

# A Logic and Decision Procedure for Predicate Abstraction of Heap-Manipulating Programs\*

Jesse Bingham and Zvonimir Rakamarić

Department of Computer Science, University of British Columbia, Canada  
jesse.d.bingham@intel.com  
zrakamar@cs.ubc.ca

**Abstract.** An important and ubiquitous class of programs are *heap-manipulating programs* (HMP), which manipulate unbounded linked data structures by following pointers and updating links. *Predicate abstraction* has proved to be an invaluable technique in the field of software model checking; this technique relies on an efficient decision procedure for the underlying logic. The expression and proof of many interesting HMP safety properties require *transitive closure* predicates; such predicates express that some node can be reached from another node by following a sequence of (zero or more) links in the data structure. Unfortunately, adding support for transitive closure often yields undecidability, so one must be careful in defining such a logic. Our primary contributions are the definition of a simple transitive closure logic for use in predicate abstraction of HMPs, and a decision procedure for this logic. Through several experimental examples, we demonstrate that our logic is expressive enough to prove interesting properties with predicate abstraction, and that our decision procedure provides us with both a time and space advantage over previous approaches.

## 1 Introduction

In recent years *software model checking* has emerged as a vibrant area of formal verification research. Much of the success of applying model checking to software has come from the use of *predicate abstraction* on the program source [16, 14, 3, 18]. In predicate abstraction, sets of states of the program and program transitions are over-approximated using a finite set of predicates over the program variables. These predicates (or boolean combinations thereof) typically express features of the program under verification such as its conditionals and relevant propositions about its variables. An integral ingredient in predicate abstraction is a decision procedure for the logic of the predicates. Since most approaches involve many queries to this decision procedure, performance is paramount.

An important class of programs are those we call *heap-manipulating programs* (HMPs), which are programs that access and modify linked data structures consisting of an unbounded number of uniform *heap nodes*. HMPs access the heap nodes through a finite number of pointers (that we call *node variables*) and following pointer fields between nodes. To apply predicate abstraction to HMPs and assert many interesting correctness properties, one must be able to express the concept of unbounded *reachability* (a.k.a. *transitive closure*) between nodes. This is done through a binary operator

---

\* The authors were supported by grants from the Natural Sciences and Research Council of Canada (NSERC). Jesse Bingham has moved to Intel Corporation, Hillsboro, Oregon, U.S.A.

that takes two node terms  $x$  and  $y$ , and asserts that the second can be reached from the first by following zero or more links; in our syntax this is written as  $f^*(x, y)$  ( $f$  is the name of the *link function*). For example,  $f^*(f(x), x)$  expresses that  $x$  is a node in a circular linked list.

Several papers have previously identified the importance of transitive closure for HMPs [30, 31, 5, 19, 2, 23]. Unfortunately, adding support for transitive closure to even relatively tame logics often yields undecidability [19]. Our first contribution is a fragment of the decidable logics that we show (through several nontrivial experiments) is still expressive enough to verify properties of interest for HMPs using predicate abstraction. Decidability of our logic follows from a small model theorem, akin to that of Benedikt et al. [5] and Balaban et al. [2], which states that if a set of predicates is satisfiable, then it is satisfiable by a heap structure with some bounded number of nodes. A naive decision procedure can thus enumerate all the (super-factorial but finite) number of structures of size up to this bound. We do not formally state or prove a small model theorem in this paper, rather, our second and most important contribution is an efficient decision procedure for our logic. We show that this procedure, though a worst case exponential time algorithm, solves the vast majority of queries sent to it during predicate abstraction very quickly. The result is an approach that can have large time and memory savings over decision procedures that enumerate all models, even when BDDs are used for this enumeration, as done by Balaban et al. [2].

The paper is organized as follows. Sect. 2 summarizes other work on verification of HMPs. Predicate abstraction and our verification framework (based on previous work), is outlined in Sect. 3. HMPs are introduced in Sect. 4. Sects. 5 and 6 respectively define our transitive closure logic and the decision procedure. We present experimental results in Sect. 7. Sect. 8 draws conclusions and discusses several important extensions to our logic and decision procedure that we believe are possible, but have been left as future work. Please note that our technical report [6] provides proofs of the theorems, additional details regarding the decision procedure, pseudocode for the example programs, and the sets of predicates needed for their verification.

## 2 Related Work

Balaban et al. [2] present an approach for shape analysis based on predicate abstraction that is similar to ours. The logic they use for describing properties of heap structures has slightly richer expressiveness than the logic we define in this paper.<sup>1</sup> The major difference between the two approaches is the way a program abstraction is computed. To compute the abstraction, they employ a small model theorem, and build BDDs representing all models up to the small model size. This is a bottleneck in both computation time and memory, since these BDDs tend to blow-up. The technique of Kesten and Pnueli [21] for establishing termination employed by Balaban et al. is likely compatible with our work also.

McPeak and Necula [28] specify heap data structures using *local equality axioms*, which constrain only a bounded fragment of the heap around some node. This enables them to describe a variety of shapes and reason about scalar values without abstracting

<sup>1</sup> Whereas our logic is unquantified, they allow restricted universal quantification.

them, while still preserving decidability. However, they can only approximate reachability between nodes (though *unreachability* is precise). When pointer disequalities are added, their decision procedure becomes incomplete. We handle both reachability and disequalities, but we can't describe such a variety of shapes. In addition, we compute an inductive invariant of a program automatically (given an appropriate set of predicates), while they require a user to provide loop invariants, which can be a significant burden.

The *Pointer Assertion Logic Engine* (PALE) [29] specifies heap structures using graph types [22], which are tree-shaped data structures augmented with extra pointers. The authors show that many common heap structures can be defined that way, some of which we cannot express. PALE relies on a decision procedure with non-elementary complexity, so, there are programs that cannot be verified in practice. Furthermore, loop invariants must be provided by the user.

