



HAL
open science

Programs with Lists are Counter Automata

Ahmed Bouajjani, Marius Bozga, Peter Habermehl, Radu Iosif, Pierre Moro,
Tomas Vojnar

► **To cite this version:**

Ahmed Bouajjani, Marius Bozga, Peter Habermehl, Radu Iosif, Pierre Moro, et al.. Programs with Lists are Counter Automata. Computer Aided Verification, Aug 2006, Seattle, United States. pp.517-531. hal-00148002

HAL Id: hal-00148002

<https://hal.science/hal-00148002>

Submitted on 21 May 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Programs with Lists are Counter Automata

Ahmed Bouajjani², Marius Bozga¹, Peter Habermehl², Radu Iosif¹, Pierre Moro², and Tomáš Vojnar³

¹ VERIMAG, 2 av. de Vignate, F-38610 Gières, e-mail:{iosif,bozga}@imag.fr

² LIAFA, Paris University 7, Case 7014, 2, place Jussieu, F-75251 Paris Cedex 05
e-mail:{Ahmed.Bouajjani,Peter.Habermehl,Pierre.Moro}@liafa.jussieu.fr

³ FIT, Brno University of Technology, Božetěchova 2, CZ-61266, Brno
e-mail:vojnar@fit.vutbr.cz

Abstract. We address the verification problem of programs manipulating one-selector linked data structures. We propose a new automated approach for checking safety and termination for these programs. Our approach is based on using counter automata as accurate abstract models: control states correspond to abstract heap graphs where list segments without sharing are collapsed, and counters are used to keep track of the number of elements in these segments. This allows to apply automatic analysis techniques and tools for counter automata in order to verify list programs. We show the effectiveness of our approach, in particular by verifying automatically termination of some sorting programs.

1 Introduction

The design of automatic verification methods for programs manipulating dynamic linked data structures is a challenging problem. Indeed, the analysis of the behaviour of such programs requires reasoning about complex transformations of data structures involving both creation and deletion of objects as well as modifications of the links between them (pointer manipulations). The heap of such programs may have in fact an arbitrary size and shape (a graph structure). There are several approaches for tackling this problem addressing different subclasses of programs and using different kinds of formalisms for representing and reasoning about infinite sets of heap structures, e.g., [21, 19, 22, 9].

We consider in this paper the class of programs manipulating linked data structures with a single data-field selector. It corresponds to programs manipulating linked lists with the possibility of sharing and circularities. We propose a new approach for the automatic verification of such programs which is mainly based on using counter automata as accurate abstract (infinite-state) models. These models can be used for checking both safety properties and termination of the considered programs using techniques such as (abstract) symbolic reachability analysis (for safety and invariance checking) and automatic generation of decreasing ranking functions (for termination checking).

Let us present in more details the proposed approach. We start from the observation that if we do not consider garbage (parts of the heap not reachable from the pointer variables of the program), the heap graph is always a finite collection of graphs of a special form close to a tree: it is either a tree (where edges are directed towards the root) or a set of trees having all their roots connected to a simple cycle. The number of such graphs is infinite, but it can be proved that for each of them, the number of vertices where sharing occurs is bounded by the number of pointer variables of the program.

Then, for data-insensitive programs (e.g., a list reversal program), a natural abstraction consists in mapping each sequence of elements between two sharing points into an

abstract sequence of some (fixed) bounded size. However, for each given value of the bound, this abstraction is obviously not precise in general. In order to define a precise abstraction, we need in fact to reason about the size of each sequence between two sharing points. This leads to the idea of using counters in order to keep this information in the abstract model (and therefore to use counter automata as abstract models).

In fact, considering counter automata-based models has several advantages. Not only it allows to define accurate abstractions, it allows us also to handle quantitative properties depending on the sizes of some parts of the heap. Thus, we can handle programs with integer variables whose value is somehow related to the contents of the lists (e.g., to their length). Moreover, it provides a powerful way for checking termination which typically requires reasoning about decreasing values (e.g., the size of the part of the list to be treated).

A first contribution of the paper is to define an abstraction mapping from data-insensitive programs to counter automata for which we prove that the (concrete) program and its abstraction are *bisimilar*. This result is interesting since it means that our abstraction preserves all properties of the class of data-insensitive programs. The control states of the built automaton correspond to abstract shapes (heap graphs where sequences between shared points are reduced to single vertices), and each transition corresponds to the execution of a program statement. It represents a modification in the shape together with a modification on the counters (attached to vertices abstracting sequences between sharing nodes).

The control structure of the built counter automata can be arbitrary in general. However, it turns out that these automata have an important property: we prove that if we consider the evolution of the sum of all counters, the effect of executing any control loop is to increment this sum by a constant which depends on the program. We use this fact to establish a new decidability result for list programs: for every given (data-insensitive) list program, if the control structure of the generated counter automaton has no nested loops, the verification problems of safety properties and termination are both decidable.

Subsequently, we go further by considering the issue of data-sensitivity. We consider the class of programs manipulating objects ranging over a potentially infinite data domain supplied with an ordering relation, and we assume that the only allowed operation on these data values is the comparison w.r.t. this ordering relation. This class of programs includes, for instance, sorting programs. We extend our previous abstraction principle to the heap graphs of these programs by taking into account (in addition to the size) some information about the order of the elements in the abstracted sequences between sharing points, and we provide a construction which associates with each program a counter automaton-based abstract model. We show that this abstraction is sound and complete w.r.t. the choice of ordering predicates.

Finally, we show the application of our approach on three examples of programs (list reversal, insertion sort, and bubble sort). We have derived systematically their counter automata models, and then we used (1) our ARMC tool [10] (and some compile-time techniques) for checking safety properties, and (2) the Terminator tool based on [13] for termination.

Related Work: Programs manipulating singly-linked lists have gained a lot of attention within the past two years, as shown by the fairly large number of recent publications on the subject [5, 7, 20, 4, 9]. Interestingly, the idea of abstracting away all the list segments

with no incoming edges is common to many of these works, even though they are independent and use different approaches and frameworks (e.g. static analysis [20], predicate abstraction [4], symbolic reachability analysis [5] and proof search [7]). The fact that the number of sharing points in abstract heap structures is bounded by the number of variables in the program is also behind the techniques proposed in [20, 9].

