

# Verifying Properties of Well-Founded Linked Lists

Shuvendu K. Lahiri

Microsoft Research  
shuvendu@microsoft.com

Shaz Qadeer

Microsoft Research  
qadeer@microsoft.com

## Abstract

We describe a novel method for verifying programs that manipulate linked lists, based on two new predicates that characterize reachability of heap cells. These predicates allow reasoning about both acyclic and cyclic lists uniformly with equal ease. The crucial insight behind our approach is that a circular list invariably contains a *distinguished head cell* that provides a handle on the list. This observation suggests a programming methodology that requires the heap of the program at each step to be *well-founded*, i.e., for any field  $f$  in the program, every sequence  $u.f, u.f.f, \dots$  contains at least one head cell. We believe that our methodology captures the most common idiom of programming with linked data structures. We enforce our methodology by automatically instrumenting the program with updates to two auxiliary variables representing these predicates and adding assertions in terms of these auxiliary variables.

To prove program properties and the instrumented assertions, we provide a first-order axiomatization of our two predicates. We also introduce a novel induction principle made possible by the well-foundedness of the heap. We use our induction principle to derive from two basic axioms a small set of additional first-order axioms that are useful for proving the correctness of several programs.

We have implemented our method in a tool and used it to verify the correctness of a variety of nontrivial programs manipulating both acyclic and cyclic singly-linked lists and doubly-linked lists. We also demonstrate the use of indexed predicate abstraction to automatically synthesize loop invariants for these examples.

**Categories and Subject Descriptors:** D.2.4: Software/program verification

**General Terms:** Reliability, security, languages, verification

**Keywords:** First-order axiomatization, well-founded linked lists, decision procedure, automated theorem proving, heap abstraction

## 1. Introduction

Program verification has made enormous progress in the last few decades. A significant contribution to this progress has been made by reasoning techniques based on first-order logic. For a number of programs, the verification condition whose validity establishes the correctness of the program is expressible in a combination of first order theories such as arithmetic, uninterpreted functions, and propositional logic. The pioneering work of Nelson and Oppen [30]

described a method for combining decision procedures for individual theories to yield a decision procedure for the combined theory. Based on their results, a number of powerful automated theorem provers [9, 4] have been developed for program verification. Many successful tools [12, 3, 17] have used these provers for proving program properties.

However, first-order reasoning typically breaks down when we want to prove properties of programs that manipulate heap-allocated linked data structures. The main problem with reasoning about a data structure, such as a linked list, is that it is impossible to express an invariant about all members of a list in first-order logic. To achieve such a specification in general requires the use of the *reachability predicate* which cannot be expressed in first-order logic. Consequently, researchers have investigated richer logics such as combination of first order logic with transitive closure [18] and monadic second-order logic [28]. These approaches are typically unable to harness the advances made in automated theorem proving based on first-order logic.

In this paper, we develop a method for verifying linked data structures based entirely on first-order logic. The idea behind our approach is to provide a first-order approximation of the reachability predicate. Nelson [32] and Lev-Ami et al. [24] have also proposed first-order axiomatizations for linked lists. Our work improves upon these previous approaches in several significant ways. Most notably, our work makes the analysis of cyclic linked lists as uniform as acyclic lists. While theoretically incomplete, we believe that our approach is complete enough for most realistic programs. We have validated this belief by building a tool to mechanize our method. We have used our tool to verify a number of programs manipulating singly-linked and doubly-linked lists.

In the next section, we illustrate the benefits of our analysis for reasoning uniformly about acyclic and cyclic lists.

### 1.1 Motivation

The novelty of our technique is best illustrated by an example. Consider the following function `acyclic_simple` that iterates over an acyclic list pointed to by the variable `head` and sets the `data` field of each list element to 0. This example is representative of a variety of programs that use linked lists as sets and iterate over them.

```
//@ requires head != null
void acyclic_simple() {
    iter = head;
    while (iter != null) {
        iter.data = 0;
        iter = iter.next;
    }
}
```

To write the loop invariant for this program, we require the reachability function  $R_{next}$ , which maps each cell  $u$  to the set containing  $u$  and all cells reachable from  $u$  by following the next field

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL'06 January 11–13, 2006, Charleston, South Carolina, USA.  
Copyright © 2006 ACM 1-59593-027-2/06/0001...\$5.00.

and excluding *null*. The cornerstone of all previous approaches to first-order reasoning about linked lists is the following axiom about  $R_{\text{next}}$ :

$$v \in R_{\text{next}}(u) \Leftrightarrow (v = u \vee (u.\text{next} \neq \text{null} \wedge v \in R_{\text{next}}(u.\text{next})))$$

This axiom and the rest of this section assumes, for simplicity but without loss of generality, that  $\text{null}.\text{next} = \text{null}$ . We have also added the precondition  $\text{head} \neq \text{null}$  to make this example similar in as many ways as the example of a cyclic list traversal described later in the section.

Using  $R_{\text{next}}$ , the loop invariant for the program can be expressed as follows:

$$\forall v \in R_{\text{next}}(\text{head}) : (\text{iter} \neq \text{null} \wedge v \in R_{\text{next}}(\text{iter})) \vee v.\text{data} = 0$$

This loop invariant proves the following desired postcondition of the function.

$$\forall v \in R_{\text{next}}(\text{head}) : v.\text{data} = 0$$

The verification condition for this program is easily constructed using the weakest-precondition transformer [10], and the verification condition can be proved from the axiom using purely first-order reasoning based on uninterpreted functions and quantifier instantiation. We encourage the reader to perform this simple but interesting calculation.

Unfortunately, this simple approach breaks down with cyclic lists which are quite common, and perhaps the most common linked data structure used in operating system kernels. To illustrate the problem, consider the following variant `cyclic_simple` of the first function that iterates over a cyclic list pointed to by `head` and sets the `data` fields of the list elements to 0. The complication, in this case, is that the `next` field of the last element of the list points to `head` rather than `null`.

```
//@ requires head != null
void cyclic_simple() {
    head.data = 0;
    iter = head.next;
    while(iter != head) {
        iter.data = 0;
        iter = iter.next;
    }
}
```

The circularity of the list results in a breakdown of the reasoning described earlier. The value of  $R_{\text{next}}(\text{head})$  and  $R_{\text{next}}(\text{iter})$  remains constant and equal to the set of all elements in the list during the entire execution of the loop. As a result, we are unable to write a simple loop invariant using  $R_{\text{next}}$ . Even proving the absence of null dereferences in `cyclic_simple` is nontrivial compared to `acyclic_simple` where it is trivial because of the loop entry condition. In the remainder of this section, we show how our approach makes verification of programs manipulating cyclic lists as simple as those manipulating acyclic lists.

The crucial insight behind our approach is that a circular list invariably contains a *distinguished head cell*. This head cell provides a handle on the list, usually marks the beginning of the list, and the last element of the list points to it. In the program above, this cell is pointed to by the program variable `head`. Moreover, it is the presence of this cell that ensures the termination of iterations over circular linked lists. Let  $H_{\text{next}}$  denote the set of head cells of a program. Then, the above program has the precondition  $\{\text{head}, \text{null}\} \subseteq H_{\text{next}}$ . We think of *null* as a special head and consequently require that  $H_{\text{next}}$  always contain *null*. The usefulness of constraining  $H_{\text{next}}$  in this fashion will become clear later in this section.

We introduce a new axiomatization of linked lists using functions  $R_{\text{next}}$  and  $B_{\text{next}}$ . In contrast to its definition earlier, the new

definition of function  $R_{\text{next}}$  maps each cell  $u$  to the set of cells containing  $u$  and all cells reachable from  $u$  by following the `next` field until a cell  $v \in H_{\text{next}}$  is reached. The cell  $v$  is not included in  $R_{\text{next}}(u)$  unless  $u = v$ . Finally, we define  $B_{\text{next}}(u)$  to be  $v$ . The cell  $v$  acts as a block to the traversal of the `next` field from  $u$  and its identity is captured as  $B_{\text{next}}(u)$ . We call  $B_{\text{next}}(u)$  the *blocking cell* for  $u$ . The axioms for  $R_{\text{next}}$  and  $B_{\text{next}}$  are as follows:

$$v \in R_{\text{next}}(u) \Leftrightarrow (u = v \vee (u.\text{next} \notin H_{\text{next}} \wedge v \in R_{\text{next}}(u.\text{next})))$$

$$B_{\text{next}}(u) = (u.\text{next} \in H_{\text{next}} ? u.\text{next} : B_{\text{next}}(u.\text{next}))$$

The first axiom is similar in spirit to the axiom for  $R_{\text{next}}$  described earlier. The second axiom says that  $B_{\text{next}}(u)$  is  $u.\text{next}$  if  $u.\text{next} \in H_{\text{next}}$  and  $B_{\text{next}}(u.\text{next})$  otherwise.

Using the definition of  $B_{\text{next}}$ , we can specify that `head` points to a circular list by adding the following precondition to the function `cyclic_simple`.

$$\text{//@ requires } B_{\text{next}}(\text{head}) = \text{head}$$

From the definition of  $R_{\text{next}}$  and  $B_{\text{next}}$ , it is clear that if  $B_{\text{next}}(u) = u$ , then  $u \in H_{\text{next}}$  and  $R_{\text{next}}(u)$  is the set of cells in a cycle in which no cell other than  $u$  is in  $H_{\text{next}}$ . Therefore, the set of cells in the circular list pointed to by `head` is given by  $R_{\text{next}}(\text{head})$ . Using the definition of  $R_{\text{next}}$  and  $B_{\text{next}}$ , we can now write the loop invariant for the loop in `cyclic_simple` as follows.

$$\forall v \in R_{\text{next}}(\text{head}) :$$

$$(\text{iter} \neq \text{head} \wedge v \in R_{\text{next}}(\text{iter})) \vee v.\text{data} = 0$$

$$\wedge B_{\text{next}}(\text{iter}) = \text{head}$$

Observe that the first conjunct of the loop invariant is similar in structure to the loop invariant for the function `acyclic_simple`. The two new axioms allow us to prove the correctness of the loop invariant, the absence of null dereferences, and the following desired postcondition for `cyclic_simple`:

$$\forall v \in R_{\text{next}}(\text{head}) : v.\text{data} = 0$$

As mentioned earlier, proving the absence of null dereferences in `cyclic_simple` is nontrivial and requires the use of the axiom about  $B_{\text{next}}$ .