The *Three Valued Logic Analyzer* (TVLA) [32, 25] extends conventional abstract interpretation with a third “uncertain” logic value, and builds so-called *3-valued logical structures* that abstract the reachable states at each program point (a.k.a. *canonical abstraction*). The abstract semantics of program statements are defined by *abstract transformers*, which can be generated by TVLA or user-defined if necessary. We cannot handle all heap structures that TVLA can, however, the abstract invariant we compute is always the most precise w.r.t. the given set of predicates. TVLA does not make such a guarantee, although some work has been done to make TLVA more precise [33]. TVLA is also employed by Manevich et al. [27], who observe that the number of shared nodes in linked lists is bounded and present a novel definition of “uninterrupted list segments”. This is used to define predicate and canonical abstractions of potentially circular singly linked lists, and enables them to verify some HMPs that we are not able to verify, though their properties tend to be simpler than ours (see Sect. 7).

Lahiri and Qadeer [23] define two new predicates to express reachability of heap nodes in linked lists. To prove properties of HMPs, they use first-order axioms over those predicates. The given set of axioms is incomplete, and they provide an induction principle that is used to derive additional axioms when necessary. Because of the purely first-order axiomatization, they are able to harness the power of available automated theorem provers; they use UCLID [8] as the underlying inference engine.

Dams and Namjoshi [11] propose an approach based on predicate abstraction and model checking. They abstract a program by iteratively calculating weakest preconditions of shape predicates, and are able to handle second-order shape properties such as reachability, cyclicity, and sharing. The algorithm doesn't use a decision procedure, and as a consequence, new predicates can be generated in every iteration. Hence, the algorithm often has to be manually provided with “approximation hints” to converge.

## 3 Verification Approach

### 3.1 Predicate Abstraction

Our approach to verifying heap programs is based on *predicate abstraction* [16], which is an instance of *abstract interpretation* [10]. In the framework of abstract interpretation, a *concrete system* (in our case an HMP) is verified by constructing a finite-state over-approximation of the concrete system called the *abstract system*. Let  $\mathcal{C}$  (the *concrete*

states) be the set of states of the concrete system. Predicate abstraction employs a finite set of predicates  $\phi_1, \dots, \phi_k$  in some logic that are assertions about concrete states. Corresponding to the predicates respectively are the *abstract boolean variables*  $b_1, \dots, b_k$ . The set of *abstract states*  $\mathcal{A}$  will be the set of assignments to the abstract boolean variables. The *abstraction function*  $\alpha : \mathcal{C} \rightarrow \mathcal{A}$  is defined such that  $\alpha(c)(b_i) = \text{true}$  if and only if  $c \models \phi_i$ . A set of concrete states  $C$  is then abstracted by

$$\alpha(C) = \{\alpha(c) \mid c \in C\}$$

Note that since  $\mathcal{A}$  is finite,  $\alpha(C)$  is always finite as well. In contrast,  $\mathcal{C}$  is often infinite; in our case the infinitude of concrete states arises from the unboundedness of the heap in HMPs.

Let  $R \subseteq \mathcal{C}$  be the set of concrete states that are reachable in the concrete system. We wish to verify that a property expressed as a state assertion  $\psi$  over the concrete states holds for all members of  $R$ , i.e. that the implication  $R \rightarrow \psi$  holds. Predicate abstraction is used to solve this problem by computing a set  $R^\alpha \subseteq \mathcal{A}$  such that  $\alpha(R) \subseteq R^\alpha$ . Verification succeeds if one can prove that  $R^\alpha \rightarrow \psi$ . A key difference in the various approaches to predicate abstraction is how  $R^\alpha$  is computed [16, 14, 12, 15, 2, 11]. This typically involves numerous queries to a decision procedure for the underlying logic and there are tradeoffs between how accurately  $R^\alpha$  approximates  $\alpha(R)$  and the number and complexity of these queries.

$R^\alpha$  is usually computed as a fixpoint of some approximation of the *abstract post image operator*  $\text{post} : 2^{\mathcal{A}} \rightarrow 2^{\mathcal{A}}$ , defined as follows. Given a set of abstract states  $A$ , let

$$\text{post}(A) = \{\alpha(c') \mid \exists c, c' \in \mathcal{C}. (c, c') \in T \wedge \alpha(c) \in A\}$$

where  $T$  is the transition relation of the concrete system.  $\text{post}(A)$  is thus the set of abstract states representing concrete states that are concrete successors of those states represented by  $A$ .

Since predicate abstraction is an incomplete approach, if it fails to verify the property, this can either happen because the concrete systems actually violates the property, or because of the loss of information inherent in the abstraction. Finding the “right” set of predicates for verification to go through can be tricky business. Many works have addressed this issue of *predicate discovery* [13, 4, 18, 11], which falls under the more general umbrella of *abstraction refinement* [9]. As in recent papers on this topic [2, 23], in our current framework, predicates are added by manual inspection of counterexample behaviors; applying automatic predicate discovery techniques is an important area of future work.

