

# A framework for describing recursive data structure topologies in Coq

Kenneth Roe and Scott Smith

The Johns Hopkins University

**Abstract.** This paper presents an axiomatic framework in Coq for verifying invariants on heap data structures such as lists and trees in a C-like language with a low-level store model. The goal of the framework is to detect common errors such as memory leaks, dangling pointers and looped data structures. The framework provides a language for expressing invariants, and a set of inference axioms for verifying them on code that manipulates the data structures. This work builds on the work done by Cook et al. which uses separation logic with recursive predicates to document data structure invariants. The key extension here is the ability to express and reason about data structures more complex than linked lists. The heap description includes a spatial component describing the basic set of lists and trees in the heap. New logical constructs are included to describe special pointer field invariants such as back pointers. We use the framework to formally prove in Coq the heap invariants of a small example program that generates a linked list representing the tree traversal of a tree. This proof guarantees the integrity of program's data structures and that common errors such as memory leaks or dangling pointer references did not arise. We define the meaning of the abstract state in terms of a simpler concrete state. We also include a number of axioms for reasoning about the abstract state that are used in the Coq verification. Proving soundness of these abstract axioms in Coq remains future work.

## 1 Introduction

Within the last few years, shape analysis has begun to make headway in verifying complex data structures [5, 16, 4, 15]. The most successful of these approaches involves the use of separation logic [14] extended to handle recursive data structures [5]. In the aforementioned work, a language for abstract description of heaps is defined which can be used to specify the set of linked lists the heap contains, and a set of cells also contained in the heap but not connected to any of the linked lists. An enhanced version of this work [8] in which doubly linked lists can also be represented was used to discover memory problems in Microsoft drivers. There are also several efforts in implementing separation logic frameworks in theorem proving environments, for example [13, 12, 1, 17, 11, 2].

This paper describes a methodology for expanding the framework of [5] to cover more complex data structures including trees and doubly linked lists. The

formalism in this paper adds considerably to the flexibility of the types of data structures that can be represented. In addition to handling back pointers as in [7], the framework can handle pointers that go from one data structure into another completely independent data structure. Also, our methodology provides a good mechanism for representing information about pointer variables which iterate through a data structure. To achieve this increased expressiveness, we use an abstract state that contains a spatial component which is similar to earlier systems, as well as a more complex predicate component in which an extension of first order logic can be used to specify heap properties.

We have implemented our framework in Coq[3] and constructed a complete formal proof of correctness for a small tree traversal program which we describe in this paper. The proof verifies that the program maintains the integrity of its data structures and does not have errors such as memory leaks or dangling pointer references. We have developed an abstract state representation which is an extension of the representation used in [5]. The Coq formalization incorporates Hoare axioms to systematically back-propagate the abstract state through the program. We have also developed folding and unfolding axioms for the recursive spatial heap predicates of separation logic. They are needed to handle join points in processing if statements as well as setting up the application of the axiom used to back propagate over a heap store. While the formal proof of the particular example is complete, the task of proving all the heap logic axioms used in the proof to be sound in Coq is a challenging task and remains future work.

The structure of the remainder of the paper is as follows. In Section 2 we present the example program and informally describe properties of the program that our formal Coq proof verifies. In Section 3 we present the abstract heap state representation, followed by Section 4 which contains the programming language syntax and semantics as well as axioms for reasoning about abstract heap states under program execution. In Section 5 we outline how portions of the proof for the example program were performed in Coq.

## 2 An example and the properties we proved

Figure 1 shows a program which performs a traversal of a tree and places the result in a linked list, pointed at by the global variable `p`. Our Coq proof verifies a number of important properties of the data structures in this program. Some properties that can be read off of the heap invariant from our formal proof include the following.

- The program maintains two well formed linked lists, the heads of which are pointed to by `n` and `p`. By well formed we mean that memory on the heap is properly allocated for the lists and there are no loops in the data structures.
- The program maintains a well formed tree pointed to by `r`.
- `t` always points to an element in the tree rooted at `r`.
- The two lists and the tree do not share any nodes.

struct list {	$\ell_{11}$	t = NULL;
struct list *n;	$\ell_{12}$	} else {
struct tree *t;	$\ell_{13}$	struct list *tmp = i->n;
int data;	$\ell_{14}$	t = i->t;
};	$\ell_{15}$	free(l);
	$\ell_{16}$	i = tmp;
	$\ell_{17}$	}
struct tree {	$\ell_{18}$	} else if (t->r==NULL) {
struct tree *l, *r;	$\ell_{19}$	t = t->l;
int value;	$\ell_{20}$	} else if (t->l==NULL) {
};	$\ell_{21}$	t = t->r;
struct list *p;	$\ell_{22}$	} else {
void build_pre_order(struct tree *r) {	$\ell_{23}$	n = i;
$\ell_1$ struct list *i = NULL, *n, *x;	$\ell_{24}$	i = malloc(
$\ell_2$ struct tree *t = r;		sizeof(struct list));
$\ell_3$ p = NULL;	$\ell_{25}$	i->n = n;
$\ell_4$ while (t) {	$\ell_{26}$	x = t->r;
$\ell_5$ n = p;	$\ell_{27}$	i->t = x;
$\ell_6$ p = malloc(sizeof(struct list));	$\ell_{28}$	t = t->l;
$\ell_7$ p->t = t;	$\ell_{29}$	}
$\ell_8$ p->n = n;	$\ell_{30}$	}
$\ell_9$ if (t->l==NULL && t->r==NULL) {	$\ell_{31}$	}
$\ell_{10}$ if (i==NULL) {		
}		
}		

Fig. 1. Example program which builds a linked list of all nodes in a tree

- Other than the memory used for the two lists and the tree, no other heap memory is allocated.
- the `t` field of every element in both list structures points to an element in the tree.

We note that some of these properties cannot be verified by existing separation logic frameworks, most notably the last one.