In [11], the authors use an abstract shape model with counters, but their concerns are mostly related to the decidability of a specification logic. The approach that is the closest to ours is [5]. However, it is rather pointed towards showing particular properties such as absence of segmentation faults and memory leak errors, than checking general safety properties, and the work does not address the problem of verifying termination. Moreover, the work reported in [5] offers less automation of the verification than ours. Recently, the same authors have started independently a work [16] on automatic construction of models based on counter automata similar to our approach.

The use of ordering predicates in order to handle sorting programs is similar to the one considered in [15, 22] based on the shape analysis approach.

Termination is tackled by works such as [23, 4]. In all of these works, ranking functions must be given manually, whereas our approach is fully automated.

2 Programs with Lists

In this section we define a model for programs manipulating dynamic list data structures. We consider that lists are implemented using reference (pointer) data types with one selector (`next`) field, as it is the case in most object-oriented imperative programming languages (e.g. Java, C, C++). For the time being we consider programs without recursion or concurrency constructs, therefore all variables are assumed to be global. In addition to the list data structures, the programs can have integer variables.

2.1 Syntactic Definitions

We consider imperative programs working with a set of pointer variables $PVar$ and a set of integer counter variables $IVar$. The pointer variables refer to list cells. Pointers can be used in assignments such as $u := \text{null}$, $u := w$ and $u := w.\text{next}$, selector updates $u.\text{next} := w$ and $u.\text{next} := \text{null}$, and new cell creation $u := \text{new}$. Counters can be incremented $i := i + 1$, decremented $i := i - 1$ and reset $i := 0$. The control structure is composed of iteration (`while`) statements and conditionals (`if-then-else`). The guards of the control constructs are pointer equality $u = w$, data comparisons $u.\text{data} <= v.\text{data}$, zero tests for counters $i = 0$ and boolean combinations of the above. An example is the list reversal program in Figure 1.

To simplify the definition of the operational semantics, we consider that all programs are precompiled, by introducing right before any pointer assignment of the form $u(u.\text{next}) := \text{new}(w, w.\text{next})$ an assignment $u(u.\text{next}) := \text{null}$. In particular, a pointer assignment of the form $u := u.\text{next}$ is turned into $v := u; u := \text{null}; u := v.\text{next}$, possibly introducing a fresh variable v .

```

1: while  $i \neq \text{null}$  do
2:    $k := i.\text{next}$ ;
3:    $i.\text{next} := j$ ; 4:  $j := i$ ;
5:    $i := k$ ;
6:   od

```

Fig. 1. List Reversal Program

2.2 Concrete Operational Semantics

In order to define the concrete semantics of programs with lists, we have to formalize the notion of *heap*. In principle, a heap is a graph in which each node has at most one successor. In addition, some nodes are designated by special labels (variables from $PVar$). If all the edges are reversed, one can imagine a heap as a set of disjoint trees, in which, for each tree there might be an extra edge from an arbitrary node back to the root.

In the rest of the paper, for a set A we denote by A_\perp the set $A \cup \{\perp\}$. The element \perp is used to denote that a (partial) function is undefined at a given point, e.g. $f(x) = \perp$. Also, for a function f we denote by $f \downarrow_A$ the projection of f on A i.e. $f \cap A \times A$.

Definition 1. Let $\langle \mathcal{D}, \preceq \rangle$ be a totally ordered set, and $PVar$ a set of pointer variables. A heap is a tuple $H = \langle N, S, V, D \rangle$, where N is a finite set of nodes, $S : N \rightarrow N_\perp$ is a successor function, $V : PVar \rightarrow N_\perp$ is a function associating nodes to variables, and $D : N \rightarrow \mathcal{D}$ is a function associating each node a data element.

The set of all heaps using variables from $PVar$ is denoted by $\mathcal{H}(PVar)$. We denote by $n_1 \xrightarrow{H} n_2$ the fact that n_2 is the successor of n_1 in H , and by $u \xrightarrow{H} n$ that the pointer variable $u \in PVar$ refers to a node whose successor is n . H might be omitted when it is clear from the context. We denote by \xrightarrow{H}^* the reflexive and transitive closure of \xrightarrow{H} . A node n is said to be a *cut point* in H , denoted as $cut_H(n)$, if either it has two predecessors or it is pointed to by a variable.⁴

The *state* of a program with lists is a triple $\langle l, \iota, H \rangle$ where $l \in Lab$ is the current program label, $\iota : IVar \rightarrow \mathbb{Z}$ is the current valuation of counter variables, and $H \in \mathcal{H}(PVar)$ is the current heap configuration. Each assignment modifies the state as follows: $\langle l, \iota, H \rangle \xrightarrow{l:s:l'} \langle l', \iota', H' \rangle$, where l' is the label of the next statement, ι' is the new valuation of counters, computed as usual, and H' is a heap configuration such that $H \xrightarrow{s} H'$, in conformance with the rules in Figure 2.2. As a result of removing a node from the heap, other nodes might become unreachable from the pointer variables. This set of nodes, whose lifetime *depends exclusively* on $n \in N$, is denoted as $dep_H(n)$. H_{err} is a special sink heap configuration, attained as the result of a null pointer dereference. A pointer equality test $u = v$ evaluates to true in a heap $H = \langle N, S, V \rangle$ if and only if $V(u) = V(v)$. Also, $u = \text{null}$ is true if and only if $V(u) = \perp$. Due to the lack of space, the rules for the assignment statements $u := w$, $u := \text{new}$, $u.\text{next} := \text{null}$ and $u.\text{next} := w$ are deferred to the long version of this paper [8].

3 Counter Automata

A counter automata with n counters is a tuple $A = \langle Q, X, \rightarrow \rangle$, where Q is a finite set of control states, $X = \{x_1, \dots, x_n\}$ are the counter variables and $\rightarrow \in Q \times \Phi \times Q$ are the transitions, where Φ is the set of Presburger formulae with free variables from $\{x_i, x'_i \mid 1 \leq i \leq n\}$. A configuration of a counter automata with n counters is a tuple $\langle q, \mathbf{v} \rangle$, where \mathbf{v} is a mapping from X to \mathbb{N} . The set of all configurations is denoted

⁴ Formally, $cut_H(n) : \exists n_1, n_2 \in N . n_1 \neq n_2 \wedge S(n_1) = S(n_2) = n \vee \exists u \in PVar . V(u) = n$.

