

Automatic Reasoning on Recursive Data Structures with Sharing

Duc-Hiep Chu
IST Austria
duc-hiep.chu@ist.ac.at

Joxan Jaffar
National University of Singapore
joxan@comp.nus.edu.sg

Abstract

We consider the problem of automatically verifying programs which manipulate arbitrary data structures. Our specification language is expressive, contains a notion of *separation*, and thus enables a precise specification of *frames*. The main contribution then is a program verification method which combines strongest postcondition reasoning in the form symbolic execution, unfolding recursive definitions of the data structure in question, and a new frame rule to achieve *local reasoning* so that proofs can be compositional. Finally, we present an implementation of our verifier, and demonstrate automation on a number of representative programs. In particular, we present the first automatic proof of a classic graph marking algorithm, paving the way for dealing with a class of programs which traverse a complex data structure.

1 Introduction

Formal reasoning about programs which manipulate dynamically allocated data structures is challenging; it is more so when data structures that (1) have “unrestricted” sharing, e.g., as in an arbitrary graph; and (2) are *recursive*, that is, their formal definition is represented by a recursively defined relation. A key technical difficulty is that “deep aliasing” prevents us from reasoning in isolation over parts of these data structures, a.k.a. *local reasoning*, because changes to one part of the structure (say, the left child of a graph) can affect other parts (the right child or its descendants) that may point into it. Consequently, there is no established systematic method than can *automatically* verify programs which manipulate such data structures, even for routine textbook graph algorithms.

In traditional Hoare logic, there is a *frame rule* (CFR), which allows an assertion that does not mention heap variables or pointers, to be “framed” through a program fragment. Augmenting this rule to accommodate recursively defined predicates has been used from as early as 1982 [5, 23].

Proposition 1.1 (Classic Frame Rule).

$$\frac{\{\phi\} P \{\psi\}}{\{\phi \wedge \pi\} P \{\psi \wedge \pi\}} \text{Mod}(P) \cap FV(\pi) = \emptyset \quad (\text{CFR})$$

where $\text{Mod}(P)$ denotes the variables that P modifies, and $FV(\pi)$ denotes the free variables of π . \square

It was Separation Logic [26, 31] (SL) which made a

$$\frac{\{\phi\} P \{\psi\}}{\{\phi * \pi\} P \{\psi * \pi\}} \quad (\text{SFR})$$

significant advance in verifying programs that manipulate recursive data structures such as linked-lists or trees. Two key ideas here are: associating a predicate with a notion of *heap*, and composing predicates with the notion of *separating conjunction* of heaps. As a result, SL has an extremely elegant frame rule (SFR): when a program fragment P is “enclosed” in some heap, then any formula π whose “footprint” is separate from this heap can be “framed” through the fragment. This notion of separation is indicated by the “separating conjunction” operator “ $*$ ” which states that the footprints of its two operands (which are logical predicates) are disjoint. Note that the validity of the triple $\{\phi\} P \{\psi\}$ entails that all heap accesses in P , read or write, are confined to the implicit heap of ϕ , or to fresh addresses. This provides for truly local reasoning, because the proof of P is done *without any prior knowledge* about the predicate π .

There have been significant advancements in automating SL and its variants to deal with recursive data-structures (without sharing) such as linked-lists or trees [8, 9, 27–30]. However, we cannot directly use SL frame rule to deal with programs that manipulate data structures with sharing because such structures cannot easily be massaged into the form $\phi * \pi$: for example, the left and right descendants of a graph node often are not disjoint.

The recent work [13] was an important contribution toward the sharing problem in SL. It introduced a new “ramification” rule which allowed for shared structures. However, the preconditions to using this rule, including the use of “magic wand” operators, are complex. Consequently, though we now have a systematic approach to formally reason about shared structures, we are not yet closer to an *automated* method. We elaborate on [13] in Section 8.

After SL, we have the method of *dynamic frames* [18] (DF), and later, the refinement to *implicit* dynamic frames (IDF) [35]. Essentially, a DF is a mathematical expression describing a (super)set of memory locations that a method refers to, a.k.a. the “footprint” of the method. It is used in the specification phase of verification: together with the provided pre and postcondition of a method, the DF represents the “modifies” property of the method. The big advantage of DF/IDF over SL is expressiveness, where frames can be specified at a higher resolution. However, while DF/IDF serve to automate local reasoning, they do not directly address *recursive data structures with sharing*. A key reason for this is that recursion typically embeds *reachability* properties, and such

properties are not typically accommodated by SMT solvers that are often used in DF/IDF systems. As circumstantial evidence for this, note that a challenge problem, marking a graph, was first systematically proved, albeit manually, in [13]. (We will show later a first systematic automatic proof.) We elaborate on DF/IDF in Section 8.

In this paper, we propose a framework that enables *fully automatic* local reasoning on recursive data structures with sharing. The key features of our framework that account for the new level of automation are as follows:

Strongest Postcondition Transform: Our assertion language, from [11], contains explicit subheaps definable by predicates, and separation is expressed by subheap disjointness. Thus we can use traditional conjunction of formulas, as opposed to separating conjunction. Again from [11], we now have a *strongest postcondition* transform (adapted and presented in Section 4) of heap formulas. What is new in this paper is to couple the transform with the reasoning of recursive predicates, and importantly, frame reasoning.

Framing of Recursive Predicates with CFR: We use a distinguished heap variable \mathcal{M} to represent the global heap memory, and extend [11] by forcing all subheaps appearing in a recursive predicate and its definition to be *ghost* variables. A subheap \mathcal{H} is explicitly constrained to be part of \mathcal{M} via a “heap reality” constraint $\mathcal{H} \sqsubseteq \mathcal{M}$. (E.g., an assertion stating a linked-list rooted at x will be $\text{list}(\mathcal{H}, x) \wedge \mathcal{H} \sqsubseteq \mathcal{M}$, instead of just $\text{list}(x)$ as in SL.) An important consequence is that, except for heap reality constraints, other constraints, e.g. $\text{list}(\mathcal{H}, x)$, are applicable to the classical frame rule (CFR).

Framing of Heap Reality Constraints: Our main contribution is a new frame rule to address framing of heap reality constraints (Section 5). Intuitively, if all the updates of the program fragment P are confined within the heap \mathcal{H}_1 (or to fresh addresses), then the heap reality constraint $\mathcal{H}_2 \sqsubseteq \mathcal{M}$ can be framed through P if \mathcal{H}_1 and \mathcal{H}_2 are separate. In other words, our frame rule is as simple to apply as the frame rule in SL. However, to achieve that we need to develop the concepts of heap “enclosure” and heap “evolution” and allow named subheaps to be explicitly nominated as frames in the specifications (of functions).

Finally, we give evidence that our verification framework has a good level of automation. In Section 6, we automatically prove one significant example for the first time: marking a graph. This example exhibits important relationships between data structures that have so far not been addressed by automatic verification: processing recursive data structures with sharing. We will present an implementation in Section 7, submitted as supplementary material for this paper, and a demonstration of automatic verification on a number of representative programs. We demonstrate the phases of specification, verification condition generation and finally theorem-proving. We stress here that we shall be using *existing* and not custom technology for the theorem-proving.

2 A Challenge Problem

In Fig. 1 we have a textbook algorithm that marks a graph. We assume a node has two successor fields `left` and `right`.

```
struct node { int m; struct node *left, *right; };

void mark(struct node *x) {
  if (!x || x->m) return;
  struct node *l = x->left, *r = x->right;
  x->m = 1; mark(l); mark(r);
}
```

Figure 1. Mark Graph Example

There are some subtle but critical points that makes the example extremely challenging. First, in order to have a well-founded recursive definition of a graph, we need some form of “history”. Yet the program itself *does not implement* any such notion. Instead it uses the mark (`m`) field for termination. Thus a challenge is to *connect* a node’s history in the specification and its mark.

A second critical point is: what is the *precise specification* of the function? It is clear that we eventually want a fully marked graph, but because the program is recursive, the pre and postcondition also act as *invariants* and are somewhat complicated. Clearly the precondition cannot be an arbitrary graph, otherwise the program may *prematurely* terminate. One intuitive precondition is that the graph is “mark successor closed”, i.e., any successor node of a marked node is itself marked. This concept covers both fully unmarked graphs as well as fully marked graphs. However, this intuitively appealing condition is, surprisingly, *too strong*. To prove this, consider the initially unmarked graph in Fig. 2 and we start the mark function at the root node 0. During processing, at the start of the second visit to node 0 (by going `left` and then `left` again) the graph no longer satisfies the proposed precondition because while 0 has been marked, one of its successors, node 2, has not yet been marked.

The actual required precondition is rather complicated, and we relegate the details to Section 6. Here it suffices to say that the precondition must state that every encountered marked node is either previously encountered, or all of its successor nodes are already marked. The take-away is that this property is not naturally expressible without using a recursive definition.

The third critical point is to achieve local reasoning. We must ensure that the first recursive call must not negate what is needed as the precondition of the second call and dually the second call should not negate the effects of the first. In other




Figure 2. A Cyclic Graph

words, we need to describe: (1) the write footprint of the first call, (2) the footprint of the precondition of the second call, (3) the footprint of the postcondition of the first call, and (4) the write footprint of the second call. The verification process needs to automatically figure out that (1) and (2) are disjoint and that (3) and (4) are also disjoint.