### 3 Abstract and concrete heap domains

In this section we define the concrete and abstract heaps, and define how a particular concrete (actual runtime) heap may satisfy a given abstract heap specification. We adapt the abstract and concrete heap domain definitions from [5]. For the concrete heap we use  $Stack \triangleq (Var \cup Var') \rightarrow Val$ ,  $Heap \triangleq Val - NULL \rightarrow_{fin} (Fld \rightarrow Val)$  and  $States \triangleq Stack \times Heap$ .  $Val$  is an integer that can be treated as either a numeric value or a pointer. `NULL` is a synonym for the integer 0. We explicitly exclude `NULL` from the domain of  $Heap$ . The cells are slightly different than those used by [5]. Rather than being a single value, they are a mapping from field names to values. Hence, the cons cell for a list or a node in a tree is a single cell in our representation.

For the abstract domain, we adapt [5]’s symbolic heaps. We replace their `ls` list operator with a  $\mathcal{R}$  operator for expressing the structure of more general recursive data structures. We have expanded their  $\Pi$  expressions to include a path operator and quantifiers in addition to equality. Both  $\Pi$  and  $\Sigma$  have been extended to allow simple arithmetic expressions to occur at various points in formulae. The grammar is given in Figure 2. More details on the novel syntax will be given below when their meaning is defined.

$v, x, y, z$	$\in Var$	integer variables
$v', x', y', z'$	$\in Var'$	primed variables
$\mathbf{f}$	$\in Fld$	fields
$F$	$\in \mathcal{P}(Fld)$	sets of fields
$X$	$::= x \mid x'$	
$E$	$::= X \mid E + E \mid E - E \mid E * E \mid Val$	expressions
$\Pi$	$::= Body \mid X \in \Sigma \wedge \Pi$	
$Body$	$::= \mathbf{true} \mid E = E \mid Body \wedge Body \mid Body \vee Body$	
	$\mid \forall Var \in \Sigma. Body \mid \exists Var \in \Sigma. Body \mid \neg Body \mid E \xrightarrow{\mathbf{f}} E$	pure formulae
$\Sigma$	$::= \mathbf{emp} \mid \Sigma * \Sigma \mid \Sigma - * \Sigma \mid X \mapsto \{\mathbf{f} : \overline{X}\}$	
	$\mid \mathcal{R}_F(X, \overline{X}) \mid \mathbf{junk}$	Spatial heaps
$H$	$::= \Pi \mid \Sigma$	Symbolic heaps

**Fig. 2.** Grammar for concrete and abstract heaps

A symbolic heap is defined to be valid in a given concrete heap via the relation  $s, h \vdash \Pi \mid \Sigma$ , where  $s \in Stack$ ,  $h \in Heap$ , and  $\Pi \mid \Sigma$  is a symbolic heap. This relationship defines the set of concrete states,  $s, h$  which are represented by the abstract state  $\Pi \mid \Sigma$ . Through the rest of this section, we will build up the definition of  $s, h \vdash \Pi \mid \Sigma$ . Note that formulae are considered well-formed only if all primed variables ( $v', x', y'$  and  $z'$ ) are existentially quantified; we may sometimes informally leave off such a quantifier but it is implicitly there.

### 3.1 The semantics of $s, h \vdash \Sigma$

In this section we give semantics for symbolic heaps, covering each  $\Sigma$  grammar component in turn.

We first introduce a number of auxiliary definitions. Note that throughout this section we will use  $h \vdash \Sigma$  to abbreviate there exists an  $s$  such that  $s, h \vdash \mathbf{true} \mid \Sigma$  and  $s, h \vdash \Sigma$  to abbreviate  $s, h \vdash \mathbf{true} \mid \Sigma$ .

**Implicit existential quantification of primed variables** Throughout the remainder of the paper  $s, h \vdash \Pi \mid \Sigma$  should be interpreted as there exists an  $s'$  which is an extension of  $s$  mapping all primed variables in  $\Pi \mid \Sigma$  such that  $s', h \vdash \Pi \mid \Sigma$ .

**Separation logic constructs** The  $*$ ,  $-*$  and  $\mapsto$  operators are adapted from separation logic [14] and we assume familiarity with them here. Our heap model is slightly different since the heap locations contain a mapping from field identifiers to values rather than just integers. The  $*$  and  $-*$  operators are standard.

**Definition 1.**  $s, h \vdash \Sigma_1 * \Sigma_2$  if and only if there exists  $h_1$  and  $h_2$  such that  $h = h_1 \cup h_2$ ,  $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$ ,  $s, h_1 \vdash \Sigma_1$  and  $s, h_2 \vdash \Sigma_2$ .

**Definition 2.**  $s, h \vdash \Sigma_1 - * \Sigma_2$  if and only if there exists  $h$  and  $h_2$  such that  $h = h_1 \cup h_2$ ,  $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$ ,  $s, h \vdash \Sigma_1$  and  $s, h_2 \vdash \Sigma_2$ .

**Singleton cell constructs** We use atomic assertion  $x \mapsto \{\overline{\mathbf{f}} : x\}$  to indicate a singleton cell mapping in the abstract state. Its meaning is defined as follows.

**Definition 3.**  $s, h \vdash l \mapsto \{\mathbf{f}_1 : v_1, \dots, \mathbf{f}_n : v_n\}$  if and only if  $h = \{s(l) \mapsto c\}$  such that  $\{\mathbf{f}_1 \mapsto s(v_1), \dots, \mathbf{f}_n \mapsto s(v_n)\} \subseteq c$ .

**Recursive data structures** The  $\mathcal{R}_F(x, \overline{X})$  predicate is used to represent recursive data structures in the spatial heap. These data structures can be either lists or trees. The  $x$  parameter is the root of the data structure. The subscripted  $F$  is the set of fields used for recursive descent. For a list, this could for example be the field labelled `next`. For a tree, there are usually two fields `left` and `right` which are pointers to the left and right subtree. The third parameter represents early stopping points in the recursion. It is used in the representation of states where the modification of a data structure is in progress. For example, consider a program that frees a leaf node but has not as yet updated the pointer from the parent to that node. Before the parent is updated, it needs to be marked as an early stopping point of the recursion.

