

Multiple Pre/Post Specifications for Heap-Manipulating Methods

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

Automated verification plays an important role for high assurance software. This typically uses a pair of pre/post conditions as a formal (but possibly partial) specification of each method before it is systematically verified. In this paper, we advocate for multiple pairs of pre/post conditions to be associated with each method which provides a way for such specification to be used in more scenarios. Multiple pre/post specifications are important for heap-manipulating programs where they can be precisely expressed using separation logic. This work highlights the importance of multiple pre/post specifications, and a methodology to capture them via set of states during proof search.

1 Introduction

In recent years, separation logic formalism has been successfully applied to analysing and verifying heap-manipulating programs. This formalism supports succinct description of shapely data structures, such as near-balanced AVL-trees [10], and has been used to analyse and verify various program properties. One important feature of separation logic is its support for accurate references into memory states which can be captured with the help of inductive shape predicate and separating conjunction. For example, a list segment of length n can be specified by the following predicate definition:

$$\text{lseg}(\text{root}, p, n) \equiv \text{root} = p \wedge n = 0 \vee \\ \exists q \cdot \text{root} \mapsto \text{node}(_, q) * \text{lseg}(q, p, n-1) \text{ inv } n \geq 0$$

where `node` is an object type declared as

```
data node { int val; node next }
```

In separation logic notation, $x \mapsto \text{node}(a, b)$ captures a distinct memory cell referenced from x , while separating conjunction $\Delta_1 * \Delta_2$ captures two *disjoint* heaps described by Δ_1 and Δ_2 , respectively. In contrast, $\Delta_3 \wedge \Delta_4$ captures two *overlapping* heaps, Δ_3 and Δ_4 . The formula $n \geq 0$ after the `inv` keyword captures a heap-independent invariant for

the predicate. This invariant is an approximation of the predicate, in the sense that whenever the predicate holds, the invariant also holds. By default, each shape predicate is expected to have a special (first) parameter, named `root`, that can transitively reach all its memory cells. This “root” parameter will be instantiated with an actual argument for each instance of its predicate that appears in formulas. Furthermore, we shall use a uniform notation where object $x \mapsto \text{node}(a, \dots)$ is written as $x :: \text{node}(a, \dots)$, and predicate $\text{pred}(x, a, \dots)$ is written as $x :: \text{pred}(a, \dots)$. We refer to them collectively as *heap nodes*. With this new notation, we can omit the `root` parameter in the head of the predicate, and define a non-empty circular list, as follows:

$$\text{clist}(n) \equiv \exists p \cdot \text{root} :: \text{node}(_, p) * p :: \text{lseg}(\text{root}, n-1) \\ \text{inv } n \geq 1$$

Shape predicates can also be used to describe more complex data structures with stronger properties. An example is the non-empty sorted list:

$$\text{sortl}(n, s, b) \equiv \text{root} :: \text{node}(s, \text{null}) \wedge s = b \wedge n = 1 \\ \vee \exists p, t \cdot \text{root} :: \text{node}(s, p) * p :: \text{sortl}(n-1, t, b) \wedge s \leq t \\ \text{inv } n > 0 \wedge s \leq b$$

Predicate `sortl`(n, s, b) ensures that all values in the list are sorted in ascending order, with s and b to capture its min and max values, respectively. The sortedness property is ensured by the presence of $s \leq t$ in the above predicate. The parameters n, s, b actually capture some derived properties of the heap data structure, as stated in the predicate definition. These parameters play a role similar to “model fields” in some specification languages, such as Spec# [9, 1] and JML [3]. As a shorthand, we may omit the existential quantifiers without ambiguity.

Though separation logic with inductive predicates can be highly expressive, current automated theorem provers hardly provide any support for this form of substructural logic. Several recent works, such as [2, 4], have attempted to address this shortcoming by building specialised solvers that work for a fixed set of predicates (e.g. `lseg` without the size parameter). Our recent work [10] has lifted one crucial limitation by supporting the automated reasoning of

user-defined predicates. To this end, we have designed a general prover that uses unfold/fold reasoning on predicates to support entailment between two heap states, and also for computing its residual heap. This is essentially a prover with “frame inferring” capability. Given two heap states Δ_a and Δ_c , our prover can check for an entailment of the form: $\Delta_a \vdash \Delta_c * \Delta_r$, where Δ_r is a residual heap from Δ_a after fully accounting for the memory heap state of Δ_c . The residue is a consequence of the frame rule from separation logic that can be determined by an incremental matching algorithm in our prover. Our entailment also handles disjunctive formulae and existential quantifiers. While our prover is sound, terminating and automated, it is incomplete.

In this paper, we propose one approach to partially overcome this incompleteness shortcoming, by providing a mechanism for specifying and handling multiple pre/post conditions in separation logic. Traditionally, each method is given a single pair of pre/post conditions which describes the expected pre-state prior to a method call and its subsequent post-state. Furthermore, even when multiple pre/post conditions are allowed in some specification languages, such as JML [3], the standard technique for handling them is to re-combine into a single pre/post condition. As an example, consider a method with two (pairs of) pre/post conditions, $(pre_1, post_1)$ and $(pre_2, post_2)$, as shown:

```
requires pre1 ensures post1
also
requires pre2 ensures post2
```

The proposed solution in JML [6] is to transform the two pre/post conditions into an integrated pre/post condition:

```
requires pre1  $\vee$  pre2
ensures old(pre1)  $\implies$  post1  $\wedge$  old(pre2)  $\implies$  post2
```