3 The Assertion Language

We assume a vanilla imperative programming language with functions but no loops (which are tacitly compiled into tail-recursive functions). Other than standard non-heap statements, the following are *heap manipulation statements*:¹

- sets x to be the value pointed to by y : $x = *y$;
- sets the value pointed to by x to be y : $*x = y$;
- points x to a freshly cell: $x = \mathbf{malloc}(1)$;
- deallocates the cell pointed to by x : $\mathbf{free}(x)$.

The heap is not explicitly mentioned in the program. Instead, it is dereferenced using the “*” notation as in the C language. (Not to be confused with the operator “*” in SL or our heap constraint language.) Since our later discussion will involve symbolic execution, we also assume that branch condition is translated to *assume*(_) statement. We first give a brief overview of Hoare and Separation Logic, then we introduce the language in [11].

3.1 Background

Hoare Logic [12] is a formal system for reasoning about program correctness. Hoare Logic is defined in terms of axioms over *triples* of the form $\{\phi\} P \{\psi\}$, where ϕ is the *precondition*, ψ is the *postcondition*, and P is some code fragment. Both ϕ and ψ are formulae over the *program variables* in P . The meaning of the triple is as follows: for all program states σ_1, σ_2 such that $\sigma_1 \models \phi$ and executing σ_1 through P derives σ_2 , then $\sigma_2 \models \psi$. For example, the triple $\{x < y\} x = x + 1 \{x \leq y\}$, x and y are integers, is *valid*. Note that under this definition, a triple is automatically valid if P is non-terminating or has undefined behavior. This is known as *partial correctness*.

Separation Logic (SL) [31] is a popular extension of Hoare Logic [12] for reasoning over *heap manipulating programs*. SL extends predicate calculus with new logical connectives – namely *empty heap* (**emp**), *singleton heap* ($p \mapsto v$), and *separating conjunction* ($F_1 * F_2$) – such that the structure of assertions reflects the structure of the underlying heap. For example, the precondition in the following valid Separation Logic triple

$$\{x \mapsto _ * y \mapsto 2\} *x = *y + 1 \{x \mapsto 3 * y \mapsto 2\}$$

represents a heap comprised of two *disjoint singleton* heaps, indicating that both x and y are *allocated* and that location y points to the value 2. In the postcondition, x points to

¹We assume (de)allocation of single heap cells; this can be easily generalized, and indeed so in our implementation.

value 3, as expected. SL also allows *recursively-defined* heaps for reasoning over data structures, such as **list** and **tree**. An SL *triple* $\{\phi\} P \{\psi\}$ additionally guarantees that any state satisfying ϕ will not cause a memory access violation in P . For example, the triple **{emp}** $*x := 1 \{x \mapsto 1\}$ is *invalid* since x is a dangling pointer in a state satisfying the precondition.

A Constraint Language of Explicit Heaps [11]: We have a set of Values (e.g. integers) and we define Heaps to be all *finite partial maps* between values, i.e., $\text{Heaps} \stackrel{\text{def}}{=} (\text{Values} \rightarrow_{\text{fin}} \text{Values})$. There is a special value *null* (“null” pointer) and a special heap **emp** (“empty” heap). Where \mathcal{V}_v and \mathcal{V}_h denote the sets of value and heap variables respectively, our *heap expressions* HE are as follows:

$$\begin{aligned} H &::= \mathcal{V}_h & v &::= \mathcal{V}_v \\ HE &::= H \mid \mathbf{emp} \mid (v \mapsto v) \mid HE * HE \end{aligned}$$

An *interpretation* \mathcal{I} maps \mathcal{V}_h to Heaps and \mathcal{V}_v to Values. Syntactically, a *heap constraint* is of the form $(HE \simeq HE)$. An interpretation \mathcal{I} satisfies a heap constraint $(HE_1 \simeq HE_2)$ iff $\mathcal{I}(HE_1) = \mathcal{I}(HE_2)$ are the same heap, and the separation properties within HE_1 and HE_2 hold.

Let $\text{dom}(H)$ be the *domain* of the heap H . As in [11], heap constraints can be normalized into three basic forms:

$$\begin{aligned} H &\simeq \mathbf{emp} \text{ (EMPTY)} & H &\simeq (p \mapsto v) \text{ (SINGLETON)} \\ H &\simeq H_1 * H_2 \text{ (SEPARATION)} \end{aligned}$$

where $H, H_1, H_2 \in \mathcal{V}_h$ and $p, v \in \mathcal{V}_v$. Here (EMPTY) constrains H to be the empty heap (i.e., $H = \emptyset$ as a set), (SINGLETON) constrains H to be the singleton heap mapping p to v (i.e., $H = \{(p, v)\}$ as sets), and (SEPARATION) constrains H to be the heap that is partitioned into two disjoint sub-heaps H_1 and H_2 (i.e., $H = H_1 \cup H_2$ as sets and $\text{dom}(H_1) \cap \text{dom}(H_2) = \emptyset$).

We will also use *sub-heap relation* ($H_1 \sqsubseteq H_2$), *domain membership* ($p \in \text{dom}(H)$), and (overloaded) for brevity, *separation relation* ($H_1 * H_2$). In fact, writing $H_1 \sqsubseteq H_2$ is equivalent to $H_2 \simeq H_1 * _, p \in \text{dom}(H)$ to $H \simeq (p \mapsto _) * _, p \notin \text{dom}(H)$ to $_ \simeq H * (p \mapsto _)$, and $H_1 * H_2$ to $_ \simeq H_1 * H_2$; where the underscore in each instance denotes a fresh variable.

Finally, we have a *recursive constraint*. This is an expression of the form $p(h_1, \dots, h_n, v_1, \dots, v_m)$ where p is a user-defined *predicate symbol*, the $h_i \in \mathcal{V}_h, 0 \leq i \leq n$ and the $v_j \in \mathcal{V}_v, 0 \leq j \leq m$. Associated with such a predicate symbol is a *recursive definition*. We use the framework of *Constraint Logic Programming* (CLP) [15] to inherit its syntax, semantics, and its built-in notions of unfolding rules, for realizing recursive definitions. The *semantics* of a set of rules is traditionally known as the “least model” semantics [15]. For brevity, we only informally explain the language. The following constitutes a recursive definition of **list**(h, x), specifying a *skeleton list* in the heap h rooted at x .

$$\begin{aligned} \mathbf{list}(h, x) &:- h \simeq \mathbf{emp}, x = \mathbf{null}. \\ \mathbf{list}(h, x) &:- h \simeq (x \mapsto y) * h_1, \mathbf{list}(h_1, y). \end{aligned}$$

Note that the comma-separated expressions in the body of each rule is either *value constraint* (e.g. $x = \mathbf{null}$), a heap

constraint (e.g. $h \simeq \text{emp}$), or a recursive constraint (e.g. $\text{list}(h_1, y)$). In this paper, our value (i.e. “pure”) constraints will either be arithmetic or basic set constraints over values.

3.2 Program Verification with Explicit Heaps

Hoare Triples: We first define an *assertion* A as a formula over $\mathcal{V}_v, \mathcal{V}_h$:

$$A ::= VF \mid HF \mid RC \mid A \wedge A \mid A \vee A$$

where VF , HF , and RC are value, heap, and recursive constraints, respectively. We next connect the interpretation of assertions with the program semantics.

Programs operate over an unbounded set of *program variables* \mathcal{V}_p , which are the *value variables*. Thus $\mathcal{V}_p \subseteq \mathcal{V}_v$. We use one distinguished heap variable $\mathcal{M} \in \mathcal{V}_h$ to represent the *global heap memory*. Variables other than the program variables and \mathcal{M} may appear in assertions; they are existential or *ghost* variables. A ghost variable of type heap will be called a *subheap*.

The subheaps serve two essential and distinct purposes: (a) to describe subheaps of the global heap \mathcal{M} at the current program point, and (b) to describe some other “existential” heap. A common instance of (b) is the heap corresponding to the global heap at some *other* program point in the past.

We use the terminology “ghost heap” in accordance to standard practice that subheaps are existential, but in assertions, they can be used to constrain the value of the global heap. Importantly, as ghost variables, their values *cannot be changed* by the program. We will see later that this is important in practice because (a) predicates in assertions often need to be defined only using ghost subheaps, and (b) it is automatic that these predicates can be “framed through” any program fragment P because P cannot change the value of a ghost variable.

Note that we shall present rules that define recursive constraints using fresh variables. Notationally, for heaps, we shall use the small letter ‘h’ in rules, while using the large letter \mathcal{H} in assertions. Also, we use “,” in assertions as shorthand for logical conjunction.

Example: see the annotated program and the definition of inc_list in Fig. 3. The program increments all the data values in an acyclic list by 1.

```

struct node { int data; struct node *next; };

inc_list(h1, h2, x) :- h1 ≃ emp, h2 ≃ emp, x = null.
inc_list(h1, h2, x) :- h1 ≃ (x ↦ (d + 1, next)) * h'1,
    h2 ≃ (x ↦ (d, next)) * h'2, inc_list(h'1, h'2, next).

{ list( $\mathcal{H}, x$ ),  $\mathcal{H} \sqsubseteq \mathcal{M}$  }
  y = x; while (y) { y->data += 1; y = y->next; }
{ inc_list( $\mathcal{H}_1, \mathcal{H}, x$ ),  $\mathcal{H}_1 \sqsubseteq \mathcal{M}$  }

```

Figure 3. Incrementing data values in an acyclic list.