We define the meaning of  $\mathcal{R}$  by induction on the concrete heap.

**Definition 4.** The meaning of  $\mathcal{R}_F(x, \overline{X})$  in a concrete state  $s, h$ , where  $x \in \text{dom}(s)$ , is defined to be the least relation that satisfies the following clauses:

$$\begin{array}{ll}
s, h \vdash \mathcal{R}_F(x, \overline{X}) \text{ iff } s, h \vdash \text{emp} & \text{where } x \in \overline{X} \\
s, h \vdash \mathcal{R}_F(x, \overline{X}) \text{ iff } s, h \vdash \text{emp} & \text{where } s(x) = \text{NULL} \\
s, h \vdash \mathcal{R}_{\mathbf{f}_1, \dots, \mathbf{f}_n}(x, \overline{X}) \text{ iff} & \text{where } x \notin \overline{X} \wedge s(x) \neq \text{NULL} \\
\quad \exists y_1, \dots, y_n, h_0, \dots, h_n. & \\
\quad s, h_0 \vdash l \mapsto \{\mathbf{f}_1 \mapsto y_1, \dots, \mathbf{f}_n \mapsto y_n\} \wedge & \\
\quad s, h_1 \vdash \mathcal{R}_{\mathbf{f}_1, \dots, \mathbf{f}_n}(y_1, \overline{X}) \wedge \dots \wedge s, h_n \vdash \mathcal{R}_{\mathbf{f}_1, \dots, \mathbf{f}_n}(y_n, \overline{X}) \wedge & \\
\quad h_0 \cup \dots \cup h_n = h \wedge \forall i, j. 0 \leq i, j \leq n \wedge i \neq j \rightarrow h_i \cap h_j = \emptyset &
\end{array}$$

Often we will abbreviate  $\mathcal{R}_F(v, \{\})$  as  $\mathcal{R}_F(v)$ .

**Some examples of  $\mathcal{R}$**  To illustrate the use of  $\mathcal{R}$  in characterizing recursive heap structures, we show some sample models that sometimes satisfy  $\mathcal{R}$ . We start with a simple linked list heap abstraction. In this abstraction, the list is rooted at  $x$  and nodes are connected together by the `next` field – this is expressed as  $\mathcal{R}_{\text{next}}(x, \emptyset)$ . The list in this first concrete heap example has two elements at locations  $x$  and  $y$ .

$$\{x \mapsto 1, y \mapsto 2\}, \{1 \mapsto (\text{next} : 2), 2 \mapsto (\text{next} : \text{NULL})\} \vdash \mathcal{R}_{\text{next}}(x, \emptyset)$$

The next example is a heap abstraction that does not model this  $\mathcal{R}$  because there is a loop.

$$\{x \mapsto 1, y \mapsto 2\}, \{1 \mapsto (\text{next} : 2), 2 \mapsto (\text{next} : 1)\} \not\vdash \mathcal{R}_{\text{next}}(x, \emptyset)$$

Another failure case is where the NULL pointer terminating the list fails to appear in the abstraction. Note that  $2 \mapsto ()$  represents a cell in which no fields are defined; this should not be confused with an unallocated cell in the formalism.

$$\{x \mapsto 1, y \mapsto 2\}, \{1 \mapsto (\text{next} : 2), 2 \mapsto ()\} \not\vdash \mathcal{R}_{\text{next}}(x, \emptyset)$$

It is still possible to partially characterize a heap in this kind of situation. In the following example we put  $z$  in the set of terminators for  $\mathcal{R}$  and hence we do not traverse past it.

$$\{x \mapsto 1, y \mapsto 2, z \mapsto 3\}, \{1 \mapsto (\text{next} : 2), 2 \mapsto (\text{next} : 3), 3 \mapsto ()\} \vdash \mathcal{R}_{\text{next}}(x, \{z\})$$

We now have all the pieces to define  $s, h \vdash \Sigma \mid \text{true}$  for all cases of  $\Sigma$ .

**Definition 5.** We say that  $s, h \vdash \Sigma$  if and only if  $\Sigma$  is

<code>emp</code>	and $\text{dom}(h) = \emptyset$
<code>junk</code>	and $\text{dom}(h) \neq \emptyset$
$\Sigma_1 * \Sigma_2$	and definition 1 holds
$\Sigma_1 - * \Sigma_2$	and definition 2 holds
$x \mapsto \{\mathbf{f}_1 \mapsto v_1 \dots \mathbf{f}_n \mapsto v_n\}$	and definition 3 holds
$\mathcal{R}_F(x, \bar{X})$	and definition 4 holds

### 3.2 The semantics of $\Pi$

Now we need to define  $s, h \vdash \Pi \mid \Sigma$ . We define each form of  $\Pi$  formulae in turn.

$\mathbf{x} \in \Sigma$  We use this construct to identify which portion of the heap contains a variable. Often a pointer variable is used to traverse a single data structure. This construct is useful in identifying which structure that variable belongs to. Also, this construct has a special case allowing the variable to be NULL. This is because most programs allow pointer variables to either point to a valid element in a structure or to be NULL. It is defined as follows.

**Definition 6.**  $s, h \vdash \Sigma_1 * \Sigma_2 \mid \mathbf{x} \in \Sigma_1 \wedge \Pi$  if and only if either  $x = \text{NULL}$  or  $s(x) \in \text{dom}(h_1)$  where  $s, h_1 \vdash \Sigma_1$  and  $h_1 \subseteq h$  and  $s, h \vdash \Pi \mid \Sigma_1 * \Sigma_2$ .