While this re-combination may work with traditional Hoare logic that is based on pure formulae, it cannot be used for an arbitrary heap state Δ of separation logic, since $old(\Delta)$ is not always determinable. As an example of multiple pre/post conditions in separation logic formula, consider the `append` method for joining two lists together:

```
void append(node x, node y)
{ if x.next  $\neq$  null then { append(x.next, y) }
  else { x.next := y } }
```

A simple specification for this method is for joining two disjoint lists into a single longer list, as illustrated by the first pre/post condition below:

```
requires x::lseg<null, n1>*y::lseg<null, n2> $\wedge$ n1>0
ensures x::lseg<null, n1+n2>
```

The $n_1 > 0$ constraint ensures that the first input is non-empty. As a result, the pointer access operation `x.next \neq null` can be proven to be safe. Another radically different view of this method is for the purpose of joining two sorted lists, whereby the largest value of the first sorted list is less than the smallest value of the second sorted list. This view is captured by a second pre/post condition below:

```
requires x::sortl<n1, a, b>*y::sortl<n2, c, d> $\wedge$ b  $\leq$  c
ensures x::sortl<n1+n2, a, d>
```

With a more specialised pre-state, our specification is able to conclude that the resulting list is a longer sorted list. Furthermore, the `sortl` predicate is non-empty which can prove that dereferencing by `x.next \neq null` is safe. Such a specification captures a different view for the same `append` method, which would be needed towards the verification of the sortedness property for the quicksort algorithm. Our thesis is that it is often futile to combine these widely different pre/post conditions into a single pre/post condition, as the translations for $old(x::lseg<null, n_1>*y::lseg<null, n_2>\wedge n_1 > 0)$ and $old(x::sortl<n_1, a, b>*y::sortl<n_2, c, d>\wedge b \leq c)$ into heap (or pure) formulae cannot be systematically determined. For example, in JML, `old(e)` is guaranteed to be safe to use if `e` denotes a primitive value, or an immutable object [7] (sec 11.4.2). Otherwise, some ambiguity in specification is possible.

In this paper, we propose a new mechanism to support the direct handling of multiple pre/post conditions. Our mechanism is based on a set of states $\{\Delta_1, \dots, \Delta_n\}$ which represent abstract states that may arise during proof search used by automated verification. As a first step, we propose to generalise the entailment procedure into a non-deterministic version that explicitly returns a set of residual states, namely: $\Delta_a \vdash \Delta_c * \{\Delta_1, \dots, \Delta_n\}$, such that for all $1 \leq i \leq n$, $\Delta_a \vdash \Delta_c * \Delta_i$ holds. In the event of an entailment failure, the set of residual states is empty. Note that “states” here are formulas, representing the abstract states of the program.

The set of residual states is meant to contain syntactically different states. This does not imply that all the states in the set are semantically different. Our approach may indeed cause some redundancy in the search for proofs, but it does not affect soundness of the entailment prover. Explicit set of states allows easier integration between the prover and other components of our verification system. In particular, its use is critical for handling multiple pre/post specifications which can now be viewed as part of proof search.

This paper makes the following contributions:

- **Multiple pre/post:** This mechanism allows users to specify more properties of a procedure. Our system automatically verifies that each user-provided specification is correct, prior to using them for systematic proof search.
- **Set of states:** Provers for separation logic typically need to derive residues (or frames), which have to be communicated to program verifiers. Our direct support for non-deterministic proof search via set of states notation makes this task both explicit and exhaustive.
- **Implementation:** We have implemented a verification

s_{pred}	$::= c\langle v^* \rangle \equiv \Phi \text{ inv } \pi_0$
Φ	$::= \bigvee (\exists v^* \cdot \kappa \wedge \pi)^*$ $\pi ::= \gamma \wedge \phi$
γ	$::= v_1 = v_2 \mid v = \text{null} \mid v_1 \neq v_2 \mid v \neq \text{null} \mid \gamma_1 \wedge \gamma_2$
κ	$::= \text{emp} \mid v :: c\langle v^* \rangle \mid \kappa_1 * \kappa_2$
Δ	$::= \Phi \mid \Delta_1 \vee \Delta_2 \mid \Delta \wedge \pi \mid \Delta_1 * \Delta_2 \mid \exists v \cdot \Delta$
ϕ	$::= \varphi \mid b \mid a \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg \phi \mid \exists v \cdot \phi \mid \forall v \cdot \phi$
b	$::= \text{true} \mid \text{false} \mid v \mid b_1 = b_2$
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	$::= v \in B \mid B_1 = B_2 \mid B_1 \sqsubset B_2 \mid \forall v \in B \cdot \phi \mid \exists v \in B \cdot \phi$
B	$::= B_1 \sqcup B_2 \mid B_1 \cap B_2 \mid B_1 - B_2 \mid \{ \} \mid \{ v \}$

Figure 1. Syntax for Formulas

system that supports multiple pre/post specifications. Initial experiments show that the cost of conducting proof search via set of states remains low, but is crucial to verify some examples. Moreover, the smaller heap formulae from using multiple pre/post specifications may sometimes result in significant improvement to the performance of automated verification.

2 Separation Logic Formalism Used

The syntax for separation logic formulae that we will use is given in Figure 1. Each shape predicate s_{pred} has a body Φ , which can be recursive as it can mention the predicate being defined. Each predicate can also be equipped with a heap-independent invariant π_0 that is user-supplied but machine-checked for its validity. The invariant is marked with a prior keyword **inv**.