The recursive constraint $\text{list}(\mathcal{H}, x)$ describes a heap \mathcal{H} which houses an acyclic list rooted at x . The constraint $\mathcal{H} \sqsubseteq \mathcal{M}$ states that it resembles a part of the global heap. The other recursive constraint $\text{inc_list}(\mathcal{H}_1, \mathcal{H}, x)$ similarly defines that x is the head of a list resides in the heap \mathcal{H}_1 . It has another argument, the ghost heap \mathcal{H} , which also appears in the precondition. This, importantly, allows us to consider the triple as a *summary*, relating values in the precondition and postcondition (using the ghost variable as an anchor value). In this case, we are stating that the final list elements are one bigger than the corresponding initial elements. Further, we are also stating that all the links (the *next* pointers) are not modified.

We conclude this section by stressing that our *interpretation* of triples follows Hoare logic: the postcondition holds *provided* the start state satisfies the precondition, *and* there is a terminating execution of the program. In contrast, in SL, a triple entails that the program is memory-safe. Though we do not provide for memory safety as an intrinsic property, we can easily enforce memory safety, e.g., by asserting that dereferences (e.g., $x \rightarrow \text{next}$) and deallocations (e.g., $\text{free}(x)$), have their arguments (x) pointing to a valid cell in the global heap ($x \in \text{dom}(\mathcal{M})$). Not enforcing memory safety up front is not a weakness. It allows us to be flexible enough to perform reasoning even when memory safety is not the property of interest. Furthermore, SL may disapprove of a memory safe program whose specifications of some functions are not sufficiently complete. In contrast, our framework can still proceed, but possibly not by means of local reasoning.

4 Symbolic Execution with Explicit Heaps

Symbolic execution of a program uses *symbolic values* as inputs, and can be used for program verification in a standard way. We start with a precondition. The output of symbolic execution on a program path is a formula representing the symbolic state obtained at the end of a path, or the *strongest postcondition* of the precondition. For a loop-free program with no function calls, symbolic execution facilitates verification by considering a disjunction of all such path postconditions, which must then imply the desired postcondition. With function calls (or loops), to achieve modular verification, we need a frame rule.

We now describe how to obtain a the strongest postcondition transform as in [11]. It suffices to consider only the four heap-manipulating primitives.

Proposition 4.1 (Strongest Postcondition). *In the following Hoare-triples, the postcondition shown is the strongest postcondition of the primitive heap operation with respect to a precondition ϕ .*

$\{ \phi \} x = \mathbf{malloc}(1) \{ \text{alloc}(\phi, x) \}$ (Heap allocation)
 $\{ \phi \} \mathbf{free}(x) \{ \text{free}(\phi, x) \}$ (Heap deallocation)
 $\{ \phi \} x = *y \{ \text{access}(\phi, y, x) \}$ (Heap access)
 $\{ \phi \} *x = y \{ \text{assign}(\phi, x, y) \}$ (Heap assignment)

where the auxiliary macros `alloc`, `free`, `access`, and `assign` expand as follows:

$\text{alloc}(\phi, x) \stackrel{\text{def}}{=} \mathcal{M} \simeq (x \mapsto v) * \mathcal{H} \wedge \phi[\mathcal{H}/\mathcal{M}, v_1/x]$
 $\text{free}(\phi, x) \stackrel{\text{def}}{=} \mathcal{H} \simeq (x \mapsto v) * \mathcal{M} \wedge \phi[\mathcal{H}/\mathcal{M}]$
 $\text{access}(\phi, y, x) \stackrel{\text{def}}{=} \mathcal{M} \simeq (y \mapsto x) * \mathcal{H} \wedge \phi[v/x]$
 $\text{assign}(\phi, x, y) \stackrel{\text{def}}{=} \mathcal{M} \simeq (x \mapsto y) * \mathcal{H}_1 \wedge \mathcal{H} \simeq (x \mapsto v) * \mathcal{H}_1 \wedge \phi[\mathcal{H}/\mathcal{M}]$

where \mathcal{H} and \mathcal{H}_1 are fresh heap variables, and v and v_1 are fresh value variables. The notation $\phi[x/y]$ means formula ϕ with variable x substituted for y . \square

```

{  $\mathcal{H}_{99} \simeq \mathcal{M}$  }
   $t_1 = *x;$ 
{  $\mathcal{M} \simeq (x \mapsto t_1) * \mathcal{H}_1, \mathcal{H}_{99} \simeq \mathcal{M}$  }
   $*x = t_1 + 1;$ 
{  $\mathcal{M} \simeq (x \mapsto t_1 + 1) * \mathcal{H}_1, \mathcal{H}_2 \simeq (x \mapsto t_1) * \mathcal{H}_1,$ 
   $\mathcal{H}_2 \simeq (x \mapsto t_1) * \mathcal{H}_1, \mathcal{H}_{99} \simeq \mathcal{H}_2$  }
   $\Downarrow$  // (simplification)
{  $\mathcal{M} \simeq (x \mapsto t_1 + 1) * \mathcal{H}_1, \mathcal{H}_{99} \simeq (x \mapsto t_1) * \mathcal{H}_1$  }
   $t_2 = *x;$ 
{  $\mathcal{M} \simeq (x \mapsto t_2) * \mathcal{H}_3, \mathcal{M} \simeq (x \mapsto t_1 + 1) * \mathcal{H}_1,$ 
   $\mathcal{H}_{99} \simeq (x \mapsto t_1) * \mathcal{H}_1$  }
   $*x = t_2 - 1;$ 
{  $\mathcal{M} \simeq (x \mapsto t_2 - 1) * \mathcal{H}_4, \mathcal{H}_5 \simeq (x \mapsto v) * \mathcal{H}_4,$ 
   $\mathcal{H}_5 \simeq (x \mapsto t_2) * \mathcal{H}_3, \mathcal{H}_5 \simeq (x \mapsto t_1 + 1) * \mathcal{H}_1,$ 
   $\mathcal{H}_{99} \simeq (x \mapsto t_1) * \mathcal{H}_1$  }
  
```

Figure 4. Demonstrating Symbolic Execution

We demonstrate the usefulness (and partly the correctness) of Proposition 4.1 with a simple example. Consider:

$\{\mathcal{H}_{99} \simeq \mathcal{M}\} *x += 1; *x -= 1; \{\mathcal{H}_{99} \simeq \mathcal{M}\}$

In other words, the heap is unchanged after an increment and then a decrement. We rewrite the program so that only one heap operation is performed per program statement; in Fig. 4 we show the rewritten program fragment together with the propagation of the formulas. (For brevity, we also perform a simplification step.) It is then easy to show that the final formula implies $\mathcal{H}_{99} \simeq \mathcal{M}$, by first establishing that $\mathcal{H}_1 \simeq \mathcal{H}_3 \simeq \mathcal{H}_4$ and $v = t_2 = t_1 + 1$. This example provides a program *summary* that the heap is the same before and after execution.

5 The Frame Rule

Recall the classic frame rule (CFR) from Section 1 where from $\{ \phi \} P \{ \psi \}$ we may infer $\{ \phi \wedge \pi \} P \{ \psi \wedge \pi \}$ with the side condition that P does not modify any free variable in

π . In our current setting where P now may contain heap references, this frame rule in fact *still* can be used if π only contains free heap variables that are *ghost*. However, because the global heap memory might be changed by P , what *cannot* be framed through with this rule, is the property that a ghost variable \mathcal{H} is consistent with the global heap memory \mathcal{M} , i.e., $\mathcal{H} \sqsubseteq \mathcal{M}$. We call such a property the “heap reality” of \mathcal{H} .

In Separation Logic, where heaps are of the main interest, a key step is that when a program fragment is “enclosed” in some heap, then any formula π whose “footprint” is separate from this heap can be framed through. Recall the SL frame rule (SFR) from Section 1 wherein the premise $\{ \phi \} P \{ \psi \}$ ensures that the implicit heap arising from the formula ϕ captures all the heap accesses, read or write, in the program fragment P . Therefore $\{ \phi * \pi \} P \{ \psi * \pi \}$ naturally follows.

In our setting of explicit heaps, the frame rule, suitably translated into this language, is simply *not valid* without additional machinery to ensure enclosure. The concept of enclosure is to have an explicit subheap (or subheaps) to contain the program heap updates. These updates are defined to be the cells that the program *writes to*, or *deallocates*. This is because the property $\mathcal{H} \sqsubseteq \mathcal{M}$, where \mathcal{H} is a ghost variable, is falsified *just in case* the program has written to or deallocated some cell in \mathcal{M} whose address is also in $\text{dom}(\mathcal{H})$. Thus, the heap reality of \mathcal{H} is lost. Note that `malloc` changes \mathcal{M} , but it does not affect what already in \mathcal{M} .