$$\begin{array}{c}
\frac{V(u) = \perp}{H \xrightarrow{u := \text{null}} H} \qquad \frac{\exists w \in PVar \setminus \{u\} . w \xrightarrow{*} V(u)}{H \xrightarrow{u := \text{null}} \langle N, S, V[u \rightarrow \perp], D \rangle} \\
\frac{V(u) = n \in N \quad N' = N \setminus \text{dep}_H(n) \quad \forall w \in PVar \setminus \{u\} . \neg w \xrightarrow{*} n}{H \xrightarrow{u := \text{null}} \langle N', S \downarrow_{N'}, V \downarrow_{N'}, D \downarrow_{N'} \rangle} \quad \frac{V(w) = n \in N}{H \xrightarrow{u := w.\text{next}} \langle N, S, V[u \rightarrow S(n)], D \rangle}
\end{array}$$

Fig. 2. Concrete Semantics of Heap Updates. $H \triangleq \langle N, S, V, D \rangle$, $\text{dep}_H(n) \triangleq \{m \in N \mid \forall u \in PVar . \neg u \xrightarrow{*} \langle N \setminus \{n\}, S \downarrow_{N \setminus \{n\}}, V \downarrow_{N \setminus \{n\}}, D \downarrow_{N \setminus \{n\}} \rangle m\}$.

by c . The transition relation $\xrightarrow{c} \subseteq c \times c$ is defined by $(q, v) \xrightarrow{c} (q', v')$ iff there exists a transition $q \xrightarrow{\phi} q'$ such that, if σ is an assignment of the free variables of ϕ ($FV(\phi)$), where $\sigma(x) = v(x)$ and $\sigma(x') = v'(x)$, we have that $\phi(FV(\phi)\sigma)$ holds and $v(x) = v'(x)$, for all variables x with $x' \notin FV(\phi)$. A *run* of A is a sequence of configurations $(q_0, v_0), (q_1, v_1), (q_2, v_2) \dots$ such that $(q_i, v_i) \xrightarrow{c} (q_{i+1}, v_{i+1})$, for each $i \geq 0$.

Definition 2. Let $A = \langle Q, X, \rightarrow \rangle$ be a counter automaton, where $X = \{x_1, \dots, x_n\}$ are counter variables, that range over positive integers. A is said to be *linear* if all its transitions are of the form: $\phi(X) \wedge \bigwedge_{1 \leq i \leq n} x'_i = f_i(X)$, where ϕ is a formula of Presburger arithmetic, and $f_i = \sum_{j=1}^n a_{ij}x_j + b_i$, $1 \leq i \leq n$ are linear functions with integer coefficients. Moreover, A is said to be *positive* if $a_{ij} \geq 0$, for all $1 \leq i, j \leq n$. A is also said to be *restrictive* if, there exists a constant $\alpha \in \mathbb{N}$ such that, for each control state $q \in Q$, on each run π that visits q , the sum of values taken by the counters, $\sum_{i=1}^n x_i$, increases by at most α between any two consecutive times when the control state is q .

The control graph of a counter automaton A is the graph having as vertices the set Q of control states, and, for any two states q and q' , there is an edge between q and q' in the control graph if and only if there exists a transition $q \xrightarrow{\phi} q'$ in A . A counter automaton is said to be *flat* if its control graph has no nested loops. We can prove:

Theorem 1. *The problems of reachability and termination for flat linear positive restrictive counter automata are decidable.*

We give a brief sketch of the proof. First notice that the transitions of a loop can be composed into one single loop of the form $l : \phi(X) \wedge X' = AX + B$. Then, we establish that for every operation $X' = AX + B$, the set $\{A^i\}_{i \geq 1}$ is finite, where $A^1 = A$ and $A^{i+1} = A^i \times A$, where \times is the product of square matrices. This implies that the reachability relation corresponding to the iteration of l is expressible in Presburger arithmetics [3, 17]. Since our automata are flat it is easy to deduce that both reachability and termination problems are decidable (it suffices to decide these problems for each loop).

4 Abstract Semantics of Programs with Lists

A common way of representing heaps compactly, consists in mapping an entire list segment with no incoming edges into a special (abstract) node. This idea constitutes also

the basis of our abstraction. Let \mathcal{N} be a set of *abstract nodes* and x be a set of *counter variables*, one for each node. We shall first define the abstract structure of heaps.

Definition 3. An abstract structure is a tuple $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle$, where:

- $\overline{N} \subseteq \mathcal{N}$ is the set of abstract nodes, and
- $\overline{S} : \overline{N} \rightarrow \overline{N}_\perp$, $\overline{V} : PVar \rightarrow \overline{N}_\perp$, are the successor and variable mappings,

An abstract structure is moreover said to be in normal form if, for each $n \in \overline{N}$, there exists $u \in PVar$ such that $u \xrightarrow[\overline{H}]{} n$, and n is a cut point in \overline{H} .

Intuitively, each abstract node corresponds to a set of concrete nodes, and the counter corresponding to each node gives the number of nodes in this set. For abstract structures in normal form, we do not allow sequences of successive abstract node that are neither pointed by a variable, nor have the indegree greater than one. This condition is needed in order to ensure that any such abstract structure defined over a finite set of variables is finite. $\overline{\mathcal{H}}(PVar)$ denotes the set of all abstract structures with variables from $PVar$.

Lemma 1. Let $PVar = \{u_1, \dots, u_n\}$ be a set of variables, and $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle$ be an abstract structure in normal form such that $\text{dom}(\overline{V}) \subseteq PVar$. Then $\|\overline{N}\| \leq 2n$. As a consequence, the number of such heaps is bounded asymptotically by n^n , and the bound is tight.

Let us define now a first abstraction function, denoted by α_s , that maps concrete heaps into abstract structures. Given a concrete heap $H = \langle N, S, V, D \rangle$, let $\triangleright_H \subseteq N \times N$ be a relation on the set of nodes, defined as: $n_1 \triangleright_H n_2 : n_1 \xrightarrow[H]{} n_2 \wedge \neg \text{cut}(n_2)$. We denote by \sim_H the reflexive, symmetric and transitive closure of \triangleright_H . The H subscript shall be further omitted for simplicity. For a node $n \in N$, we denote by $[n]$ the equivalence class of n with respect to \sim , also referred to a *list segment*. The *quotient heap* $H_{/\sim} = \langle N_{/\sim}, S_{/\sim}, V_{/\sim} \rangle$ is defined as follows:

- $N_{/\sim} = \{[n] \mid n \in N\}$,
- for all $n, m \in N$, $S_{/\sim}([n]) = [m]$ iff $\exists n_0 \in [n] \exists m_0 \in [m] . S(n_0) = m_0 \wedge \text{cut}_H(m_0)$,
- for all $u \in PVar$, $n \in N$, $V_{/\sim}(u) = [n]$ iff $V(u) \in [n]$, and
- $S_{/\sim}$ and $V_{/\sim}$ are undefined, otherwise.