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 restricted to pointer equality/disequality (to facilitate precise aliasing and non-aliasing), Presburger arithmetic and set/bag constraints that can capture a collection of reachable values/addresses. Furthermore, Δ denotes a composite formula that could always be normalised into the Φ form (see Fig. 5 in Sec 4.1).

Separation constraints are used in pre/post conditions and shape definitions. We will also use another special variable “**res**” to denote the returned value of an expression (including the method body). In order to handle them correctly without running into unmatched residual heap nodes, we require all constraints used in specifications be well-formed.

Definition 2.1 (Accessible) *A variable in a specification is said to be accessible if it is a method parameter or a special variable, i.e. root or res.*

Definition 2.2 (Reachable) *Given a heap constraint κ such that $\kappa = p :: c\langle v^* \rangle * \kappa_1$, node $p :: c\langle v^* \rangle$ is reachable from a variable q iff they satisfy the following relation:*

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

Definition 2.3 (Well-Formed Constraint) *A separation constraint Φ is well-formed if (i) all objects and predicate instances in the constraint are reachable from accessible variables, (ii) Φ is in a disjunctive normal form $\bigvee (\exists v^* \cdot \kappa \wedge \pi)^*$ where κ denotes heap constraints, π pointer, arithmetic and set/bag constraints.*

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 proving procedure to correctly match nodes from the consequent with nodes from the antecedent of an entailment relation. Furthermore, arbitrary recursive shape relation can lead to non-termination in unfold/fold reasoning. We avoid this pitfall by using only *well-founded* shape predicates:

Definition 2.4 (Well-Founded Predicate) *A shape predicate is said to be well-founded if its body satisfies three conditions, namely: (i) it is a well-formed constraint, (ii) the parameter root may not be bound to a predicate instance, (iii) there is at most one object in each disjunct.*

3 Multiple Pre/Post Specifications

Our approach currently expects pre/post conditions to be specified for each method. With the rich variety of shapes that can be specified, there are often multiple ways of viewing a method’s intended behaviour. In this section, we explore the contexts in which multiple pre/post specifications are useful for heap-manipulating methods. For simplicity, we shall illustrate via our running `append` method, that is reproduced below.

```
void append(node x, node y)
{ if x.next ≠ null then { append(x.next, y) }
  else { x.next := y } }
```

Firstly, multiple pre/post specifications are helpful for capturing *strong aliasing* properties in separation logic. Our previous specification was for two disjoint input lists. We can also provide a pre/post specification in which `x` and `y` parameters are aliased, as follows:

```
requires x :: lseg⟨null, n⟩ ∧ n > 0 ∧ x = y
ensures x :: clist⟨n⟩
```

The aliasing by `x=y` causes the `append` function to return a circular list instead! To successfully verify the recursive call of `append`, we also require another pre/post specification that returns a list segment, as follows:

```
requires x :: lseg⟨null, n⟩ ∧ n > 0
ensures x :: lseg⟨n, y⟩
```

Secondly, we may require pre/post specification to handle *missing cases* that are absent in our predicates. Our earlier example made use of a non-empty `sortl` predicate which was essential for the first `x` parameter, but not required for the second `y` parameter. To allow the `y` parameter to be possibly empty, we shall provide an extra pre/post specification, as follows:

```
requires x::sortl⟨n1, i1, j1⟩ ∧ y=null
ensures x::sortl⟨n1, i1, j1⟩
```

Thirdly, we may use smaller predicates to specify additional properties of our method. This is helpful towards better functional correctness and also *scalable verification*. For example, we can further introduce another predicate that captures the reachability of a linked-list as a bag of its values, as follows:

$$\text{llR}\langle B \rangle \equiv (\text{root}=\text{null} \wedge B=\{\}) \vee (\exists v, B_1 \cdot \text{root}::\text{node}\langle v, r \rangle * r::\text{llR}\langle B_1 \rangle \wedge B=\{v\} \cup B_1)$$

This new predicate captures the reachable values in a linked-list but requires a bag constraint solver. It is useful as it can help specify the reachability property by the `append` method, as captured by the following extra pre/post specification:

```
requires x::llR⟨B1⟩ * y::llR⟨B2⟩ ∧ x≠null
ensures x::llR⟨B1 ∪ B2⟩
```

Multiple pre/post specifications allow us to freely use smaller predicates whereby different properties can be separately proven. As illustrated later by our experiments, this can help support scalable verification.

As stated before, the concept of *multiple pre/post* specification is not new. Conceptually, it is also related to the notion of intersection types [11], especially when it is compared with a dependent type system with effects on heap states. While the concept of multiple pre/post is not new, its use in separation logic is novel and provides a fresh challenge for automated verification.

In this paper, we propose the concept of *set of states* to handle multiple pre/post specification, and as a means towards systematic proof search. Our solution avoids the need to transform multiple pre/post specification into a single pre/post specification. We chose this path for the following reasons: (i) it is more concise, (ii) it allows pre/post conditions to be decomposed in a modular fashion, (iii) it integrates well with set of states, and (iv) its use results in smaller heap states. As smaller heap states are likely to be faster to verify, we expect multiple pre/post to give better support for scalable verification, whereby smaller predicates are strongly encouraged. This facilitates modular reuse of the verification processes, in addition to support for systematic proof search. The next two sections highlight the techniques we have formulated to support multiple pre/post specification in separation logic, and some experiments we conducted to validate its utility.