Definition 5.1 (Heap Update). Given an address value v , a *heap update* to location v is defined as a statement that either *writes to* or *deallocates* the location v . \square

Before formalizing our notion of “enclosure”, however, we first need a concept of heap “evolution”. Let us use the notation $\tilde{\mathcal{H}}$ to denote the union $\bigcup_i \mathcal{H}_i$ of a collection of subheaps $\mathcal{H}_1, \dots, \mathcal{H}_n, n \geq 2$. Thus for example, $\tilde{\mathcal{H}} \sqsubseteq \mathcal{M}$ simply abbreviates $\mathcal{H}_1 \sqsubseteq \mathcal{M} \wedge \dots \wedge \mathcal{H}_n \sqsubseteq \mathcal{M}$.

Definition 5.2 (Evolution). Given a valid triple $\{ \phi \} P \{ \psi \}$, we say that a collection $\tilde{\mathcal{H}}$ in ϕ , where $\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M}$, *evolves* to a collection $\tilde{\mathcal{H}}'$ in ψ , where $\psi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M}$, if for each model \mathcal{I} of ϕ , executing P from \mathcal{I} will result in \mathcal{I}' , such that for any (address) value $v, v \in (\text{dom}(\mathcal{I}(\mathcal{M})) \setminus \text{dom}(\mathcal{I}(\tilde{\mathcal{H}})))$ implies $v \notin \text{dom}(\mathcal{I}'(\tilde{\mathcal{H}}'))$.

We shall use the notation $\{ \phi \} P \{ \psi \} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ to denote such evolution. \square

Intuitively, $\{ \phi \} P \{ \psi \} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ means that the largest $\tilde{\mathcal{H}}'$ can be is $\tilde{\mathcal{H}}$ plus any new cells allocated by P , and minus any that are freed by P . Note also that because the triple is valid, \mathcal{I}' will be a model of ψ . One important usage of the evolution concept is as follows: any heap \mathcal{H}_i such that $\mathcal{H}_i * \tilde{\mathcal{H}}$ and $\mathcal{H}_i \sqsubseteq \mathcal{M}$ at the point of the precondition ϕ (i.e., before P is executed), \mathcal{H}_i will be separate from $\tilde{\mathcal{H}}'$ at the point of the postcondition (i.e., after P is executed).

Consider the (linked-list) node defined in Section 3 (Fig. 3) and the triple shown below.

```
{ list( $\mathcal{H}_1, x$ ),  $\mathcal{H}_1 \sqsubseteq \mathcal{M}$  }
  z = malloc(sizeof(struct node));
  z->next = x;
{ list( $\mathcal{H}'_1, z$ ),  $\mathcal{H}'_1 \sqsubseteq \mathcal{M}$  }
```

We say that \mathcal{H}'_1 is an *evolution* of \mathcal{H}_1 , or $\text{evolve}(\mathcal{H}_1, \mathcal{H}'_1)$, notationally. Now assume that the triple represents only a local proof (i.e., we are also interested in other parts of \mathcal{M}). How should we compose this local triple to obtain a new triple? Formally, we have the following:

$$\frac{\{\phi\} P \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')}{\{\phi \wedge \tilde{\mathcal{H}} * \mathcal{H}_0 \wedge \mathcal{H}_0 \sqsubseteq \mathcal{M}\} P \{\psi \wedge \tilde{\mathcal{H}}' * \mathcal{H}_0\}} \quad (\text{EV})$$

Theorem 5.3 (Propagation of Separation). *The rule (EV) is correct. \square*

Proof Sketch 1. Let \mathcal{I} be a model of ϕ that is also a model of $\tilde{\mathcal{H}} * \mathcal{H}_0 \wedge \mathcal{H}_0 \sqsubseteq \mathcal{M}$. Let \mathcal{I}' be the result of executing P from \mathcal{I} . For each address $v \in \text{dom}(\mathcal{I}'(\mathcal{H}_0))$, because \mathcal{H}_0 is a ghost variable, i.e., its domain is not affected by executing P , we also have $v \in \text{dom}(\mathcal{I}(\mathcal{H}_0))$. It follows that $v \in (\text{dom}(\mathcal{I}(\mathcal{M})) \setminus \text{dom}(\mathcal{I}(\tilde{\mathcal{H}})))$. Directly from the definition of evolution, we deduce $v \notin \text{dom}(\mathcal{I}'(\tilde{\mathcal{H}}'))$ must hold. As a result, \mathcal{I}' also satisfies $\tilde{\mathcal{H}}' * \mathcal{H}_0$. \square

We are now ready to describe our notion of enclosure. We wish to describe, given a program P and a heap collection $\tilde{\mathcal{H}}$ in a precondition description ϕ , that all heap updates (heap assignments or deallocations) in P , are confined to an evolution of $\tilde{\mathcal{H}}$. The following definition, intuitively, is about one aspect of memory-safety: the heap updates are safe.

Definition 5.4 (Enclose). Suppose we have a valid triple $T = \{\phi\} P \{_ \}$, $\tilde{\mathcal{H}}$ appears in ϕ , and that $\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M}$. We say $\tilde{\mathcal{H}}$ *encloses* all heap updates of P if for any model \mathcal{I} of ϕ and for any execution path of P of the form $P_1; s; P_2$ where s is a heap update to a location v , it follows that there exists $\tilde{\mathcal{H}}'$ s.t. $\{\phi\} P_1 \{_ \} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ and $v \in \text{dom}(\mathcal{I}'(\tilde{\mathcal{H}}'))$ hold, where \mathcal{I}' is the result of executing P_1 from \mathcal{I} .

We shall use the notation $T \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})$ to denote that $\tilde{\mathcal{H}}$ encloses all the updates of P wrt. T . \square

We now can introduce our frame rule. It is in fact all about “preserving the heap reality”. Recall that a recursive constraint, which satisfies the standard side condition and of which the heap variables are all ghost (and this is a common situation), *remains true* from precondition to postcondition. What may no longer hold in the postcondition is the heap reality of some \mathcal{H}_0 . That is, $\mathcal{H}_0 \sqsubseteq \mathcal{M}$ may hold at the precondition, but no longer so at the postcondition. In other words, given local reasoning for a code fragment P and the fact that $\mathcal{H}_0 \sqsubseteq \mathcal{M}$ holds before executing P , how would we preserve this heap reality, without the need to reconsider the code

fragment P ? Our answer is the following Hoare-style rule, our new frame rule:

$$\frac{\{\phi\} P \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}{\{\phi \wedge \tilde{\mathcal{H}} * \mathcal{H}_0 \wedge \mathcal{H}_0 \sqsubseteq \mathcal{M}\} P \{\psi \wedge \mathcal{H}_0 \sqsubseteq \mathcal{M}\}} \quad (\text{FR})$$

Theorem 5.5 (Frame Rule). *The rule (FR) is correct. \square*

Proof Sketch 2. We prove by contradiction. Assume it is not the case, meaning that there is model \mathcal{I} of ϕ that is also a model of $\tilde{\mathcal{H}} * \mathcal{H}_0 \wedge \mathcal{H}_0 \sqsubseteq \mathcal{M}$ and \mathcal{I}' is the result of executing P from \mathcal{I} , but \mathcal{I}' does not satisfy $\mathcal{H}_0 \sqsubseteq \mathcal{M}$. Thus there must be a cell ($v \mapsto _$) that belongs to $\mathcal{I}'(\mathcal{H}_0)$ but not $\mathcal{I}'(\mathcal{M})$. Because $\mathcal{I}(\mathcal{H}_0) \sqsubseteq \mathcal{I}(\mathcal{M})$, the fragment P must have updated the location v . Therefore, there must be an execution path of P which is of the form $P_1; s; P_2$, where s is a heap update to the location v . Let $\bar{\mathcal{I}}$ be the result of executing P_1 from \mathcal{I} . By the definition of enclosure, assume $\{\phi\} P_1 \{_ \} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ and $v \in \text{dom}(\bar{\mathcal{I}}(\tilde{\mathcal{H}}'))$ hold. By (EV) rule, we have $\bar{\mathcal{I}}$ satisfies $\tilde{\mathcal{H}}' * \mathcal{H}_0$. Since \mathcal{H}_0 is a ghost variable, its domain is not affected by executing P_1 , i.e., $v \in \text{dom}(\bar{\mathcal{I}}(\mathcal{H}_0))$ holds. This is a contradiction. \square

Let us demonstrate the use of the two theorems on a very simple example. Consider the triple:

$$\{((x \mapsto _) * \mathcal{H}) \sqsubseteq \mathcal{M}\} * x = 1; \{((x \mapsto 1) * \mathcal{H}) \sqsubseteq \mathcal{M}\}$$

We could follow the symbolic execution rules presented in Section 4 and also be able to prove this triple. But, for the sake of discussion, we consider local reasoning over triple T :

$$\{(x \mapsto _) \sqsubseteq \mathcal{M}\} * x = 1; \{(x \mapsto 1) \sqsubseteq \mathcal{M}\},$$

which holds trivially. Also, we can clearly see that both $T \rightsquigarrow \text{evolve}((x \mapsto _), (x \mapsto 1))$ and $T \rightsquigarrow \text{enclose}((x \mapsto _))$ hold. Applying the rule (EV), we deduce that $(x \mapsto 1) * \mathcal{H}$ holds after executing the code fragment. Furthermore, applying the frame rule, rule (FR), we deduce that $\mathcal{H} \sqsubseteq \mathcal{M}$ remains true, i.e., the heap reality of \mathcal{H} is preserved. Putting the pieces together, we can establish the truth of the original triple by making use of the two theorems.