Our new axiomatization is a generalization of the first axiomatization that worked only for acyclic lists. If  $H_{\text{next}} = \{\text{null}\}$ , the new axiom about  $R_{\text{next}}$  reduces to the first axiom about  $R_{\text{next}}$ . Indeed, we can verify the first program using the new axiomatization by adding the precondition  $H_{\text{next}} = \{\text{null}\}$  to the function `acyclic_simple`. Thus, our new axiom system can verify programs manipulating acyclic as well as cyclic lists with equal ease.

## 1.2 Contributions

The main technical contribution of this paper is a novel method for verifying linked lists based on two new predicates that characterize reachability of heap cells. These predicates allow reasoning about both acyclic and cyclic lists uniformly with equal ease. The crucial insight behind our approach is that a circular list invariably contains a *distinguished head cell* that provides a handle on the list. This observation suggests a programming methodology that requires the heap of the program at each step to be *well-founded*, i.e., for any field  $f$  in the program, every sequence  $u.f, u.f.f, \dots$  contains at least one head cell. The set of head cells is identified by a new variable added to the program by the programmer.

We believe that our methodology captures the most common idiom of programming with linked data structures. We enforce our methodology by automatically instrumenting the program with updates to two auxiliary variables representing these predicates and adding assertions in terms of these auxiliary variables. Our instrumentation captures well-foundedness *precisely* — the instrumented program fails one of these assertions if and only if the original program reaches a state containing a heap that is not well-founded.

```

/*@ requires a != null && b != null
/*@ requires {null, a, b} ⊆ Hnext
/*@ requires Bnext(a) == a && Bnext(b) == b
/*@ requires Rnext(a) ∩ Rnext(b) == {}
/*@ requires ∀u ∈ Rnext(b) : u.r == b.r

/*@ ensures Bnext(a) == b && Bnext(b) == a
/*@ ensures ∀u ∈ (Rnext(a) ∪ Rnext(b)) : u.r == b.r
/*@ ensures
  (Rnext(a) ∪ Rnext(b)) == (old(Rnext)(a) ∪ old(Rnext)(b))
/*@ ensures ∀u ∉ old(Rnext)(a) : u.r == old(u.r)

void union(Cell a, Cell b) {
  a.r = b.r;
  Cell curr = a.next;
  while (curr != a) {
    curr.r = b.r;
    curr = curr.next;
  }
  Cell tmp = a.next;
  a.next = b.next;
  b.next = tmp;
}

```

**Figure 1.** Performing the union of two circular lists

To prove program properties and the instrumented assertions, we provide a first-order axiomatization of our two predicates. Our axiomatization consists of the two basic first-order axioms introduced in Section 1.1. We also introduce a novel induction principle made possible by the well-foundedness of the heap. We use our induction principle to derive from the two basic axioms a small set of additional first-order axioms that are useful for proving the correctness of several programs. All of these derived axioms are intuitive and natural and we state them precisely in Section 5.

We also explore an alternative verification approach based on a decision procedure rather than a first-order axiomatization. In Section 6, we show that the quantifier-free fragment of the theory of  $R_{\text{next}}$  and  $B_{\text{next}}$  is NP-complete and propose a decision procedure for this fragment based on a small-model encoding.

We have implemented our method in a tool and used it to verify the correctness of a variety of nontrivial programs manipulating both acyclic and cyclic singly-linked lists and doubly-linked lists. We also demonstrate the use of indexed predicate abstraction [13, 21] for automatically synthesizing the loop invariants for these examples. We describe our preliminary experience with this tool in Section 7. For lack of space, we have omitted the proofs of the various theorems and lemmas in the paper; they can be found in the full version of the paper [23].

## 2. Examples

In Section 1.1, we illustrated our method by verifying the functions `acyclic_simple` and `cyclic_simple`. These functions were interesting and illustrative but comparatively simple. In this section, we illustrate our method on three more complex examples. The first example in Section 2.1 modifies the link structure of the heap; the second example in Section 2.2 uses arithmetic in addition to heap manipulation; the third example in Section 2.3 manipulates a doubly-linked list. We note that the examples in Section 1.1 and all examples of this section have been verified by the tool whose details we provide in Section 7.

### 2.1 Set-union

Our first example is the function `union` shown in Figure 1. This example is taken from a paper by Nelson [32]. The function `union` takes two circular linked lists, `a` and `b`, as arguments. Each list represents a set; the field `r` in a cell contains the identifier of the unique set to which it belongs.

The function `union` has a number of preconditions also stated in Figure 1. The second precondition says that both `a` and `b` are pointing to head cells. The third precondition says that `a` and `b` are pointing to circular lists. The fourth precondition says that the lists pointed to by `a` and `b` are disjoint.

The objective of `union` is to merge the list pointed to by `a` into the list pointed to by `b`. The function `union` first sets the `r` field of each cell in `a` to `b.r`. Finally, the contents of `a.next` and `b.next` are swapped to merge the two lists.

Our tool automatically instruments `union` with updates to the functions  $B_{\text{next}}$  and  $R_{\text{next}}$ . This instrumentation is introduced only for the last two statements, which are the only statements that update the `next` field. For each of these statements, the instrumentation also checks, by inserting an assertion, that the update of `next` leaves the heap well-founded. For example, the assertion introduced just before the last statement of the function is as follows:

$$\text{assert}(b \in R_{\text{next}}(\text{tmp}) \Rightarrow \text{tmp} \in H_{\text{next}})$$

This assertion fails precisely if there is a chain of cells connected by `next` from `tmp` to `b` in which no cell is in  $H_{\text{next}}$ . Such a chain causes an ill-founded cycle of cells upon execution of the statement `b.next = tmp`. Furthermore, this statement modifies  $B_{\text{next}}(b)$  and  $R_{\text{next}}(b)$  as follows:

$$\begin{aligned} B_{\text{next}}(b) &= (\text{tmp} \in H_{\text{next}}) ? \text{tmp} : B_{\text{next}}(\text{tmp}) \\ R_{\text{next}}(b) &= (\text{tmp} \in H_{\text{next}}) ? \{b\} : \{b\} \cup R_{\text{next}}(\text{tmp}) \end{aligned}$$

A precise and complete description of both  $B_{\text{next}}$  and  $R_{\text{next}}$  is given in Section 4.

For `union`, we prove the absence of null dereferences, the assertions described above to check that the heap remains well-founded, and four postconditions. The first postcondition is particularly interesting. It states that the blocking cell for `a` is `b` and the blocking cell for `b` is `a`. These two facts together mean that cells `a` and `b` are the only two head cells in a cycle, which indicates indirectly that indeed the union of the two lists has been created. In the third postcondition, `old(Rnext)` refers to the value of  $R_{\text{next}}$  at the beginning of the function. The final postcondition says that the `r` field remains unchanged for all those cells which do not belong to  $R_{\text{next}}(a)$  initially.

So far, the set of head cells has been a constant. The first postcondition also motivates the need to modify the set of head cells. It would be intuitively more satisfactory if the programmer can remove the cell `a` from  $H_{\text{next}}$  once the list `a` has been merged into `b`. Therefore, we allow the programmer to remove a cell from  $H_{\text{next}}$  by using the `Remove` operation. For example, we could add the statement

$$H_{\text{next}}.\text{Remove}(a);$$

at the very end of `union`. With this modification, our tool proves the following more pleasing postconditions.

```

/*@ ensures Bnext(b) == b && Bnext(a) == b
/*@ ensures ∀u ∈ Rnext(b) : u.r == b.r
/*@ ensures Rnext(b) == (old(Rnext)(a) ∪ old(Rnext)(b))

```

Since the values of the variables  $B_{\text{next}}$  and  $R_{\text{next}}$  depend on  $H_{\text{next}}$ , the instrumentation for this statement is nontrivial. Since  $H_{\text{next}}$  becomes smaller, the heap might not remain well-founded. The instrumentation not only checks via an assertion that the heap remains well-founded, but also updates the values of  $B_{\text{next}}$  and  $R_{\text{next}}$  appropriately. We give a precise description of the instrumentation

```

/*@ requires {null} ⊆ Hnext
/*@ requires Bnext(l) == null
/*@ requires l != null && p != null
/*@ requires p ∉ Hnext
/*@ requires p.data > l.data
/*@ requires sorted(l)
/*@ requires p ∉ Rnext(l)

/*@ ensures sorted(l)
/*@ ensures p ∈ Rnext(l)
/*@ ensures
  ∀u : u ∈ Rnext(l) <=> (u ∈ old(Rnext)(l) || u == p)

void insert(Cell l, Cell p) {
  Cell curr = l;
  Cell succ = l.next;
  while (succ != null) {
    if (p.data > succ.data) {
      curr = succ;
      succ = curr.next;
    }
    else
      break;
  }
  p.next = succ;
  curr.next = p;
}

```

**Figure 2.** Inserting an element into a sorted acyclic list

for this statement as well as for  $H_{next}.Add(x)$ , the converse of this statement, in Section 4.

## 2.2 Insert

Our second example is the function `insert` shown in Figure 2. The function `insert` takes an acyclic list `l` and a cell `p` as arguments. The cells in the list `l` are in sorted order based on the values of the field `data` of the cells. The predicate `sorted(l)` is defined as follows:

```

sorted(l) =
  ∀u ∈ Rnext(l) : u.next == null || u.data <= u.next.data

```

The objective of `insert` is to insert the cell `p` in the appropriate place in `l` so that `l` remains sorted. The fourth precondition requires that the cell `p` is not in  $H_{next}$ . If  $p \in H_{next}$ , then on return,  $R_{next}(l)$  contains all the cells from `l` up to but not including `p`, and thus violates the the second postcondition  $p \in R_{next}(l)$ . The fifth precondition of `insert` is also worth explaining. To simplify the coding of `insert`, it is expected that the first element of `l` is a dummy whose `data` field is guaranteed to be less than any value that might be inserted in the list.

