

Automated Verification of Shape, Size and Bag Properties

Wei-Ngan Chin^{1,2} Cristina David¹ Huu Hai Nguyen² Shengchao Qin³

¹ Department of Computer Science, National University of Singapore

² Computer Science Programme, Singapore-MIT Alliance

³ Department of Computer Science, Durham University

{chinwn,davidcri,nguyenh2}@comp.nus.edu.sg shengchao.qin@durham.ac.uk

Abstract

In recent years, separation logic has emerged as a contender for formal reasoning of heap-manipulating imperative programs. Recent works have focused on specialised provers that are mostly based on fixed sets of predicates. To improve expressivity, we have proposed a prover that can automatically handle user-defined predicates. These shape predicates allow programmers to describe a wide range of data structures with their associated size properties. In the current work, we shall enhance this prover by providing support for a new type of constraints, namely bag (multi-set) constraints. With this extension, we can capture the reachable nodes (or values) inside a heap predicate as a bag constraint. Consequently, we are able to prove properties about the actual values stored inside a data structure.

1 Introduction

Separation logic supports reasoning about shared mutable data structures. Therefore, specification of memory operations with advanced pointer usage is enabled. While the foundations of separation logic have been laid in seminal papers by Reynolds [23] and Isthiaq and O’Hearn [11], new automated reasoning tools based on separation logic, such as [2, 9], are beginning to appear. Several recent works, such as [1, 8], have developed specialised solvers that work for a fixed set of predicates (e.g. `lseg`). The level of expressivity for an automated verification system was raised by our previous work [20] where we support automated reasoning of user-defined predicates. The main goal of our current paper is to enhance the prover proposed in [20] by providing support for bag constraints. We make the following technical contributions towards this overall goal :

- We provide a *shape predicate specification* mechanism that can capture a wide range of data structures together with size and bag properties, such as various

height-balanced trees, priority heap, sorted list, etc. Moreover, we provide a mechanism to soundly approximate each shape predicate by a heap-independent *invariant* which plays an important role in entailment checking (Secs 2 and 4.1).

- We improve the expressiveness of our automatic verification tool by allowing it to capture both the size properties and the bag of values from each shape predicate. The numeric properties capture sophisticated data structure invariants, such as orderedness (for sorted list/trees) and also balanced height properties (for AVL-trees). The bag constraints enable expressing reachability properties, such as the preservation of the elements inside a list after sorting.
- We design a new procedure to check entailment of separation heap constraints. This procedure uses *unfold/fold* reasoning to deal with shape definitions. While the unfold/fold mechanism is not new, we have identified sufficient conditions for soundness and termination of the procedure in the presence of recursive user-defined shape predicates. (Secs 3.1, 4 and 5)
- We have implemented a prototype verification system with the above features and have also proven both its soundness and termination (Secs 6 and 7).

2 User-Definable Shape Predicates

Separation logic [23, 11] extends Hoare logic to support reasoning about shared mutable data structures. It adds two more connectives to classical logic : separating conjunction $*$, and separating implication $-*$. $h_1 * h_2$ asserts that two heaps described by h_1 and h_2 are domain-disjoint. $h_1 -* h_2$ asserts that if the current heap is extended with a disjoint heap described by h_1 , then h_2 holds in the extended heap. In this paper we use only separating conjunction.

We propose an intuitive mechanism based on inductive predicates (or relations) to allow user specification of

shapely data structures with size and reachability properties. Our shape specification is based on separation logic with support for disjunctive heap states. Furthermore, each shape predicate may have pointer, integer or bag parameters to capture relevant properties of data structures. We use the following data node declarations for the examples in the paper. They are recursive data declarations with different number of fields.

```
data node { int val; node next }
data node2 { int val; node2 prev; node2 next }
data node3 { int val; node3 left; node3 right;
             node3 parent }
```

We use $p::c\langle v^* \rangle$ to denote two things in our system. When c is a data name, $p::c\langle v^* \rangle$ stands for singleton heap $p \mapsto [(f:v)^*]$ where f^* are fields of data declaration c . When c is a predicate name, $p::c\langle v^* \rangle$ stands for the formula $c(p, v^*)$. The reason we distinguish the first parameter from the rest is that each predicate has an implicit parameter `self` as the first one. Effectively, `self` is a “root” pointer to the specified data structure that guides data traversal and facilitates the definition of *well-founded* predicates (Sec 3.1). As an example, a singly linked list with length n is described by :

$$ll\langle n \rangle \equiv (\text{self} = \text{null} \wedge n = 0) \vee (\exists i, m, q. \text{self}::\text{node}\langle i, q \rangle * q::ll\langle m \rangle \wedge n = m + 1) \quad \text{inv } n \geq 0$$

Note that the parameter n captures a *derived* value. The above definition asserts that an `ll` list can be empty (the base case `self=null`) or consists of a head data node (specified by `self::node⟨i, q⟩`) and a separate tail data structure which is also an `ll` list (`q::ll⟨m⟩`). The `*` connector ensures that the head node and the tail reside in disjoint heaps. We also specify a default invariant $n \geq 0$ that holds for all `ll` lists. Our predicate uses existential quantifiers for local values and pointers, such as i, m, q .

A more complex shape, doubly linked-list with length n , is described by :

$$dll\langle p, n \rangle \equiv (\text{self} = \text{null} \wedge n = 0) \vee (\text{self}::\text{node2}\langle _, p, q \rangle * q::dll\langle \text{self}, n-1 \rangle) \quad \text{inv } n \geq 0$$

The `dll` shape predicate has a parameter p that represents the `prev` field of the first node of the doubly linked-list. It captures a chain of nodes that are to be traversed via the `next` field starting from the current node `self`. The nodes accessible via the `prev` field of the `self` node are not part of the `dll` list. This example also highlights some shortcuts we may use to make shape specification shorter. We use underscore `_` to denote an anonymous variable. Non-parameter variables in the RHS of the shape definition, such as q , are considered existentially quantified. Furthermore, terms may be directly written as arguments of shape predicate or data node.

User-definable shape predicates provide us with more flexibility than some recent automated reasoning systems [1, 3] that are designed to work with only a small set of fixed predicates. Furthermore, our shape predicates can describe not only the *shape* of data structures, but also their *size* and *bag* properties. This capability enables many applications, especially to support data structures with sophisticated invariants. For example, we may define a non-empty sorted list as below. The predicate also tracks the length, the minimum and maximum elements of the list.

$$\begin{aligned} \text{sortl}\langle n, \text{min}, \text{max} \rangle &\equiv \\ &(\text{self}::\text{node}\langle \text{min}, \text{null} \rangle \wedge \text{min} = \text{max} \wedge n = 1) \\ &\vee (\text{self}::\text{node}\langle \text{min}, q \rangle * q::\text{sortl}\langle n-1, k, \text{max} \rangle \wedge \text{min} \leq k) \\ &\quad \text{inv } \text{min} \leq \text{max} \wedge n \geq 1 \end{aligned}$$

The constraint $\text{min} \leq k$ guarantees that sortedness property is adhered between any two adjacent nodes in the list. We may now specify (and then verify) the following insertion sort algorithm :

```
node insert(node x, node vn) where
  x::sortl⟨n, mi, ma⟩ * vn::node⟨v, -⟩ *→
  res::sortl⟨n+1, min(v, mi), max(v, ma)⟩
  { if (vn.val ≤ x.val) then { vn.next:=x; vn }
    else if (x.next=null) then
      { x.next:=vn; vn.next:=null; x }
    else { x.next:=insert(x.next, vn); x }}
```

```
node insertion_sort(node y) where
  y::ll⟨n⟩ ∧ n > 0 *→ res::sortl⟨n, -, -⟩
  { if (y.next=null) then y
    else { y.next:=insertion_sort(y.next);
          insert(y.next, y) }}
```

We use the notation $\Phi_{pr} * \rightarrow \Phi_{po}$ to capture a precondition Φ_{pr} and a postcondition Φ_{po} of a method. Note that we use an expression-oriented language where the last sub-expression (e.g. e_2 from $e_1; e_2$) denotes the result of an expression. A special identifier `res` is used in the postcondition to denote the result of a method. The postcondition of `insertion_sort` shows that the output list is sorted and has the same number of nodes as the input list.