Recall that we use traditional conjunction, as opposed to separating conjunction in SL. We thus emphasize that all the rules presented above (CFR, EV and FR in particular) can be used in combination because in our framework: $\{\phi\} P \{\psi_1\}$ and $\{\phi\} P \{\psi_2\}$ imply $\{\phi\} P \{\psi_1 \wedge \psi_2\}$.

Our frame rules vs. SL's frame rule: We now elaborate the connection of our two rules (EV) and (FR) with the traditional frame rule in Separation Logic (SL). First, why do we have two rules while SL has one, as introduced in the beginning of this section? The reason is that SL, succinctly, captures *two* important properties: that

- π can be added to precondition ϕ and it *remains true* in the postcondition;

$\frac{\text{[MALLOC]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M} \quad \psi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M} \quad \psi \models \text{dom}(\tilde{\mathcal{H}}') \subseteq \text{dom}(\tilde{\mathcal{H}}) \cup \{x\}}{\{\phi\} x = \text{malloc}(1) \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')}$
$\frac{\text{[FREE]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M} \quad \psi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M} \quad \psi \models \text{dom}(\tilde{\mathcal{H}}') \subseteq \text{dom}(\tilde{\mathcal{H}}) \setminus \{x\}}{\{\phi\} \text{free}(x) \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')}$
$\frac{\text{[OTHER-STATEMENTS]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M} \quad \psi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M} \quad \psi \models \text{dom}(\tilde{\mathcal{H}}') \subseteq \text{dom}(\tilde{\mathcal{H}})}{\{\phi\} s \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')}$
$\frac{\text{[SEQ-COMPOSITION]}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}') \quad \{\psi\} Q \{\gamma\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}', \tilde{\mathcal{H}}'')}{\{\phi\} P; Q \{\gamma\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}'')}$
$\frac{\text{[CALL]}}{[\{\phi\} f() \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')] \in \text{Specs} \quad \phi' \models \phi}{\{\phi'\} \text{call } f() \{_ \} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')}$
$\frac{\text{[COMPOSITION]}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}_1, \tilde{\mathcal{H}}'_1) \quad \{\phi\} P \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}}) \quad \phi \models \tilde{\mathcal{H}} * \tilde{\mathcal{H}}_2 \wedge \tilde{\mathcal{H}}_2 \sqsubseteq \mathcal{M}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}_1 \cup \tilde{\mathcal{H}}_2, \tilde{\mathcal{H}}'_1 \cup \tilde{\mathcal{H}}'_2)}$

Figure 5. Hoare-style Rules for Evolution. OTHER-STATEMENTS applies to s not of the kind covered by the rules above.

- π retains its *separateness*, from precondition ϕ to postcondition ψ .

The second property is important for *successive* uses of the frame rules. Our rule (FR) above only provides for the first property. We accommodate the second property with the other rule (EV), i.e., the “propagation of separation” rule.

The two concepts of evolution and enclosure in fact exist in SL, *implicitly*. Given the triple $T = \{\phi\} P \{\psi\}$, assume that \mathcal{H} is the heap housing the precondition ϕ and \mathcal{H}' is the heap housing the postcondition ψ . In SL, the frame rule also requires that $T \rightsquigarrow \text{evolve}(\mathcal{H}, \mathcal{H}')$ and that $T \rightsquigarrow \text{enclose}(\mathcal{H})$. In short, this means that whenever the traditional frame rule in SL² is applicable, our frame rules are also applicable without any additional complexity.

²We assume an SL fragment without magic wands.

$\frac{\text{[HEAP-ASSIGN]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M} \quad x \in \text{dom}(\tilde{\mathcal{H}})}{\{\phi\} *x = y \{_ \} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}$
$\frac{\text{[FREE]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M} \quad x \in \text{dom}(\tilde{\mathcal{H}})}{\{\phi\} \text{free}(x) \{_ \} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}$
$\frac{\text{[OTHER-STATEMENTS]}}{\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M}}{\{\phi\} s \{_ \} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}$
$\frac{\text{[SEQ-COMPOSITION]}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}_1, \tilde{\mathcal{H}}'_1) \quad \{\phi\} P \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}}) \quad \{\psi\} Q \{\gamma\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}}')}{\{\phi\} P; Q \{\gamma\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}$
$\frac{\text{[CALL]}}{[\{\phi\} f() \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})] \in \text{Specs} \quad \phi' \models \phi \wedge \mathcal{H} \sqsubseteq \mathcal{M}}{\{\phi'\} \text{call } f() \{_ \} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})}$
$\frac{\text{[COMPOSITION]}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}}) \quad \phi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M}}{\{\phi\} P \{\psi\} \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}} \cup \tilde{\mathcal{H}}')}$

Figure 6. Hoare-style Rules for Enclosure. OTHER-STATEMENTS applies to s not of the kind covered by the rules above.

However, in general our assertion language allows for multiple subheaps, which entails more expressive power, but at the cost that we no longer can resort to the abovementioned default. For this paper, we require the specifications to also nominate the subheaps participating in the evolution and/or enclosure relations, stated under the keyword **frame**, following the typical **requires** and **ensures** keywords. We demonstrate this in Section 6 with our driving example.

Proving the Evolution and Enclosure relations. The next question of interest is how the evolution and enclosure relations are practically checked. For evolution, we use the rules in Fig. 5. In the rule [CALL],

$$[\{\phi\} f() \{\psi\} \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')] \in \text{Specs}$$

means that we have nominated $\text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ the specifications of function f . Similarly for enclosure relation, which can be effectively checked using the rules presented in Fig. 6. Checking evolution and enclosure relations is also performed modularly. Specifically, at call sites, we make use of the rule

[CALL] and then achieve compositional reasoning with the rule [COMPOSITION].

We finally conclude this section with two Lemmas about the correctness of the rules presented in Figures 5 and 6. The proofs of the two lemmas follow similar (but more tedious) steps as in proving our two main theorems. For brevity, we omit the details.

Lemma 5.6 (Evolution). *Given a valid triple $T = \{\phi\} P \{\psi\}$ where $\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M}$ and $\psi \models \tilde{\mathcal{H}}' \sqsubseteq \mathcal{M}$, $T \rightsquigarrow \text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ holds if it follows from the rules in Fig. 5. \square*

Lemma 5.7 (Enclose). *Given a valid triple $T = \{\phi\} P \{_ \}$ where $\phi \models \tilde{\mathcal{H}} \sqsubseteq \mathcal{M}$, $T \rightsquigarrow \text{enclose}(\tilde{\mathcal{H}})$ holds if it follows from the rules in Fig. 6. \square*

6 Automatic Proof of the Challenge Example

```

mgraph( $h, x, t^{in}$ ) :-
   $h \doteq \text{emp}, x = \text{null}$ .
mgraph( $h, x, t^{in}$ ) :-
   $(x \mapsto (1, \_, \_)) \sqsubseteq t^{in}, h \doteq \text{emp}$ .
mgraph( $h, x, t^{in}$ ) :-
   $h_x \doteq (x \mapsto (1, l, r)), t_l \doteq h_x * t^{in},$ 
  mgraph( $h_l, l, t_l$ ),  $t_r \doteq h_l * t_l,$ 
  mgraph( $h_r, r, t_r$ ),  $h \doteq h_x * h_l * h_r, h * t^{in}$ .

pmg( $h, x, t^{in}, t^{out}$ ) :-
   $h \doteq \text{emp}, x = \text{null}, t^{out} \doteq t^{in}$ .
pmg( $h, x, t^{in}, t^{out}$ ) :-
   $(x \mapsto (1, \_, \_)) \sqsubseteq t^{in}, h \doteq \text{emp}, t^{out} \doteq t^{in}$ .
pmg( $h, x, t^{in}, t^{out}$ ) :-
   $h_x \doteq (x \mapsto (1, l, r)), t_l \doteq h_x * t^{in},$ 
  mgraph( $h_l, l, t_l$ ),  $t_r \doteq h_l * t_l,$  mgraph( $h_r, r, t_r$ ),
   $t_{out} \doteq h_r * t_r, h \doteq h_x * h_l * h_r, h * t^{in}$ .
pmg( $h, x, t^{in}, t^{out}$ ) :-
   $h_x \doteq (x \mapsto (0, l, r)), t_l^{in} \doteq (x \mapsto (1, l, r)) * t^{in},$ 
  pmg( $h_l, l, t_l^{in}, t_r^{in}$ ), pmg( $h_r, r, t_r^{in}, t_{out}^{in}$ ),
   $h \doteq h_x * h_l * h_r, h * t^{in}$ .

```

Figure 7. Definitions of mgraph and pmg

Recall the graph marking algorithm in Fig. 1, which exhibits precisely four scenarios. (1) The function terminates upon seeing a null pointer. The function also terminates upon encountering a marked node. For this there are two possibilities: (2) the current node has been encountered before (in the history); or (3) the subgraph rooted at the current node had already been fully marked (modulo the history). Finally, when encountering an unmarked node (4), the function first marks the node, then invokes two recursive calls to deal with the left and right subgraphs. This last scenario poses a technical challenge, concerning separation of the two recursive calls, so that a frame rule can be used to protect

the effects of the first call from the that of the second. In actual fact, the second call can refer to a portion of the heap modified by the first call. The important point however is the second call does not *write* to this subheap.