$\forall v \in \Sigma.Body$  and  $\exists v \in \Sigma.Body$  These are the standard universal and existential quantifiers. However the variable  $v$  is restricted to the specified portion of  $\Sigma$ . Formally,

**Definition 7.**  $s, h \vdash \forall v \in \Sigma.Body | \Sigma * \Sigma'$  holds if and only if  $s, h \vdash Body[v \mapsto n] | \Sigma * \Sigma'$  holds for all  $n \in dom(h_1)$  where  $h_1 \vdash \Sigma$  and  $h_1 \subseteq h$  and  $h \vdash \Sigma * \Sigma'$ .

and

**Definition 8.**  $s, h \vdash \exists v \in \Sigma.Body | \Sigma * \Sigma'$  holds if and only if there exists an  $n$  such that  $s, h \vdash Body[v \mapsto n] | \Sigma * \Sigma'$  holds and  $n \in dom(h_1)$  where  $h_1 \vdash \Sigma$  and  $h_1 \subseteq h$  and  $h \vdash \Sigma * \Sigma'$ .

**Expression evaluation** We now define the meaning of simple arithmetic expressions that may appear in many of the atomic formulae, defining an  $\mathcal{E}$  operator that gives the meaning of these arithmetic expressions.

**Definition 9.**  $\mathcal{E}_s(E)$  where  $E$  is an expression and  $s \in Stack$  is defined as follows:

$$\begin{aligned} \mathcal{E}_s(E_1 + E_2) &= \mathcal{E}_s(E_1) + \mathcal{E}_s(E_2) \\ \mathcal{E}_s(E_1 - E_2) &= \mathcal{E}_s(E_1) - \mathcal{E}_s(E_2) \\ \mathcal{E}_s(E_1 * E_2) &= \mathcal{E}_s(E_1) * \mathcal{E}_s(E_2) \\ \mathcal{E}_s(l) &= l && \text{where } l \in Val \\ \mathcal{E}_s(x) &= s(x) && \text{where } x \in Var \end{aligned}$$

**The  $E \xrightarrow{f} E'$  predicate** This atomic formula asserts that dereferencing  $e$  via field  $f$  yields heap location  $e'$ . Its meaning is defined as follows:

**Definition 10.**  $h, s \vdash E \xrightarrow{f} E'$  holds if and only if  $h(\mathcal{E}_s(E))f = \mathcal{E}_s(E')$

With the above definitions we can put everything together.

**Definition 11.** We say that  $s, h \vdash \Pi | \Sigma$  if and only if  $s, h \vdash \Sigma$  and recursively if  $\Pi$  is

$$\begin{aligned} x \in \Sigma \wedge \Pi_1 & \text{ and definition 6 holds} \\ \Pi_1 \wedge \Pi_2 & \text{ and } s, h \vdash \Pi_1 | \Sigma \text{ and } s, h \vdash \Pi_2 | \Sigma \\ \Pi_1 \vee \Pi_2 & \text{ and } s, h \vdash \Pi_1 | \Sigma \text{ or } s, h \vdash \Pi_2 | \Sigma \\ \neg \Pi & \text{ and } s, h \not\vdash \Pi | \Sigma \\ \forall Var \in \Sigma. \Pi & \text{ and definition 7 holds} \\ \exists Var \in \Sigma. \Pi & \text{ and definition 8 holds} \\ E_1 = E_2 & \text{ and } \mathcal{E}_s(E_1) = \mathcal{E}_s(E_2) \\ x \xrightarrow{f} y & \text{ and definition 10 holds} \\ \text{true} & \end{aligned}$$

## 4 Semantics of Programs

We now give semantics for a small imperative programming language, first via a definition of evaluation with respect to a concrete heap, and then by a set of principles axiomatizing program execution with respect to an abstract heap.

The grammar for the programming language is similar to that of [5]. However, since our locations can hold an entire structure rather than a single value, we need to introduce a field dereferencing operator,  $E \rightarrow \mathbf{f}$ . We give semantics for a language with the following grammar:

$$\begin{aligned} b &::= E = E \mid E \neq E \\ p &::= x := E \mid x := E \rightarrow \mathbf{f} \mid E \rightarrow \mathbf{f} := x \mid x := \mathbf{new}(\bar{\mathbf{f}}) \mid \mathbf{delete}(E) \\ c &::= \mathbf{skip} \mid p \mid c; c \mid \mathbf{if} \ b \ \mathbf{then} \ c \ \mathbf{else} \ c \mid \mathbf{while} \ b \ \mathbf{do} \ c \end{aligned}$$

### 4.1 Concrete semantics

We now define a relation which represents a single step of execution in the concrete heap.