2.1 Bag of Values/Addresses

The earlier specification of sorting captures neither the in-situ reuse of memory cells nor the fact that all the elements of the list are preserved by sorting. The reason is that the shape predicate captures only pointers and numbers but does not capture the set of reachable nodes in a heap predicate. A possible solution to this problem is to extend our specification mechanism to capture either a set or a bag of values. For generality and simplicity, we propose to use only the bag (or multi-set) notation that permits

duplicates, though set notation could also be supported. In the rest of the paper, we will use the following bag operators: bag union \sqcup , bag intersection \sqcap , bag subsumption \sqsubseteq , and bag cardinality $|B|$. The shape specifications from the previous section are revised as follows:

$$\begin{aligned} \text{ll2}\langle n, B \rangle &\equiv (\text{self}=\text{null} \wedge n=0 \wedge B=\{\}) \\ &\vee (\text{self}::\text{node}\langle _, q \rangle * \text{q}::\text{ll2}\langle n-1, B_1 \rangle \wedge B=B_1 \sqcup \{\text{self}\}) \\ \text{inv } n \geq 0 \wedge |B| &= n \end{aligned}$$

$$\begin{aligned} \text{sortl2}\langle B, \text{mi}, \text{ma} \rangle &\equiv \\ &(\text{self}::\text{node}\langle \text{mi}, \text{null} \rangle \wedge \text{mi}=\text{ma} \wedge B=\{\text{self}\}) \\ &\vee (\text{self}::\text{node}\langle \text{mi}, \text{q} \rangle * \text{q}::\text{sortl2}\langle B_1, k, \text{ma} \rangle \\ &\quad \wedge B=B_1 \sqcup \{\text{self}\} \wedge \text{mi} \leq k) \\ \text{inv } \text{mi} \leq \text{ma} \wedge B &\neq \{\} \end{aligned}$$

Each predicate of the form $\text{ll2}\langle n, B \rangle$ or $\text{sortl2}\langle B, \text{mi}, \text{ma} \rangle$ now captures a bag of addresses B for all the data nodes of its data structure (or heap predicate). With this extension, we can provide a more comprehensive specification for in-situ sorting, as follows :

$$\begin{aligned} \text{node insert}(\text{node } x, \text{node } v_n) &\text{ where} \\ x::\text{sortl2}\langle B, \text{mi}, \text{ma} \rangle * v_n::\text{node}\langle v, _ \rangle &* \rightarrow \\ \text{res}::\text{sortl2}\langle B \sqcup \{v_n\}, \min(v, \text{mi}), \max(v, \text{ma}) \rangle & \\ \{ \dots \} & \end{aligned}$$

$$\begin{aligned} \text{node insertion_sort}(\text{node } y) &\text{ where} \\ y::\text{ll2}\langle n, B \rangle \wedge B \neq \{\} &* \rightarrow \text{res}::\text{sortl2}\langle B, _, _ \rangle \\ \{ \dots \} & \end{aligned}$$

We stress that the bag mechanism used to capture the reachable nodes in a shape predicate is quite general. Instead of heap addresses, we may also revise our views to capture a bag of reachable values, as well as the length. For example:

$$\begin{aligned} \text{ll3}\langle n, B \rangle &\equiv (\text{self}=\text{null} \wedge n=0 \wedge B=\{\}) \vee \\ &(\text{self}::\text{node}\langle a, q \rangle * \text{q}::\text{ll3}\langle n-1, B_1 \rangle \wedge B=B_1 \sqcup \{a\}) \\ \text{inv } n \geq 0 \wedge |B| &= n \end{aligned}$$

Capturing a bag of values allows us to reason about the collection of values in a data structure, and permits relevant properties to be automatically verified. Both universal and existential properties over bags may be expressed and proven, when implemented with an appropriate constraint

solver, as highlighted below :

```
data pair{node v1; node v2}
pair partition(node x, int p) where
x::ll3⟨n, A⟩ *→ res::pair⟨r1, r2⟩ * r1::ll3⟨n1, B1⟩ *
r2::ll3⟨n2, B2⟩ ∧ A=B1⊔B2 ∧ n=n1 + n2 ∧ (∀a∈B1·a≤p)
∧ (∀a∈B2·a>p)
{ if (x=null) then new pair(null, null)
  else { pair t; t:=partition(x.next, p);
        if (x.val≤p) then { x.next:=t.v1; t.v1:=x }
        else { x.next:=t.v2; t.v2:=x }
        t } }
```

```
bool allPos(node x) where
x::ll3⟨n, B⟩ *→ x::ll3⟨n, B⟩ ∧ ((∀a∈B·a≥0) ∧ res
∨ (∃a∈B·a<0) ∧ ¬res)
{ if (x=null) then true
  else if (x.val<0) then false
  else allPos(x.next) }
```

The first example returns a pair of lists that have been partitioned from a single input list according to an integer pivot. This partition function and its pre/post specification can be used to prove the total correctness of quicksort algorithm. The second example uses existentially and universally quantified formulae to determine if at least one negative number is present in an input list, or not.

3 Automated Verification

In this section, we first introduce a core object-based imperative language and then propose a set of forward verification rules to systematically check that preconditions are satisfied at call sites, and that the declared postcondition is successfully verified (assuming the precondition) for each method definition.

3.1 Language

We provide a simple imperative language in Figure 1. Our language is strongly typed and we assume programs and constraints are well-typed. The language supports data type declaration via *datat*, and shape predicate definition via *spred*. For each shape definition *spred*, we also declare a heap-independent invariant π_0 over the parameters $\{\text{self}, v^*\}$ that holds for each instance of the predicate.

Each method *meth* and *while* loop is declared with pre- and post-conditions of the form $\Phi_{pr} * \rightarrow \Phi_{po}$. For simplicity, we assume that variable names declared in each method are all distinct and that parameters are passed by-value. Primed notation is used to denote the latest value of variables and may appear in the postcondition of loops. For example, a simple loop with pre/post conditions is shown below:

P	$::= tdecl^* meth^*$	$tdecl ::= datat \mid spread$
$datat$	$::= data\ c \{ field^* \}$	$field ::= tv \quad t ::= c \mid \tau$
τ	$::= int \mid bool \mid float \mid void$	
$spread$	$::= c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0$	
$meth$	$::= t\ mn\ ((tv)^*) \text{ where } \bigwedge (\Phi_{pr} \rightsquigarrow \Phi_{po})^* \{e\}$	
e	$::= null \mid k^\tau \mid v \mid v.f \mid v:=e \mid v_1.f:=v_2 \mid \text{new } c\langle v^* \rangle$ $\mid e_1; e_2 \mid tv; e \mid mn\langle v^* \rangle \mid \text{if } v \text{ then } e_1 \text{ else } e_2$ $\mid \text{while } v \text{ where } \bigwedge (\Phi_{pr} \rightsquigarrow \Phi_{po})^* \text{ do } e$	
Φ	$::= \bigvee (\exists v^*. \kappa \wedge \pi)^*$	$\pi ::= \gamma \wedge \phi$
γ	$::= v_1=v_2 \mid v=null \mid v_1 \neq v_2 \mid v \neq null \mid \gamma_1 \wedge \gamma_2$	
κ	$::= emp \mid v::c\langle v^* \rangle \mid \kappa_1 * \kappa_2$	
Δ	$::= \Phi \mid \Delta_1 \vee \Delta_2 \mid \Delta \wedge \pi \mid \Delta_1 * \Delta_2 \mid \exists v. \Delta$	
ϕ	$::= \varphi \mid b \mid a \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg \phi \mid \exists v. \phi \mid \forall v. \phi$	
b	$::= true \mid false \mid v \mid b_1 = b_2$	
a	$::= s_1 = s_2 \mid s_1 \leq s_2$	
s	$::= k^{\text{int}} \mid v \mid k^{\text{int}} \times s \mid s_1 + s_2 \mid -s \mid \max(s_1, s_2)$ $\mid \min(s_1, s_2) \mid B $	
φ	$::= v \in B \mid B_1 = B_2 \mid B_1 \sqsubset B_2 \mid \forall v \in B. \phi \mid \exists v \in B. \phi$	
B	$::= B_1 \sqcup B_2 \mid B_1 \cap B_2 \mid B_1 - B_2 \mid \{ \} \mid \{ v \}$	