4 Our Approach

In our approach to verification, we expect users to supply two things that can assist in automatic verification of their code, namely shape predicate definitions and multiple pre/post specifications for each method. Our system has two major components, namely (i) Hoare-style verifier and (ii) entailment prover for separation logic, as shown in Figure 2.

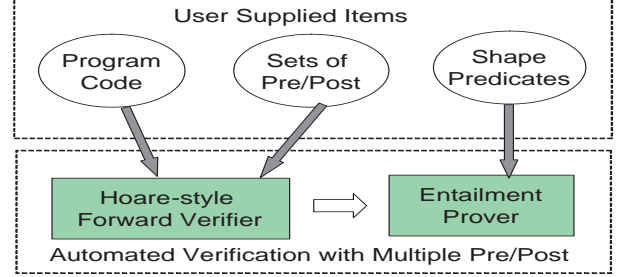


Figure 2. Our Approach to Verification

4.1 Forward Verifier

Our Hoare-style verifier is defined for an imperative object-based language with syntax in Figure 3. A program P consists of declarations $tdecl$ and methods $meth$. Declarations can be shape predicates $spre$ d or object types $objt$. Each method is decorated with the specification $\{\Phi_{pr}^i \mapsto \Phi_{po}^i\}_{i=1}^p$ that is made up of a collection of pre- and post-condition pairs. This more precise syntax for multiple pre/post pairs shall be used in our formalization. The intended meaning is that whenever the method is called in a program state satisfying precondition Φ_{pr}^i , and if the method terminates, the resulting state will satisfy the corresponding postcondition Φ_{po}^i . We handle `while` loop in a similar way. Other constructs are standard.

Hoare-style code verifier is based on forward rules of the form $\{\Delta_1\}code\{\Delta_2\}$, where Δ_1 is a prestate in separation formula that is given, while Δ_2 is a poststate that shall be computed by forward reasoning. To capture proof search, we generalize the forward rule to the form $\{\Delta\}code\{S\}$ where S is a set of heap states, discovered by a search-based

P	$::= tdecl^* meth^*$	$tdecl ::= objt \mid spre$
$objt$	$::= data \ c \ \{ field^* \}$	$field ::= t \ v$
t	$::= c \mid \tau$	$\tau ::= int \mid bool$
$meth$	$::= (t \mid void) \ mn \ ((t \ v)^*) \ where \ mspec \ \{e\}$	
$mspec$	$::= \{\Phi_{pr}^i \mapsto \Phi_{po}^i\}_{i=1}^p$	
e	$::= null \mid k^\tau \mid v \mid v.f \mid v:=e \mid v_1.f:=v_2$ $\mid new \ c(v^*) \mid e_1; e_2 \mid t \ v; \ e \mid mn(v^*)$ $\mid if \ v \ then \ e_1 \ else \ e_2$ $\mid while \ v \ where \ mspec \ do \ e$	
c, v	$::= identifiers$	

Figure 3. A Core Imperative Language

verification process. When S is empty, the forward verification is said to have failed for Δ as prestate. For convenience, we also provide lifted variant of the forward verifier to take a set of prestates. Verification in such a case succeeds if any of the prestates gives rise to a successful verification, that is if at least one of the S_i is non-empty. This rule is useful when the forward verifier has processed at least one sub-expression, potentially giving rise to a set of residual states.

$$\frac{\forall i \in 1..n. \{\Delta_i\} \text{ code } \{S_i\}}{\vdash \{\{\Delta_1, \dots, \Delta_n\}\} \text{ code } \{\bigcup_{i=1}^n S_i\}}$$

Verification of a method starts with each precondition, and proves that the corresponding postcondition is guaranteed at the end of the method. The verification is formalized in the following rule:

$$\frac{\begin{array}{c} \text{[FV-METH]} \\ V = \{v_1..v_n\} \quad W = \text{prime}(V) \\ \forall i = 1, \dots, p. (\vdash \{\Phi_{pr}^i \wedge \text{nochange}(V)\} e \{S_i^i\} \\ (\exists W \cdot S_1^i) \vdash \Phi_{po}^i * S_2^i \quad S_2^i \neq \{\}) \end{array}}{\vdash t_0 \text{ mn}(t_1 v_1, \dots, t_n v_n) \text{ where } \{\Phi_{pr}^i \rightsquigarrow \Phi_{po}^i\}_{i=1}^p \{e\}}$$

The function $\text{prime}(V)$ returns $\{v' \mid v \in V\}$. The predicate $\text{nochange}(V)$ returns $\bigwedge_{v \in V} (v = v')$. If $V = \{\}$, $\text{nochange}(V) = \text{true}$. $\exists W \cdot S$ returns $\{\exists W \cdot S_i \mid S_i \in S\}$. The entailment $(\exists W \cdot S_1^i) \vdash \Phi_{po}^i * S_2^i$ is discharged by the entailment prover described in the next subsection.