**Definition 12.** We define the relation  $s, h \xrightarrow{c} s', h'$  as follows:

- $s, h \xrightarrow{\mathbf{skip}} s, h$
- $s, h \xrightarrow{x:=E} s[x \rightarrow \mathcal{E}_s(E)], h$
- if  $l \notin \text{dom}(h)$  and  $l \neq 0$  then  $s, h \xrightarrow{x:=\mathbf{new}} s[x \rightarrow l], h[l \rightarrow \{\dots\}]$
- if  $l \in \text{dom}(h)$  then  $s, h \xrightarrow{\mathbf{delete} \ l} s, h - \{l \rightarrow \dots\}$
- if  $\mathcal{E}_s(E) \in \text{dom}(h)$  and  $\mathbf{f} \in \text{dom}(h(\mathcal{E}_s(E)))$  then  $s, h \xrightarrow{x:=E \rightarrow \mathbf{f}} s[x \rightarrow h(\mathcal{E}_s(E))\mathbf{f}], h$
- if  $\mathcal{E}_s(E_1) \in \text{dom}(h)$  then  $s, h \xrightarrow{E_1 \rightarrow \mathbf{f} = E_2} s, h[\mathcal{E}_s(E_1) \rightarrow h(\mathcal{E}_s(E_1))\mathbf{f} \rightarrow \mathcal{E}_s(E_2)]$
- if  $s, h \xrightarrow{c_1} s', h'$  and  $s', h' \xrightarrow{c_2} s'', h''$  then  $s, h \xrightarrow{c_1; c_2} s'', h''$ .
- if  $\mathcal{E}_s(E_1) = \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c_1} s', h'$  then  $s, h \xrightarrow{\mathbf{if} \ E_1 = E_2 \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2} s', h'$
- if  $\mathcal{E}_s(E_1) \neq \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c_2} s', h'$  then  $s, h \xrightarrow{\mathbf{if} \ E_1 = E_2 \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2} s', h'$
- if  $\mathcal{E}_s(E_1) \neq \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c_1} s', h'$  then  $s, h \xrightarrow{\mathbf{if} \ E_1 \neq E_2 \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2} s', h'$
- if  $\mathcal{E}_s(E_1) \neq \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c_2} s', h'$  then  $s, h \xrightarrow{\mathbf{if} \ E_1 \neq E_2 \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2} s', h'$
- if  $\mathcal{E}_s(E_1) = \mathcal{E}_s(E_2)$  then  $s, h \xrightarrow{\mathbf{while} \ E_1 = E_2 \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2} s, h$
- if  $\mathcal{E}_s(E_1) = \mathcal{E}_s(E_2)$  then  $s, h \xrightarrow{\mathbf{while} \ E_1 \neq E_2 \ \mathbf{do} \ c} s, h$
- if  $\mathcal{E}_s(E_1) = \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c} s', h'$  and  $s', h' \xrightarrow{\mathbf{while} \ E_1 = E_2 \ \mathbf{do} \ c} s'', h''$  then  $s, h \xrightarrow{\mathbf{while} \ E_1 = E_2 \ \mathbf{do} \ c} s'', h''$
- if  $\mathcal{E}_s(E_1) \neq \mathcal{E}_s(E_2)$  and  $s, h \xrightarrow{c} s', h'$  and  $s', h' \xrightarrow{\mathbf{while} \ E_1 \neq E_2 \ \mathbf{do} \ c} s'', h''$  then  $s, h \xrightarrow{\mathbf{while} \ E_1 \neq E_2 \ \mathbf{do} \ c} s'', h''$

### 4.2 Abstract semantics

With the meaning of a concrete step defined, we can easily define when a step soundly takes one abstract heap to another.

**Definition 13.**  $\Pi \mid \Sigma \xrightarrow{c} \Pi' \mid \Sigma'$  iff for all  $s, h$  such that  $s, h \vdash \Pi \mid \Sigma$  and  $s, h \xrightarrow{c} s', h'$  for some  $s', h'$ , it follows that  $s', h' \vdash \Pi' \mid \Sigma'$ .



In order to do symbolic program analysis, a number of rules need to be introduced. The rules are all designed to facilitate backward chaining. If for some step  $\Pi|\Sigma \xRightarrow{c} \Pi'|\Sigma'$ , a symbolic expression already exists for  $\Pi'|\Sigma'$ , the rules will make it easy to generate an expression for  $\Pi|\Sigma$  so that the relationship holds. We do not necessarily guarantee generation of the weakest pre-condition. We do, however, generate a pre-condition that contains enough useful information for program verification. Each rule defines how to transform the abstract state for a specific statement. Note that the rules have not yet been proved sound with respect to the semantics – this is an axiomatic programming logic. Many of the rules are in fact very easy to prove sound, but a few are quite difficult and we leave overall soundness to future work.

**While, if, sequence and skip** We use the following rules to reason about these statements:

$$\begin{array}{c}
\frac{\Pi_1|\Sigma_1 \wedge b \xRightarrow{c_1} \Pi'|\Sigma' \quad \Pi_2|\Sigma_2 \wedge \neg b \xRightarrow{c_2} \Pi'|\Sigma' \quad \Pi|\Sigma \in \text{join}(\Pi_1|\Sigma_1, \Pi_2|\Sigma_2)}{\Pi|\Sigma \text{ if } b \text{ then } c_1 \text{ else } c_2 \xRightarrow{} \Pi'|\Sigma'} \text{(absIf)} \\
\\
\frac{\Sigma^{Inv}|\Pi^{Inv} \xRightarrow{c} \Sigma^{Inv}|\Pi^{Inv} \quad \Pi|\Sigma \wedge b \xRightarrow{} \Sigma^{Inv}|\Pi^{Inv} \quad \Sigma^{Inv}|\Pi^{Inv} \wedge \neg b \xRightarrow{} \Pi'|\Sigma'}{\Pi|\Sigma \text{ while } b \text{ do } c \xRightarrow{} \Pi'|\Sigma' \wedge \neg b} \text{(absWhile)} \\
\\
\frac{}{\Pi|\Sigma \xRightarrow{\text{skip}} \Pi|\Sigma} \text{(absSkip)} \\
\\
\frac{\Pi|\Sigma \xRightarrow{c_1} \Pi'|\Sigma' \quad \Pi'|\Sigma' \xRightarrow{c_2} \Pi''|\Sigma''}{\Pi|\Sigma \xRightarrow{c_1;c_2} \Pi''|\Sigma''} \text{(absSeq)}
\end{array}$$

In the first two definitions above,  $b$  is used both in the program code and the predicate. When applying this rule, the AST of the  $b$  in the program code is mapped to a corresponding AST using the constructs in  $\Pi$ .  $\Sigma^{Inv}|\Pi^{Inv}$  is the loop invariant for the while rule. In the first rule, used for processing if statements, generally, one is using them in a backchaining fashion. The  $\Pi|\Sigma$  precondition needs to be generated as a join of  $\Pi_1|\Sigma_1$  and  $\Pi_2|\Sigma_2$ . It is defined formally as

**Definition 14.**  $\Pi|\Sigma \in \text{join}(\Pi_1|\Sigma_1, \Pi_2|\Sigma_2)$  if and only if for all  $s, h$  it is the case that  $s, h \vdash \Pi|\Sigma$  implies  $s, h \vdash \Pi_1|\Sigma_1$  and  $s, h \vdash \Pi|\Sigma$  implies  $s, h \vdash \Pi_2|\Sigma_2$

We discuss how to do this join in section 4.4.