The four rules in our definition of “partially-marked-graph” $\text{pmg}(h, x, t^{in}, t^{out})$ correspond to the four scenarios identified above. We address the technical challenges by having: (a) h encloses the *write* footprint of the code while precisely excludes the nodes that had been visited in the history; (b) t^{in} captures the history, i.e. nodes visited starting from the root node to the current node x ; and importantly, (c) t^{out} captures the output history, which would be the set of visited nodes right after the function `mark` finishes processing the subgraph rooted at x . The use of t^{out} resembles a form of “continuation passing”.

The importance of t^{out} can be understood by investigating the 4th scenario identified above. Encountering the node x that is unmarked, the function first marks it before recursively processing the left subgraph and then the right subgraph. What then should be used as the histories for these recursive calls? The history used for the first call can be easily constructed by conjoining the history of the call to x with the updated node x (the mark field has been set). However, the actual history used for the second call very much depends on the shape of the original graph. We choose to construct t^{out} recursively, thus the output history of the first recursive call can be conveniently used as the input history for the second call. The 4th rule in the definition of `pmg` closely follows these intuitions.

Before proceeding, we contrast here our use of the predicate `pmg` with the way predicates are used in SL. In SL, a predicate describes (a part of) the current heap; in `pmg`, we simultaneously describe *three heaps* corresponding to different stages of computation.

In Fig. 8 we show the specification of the function `mark` and the proof for the most interesting case: x is not null and its `m` field has not been marked. In the precondition, \mathcal{H} , the first component of the definition of `pmg`, appropriately encloses the *write* footprints of the function. It is thus easy to derive, $\text{ENCLOSE}(\mathcal{H})$. Proving that $\text{EVOLVE}(\mathcal{H}, \mathcal{H}')$ is also standard, thus we will not elaborate on this. Instead, let us focus the discussion on how the frame rules are used.

The assertion after step 1 is obtained by unfolding the definition of `pmg` using the fourth rule and instantiating the values of l and r . Note that this unfolding is triggered since the footprint of x is touched. (Using the other rules will lead to a conflict with the guard `assume(x && x->m != 1)`.) At the recursive call `mark(1)` (point 3), we need to prove that the assertion after program point 2 implies the precondition of the function `mark`. In this context, the precondition is:

$\text{pmg}(\mathcal{H}^l, l, t^{in^l}, t^{out^l}), \mathcal{H}^l \sqsubseteq \mathcal{M}, t^{in^l} \sqsubseteq \mathcal{M}$. Such a proof can be achieved simply by matching \mathcal{H}^l with \mathcal{H}_l , t^{in^l} with t_l^{in} , and t^{out^l} with t_r^{in} .


```

requires:   pmg( $\mathcal{H}, x, t^{in}, t^{out}$ ),  $\mathcal{H} \sqsubseteq \mathcal{M}$ ,  $t^{in} \sqsubseteq \mathcal{M}$ 
ensures:   mgraph( $\mathcal{H}', x, t^{in}$ ),  $\mathcal{H}' \sqsubseteq \mathcal{M}$ ,  $t^{out} \simeq t^{in} * \mathcal{H}'$ ,
frame:     enclose( $\mathcal{H}$ ), evolve( $\mathcal{H}, \mathcal{H}'$ )

void mark(struct node *x) {
  { pmg( $\mathcal{H}, x, t^{in}, t^{out}$ ),  $\mathcal{H} \sqsubseteq \mathcal{M}$ ,  $t^{in} \sqsubseteq \mathcal{M}$  }
  1 assume(x && x->m != 1); struct node *l = x->left, *r = x->right;
  {  $\mathcal{H}_x \simeq (x \mapsto (0, l, r))$ ,  $t_l^{in} \simeq (x \mapsto (1, l, r)) * t^{in}$ , pmg( $\mathcal{H}_l, l, t_l^{in}, t_r^{in}$ ), pmg( $\mathcal{H}_r, r, t_r^{in}, t^{out}$ ),
     $\mathcal{H} \simeq \mathcal{H}_x * \mathcal{H}_l * \mathcal{H}_r$ ,  $\mathcal{H} * t^{in}$ ,  $\mathcal{H} \sqsubseteq \mathcal{M}$ ,  $t^{in} \sqsubseteq \mathcal{M}$  }
  2 x->m = 1;
  { pmg( $\mathcal{H}_l, l, t_l^{in}, t_r^{in}$ ),  $t_l^{in} \simeq (x \mapsto (1, l, r)) * t^{in}$ ,  $\mathcal{H}_l * \mathcal{H}_r * t_l^{in} \sqsubseteq \mathcal{M}$ , pmg( $\mathcal{H}_r, r, t_r^{in}, t^{out}$ ) }
  3 mark(l);
  { mgraph( $\mathcal{H}'_l, l, t_l^{in}$ ),  $\mathcal{H}'_l \sqsubseteq \mathcal{M}$ ,  $t_r^{in} \simeq t_l^{in} * \mathcal{H}'_l$ , // postcondition
     $t_l^{in} \simeq (x \mapsto (1, l, r)) * t^{in}$ , pmg( $\mathcal{H}_r, r, t_r^{in}, t^{out}$ ), // (CFR)
     $\mathcal{H}'_l * \mathcal{H}_r * t_l^{in}$ , // (EV)
     $\mathcal{H}_r * t_l^{in} \sqsubseteq \mathcal{M}$  } // (FR)
  4 mark(r);
  { mgraph( $\mathcal{H}'_l, l, t_l^{in}$ ),  $t_r^{in} \simeq t_l^{in} * \mathcal{H}'_l$ ,  $t_l^{in} \simeq (x \mapsto (1, l, r)) * t^{in}$ , // (CFR)
    mgraph( $\mathcal{H}'_r, r, t_r^{in}$ ),  $\mathcal{H}'_r \sqsubseteq \mathcal{M}$ ,  $t^{out} \simeq t_r^{in} * \mathcal{H}'_r$ , // postcondition
     $\mathcal{H}'_l * \mathcal{H}'_r * t_l^{in}$ , // (EV)
     $\mathcal{H}'_l \sqsubseteq \mathcal{M}$ ,  $t_l^{in} \sqsubseteq \mathcal{M}$  } // (FR)
}

```

Figure 8. Mark Graph Example

The assertion after this call (step 3) is then obtained by application of framing. First we use the specification to replace the first occurrence of `pmg` by `mgraph`. What we would like to focus on here is the shaded heap formula. First, applying rule (FR), we frame $\mathcal{H}_r * t_l^{in} \sqsubseteq \mathcal{M}$ through the step 3 because the heaps \mathcal{H}_r and t_l^{in} lie outside the updates of the recursive call `mark(l)`; note that $\mathcal{H}_l \sqsubseteq \mathcal{M}$, however, no longer holds and is removed. Second, \mathcal{H} evolves into \mathcal{H}' , so a heap's separation from \mathcal{H} before the step was propagated into its separation from \mathcal{H}' after the step, shown as the application of rule (EV).

This explanation is easily adapted for the call at program point 4. Finally, the postcondition is proved by unfolding `mgraph(\mathcal{H}', x, t^{in})` using the third rule, followed by appropriate variable matching.

In our graph marking example, our “invariant” precondition involves the predicate `pmg` while the final postcondition involves the predicate `mgraph`. The fact that `pmg` resembles the code is coincidental but unsurprising, since it needs to describe the subheaps relevant to the two recursive calls. One might argue that the top-level specification `mgraph` is contrived so as to be similar to `pmg`. One could notice that the former definition is “left-askew”, as the “history” used for the right subgraph is computed by conjoining the footprint and the history of the left subgraph. If instead we had used a “right-askew” definition, the final entailment may become very hard to prove. In the end, this paper is ultimately about

automation, and not about how we can hide implementation details and use highly declarative specifications.

Remark: There is a recently published proof [24] that considers a similar graph marking algorithm. By supporting the construct of *iterated separating conjunction* [31], they managed to “verify challenging examples such as graph-marking algorithms that so far were beyond the scope of automated verifiers based on permission logics [such as SL and IDFs]”. The critical difference is that their method precondition does not require that the input graph to be “properly marked”. For example, it does not allow a fully marked graph to be a valid input. More specifically, it requires that the root node must be unmarked. Furthermore, the postcondition on its own *does not imply* the final graph is completely marked. The crucial point here is that the proof in [24] does not prove the same thing as we do. As an aside, that proof is not about local reasoning; it does not use framing at all. Indeed, the specification even refers to addresses *outside* its code footprint. The dynamic frame of the method, *and* those of its recursive method calls, are all the same: it represents the one global graph. In the end, [24] does not contribute to the reasoning of recursive predicates.

7 A Prototype Implementation

We implemented a prototype in `CLP(\mathcal{R})` [16], submitted as supplementary material for this paper.

Our prototype takes as input a C program with proper function specifications. We support only a small subset of the

C programming language. For example, we support recursive functions but not loops, and expect loops to be manually compiled into tail-recursive functions³. (We also disallow nested expressions, and instead use more temporary variables.)