Figure 1. A Core Imperative Language

```
while x<0 where true  $\rightsquigarrow$  (x>0  $\wedge$  x'=x)  $\vee$  (x≤0  $\wedge$  x'=0)
do { x:=x+1 }
```

Here x and x' denote the values of variable x at the entry and exit of the loop, respectively.

The separation constraints we use are in a disjunctive normal form Φ . Each disjunct consists of a $*$ -separated heap constraint κ , referred to as *heap part*, and a heap-independent formula π , referred to as *pure part*. The pure part does not contain any heap nodes and is presently restricted to pointer equality/disequality γ , Presburger arithmetic ϕ and bag constraint φ . Furthermore, Δ denotes a composite formula that could always be normalised into the Φ form (see Figure 3). The semantic model for the separation constraints is left in the technical report [19].

Separation constraints are used in pre/post conditions and shape definitions. In order to handle them correctly without running into unmatched residual heap nodes, we require each separation constraint to be well-formed, as given by the following definitions:

Definition 3.1 (Accessible) A variable is said to be accessible w.r.t. a shape predicate if it is a parameter or it is a special variable, either `self` or `res`.

Definition 3.2 (Reachable) Given a heap constraint $\kappa = p::c\langle v^* \rangle * \kappa_1$, node $p::c\langle v^* \rangle$ is reachable from a variable q if and only if the following relation holds:

$$\text{reach}(\kappa, q, p::c\langle v^* \rangle) \stackrel{\text{df}}{=} (p=q) \vee (\kappa_1 = q::c\langle \dots, r, \dots \rangle * \kappa_2 \wedge \text{reach}(\kappa_2, r, p::c\langle v^* \rangle))$$

Definition 3.3 (Well-Formed Constraint) A separation constraint Φ is well-formed if (i) every data node and shape predicate are reachable from their accessible variables, (ii) it is in a disjunctive normal form $\bigvee (\exists v^*. \kappa \wedge \gamma \wedge \phi)^*$ where κ is for heap nodes, γ is for pointer constraint, and ϕ is for arithmetic formula.

The primary significance of the *well-formed* condition is that all heap nodes of a heap constraint are reachable from accessible variables. This allows the entailment checking procedure to correctly match nodes from the consequent with nodes from the antecedent of an entailment relation.

Arbitrary recursive shape relation can lead to non-termination in unfold/fold reasoning. To avoid that problem, we propose to use only *well-founded* shape predicates in our framework.

Definition 3.4 (Well-Founded Predicate) A shape predicate is said to be well-founded if it satisfies four conditions, namely: (i) it is a well-formed constraint, (ii) the parameter `self` may only be bound to a data node and not a predicate, (iii) only `self` is allowed to be bound to a data node and (iv) every predicate is reachable from `self`.

Note that the definitions above are syntactic and can easily be enforced. An example of well-founded shape predicates is `avl` - binary tree with near balanced heights, as follows :

$$\begin{aligned} \text{avl}\langle n, h \rangle &\equiv (\text{self} = \text{null} \wedge n=0 \wedge h=0) \\ &\vee (\text{self}::\text{node2}\langle _, p, q \rangle * p::\text{avl}\langle n_1, h_1 \rangle * q::\text{avl}\langle n_2, h_2 \rangle \\ &\wedge n=1+n_1+n_2 \wedge h=1+\max(h_1, h_2) \wedge -1 \leq h_1 - h_2 \leq 1) \\ &\text{inv } n, h \geq 0 \end{aligned}$$

In contrast, the following three shape definitions are not well-founded.

$$\begin{aligned} \text{foo}\langle n \rangle &\equiv \text{self}::\text{foo}\langle m \rangle \wedge n=m+1 \\ \text{goo}\langle \rangle &\equiv \text{self}::\text{node}\langle _, _ \rangle * q::\text{goo}\langle \rangle \\ \text{too}\langle \rangle &\equiv \text{self}::\text{node}\langle _, q \rangle * q::\text{node}\langle _, _ \rangle \end{aligned}$$

For `foo`, the `self` identifier is bound to a shape predicate. For `goo`, the heap node pointed by `q` is *not* reachable from variable `self`. For `too`, an extra data node is bound to a non-`self` variable. The first example may cause infinite unfolding, while the second example captures an unreachable (junk) heap that cannot be located by our entailment procedure. The last example is just a syntactic restriction to facilitate termination of proof reasoning, and can be easily overcome by introducing intermediate predicates.

3.2 Forward Verification

We use P to denote the program being checked. With pre/post conditions declared for each method in P , we can

$\frac{\boxed{\text{FV-PRED}} \quad XPure_0(\Phi) \implies [0/\text{null}]\pi_0}{\vdash c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0}$	$\frac{\boxed{\text{FV-VAR}} \quad \Delta_1 = (\Delta \wedge \text{res} = v')}{\vdash \{\Delta\} v \{\Delta_1\}}$	$\frac{\boxed{\text{FV-NEW}} \quad \Delta_1 = (\Delta * \text{res} :: c\langle v'_1, \dots, v'_n \rangle)}{\vdash \{\Delta\} \text{new } c\langle v_1, \dots, v_n \rangle \{\Delta_1\}}$	$\frac{\boxed{\text{FV-ASSIGN}} \quad \vdash \{\Delta\} e \{\Delta_1\} \quad \Delta_2 = \exists \text{res} \cdot (\Delta_1 \wedge \{v\} v' = \text{res})}{\vdash \{\Delta\} v := e \{\Delta_2\}}$
$\frac{\boxed{\text{FV-CALL}} \quad t \text{ mn}((t_i v_i)_{i=1}^n) \text{ where } \Phi_{pr} * \Phi_{po} \{..\} \in P \quad \rho = [v'_i/v_i] \quad \Delta \vdash \rho \Phi_{pr} * \Delta_1 \quad \Delta_2 = (\Delta_1 * \Phi_{po})}{\vdash \{\Delta\} \text{mn}(v_1..v_n) \{\Delta_2\}}$		$\frac{\boxed{\text{FV-METH}} \quad V = \{v_1..v_n\} \quad W = \text{prime}(V) \quad \Delta = \Phi_{pr} \wedge \text{nochange}(V) \quad \vdash \{\Delta\} e \{\Delta_1\} \quad (\exists W \cdot \Delta_1) \vdash \Phi_{po} * \Delta_2}{\vdash t_0 \text{ mn}(t_1 v_1, \dots, t_n v_n) \text{ where } \Phi_{pr} * \Phi_{po} \{e\}}$	

Figure 2. Some Forward Verification Rules

apply modular verification to a method’s body using Hoare-style triples $\vdash \{\Delta_1\} e \{\Delta_2\}$. These are *forward verification* rules that expect Δ_1 to be given before computing Δ_2 . Some rules are given in Fig 2 while others are left in the technical report [19]. They are used to track heap states as accurately as possible with path-, flow-, and context-sensitivity. For each call site, [FV-CALL] ensures that the callee’s precondition is satisfied. For each method definition, [FV-METH] checks that its postcondition holds for the method body assuming its precondition. A method postcondition may capture only part of the heap at the end of the method, leaving the residue heap nodes in Δ_2 . For each shape definition, [FV-PRED] checks that its given invariant is a consequence of the well-founded heap formula. The soundness of the forward verification is also left in the technical report.