This example illustrates an important advantage of axiomatizing linked lists in first-order logic. The specification of `insert` uses a combination of facts about reachability via the `next` field, uninterpreted functions, arithmetic, and propositional logic. The axiomatization of linked lists in first-order logic immediately allows us to use any one of a number of theorem provers that deal with a combination of first-order theories. For our implementation, we are using the UCLID theorem prover [7]. Although this example does not use any arrays, adding them is a simple matter because they can be modeled easily with uninterpreted functions and the well-known select-update axioms.

```

/*@ requires wf_dlist_head(hd)
/*@ requires p != null
/*@ requires p ∈ Rnext(hd)
/*@ requires p != hd
/*@ requires Bnext(p) == hd
/*@ requires p ∉ Hnext && p ∉ Hprev

/*@ ensures p ∉ Rnext(hd)
/*@ ensures wf_dlist_head(hd)
/*@ ensures
  ∀u : (u ∈ Rnext(hd) || u == p) <=> u ∈ old(Rnext)(hd)

void dlist_remove(Cell hd, Cell p) {
  Cell tp = p.prev;
  Cell tn = p.next;
  tp.next = tn;
  tn.prev = tp;
}

```

**Figure 3.** Removing an element from a doubly-linked list

## 2.3 Remove

Our third example is the function `dlist_remove` given in Figure 3. This function removes a cell `p` from a cyclic doubly-linked list with head `hd`. This example illustrates that our technique handles doubly-linked lists just as cleanly as singly-linked lists.

Our instrumentation adds variables  $R_{next}$  and  $B_{next}$  for the linking field `next`, and  $R_{prev}$  and  $B_{prev}$  for the linking field `prev`. The instrumentation happens just as before for each linking field as if they are independent. However since the linking field pair (`next`, `prev`) forms a doubly linked list, we need to define additional preconditions to relate the two fields and their auxiliary variables.

We define a predicate `wf_dlist_head(hd)` to denote that `hd` points to a well formed doubly linked list. This predicate is a conjunction of the following predicates:

1.  $hd \neq null$
2.  $hd \in H_{next} \ \&\& \ hd \in H_{prev}$
3.  $\forall u \in R_{next}(hd) : u.next.prev == u \ \&\& \ u.prev.next == u$
4.  $R_{next}(hd) == R_{prev}(hd)$
5.  $B_{next}(hd) == hd \ \&\& \ B_{prev}(hd) == hd$

The last conjunct indicates that `hd` is the unique head cell in  $H_{next}$  and  $H_{prev}$  present in the cyclic list. The predicate `wf_dlist_head(hd)` is particularly useful when the routine `dlist_remove` is invoked in a loop to, for example, remove all the cells from the list that satisfy some property. Then this predicate will be a conjunct in the loop invariant.

The first precondition of the routine requires that the list is well formed. The next two preconditions state that `p` points to a cell in the list. The next precondition states that `p` is different from `hd` and is required otherwise the list becomes ill-founded after the remove operation. The postconditions assert that `p` is removed from the list and the list remains well-formed.

## 3. Programs

The set *Var* is a set of program variables. A variable in *Var* may have one of three types: *Boolean*, *Integer*, or *Cell*. The set *Cell* contains the addresses of heap objects, each of which may have fields from the set *Field*. A field may also have one of the three aforementioned types. A field of type *Boolean* is a map from *Cell*

$Boolean$	$=$	$\{false, true\}$
$Integer$	$=$	$\{\dots, -2, -1, 0, 1, 2, \dots\}$
$u, v, w \in Cell$	$=$	$\{null, \dots\}$
$x, y, z \in Var$		
$f, g, h \in Field$		
$c \in Const$	$=$	$Boolean \cup Integer \cup \{null\}$
$Op$	$=$	$\{+, -, ==, !=, <, \&\&,   \}$
$s \in Stmt$	$::=$	$x = c \mid x = y \mid x = y \ Op \ z \mid$ $assume(x) \mid assert(x) \mid$ $x = new \mid x = y.f \mid x.f = y$

**Figure 4.** Program syntax

to *Boolean*, a field of type *Integer* is a map from *Cell* to *Integer*, and a field of type *Cell* is a map from *Cell* to *Cell*. We refer to the last category of fields as reference fields. Finally, *Const* is the set of constants that may appear in the program. The constant *null* is a special constant of type *Cell*.

A program is a *control flow graph*  $(PC, E, pc_i, pc_f, L)$  with five components.  $PC$  is a finite set of program locations,  $E \subseteq PC \times PC$  is the set of control-flow edges,  $pc_i \in PC$  is the initial location where program execution begins,  $pc_f \in PC$  is the final location where program execution terminates, and  $L$  is a function that maps each edge in  $E$  to a statement in *Stmt*.

The restricted syntax of the statements in *Stmt* does not result in any loss of generality, because more complex statements can be reduced to simple statements by the introduction of temporary variables. The statement `assume(x)` together with nondeterminism in the control flow graph can be used to encode *if-then-else* and *while* statements. Moreover, other operations can be encoded with the operations in *Op*. For example, boolean negation of  $x$  can be encoded as  $x == false$ . We also assume that the program is free of typing errors.

### 3.1 Programming with heads

We allow programmers to specify a subset of the reference fields as *linking* fields. These fields are expected to provide the links in data structures such as singly-linked and doubly-linked lists. For each linking field  $f$ , the programmer is allowed access to a variable  $H_f$  of type  $Set(Cell)$ . The set  $H_f$  contains a subset of the set of allocated cells that together make all lists linked by  $f$  well-founded. The set  $H_f$  is required to always contain the cell *null*. We provide two operations to update  $H_f$ .

1.  $H_f.Add(x)$ : The variable  $x$  is required to be of type *Cell*. This statement adds the cell obtained by evaluating  $x$  to the set  $H_f$ . This statement provides the fundamental mechanism for making circular lists linked by the field  $f$  well-founded. A programmer creates such a list by first creating the head cell and then adding it to the set  $H_f$ .
2.  $H_f.Remove(e)$ : Again, the variable  $x$  is required to be of type *Cell*. This operation requires that  $x$  be different from *null*. This statement removes the cell obtained by evaluating  $x$  from the set  $H_f$ . The precondition ensures that we never remove *null* from  $H_f$ . We have already illustrated the utility of this operation in Section 2.1. This operation is used to remove the head cells of those circular linked lists that are merged into other lists.

In addition to the usual statements, we also allow the edges of the control flow graph to be labelled with the above operations on  $H_f$  for each linking field  $f$ . In Section 4, we describe how we ensure well-foundedness of the heap by automatically instrumenting the program with assertions at every occurrence of an update to the linking fields and every occurrence of the two operations described here.

### 3.2 Semantics

The state  $\sigma$  of the program contains a program counter in *PC* and a valuation of the variables in *Var*, the fields in *Field*, and the head variables  $H_f$  for each linking field  $f$ . Variables of type *Boolean* are initialized to *false*, variables of type *Integer* are initialized to 0, and variables of type *Cell* are initialized to *null*.

The state also contains a special variable `Alloc` of type  $Set(Cell)$  to model memory allocation via `new`. The variable `Alloc` is initialized to  $\{\}$ . The statement  $x = new$  removes a nondeterministically chosen cell that is not equal to *null* from  $Cell \setminus Alloc$ , assigns it to  $x$ , and adds it to `Alloc`. Thus, the statement  $x = new$  is desugared to the following statements:

```
x = choose(Cell \ {null} \ Alloc);
Alloc = Alloc ∪ {x};
```

When the program executes, its state changes in accordance with the standard operational semantics of its control flow graph. To model the program misbehaving by performing a dereference of *null* or by failing an assertion, we add a special state *wrong* with no transition out of it. For example, if the value of the program counter in  $\sigma$  is  $l$  and the edge  $(l, l') \in E$  is labelled with  $x.f = y$ , then one of the following may happen.

1.  $\sigma(x)$  is *null* and the program moves to the state *wrong*.
2.  $\sigma(x)$  is not *null* and the program makes a transition to a state  $\sigma'$  in which the program counter is  $l'$ , the map for  $f$  is modified only at cell  $\sigma(x)$  to  $\sigma(y)$ , and otherwise the state remains unchanged.

The operational semantics of the other statements can be defined similarly.

## 4. Program instrumentation

In this section, we show how to automatically instrument a program to ensure that the linked lists in the heap always remain well-founded. The instrumentation is performed with respect to a subset of reference fields called *linking* fields that act as the links in a linked structure. We automatically instrument the program with two auxiliary variables for every linking field  $f$ . These variables, called  $R_f$  and  $B_f$ , record information about the shape of the heap graph and are essential to our verification method. The variable  $R_f$  is a map from *Cell* to  $Set(Cell)$ . The variable  $B_f$  is a map from *Cell* to *Cell*. Intuitively, for any  $u \in Cell$ ,  $R_f(u)$  is the set of heap cells containing  $u$  and every cell obtained by one or more applications of  $f$  to  $u$  until a cell  $v \in H_f$  is reached and  $B_f(u) = v$ . Note that the final cell  $v$  is not a member of  $R_f(u)$ . We will often write  $R_f(u, v)$  to denote  $v \in R_f(u)$ .

Although the variables  $R_f$  and  $B_f$  can be defined as a mathematical function of the program state, this definition uses transitive closure and therefore cannot be expressed in first-order logic. An important insight of our work is that even though these variables cannot be defined without using transitive closure, the *updates* to them as the program state changes can be defined entirely in first-order logic. We automatically instrument the program to record these updates and are thus able to state the assertions which ensure that the heap remains well-founded.