At a method call, each of the method's precondition is checked. The combination of the residue S_i and the postcondition is added to the poststate. If a precondition is not entailed by the program state Δ , the corresponding residue is not added to the set of states. The test $S \neq \{\}$ ensures that at least one precondition is satisfied.

$$\frac{\begin{array}{c} \text{[FV-CALL]} \\ t_0 \text{ m}(t_1 v_1, \dots, t_n v_n) \text{ where } \{\Phi_{pr}^i \rightsquigarrow \Phi_{po}^i\}_{i=1}^p \{e\} \in P \\ \rho = [v_j'/v_j]_{j=1}^n \quad \Delta \vdash \rho \Phi_{pr}^i * S_i \quad \forall i = 1, \dots, p \\ S = \bigcup_{i=1}^p S_i * \Phi_{po}^i \quad S \neq \{\} \end{array}}{\vdash \{\Delta\} \text{ m}(v_1..v_n) \{S\}}$$

Note that the verification rule also invokes the entailment prover to discharge $\Delta \vdash \rho \Phi_{pr}^i * S_i$, where ρ represents a substitution of v_j by v_j' , for all $j = 1, \dots, n$. The lifted separation conjunction $*$ over a set (i.e., $S_i * \Phi_{po}^i$) is defined in Fig. 5 in the next subsection.

Our verifier also ensures that each field access is safe from null dereferencing. This is shown in the field access rules in Fig. 4 which also includes other forward verification rules for the language. The verification rules attempt to track heap states, as accurately as possible, with path-sensitivity captured by [FV-IF] rule, flow-sensitivity by [FV-SEQ] rule and context sensitivity by the [FV-CALL] rule. In a nutshell, verification is carried out at three places. For each call site, the [FV-CALL] rule (mentioned earlier) ensures that at least one of its method's preconditions is

satisfied. At each method definition, the [FV-METH] rule checks that every postcondition holds for the method body assuming its respective precondition. At each shape definition, [FV-SPRED] checks that its given invariant π_{inv} is sound w.r.t. (i.e. semantic consequence of) the well-formed heap formula Φ . (The rule for while loop is omitted but is essentially similar to the mechanics for handling tail-recursive methods.) The function $XPure_0(\Phi)$ generates a sound and heap-independent approximation of the heap constraint Φ . For instance,

$$\begin{array}{l} XPure_0(x::\text{node}(_, _)) \equiv x > 0 \\ XPure_0(x::\text{node}(_, _) * y::\text{node}(_, _)) \equiv x > 0 \wedge y > 0 \wedge x \neq y \\ XPure_0(x::\text{lseg}(p, n)) \equiv n \geq 0 \end{array}$$

For the shape predicate case above, we can get a more precise approximation by unrolling the predicate definition once, for example:

$$XPure_1(x::\text{lseg}(p, n)) \equiv (x = p \wedge n = 0 \vee x > 0 \wedge n > 0)$$

The definition for the general approximation procedure $XPure_n(\Phi)$ (also used in the entailment prover) can be found in [10] where n denotes the number of unrollings done on the shape predicates.

The operators $\wedge_{\{v\}}$ (in assignment rule) and $*_W$ (in while rule) are *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:

$$\begin{array}{l} \Delta_1 \oplus_X \Delta_2 =_{df} \exists r_1..r_n. \rho_1 \Delta_1 \oplus \rho_2 \Delta_2 \\ \text{where } r_1, \dots, r_n \text{ are fresh variables;} \\ \rho_1 = [r_i/x_i]_{i=1}^n; \rho_2 = [r_i/x_i]_{i=1}^n \end{array}$$

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 $*$.

Normalization rules for separation constraints and lifted operators over sets of states are given in Fig. 5.

4.2 Entailment Prover

The other major component of our system is the entailment prover. This prover directly uses the definitions of shape predicates to reason about entailment between two heap formulae. A key novelty is the use of a set of heap states to support non-deterministic entailment. By non-determinism, we mean a search process that returns multiple answers, any one of which indicates a successful verification. Our entailment prover has the form $\Delta_a \vdash \Delta_c * S$ where S is a set of possible residual poststates. The entailment succeeds when S is non-empty, otherwise it is deemed to have failed. When S captures multiple residual states, they signify different search outcomes during proving.

During entailment, each pair of aliased objects/predicates from Δ_a and Δ_c are matched up, whenever

$\frac{[FV-SPRED] \quad XPure_0(\Phi) \implies [0/null](\pi_{inv})}{\vdash c(v^*) \equiv \Phi \text{ inv } \pi_{inv}}$	$\frac{[FV-VAR] \quad S = \{\Delta \wedge \text{res} = v'\}}{\vdash \{\Delta\} v \{S\}}$	$\frac{[FV-CONST] \quad S = \{\Delta \wedge eq_\tau(\text{res}, k)\}}{\vdash \{\Delta\} k^\tau \{S\}}$	$\frac{[FV-LOCAL] \quad \vdash \{\Delta\} e \{S\}}{\vdash \{\Delta\} \{t v; e\} \{\exists v, v'. S\}}$
$\frac{[FV-IF] \quad \vdash \{\Delta \wedge v'\} e_1 \{S_1\} \quad \vdash \{\Delta \wedge \neg v'\} e_2 \{S_2\}}{\vdash \{\Delta\} \text{if } v \text{ then } e_1 \text{ else } e_2 \{S_1 \vee S_2\}}$	$\frac{[FV-NEW] \quad S = \{\Delta * \text{res} :: c(v'_1, \dots, v'_n)\}}{\vdash \{\Delta\} \text{new } c(v_1, \dots, v_n) \{S\}}$	$\frac{[FV-SEQ] \quad \vdash \{\Delta\} e_1 \{S_1\} \quad \vdash \{S_1\} e_2 \{S_2\}}{\vdash \{\Delta\} e_1; e_2 \{S_2\}}$	
$\frac{[FV-ASSIGN] \quad \vdash \{\Delta\} e \{S_1\} \quad S_2 = \exists \text{res}. (S_1 \wedge \{v\} v' = \text{res})}{\vdash \{\Delta\} v := e \{S_2\}}$	$\frac{[FV-FIELD-READ] \quad \Delta \vdash v' :: c(v_{1..n}) * S_1 \quad S_1 \neq \{\} \quad \text{fresh } v_1..v_n \quad S_2 = \exists v_1..v_n. (S_1 * v' :: c(v_{1..n}) \wedge \text{res} = v_i)}{\vdash \{\Delta\} v.f_i \{S_2\}}$	$\frac{[FV-FIELD-UPDATE] \quad \Delta \vdash v' :: c(v_{1..n}) * S_1 \quad S_1 \neq \{\} \quad \text{fresh } v_1..v_n \quad S_2 = \exists v_1..v_n. (S_1 * v' :: \{v'_0/v_i\} c(v_{1..n}))}{\vdash \{\Delta\} v.f_i := v_0 \{S_2\}}$	