We now explain the operators/functions used in our verification rules. The operator $\wedge_{\{v\}}$ used in the assignment rule is an instance of *composition with update* operators. Given a state Δ_1 , a state change Δ_2 , and a set of variables to be updated $X = \{x_1, \dots, x_n\}$, the composition operator \oplus_X is defined as:

$$\Delta_1 \oplus_X \Delta_2 =_{df} \exists r_1..r_n \cdot \rho_1 \Delta_1 \oplus \rho_2 \Delta_2$$

where r_1, \dots, r_n are fresh variables;
 $\rho_1 = [r_i/x'_i]_{i=1}^n$; $\rho_2 = [r_i/x_i]_{i=1}^n$

Note that ρ_1 and ρ_2 are substitutions that link each latest value of x'_i in Δ_1 with the corresponding initial value x_i in Δ_2 via a fresh variable r_i . The binary operator \oplus is either \wedge or $*$. The function *nochange*(V) returns a formula asserting that the unprimed and primed versions of each variable in V are equal; *prime*(V) returns the primed form of all variables in V . $[e^*/v^*]$ represents substitutions of v^* by e^* . A special case is $[0/\text{null}]$, which denotes replacement of null by 0. Normalization rules for separation constraints are given in Figure 3. $XPure$ is described in the next section.

3.3 Forward Verification Example

We present the detailed verification of the first branch of the `insert` function from Sec 2. Note that program variables appear primed in formulae whereas logical variables unprimed. The proof is straightforward, except for the last step where a disjunctive heap state is folded to form a shape predicate. The procedure to perform the folding step is presented in Sec 4.

```

{x'::sortl⟨n, mi, ma⟩ * vn'::node⟨v, -⟩} // precondition
  if (vn.val ≤ x.val) then {
    {(x'::node⟨mi, null⟩ * vn'::node⟨v, -⟩ ∧ mi=ma ∧ n=1
    ∧ v≤mi) ∨ (∃q, k · x'::node⟨mi, q⟩ * q::sortl⟨n-1, k, ma⟩
    * vn'::node⟨v, -⟩ ∧ mi≤k ∧ mi≤ma ∧ n≥2 ∧ v≤mi)}
    // unfold and conditional
    vn.next := x;
    {(x'::node⟨mi, null⟩ * vn'::node⟨v, x'⟩ ∧ mi=ma ∧ n=1
    ∧ v≤mi) ∨ (∃q, k · x'::node⟨mi, q⟩ * q::sortl⟨n-1, k, ma⟩
    * vn'::node⟨v, x'⟩ ∧ mi≤k ∧ mi≤ma ∧ n≥2 ∧ v≤mi)}
    // field update
    vn
    {(x'::node⟨mi, null⟩ * vn'::node⟨v, x'⟩ ∧ mi=ma ∧ n=1
    ∧ v≤mi ∧ res=vn') ∨ (∃q, k · x'::node⟨mi, q⟩
    * q::sortl⟨n-1, k, ma⟩ * vn'::node⟨v, x'⟩ ∧ mi≤k ∧ mi≤ma
    ∧ n≥2 ∧ v≤mi ∧ res=vn')} // returned value
  }
  {res::sortl⟨n+1, min(v, mi), max(v, ma)⟩}
  // fold to postcondition

```

4 Entailment

We present in this section the entailment checking rules for the class of constraints used by our verification system.

4.1 Separation Constraint Approximation

Entailment between separation formulae (detailed in section 4.2) is reduced to entailment between pure formulae by successively removing heap nodes from the consequent until only a pure formula remains. When the consequent is

$(\Delta_1 \vee \Delta_2) \wedge \pi$	$\rightsquigarrow (\Delta_1 \wedge \pi) \vee (\Delta_2 \wedge \pi)$	$(\gamma_1 \wedge \phi_1) \wedge (\gamma_2 \wedge \phi_2)$	$\rightsquigarrow (\gamma_1 \wedge \gamma_2) \wedge (\phi_1 \wedge \phi_2)$
$(\Delta_1 \vee \Delta_2) * \Delta$	$\rightsquigarrow (\Delta_1 * \Delta) \vee (\Delta_2 * \Delta)$	$(\exists x \cdot \Delta) \wedge \pi$	$\rightsquigarrow \exists y \cdot ([y/x]\Delta \wedge \pi)$
$(\kappa_1 \wedge \pi_1) * (\kappa_2 \wedge \pi_2)$	$\rightsquigarrow (\kappa_1 * \kappa_2) \wedge (\pi_1 \wedge \pi_2)$	$(\exists x \cdot \Delta_1) * \Delta_2$	$\rightsquigarrow \exists y \cdot ([y/x]\Delta_1 * \Delta_2)$
$(\kappa_1 \wedge \pi_1) \wedge (\pi_2)$	$\rightsquigarrow \kappa_1 \wedge (\pi_1 \wedge \pi_2)$		

Figure 3. Normalization Rules

pure, the heap formula in the antecedent is soundly approximated by function $XPure_n$. The function $XPure_n(\Phi)$, whose definition is given in Fig 4, returns a sound approximation of Φ as formula $\text{ex } i^* \cdot \bigvee (\exists v^* \cdot \pi)^*$ where i^* are (non-null) distinct symbolic addresses of heap nodes of Φ . The function $IsData(c)$ returns `true` if c is a data node, while $IsPred(c)$ returns `true` if c is a shape predicate.

We illustrate how these functions work by computing $XPure_1(p::ll\langle n \rangle)$. Let Φ be the body of the `ll` predicate.

$$\begin{aligned}
Inv_0(p::ll\langle n \rangle) &=_{df} n \geq 0 \\
XPure_0(\Phi) &=_{df} \text{ex } j \cdot (\text{self} = 0 \wedge n = 0) \\
&\quad \vee (\text{self} = j \wedge j > 0 \wedge n - 1 \geq 0) \\
Inv_1(p::ll\langle n \rangle) &=_{df} [p/\text{self}]XPure_0(\Phi) \\
XPure_1(p::ll\langle n \rangle) &=_{df} \text{ex } i \cdot [i/j]Inv_1(p::ll\langle n \rangle) \\
&\rightsquigarrow \exists i \cdot (p = 0 \wedge n = 0) \vee (p = i \wedge i > 0 \wedge n > 0) \\
&= (p = 0 \wedge n = 0) \vee (p > 0 \wedge n > 0)
\end{aligned}$$

The following normalization rules are also used :

$$\begin{aligned}
(\text{ex } I \cdot \phi_1) \vee (\text{ex } J \cdot \phi_2) &\rightsquigarrow \text{ex } I \cup J \cdot (\phi_1 \vee \phi_2) \\
\exists v \cdot (\text{ex } I \cdot \phi) &\rightsquigarrow \text{ex } I \cdot (\exists v \cdot \phi) \\
(\text{ex } I \cdot \phi_1) \wedge (\text{ex } J \cdot \phi_2) &\rightsquigarrow \text{ex } I \cup J \cdot \phi_1 \wedge \phi_2 \wedge \bigwedge_{i \in I, j \in J, i \neq j}
\end{aligned}$$

The $\text{ex } i^*$ construct is converted to $\exists i^*$ when the formula is used as a pure formula. The soundness of $XPure_n$ is formalized by :

Lemma 4.1 (Sound Invariant) *Given a shape predicate $c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0$, we have $\Phi \models Inv_n(\text{self}::c\langle v^* \rangle)$ if $XPure_0(\Phi) \implies [0/\text{null}]\pi_0$. π_0 is said to be sound.*

Proof: By structural induction on Φ .