Note that $S_{/\sim}$ and $V_{/\sim}$ are well defined partial functions.

Definition 4. Let $H = \langle N, S, V, D \rangle$ be a concrete heap and $H_{/\sim} = \langle N_{/\sim}, S_{/\sim}, V_{/\sim} \rangle$ its quotient. An abstract structure $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle$ is said to be a structural abstraction of H if and only if there exists a bijective function $\beta : N_{/\sim} \cup \{\perp\} \rightarrow \overline{N} \cup \{\perp\}$ such that, $\beta(\perp) = \perp$, and for all $u \in PVar$:

- $\overline{S}(\beta([n])) = \beta(S_{/\sim}([n]))$, and $\overline{V}(u) = \beta(V_{/\sim}(u))$.

Two abstract structures that differ only in the naming of nodes and counter variables are semantically equivalent, in the sense that they are abstractions of the same set of concrete heaps. In practice, this increases the number of abstract structures generated by a symbolic state exploration tool. This problem can be overcome by choosing a canonical representation of abstract structures, as described in e.g. [18]. We define the structural abstraction function $\alpha_s : \overline{\mathcal{H}}(PVar) \rightarrow \overline{\mathcal{H}}(PVar)$, $\alpha_s(H) = \overline{H}$, iff \overline{H} is the canonical representative of a structural abstraction of H .

4.1 Data Insensitive Programs

This section is devoted to the description of counter automata that abstract the behaviour of the programs with lists. We formalize the correctness of our construction by proving bisimulation between the semantics of a list program and the semantics of a counter automaton. This entails the strong preservation of temporal logic properties. In particular, safety and termination are strongly preserved by the counter automaton, meaning that one can accept and/or refute them based on the behaviour of the latter.

Consider a list program with k pointer variables and l counter variables, i.e. $\|PVar\| = k$ and $\|IVar\| = l$. We construct a counter automaton $A = \langle Q, X, \xrightarrow{s} \rangle$ with $2k + l$ counters as follows. The control states Q of the counter automaton are elements of the set $Lab \times (\overline{\mathcal{H}}(PVar) \cup \{H_{err}\})$. Let $\mathcal{X} = \bigcup \{ \overline{N} \mid \langle \overline{N}, \overline{S}, \overline{V} \rangle \in \overline{\mathcal{H}}(PVar) \}$ be the set of nodes used in the structural abstraction. The counters are $X = \{x_n \mid n \in \mathcal{X}\} \cup IVar$, one for each node, and including the counter variables from the original program. The transitions are given by the triples $q \xrightarrow{\phi} q'$ with $q = \langle l, \overline{H} \rangle$, $q' = \langle l', \overline{H}' \rangle$ such that there is a statement $l : s; l'$ in the program and the relation $\overline{H} \xrightarrow[\phi]{s} \overline{H}'$ is described by the structural rules in Fig. 3. Due to a lack of space, the rules for the assignment statements $u := w$, $u := \text{new}$, and $u.\text{next} := \text{null}$ are deferred to the long version of this paper [8].

In order to simplify the treatment of the different cases, we have introduced two low-level operations, that perform merging and splitting of abstract nodes (Figure 3). Intuitively, we need to perform merging of two abstract nodes n and m ($\mu(\overline{H}, n, m)$) in order to re-normalize the abstract structure, after a destructive update. In the case of $u := w.\text{next}$, we need to split ($\sigma(\overline{H}, n, m)$) the abstract node n referred to by w in \overline{H} , into two nodes n and m , based on whether the value of its corresponding counter is greater than one or one ($x_n = 1$, $x_n > 1$). The semantics of conditional tests ($u = v$ and $u = \text{null}$) is similar to the concrete case. For more details concerning the translation, the reader is referred to the list reversal example in Figure 4.

Now we can state the main theorem of this section. Given a data insensitive program P , let $\langle s, \xrightarrow{c} \rangle$ be its concrete semantics with set of states $s = Lab \times (IVar \rightarrow \mathbb{Z}) \times \mathcal{H}(PVar)$ and \xrightarrow{c} its transition relation. Let $\overline{s} = Q \times (X \rightarrow \mathbb{Z})$ be the set of all configurations of the corresponding counter automaton and \xrightarrow{s} its transition relation.

Theorem 2. $\langle s, \xrightarrow{c} \rangle$ and $\langle \overline{s}, \xrightarrow{s} \rangle$ are bisimilar.

List Reversal Example Figure 4 shows the counter automaton for the list reversal program from Figure 1, started with a non-circular list pointed to by i , as input. The counter variable corresponding to each abstract node is depicted inside the node itself.

4.2 Ordered Data Programs

In this section we complete the definition of abstraction for programs with lists, by introducing an abstraction for heaps containing data from an ordered domain $\langle \mathcal{D}, \preceq \rangle$. More precisely, we need to abstract the order relations that may occur inside a list segment, and between two list segments. We shall therefore consider five predicates,

$$\begin{array}{c}
\frac{\exists w \in \text{Var} \setminus \{u\} \quad \overline{V}(w) = \overline{V}(u) \neq \perp}{\overline{H} \xrightarrow[u:=\text{null}]{\text{true}} \langle \overline{N}, \overline{S}, \overline{V}[u \rightarrow \perp] \rangle} \quad \frac{\overline{V}(u) = n \in \overline{N} \quad \forall w \in \text{Var} \setminus \{u\} . \overline{V}(w) \neq n}{\overline{H} \xrightarrow[u:=\text{null}]{\text{true}} \langle \overline{N}, \overline{S}, \overline{V}[u \rightarrow \perp] \rangle} \\
\frac{\overline{V}(u) = n \in \overline{N} \quad \forall w \in \text{Var} \setminus \{u\} . \overline{V}(w) \neq n \quad \exists m \in \overline{N} \setminus \{n\} . \overline{S}(m) = n \quad \forall p \in \overline{N} \setminus \{n\} . \overline{S}(p) \neq n}{\overline{H} \xrightarrow[u:=\text{null}]{x'_m = x_m + x_n} \mu(\langle \overline{N}, \overline{S}, \overline{V}[u \rightarrow \perp] \rangle, m, n)} \quad \frac{\overline{V}(u) = n \in \overline{N} \quad \forall w \in \text{Var} \setminus \{u\} . w \xrightarrow[\overline{H}]{*} n \quad \overline{S}(n) = m \in \overline{N} \setminus \{n\} \quad \forall w \in \text{Var} \setminus \{u\} . \overline{V}(w) \neq m \quad \exists p \in \overline{N} \setminus \{n, m\} . \overline{S}(p) = m \quad \forall q \in \overline{N} \setminus \{n, p\} . \overline{S}(q) \neq m \quad \overline{N}' = \overline{N} \setminus \{n\}}{\overline{H} \xrightarrow[u:=\text{null}]{x'_p = x_p + x_m} \mu(\langle \overline{N}', \overline{S} \downarrow_{\overline{N}'}, \overline{V} \downarrow_{\overline{N}'} \rangle, p, m)} \\
\frac{\overline{V}(w) = n \in \overline{N}}{\overline{H} \xrightarrow[u:=w.\text{next}]{x_n = 1} \langle \overline{N}, \overline{S}, \overline{V}[u \rightarrow \overline{S}(n)] \rangle} \quad \frac{\overline{V}(w) = n \in \overline{N} \quad m \in \mathcal{X} \setminus \overline{N}'}{\overline{H} \xrightarrow[u:=w.\text{next}]{x_n > 1 \wedge x'_m = x_n - 1} \sigma(\langle \overline{N}, \overline{S}, \overline{V}[u \rightarrow m] \rangle, n, m)}
\end{array}$$