### 3.2 Computing post

Our tool computes  $\text{post}$  precisely; the algorithm can be viewed as an improvement over the following naive algorithm. Since  $\text{post}$  distributes over disjunction,<sup>2</sup> computing  $\text{post}(A)$  is reducible to computing  $\text{post}(\rho)$  for each cube  $\rho$  in some disjunctive normal form decomposition of  $A$ . Here, *cube* means a partial boolean assignment to

<sup>2</sup> Meaning that  $\text{post}(A_1 \vee A_2) = \text{post}(A_1) \vee \text{post}(A_2)$ .

the abstract variables, and represents all abstract states that agree on this subset of the abstract variables.<sup>3</sup> By using a BDD [7] to represent  $A$ , we can easily obtain such a decomposition. The naive algorithm cycles through all  $2^k$  abstract states  $a$ , and checks if  $a \in \text{post}(\rho)$ ;  $\text{post}(\rho)$  is then the BDD representing the disjunction of all such  $a$ . Each check of  $a \in \text{post}(\rho)$  involves a call to the decision procedure to determine if the following formula is satisfiable:

$$\gamma(\rho) \wedge \text{wp}(\gamma(a)) \quad (1)$$

where  $\gamma$  is the *concretization function*, and  $\text{wp}$  is the *weakest precondition* operator [17]. Intuitively,  $\gamma$  maps a cube to a logic formula that denotes the set of concrete states represented by the cube. Formally, for a cube  $\mu$  let  $P(\mu)$  (resp.  $N(\mu)$ ) denote the set  $\{i \mid \mu(b_i) = \text{true}\}$  (resp.  $\{i \mid \mu(b_i) = \text{false}\}$ ). Then define

$$\gamma(\mu) = \bigwedge_{i \in P(\mu)} \phi_i \wedge \bigwedge_{i \in N(\mu)} \neg \phi_i$$

The weakest precondition operator  $\text{wp}$  is a syntactic transformation on logic formulas that depends on the program statement under consideration [17]. For example, for an assignment statement  $x := e$ , where  $x$  is a variable and  $e$  is some expression,  $\text{wp}(\pi)$  is constructed by syntactically replacing all occurrences of  $x$  with  $e$  in the formula  $\pi$ .<sup>4</sup> Our approach applies  $\text{wp}$  at the granularity of individual program statements when performing predicate abstraction.

Das et al.’s computation of  $\text{post}$  that we employ uses several straightforward optimizations over this naive algorithm [14]. First, if (1) contains a syntactic contradiction, meaning the existence of a predicate and its negation, then clearly the formula is not satisfiable. In such circumstances there is no need to call the decision procedure. When computing  $\text{post}(\rho)$ , our implementation initially computes a BDD  $C$  representing the set of all  $a$  that won’t yield such a contradiction. Second, rather than enumerating all  $a \in C$ , we do recursive case-splitting on the abstract variables, which allows for pruning of large portions of  $C$ . For example, let  $\mu$  be the cube that assigns true to  $b_1$  and leaves all other variables unconstrained. Then if  $\gamma(\rho) \wedge \text{wp}(\gamma(\mu))$  is unsatisfiable, then so too is  $\gamma(\rho) \wedge \text{wp}(\gamma(a))$  for *any* abstract state  $a$  that has  $b_1$  equal to true. Hence, our algorithm would only explore those abstract states having  $b_1$  false.

## 4 Heap-Manipulating Programs

In our framework, the *heap* consists of an unbounded number of *nodes*. HMPs allow for node variables (pointers), data fields for nodes, a link field  $f$  for nodes, and all other variables are modelled (or encoded as) booleans.

In lieu of a formal presentation of HMPs, we give an example called ND-INSERT in Fig. 1 that captures most of the interesting features. This program takes a node *head* and a node *item*, and inserts *item* into the linked list pointed to by *head* at a position

<sup>3</sup> A *partial boolean assignment* maps each variable  $b_i$  to an element of  $\{\text{true}, \text{false}, \text{undef}\}$ .

<sup>4</sup> This only works under the assumption that  $x$  cannot be aliased.

```

1: procedure ND-INSERT(head, item)
2:   assume  $\neg f^*(head, item) \wedge f^*(head, nil) \wedge \neg head = nil \wedge f(item) = nil \wedge p = head$ 
3:   while true do
4:     if  $ND \vee f(p) = nil$  then
5:        $f(item) := f(p)$ ;
6:        $f(p) := item$ ;
7:       break
8:     else
9:        $p := f(p)$ ;
10:    end if
11:  end while
12:  assert  $f^*(head, item) \wedge f^*(head, nil)$ 
13: end procedure

```

**Fig. 1.** A program that nondeterministically inserts a node *item* into the list pointed to by *head*. Here *ND* is a boolean value that is nondeterministically true or false.

selected nondeterministically. *head* is assumed to be non-nil and to point to an acyclic linked list that does not contain *item*. These assumptions are formalized by the **assume** statement on line 2 of the program. In the **assume** statement, and also in the **assert** statement, the subformulas of the form  $f^*(x, y)$  express that node *y* is reachable from node *x* by following a sequence of *f* links of any length; we will formally define these predicates in Sect. 5. The fact that nil is reachable from *head* enforces the acyclicity assumption.<sup>5</sup>

The body of ND-INSERT is straightforward; a pointer *p* walks the list, and *item* is inserted at some point. The loop breaks once the insertion has occurred. The expression *ND* represents a nondeterministic boolean value. *item* is inserted when either  $ND = \text{true}$ , or the end of the list is reached (detected by the disjunct  $f(p) = \text{nil}$  on line 4). The specification is expressed by the **assert** statement on line 12, and indicates that whenever line 12 is reached, *head* must point to an acyclic list that contains *item*.

The verification problem we wish to solve can be stated as follows: given an HMP, determine whether it is the case that all executions that satisfy all **assume** statements also satisfy all **assert** statements. Since the number of nodes in the heap is unbounded, HMPs are generally infinite state, thus one cannot directly apply finite-state model checking to this problem without using abstraction.

## 5 A Simple Transitive Closure Logic

Our logic assumes finite sets of *node* variables *V*, *boolean* variables *B*, *data function* variables *D*, and a single *link function* symbol *f*. The *term*, *atom*, and *literal* syntactic entities are given in Fig. 2. Literals of the form  $x = y$ ,  $\neg x = y$ ,  $f^*(x, y)$ , and  $\neg f^*(x, y)$  (where *x* and *y* are terms) are called *equality*, *inequality*, *reachability*, and *unreachability* literals, respectively. Literals of the form  $d(x)$  or  $\neg d(x)$ , where  $d \in D$ , are called *data* literals, while those of the form *b* or  $\neg b$  are called simply *boolean variable* literals.