Figure 4. Forward Verification Rules with Non-Determinism

$(\Delta_1 \vee \Delta_2) \wedge \pi \quad \rightsquigarrow (\Delta_1 \wedge \pi) \vee (\Delta_2 \wedge \pi)$	$(\exists x. \Delta_1) * \Delta_2 \quad \rightsquigarrow \exists y. ([y/x]\Delta_1 * \Delta_2)$
$(\Delta_1 \vee \Delta_2) * \Delta \quad \rightsquigarrow (\Delta_1 * \Delta) \vee (\Delta_2 * \Delta)$	$(S_1 \vee S_2) \quad \rightsquigarrow \{\Delta_1 \vee \Delta_2 \mid \Delta_1 \in S_1, \Delta_2 \in S_2\}$
$(\kappa_1 \wedge \pi_1) * (\kappa_2 \wedge \pi_2) \quad \rightsquigarrow (\kappa_1 * \kappa_2) \wedge (\pi_1 \wedge \pi_2)$	$F(S) \quad \rightsquigarrow \{F(\Delta) \mid \Delta \in S\}$
$(\kappa_1 \wedge \pi_1) \wedge (\pi_2) \quad \rightsquigarrow \kappa_1 \wedge (\pi_1 \wedge \pi_2)$	where
$(\gamma_1 \wedge \phi_1) \wedge (\gamma_2 \wedge \phi_2) \quad \rightsquigarrow (\gamma_1 \wedge \gamma_2) \wedge (\phi_1 \wedge \phi_2)$	$F(\mathcal{A}) ::= \mathcal{A} \wedge \pi \mid \mathcal{A} \wedge_W \pi \mid \mathcal{A} * \Delta \mid \mathcal{A} *_W \Delta \mid \exists x. \mathcal{A}$
$(\exists x. \Delta) \wedge \pi \quad \rightsquigarrow \exists y. ([y/x]\Delta \wedge \pi)$	y denotes fresh variable

Figure 5. Normalization Rules for Separation Constraints and with operators lifted to a Set

they are proven identical. The formal rule for matching is:

$$\frac{XPure_n(p_1 :: c(v_1^*) * \kappa_1 \wedge \pi_1) \implies p_1 = p_2 \quad \rho = [v_1^*/v_2^*] \quad \kappa_1 \wedge \pi_1 \wedge \text{freeEqn}(\rho, V) \vdash_{V - \{v_2^*\}} \rho(\kappa_2 \wedge \pi_2) * S}{p_1 :: c(v_1^*) * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c(v_2^*) * \kappa_2 \wedge \pi_2) * S}$$

Note that the complete form for the entailment relation is $\Delta_a \vdash_V^\kappa \Delta_c * S$ which denotes $\kappa * \Delta_a \vdash \exists V. (\kappa * \Delta_c) * S$. As mentioned earlier, the purpose of the entailment prover 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 set of possible residual heap states S (which is empty if the entailment search fails). 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 κ and V are derived. The entailment prover is invoked with $\kappa = \text{emp}$ and $V = \emptyset$. When a match occurs, the bindings between free variables from the matched node in the consequent and the corresponding variables from the antecedent are generated and kept in the antecedent via $\text{freeEqn}(\rho, V)$ which is defined as follows:

$$\text{freeEqn}(\{u_i/v_i\}_{i=1}^n, V) =_{df}$$

$$\text{let } \pi_i = (\text{if } v_i \in V \text{ then true else } v_i = u_i) \text{ in } \bigwedge_{i=1}^n \pi_i$$

If no immediate match can be identified, an unfold/fold operation may be invoked. The rule for unfolding given below is to unfold a predicate in the antecedent Δ_a so as to match up with an object in the consequent Δ_c :

$$\frac{XPure_n(p_1 :: c(v_1^*) * \kappa_1 \wedge \pi_1) \implies p_1 = p_2 \quad \text{IsPred}(c_1) \wedge \text{IsObj}(c_2)}{\text{unfold}(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * S) \quad p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * S}$$

The test $\text{IsPred}(c)$ (resp. $\text{IsObj}(c)$) returns true if c is defined as a shape predicate (resp. an object).