Fig. 3. The Counter Automaton for Data Insensitive Programs Let $\overline{H} \triangleq \langle \overline{N}, \overline{S}, \overline{V} \rangle$. The merging function is $\mu : \mathcal{H}(\text{Var}) \times \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{H}(\text{Var})$ given by $\mu(\overline{H}, n, m) = \langle \overline{N}', \overline{S} \downarrow_{\overline{N}'} [n \rightarrow \overline{S}(m)], \overline{V} \rangle$ where $\overline{N}' = \overline{N} \setminus \{m\}$. The splitting function is $\sigma : \mathcal{H}(\text{Var}) \times \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{H}(\text{Var})$ given by $\sigma(\overline{H}, n, m) = \langle \overline{N} \cup \{m\}, \overline{S}', \overline{V} \rangle$ where $\overline{S}' = (\overline{S} \setminus \{(n, p) \mid n \xrightarrow{\overline{H}} p\}) \cup \{(m, p) \mid n \xrightarrow{\overline{H}} p\} \cup \{(n, m)\}$.

with the following meanings. For each predicate P , we denote by P^c its semantics, i.e. the relation it induces on a set of nodes.

Let $H = \langle N, S, V, D \rangle$ be a concrete heap and H/\sim its quotient as defined above. Then, $\text{o}^c([n])$ is true for a list segment $[n]$ iff all its elements are ordered w.r.t. \preceq . Similarly, $[n] \preceq_\diamond [m]$ for $\diamond \in \{ff, fa, af, aa\}$ iff the first (all) element(s) of $[n]$ is (are) less than the first (all) element(s) of $[m]$.

Definition 5. An abstract heap is a tuple $\tilde{H} = \langle \overline{H}, \text{o}, \preceq_{ff}, \preceq_{fa}, \preceq_{af}, \preceq_{aa} \rangle$, where $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle$ is an abstract structure, $\text{o} \subseteq \overline{N}$ is a unary ordering predicate, and $\preceq_{ff, fa, af, aa} \subseteq \overline{N} \times \overline{N}$ are binary ordering predicates.

An abstract heap $\tilde{H} = \langle \overline{H}, \text{o}, \preceq_{ff}, \preceq_{fa}, \preceq_{af}, \preceq_{aa} \rangle$ sharing the same structure $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle$ as another abstract heap $\tilde{H}' = \langle \overline{H}, \text{o}', \preceq'_{ff}, \preceq'_{fa}, \preceq'_{af}, \preceq'_{aa} \rangle$, is said to be *more precise*, denoted as $\tilde{H} \sqsubseteq \tilde{H}'$, if and only if, for each $n, m \in \overline{N}$ we have $\text{o}(n) \Leftarrow \text{o}'(n)$ and $n \preceq_\diamond m \Leftarrow n \preceq'_\diamond m$, for all $\diamond \in \{ff, fa, af, aa\}$. Intuitively, the absence of a predicate indicates incertitude w.r.t. the concrete ordering configuration. For instance if $\text{o}(n)$ does not hold, this means that, in the concrete setting, n "represents" a list segment that may or may not be ordered.

Given a set S of abstract heaps sharing the same structure, we denote by $\sqcup S$ the least upper bound, and by $\sqcap S$ the greatest lower bound of S , with respect to \sqsubseteq . Note that \sqcup and \sqcap are undefined for sets of abstract heaps that have different structures. The domain of abstract heaps is denoted by $\langle \mathcal{H}(\text{PVar}), \sqsubseteq \rangle$.

Definition 6. Let $H = \langle N, S, V, D \rangle$ be a concrete heap with data from the ordered domain $\langle \mathcal{D}, \preceq \rangle$ and $H/\sim = \langle N/\sim, S/\sim, V/\sim \rangle$ its quotient. An abstract heap $\tilde{H} = \langle \overline{H}, \text{o}, \preceq_{ff}$

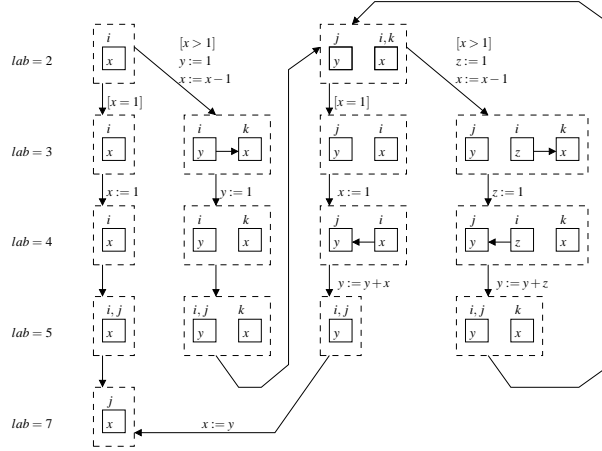


Fig. 4. Non-circular List Reversal

, \preceq_{fa} , \preceq_{af} , \preceq_{aa}) is said to be an abstraction of H if and only if $\alpha_s(H) = \overline{H}$, where β is the bijection from Definition 4, and for all $[n], [m] \in N_{/\sim}$, $\diamond \in \{ff, fa, af, aa\}$:

- $\circ(\beta([n])) \Rightarrow \circ^c([n])$ and $\beta([n]) \preceq_\diamond \beta([m]) \Rightarrow [n] \preceq_\diamond^c [m]$.