**Assignment, allocation and deallocation** The following rules allow for reasoning about assignment, allocation and deallocation commands. The assignment operation is divided into three cases. The first case is where no heap information is involved. The second is an assignment of the form  $v := E \rightarrow \mathbf{f}$  which

involves retrieving a single value from the heap. The third is an assignment of the form  $v \rightarrow \mathbf{f} := E$  where a single value is stored back to the heap. Unlike C or other popular programming languages we do not allow assignments that perform multiple stores or retrievals of values. This is to simplify the presentation; many common preprocessors such as CIL [6] can break down assignments into their atomic store components.

$$\{II[x/v] \wedge x = E \mid \Sigma[x/v]\} \xrightarrow{v := E} \{II \mid \Sigma\} \quad (\text{as1})$$

$$\{v \in \Sigma' \wedge x' \in \Sigma' \wedge II[x'/v] \wedge E \xrightarrow{\mathbf{f}} x' \mid \Sigma[x'/v]\} \xrightarrow{v := E \rightarrow \mathbf{f}} \{v \in \Sigma' \wedge II \mid \Sigma\} \quad (\text{as2})$$

$$\{II' \mid \Sigma * x \mapsto \{\overline{\mathbf{f} : v}\}\} \xrightarrow{x \rightarrow \mathbf{f} := G} \{II \mid \Sigma * x \mapsto \{\mathbf{f} \mapsto G, \overline{\mathbf{f} : v}\}\} \quad (\text{as3})$$

$$II' = \mathcal{T}(II, \Sigma, x \mapsto \{\mathbf{f} : G, \overline{\mathbf{f} : v}\}, x \mapsto \{\overline{\mathbf{f} : v}\}) \\ \{II \mid \Sigma\} \xrightarrow{x := \text{new}} \{(II \mid \Sigma) * x \mapsto \{\}\} \quad (\text{as4})$$

$$x \text{ not free in } II \text{ or } \Sigma \\ \{II' \mid \Sigma * E \mapsto \{\overline{\mathbf{f} : v}\}\} \xrightarrow{\text{delete}(E)} \{(II \mid \Sigma)\} \quad (\text{as5})$$

The assignment rule **as3** requires  $II$  be transformed to  $II'$  via a  $\mathcal{T}$  operator, defined below. Since the assignment occurring in this rule alters the heap, it changes the meaning of atomic formulae  $l_1 \xrightarrow{\mathbf{f}} l_2$  in  $II$  and  $\mathcal{T}$  updates any such atomic assertions affected by this particular assignment. Note that  $\mathcal{T}$  is not a total function; in cases where it is not defined, the rule cannot apply.

Before defining  $\mathcal{T}$ , we need to establish two Lemmas. Many of the  $l_1 \xrightarrow{\mathbf{f}} l_2$  in  $II$  are in a context where  $l_1$  will never refer to the cell being modified. The following lemma allows us to simply carry these predicates from  $II$  to  $II'$ .

**Lemma 1.** *if  $s, h_1 \vdash \Sigma_1$ ,  $s, h_2 \vdash \Sigma_2$  and  $s, h'_2 \vdash \Sigma'_2$  then  $s, h_1 * h_2 \vdash l_1 \in \Sigma_1 \wedge l_1 \xrightarrow{\mathbf{f}} l_2 \mid \Sigma_1 * \Sigma_2$  if and only if  $s, h_1 * h'_2 \vdash l_1 \in \Sigma_1 \wedge l_1 \xrightarrow{\mathbf{f}} l_2 \mid \Sigma_1 * \Sigma'_2$*

Basically, the above Lemma states that if we know that  $l_1$  is outside of the portion of the heap being modified that the truth of the  $l_1 \xrightarrow{\mathbf{f}} l_2$  predicate remains unchanged.

For **as3**, there is a special case where we may have a predicate of the form  $x \xrightarrow{\mathbf{f}'} l_2$  in  $II$ . If  $\mathbf{f} = \mathbf{f}'$ , then we know the predicate is true if and only if  $s(l_2) = \mathcal{E}_s(G)$  and if  $\mathbf{f} \neq \mathbf{f}'$  then the predicate propagates from the post- to the pre- condition. The following lemma captures the case where  $\mathbf{f} \neq \mathbf{f}'$ .

**Lemma 2.** *if  $s, h_1 \vdash \Sigma$ ,  $s, h_2 \vdash x \mapsto \{\overline{\mathbf{f} : v}\}$ ,  $s, h'_2 \vdash x \mapsto \{\mathbf{f} : G, \dots\}$ , and  $\mathbf{f} \neq \mathbf{f}'$  then  $s, h_1 * h_2 \vdash x \xrightarrow{\mathbf{f}'} l \mid \Sigma_1 * x \mapsto \{\overline{\mathbf{f} : v}\}$  if and only if  $s, h_1 * h'_2 \vdash x \xrightarrow{\mathbf{f}'} l \mid \Sigma_1 * x \mapsto \{\mathbf{f} : G, \overline{\mathbf{f} : v}\}$ .*

Using the two Lemmas above, the following definition of  $\mathcal{T}$  can be justified as properly fixing any embedded assertions  $l_1 \xrightarrow{\mathbf{f}} l_2$  in  $II$ .

**Definition 15.**  $II' = \mathcal{T}(II, \Sigma, \Sigma_A, \Sigma_B)$  if and only if  $II' = \mathcal{T}'(II, \Sigma, \Sigma_A, \Sigma_B, \emptyset, \emptyset)$ . The two tail arguments here are members of  $\text{Var} \cup \text{Var}' \rightarrow \{\Sigma_1, \dots, \Sigma_n\}$ , for some fixed  $\Sigma_1 * \dots * \Sigma_n = \Sigma$ .  $II' = \mathcal{T}'(II, \Sigma, \Sigma_A, \Sigma_B, b, b')$  is the function defined inductively on the structure of  $II$  as follows.