<sup>5</sup> In our logical framework, nil is modelled simply as a node having  $f(\text{nil}) = \text{nil}$ .

$$\begin{aligned}
v &\in V \\
d &\in D \\
b &\in B \\
\text{term} &::= v \mid f(\text{term}) \\
\text{atom} &::= f^*(\text{term}, \text{term}) \mid \text{term} = \text{term} \mid d(\text{term}) \mid b \\
\text{literal} &::= \text{atom} \mid \neg \text{atom}
\end{aligned}$$

**Fig. 2.** The syntax of our simple transitive closure logic

The structures over which the semantics of our logic is defined are called *heap structures*. A heap structure  $H = (N, \Theta)$  involves a finite set of *nodes*  $N$  and a function  $\Theta$  that interprets each symbol  $\sigma$  in  $V \cup B \cup D \cup \{f\}$  such that

$$\begin{aligned}
\Theta(\sigma) &\in N && \text{if } \sigma \in V \\
\Theta(\sigma) &\in \{\text{true}, \text{false}\} && \text{if } \sigma \in B \\
\Theta(\sigma) &\in N \rightarrow \{\text{true}, \text{false}\} && \text{if } \sigma \in D \\
\Theta(\sigma) &\in N \rightarrow N && \text{if } \sigma = f
\end{aligned}$$

Thus  $\Theta$  interprets each node variable as a node, each boolean variable as a boolean value, each data function variable as a function that maps nodes to booleans, and the link function  $f$  is interpreted as a mapping from nodes to nodes. Heap structures naturally model a linked data structure of nodes, each node having a single pointer to another node and some finite number of boolean-valued fields. The *size* of  $H$  is defined to be  $|N|$ . The variables of  $V$  model program variables that point to nodes in the data structure, while the variables of  $B$  model program variables of boolean type. Clearly, program variables or node fields of any finite enumerated type can be encoded using the booleans accommodated by our logic.

We extend  $\Theta$  to  $\Theta^e$ , which interprets any term or atom in the obvious way, formally defined here. The interpretation of a term  $\tau$  is defined inductively by:

$$\Theta^e(\tau) = \begin{cases} \Theta(\tau) & \text{if } \tau \in V \\ \Theta(f)(\Theta^e(\tau')) & \text{if } \tau \text{ has the form } f(\tau') \text{ for some term } \tau' \end{cases}$$

$\Theta^e$  interprets atoms as boolean values. An equality atom  $\tau_1 = \tau_2$  is interpreted as true by  $\Theta^e$  iff  $\Theta^e(\tau_1) = \Theta^e(\tau_2)$ . A data atom is interpreted by defining  $\Theta^e(d(\tau)) = \Theta(d)(\Theta^e(\tau))$ . A reachability atom  $f^*(\tau_1, \tau_2)$  is interpreted as true iff there exists some  $n \geq 0$  such that  $\Theta(f)^n(\Theta^e(\tau_1)) = \Theta^e(\tau_2)$ .<sup>6</sup> Finally, a literal that is not an atom is of the form  $\neg\phi$  where  $\phi$  is an atom, and we simply define  $\Theta^e(\neg\phi) = \neg\Theta^e(\phi)$ .

Sticking to the usual notation, given a heap structure  $H = (N, \Theta)$  and a literal  $\phi$ , we write  $H \models \phi$  iff  $\Theta^e(\phi) = \text{true}$ . For a set of literals  $\Phi$ , we write  $H \models \Phi$  iff  $H \models \phi$  for all  $\phi \in \Phi$ .

## 6 Decision Procedure

The decision problem we aim to solve with our decision procedure is this: given a finite set of literals  $\Phi$ , does there exist a heap structure  $H$  such that  $H \models \Phi$ ? If there is such an

<sup>6</sup> Here, function exponentiation represents iterative application: for a function  $g$  and an element  $x$  in its domain,  $g^0(x) = x$ , and  $g^n(x) = g(g^{n-1}(x))$  for all  $n \geq 1$ .

$$\begin{array}{c}
 \frac{}{v=v}\text{IDENT} \qquad \frac{}{f^*(v,v)}\text{REFLEX} \qquad \frac{f(x)=y}{f^*(x,y)}\text{TRANS1} \\
 \\
 \frac{f^*(x,y) \quad f^*(y,z)}{f^*(x,z)}\text{TRANS2} \qquad \frac{f(x)=y \quad f^*(x,z)}{x=z \quad f^*(y,z)}\text{FUNC} \\
 \\
 \frac{f(x_1)=x_2 \quad f(x_2)=x_3 \quad \cdots \quad f(x_k)=x_1 \quad f^*(x_1,y)}{y=x_1 \quad y=x_2 \quad \cdots \quad y=x_k}\text{CYCLE}_k \\
 \\
 \frac{f^*(x,y) \quad f^*(y,x) \quad f^*(x,z)}{x=y \quad f^*(z,x)}\text{SCC} \qquad \frac{f^*(x,y) \quad f^*(x,z)}{f^*(y,z) \quad f^*(z,y)}\text{TOTAL} \\
 \\
 \frac{f(x)=z \quad f(y)=z \quad f^*(x,y) \quad f^*(y,x)}{x=y}\text{SHARE}
 \end{array}$$

**Fig. 3.** The set of inference rules. Here  $x, y,$  and  $z$  range over (not necessarily distinct) terms. In the rules IDENT and REFLEX,  $v$  is restricted to be a variable that is already mentioned; this restriction prevents either of these rules from introducing new terms.  $\text{CYCLE}_k$  actually defines a separate rule for each  $k \geq 1$ .

$H$ , then we say that  $\Phi$  is *satisfiable*, otherwise  $\Phi$  is *unsatisfiable*. Clearly, any algorithm for this problem can be used to decide the satisfiability of a conjunction-of-literals (1) by simply taking  $\Phi$  to be the set of its conjuncts.