We define $\alpha : \mathcal{H}(PVar) \rightarrow \tilde{\mathcal{H}}(PVar)$ as $\alpha(H) = \sqcap \{ \tilde{H} \mid \tilde{H} \text{ is an abstraction of } H \}$. Note that all abstract heaps that are abstractions of H share the same structure, hence \sqcap is defined for this set. The concretization function is $\gamma : \tilde{\mathcal{H}}(PVar) \rightarrow \mathcal{P}(\mathcal{H}(PVar))$, defined as $\gamma(\tilde{H}) = \{ H \mid \alpha(H) \sqsubseteq \tilde{H} \}$. Clearly, $\gamma(\tilde{H}_1) \subseteq \gamma(\tilde{H}_2)$ if $\tilde{H}_1 \sqsubseteq \tilde{H}_2$, but the dual does not necessarily hold.

4.3 Counter Automata Semantics with Ordering Predicates

Taking ordering predicates $\circ, \preceq_{ff, fa, af, aa}$ into account refines our notion of counter automaton, previously introduced. The counter automaton defined in this section keeps track of the ordering information, allowing one to verify properties related to the ordering of lists, as it is the case for sorting programs, e.g. insertsort, bubblesort, etc.

A counter automaton with ordering predicates is $A_a = \langle Q_a, X, \xrightarrow{a} \rangle$. The set of control states is defined now as $Q_a = Lab \times (\tilde{\mathcal{H}}(PVar) \cup \{H_{err}\})$, and the set of configurations is $s_a = Q_a \times (X \rightarrow \mathbb{Z})$, with the usual notation. In addition to updating the abstract structure, the transition relation \xrightarrow{a} has to also update the ordering predicates. Our goal is to define the "best transformer" in the sense of [14]. More precisely, our loss of information is due only to the choice of ordering predicates, the definition of \xrightarrow{a} does not introduce further imprecision. Theorem 4 below formalizes this statement.

In order to achieve completeness of the abstract operational semantics, we have designed our abstract state transformer function in two stages. The first stage yields the actual change of the predicates, and the second one is an operation of "saturation" whose goal is to add all the predicates that can be derived from the existing ones, on

Weakening	Transitivity
1. $n \preceq_{aa} m \Rightarrow n \preceq_{af} m$	5. $n \preceq_{ff} m \wedge m \preceq_{ff} p \Rightarrow n \preceq_{ff} p$
2. $n \preceq_{aa} m \Rightarrow n \preceq_{fa} m$	6. $n \preceq_{af} m \wedge m \preceq_{ff} p \Rightarrow n \preceq_{af} p$
3. $n \preceq_{af} m \Rightarrow n \preceq_{ff} m$	7. $n \preceq_{ff} m \wedge m \preceq_{fa} p \Rightarrow n \preceq_{fa} p$
4. $n \preceq_{fa} m \Rightarrow n \preceq_{ff} m$	8. $n \preceq_{aa} m \wedge m \preceq_{aa} p \Rightarrow n \preceq_{aa} p$
Reflexivity	9. $n \preceq_{aa} n \wedge n \preceq_{fa} m \Rightarrow n \preceq_{aa} m$
12. $n \preceq_{ff} n$	10. $n \preceq_{ff} m \wedge o(m) \Rightarrow n \preceq_{fa} m$
13. $n \preceq_{aa} n \Rightarrow o(n)$	11. $n \preceq_{af} m \wedge o(m) \Rightarrow n \preceq_{aa} m$
14. $o(n) \Rightarrow n \preceq_{fa} n$	Symmetry
	15. $n \preceq_{af} n \wedge n \preceq_{fa} n \Rightarrow n \preceq_{aa} n$

Fig. 5. Saturation rules

a given abstract heap, without changing the corresponding set of concrete heaps. For the remainder of this section, we fix an abstract heap $\tilde{H} = \langle \bar{H}, o, \preceq_{ff}, \preceq_{fa}, \preceq_{af}, \preceq_{aa} \rangle$, with its abstract structure $\bar{H} = \langle \bar{N}, \bar{S}, \bar{V} \rangle$, and let \tilde{H}' be just like \tilde{H} , except that all the components of the tuples are primed.

Let us begin by the presentation of the second stage. Given an abstract heap \tilde{H} , we define the *saturation* of \tilde{H} to be the most precise abstract heap whose concretization is the concretization of \tilde{H} . This notion of saturation is necessary to define an abstract operational semantics that is most precise w.r.t. the ordering abstraction, induced by the o and $\preceq_{ff, fa, af, aa}$ predicates. Unfortunately, this definition does not allow one to effectively check that \tilde{H}' is the saturation of \tilde{H} for arbitrary abstract heaps. The problem is that the set $\gamma(\tilde{H})$ is infinite. To overcome this problem, we introduce “syntactical” saturation rules, given in Figure 5. The closure of an abstract heap \tilde{H} w.r.t. the rules in Figure 5 is denoted as $sat(\tilde{H})$.

The saturation rules need to be applied with the following premise: if (\tilde{H}, v) is a configuration of the counter automaton, and n is an abstract node of \tilde{H} such that $v(x_n) = 1$, then it must be the case that $o(n)$ and $n \preceq_{\diamond} n$, $\diamond \in \{ff, fa, af, aa\}$ all hold in \tilde{H} . The reason is that, list segments of size one are ordered, and in all possible ordering relations with themselves. The generated counter automaton will test, at each step, for each node $n \in \bar{N}$, that $x_n = 1$ and update the ordering predicates accordingly. The next Theorem shows the soundness and completeness of the saturation rules.

Theorem 3. *Given an abstract heap \tilde{H} , we have $sat(\tilde{H}) = \sqcap \{ \tilde{H}' \mid \gamma(\tilde{H}') = \gamma(\tilde{H}) \}$.*

We define now how the change of abstract predicates is being performed. Most of the rules that affect only the abstract structure of the state are very similar with the data insensitive case. To be more precise, all rules from Figure 3, with the exception of the ones that use the merging (μ) or the splitting (σ) functions, will simply maintain the same predicates between the source and destination of the transition. For example, if we had $\bar{V}(u) = \bar{V}(w) = n$ and $n \preceq_{fa} m$, then the result of applying the statement $u := \text{null}$ is $\bar{V}' = \bar{V}[u \rightarrow \perp]$ and $n \preceq'_{fa} m$. The remaining rules are dealt with by introducing *ordered* versions of the merging and splitting functions, called μ_o and σ_o , respectively. As a general rule, the new merging and splitting operations are performed on saturated abstract heaps, and another saturation is applied to the result, in order to maintain the desired precision.