Let  $P$  be a program with control flow graph  $(PC, E, pc_i, pc_f, L)$ . The instrumented program  $P^\#$  is obtained by instrumenting each individual statement  $L(e)$  for each edge  $e \in E$ . We now define the instrumentation for a given statement  $s$ . In many cases shown below, the instrumented statement contains control flow. Such a complex statement is used only for clarity of presentation, and can easily be translated into a simple statement without control flow.

1. If the statement  $s$  is of the form  $x = c$ ,  $x = y$ ,  $x = y$  *Op*  $z$ , `assume(x)`, `assert(x)`,  $x = \text{new}$ , or  $x = y.f$ , then the instrumented statement is identical to  $s$ . The definition of  $R_f$  and  $B_f$  depends only on the value of `map f` and the set  $H_f$ . Since neither of those are changed by the statement, we do not need to update the instrumentation variables.
2. If the statement is of the form  $x.f = y$  and  $f$  is a linking field, then the instrumented statement is

```

assert(x != null);
assert(x ∈ R_f(y) => y ∈ H_f);
if (y ∈ H_f) {
    B_f = λi. {
        y      if x ∈ R_f(i)
        B_f(i) otherwise
    }
    R_f = λi. {
        (R_f(i) \ R_f(x)) ∪ {x}  if x ∈ R_f(i)
        R_f(i)                  otherwise
    }
} else {
    B_f = λi. {
        B_f(y)  if x ∈ R_f(i)
        B_f(i)  otherwise
    }
    R_f = λi. {
        (R_f(i) \ R_f(x)) ∪ {x} ∪ R_f(y)  if x ∈ R_f(i)
        R_f(i)                            otherwise
    }
}
x.f = y;

```

This statement is the only one that updates the link structure of the heap. First, the instrumentation checks via the assertion `assert(x ∈ R_f(y) => y ∈ H_f)` that the heap remains well-founded, that is, no cycle of cells unbroken by a member of  $H_f$  is created as a result of updating the  $f$  field of  $x$  to  $y$ . If such a cycle is created, then there must be a path from  $y$  to  $x$  in which no cell, including  $y$  and  $x$ , is in  $H_f$ , which results in a violation of this assertion.

Second, the two instrumentation variables  $B_f$  and  $R_f$  are updated. The values of these functions is updated at a particular cell  $i$  only if  $x \in R_f(i)$ , that is, there is a path from  $i$  to  $x$  not broken by any member of  $H_f$ . If  $x \in R_f(i)$ , the update is split into two cases. If  $y \in H_f$ , then  $B_f(i)$  becomes  $y$  and we remove from  $R_f(i)$  everything that is reachable from  $x$  without hitting a member of  $H_f$  but then add  $x$  itself. If  $y \notin H_f$ , then  $B_f(i)$  becomes  $B_f(y)$  and we remove from  $R_f(i)$  everything that is reachable from  $x$  but then add  $x$  itself and everything that is reachable from  $y$ .

3. If the statement is of the form  $H_f.\text{Add}(x)$ , then the instrumented statement is

```

if (x ∉ H_f) {
    B_f = λi. {
        x      if x ∈ R_f(i) && x != i
        B_f(i) otherwise
    }
    R_f = λi. {
        R_f(i) \ R_f(x)  if x ∈ R_f(i) && x != i
        R_f(i)          otherwise
    }
    H_f = H_f ∪ {x};
}

```

This statement does not update the linking structure of the heap; it only changes the value of  $H_f$ . Since the values of the instrumentation variables  $B_f$  and  $R_f$  depend on  $H_f$ , these variables must be updated if  $H_f$  changes in case  $x \notin H_f$ . The

values of  $B_f$  and  $R_f$  are updated at a particular cell  $i$  only if  $x \in R_f(i) \wedge x \neq i$ , that is,  $x$  becomes new blocking cell for  $i$ . In this case,  $B_f(i)$  becomes  $x$  and we remove from  $R_f(i)$  everything that is reachable from  $x$  without hitting a member of  $H_f$ .

4. If the statement is of the form  $H_f.\text{Remove}(x)$ , then the instrumented statement is

```

if (x ∈ H_f) {
    assert(B_f(x) != x);

    B_f = λi. {
        B_f(x)  if B_f(i) == x
        B_f(i)  otherwise
    }

    R_f = λi. {
        R_f(i) ∪ R_f(x)  if B_f(i) == x
        R_f(i)          otherwise
    }

    H_f = H_f \ {x};
}

```

This statement, just like  $H_f.\text{Add}(x)$ , does not update the linking structure of the heap; it only changes the value of  $H_f$ . The values of the instrumentation variables  $B_f$  and  $R_f$  must be updated if  $H_f$  changes in case  $x \in H_f$ . First, the instrumentation checks via an assertion that the removal of  $x$  from the set of head cells does not result in a heap that is not well-founded. A bad heap can result if  $x$  is the only head cell in a cycle, a condition captured by  $B_f(x) = x$  and the assertion checks for precisely the negation of this condition. Note that since  $B_f(\text{null}) = \text{null}$  by definition, this assertion also checks the precondition of this operation that  $x$  is nonnull.

Second, the two instrumentation variables  $B_f$  and  $R_f$  are updated. The values of  $B_f$  and  $R_f$  are updated at a particular cell  $i$  only if  $B_f(i) = x$ , that is,  $x$  is the blocking cell for  $i$ . In this case,  $B_f(i)$  becomes  $B_f(x)$  and we add to  $R_f(i)$  everything that is reachable from  $x$  without hitting a member of  $H_f$ .

#### 4.1 Correctness

In this section, we formalize the correctness of our instrumentation. The instrumented state  $\theta$  of the program is a valuation of the variables in  $R_f$  and  $B_f$  for each linking field  $f$ . While the original program makes transitions of the form  $\sigma \rightarrow \sigma'$ , the instrumented program makes transitions of the form  $(\sigma, \theta) \rightarrow (\sigma', \theta')$ . To state our correctness theorem, we first need to define a well-formed state.

**DEFINITION 1 (Well-founded function).** A function  $f : \text{Cell} \rightarrow \text{Cell}$  is well-founded with respect to a set of cells  $H$ , if for any cell  $u \in \text{Cell}$ , there exists  $n > 0$ , such that  $f^n(u) \in H$ .

If  $f$  is well-founded with respect to  $h$ , we can define the functions  $B_H^f : \text{Cell} \rightarrow \text{Cell}$  and  $R_H^f : \text{Cell} \rightarrow \text{Set}(\text{Cell})$  as follows.

$$\begin{aligned}
 R_H^f(u) &= \{v \in \text{Cell} \mid v = u \vee \\
 &\quad \exists n : 0 < n : (v = f^n(u) \wedge \\
 &\quad \quad \forall m : 0 < m \leq n : f^m(u) \notin H)\} \\
 B_H^f(u) &= f^n(u), \text{ where } n = \min\{m \mid 0 < m \wedge f^m(u) \in H\}
 \end{aligned}$$

**DEFINITION 2 (Well-founded state).** A state  $\sigma$  is well-founded if for every linking field  $f$ , the function  $\sigma(f)$  is well-founded with respect to  $\sigma(H_f)$ .

**DEFINITION 3 (Well-formed state).** A state  $\sigma$  is well-formed if the following conditions are satisfied:

1.  $\sigma$  is well-founded.
2.  $\text{null} \notin \text{Alloc}$ .

3. For every variable  $x$  of type  $Cell$ ,  $\sigma(x) \in \sigma(\text{Alloc})$  or  $\sigma(x) = \text{null}$ .
4. For every cell  $u \in \sigma(\text{Alloc})$  and reference field  $\mathbf{f} \in \text{Field}$ ,  $\sigma(\mathbf{f})(u) \in \sigma(\text{Alloc})$  or  $\sigma(\mathbf{f})(u) = \text{null}$ .
5. For every cell  $u \in Cell \setminus \sigma(\text{Alloc})$  and reference field  $\mathbf{f} \in \text{Field}$ ,  $\sigma(\mathbf{f})(u) = \text{null}$ .
6. For every linking field  $\mathbf{f}$ ,  $\text{null} \in \sigma(\mathbb{H}_{\mathbf{f}})$ .

A state is *ill-formed* if it is not well-formed.

We define an instrumentation function  $I$  on well-formed states. For every program state  $\sigma$ , the application  $I(\sigma)$  yields an instrumented state  $\theta$  such that  $\theta(\mathbb{B}_{\mathbf{f}}) = B_H^{\mathbf{f}}$  and  $\theta(\mathbb{R}_{\mathbf{f}}) = R_H^{\mathbf{f}}$ , where  $f = \sigma(\mathbf{f})$  and  $H = \sigma(\mathbb{H}_{\mathbf{f}})$ . Now, we can state a theorem that characterizes the precision of our instrumentation.

**THEOREM 1.** *Suppose  $\sigma$  and  $\sigma'$  are well-formed states of a program  $P$  and  $P^{\#}$  is the instrumented version of  $P$ . Then the following statements are true.*

1.  $P$  can make a transition from  $\sigma$  to  $\sigma'$  iff  $P^{\#}$  can make a transition from  $(\sigma, I(\sigma))$  to  $(\sigma', I(\sigma'))$ .
2.  $P$  makes a transition from  $\sigma$  to either an ill-formed state or wrong iff  $P^{\#}$  makes a transition from  $(\sigma, I(\sigma))$  to wrong.

The proof of this theorem is performed by a case analysis over the various types of statements and is given in the full version of the paper [23].

## 5. First-order axiomatization of well-foundedness

Our checker verifies a program by checking the existence of a well-formed state satisfying a verification condition. The definition of a well-formed state (from Section 4.1) has six constraints. It is easy to write simple first-order axioms to model constraint 2–6. But the first constraint, which requires the state to be well-founded, involves a nontrivial relationship between the variables  $\mathbf{f}$ ,  $\mathbb{R}_{\mathbf{f}}$ ,  $\mathbb{B}_{\mathbf{f}}$  and  $\mathbb{H}_{\mathbf{f}}$  for any linking field  $\mathbf{f}$ . In fact, well-foundedness cannot be precisely expressed in first-order logic.