If $\Pi$ is	then $\Pi'$ is	if the following condition holds:
$v \in \Sigma_i \wedge \Pi_1$	$v \in \Sigma_i \wedge \Pi'_1$	$\Pi'_1 = \mathcal{T}'(\Pi_1, \Sigma, \Sigma_A, \Sigma_B, b \cup (v \mapsto \Sigma_i), b' \cup (v \mapsto \Sigma_i))$
$\Pi_1 \wedge \Pi_2$	$\Pi'_1 \wedge \Pi'_2$	$\Pi'_1 = \mathcal{T}'(\Pi_1, \Sigma, \Sigma_A, \Sigma_B, b, b')$ and
$\Pi_1 \vee \Pi_2$	$\Pi'_1 \vee \Pi'_2$	$\Pi'_2 = \mathcal{T}'(\Pi_2, \Sigma, \Sigma_A, \Sigma_B, b, b')$
$\Pi = \neg \Pi_1$	$\Pi' = \neg \Pi'_1$	$\Pi'_1 = \mathcal{T}'(\Pi_1, \Sigma, \Sigma_A, \Sigma_B, b, b')$
$E_1 = E_2$	$E_1 = E_2$	
$\forall v \in \Sigma_i. \Pi_1$	$\forall v \in \Sigma_i. \Pi'_1$	$\Pi'_1 = \mathcal{T}'(\Pi_1, \Sigma, \Sigma_A, \Sigma_B, (v \mapsto \Sigma_i) \cup b / \text{dom}(b) - v, (v \mapsto \Sigma_i) \cup b' / \text{dom}(b') - v)$
$\exists v. \Pi_1$	$\exists v. \Pi'_1$	
$x_1 \xrightarrow{\mathbf{f}} x_2$	$x_1 \xrightarrow{\mathbf{f}} x_2$	$x_1 \in \text{dom}(b)$ and $b(x_1) = \Sigma_i$ for some $\Sigma_i \in \Sigma$ $b(x_1) = \Sigma_A$ and $b'(x_1) = \Sigma_B$ and $\Sigma_A = x_1 \mapsto \{\mathbf{f} : v, \mathbf{f}' : v'\}$ and $\Sigma_B = x_1 \mapsto \{\mathbf{f} : v\}$ and $\mathbf{f} \neq \mathbf{f}'$
$x_1 \xrightarrow{\mathbf{f}} x_2$	$x_2 = v'$	$\Sigma_A = x_1 \mapsto \{\mathbf{f} : v, \mathbf{f}' : v'\}$ and $\Sigma_B = x_1 \mapsto \{\mathbf{f} : v\}$ and $b(x_1) = \Sigma_A$ and $b'(x_1) = \Sigma_B$

The two auxiliary parameters in  $\mathcal{T}'$  are mappings representing the bindings of variables in  $\Pi$  and  $\Pi'$  respectively. Also note that the two conditions for the second to last row in the table above are justified by Lemmas 1 and 2 above, respectively.

### 4.3 Unfold rule

The **as3** rule described above often requires the mutated heap component to be unfolded. This is accomplished using the **unfold** rule defined below.

$$\frac{\Pi'_0 | \Sigma_0 * v_0 \mapsto (\overline{\mathbf{f} : v}) * \mathcal{R}_F(v_1, \bar{t}) * \dots * \mathcal{R}_F(v_n, \bar{t})}{\Pi_0 | \Sigma_0 * \mathcal{R}_F(v_0, \bar{t})} (\text{unfold})$$

where  $\Pi'_0 = \text{unf}(\Pi_0, \Sigma_0, \mathcal{R}_F(t, \bar{t}), t \mapsto (\overline{\mathbf{f} : v}) * \mathcal{R}_F(v_1, \bar{t}) * \dots * \mathcal{R}_F(v_n, \bar{t}))$ . We define *unf* in Figure 4.4. The purpose of *unf* is to modify the  $\Pi$  portion of an abstract state to match the changes in  $\Sigma$ . The key issue in making these changes is to modify the  $x \in \mathcal{R}_F(v_0, \bar{t}), \forall x \in \mathcal{R}_F(v_0, \bar{t}). \text{body}$  and  $\exists x \in \mathcal{R}_F(v_0, \bar{t}). \text{body}$  which exist in  $\Pi$ . The  $\mathcal{R}_F(v_0, \bar{t})$  needs to be replaced with its unfolded components and often subterms need to be duplicated.

### 4.4 Merging

The inference rule for if statements in Section 4.2 requires the pre-condition abstract states from their two branches to be identical. A  $\Pi | \Sigma$  pair thus needs to be generated such that  $\Pi | \Sigma$  implies  $\Pi_1 | \Sigma_1, \Pi_2 | \Sigma_2$ . Since the  $\Pi$  portion of the abstract state is simply first order logic, if  $\Sigma_1 = \Sigma_2$ , then we can generate  $\Pi = \Pi_1 \wedge \Pi_2$ .

The key issue that needs to be resolved is how to transform either  $\Sigma_1$  or  $\Sigma_2$  if they are not equal. In our example program the one rule we needed to introduce is a folding rule; other merge rules may also be needed in the future. The rule needed is as follows.

$$\frac{\Sigma * v \mapsto \{\overline{\mathbf{f} : x}\} * \mathcal{R}_S(x_1) * \dots * \mathcal{R}_S(x_n) | \Pi}{\Sigma * \mathcal{R}_S(v) | \Pi'} (\text{fold})$$