Alternatively, a predicate in the consequent that is aliased with an object in the antecedent is handled by folding. The folding rule given next is a recursive invocation of the entailment procedure for a predicate from the consequent, in order to identify a set of heap nodes that match with that predicate's definition. It is different from unfolding a predicate in the antecedent as we allow bindings on free variables to transfer into the antecedent at the end of folding. This is critical for capturing free variables from preconditions that are used in postconditions.

$$\frac{\text{IsPred}(c_2) \wedge \text{IsObj}(c_1) \quad \{(\Delta_i, \kappa_i^f, \pi_i^f)\}_{i=1}^n = \text{fold}^\kappa(p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1, p_2 :: c_2 \langle v_2^* \rangle) \quad XPure_n(p_1 :: c(v_1^*) * \kappa_1 \wedge \pi_1) \implies p_1 = p_2 \quad (\pi_i^a, \pi_i^c) = \text{split}_V^{\{v_2^*\}}(\pi_i^f) \quad \Delta_i \wedge \pi_i^a \vdash_V^{\kappa_i^f} \kappa_2 \wedge (\pi_2 \wedge \pi_i^c) * S_i}{p_1 :: c_1 \langle v_1^* \rangle * \kappa_1 \wedge \pi_1 \vdash_V^\kappa (p_2 :: c_2 \langle v_2^* \rangle * \kappa_2 \wedge \pi_2) * \bigcup_{i=1}^n S_i}$$

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 them based on free variables. For instance, let us consider that the parameters of the folded predicate are n and B , and the bindings introduced by the folding process are $n=0 \wedge B = \{1\}$. If B is a free variable and n is bound, then the split function differentiates the given constraints as follows:

$$\text{split}_{\{n\}}^{\{n, B\}}(n=0 \wedge B = \{1\}) = (B = \{1\}, n=0)$$

The binding $B = \{1\}$ on the free variable B is to be transferred to the antecedent, while $n=0$ will be kept in the consequent.

The folding operation requires a special version of entailment that returns three extra things: (i) consumed heap nodes, κ_i , (ii) existential variables used, V_i , and (iii) final consequent, π_i . 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 definition:

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

5 Implementation

We have built a prototype system using Objective Caml. The proof obligations generated by our verification are discharged by our entailment proving procedure with the help of Omega Calculator [12] for arithmetic constraints of Presburger form and MONA [5] for set constraints.

Our system uses the set-of-states technique to implement multiple pre/post specification, but this technique is orthogonal to multiple pre/post. Set-of-states may also arise from disjunction in the consequent from entailment proving, even when single pre/post specifications are used. Our first set of experiments, in Figure 6, was designed to evaluate the cost of supporting set-of-states. We provide the time taken (in seconds) to verify each of those programs. Most of our examples can be verified within 5 seconds despite the use of a Presburger solver. Verification time of a function includes time to verify all functions that it calls.

The examples show that proof search via set of states does not incur much overhead, since most of the time there is only one state. The average overhead introduced by the set of states for our examples is around 0.05 seconds. However, set of states is crucial to verify the `count` or `delete` methods of circular list, which is based on the `lseg` predicate, and in those examples that have been marked with a *failed* (to verify). From the experiments presented in Figure 6, we may conclude that the cost of supporting set of states is quite low. However, set of states may be crucial for some examples, as it may be required when handling disjunctive consequent or multiple pre/post.

Even if multiple pre/post conditions are not critically needed, there may be occasions when they are desirable due to the possibility of relying on smaller and simpler constraints during verification. In order to highlight the performance gains due to the use of multiple pre/post, Figure 7 compares the timings obtained for some programs with single pre/post (capturing the reachability property) that have been translated to equivalent programs with multiple pre/post. This translation exploits the use of smaller and simpler predicates, where possible, which resulted in faster overall verification. As the proof obligations generated for the examples in Figure 7 contain set constraints,

Programs	Timing (in secs) with Omega solver	
	set-of-states	one-state
Linked List	verifies size/length	
delete	0.08	0.06
reverse	0.06	0.05
Circular Linked List	verifies size + cyclic structure	
delete	0.2	<i>failed</i>
count	0.24	<i>failed</i>
Doubly Linked List	verifies size + double links	
append	0.16	0.12
flatten (from tree)	0.35	0.33
Sorted List	verifies size + min + max + sortedness	
delete	0.16	0.18
insertion_sort	0.5	0.48
selection_sort	0.37	0.33
merge_sort	0.74	0.72
quick_sort	0.82	0.82
AVL Tree	verifies size + height + height-balanced	
insert	5.06	5.00
Red-Black Tree	verifies size + black-height properties	
insert	1.53	1.39
delete	17.44	14.72
2-3 Tree	verifies height-balanced	
insert	24.41	<i>failed</i>
Perfect Tree	verifies perfectness	
insert	0.28	0.24
Complete Tree	verifies completeness	
insert	1.62	1.49

Figure 6. Comparing *one-state* vs *set-of-states*.

we used MONA to discharge them when needed. However, for some examples with single pre/post, the complexity of the constraints caused an out-of-memory error.