Decidability of this problem follows from a small model theorem enjoyed by our logic, akin to other transitive closure logics [2, 5]. Our small model theorem states that  $\Phi$  is satisfiable if and only if there exists  $H$  of size at most  $n$  such that  $H \models \Phi$ , where  $n$  is the number of distinct terms mentioned in  $\Phi$ . Hence, a decision procedure can simply enumerate the finite set of such  $H$ , and for each one check if  $H \models \Phi$ . However, since the number of such heap structures is at least  $n^n$ , this approach is impractical. Employing BDDs [7] to represent the set of heap structures that satisfy  $\Phi$  [2] is also memory-intensive; building a BDD for the literal  $f^*(x,y)$  over just 8 nodes cannot be done in 2 GB of memory. This stems from the fact that such a BDD must represent the multitude of different paths that could exist between the nodes  $\Theta^e(x)$  and  $\Theta^e(y)$ .

Our approach has relatively small memory requirements, and is based on a set of inference rules (IRs) with the property that  $\Phi$  is satisfiable if and only if their exhaustive application does not introduce a contradiction. Here *contradiction* means the inference of both an atom  $\phi$  and its negation  $\neg\phi$ . The IRs are presented in Fig. 3. For an IR  $r$ , the *antecedents* of  $r$  are the literals appearing above the line, while the *consequents* are those appearing below the line. We say that an IR  $r$  is *applicable* (to  $\Phi$ ) if there are terms appearing in  $\Phi$  such that when these terms are substituted for the term placeholders of  $r$  (i.e.  $x, y, z, x_1,$  etc.), *all* of  $r$ 's antecedents appear in  $\Phi$ , and *none* of  $r$ 's consequents appear in  $\Phi$ .

We now explain each IR of Fig. 3. IDENT states that each node variable is equal to itself, while REFLEX enforces that any node variable is reachable from itself. TRANS1 states that the transitive closure  $f^*$  must extend the function  $f$ . TRANS2 simply enforces that  $f^*$  is transitive. FUNC asserts that if  $f(x)=y$  and  $z$  is reachable from  $x$ , then  $z$  must also be reachable from  $y$ , unless  $x=z$ . If there is a cycle of length  $k \geq 1$  in  $f$ , then it

follows that any node  $y$  reachable from a node on the cycle must be on the cycle as well; this is formalized by  $\text{CYCLE}_k$ . Similar to  $\text{FUNC}$  is  $\text{SCC}$ , which states that if  $x$  and  $y$  are distinct and mutually reachable from each other, and  $z$  is reachable from  $x$ , then  $x$  is reachable from  $z$  (since  $x$  must lie on a cycle of  $f$ ).  $\text{TOTAL}$  requires that if  $y$  and  $z$  are both reachable from another node  $x$ , then there must exist some reachability relationship between  $y$  and  $z$ . The fact that in a cycle of  $f$ , no two distinct nodes  $x$  and  $y$  can have  $f(x) = f(y)$  is captured by  $\text{SHARE}$ . Given the preceding intuition, it is easy to prove the following.

**Theorem 1.** *The inference rules of Fig. 3 are sound.*

Theorem 1 tells us that if iterative application of the IRs yields a contradiction, then we can conclude that the original set of literals is unsatisfiable. Conversely, we have proven our IRs to be complete with respect to sets of literals in a certain normal form, and Theorem 2 below states that it is sufficient to restrict attention to such sets. Let  $\text{Vars}(\Phi)$  denote the subset of the node variables  $V$  appearing in  $\Phi$ .

**Definition 1 (normal).** *A set of literals  $\Phi$  is said to be normal if*

1. *For each  $v_i \in \text{Vars}(\Phi)$ , there exists*
  - (a) *at most one equality literal of the form  $f(v_i) = v_j$ , where  $v_j \in \text{Vars}(\Phi)$ , and*
  - (b) *the literal  $v_i = v_i$ .**All equality literals of  $\Phi$  are required to be of one of the forms (a) or (b).*
2. *All inequality literals are of the form  $\neg v_i = v_j$ , where  $v_i, v_j \in \text{Vars}(\Phi)$ .*
3. *All reachability literals are of the form  $f^*(v_i, v_j)$ , where  $v_i, v_j \in \text{Vars}(\Phi)$ .*
4. *All unreachability literals are of the form  $\neg f^*(v_i, v_j)$ , where  $v_i, v_j \in \text{Vars}(\Phi)$ .*
5. *There exist no data or boolean variable literals in  $\Phi$ .*

**Theorem 2.** *There exists a polynomial-time algorithm that transforms any set  $\Phi$  into a normal set  $\Phi'$  such that  $\Phi'$  is satisfiable if and only if  $\Phi$  is satisfiable.*

Thanks to Theorem 2, our decision procedure can without loss of generality assume that  $\Phi$  is normal. Let us call a set of literals  $\Phi$  *consistent* if it does not contain a contradiction, and call  $\Phi$  *closed* if none of the IRs of Fig. 3 are applicable. The following completeness result is the crux of our decision procedure.

**Theorem 3.** *If  $\Phi$  is consistent, closed, and normal, then  $\Phi$  is satisfiable.*

The proof of Theorem 3 is quite technical, and involves reasoning about the dependencies between digraphs of partial functions and the digraphs of their transitive closures. For details, please see our technical report [6].

Viewed from a high level, our decision procedure first applies the transformation of Theorem 2, and then repeatedly searches for an applicable IR, applies it (i.e. adds a consequent to the set), and recurses. The recursion is necessary for those IRs that *branch*, i.e. have multiple consequents. If the procedure ever infers a contradiction, it backtracks to the last branching IR with an unexplored consequent, or returns *unsatisfiable* if there is no such IR. If the procedure reaches a point where there are no applicable IRs and no contradictions, then the inferred set of literals is consistent, closed, and normal. Hence, by Theorem 3, it may correctly return *satisfiable*. Our technical report [6] provides a more formal presentation of the decision procedure. We note that our decision procedure is guaranteed to terminate because none of the IRs introduce new terms.