Algorithm 1: Modular Verification of Pointer Programs

```

1 function VERIFY
  Input: Function  $\mathcal{P}$  as a transition system
  Output: True if successfully verified; otherwise False
2  $\phi \leftarrow$  the precondition of  $\mathcal{P}$  ;
3  $\ell_0 \leftarrow$  the starting location of  $\mathcal{P}$  ;
4 States  $\leftarrow \{ \langle \ell_0, \phi \rangle \}$  ;
5 while (States is not empty) do
6    $\langle \ell, \pi \rangle \leftarrow$  pop a state from States ;
7   if ( $\ell$  is ending location of  $\mathcal{P}$ ) then
8      $\psi \leftarrow$  the postcondition of  $\mathcal{P}$  ;
9     if (not ENTAIL( $\pi, \psi$ )) return False;
10  else
11    for each transition  $\ell \xrightarrow{\text{stmt}} \ell'$  in  $\mathcal{P}$  do
12      if stmt is “call  $Q$ ” then
13         $\phi' \leftarrow$  the precondition of  $Q$  ;
14        if (not ENTAIL( $\pi, \phi'$ )) return False;
15         $\psi' \leftarrow$  the postcondition of  $Q$  ;
16         $\pi' \leftarrow$  FRAMERULE( $\pi, \phi', \psi'$ ) ;
17      else
18         $\pi' \leftarrow$  SYMBOLICEXEC( $\pi, \text{stmt}$ ) ;
19      end
20      States  $\leftarrow$  States  $\cup \{ \langle \ell', \pi' \rangle \}$  ;
21    end
22  end
23 end
24 return True;

```

Our prototype compiles a C program into a transition system, which will be fed into the main algorithm, presented in **Algorithm 1**.

For presentation purpose, in **Algorithm 1**, symbolic state only consists of the current program point and a *state formula* that is in our assertion language presented in Section 3.

We maintain a worklist of symbolic states (States) and consider one symbolic path at a time (the while loop and line 6). If we reach the function end point, we ensure that the state formula implies the postcondition; if the entailment does not hold, we can immediately terminate the verification process with failure (line 8-9). Otherwise, for each transition emanating from the current program location ℓ , we compute the successor states and add them into the worklist.

We elaborate on the computation of successor states. There are three types of transitions due to different statement types:

³We manually translate the example in Fig. 3 to a recursive function, used later as a benchmark program increment in Table 1.

(a) a function call; (b) a statement which manipulates only the *stack* memory; and (c) a heap-manipulating statement as identified in Section 3.

- **For a function call:** We first check that the precondition of the corresponding function (Q) is met, otherwise we immediately terminate with failure (line 14). We then apply the frame rules in Section 5 in combination with the classical frame rule (CFR) to achieve compositional reasoning, denoted by the helper function FRAMERULE.
- **For a statement that is not a function call:** We just apply symbolic execution, denoted by calling the helper function SYMBOLICEXEC. Note that for (b), standard symbolic execution is well-understood; for (c) we used the rules presented in Proposition 4.1.

For presentation purpose, we omit the details on the handling of enclosure and evolution relations from our high-level algorithm. In fact, the rules in Fig. 5 and Fig. 6 are incorporated into our implementation to work in tandem with our symbolic execution rules and frame rules. For example, to prove $\text{evolve}(\tilde{\mathcal{H}}, \tilde{\mathcal{H}}')$ for a symbolic path, at any point in the path we would track the largest possible subheap $\bar{\mathcal{H}}$ such that $\text{evolve}(\tilde{\mathcal{H}}, \bar{\mathcal{H}})$. In the end, the remaining obligation is to prove that $\tilde{\mathcal{H}}' \sqsubseteq \bar{\mathcal{H}}$. For this reason, our implementation does not suffer from any degree of non-determinism when dealing with the [SEQ-COMPOSITION] rules in Figures 5 and 6.

We now comment on the proofs of *entailments* between recursive definitions at call sites and at the end of a function, i.e., the helper function ENTAIL. To demonstrate full automation, our prototype adapted an entailment check procedure from [8, 30]. There they use a general strategy of unfolding a predicate in both the premise and conclusion until the entailment becomes obvious; [9] describes this strategy as “unfold-and-match” (U+M) and we will follow this terminology. In particular:

- We unfold a recursive constraint on a pointer x when its “footprint” (e.g., $x \rightarrow \text{next}$) is touched by the code [30]. This step is performed during symbolic execution.
- At a call site or the end of a function, we deal with obligations of the form $\mathcal{L} \models \mathcal{R}$, performing a sequence of left unfolds (unfolding \mathcal{L}) and/or right unfolds (unfolding \mathcal{R}) until the proof obligation is simple enough such that a “proof by matching” is successful. At this point, recursive predicates are treated as *uninterpreted*. After dealing with with the heap constraints, an SMT solver – Z3 [10] – can be employed to discharge the obligation.

In summary, our entailment check procedure directly deals with user-defined recursive predicates, yet does *not* employ any user-defined lemmas or axioms. Neither does it involve newer technology such as automatic induction [9]. The important point here is that our entailment check procedure is obtained not from a *custom* theorem-proving method, but rather from established methods.

Table 1. Benchmarking Our Prototype Implementation. # VCs denotes the number of entailment checks; # Z3 calls denotes the number of calls to Z3.

Group	Program	Time (s)	# VCs	# Z3 calls
simple	ex1 (Fig. 4)	0.2	1	11
	*buggy-ex1	0.2	1	11
	other examples	0.3	4	(total) 11
sll	append	0.9	3	86
	copy	8.4	3	66
	filter	2.0	4	74
	increment	0.7	3	58
	insert	0.4	1	19
	insert-end	2.1	3	74
	length	0.1	3	23
	*buggy-length	0.2	3	32
	remove	0.1	2	15
	traverse	0.1	1	10
	zero	0.9	3	60
	tree	bst-search	1.0	6
isocopy		12.9	4	75
*buggy-isocopy		0.1	0	23
traverse		0.5	4	38
graph	markgraph	6.8	6	174
	*buggy-mark1	33.7	4	303
	*buggy-mark2	11.9	4	161

Further, it should also be noted that there are example programs manipulating linked-lists, whose generated VCs would require *induction* to prove (see [9]). Thus to build a comprehensive verifier, we should not be content with just “unfold-and-match”. Instead, we should incorporate other advances in the area of theorem proving (e.g. [9, 28, 29]) into the entailment check procedure ENTAIL.

Experimental Results. To demonstrate the applicability of our framework, other than our breakthrough example and examples presented throughout this paper, we have also selected a number of example programs from the GRASSHOPPER system [28]. As sanity checks, we also introduce a number of buggy variants (prefixed by *buggy-) which, as expected, our prototype will fail to verify. We used a 2.3 GHz machine running Ubuntu 14.04.3 LTS, with 4GB memory. Results appear in Table 1.

Our benchmarks are in four categories:

- *heap manipulations* (simple). The properties to be proved do not involve recursive constraints.

- *singly-linked lists* (sll for short). The properties (collectively) involve reasoning about the shape, data, and size of a list.
- *trees*. Programs here traverse a binary and binary search tree. We also have a distinguished example isocopy which has not been verified before in as general a manner.
- The last group is about our driving example: graph marking, and two buggy variants. The purpose here is simply to present some performance metrics.

Proving isomorphic trees. Consider isocopy, which is about the classic problem of copying a tree. This program has been previously used by [4] to demonstrate symbolic execution with Separation Logic (SL). However, [4] simply proves that the new tree is separate from the original one; here we prove a more challenging property, that the copy, also a tree, is *isomorphic* to the original tree. Specifying such property is easy using our framework of explicit heaps, because we can simultaneously describe different heaps corresponding to different stages of computation.

On buggy examples. We have deliberately injected a number of different bugs into originally safe programs. To name a few: wrongly specified “enclosure” heap (*buggy-isocopy), buggy recursive definitions (*buggy-mark2), buggy stack manipulating statements (*buggy-length), and buggy heap-manipulating statements (*buggy-mark1). For these cases, the performance of our verifier can diverge significantly. For most examples, we fail and terminate quickly. Notably, however, for the case of *buggy-mark1, our entailment check procedure exhausts its options without being able to find a successful proof.

8 Further Related Work and Discussion

It is possible, but very difficult, to reason in Hoare logic about programs with pointers; [5, 23] explore this direction. The resulting proofs are inelegant and remain too low-level to be widely applicable, let alone being automated.

Separation Logic (SL) [26, 31] was a significant advance with local reasoning via a frame rule, influencing modern verification tools. For example, [3, 7, 14] implement SL-based symbolic execution, as described in [4]. But there was a problem in accommodating data structures with *sharing*.

Bornat et al. [6] present a pioneering SL-based approach for reasoning about data structures with intrinsic sharing. The attempt results in “dauntingly subtle” [6] definitions and verifications. Thus it is unclear how to automate such proofs.

Explicit naming of heaps naturally emerged as extensions of SL [11]. Reynolds [32] conjectured that referring explicitly to the current heap in specifications would allow better handles on data structures with sharing. We advance [11] with a proof method for propagating and reasoning about recursive definitions. More specifically, we now considered

entailments between such definitions, whereas [11] only considered simple safety properties, which can be translated to the satisfiability problem restricted to non-recursive definitions. But more importantly, it is this current paper that fully realizes Reynolds' conjecture by connecting the explicit subheaps to the global heap (\mathcal{M}) with the concept of heap reality and formalizing the concepts of "evolution" and "enclosure". This leads to a new frame rule, and consequently enables local reasoning.