Lemma 4.2 (Sound Abstraction) *Given a separation constraint Φ where the invariants of the predicates appearing in Φ are sound, we have $\Phi \models XPure_n(\Phi)$.*

Proof : By structural induction on Φ .

Lemma 4.1 ensures that a supplied invariant that passes [FV-PRED] is a semantic consequence of the predicate. Lemma 4.2 asserts that it is safe to approximate an antecedent by using $XPure$ if all the predicate invariants are sound. They also allow the possibility of obtaining a more precise invariant by applying $XPure$ one or more times. For example, when given a pure invariant $n \geq 0$ for the predicate `ll` $\langle n \rangle$, a single application returns $\text{ex } i \cdot (\text{self} = 0 \wedge n = 0 \vee \text{self} = i \wedge i > 0 \wedge n > 0)$ which is sound and more precise, as it relates the nullness of the `self` pointer with the size n of the list.

$\frac{(c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0) \in P}{Inv_0(p::c\langle v^* \rangle) = [p/\text{self}, 0/\text{null}]\pi_0}$
$\frac{(c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0) \in P}{Inv_n(p::c\langle v^* \rangle) = [p/\text{self}, 0/\text{null}]XPure_{n-1}(\Phi)}$
$XPure_n(\bigvee (\exists v^* \cdot \kappa \wedge \pi)^*) =_{df} \bigvee (\exists v^* \cdot XPure_n(\kappa) \wedge [0/\text{null}]\pi)^*$
$XPure_n(\text{emp}) =_{df} \text{true}$
$\frac{IsData(c) \text{ fresh } i}{XPure_n(p::c\langle v^* \rangle) =_{df} \text{ex } i \cdot (p = i \wedge i > 0)}$
$\frac{IsPred(c) \text{ fresh } i^* \quad Inv_n(p::c\langle v^* \rangle) = \text{ex } j^* \cdot \bigvee (\exists u^* \cdot \pi)^*}{XPure_n(p::c\langle v^* \rangle) =_{df} \text{ex } i^* \cdot [i^*/j^*] \bigvee (\exists u^* \cdot \pi)^*}$
$XPure_n(\kappa_1 * \kappa_2) =_{df} XPure_n(\kappa_1) \wedge XPure_n(\kappa_2)$

Figure 4. $XPure$: Translating to Pure Form

The invariants associated with shape predicates play an important role in our system. Without the knowledge $m \geq 0$, the entailment $x::\text{node}(_, y) * y::ll\langle m \rangle \vdash x::ll\langle n \rangle \wedge n \geq 1$ would not have succeeded due to $n \geq 1$. Without the more precise derived invariant using $XPure$ for predicate `ll`, the entailment $x::ll\langle n \rangle \wedge n > 0 \vdash x \neq \text{null}$ would not have succeeded either.

4.2 Separation Constraint Entailment

We express the main procedure for heap entailment by the relation

$$\Delta_A \vdash_V^{\kappa} \Delta_C * \Delta_R$$

which denotes $\kappa * \Delta_A \vdash \exists V \cdot (\kappa * \Delta_C) * \Delta_R$.

The purpose of heap entailment is to check that heap nodes in the antecedent Δ_A are sufficiently precise to cover all nodes from the consequent Δ_C , and to compute a residual heap state Δ_R . κ is the history of nodes from the antecedent that have been used to match nodes from the consequent, V is the list of existentially quantified variables from the consequent. Note that k and V are derived. The entailment checking procedure is invoked with $\kappa = \text{emp}$ and $V = \emptyset$. The entailment checking rules are given in Fig 5. We discuss the matching rule in what follows, and leave unfold/fold rules to Sec 5.

The procedure works by successively matching up heap nodes that can be proven aliased. As the matching process is incremental, we keep the successfully matched nodes from

$\frac{\boxed{\text{ENT-EMP}} \quad \rho = [0/\text{null}] \quad \mathit{XPure}_n(\kappa_1 * \kappa) \wedge \rho \pi_1 \implies \rho \exists V. \pi_2}{\kappa_1 \wedge \pi_1 \vdash_V^\kappa \pi_2 * (\kappa_1 \wedge \pi_1)}$	$\frac{\boxed{\text{ENT-MATCH}} \quad \mathit{XPure}_n(p_1 :: c \langle v_1^* \rangle * \kappa_1 * \pi_1) \implies p_1 = p_2 \quad \rho = [v_1^*/v_2^*] \quad \kappa_1 \wedge \pi_1 \wedge \mathit{freeEqn}(\rho, V) \vdash_{V - \{v_2^*\}}^{\kappa * p_1 :: c \langle v_1^* \rangle} \rho(\kappa_2 \wedge \pi_2) * \Delta}{p_1 :: c \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * \Delta}$	$\frac{\boxed{\text{ENT-RHS-EX}} \quad \Delta_1 \vdash_V^\kappa \{w\} ([w/v] \Delta_2) * \Delta_3 \quad \mathit{fresh} \ w \quad \Delta = \exists w. \Delta_3}{\Delta_1 \vdash_V^\kappa (\exists w. \Delta_2) * \Delta}$
$\frac{\boxed{\text{ENT-FOLD}} \quad \mathit{IsPred}(c_2) \wedge \mathit{IsData}(c_1) \quad (\Delta^r, \kappa^r, \pi^r) \in \mathit{fold}^\kappa(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1, p_2 :: c_2 \langle v_2^* \rangle) \quad \mathit{XPure}_n(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 * \pi_1) \implies p_1 = p_2 \quad (\pi^a, \pi^c) = \mathit{split}_V^{\{v_2^*\}}(\pi^r) \quad \Delta^r \wedge \pi^a \vdash_V^{\kappa^r} (\kappa_2 \wedge \pi_2 \wedge \pi^c) * \Delta}{p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * \Delta}$		$\frac{\boxed{\text{ENT-LHS-EX}} \quad [w/v] \Delta_1 \vdash_V^\kappa \Delta_2 * \Delta \quad \mathit{fresh} \ w}{\exists v. \Delta_1 \vdash_V^\kappa \Delta_2 * \Delta}$
$\frac{\boxed{\text{ENT-UNFOLD}} \quad \mathit{XPure}_n(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 * \pi_1) \implies p_1 = p_2 \quad \mathit{IsPred}(c_1) \wedge \mathit{IsData}(c_2) \quad \mathit{unfold}(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * \Delta)}{p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * \Delta}$	$\frac{\boxed{\text{ENT-LHS-OR}} \quad \Delta_1 \vdash_V^\kappa \Delta_3 * \Delta_4 \quad \Delta_2 \vdash_V^\kappa \Delta_3 * \Delta_5}{\Delta_1 \vee \Delta_2 \vdash_V^\kappa \Delta_3 * (\Delta_4 \vee \Delta_5)}$	$\frac{\boxed{\text{ENT-RHS-OR}} \quad \Delta_1 \vdash_V^\kappa \Delta_i * \Delta_i^R}{\Delta_1 \vdash_V^\kappa (\Delta_2 \vee \Delta_3) * \Delta_i^R} \quad i \in \{2, 3\}$

Figure 5. Separation Constraint Entailment

antecedent in κ for better precision. For example, consider the following (valid) proof:

$$\frac{\frac{((p = \text{null} \wedge n = 0) \vee (p \neq \text{null} \wedge n > 0)) \wedge n > 0 \wedge m = n \implies p \neq \text{null}}{\Delta_R = (n > 0 \wedge m = n)}}{n > 0 \wedge m = n \vdash^{p::\text{ll}\langle n \rangle} p \neq \text{null} * \Delta_R}}{p::\text{ll}\langle n \rangle \wedge n > 0 \vdash p::\text{ll}\langle m \rangle \wedge p \neq \text{null} * \Delta_R}$$

Had the predicate $p::\text{ll}\langle n \rangle$ not been kept and used, the proof would not have succeeded. Such an entailment would be useful when, for example, a list with positive length n is used as input for a function that requires a non-empty list.