Let $n, m \in \bar{N}$ be such that $\bar{S}(n) = m$ and m is not a cut point in \bar{H} . We recall that the result of $\mu(\bar{H}, n, m)$ in this case is the abstract structure in which n takes the place of both

n and m . Then $\mu_o(\tilde{H}, n, m) = \langle \mu(\overline{H}, n, m), o', \preceq'_{ff}, \preceq'_{fa}, \preceq'_{af}, \preceq'_{aa} \rangle$ where $o', \preceq'_{ff, fa, af, aa}$ are the (unique) relations on \overline{N} and $\overline{N} \times \overline{N}$ satisfying the following constraints, for all $p \in \overline{N}, q, r \in \overline{N} \setminus \{p, n\}$ and $\diamond \in \{ff, fa, af, aa\}$:

$$\begin{array}{ll} o(n) \wedge o(m) \wedge n \preceq_{aa} m \Leftrightarrow o'(n) & n \preceq_{af} p \wedge m \preceq_{af} p \Leftrightarrow n \preceq'_{af} p \\ n \preceq_{aa} p \wedge m \preceq_{aa} p \Leftrightarrow n \preceq'_{aa} p & p \preceq_{fa} n \wedge p \preceq_{fa} m \Leftrightarrow p \preceq'_{fa} n \\ p \preceq_{aa} n \wedge p \preceq_{aa} m \Leftrightarrow p \preceq'_{aa} n & o(q) \Leftrightarrow o'(q) \text{ and } q \preceq_{\diamond} r \Leftrightarrow q \preceq'_{\diamond} r \end{array}$$

The following Lemma shows that no information is lost by an application of μ_o , provided that the source of the transition was a saturated abstract heap. The intuition is that, by merging two abstract nodes, where one of them is not a cut point, the set of concretizations is preserved.

Lemma 2. *Let $\tilde{H} = \langle \overline{H}, o, \preceq_{ff}, \preceq_{fa}, \preceq_{af}, \preceq_{aa} \rangle \in \tilde{\mathcal{H}}(PVar)$ where $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle \in \overline{\mathcal{H}}(PVar)$, and $n, m \in \overline{N}$ such that $\overline{S}(n) = m$ and m is not a cut point in \overline{H} . Then $\gamma(\tilde{H}) = \gamma(\mu_o(\text{sat}(\tilde{H}), n, m))$.*

The splitting operation on abstract structures replaces one node n with two nodes n and m , such that m becomes the successor of n and the previous successor of n becomes the successor of m . In addition, the effect of the split operation on the ordering predicates is modeled by the rules given in the following. Formally, $\sigma_o(\tilde{H}, n, m) = \langle \sigma(\overline{H}, n, m), o', \preceq'_{ff}, \preceq'_{fa}, \preceq'_{af}, \preceq'_{aa} \rangle$, where $o', \preceq'_{ff, fa, af, aa}$ are the (unique) relations on \overline{N} and $\overline{N} \times \overline{N}$ that satisfy the following constraints, for all $p \in \overline{N} \setminus \{n\}, q, r \in \overline{N} \setminus \{p, n\}$, and all $\diamond \in \{ff, fa, af, aa\}$:

$$\begin{array}{lll} o(n) \Leftrightarrow n \preceq'_{aa} m & p \preceq_{fa} n \Leftrightarrow p \preceq'_{fa} m & p \preceq_{aa} n \Leftrightarrow p \preceq'_{aa} m \\ o(n) \Leftrightarrow o'(m) & n \preceq_{aa} n \Leftrightarrow m \preceq'_{aa} m & n \preceq_{aa} p \Leftrightarrow m \preceq'_{aa} p \\ n \preceq_{af} p \Leftrightarrow m \preceq'_{af} p & n \preceq_{aa} n \Leftrightarrow m \preceq'_{aa} n & o(q) \Leftrightarrow o'(q) \text{ and } q \preceq_{\diamond} r \Leftrightarrow q \preceq'_{\diamond} r \end{array}$$

The following Lemma formalizes the correctness σ_o :

Lemma 3. *Let $\tilde{H} = \langle \overline{H}, o, \preceq_{ff}, \preceq_{fa}, \preceq_{af}, \preceq_{aa} \rangle \in \tilde{\mathcal{H}}(PVar)$, where $\overline{H} = \langle \overline{N}, \overline{S}, \overline{V} \rangle \in \overline{\mathcal{H}}(PVar)$, $n \in \overline{N}$ and $m \notin \overline{N}'$. Then $\gamma(\tilde{H}) = \gamma(\sigma_o(\text{sat}(\tilde{H}), n, m))$.*

A conditional test involving data $u.\text{data} \leq w.\text{data}$ evaluates true in the abstract heap \tilde{H} if and only if $\overline{V}(u) \preceq_{ff} \overline{V}(w)$ holds on $\text{sat}(\tilde{H})$. Otherwise, such tests introduce non-determinism in the generated counter automaton. Therefore, the semantics of the counter automaton is a simulation of the semantics of the original program, but not a bisimulation anymore.

Theorem 4. *Let $\langle l, \iota, H \rangle \in s$ be a concrete program state. Then there exists $\langle l', \iota', H' \rangle \in s$ such that $\langle l, \iota, H \rangle \xrightarrow{c} \langle l', \iota', H' \rangle$ if and only if there exists an abstract state $\langle l, \tilde{H}', \nu' \rangle \in s_a$ such that $\langle l, \alpha(H), \nu \rangle \xrightarrow{a} \langle l', \tilde{H}', \nu' \rangle$ and $H' \in \gamma(\tilde{H}')$.*

The following is a consequence of Theorems 1, 2 and 4.

Corollary 1. *For every program with lists, if its counter automaton is flat, then safety and termination are decidable properties.*

Notice that the number of objects created by a flat list program is always bounded by a constant, therefore its counter automaton is linear positive and restrictive (but not necessarily flat). If this automaton is moreover flat, we can apply Theorem 1.

5 Experimental Results

In order to obtain experimental evidence about how our techniques behave in practice, we have applied them to several non-trivial procedures manipulating singly-linked lists. In particular, we have considered a procedure for *reversing lists*, whose behaviour we have studied both for an *acyclic* as well as *cyclic* input, and then two procedures for sorting lists, namely *InsertSort* and *BubbleSort*.