In this section, we first provide two first-order axioms (called base axioms) to capture the relationships between  $\mathbb{R}_{\mathbf{f}}$ ,  $\mathbb{B}_{\mathbf{f}}$  and  $\mathbf{f}$ . We then provide an induction principle (*IND-WF*) that enables us to derive additional first-order axioms from the base axioms. For convenience, we will use the notation  $\mathbf{f}(u)$  to denote  $u.\mathbf{f}$  in this section.

### 5.1 Base axioms

The following two fundamental axioms characterizes the relationship between the predicates  $\mathbb{R}_{\mathbf{f}}$ ,  $\mathbb{B}_{\mathbf{f}}$ ,  $\mathbb{H}_{\mathbf{f}}$  and  $\mathbf{f}$  in any state  $\Sigma$  of the program. In all these axioms, the cells  $u \in Cell$ ,  $v \in Cell$  are implicitly universally quantified out.

1. The first axiom specifies that  $\mathbb{R}_{\mathbf{f}}(u, v)$  is true if and only if either (i) there is a zero-length path from  $u$  to  $v$  (when  $u = v$ ), or (ii) there is a path of length one or more from  $u$  to  $v$  without encountering any cells from  $\mathbb{H}_{\mathbf{f}}$ .

$$\mathbb{R}_{\mathbf{f}}(u, v) \Leftrightarrow (u = v \vee (\mathbf{f}(u) \notin \mathbb{H}_{\mathbf{f}} \wedge \mathbb{R}_{\mathbf{f}}(\mathbf{f}(u), v))) \quad (\text{AX1})$$

2. The second axiom relates  $\mathbb{B}_{\mathbf{f}}(u)$  with  $\mathbb{B}_{\mathbf{f}}(\mathbf{f}(u))$  when  $\mathbf{f}(u) \notin \mathbb{H}_{\mathbf{f}}$ . If  $\mathbf{f}(u) \in H$ , then  $\mathbb{B}_{\mathbf{f}}(u) = \mathbf{f}(u)$ .

$$\mathbb{B}_{\mathbf{f}}(u) = (\mathbf{f}(u) \in \mathbb{H}_{\mathbf{f}}) ? \mathbf{f}(u) : \mathbb{B}_{\mathbf{f}}(\mathbf{f}(u)) \quad (\text{AX2})$$

The above axioms follow from the definition of the two predicates from Section 4.1.

We define a state  $\Sigma$  to be *finite* if the set  $\text{Alloc}$  is finite in  $\Sigma$ . Also, recall (from Section 4.1) that for any function  $f : Cell \rightarrow Cell$  and a set  $H \subseteq Cell$ ,  $R_H^f(u)$  defines the set of all cells

$u, f(u), \dots$  until the first cell from  $H$  is encountered, and  $B_H^f(u)$  denotes the identify of the first cell in  $H$  encountered.

The base axioms intuitively capture the definition of the fields  $\mathbb{R}_{\mathbf{f}}$  and  $\mathbb{B}_{\mathbf{f}}$ . The following theorem, which is a natural generalization of the theorem for acyclic case [24], serves to make this intuition precise.

**THEOREM 2.** *For any finite and well-formed state  $\Sigma$  satisfying axioms AX1 and AX2,  $\Sigma(\mathbb{R}_{\mathbf{f}}) = R_H^{\mathbf{f}}$ , and  $\Sigma(\mathbb{B}_{\mathbf{f}}) = B_H^{\mathbf{f}}$ , where  $f = \Sigma(\mathbf{f})$  and  $H = \Sigma(\mathbb{H}_{\mathbf{f}})$ .*

### 5.2 Induction principle

Even though the axioms AX1 and AX2 capture fundamental properties of  $R_H^{\mathbf{f}}$  and  $B_H^{\mathbf{f}}$ , they are not complete. As noted by Lev-Ami et al. [24], at least one of the limitation of the above axioms (and first-order logic in general) is that there is no complete axiomatization of “finiteness”.

A consequence of this limitation is that above axioms might imply other derived axioms in any finite state, but these additional axioms can’t be derived from the two axioms above solely by first order reasoning. These derived axioms are required to further constrain the set of states to give a meaningful assignments to  $\mathbb{R}_{\mathbf{f}}$ ,  $\mathbb{B}_{\mathbf{f}}$  and  $\mathbf{f}$ .

In this section, we present an induction principle (*IND-WF*) that can be used to derive other theorems from the base axioms AX1 and AX2.

**DEFINITION 4 (IND-WF).** *Consider a well-formed state  $\Sigma$  where  $f = \Sigma(\mathbf{f})$  and  $H = \Sigma(\mathbb{H}_{\mathbf{f}})$ , for the linking field  $\mathbf{f}$ . To show that a property  $P(u)$  holds for all  $u \in Cell$ , it is sufficient to establish two cases:*

1. *Base Case: Establishes the property for all cells  $u$  such that  $f(u) \in H$ :*

$$f(u) \in H \Rightarrow P(u) \quad (1)$$

2. *Induction Step: If  $P$  holds for a cell  $v$  and  $v \notin H$ , then establish that  $P$  holds for all the cells  $u$  such that  $f(u) = v$ :*

$$(P(v) \wedge f(u) = v \wedge v \notin H) \Rightarrow P(u) \quad (2)$$

We show that if the above induction principle establishes that a predicate  $P$  is true for all the cell  $u \in Cell$  in a state  $\Sigma$ , then  $\Sigma$  satisfies  $\forall u : P(u)$ .

**THEOREM 3.** *The induction principle IND-WF is sound.*

### 5.3 Derived axioms

In this section, we present a small set of useful first-order axioms that have been required in the various proof efforts we have undertaken. These axioms can be derived from the base axioms using the induction principle described above. It is a challenge to identify a “core” subset of axioms that is not only useful in practice, but also fairly intuitive to understand. We have identified a small set of such axioms that seem to be sufficient for the set of examples handled in this paper. Of course, we can’t make any claims whether this set will suffice for other programs too, since our experience is limited to the set of programs we have handled. But we believe that this set captures the interesting cases in dealing with most linked list programs.

1. **Transitivity:** The relation  $\mathbb{R}_{\mathbf{f}}$  enjoys the transitivity property:

$$(\mathbb{R}_{\mathbf{f}}(u, v) \wedge \mathbb{R}_{\mathbf{f}}(v, w)) \Rightarrow \mathbb{R}_{\mathbf{f}}(u, w) \quad (\text{TR})$$

2. **Antisymmetry:** This property is key to breaking the symmetry in a cycle.

$$(\mathbb{R}_{\mathbf{f}}(u, v) \wedge \mathbb{R}_{\mathbf{f}}(v, u)) \Rightarrow u = v \quad (\text{AS})$$

3. **Bounded distinctness:** For any heap cell  $u$ , if none of the cells  $\mathbf{f}(u), \mathbf{f}^2(u), \dots, \mathbf{f}^k(u)$  intersect with the set  $H_{\mathbf{f}}$ , then all the cells in this sequence are distinct from each other. Let  $F_k \doteq \{\mathbf{f}(u), \dots, \mathbf{f}^k(u)\}$ , and  $DISTINCT(S)$  denotes that all the members in a set  $S$  are pairwise distinct. We can derive a parameterized system of theorems  $DT_k$  for different finite values of  $k$ :

$$F_k \cap H = \{\} \Rightarrow DISTINCT(\{u\} \cup F_k) \quad (DT_k)$$

We have found that the case for  $k = 1$  to be useful in proving properties of many linked list programs (e.g. *reverse* of an acyclic linked list):

$$(\mathbf{f}(u) \notin H) \Rightarrow u \neq \mathbf{f}(u) \quad (DT_1)$$

These axioms can be proved easily from the base axioms AX1 and AX2, and the induction principle *IND-WF*. We illustrate this by working through the proof of  $DT_1$ :

*Proof:* We will substitute the expression  $(\mathbf{f}(u) \notin H) \Rightarrow u \neq \mathbf{f}(u)$ , in place of  $P(u)$  in Equation 1 and Equation 2. Then we derive contradiction from the negation of the formula.

- *Base case:* Since  $\mathbf{f}(u) \notin H$ , the formula for the base case is unsatisfiable.
- *Induction step:* Substituting the definition of  $P$  and rewriting  $v$  with  $\mathbf{f}(u)$  in Equation 2, we get:

$$\begin{aligned} &(\mathbf{f}(\mathbf{f}(u)) \in H \vee \mathbf{f}(u) \neq \mathbf{f}(\mathbf{f}(u))) \wedge \\ &\mathbf{f}(\mathbf{f}(u)) \notin H \wedge \mathbf{f}(u) \notin H \wedge u = \mathbf{f}(u) \end{aligned}$$

Rewriting the formula after replacing  $\mathbf{f}(u)$  with  $u$  (since  $u = \mathbf{f}(u)$ ), and removing  $\mathbf{f}(u) = \mathbf{f}(u)$ , we get

$$(\mathbf{f}(u) \in H \vee \mathbf{f}(u) \neq \mathbf{f}(u)) \wedge \mathbf{f}(u) \notin H$$

which is a contradiction.

□

## 5.4 Optimization

In addition to the instrumentation discussed in Section 4, we also introduce an `assume` statement immediately after each statement of the form  $\mathbf{x} = \mathbf{y}.\mathbf{f}$ :

$$\text{assume}(\mathbf{x} \notin H_{\mathbf{f}} \Rightarrow R_{\mathbf{f}}(\mathbf{x}) = R_{\mathbf{f}}(\mathbf{y}) \setminus \{\mathbf{y}\});$$

The constraint described by the `assume` statement is an instance of the following theorem (that can again be derived from AX1 and the induction principle):

$$\mathbf{f}(u) \notin H_{\mathbf{f}} \Rightarrow R_{\mathbf{f}}(\mathbf{f}(u)) = R_{\mathbf{f}}(u) \setminus \{u\} \quad (T1)$$

The addition of the constraint eagerly introduces an appropriate instance of Theorem T1 during the instrumentation phase. Since this axiom is not added explicitly to the set of derived axioms used in the proofs, the theorem prover is later prevented from searching for instances of this axiom.