Another feature of the entailment procedure is exemplified by the transfer of $m = n$ to the antecedent (and subsequently to the residue). In general, when a match occurs (rule `[ENT-MATCH]`) and an argument of the heap node coming from the consequent is free, the entailment procedure binds the argument to the corresponding variable from the antecedent and moves the equality to the antecedent. In our system, free variables in consequent are variables from method preconditions. Hence these bindings act as substitutions that have to be kept in antecedent to allow subsequent program state (from residual heap) to be aware of their values. This process is formalized by the function $\mathit{freeEqn}$ below, where V is the set of existentially quantified variables:

$$\mathit{freeEqn}([u_i/v_i]_{i=1}^n, V) =_{df} \text{let } \pi_i = (\text{if } v_i \in V \text{ then true else } v_i = u_i) \text{ in } \bigwedge_{i=1}^n \pi_i$$

For soundness, we perform a preprocessing step to ensure that variables appearing as arguments of heap nodes and predicates are i) distinct and ii) if they are free, they do not appear in the antecedent by adding (existentially quantified) fresh variables and equalities. This guarantees that the generated substitutions are well-defined. It also guar-

antees that the formula generated by $\mathit{freeEqn}$ does not introduce any additional constraints over existing variables in the antecedent, as one side of each equation does not appear anywhere else in the antecedent.

5 Unfold/Fold Mechanism

Unfold/fold operations can be used to handle well-founded inductive predicates in a deductive manner. In particular, we can unfold a predicate that appears in the antecedent if it matches with a data node in the consequent. Correspondingly, we fold a predicate that appears in the consequent if it matches with a data node in the antecedent. The well-founded condition is sufficient to ensure termination.

5.1 Unfolding a Shape Predicate in the Antecedent

We apply an unfold operation on a predicate in the antecedent that matches with a data node in the consequent. Consider :

$$x::\text{ll3}\langle n, B \rangle \wedge n > 2 \vdash (\exists r. x::\text{node}\langle r, y \rangle \wedge y \neq \text{null} \wedge r \in B) * \Delta_R$$

where Δ_R captures the residue of entailment. For the entailment to succeed, we would first unfold the $\text{ll3}\langle n, B \rangle$ predicate in the antecedent:

$$\exists q_1, v. x::\text{node}\langle v, q_1 \rangle * q_1::\text{ll3}\langle n-1, B_1 \rangle \wedge n > 2 \wedge B = B_1 \cup \{v\} \vdash (\exists r. x::\text{node}\langle r, y \rangle \wedge y \neq \text{null} \wedge r \in B) * \Delta_R$$

After removing the existential quantifiers, we obtain:

$$x::\text{node}\langle v, q_1 \rangle * q_1::\text{ll3}\langle n-1, B_1 \rangle \wedge n > 2 \wedge B = B_1 \cup \{v\} \vdash (x::\text{node}\langle r, y \rangle \wedge y \neq \text{null} \wedge r \in B) * \Delta_R$$

The data node in the consequent is then matched up, giving:

$$\begin{aligned} & q_1 :: \text{ll3}\langle n-1, B_1 \rangle \wedge n > 2 \wedge B = B_1 \cup \{v\} \wedge q_1 = y \\ & \vdash (q_1 \neq \text{null} \wedge v \in B) * \Delta_R \end{aligned}$$

Due to the well-founded condition, each unfolding exposes a data node that matches the data node in the consequent. Thus a reduction of the consequent immediately follows, which contributes to the termination of the entailment check. A formal definition of unfolding is given by the [UNFOLDING] rule:

$$\frac{\text{[UNFOLDING]} \quad c\langle v^* \rangle \equiv \Phi \in P}{\text{unfold}(p :: c\langle v^* \rangle) =_{df} [p/\text{self}]\Phi}$$

5.2 Folding a Shape Predicate in the Consequent

We apply a fold operation when a data node in the antecedent matches with a predicate in the consequent. An example is :

$$\begin{aligned} & x :: \text{node}\langle 1, q_1 \rangle * q_1 :: \text{node}\langle 2, \text{null} \rangle * y :: \text{node}\langle 3, \text{null} \rangle \\ & \vdash (x :: \text{ll3}\langle n, B \rangle \wedge n > 1 \wedge 1 \in B) * \Delta_R \end{aligned}$$

The fold step may be recursively applied but is guaranteed to terminate for well-founded predicate as it will reduce a data node in the antecedent for each recursive invocation. This reduction in the antecedent cannot go on forever. Furthermore, the fold operation may introduce bindings for the parameters of the folded predicate. In the above, we obtain $\exists n_1, n_2 \cdot n = n_1 + 1 \wedge n_1 = n_2 + 1 \wedge n_2 = 0$ and $\exists B_1, B_2 \cdot B = B_1 \cup \{2\} \wedge B_1 = \{1\} \cup B_2 \wedge B_2 = \{\}$, where n_1, n_2, B_1, B_2 are existential variables introduced by the folding process, and are subsequently eliminated. Each of the previous formulae may be transferred to the antecedent if n and B are free, respectively. Otherwise, they will be kept in the consequent. Since n and B are indeed free, our folding step would finally derive:

$$\begin{aligned} & y :: \text{node}\langle 3, \text{null} \rangle \wedge n = 2 \wedge B = \{1, 2\} \\ & \vdash (n > 1 \wedge 1 \in B) * \Delta_R \end{aligned}$$

The effects of folding may seem similar to unfolding the predicate in the consequent. However, there is a subtle difference in their handling of bindings for free derived variables. If we choose to use unfolding on the consequent instead, these bindings may not be transferred to the antecedent. Consider the example below where n is free :

$$z = \text{null} \vdash z :: \text{ll3}\langle n, B \rangle \wedge n > -1 * \Delta_R$$

By unfolding the predicate $\text{ll}\langle n \rangle$ in the consequent, we obtain :

$$\begin{aligned} & z = \text{null} \vdash (z = \text{null} \wedge n = 0 \wedge B = \{\} \wedge n > -1) \\ & \quad \vee (\exists q, v \cdot z :: \text{node}\langle v, q \rangle * q :: \text{ll3}\langle n-1, B_1 \rangle \\ & \quad \wedge B = B_1 \cup \{v\} \wedge n > -1) * \Delta_R \end{aligned}$$

There are now two disjuncts in the consequent. The second one fails because it mismatches. The first one matches but still fails as the derived binding $n=0$ was not transferred to the antecedent.