## 6.1 An Extension

In order to handle program assignments that mutate the links in the heap, i.e. modify  $f$ , we must extend our logic and decision procedure to support simultaneous reference to  $f$  and  $f'$ , which respectively model the link function before and after the assignment. Such an assignment has the general form  $f(\tau_1) := \tau_2$ , where  $\tau_1$  and  $\tau_2$  are arbitrary terms. Lines 5 and 6 of the HMP of Fig. 1 are examples of such assignments. The semantic relationship between  $f$  and  $f'$  can be expressed using the well-known update operator:<sup>7</sup>

$$\Theta^e(f') = \text{update}(\Theta^e(f), \Theta^e(\tau_1), \Theta^e(\tau_2)) \quad (2)$$

Rather than support update as an interpreted second order function symbol in the logic, we add inference rules that implicitly enforce the constraint (2). For each of the eight IRs of Fig. 3 that mention  $f$ , we add an analogous IR with  $f$  replaced with  $f'$ ; these enforce analogous constraints between  $f'^*$ ,  $f'$ , and  $=$  as are enforced by the unmodified IRs of Fig. 3 between  $f^*$ ,  $f$ , and  $=$ . Furthermore, to enforce the constraint (2), the seven IRs of Fig. 4 are also included. The IRs introduce a fresh variable  $w$  that is forced to be equal to  $f(\tau_1)$ . This allows us to state that  $\Theta^e(f) = \text{update}(\Theta^e(f'), \Theta^e(\tau_1), \Theta^e(w))$ , and hence the symmetry between the IRs UPDFUNC1 and UPDFUNC2, between UPDTRANS1 and UPDTRANS2, and between UPDTRANS3 and UPDTRANS4. Note that these IRs can introduce new terms, however, given a normal set of literals, the number of new terms is bounded. This implies that the extended decision procedure always terminates.

$$\frac{}{f'(\tau_1) = \tau_2} \text{UPDATE}$$

$$f(\tau_1) = w$$

$$\frac{f(x) = y}{x = \tau_1 \quad y = w} f'(x) = y \text{ UPDFUNC1} \qquad \frac{f'(x) = y}{x = \tau_1 \quad y = \tau_2} f(x) = y \text{ UPDFUNC2}$$

$$\frac{f^*(x, y)}{f'^*(x, \tau_1) \quad f'^*(w, y)} f'^*(x, y) \text{ UPDTRANS1} \qquad \frac{f'^*(x, y)}{f'^*(x, \tau_1) \quad f'^*(\tau_2, y)} f'^*(x, y) \text{ UPDTRANS2}$$

$$\frac{f^*(x, \tau_1) \quad f'^*(x, y)}{f'^*(x, y) \quad f'^*(\tau_1, y)} \text{UPDTRANS3} \qquad \frac{f'^*(x, \tau_1) \quad f'^*(x, y)}{f'^*(x, y) \quad f'^*(\tau_1, y)} \text{UPDTRANS4}$$

**Fig. 4.** The update inference rules, which are used to extend our logic to support a second function symbol  $f'$ , with the implicit constraint  $f' = \text{update}(f, \tau_1, \tau_2)$ , where  $\tau_1$  and  $\tau_2$  are fixed but arbitrary terms, and  $w$  is a fresh variable used to capture  $f(\tau_1)$ . Note that the rule UPDATE can introduce literals that violate normalcy (Def. 1) in the case that  $\tau_1$  or  $\tau_2$  are not variables. However, this can be remedied by the addition of a new variable and equality literal for each sub-term of  $\tau_1$  and  $\tau_2$ .

<sup>7</sup> If  $g$  is a function,  $a$  is an element in  $g$ 's domain, and  $b$  is an element in  $g$ 's codomain, then  $\text{update}(g, a, b)$  is defined to be the function  $\lambda x. (\text{if } x = a \text{ then } b \text{ else } g(x))$ .

**Theorem 4.** *The inference rules of Fig. 4 are sound.*

The proof of this theorem is provided in our technical report [6]. We have yet to flesh out the details of a proof of a conjecture analogous to Theorem 3 stating that this extended set of IRs is complete. However, we have empirical support for this conjecture: in conducting our experiments of Sect. 7, we never found any property violations caused by the extended decision procedure erroneously concluding that a set of literals was satisfiable. Of course, not having such a theorem *does not* compromise the soundness of verification by predicate abstraction.

## 7 Experiments

We have tested our tool on a number of HMP examples and summarized the results in Table 1. We ran the experiments on a Pentium 4 2.6 GHz machine. The safety properties we checked (when applicable) at the end of the HMP are:

- *no leaks* (NL) – all nodes reachable from the head of the list at the beginning of the program are also reachable at the end of the program.
- *insertion* (IN) – a distinguished node that is to be inserted into a list is actually reachable from the head of the list, i.e. the insertion “worked”.
- *acyclic* (AC) – the final list is acyclic, i.e. nil is reachable from the head of the list.
- *cyclic* (CY) – list is a singly linked circular list, i.e. the head of the list is reachable from its successor.
- *sorted* (SO) – list is a sorted linked list, i.e. each node’s data field is less than or equal to its successor’s.
- *remove elements* (RE) – for examples that remove node(s), this states that the node(s) was (were) actually removed. For the program REMOVE-ELEMENTS, RE also asserts that the data field of all removed elements is false.