## 6. Decision procedure for well-foundedness

Many automated theorem provers including Simplify [9] and UCLID [7] use a combination of heuristic quantifier instantiation and decision procedures for ground theories to reason about quantified formulas over certain theories. To take advantage of this approach, we explore decision procedures for quantifier-free formulas containing  $R_{\mathbf{f}}$  and  $B_{\mathbf{f}}$ . A decision procedure provides an alternative to the methodology presented in Section 5, where we provide a set of (incomplete) first-order axioms about transitive closure to reason about formulas containing  $R_{\mathbf{f}}$  and  $B_{\mathbf{f}}$ . We show that the complexity of the decision problem for the quantifier-free theory with  $R_{\mathbf{f}}$  and

$$\begin{aligned} l \in \text{Literal} & ::= \mathbf{x} \neq \mathbf{y} \mid \mathbf{f}(\mathbf{x}) = \mathbf{y} \mid \mathbf{x} = \text{null} \mid \\ & R_{\mathbf{f}}(\mathbf{x}, \mathbf{y}) \mid \neg R_{\mathbf{f}}(\mathbf{x}, \mathbf{y}) \mid B_{\mathbf{f}}(\mathbf{x}) = \mathbf{y} \mid \\ & \mathbf{x} \in H_{\mathbf{f}} \mid \mathbf{x} \notin H_{\mathbf{f}} \mid \\ & \mathbf{x} \in \text{Alloc} \mid \mathbf{x} \notin \text{Alloc} \\ \varphi \in \text{Formula} & ::= l \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \end{aligned}$$

Figure 5. Quantifier-free logic with  $R_{\mathbf{f}}$  and  $B_{\mathbf{f}}$ .

$B_{\mathbf{f}}$  is NP-complete. We also provide a naïve decision procedure for the fragment based on a small-model encoding.

Consider a quantifier-free formula  $\varphi \in \text{Formula}$  in the logic described in Figure 5. It is easy to see that one can reduce a formula containing nested applications of  $\mathbf{f}$  or  $B_{\mathbf{f}}$  (such as  $\mathbf{f}(\mathbf{f}(\mathbf{x}))$ ) to a formula in this logic by introducing additional *skolem constants*. For simplicity of presentation, let us take the liberty to include these additional constants in the set of program variables  $\text{Var}$ . Accordingly, a program state  $\sigma$  includes an evaluation for these additional symbols. For any state  $\sigma$  and a formula  $\varphi \in \text{Formula}$ , we can define the meaning of  $\varphi$  under  $\sigma$  ( $\sigma \models \varphi$ ) inductively over the structure of  $\varphi$  in the usual way.

Suppose  $L \subseteq \text{Literal}$  is a finite set of literals. We write  $\sigma \models L$  if  $\sigma \models \bigwedge l$  for all  $l \in L$ . The set  $L$  is *satisfiable* if there exists a well-formed state  $\sigma$  such that  $\sigma \models L$ . The decision problem *WFSAT* is the following: Is a given finite set  $L$  of literals satisfiable?

LEMMA 1. *The problem WFSAT is NP-hard.*

The proof of this lemma is by reduction from the 3-COLOR problem. In fact, the decision problem *WFSAT* remains NP-hard even when we restrict ourselves to literals of the form  $\mathbf{x} \neq \mathbf{y}$ ,  $\mathbf{f}(\mathbf{x}) = \mathbf{y}$ ,  $\mathbf{x} = \text{null}$  and  $R_{\mathbf{f}}(\mathbf{x}, \mathbf{y})$ .

A state  $\sigma$  is *finite* if  $\sigma(\text{Alloc})$  is a finite subset of  $\text{Cell}$ . If  $\sigma$  is finite, we define  $|\sigma| = |\sigma(\text{Alloc})|$ . We show that the logic described in Figure 5 above enjoys a *small-model* property. Given a finite set  $L$  of literals, let  $N\text{Var}(L)$  be the number of distinct variables that occur in the literals in  $L$ . Then we can prove the following:

LEMMA 2. *If a finite set of literals  $L$  is satisfiable, then there is a finite well-formed state  $\sigma$  such that  $\sigma \models L$  and  $|\sigma| \leq N\text{Var}(L) + 1$ .*

The following theorem follows easily from Lemmas 1 and 2.

THEOREM 4. *The problem WFSAT is NP-complete.*

The small-model property of Lemma 2 suggests a straightforward decision procedure for the logic. Given a formula  $\varphi$  in this logic, we can replace each occurrence of  $R_{\mathbf{f}}$  and  $B_{\mathbf{f}}$  with their definitions. For example, if  $k \doteq N\text{Var}(L) + 1$ , we can replace  $B_{\mathbf{f}}(\mathbf{x})$  with the formula:

$$\mathbf{f}(\mathbf{x}) \in H_{\mathbf{f}} \ ? \ \mathbf{f}(\mathbf{x}) : (\mathbf{f}^2(\mathbf{x}) \in H_{\mathbf{f}} \ ? \ \mathbf{f}^2(\mathbf{x}) : \dots \text{k times} \dots)$$

The resultant formula can be solved using a (combined) decision procedure for equality with uninterpreted functions (EUF) and linear arithmetic. In Section 7, we briefly describe our initial experience with the decision procedure.

## 7. Experiments

We have implemented a prototype tool to mechanically verify properties of programs containing linked lists. In this section, we briefly describe the components of the prototype and some preliminary results on a set of programs manipulating singly or doubly linked lists.

The source program is written in the Zing [1] programming language, an imperative Java-like language but without inheritance.

In addition to the programming language discussed in the paper, it can also support arrays of *Boolean*, *Integer* and *Cell*.

The user specifies a subset of fields in the program as linking fields. For each field  $f$ , the tool automatically instruments the Zing program with the auxiliary variables  $R_f$ ,  $B_f$ , and adds the updates to these variables. The instrumented program also contains the asserts for null-dereference and the asserts to ensure that the fields are well-founded. The instrumented program is then translated to a guarded transition system and is fed to the UCLID verification system [7]. Let us briefly describe the UCLID system and the features that are used in this paper.

## 7.1 UCLID

The input language of UCLID supports variables of type Booleans, integers, and functions (and predicates) over integers. The functions can have any finite arity. Each variable of type *Cell* in the source program is mapped to an integer variable. Each field  $f$  is mapped to a function variable from integers to integers or Booleans. Set-valued variables such as  $R_f$  (respectively  $H_f$ ) are mapped to a predicate variable of arity two (respectively one). Similarly,  $B_f$  is mapped to a function variable of arity one.

The updates to function and predicate variables are modeled using a  $\lambda$  notation used in Section 4 of this paper. The  $\lambda$  notation generalizes the theory of arrays, and allows us to modify an arbitrary number of entries in an array in a single step. This is convenient for expressing the updates of  $R_f$  and  $B_f$  variables.

The user can add the necessary preconditions, loop invariants and the postconditions to check in the UCLID file. Recall that some of the asserts are automatically generated by our translator and are part of the list of properties to check in the UCLID file. The file also includes the set of base axioms and derived axioms described in Section 5.

UCLID supports specification of first-order axioms and properties of the form  $\forall X : \phi(X)$ , where  $\phi$  is a quantifier-free expression over the state variables in the program. This is sufficient to express the axioms, preconditions, loop invariants, and the postconditions for all the examples that we have encountered in this work. Set operations are modeled using quantified expressions (e.g.  $S = S_1 \cup S_2$  gets translated to  $\forall x : S(x) \Leftrightarrow (S_1(x) \vee S_2(x))$ ).

The tool can be used in two different ways:

- **Proving verification conditions (VC):** Given the preconditions, loop invariants and the postconditions, the tool generates a VC that is checked using the theorem prover in UCLID. The theorem prover uses quantifier instantiation to eliminate quantifiers in the formula and then uses a Boolean Satisfiability (SAT) based decision procedure for the theories of uninterpreted functions and linear arithmetic. Failed VCs result in a concrete counterexample that is immensely useful for strengthening the loop invariants, adding more preconditions or (in rare cases) adding new axioms.
- **Synthesizing loop invariants using indexed predicate abstraction:** The tool can automatically construct universally quantified invariants using simple *indexed* predicates from a set  $\mathcal{P}$ . Each predicate  $p \in \mathcal{P}$  is a Boolean expression over the state variables and a set of *index* symbols  $\mathcal{X}$  of type integers (recall *Cell* are identified as integers in UCLID). The tool then constructs the strongest loop invariant of the form  $\forall \mathcal{X} : \phi(\mathcal{X})$ , where  $\phi(\mathcal{X})$  is a Boolean combination of the predicates in  $\mathcal{P}$  [21, 13].

For instance, given the set of predicates  $\mathcal{P} = \{u = \text{next}(v), v = \text{prev}(u), R_{\text{next}}(\text{hd}, u), R_{\text{next}}(\text{hd}, v)\}$ , with  $\mathcal{X} = \{u, v\}$ , the tool can compose the predicates to construct an invariant (say)

$$\forall u, v : \quad (R_{\text{next}}(\text{hd}, u) \wedge R_{\text{next}}(\text{hd}, v)) \Rightarrow \\ (v = \text{prev}(u) \Leftrightarrow u = \text{next}(v))$$

Example	VC Size	# Instants	Time Taken in UCLID (sec)
<code>cyclic_simple</code>	991 (45)	0	2.1
<code>reverse_a</code>	1896 (72)	0	9.5
<code>setunion_S1</code>	1134 (49)	0	3.9
<code>setunion</code>	2960 (104)	1	11.4
<code>insert</code>	5950(110)	2	37.5
<code>dlist_remove</code>	3362 (115)	1	30.1

**Figure 6.** Results of verifying linked list programs. “VC Size” denotes the number of nodes in the VC formula after instantiation; the numbers in the parenthesis denote the number of integer-valued terms in the formula. The column “# Instants” denotes the number of manual instantiations that had to be done.