When a fold to a predicate $p_2 :: c_2\langle v_2^* \rangle$ is performed, the constraints related to variables v_2^* are important. The *split* function projects these constraints out and differentiates those constraints based on free variables.

$$\begin{aligned} & \text{split}_V^{\{v_2^*\}}(\bigwedge_{i=1}^n \pi_i^r) = \\ & \quad \text{let } \pi_i^a, \pi_i^c = \text{if } FV(\pi_i^r) \cap v_2^* = \emptyset \text{ then } (\text{true}, \text{true}) \\ & \quad \quad \text{else if } FV(\pi_i^r) \cap V = \emptyset \text{ then } (\pi_i^r, \text{true}) \\ & \quad \quad \quad \text{else } (\text{true}, \pi_i^r) \\ & \quad \text{in } (\bigwedge_{i=1}^n \pi_i^a, \bigwedge_{i=1}^n \pi_i^c) \end{aligned}$$

$$\frac{\text{[FOLDING]} \quad c\langle v^* \rangle \equiv \Phi \in P \quad W_i = V_i - \{v^*, p\}}{\text{fold}^{\kappa'}(\kappa \wedge \pi_{\{p, v^*\}}^{\kappa'} [p/\text{self}]\Phi * \{(\Delta_i, \kappa_i, V_i, \pi_i)\}_{i=1}^n) =_{df} \{(\Delta_i, \kappa_i, \exists W_i \cdot \pi_i)\}_{i=1}^n}$$

A formal definition of folding is specified by rule [FOLDING]. Some heap nodes from κ are removed by the entailment procedure so as to match with the heap formula of predicate $p :: c\langle v^* \rangle$. This requires a special version of entailment that returns three extra things: (i) consumed heap nodes, (ii) existential variables used, and (iii) final consequent. The final consequent is used to return a constraint for $\{v^*\}$ via $\exists W_i \cdot \pi_i$. A set of answers is returned by the fold step as we allow it to explore multiple ways of matching up with its disjunctive heap state. Our entailment also handles empty predicates correctly with a couple of specialised rules.

6 Soundness of Entailment

The following theorems state that our entailment check procedure (given in Fig. 5) is sound and terminating. Proofs are given in the technical report [19].

Theorem 6.1 (Soundness) *If entailment check $\Delta_1 \vdash \Delta_2 * \Delta$ succeeds, we have: for all s, h , if $s, h \models \Delta_1$ then $s, h \models \Delta_2 * \Delta$.*

Theorem 6.2 (Termination) *The entailment check $\Delta_1 \vdash \Delta_2 * \Delta$ always terminates.*

7 Implementation

We have built a prototype system using Objective Caml. The proof obligations generated by our verification are discharged using either a constraint solver or a theorem prover. This is organised as an option in our system and currently

covers automatic provers of Omega Calculator [22], Isabelle [21], and MONA [13].

Figure 6 summarizes a suite of programs tested. These examples use complicated recursion and data structures with sophisticated shape, size and bag properties. They show that our approach is general enough to handle interesting data structures such as sorted lists, sorted trees, priority queues, balanced trees, etc. in a uniform way. Verification time of a function includes time to verify all functions that it calls.

Arithmetic Constraints. The time required for shape and size verification is mostly within a couple of seconds when using the Omega Calculator to discharge the proof obligations (third column). In order to have a reference point for the Omega timings, we tried solving the same constraints with two other theorem provers: Isabelle (fourth column) and MONA (fifth column). From the former, we only used an automatic but incomplete prover, which caused some proofs to fail. The latter is an implementation of the weak monadic second-order logics WS1S and WS2S. Therefore, first-order variables can be compared and subjected only to addition with constants. As Presburger arithmetic allows the addition of arbitrary linear arithmetic terms, we converted its formulas into WS1S by encoding naturals as Base-2 bit strings. MONA translates WS1S and WS2S formulas into minimum DFAs (Deterministic Finite Automata) and GTAs (Guided Tree Automata), respectively. However, this translation may cause a state-space explosion. In our case, we confronted such a problem when verifying the `quick_sort` method, where the size of its corresponding automaton exceeded the available space.

Bag/Set Constraints. Bag constraints were solved using the multiset theory of Isabelle (sixth column), while weak monadic second-order theory of 1 successor WS1S from MONA was used to handle set constraints (seventh column). Due to the incompleteness of the automatic prover that we used from Isabelle, some of the proofs failed (`quick_sort`, `delete_max` from priority queue heap and `insert` into a red-black tree). On the other hand, MONA seemed to provide satisfactory results when handling set constraints.

From our experiments, we may conclude that verification process is dominated by entailment proving of pure formulas that are fast with specialised solvers, such as Omega for Presburger constraint and MONA for set constraints. The timings for verifying shapes only (without size/bag proving) are benign, as reflected in the second column. An important future work is to design a combination strategy that allows specialised solvers to work well with mixed constraints and to handle large programs. We expect *code modularity*, *decomposed shape views* and *multi-core parallelism* to be important techniques in performance engineering of automated verification system.

8 Related Work

Separation Logic. The general framework of separation logic [23, 11] is highly expressive but undecidable. Likewise, [17] formalised the proof rules for handling abstract predicates (with scopes on visibility of predicates) but provided no automated procedure for checking the user supplied specifications. In the search for a decidable fragment of separation logic for automated verification, Berdine *et al.* [1] supports only a limited set of predicates *without* size properties, disjunctions and existential quantifiers. Similarly, Jia and Walker [12] postponed the handling of recursive predicates in their recent work on automated reasoning of pointer programs. Our approach is more pragmatic as we aim for a sound and terminating formulation of automated verification via separation logic but do not aim for completeness in the expressive fragment that we handle. On the inference front, Lee *et al.* [16] has conducted an intraprocedural analysis for loop invariants using grammar approximation under separation logic. Their analysis can handle a wide range of shape predicates with local sharing but is restricted to predicates with two parameters and without size properties. A recent work [9] has also formulated interprocedural shape inference but is restricted to just the list segment shape predicate. Sims [26] extends separation logic with fixpoint connectives and postponed substitution to express recursively defined formulae to model the analysis of while-loops. However, it is unclear how to check for entailment in their extended separation logic. While our work does not address the inference/analysis challenge, we have succeeded in providing direct support for automated verification via an expressive shape and size specification mechanism.

Shape Checking/Analysis. Many formalisms for shape analysis have been proposed for checking user programs' intricate manipulations of shapely data structures. One well-known work is Pointer Assertion Logic [18] by Moeller and Schwartzbach where shape specifications in monadic second-order logic are given by programmers for loop invariants and method pre/post conditions, and checked by their MONA tool. For shape inference, Sagiv *et al.* [25] presented a parameterised framework, called TVLA, using 3-valued logic formulae and abstract interpretation. Based on the properties expected of data structures, programmers must supply a set of predicates to the framework which are then used to analyse that certain shape invariants are maintained. However, most of these techniques were focused on analysing shape invariants, and did not attempt to track the size properties of complex data structures. An exception is the quantitative shape analysis of Rugina [24] where a data flow analysis was proposed to compute quantitative information for programs with destructive updates. By tracking unique points-to reference and its height

Programs	Without size/bag	Omega Calculator	Isabelle Prover	MONA Prover	Isabelle Prover	MONA Prover
Linked List		verifies size/length			verifies bag/set	
delete	0.02	0.09	8.35	0.33	5.00	0.34
reverse	0.02	0.07	3.28	0.21	3.01	0.20
Circular Linked List		verifies size + cyclic structure			verifies bag/set + cyclic structure	
delete (first)	0.01	0.09	5.46	0.26	7.17	0.40
count	0.04	0.16	14.99	0.71	21.01	2.29
Doubly Linked List		verifies size + double links			verifies bag/set + double links	
append	0.05	0.16	28.18	0.83	23.73	0.93
flatten (from tree)	0.08	0.30	158.3	6.65	55.78	2.03
Sorted List		verifies size + min + max + sortedness			verifies bag/set + sortedness	
delete	0.02	0.13	34.09	26.68	51.39	0.60
insertion_sort	0.07	0.27	41.17	18.22	27.34	0.73
selection_sort	0.10	0.41	79.08	20.62	221.7	1.10
bubble_sort	0.16	0.64	358.7	9.36	221.2	2.84
merge_sort	0.11	0.61	342.9	105.1	150.1	21.75
quick_sort	0.19	0.59	642.0	<i>out of memory</i>	<i>failed</i>	3.40
Binary Search Tree		verifies min + max + sortedness			verifies bag/set + sortedness	
insert	0.03	0.20	<i>failed</i>	11.92	99.57	0.95
delete	0.06	0.38	97.5	6.86	943.5	3.03
Priority Queue Heap		verifies size + height + max-heap			verifies bag/set + size + max-heap	
insert	0.15	0.45	520.8	41.55	416.2	6.45
delete_max	0.55	7.17	<i>failed</i>	290.7	<i>failed</i>	626.1
AVL Tree		verifies size + height + height-balanced			verifies bag/set + height + height-balanced	
insert	1.04	5.06	<i>failed</i>	36.02	1973	7.38
Red-Black Tree		verifies size + black-height + height-balanced			verifies bag/set + black-height + height-balanced	
insert	0.44	1.53	2992	352.4	<i>failed</i>	392.8

Figure 6. Verification Times (in seconds) for Data Structures with Arithmetic and Bag/Set Constraints

property, their algorithm is able to handle AVL-like tree structures. Even then, the author acknowledged the lack of a general specification mechanism for handling arbitrary shape/size properties.