Often, the properties one is interested in verifying for HMPs involve universal quantification over the heap nodes. For example, to assert the property NL, we must express that for all nodes  $t$ , if  $t$  is reachable from *head* initially, then  $t$  is also reachable from *head* (or some other node) at the end of the program. Since our logic doesn’t support quantification, we use the trick of introducing a Skolem constant  $t$  [15, 2] to represent a universally quantified variable. Here,  $t$  is a new node variable that is initially assumed to satisfy the antecedent of our property, and is otherwise unmodified by the program. For the example program of Fig. 1, we can express NL by conjoining  $\neg t = \text{nil} \wedge f^*(\text{head}, t)$  to the **assume** statement on line 2, and conjoining  $f^*(\text{head}, t)$  to the assertion on line 12. Since  $t$  can be any non-nil node reachable from *head*, if the assertion is never violated, we have proven NL.

Our example programs are the following:

LIST-REVERSE – a classical HMP example that performs in-place reversal of a linked list.

LIST-ADD – a linked list is traversed, and the end of the list is reached. Then, a node is added to the end of the list.

ND-INSERT – pseudocode for this example is given in Fig. 1.

**Table 1.** Results of verifying HMPs. “property” specifies the verified property; “CFG edges” denotes the number of edges in the control-flow graph of the program; “preds” is the number of predicates required for verification; “time” is the average execution time over five runs to prove the properties; “DP calls” is the number of decision procedure queries. The largest memory usage for all these examples was 125 MB.

program	property	CFG edges	preds	time (sec)	DP calls
LIST-REVERSE	NL	6	8	1.1	173
LIST-ADD	NL $\wedge$ AC $\wedge$ IN	7	8	0.8	66
ND-INSERT	NL $\wedge$ AC $\wedge$ IN	5	13	7.9	258
ND-REMOVE	NL $\wedge$ AC $\wedge$ RE	5	12	19.3	377
ZIP	NL $\wedge$ AC	20	22	280.7	9185
SORTED-ZIP	NL $\wedge$ SO $\wedge$ IN	28	22	249.9	13760
SORTED-INSERT	NL $\wedge$ AC $\wedge$ SO	10	20	217.6	8476
BUBBLE-SORT	NL $\wedge$ AC	21	24	204.6	5930
BUBBLE-SORT	NL $\wedge$ AC $\wedge$ SO	21	27	441.0	6324
REMOVE-ELEMENTS	NL $\wedge$ CY $\wedge$ RE	15	17	1263.7	26721

ND-REMOVE – similar to ND-INSERT, except that instead of inserting a node, a node is nondeterministically chosen and removed from the list.

ZIP – zips two linked lists, shuffling the elements of both list into one. Then, the tail of the longer list is appended to the resulting list. This example is taken from a paper by Jensen et al. [20].

SORTED-ZIP – joins the elements of two sorted lists into one, also sorted. Here the data elements are simply booleans, so “sorted” means that all nodes with false fields come before nodes with true fields.

SORTED-INSERT – inserts a node into a sorted linked list so that sortedness is preserved. This is a modification of the example from a technical report by Lahiri and Qadeer [23].<sup>8</sup>

BUBBLE-SORT – The bubble sort example sorts elements of a linked list using the bubble sort algorithm. It is taken from a paper by Balaban et al. [2]. The data fields on which we sort are again booleans.

REMOVE-ELEMENTS – removes from a circular list all elements whose data field is false.

Our technical report [6] provides pseudocode and lists the required predicates for these examples.

As Table 1 shows, we were successful in verifying interesting properties of many examples in reasonable amounts of time. Of special note is our verification of sortedness for BUBBLE-SORT. This example is from Balaban et al. [2]; because of the BDD blow-up inherent in their decision procedure, their tool spaced out for the small model bound necessary for sound verification [1]. In contrast, our trading of space for time appears to be quite advantageous here.

<sup>8</sup> To simplify things, they require that the input list starts with a dummy element whose data field value has to be less than all possible values of that data field. We don’t have such requirements in our example, which makes it slightly more complicated.

Running time of TVLA on the bubble sort example is somewhat faster than ours, although they are using a slower machine [24]. The recent experimental results of Manevich et al. [27] are significantly faster, in spite of the fact they were executed on a slower machine. For most of their examples, however, they only verify the simple property of no null dereferences (they also verify cyclicity for two examples). We are verifying more complicated properties, for instance SO. Very recently, Loginov et al. [26] have used TVLA to fully automatically verify the bubblesort example.

For the two examples in common with Lahiri and Qadeer [23],<sup>9</sup> LIST-REVERSE and SORTED-INSERT, we are significantly faster at verifying the same properties, with respective speed-ups of  $75\times$  and  $6\times$ . It should be noted, however, that we used a slightly faster machine, and also that for SORTED-INSERT, our data fields are merely booleans, while theirs are the full integers.

## 8 Future Work and Conclusions

Despite the fact that this work is in its early stages, our experiments demonstrate its effectiveness for verification of heap-manipulating programs. There are many directions for future research, which are outlined here.

We have identified the following issues related to the expressiveness of the simple transitive closure logic presented in this paper:

- This paper only supports a single link function  $f$ , yet clearly many heap-manipulating programs involve multiple link fields.
- We have found that even minimal support for universally quantified variables (as in the logic of Balaban et al. [2]) would allow expression of common heap structure attributes. For example, the current logic cannot assert that two terms  $x$  and  $y$  point to disjoint linked lists; a single universally quantified variable would allow for this property (see Nelson [30–page 22]). We found that capturing disjointedness is necessary for verifying that LIST-REVERSE always produces an acyclic list; hence we were unable to verify this property.
- Another situation that cannot be characterized relates to term ordering in circularly linked lists. Suppose  $x$ ,  $y$ , and  $z$  are terms in such a list; we would like to express that  $y$  does or does not “come between”  $x$  and  $z$  in the list. Nelson [31] and Manevich et al. [27] have previously recognized the importance of such properties.