For all the examples, we generated (by hand—an implementation of the translation procedure is our future work) the corresponding counter automata. Sizes of the automata—after some trivial simplifications joining sequences of states with no variation in the underlying heap graph—varied as follows: (1) 15 states and 3 counters for reversing acyclic lists (no optimizations were used in this case), (2) 11 states and 3 counters for reversing cyclic lists, (3) 88 states and 6 counters for *InsertSort*, and (4) 149 states and 7 counters for *BubbleSort* (we considered an optimized version of the sort with a pointer remembering the already sorted part of the list). For list reversing, no ordering predicates were used.

As for the *safety properties* of the considered programs, we checked that there are no null pointer assignments, no elements are lost, the shape is preserved, and—in the case of the sorting algorithms—that the result is sorted. These properties may be checked by generating a symbolically encoded set of the reachable configurations of the counter automaton corresponding to the program. Using an implementation of the abstract regular model checking technique [10] based on LASH automata libraries [1], the verification took 10 sec for the acyclic list reversion case study and 0.5 sec for cyclic list reversion on a Pentium 4 machine with a 2.6 GHz processor.

Moreover, let us note that all the above properties may often be checked already at the *counter automaton extraction phase*. The checking is mostly straightforward. A slight complication is just checking that no elements of the list are lost via the $u.next := w$ operations. However, even here a simple (fully automatable) heuristic may be used. When we generate a counter automaton state containing a new abstract heap and we can grant that some of its nodes have size one (e.g., after a $u := w.next$ statement), we remember this fact. Later when we again encounter such a heap and we cannot statically guarantee that the appropriate nodes have size one, we may drop the information. Then, when we see that an $u.next := w$ operation is performed on a node for which we remembered that its size is one, we know that we do not lose any list elements. If this is not the case, we have to analyse the dynamic behaviour of the counter automaton and check whether it may actually happen that we lose some elements. In our examples, however, we were able to perform all the checks statically.

In addition to checking safety properties, we have also fully-automatically checked that all the considered programs *terminate*. For checking termination, we analysed the generated counter automata using the tool described in [13]. On the same machine as above, we were able to check termination in 4 sec for reversing acyclic lists, 1.5 sec for reversing cyclic lists, 90 sec for *InsertSort*, and 150 sec for *BubbleSort*.

6 Conclusion

We have presented an approach for automatic verification of programs with 1-selector dynamic linked structures. It is based on using counter automata as accurate abstract models for such programs. These infinite-state models can be handled using various

advanced techniques and tools which have been designed recently for their automatic analysis (e.g., [1, 2, 6]), and in particular concerning checking termination and liveness properties (e.g., [13, 12]). Indeed, using counters referring to the sizes of parts of the heap structure (e.g., list segments) of a program is a powerful means for dealing with quantitative reasoning about programs, and in particular about their termination. Our future work naturally includes extending this approach to more general linked structures such as doubly linked lists, tree-like structures, etc.

References

1. The LASH toolset. <http://www.montefiore.ulg.ac.be/~boigelot/research/lash/>.
2. A. Bouajjani, A. Annichini and M. Sighireanu. TReX: A Tool for Reachability Analysis of Complex Systems. In *CAV*, 2001.
3. B. Boigelot and P. Wolper. Symbolic Verification with Periodic Sets. In *CAV*, 1994.
4. I. Balaban, A. Pnueli, and L. Zuck. Shape Analysis by Predicate Abstraction. In *VMCAI*, volume 3385 of *LNCS*, 2005.
5. S. Bardin, A. Finkel, and D. Nowak. Toward Symbolic Verification of Programs Handling Pointers. In *AVIS*, Barcelona, Spain, 2004.
6. S. Bardin, A. Finkel, J. Leroux, and L. Petrucci. FAST: Fast Acceleration of Symbolic Transition systems. In *CAV*, volume 2725, pages 118–121, 2003.
7. J. Berdine, C. Calcagno, and P. O’Hearn. A Decidable Fragment of Separation Logic. In *FSTTCS*, volume 3328 of *LNCS*, 2004.
8. A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, and T. Vojnar. Programs with list are Counter Automata. Technical report, 2006. to appear.
9. A. Bouajjani, P. Habermehl, P. Moro, and T. Vojnar. Verifying Programs with Dynamic 1-Selector-Linked Structures in Regular Model Checking. In *TACAS*, volume 3440 of *LNCS*, 2005.
10. A. Bouajjani, P. Habermehl, and T. Vojnar. Abstract Regular Model Checking. In *CAV*, volume 3114 of *LNCS*, 2004.
11. M. Bozga and R. Iosif. Quantitative Verification of Programs with Lists. Technical report, 2005.
12. A. Bradley, Z. Manna, and H. Sipma. Termination Analysis of Integer Linear Loops. In *CONCUR*, volume 3653 of *LNCS*, 2005.
13. B. Cook, A. Podelski, and A. Rybalchenko. Abstraction Refinement for Termination. In *SAS*, volume 3672 of *LNCS*, 2005.
14. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, 1977.
15. N. Dor, M. Rodeh, and S. Sagiv. Checking Cleanness in Linked Lists. In *SAS*, volume 1824 of *LNCS*, 2000.
16. A. Finkel, 2006. Personal communication.
17. A. Finkel and J. Leroux. How to Compose Presburger-Accelerations: Applications to Broadcast Protocols. In *FSTTCS*, volume 2556 of *LNCS*, 2002.
18. R. Iosif. Symmetry Reductions for Model Checking of Concurrent Dynamic Software. *STTT*, pages 302–319, 2004.
19. S. Ishtiaq and P. O’Hearn. BI as an assertion language for mutable data structures. In *POPL*, 2001.
20. R. Manevich, E. Yahav, G. Ramalingam, and M. Sagiv. Predicate Abstraction and Canonical Abstraction for Singly-Linked Lists. In *VMCAI*, volume 3385 of *LNCS*, 2005.
21. A. Møller and M.I. Schwartzbach. The Pointer Assertion Logic Engine. In *PLDI*, 2001.
22. S. Sagiv, T.W. Reps, and R. Wilhelm. Parametric Shape Analysis via 3-Valued Logic. *TOPLAS*, 2002.
23. E. Yahav, T. Reps, M. Sagiv, and R. Wilhelm. Verifying Temporal Heap Properties Specified via Evolution Logic. In *ESOP*, 2003.