Size Properties. In another direction of research, size properties have been most explored for declarative languages [10, 28, 6] as the immutability property makes their data structures easier to analyse statically. Size analysis was later extended to object-based programs [7] but was restricted to tracking either size-immutable objects that can be aliased and size-mutable objects that are unaliased, with no support for complex shapes. The Applied Type System (ATS) [5] was proposed for combining programs with proofs. In ATS, dependent types for capturing program invariants are extremely expressive and can capture many program properties with the help of accompanying proofs. Using linear logic, ATS may also handle mutable data structures with sharing. However, users must supply all expected properties, and precisely state where they are to be applied, with ATS playing the role of a proof-checker. Comparatively, we use a more limited class of constraint for shape and size analysis but supports automated modular verification.

Set/Bag Properties. Set-based analysis has been proposed to verify data structure consistency properties in [14], where a decision procedure is given for a first order theory that combines set and Presburger arithmetic. This result may be used to build a specialised mixed constraint solver but currently has high algorithmic complexity. In [15], Lahiri and Qadeer reported an intra-procedural reachability analysis for well-founded linked lists using first-order axiomatization. Reachability analysis is related to set/bag property that we capture but implemented by transitive closure at the predicate level.

Unfold/Fold Mechanism. Unfold/fold techniques were originally used for program transformation [4] on purely functional programs. A similar technique called unroll/roll was later used in alias types [27] to *manually* witness the isomorphism between a recursive type and its unfolding. Here, each unroll/roll step must be manually specified by programmer, in contrast to our approach which applies these steps automatically during entailment checking. In [1], an automated procedure that uses unroll/roll was given but it was hardwired to work for only `1seg` and `tree` predicates. Furthermore, it performs rolling by unfolding a pred-

icate in the consequent which would miss bindings on free variables. Our unfold/fold mechanism is general, automatic and terminates for heap entailment checking.

9 Conclusion

We have presented a new approach to verifying pointer-based programs that can precisely track shape, size and bag properties. Our approach is built on well-founded shape relations and well-formed separation constraints from which we have designed a sound procedure for heap entailment. Our automated deduction mechanism is based on the unfold/fold reasoning of user-definable predicates that has been proven to be sound and terminating.

Acknowledgement

This work is supported by the Singapore-MIT Alliance and NUS research grant R-252-000-213-112.

References

- [1] J. Berdine, C. Calcagno, and P. W. O'Hearn. Symbolic Execution with Separation Logic. In *APLAS*. Springer-Verlag, November 2005.
- [2] J. Berdine, C. Calcagno, and P. W. O'Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, Springer LNCS 4111, 2006.
- [3] J. Bingham and Z. Rakamaric. A Logic and Decision Procedure for Predicate Abstraction of Heap-Manipulating Programs. In *VMCAI*, Springer LNCS 3855, pages 207–221, Charleston, U.S.A, January 2006.
- [4] R.M. Burstall and J. Darlington. A transformation system for developing recursive programs. *Journal of ACM*, 24(1):44–67, January 1977.
- [5] C. Chen and H. Xi. Combining Programming with Theorem Proving. In *ACM SIGPLAN ICFP*, Tallinn, Estonia, September 2005.
- [6] W.N. Chin and S.C. Khoo. Calculating sized types. In *ACM SIGPLAN PEPM*, pages 62–72, Boston, United States, January 2000.
- [7] W.N. Chin, S.C. Khoo, S.C. Qin, C. Popeea, and H.H. Nguyen. Verifying Safety Policies with Size Properties and Alias Controls. In *ACM SIGSOFT ICSE*, St. Louis, Missouri, May 2005.
- [8] D. Distefano, P. W. O'Hearn, and H. Yang. A Local Shape Analysis based on Separation Logic. In *TACAS*. Springer-Verlag, March 2006.
- [9] A. Gotsman, J. Berdine, and B. Cook. Interprocedural Shape Analysis with Separated Heap Abstractions. In *SAS*, Springer LNCS, Seoul, Korea, August 2006.
- [10] J. Hughes, L. Pareto, and A. Sabry. Proving the correctness of reactive systems using sized types. In *ACM POPL*, pages 410–423. ACM Press, January 1996.
- [11] S. Isthiaq and P.W. O'Hearn. BI as an assertion language for mutable data structures. In *ACM POPL*, London, January 2001.
- [12] L. Jia and D. Walker. ILC: A foundation for automated reasoning about pointer programs. In *15th ESOP*, March 2006.
- [13] Nils Klarlund and Anders Mller. Mona version 1.4 - user manual.
- [14] V. Kuncak, H. H. Nguyen, and M. Rinard. An algorithm for deciding bapa: Boolean algebra with presburger arithmetic. In *20th International Conference on Automated Deduction (CADE-20)*, Tallinn, Estonia, Jul 2005.
- [15] S. Lahiri and S. Qadeer. Verifying Properties of Well-Founded Linked Lists. In *ACM POPL*, South Carolina, January 2006.
- [16] O. Lee, H. Yang, and K. Yi. Automatic verification of pointer programs using grammar-based shape analysis. In *ESOP*. Springer Verlag, April 2005.
- [17] M.J.Parkinson and G.M.Bierman. Separation logic and abstraction. In *ACM POPL*, pages 247–258, 2005.
- [18] A. Moeller and M. I. Schwartzbach. The Pointer Assertion Logic Engine. In *ACM PLDI*, June 2001.
- [19] H.H. Nguyen, C. David, S.C. Qin, and W.N. Chin. Automated Verification of Shape, Size and Bag Properties via Separation Logic. Technical report, SoC, Natl Univ. of Singapore, July 2006. avail. at <http://www.comp.nus.edu.sg/~chinwn/papers/verify-report.pdf>.
- [20] H.H. Nguyen, C. David, S.C. Qin, and W.N. Chin. Automated Verification of Shape And Size Properties via Separation Logic. In *VMCAI*, Nice, France, January 2007.
- [21] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [22] W. Pugh. The Omega Test: A fast practical integer programming algorithm for dependence analysis. *Communications of the ACM*, 8:102–114, 1992.
- [23] J. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *IEEE LICS*, Copenhagen, Denmark, July 2002.
- [24] R. Rugina. Quantitative Shape Analysis. In *SAS*, Springer LNCS, Verona, Italy, August 2004.
- [25] S. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *ACM TOPLAS*, 24(3), May 2002.
- [26] É.-J. Sims. Extending separation logic with fixpoints and postponed substitution. *Theoretical Computer Science*, 351(2):258–275, 2006.
- [27] D. Walker and G. Morrisett. Alias Types for Recursive Data Structures. In *TIC*, Springer LNCS 2071, pages 177–206, 2000.
- [28] H. Xi. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, 1998.