. Interprocedural shape analysis with separated heap abstractions. In *SAS '06*, volume 4134 of *LNCS*, pages 240–260. Springer-Verlag, 2006.

- [25] S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *CAV '97*, volume 1254 of *LNCS*, pages 72–83. Springer–Verlag, 1997.
- [26] B. Hackett and R. Rugina. Region–based shape analysis with tracked locations. In *POPL '05*, pages 310–323. ACM Press, 2005.
- [27] N. Heintze and O. Tardieu. Demand–driven pointer analysis. *ACM SIGPLAN Not.*, 36(5):24–34, 2001.
- [28] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Software verification with BLAST. In *SPIN '03*, volume 2648 of *LNCS*, pages 235–239. Springer–Verlag, 2003.
- [29] G. Holzmann. *The Spin Model Checker: Primer and Reference Manual*. Addison–Wesley, 2003.
- [30] B. Jeannet, A. Loginov, T. W. Reps, and S. Sagiv. A relational approach to interprocedural shape analysis. In *SAS '04*, volume 3148 of *LNCS*, pages 246–264. Springer–Verlag, 2004.
- [31] O. Lee, H. Yang, and K. Yi. Automatic verification of pointer programs using grammar–based shape analysis. In *ESOP '05*, volume 3444 of *LNCS*, pages 124–140. Springer–Verlag, 2005.
- [32] T. Lev–Ami, N. Immerman, T. W. Reps, S. Sagiv, S. Srivastava, and G. Yorsh. Simulating reachability using first–order logic with applications to verification of linked data structures. In *CADE '05*, volume 3632 of *LNCS*, pages 99–115. Springer–Verlag, 2005.
- [33] T. Lev–Ami, N. Immerman, and S. Sagiv. Abstraction for shape analysis with fast and precise transformers. In *CAV '06*, volume 4144 of *LNCS*, pages 547–561. Springer–Verlag, 2006.
- [34] T. Lev–Ami and S. Sagiv. TVLA: A system for implementing static analyses. In *SAS '00*, volume 1824 of *LNCS*, pages 280–302. Springer–Verlag, 2000.
- [35] A. Loginov, T. W. Reps, and S. Sagiv. Abstraction refinement via inductive learning. In *CAV '05*, volume 3576 of *LNCS*, pages 519–533. Springer–Verlag, 2005.
- [36] R. Manevich, E. Yahav, G. Ramalingam, and M. Sagiv. Predicate abstraction and canonical abstraction for singly–linked lists. In *VMCAI '05*, volume 3385 of *LNCS*, pages 181–198. Springer–Verlag, 2005.
- [37] A. Møller and M. I. Schwartzbach. The pointer assertion logic engine. In *PLDI '01*, pages 221–231. ACM Press, 2001.
- [38] G. Nelson. Verifying reachability invariants of linked structures. In *POPL '83*, pages 38–47. ACM Press, 1983.
- [39] E. M. Nystrom, H.-S. Kim, and W. mei W. Hwu. Bottom–up and top–down context–sensitive summary–based pointer analysis. In *SAS '04*, volume 3148 of *LNCS*, pages 165–180. Springer–Verlag, 2004.
- [40] P. W. O’Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *POPL '04*, pages 268–280. ACM Press, 2004.
- [41] A. Podelski and T. Wies. Boolean heaps. In *SAS '05*, volume 3672 of *LNCS*, pages 268–283. Springer–Verlag, 2005.
- [42] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS '02*, pages 55–74. IEEE Computer Society, 2002.
- [43] N. Rinetzky, J. Bauer, T. Reps, M. Sagiv, and R. Wilhelm. A semantics for procedure local heaps and its abstractions. In *POPL '05*, pages 296–309. ACM Press, 2005.
- [44] N. Rinetzky, M. Sagiv, and E. Yahav. Interprocedural shape analysis for cutpoint–free programs. In *SAS '05*, volume 3672 of *LNCS*, pages 284–302. Springer–Verlag, 2005.
- [45] M. Sagiv, T. Reps, and R. Wilhelm. Solving shape–analysis problems in languages with destructive updating. *ACM Trans. Program. Lang. Syst.*, 20(1):1–50, 1998.
- [46] M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3–valued logic. *ACM Trans. Program. Lang. Syst.*, 24(3):217–298, 2002.
- [47] J. Whaley and M. S. Lam. An efficient inclusion–based points–to analysis for strictly–typed languages. In *SAS '02*, volume 2477 of *LNCS*, pages 180–195. Springer–Verlag, 2002.
- [48] E. Yahav and G. Ramalingam. Verifying safety properties using separation and heterogeneous abstractions. In *PLDI '04*, pages 25–34. ACM Press, 2004.
- [49] E. Yahav, T. Reps, M. Sagiv, and R. Wilhelm. Verifying temporal heap properties specified via evolution logic. In *ESOP '03*, volume 2618 of *LNCS*, pages 204–222. Springer–Verlag, 2003.
- [50] S. H. Yong and S. Horwitz. Pointer–range analysis. In *SAS '04*, volume 3148 of *LNCS*, pages 133–148. Springer–Verlag, 2004.
- [51] J. Zhu and S. Calman. Symbolic pointer analysis revisited. In *PLDI '04*, pages 145–157. ACM Press, 2004.