The ability to generate complex quantified invariants from simple predicates often relieves the user from writing down well-formed loop invariants. Various heuristics are also provided to often infer most predicates [22].

In the next two subsections, we describe our initial experience with using the two features of UCLID to verify a set of programs manipulating linked lists.

## 7.2 Benchmarks

We have currently handled the following set of examples using the two options of UCLID mentioned above:

- `cyclic_simple`: This the example from Section 1 where all the elements of a cyclic list are initialized to 0.
- `reverse_a`: This example performs an in-place destructive update of an acyclic linked list pointed to by a variable  $l$ .
- `setunion`, `setunion_S1`: This is the example described in Figure 1. The example `setunion_S1` denotes the portion of the example before the destructive updates to combine the two lists into one. We have included both `setunion` and `setunion_S1` to demonstrate the increase in complexity in the presence of destructive updates.
- `insert`: This is the example of insertion into a sorted list described in Figure 2 in Section 2.
- `dlist_remove`: This example of a doubly linked list is also described in Section 2.

## 7.3 Proving verification conditions

In this section, we describe the set of examples for which we provided the loop invariant manually and used the theorem prover to prove the verification conditions.

Figure 6 describe the results of verifying these examples on a 2 GHz machine running Linux and 1GB memory. ZChaff [29] was used as the SAT solver inside UCLID. We have measured the complexity of the VC generated (VC Size) and the total time taken by the theorem prover to prove the formula. We have also included the number of cases where we had to manually instantiate an axiom.

Observe the difference between the verification of `setunion_S1` and `setunion`. In the latter case, adding the destructive updates requires one manual instantiation as well as increased CPU time. Also, the complexity of verifying `dlist_remove` comes from two sources even though the program is just four lines long: first, the precondition of the method is complex, and secondly, the set of axioms double and number terms to instantiate increases because of the presence of two fields (`next`, `prev`) in the program.

Example	# Predicates	# Iterations	Time Taken (sec)
cyclic_simple	15(1)	4	11.4
reverse_a	16(1)	6	85.45
set_union	24 (1)	5	79.79
insert	21(2)	9	1404.07
dlist_remove	-	-	30.1

**Figure 7.** Results of verifying linked list programs using indexed predicate abstraction. “# Predicates” denotes the number of predicates requires; the number in the parenthesis denotes the number of index variables in the predicates. “# Iterations” is the number of iterations of abstract reachability computation to compute the invariant. “Time taken” is the time taken to construct the invariant and prove the postconditions. The example `dlist_remove` did not require any loop invariant.

The need for manual instantiation of quantifiers in the axioms have come mainly from two cases so far:

- The quantifier instantiation heuristics fail to infer relevant terms to instantiate. It only happened once for the example `insert`, where we had to instantiate `AX2` with a concrete term.
- Since the axiom `TR` relates three bound variables  $u$ ,  $v$  and  $w$ , and the number of combinations to instantiate grows exponentially with the number bound variables resulting in a very large VC, we always use the axiom after instantiating the first argument  $u$  with one of the program variables. This was the source of one instantiation for each of `setunion`, `insert` and `dlist_remove` examples.

Both the above problems can be mitigated by better quantifier instantiation strategies, and we are working towards improving the heuristics in UCLID. We do not use Simplify to discharge the VCs because the concrete counterexample facility in UCLID offsets the value of more advanced instantiation heuristics present in Simplify.

We also implemented a simple decision procedure for the quantifier-free fragment (described in Section 6), that is based on the small-model property (Lemma 2) of the logic. However, this simple scheme did not yield an efficient decision procedure for this fragment. For the set of examples in the paper (where  $NVar(\psi)$  ranges between 10 and 30), expanding  $R_f$  and  $B_f$  resulted in a huge blowup. Apart from the simplest example `cyclic_simple`, the decision procedure did not terminate on the other examples within a 1000 seconds time limit.

#### 7.4 Invariant synthesis

We have also leveraged the predicate abstraction engine in UCLID to infer loop invariants for some of the examples. Our initial experience has been encouraging, and we report some preliminary results in this section. At present, we have managed to construct the loop invariants for the `cyclic_simple`, `reverse_a`, `set_union` and `insert` examples, given a set of indexed predicates.

Figure 7 illustrates the result of synthesizing loop invariants using indexed predicate abstraction for a subset of the examples. Currently the tool suffers from two bottlenecks that results in significant time to construct the invariants. This also explains the time taken to prove the verification conditions in the last section.

- UCLID does not maintain the control flow graph explicitly and encodes the program counter as a variable — the entire program is encoded as a single transition relation. This is because the tool was primarily built for analyzing distributed protocols and systems. This results in a large blowup in the formulas that the theorem prover or the predicate abstraction engine gets, and

```

/*@ requires null ∈ Hnext
    */
/*@ requires Bnext(l) == null

/*@ ensures Rnext(res) == old(Rnext)(l)

Cell reverse (Cell l) {
    Cell curr = l;
    Cell res = null;
    while (curr != null) {
        Cell tmp = curr.next;
        curr.next = res;
        res = curr;
        curr = tmp;
    }
    return res;
}

```

#### Set of Predicates

$$\mathcal{X} = \{u\}$$

$$\mathcal{P} = \{u = null, u = curr, u = res, u = old(l), R_{next}(curr, u), R_{next}(res, u), H_{next}(u), old(R_{next})(old(l), u), l = old(l), R_{next}(old(l), u), curr = null, B_{next}(u) = null\}$$

**Figure 8.** Reversing an acyclic list. The source program along with the set of predicates. We use `old(x)` to denote the value of  $x$  at the method entry.

consequently slows the analysis. We are currently working on incorporating explicit control flow into the tool.

- The quantifier instantiation engine generates a large number of (often redundant) terms to instantiate. Moreover, the number of combinations to instantiate grows exponentially with the number of index variables in  $\mathcal{X}$ . For instance, with two index variables, the number of combinations to instantiate went up to 81 for the `insert` example. In some cases, SAT often saw formulas with more than 300K clauses in them.

To mitigate these problems, we are exploring alternate quantifier instantiation and predicate abstraction strategies. Besides, one can often trade off the precision of the predicate abstraction to construct less precise loop invariants more efficiently. We plan to investigate if such loop invariants suffice to prove the properties of interest for these candidate programs.

Figure 8 describe the `reverse_a` example along with the set of predicates required to prove the postcondition. The example required a single index symbol (denoting a heap cell)  $u$  to construct the loop invariant. The tool requires 85.45 seconds to construct the loop invariant and prove the property. The set of predicates for this example was supplied manually. We are currently working on automating the process of predicate discovery.

## 8. Related work

There is a rich body of work in reasoning about programs that perform destructive updates of heap allocated data structures. Work in this area can be divided into the following often overlapping categories.

**First-order axiomatization of reachability.** The work in this category is closest to our work. Nelson [32] proposed the *ternary* reachability predicate  $u \xrightarrow{f} v$  to define that  $u$  can reach  $v$  through applications of  $f$  without encountering  $x$ , and provided a set of

first-order axioms to capture this predicate. Our axiomatization is based on two *binary* predicates. We believe that our predicates are more intuitive to a programmer and our axiomatization yields simpler correctness proofs. For example, Nelson’s proof of the set-union example from Section 2.1 required eight axioms whereas our proof requires only the transitivity axiom in addition to the two base axioms. Lev-Ami et al. [24] proposed another set of axioms for characterizing the reachability predicate. Their approach works only for acyclic lists and will not allow them to express loop invariants for programs manipulating cyclic lists. McPeak [27] proposes a methodology for writing specifications of data structures in first-order logic. Since he does not attempt to axiomatize reachability, his method, by itself, is not sufficient to express reachability or disjointness properties of linked lists. To express such facts, the programmer has to manually introduce and update ghost variables in a program specific way. His work, however, provides new heuristics for quantifier instantiation targeted towards reasoning for linked data structures. This work should complement our work which depends crucially on first-order theorem provers. Finally, our instrumentation for updating  $R_\ell$  and  $B_\ell$  is similar to the work by Dong and Su [11] on incrementally updating the transitive closure relation of a graph due to unit changes in it.

**Shape analysis.** Work in this category attempts to reason about the shape properties of the heap, such as acyclicity and sharing, in the presence of destructive updates to the heap. Ghiya and Hendren [14] propose the use of two Boolean matrices to record if two pointers are reachable from each other and if they share any heap cell. They provide conservative updates to these matrices for various statements in the program. Sagiv et al. [38, 39] use a 3-valued logic to represent abstractions of the heap graph. For improved precision, their approach requires *instrumentation predicates*, which usually use a reachability predicate as a building block. The updates for these predicates are either supplied by the user or constructed conservatively in some cases using finite-differencing [36]. The idea is implemented in the TVLA [25] system and has been used to infer loop invariants for various linked data structures. In contrast to this work, our approach depends on purely first-order reasoning. Recently, Hackett and Rugina [16] have proposed a method that uses points-to information to construct an abstraction of the heap. Their approach cannot encode relationships (e.g. equality) between program expressions and cannot describe doubly-linked lists. One advantage of our approach compared to all of the above is the ease with which we can combine reasoning about linked lists with reasoning about arithmetic and arrays.

**Decidable logics.** Various decidable logic fragments have been also proposed to express properties of linked data structures. Nelson and Yao [31] provide a polynomial decision procedure for the quantifier-free fragment with  $R_\ell(x, y)$  and  $\neg R_\ell(x, y)$ . Ranise and Zarba [35] provide a decidable logic with NP-complete decision problem for reasoning about linked lists. Our contribution in this paper has been to show the minimal extension to the logic in [31] to obtain NP-hardness. PALE [28] uses monadic second-order logic to express properties of lists, trees, and graphs. The work on *graph types* by Klarlund et al. [20] allows common shapes, such doubly-linked lists and threaded trees, to be concisely expressed. They also provide a decidable monadic second-order logic to check the well-formedness of graph type specifications. The logic  $L_r$  [5] was proposed to reason about reachability. However, the decidability results quickly break down in the presence of complex shapes and scalar values in the program.