We believe that our decision procedure can be enhanced to handle each of these three cases. A final expressiveness deficiency, that we see no immediate solution to, is the expression of more involved heap structure properties, in particular trees. Though our logic cannot capture “ $x$  points to a tree”, we believe that it is possible that an extension could be used to verify simple properties of programs that manipulate trees, for example that there are no memory leaks.

We also plan on investigating how existing techniques for predicate discovery and more advanced predicate abstraction algorithms mesh with our decision procedure. Our approach appears to be very promising, despite the fact that we have yet to harness the recent innovations in these areas.

<sup>9</sup> We were unable to run our tool on four of Lahiri and Qadeer’s [23] examples because we have yet to implement support for data field mutations.

## Acknowledgement

We acknowledge our mutual supervisor Alan J. Hu for his support during this project and Ittai Balaban, Shuvendu Lahiri, and Shaz Qadeer for answering our questions; we also thank Shaz for suggesting this research problem to us.

## References

1. I. Balaban, 2005. personal correspondence.
2. I. Balaban, A. Pnueli, and L. Zuck. Shape analysis by predicate abstraction. In *Conf. on Verification, Model Checking and Abstract Interpretation (VMCAI)*, 2005.
3. T. Ball, R. Majumdar, T. D. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *Conf. on Programming Language Design and Implementation (PLDI)*, pages 203–213, 2001.
4. T. Ball, A. Podelski, and S.K. Rajamani. Relative completeness of abstraction refinement for software model checking. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2002.
5. M. Benedikt, T. Reps, and M. Sagiv. A decidable logic for describing linked data structures. In *European Symposium on Programming (ESOP)*, 1999.
6. J. Bingham and Z. Rakamarić. A logic and decision procedure for predicate abstraction of heap-manipulating programs, 2005. UBC Dept. Comp. Sci. Tech Report TR-2005-19, <http://www.cs.ubc.ca/cgi-bin/tr/2005/TR-2005-19>.
7. R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. on Computers*, C-35(8):677–691, August 1986.
8. 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 *Conf. on Computer Aided Verification (CAV)*, pages 78 – 92, 2002.
9. E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *Conf. on Computer Aided Verification (CAV)*, pages 154–169, 2000.
10. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Symp. on Principles of Programming Languages (POPL)*, pages 238–252, 1977.
11. D. Dams and K. S. Namjoshi. Shape analysis through predicate abstraction and model checking. In *Conf. on Verification, Model Checking and Abstract Interpretation (VMCAI)*, pages 310–323, 2003.
12. S. Das and D. L. Dill. Successive approximation of abstract transition relations,. In *IEEE Symp. on Logic in Computer Science (LICS)*, 2001.
13. S. Das and D. L. Dill. Counter-example based predicate discovery in predicate abstraction. In *Formal Methods in Computer-Aided Design (FMCAD)*, 2002.
14. S. Das, D. L. Dill, and S. Park. Experience with predicate abstraction. In *Conf. on Computer Aided Verification (CAV)*, 1999.
15. C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *Symp. on Principles of Programming Languages (POPL)*, pages 191–202, 2002.
16. S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *Conf. on Computer Aided Verification (CAV)*, 1997.
17. D. Gries. *The Science of Programming*. Springer-Verlag, New York, 1981.
18. T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *Symp. on Principles of Programming Languages (POPL)*, pages 58–70, 2002.

19. N. Immerman, A. Rabinovich, T. Reps, M. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive closure logics. In *Workshop on Computer Science Logic (CSL)*, pages 160–174, 2004.
20. J. L. Jensen, M. E. Jørgensen, N. Klarlund, and M. I. Schwartzbach. Automatic verification of pointer programs using monadic second-order logic. In *Conf. on Programming Language Design and Implementation (PLDI)*, pages 226–236, 1997.
21. Y. Kesten and A. Pnueli. Verification by augmented finitary abstraction. *Information and Computation*, 163(1):203–243, 2000.
22. N. Klarlund and M. I. Schwartzbach. Graph types. In *Symp. on Principles of Programming Languages (POPL)*, pages 196–205, 1993.
23. S. K. Lahiri and S. Qadeer. Verifying properties of well-founded linked lists, 2005. Microsoft Research Tech Report MSR-TR-2005-97.
24. T. Lev-Ami, T. Reps, M. Sagiv, and R. Wilhelm. Putting static analysis to work for verification: A case study. In *Intl. Symp. on Software Testing and Analysis*, pages 26–38, 2000.
25. T. Lev-Ami and M. Sagiv. TVLA: A system for implementing static analyses. In *Static Analysis Symposium (SAS'00)*, pages 280–301, 2000.
26. A. Loginov, T. W. Reps, and S. Sagiv. Abstraction refinement via inductive learning. In *Conf. on Computer Aided Verification (CAV)*, pages 519–533, 2005.
27. R. Manevich, E. Yahav, G. Ramalingam, and M. Sagiv. Predicate abstraction and canonical abstraction for singly-linked lists. In *Conf. on Verification, Model Checking and Abstract Interpretation (VMCAI)*, pages 181–198, 2005.
28. S. McPeak and G. C. Necula. Data structure specifications via local equality axioms. In *Conf. on Computer Aided Verification (CAV)*, pages 476–490, 2005.
29. A. Møller and M. I. Schwartzbach. The pointer assertion logic engine. In *Conf. on Programming Language Design and Implementation (PLDI)*, pages 221–231, 2001.
30. G. Nelson. *Techniques for program verification*. PhD thesis, Stanford University, 1979.
31. G. Nelson. Verifying reachability invariants of linked structures. In *Symp. on Principles of Programming Languages (POPL)*, pages 38–47, 1983.
32. M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *ACM Trans. on Programming Languages and Systems*, 24(3):217–298, 2002.
33. G. Yorsh, T. Reps, and M. Sagiv. Symbolically computing most-precise abstract operations for shape analysis. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 530–545, 2004.