As an example of breaking a single specification into multiple pre/post, let us consider the following predicate which describes a non-empty list sorted in ascending order. The predicate tracks the length of the list, n , the minimum value, m , and with the entire set/bag of values stored in the list, B . The sorting property is ensured by $m \leq m_1$.

$$\begin{aligned} \text{sllm}\langle n, m, B \rangle &\equiv \text{root}::\text{node}\langle m, \text{null} \rangle \wedge B = \{m\} \wedge n = 1 \\ &\vee \exists p \cdot \text{root}::\text{node}\langle m, p \rangle * p::\text{sllm}\langle n-1, m_1, B_1 \rangle \\ &\wedge m \leq m_1 \wedge B = B_1 \cup \{m\} \quad \text{inv } n > 0 \end{aligned}$$

In order to transform an example using the above predicate into multiple pre/post, `sllm` can be split into three smaller predicates, each of them specialized on a certain property: `llr` which captures the set/bag of values and was already defined in Section 3, `sllm` which captures the sortedness and `ll` which captures the size.

$$\begin{aligned} \text{sllm}\langle m \rangle &\equiv \text{root}::\text{node}\langle m, \text{null} \rangle \\ &\vee \exists p \cdot \text{root}::\text{node}\langle m, p \rangle * p::\text{sllm}\langle m_1 \rangle \wedge m \leq m_1 \\ \text{ll}\langle n \rangle &\equiv \text{root} = \text{null} \wedge n = 0 \\ &\vee \exists p \cdot \text{root}::\text{node}\langle _, p \rangle * p::\text{ll}\langle n-1 \rangle \quad \text{inv } n \geq 0 \end{aligned}$$

By applying the above splitting to the insertion sort algorithm to obtain three smaller pairs of pre/post specification, we have managed to verify it in 3.96 seconds, while with a

Programs	Method	Timing (in seconds) with Omega & MONA	
		single pre/post	multiple pre/post
Linked List	append	2.68	0.42
	insert	0.95	0.4
	reverse	2.7	0.41
Circular Linked List	insert	0.47	0.22
	delete	2.51	0.67
Doubly Linked List	insert	1.51	0.55
	append	29.75	1.8
	flatten (from tree)	<i>out-of-mem</i>	5.49
Sorted List	insertion_sort	<i>out-of-mem</i>	3.96
	selection_sort	<i>out-of-mem</i>	2.6
Binary Search Tree	insert	<i>out-of-mem</i>	2.04
	delete	<i>out-of-mem</i>	4.7

Figure 7. Comparing single vs multiple pre/post

single pre/post we obtained an out-of-memory error.

6 Related Work

Recently, separation logic has been advocated to reason about heap-manipulating programs [2], but most of these early works support only a limited set of predicates. Our recent work [10] allowed size properties to be defined in user-supplied recursive predicates in separation logic, which are then automatically verified via a sound, terminating but incomplete proof system. Building on this work, the current paper advocates the use of multiple pre/post for better expressivity in method specifications using separation logic. We also propose a non-deterministic proof search procedure by carrying over sets of abstract states, which proves to be crucial in verifying certain complicated properties for heap manipulating programs.

On the inference front, Lee et al. [8] 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/bag properties. A recent work [4] has also formulated interprocedural shape inference but is restricted to just the list segment shape predicate. While our system does not perform inference of pre/post specification, we provide better support for automated verification via an expressive specification mechanism. For example, data structures with strong invariants, such as balanced heights and sortedness, are captured by our specification mechanism (with the help of multiple pre/post) prior to automatic verification.

7 Conclusion

We have introduced set of states as a way to support multiple pre/post specifications in separation logic. This approach allows proof search to be captured explicitly, permitting a simple automated search strategy to be developed. Our current search strategy is exhaustive but *directed* and is guaranteed to *terminate*.

Multiple pre/post method specifications enhances proof search in separation logic. This feature puts creative control back into users' hands. Nevertheless, we provide machine support for automatically checking and then applying each high-level specification from the user. We believe that multiple pre/post specifications can greatly enhance both the expressivity and performance of automated verification via separation logic. As shown by our initial experiments, it allows more verifications to be carried out successfully and with potential improvement to performance.

Acknowledgement

This work is supported by A*STAR research grant R-252-000-233-205, and Shengchao Qin is supported in part by EPSRC grant EP/E021948/1.

References

- [1] M. Barnett, K. R. M. Leino, and W. Schulte. The Spec# programming system: An overview. In *Int'l Workshop on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices*. Springer-Verlag, LNCS, 2004.
- [2] J. Berdine, C. Calcagno, and P. W. O'Hearn. Symbolic Execution with Separation Logic. In *APLAS*. Springer-Verlag, November 2005.
- [3] L. Burdy, Y. Cheon, D. R. Cok, M. D. Ernst, J. R. Kiniry, G. T. Leavens, K. R. M. Leino, and E. Poll. An overview of JML tools and applications. *Software Tools for Technology Transfer*, 2005.
- [4] A. Gotsman, J. Berdine, and B. Cook. Interprocedural Shape Analysis with Separated Heap Abstractions. In *SAS*. Springer LNCS, Seoul, Korea, August 2006.
- [5] Nils Klarlund and Anders Miller. Mona version 1.4 - user manual.
- [6] G. T. Leavens. JML's Rich, Inherited Specifications for Behavioral Subtypes. In *ICFEM*, Macao, China, November 2006. Springer-Verlag.
- [7] G. T. Leavens, E. Poll, C. Clifton, Y. Cheon, C. Ruby, D. Cok, P. Miller, and J. Kiniry. JML Reference Manual (DRAFT), February 2007.
- [8] O. Lee, H. Yang, and K. Yi. Automatic verification of pointer programs using grammar-based shape analysis. In *ESOP*. Springer Verlag, April 2005.
- [9] K.R.M. Leino and P. Muller. A verification methodology for model fields. In *15th ESOP*, March 2006.
- [10] 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.
- [11] Benjamin C. Pierce. *Programming with Intersection Types and Bounded Polymorphism*. PhD thesis, Carnegie Mellon University, 1991.
- [12] W. Pugh. The Omega Test: A fast practical integer programming algorithm for dependence analysis. *Communications of the ACM*, 8:102-114, 1992.