Next consider [13] which addressed sharing (but not automation). Recall the mark function, but now consider its application on a DAG, Fig. 9. The "ramify" rule in [13] would isolate the shaded heap portion 1 and prove that the portion 1 has all been marked. With the help of the




Figure 9. mark DAG

magic wand, this seems general. Its application, however, is counter-intuitive and hard to automate, because the portion 1 is *artificial*: it does not correspond to the actual traversal of the code.

We now elaborate on our earlier comments about DF/IDF. In DF, the method footprint is described by a distinguished variable in an expression containing program variables, and additional ghost variables. To have *footprint compliance*, code is annotated with ghost variable statements. An advantage of DF is in its expressivity of the footprint. IDF is a refinement of DF that avoids the need in DF to explicitly specify and then verify frame properties. Instead, these properties are *inferred* from assertions about accessibility, and these are typically written in the pre and post conditions of methods. The concept of separation is then covered by reverting back to SL, ie. using separating conjunction. As a result, IDF can concisely reason about frames.

Some prominent verifiers that use DF/IDF are Vericool [34], Verifast [14], Dafny [20], Chalice [21] and Viper [1]. None of these systems directly support recursive data structures in a general way. As [17] remarks, "it is not clear how to prove properties involving footprints ... , e.g., footprint compliance and self-framing, without an *exact* non-recursive definition of dynamic frames. Such a definition however would in general require reachability predicates ... Reachability predicates are not expressible in first order logic, and therefore cannot be used with an SMT-based tool".

The work [25] shows that by choosing less straightforward definitions of heaps and of heap union in Coq, we can obtain effective reasoning with abstract heap variables, and hence support full separation logic without resulting in excessive proof obligations. As a result, proofs of a number of simple but realistic programs have been successfully mechanized. Similarly, the work [33], which described a mechanized proof of a concurrent in-place spanning tree construction algorithm, bears resemblance to our graph marking example. This is because they traverse via two recursive calls (but they are unconcerned about their relative order). Therefore this work does address the challenge of dealing with the interaction of two recursive calls. Both these works [25, 33] do not address the automation of local reasoning.

Next we briefly mention some recent work on Region Logic, see e.g. [2]. This work is related to DF: it is essentially a form of Hoare logic for object-based programs. A region, like a dynamic frame, is an expression to describe the footprint of a function. Finally, we also mention [19, 22] that address sharing in the context of shape analysis. In contrast, our current paper focuses on functional verification, thus the ability to perform strongest postcondition propagation is crucial. In this context, automated framing of such assertion formula is a non-trivial task.

On "Proof by Framing": This paper advances local reasoning when dealing with data structures with sharing. However, it should be noted that local reasoning might not always be applicable. Consider the following example: a modification of the mark example, but instead working on DAGs.

```
void countpath(struct node *x) {
    if (!x) return;
    struct node *l = x->left, *r = x->right;
    x->mark = x->mark + 1;
    countpath(l); countpath(r);
}
```

This program, intuitively, counts the number of "paths" from the root to each node in the DAG. It cannot be verified using our frame rule(s), simply because the sets of cells modified by left and right recursive calls overlap: what established by the first cannot be framed over the subsequent fragment. In this case, it is clear that "local reasoning" with framing is *not* the way to proceed. (It does not mean that we cannot prove such program using a manual or a non-compositional method.)

9 Conclusion

We considered the problem of automatically verifying programs which manipulate complex data structures. These structures, which may exhibit unrestricted sharing including being cyclic, are defined in our specification language which uses recursive definitions. A key feature of our definitions is their use of explicit heaps in order to frame away constituent substructures, in preparation for local reasoning.

Our main contribution is then, given a program which has been annotated with preconditions and postconditions, a method to:

- generate verification conditions via symbolic execution (which realizes strongest postcondition reasoning) and framing; and
- discharge the verification conditions by using a standard method of unfolding recursive definitions.

Finally, we presented a prototype implementation and demonstrated it over a number of representative programs. In particular, we focused on a graph marking program and presented its first verification by systematic means. An practical outcome of this important example is that its proof provides a template for the formal proof of a class of programs which traverse possibly cyclic data structures.

References

- [1] <http://www.pm.inf.ethz.ch/research/viper.html>.
- [2] A. Banerjee, D. Naumann, and S. Rosenberg. Regional logic for local reasoning about global invariants. In *ECOOP*, pages 387–411, 2008.
- [3] J. Berdine, C. Calcagno, and P. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, pages 115–137, 2005.
- [4] J. Berdine, C. Calcagno, and P. O’Hearn. Symbolic execution with separation logic. In *APLAS*, pages 52–68, 2005.
- [5] R. Bornat. Proving Pointer Programs in Hoare Logic. In *Mathematics of Program Construction*, pages 102–126, 2000.
- [6] R. Bornat, C. Calcagno, and P. O’Hearn. Local reasoning, separation, and aliasing. In *Proc. 2nd workshop on Semantics, Program Analysis and Computing Environments for Memory Management*, 2004.
- [7] M. Botinčan, M. Parkinson, and W. Schulte. Separation logic verification of C programs with an SMT solver. *Electronic Notes in Theoretical Computer Science*, 254:5–23, October 2009.
- [8] W.-N. Chin, C. David, H. H. Nguyen, and S. Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. In *Science of Computer Programming*, 77(9), pages 1006–1036, 2012.
- [9] D. H. Chu, J. Jaffar, and M. T. Trinh. Automatic induction proofs of data-structures in imperative programs. In *PLDI*, pages 457–466, 2015.
- [10] L. De Moura and N. Bjørner. Z3: an efficient smt solver. In *TACAS*, 2008.
- [11] G. Duck, J. Jaffar, and Nicolas Koh. Constraint-based program reasoning with heaps and separation. In *CP*, pages 282–298, 2013.
- [12] C. A. R. Hoare. An axiomatic basis for computer programming. *Comm. ACM*, 1969.
- [13] A. Hobor and J. Villard. The ramifications of sharing in data structures. In *POPL*, pages 523–536, 2013.
- [14] B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens. Verifast: A powerful, sound, predictable, fast verifier for c and java. In *NFM*, pages 41–55, 2011.
- [15] J. Jaffar and J.-L. Lassez. Constraint Logic Programming. In *POPL*, pages 111–119, 1987.
- [16] J. Jaffar, S. Michaylov, P. J. Stuckey, and R. H. C. Yap. The CLP(R) language and system. *ACM TOPLAS*, 14(3):339–395, 1992.
- [17] D. Kassios. Dynamic frames and automated verification – a tutorial, 2011.
- [18] I. Kassios. Dynamic frames: Support for framing, dependencies and sharing without restrictions. In *FM*, pages 268–283, 2006.
- [19] O. Lee, H. Yang, and R. Petersen. Program analysis for overlaid data structures. In *CAV*, pages 592–608, 2011.
- [20] K. R. Leino. Dafny: An automatic program verifier for functional correctness. In *LPAR’10*, 2010.
- [21] K. R. M. Leino and P. Mueller. A basis for verifying multi-threaded programs. In *ESOP*, volume 5502, pages 378–393. Springer, 2009.
- [22] H. Li, B.-Y. E. Chang, and X. Rival. Shape analysis for unstructured sharing. In *SAS*, pages 90–108, 2015.
- [23] J. M. Morris. A general axiom of assignment. assignment and linked data structures. a proof of the schorr-waite algorithm. In *Theoretical Foundations of Programming Methodology*, 1982.
- [24] Peter Müller, Malte Schwerhoff, and Alexander J. Summers. Automatic verification of iterated separating conjunctions using symbolic execution. In *CAV*, 2016.
- [25] A. Nanevski, V. Vafeiadis, and J. Berdine. Structuring the verification of heap-manipulating programs. In *Symposium on Principles of Programming Languages*, January 2010.
- [26] P. O’Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Computer Science Logic*, pages 1–19, 2001.
- [27] E. Pek, X. Qiu, and P. Madhusudan. Natural proofs for data structure manipulation in C using separation logic. In *PLDI*, pages 440–451, 2014.
- [28] R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using smt. In *CAV*, 2013.
- [29] R. Piskac, T. Wies, and D. Zufferey. Automating separation logic with trees and data. In *CAV*, pages 711–728, 2014.
- [30] X. K. Qiu, P. Garg, A. Stefanescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In *PLDI*, pages 231–242, 2013.
- [31] J. C. Reynolds. Separation logic: A logic for shared mutable data objects. In *LICS*, pages 55–74, 2002.
- [32] J. C. Reynolds. A short course on separation logic. <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/fox-19/member/jcr/wwwaac2003/aac.html>, 2003.
- [33] I. Sergey, A. Nanevski, and A. Banerjee. Mechanized verification of fine-grained concurrent programs. In *PLDI*, pages 77–87, 2015.
- [34] J. Smans, F. Piessens B. Jacobs, and Wolfram Schulte. An automatic verifier for java-like programs based on dynamic frames. In *FASE*, pages 261–275, 2008.
- [35] J. Smans, B. Jacobs, and F. Piessens. Implicit dynamic frames: Combining dynamic frames and separation logic. In *ECOOP*, pages 148–172, 2009.