**Predicate abstraction for discovering invariants.** Recently, predicate abstraction [15] has been extended to construct invariants for linked-list programs. Dams and Namjoshi [8] proposed the use of reachability predicate to construct an abstraction of the concrete system. They provide a heuristic for predicate discovery based

on constructing the weakest precondition of the reachability predicate. Manevich et al. [26] observe that the number of shared nodes in the heap consisting of singly-linked lists is statically bounded, and propose a family of predicates to exploit this observation. The idea has been implemented in TVLA. However, the method targets shape properties of lists and can’t verify properties regarding the contents of a list as we do for the `cyclic_simple` and `insert` examples. Balaban et al. [2] provide a decision procedure for restricted formulae involving the reachability predicate and provide a method to compute the abstraction. The restrictions on the reachability predicate does not allow them, for example, to express the loop invariant for the `cyclic_simple` program from Section 1.1. In comparison to all of the above, our ability to harness invariant inference methods for first-order logic (e.g. indexed predicate abstraction) provides us with appreciable automation for programs that manipulate linked structures, arrays, and other scalar values.

**Separation logic.** Separation logic [19, 37] is a promising idea for local reasoning of heap-manipulating programs. Separation logic naturally extends traditional Hoare-style reasoning to bear upon such programs. However, most of the work on separation logic focused on proving programs manually [33, 34]. Recently, Berdine et al. [6] have developed decision procedures for fragments of separation logic. However, we are not aware of any automated tools for program verification based on these decision procedures. Proofs of linked-list programs using separation logic use the reachability predicate; our axiomatization of reachability could potentially help with mechanizing such proofs. Alternatively, our work can benefit from separation logic specifications in future as we extend our work to the inter-procedural setting.

## 9. Conclusions

Programs such as kernels of operating systems and device drivers manipulate a variety of data structures, such as arrays, singly and doubly linked lists, and hashables. Current verification tools focus on control-dominated properties and are consequently unable to handle such programs in which the control flow interacts with the data in subtle ways. This paper is our first step towards the goal of building a scalable checker for verifying low-level data-rich systems software.

In this paper, we presented a novel method for verifying linked data structures based on a first-order axiomatization of reachability with respect to a set of head cells. This axiomatization is based on the idea of a well-founded heap and has the advantage that acyclic and cyclic lists are handled uniformly with equal ease. We have implemented our method in a tool and used it to verify the correctness of a variety of nontrivial programs manipulating both acyclic and cyclic singly-linked lists and doubly-linked lists.

There are several immediate directions for future work. We would like to verify more examples that use doubly-linked lists and arrays to evaluate the overhead of using our methodology and the adequacy of our list of derived first-order axioms. We have access to several programs from the Windows kernel that use rich data structures; we intend to evaluate our method on these programs. We would like to automate the inference of invariants more by developing heuristics for predicate discovery targeted towards linked lists. Finally, we would like to extend our work to deal with procedure calls and develop techniques to summarize procedures.

## Acknowledgements

We are grateful to Venkatesan Guruswami for his help in proving Lemma 1 and to Thomas Ball for his insightful comments on this paper.

## References

- [1] T. Andrews, S. Qadeer, S. K. Rajamani, J. Rehof, and Y. Xie. Zing: Exploiting program structure for model checking concurrent software. In *CONCUR 04: 15th International Conference on Concurrency Theory*, volume 3170 of *LNCS*, pages 1–15. Springer-Verlag, 2004.
- [2] I. Balaban, A. Pnueli, and L. D. Zuck. Shape analysis by predicate abstraction. In *VMCAI 05: Verification, Model checking, and Abstract Interpretation*, volume 3385 of *LNCS*, pages 164–180. Springer-Verlag, 2005.
- [3] T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *PLDI 01: Programming Language Design and Implementation*, pages 203–213, 2001.
- [4] C. Barrett and S. Berezin. CVC Lite: A new implementation of the cooperating validity checker. In *Computer Aided Verification (CAV '04)*, volume 3114 of *LNCS*, pages 515–518. Springer-Verlag, 2004.
- [5] M. Benedikt, Thomas W. Reps, and S. Sagiv. A decidable logic for describing linked data structures. In *ESOP 99: European Symposium on Programming*, volume 1576 of *LNCS*, pages 2–19. Springer-Verlag, 1999.
- [6] J. Berdine, C. Calcagno, and P. W. O'Hearn. A decidable fragment of separation logic. In *FSTTCS 04: Foundations of Software Technology and Theoretical Computer Science*, volume 3328 of *LNCS*, pages 97–109. Springer-Verlag, 2004.
- [7] R. E. Bryant, S. K. Lahiri, and S. A. Seshia. Modeling and verifying systems using a logic of counter arithmetic with lambda expressions and uninterpreted functions. In *Computer-Aided Verification (CAV '02)*, volume 2404 of *LNCS*, pages 78–92, July 2002.
- [8] D. Dams and K. S. Namjoshi. Shape analysis through predicate abstraction and model checking. In *VMCAI 03: Verification, Model checking, and Abstract Interpretation*, volume 2575 of *LNCS*, pages 310–324. Springer-Verlag, 2003.
- [9] D. L. Detlefs, G. Nelson, and J. B. Saxe. Simplify: A theorem prover for program checking. Technical report, HPL-2003-148, 2003.
- [10] E.W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
- [11] G. Dong and J. Su. Incremental and decremental evaluation of transitive closure by first-order queries. *Information and Computation*, 120(1):101–106, 1995.
- [12] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In *PLDI 02: Programming Language Design and Implementation*, pages 234–245. ACM Press, 2002.
- [13] C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *POPL 02: Principles of Programming Languages*, pages 191–202. ACM Press, 2002.
- [14] R. Ghiya and L. J. Hendren. Is it a tree, a DAG, or a cyclic graph? A shape analysis for heap-directed pointers in C. In *POPL 96: Principles of Programming Languages*, pages 1–15. ACM Press, 1996.
- [15] S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *Computer-Aided Verification (CAV '97)*, volume 1254 of *LNCS*, pages 72–83. Springer-Verlag, June 1997.
- [16] B. Hackett and R. Rugina. Region-based shape analysis with tracked locations. In *POPL 05: Principles of Programming Languages*, pages 310–323. ACM Press, 2005.
- [17] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *POPL 02: Principles of Programming Languages*, pages 58–70. ACM Press, 2002.
- [18] N. Immerman, A. M. Rabinovich, T. W. Reps, S. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *CSL 04: Computer Science Logic*, volume 3210 of *LNCS*, pages 160–174. Springer-Verlag, 2004.
- [19] S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *POPL 01: Principles of Programming Languages*, pages 14–26. ACM Press, 2001.
- [20] N. Klarlund and M. I. Schwartzbach. Graph types. In *POPL 93: Principles of Programming Languages*, pages 196–205. ACM Press, 1993.
- [21] S. K. Lahiri and R. E. Bryant. Constructing quantified invariants via predicate abstraction. In *VMCAI 04: Verification, Model Checking and Abstract Interpretation*, volume 2937 of *LNCS*, pages 267–281. Springer-Verlag, 2004.
- [22] S. K. Lahiri and R. E. Bryant. Indexed predicate discovery for unbounded system verification. In *Computer Aided Verification (CAV '04)*, volume 3114 of *LNCS*, pages 135–147. Springer-Verlag, 2004.
- [23] S. K. Lahiri and S. Qadeer. Verifying properties of well-founded linked lists. Technical Report MSR-TR-2005-97, Microsoft Research, 2005.
- [24] 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: Conference on Automated Deduction*, volume 3632 of *LNCS*, pages 99–115. Springer-Verlag, 2005.
- [25] T. Lev-Ami and S. Sagiv. TVLA: A system for implementing static analyses. In *SAS 00: Static Analysis Symposium*, volume 1824 of *LNCS*, pages 280–301. Springer-Verlag, 2000.
- [26] R. Manevich, E. Yahav, G. Ramalingam, and M. Sagiv. Predicate abstraction and canonical abstraction for singly-linked lists. In *VMCAI 05: Verification, Model Checking and Abstract Interpretation*, volume 3148 of *LNCS*, pages 181–198. Springer-Verlag, 2005.
- [27] S. McPeak and G. C. Necula. Data structure specifications via local equality axioms. In *Computer-Aided Verification (CAV '05)*, volume 3576 of *LNCS*, pages 476–490. Springer-Verlag, 2005.
- [28] A. Møller and M. I. Schwartzbach. The pointer assertion logic engine. In *PLDI 01: Programming Language Design and Implementation*, pages 221–231, 2001.
- [29] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *38th Design Automation Conference (DAC '01)*, 2001.
- [30] G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 2(1):245–257, 1979.
- [31] G. Nelson and F. F. Yao. Solving reachability constraints for linear lists, 1982. Unpublished manuscript.
- [32] G. Nelson. Verifying reachability invariants of linked structures. In *POPL 83: Principles of Programming Languages*, pages 38–47. ACM Press, 1983.
- [33] P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL 01: 15th International Workshop on Computer Science Logic*, volume 2142 of *LNCS*, pages 1–19. Springer-Verlag, 2001.
- [34] P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *POPL 04: Principles of Programming Languages*, pages 268–280. ACM Press, 2004.
- [35] S. Ranise and C. Zarba. A decidable logic for pointer programs manipulating linked lists, 2004. Unpublished manuscript.
- [36] T. Reps, M. Sagiv, and A. Loginov. Finite differencing of logical formulas for static analysis. In *ESOP 03: European Symposium on Programming*, volume 2618 of *LNCS*, pages 380–398. Springer-Verlag, 2003.
- [37] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS 02: Logic in Computer Science*, pages 55–74. IEEE Computer Society Press, 2002.
- [38] R. Wilhelm S. Sagiv, T. W. Reps. Solving shape-analysis problems in languages with destructive updating. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 20(1):1–50, 1998.
- [39] M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. In *POPL 99: Principles of Programming Languages*, pages 105–118. ACM Press, 1999.