$\Pi'$	$=$	$unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$	is	defined	inductively	as	follows:
This holds				if	the following	holds	
$unf(\Pi_0 \wedge \Pi_1, \Sigma_u, \Sigma_t, \Sigma_e) = \Pi'_0 \wedge \Pi'_1$				$\Pi'_0 = unf(\Pi_0, \Sigma_u, \Sigma_t, \Sigma_e)$	and		
$unf(\Pi_0 \vee \Pi_1, \Sigma_u, \Sigma_t, \Sigma_e) = \Pi'_0 \vee \Pi'_1$				$\Pi'_1 = unf(\Pi_1, \Sigma_u, \Sigma_t, \Sigma_e)$			
$unf(\neg \Pi, \Sigma_u, \Sigma_t, \Sigma_e) = \neg \Pi'$				$\Pi' = unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$			
$unf(\forall v \in \Sigma'_t, \Pi, \Sigma_u, \Sigma_t, \Sigma_e) = \forall v \in \Sigma'_t, \Pi'$				$\Pi' = unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$	and	$\Sigma'_t \neq \Sigma_t$	
$unf(\exists v \in \Sigma'_t, \Pi, \Sigma_u, \Sigma_t, \Sigma_e) = \exists v \in \Sigma'_t, \Pi'$							
$unf(l_1 \xrightarrow{f} l_2, \Sigma_u, \Sigma_t, \Sigma_e) = l_1 \xrightarrow{f} l_2$							
$unf(E_1 = E_2, \Sigma_u, \Sigma_t, \Sigma_e) = E_1 = E_2$							
$unf(v \in \Sigma_i \wedge \Pi, \Sigma_u, \Sigma_t, \Sigma_e) = v \in \Sigma_i \wedge \Pi'$				$\Sigma = \dots * \Sigma_i * \dots$	and	$\Pi' = unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$	
$unf(x' \in \Sigma_t \wedge \Pi, \Sigma_u, \Sigma_t, \Sigma_e) =$ $x \in (v_0 \mapsto \{f : v\} * R_F(v_1, \bar{t}) * \dots * R_F(v_n, \bar{t})) \wedge \Pi'$				$\Pi' = unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$			
$unf(x \in \Sigma_t \wedge \Pi, \Sigma_u, \Sigma_t, \Sigma_e) =$ $x_0 \in v_0 \mapsto \{f : v\} \wedge$ $x_1 \in R_F(v_1, \bar{t}) \wedge \dots \wedge$ $x_n \in R_F(v_n, \bar{t}) \wedge$ $(\Pi' [x_0/x] \wedge \dots \wedge \Pi' [x_n/x])$				$x$ is not a primed variable.			
$unf(\forall x \in \Sigma_t, \Pi, \Sigma_u, \Sigma_t, \Sigma_e) =$ $\Pi' [x_0/x] \wedge$ $\forall x \in R_F(v_1, \bar{t}), \Pi' \wedge$ $\dots$ $\forall x \in R_F(v_n, \bar{t}), \Pi'$				$\Pi' = unf(\Pi, \Sigma_u, \Sigma_t, \Sigma_e)$			
$unf(\exists x \in \Sigma_t, \Pi, \Sigma_u, \Sigma_t, \Sigma_e) =$ $\Pi' [x_0/x] \vee$ $\exists x \in R_F(v_1, \bar{t}), \Pi' \vee$ $\dots$ $\exists x \in R_F(v_n, \bar{t}), \Pi'$				$x_0, \dots, x_n$ are fresh			
				$\Sigma_e = v_0 \mapsto \{f : v\} * R_F(v_1, \bar{t}) * \dots * R_F(v_n, \bar{t})$			

**Fig. 3.** The definition of  $unf$

where  $\bar{f} = S$  and  $\Pi' = foldPi(\Pi, \Pi, \mathcal{R}_S(v), \{\mathcal{R}_S(x_1), \dots, \mathcal{R}_S(x_n)\})$ . The auxiliary  $foldPi$  function is now given.

**Definition 16.** The function  $\Pi' = foldPi(\Pi, \Pi^r, \mathcal{R}_S(v), \bar{\Sigma})$  is defined by the table below:

If $\Pi$ is	then $\Pi'$ is	if the following condition holds:
$\Pi_1 \wedge \Pi_2$	$\Pi'_1 \wedge \Pi'_2$	$\Pi'_1 = foldPi(\Pi_1, \Pi^r, \mathcal{R}_S(v), \bar{\Sigma})$ and
$\Pi_1 \vee \Pi_2$	$\Pi'_1 \vee \Pi'_2$	$\Pi'_2 = foldPi(\Pi_2, \Pi^r, \mathcal{R}_S(v), \bar{\Sigma})$
$\neg \Pi_1$	$\neg \Pi'_1$	$\Pi'_1 = foldPi(\Pi_1, \Pi^r, \mathcal{R}_S(v), \bar{\Sigma})$
$E_1 = E_2$	$E_1 = E_2$	
$v \in \Sigma' \wedge \Pi_1$	$v \in \Sigma' \wedge \Pi'_1$	$\Sigma' \notin \bar{\Sigma}$
$\forall v \in \Sigma', \Pi_1$	$\forall v \in \Sigma', \Pi'_1$	$\Sigma' \notin \bar{\Sigma}$
$\exists v \in \Sigma', \Pi_1$	$\exists v \in \Sigma', \Pi'_1$	$\Sigma' \notin \bar{\Sigma}$
$\forall v \in \Sigma', \Pi_1$	$\forall v \in \mathcal{R}_S(v), \Pi'_1$	$\Sigma' \in \bar{\Sigma}$ and $\Pi^r \rightarrow \forall v \in \mathcal{R}_S(v), \Pi'_1$ is valid
$\exists v \in \Sigma', \Pi_1$	$\exists v \in \mathcal{R}_S(v), \Pi'_1$	$\Sigma' \in \bar{\Sigma}$ and $\Pi^r \rightarrow \exists v \in \mathcal{R}_S(v), \Pi'_1$ is valid
$x_1 \xrightarrow{f} x_2$	$x_1 \xrightarrow{f} x_2$	

## 4.5 Simplification and strengthening rules

Beyond the above rules there are a large number of rules needed for more basic reasoning, too many to put in this paper (but of course all are in the Coq development). We for example have the following axioms to simplify spatial heap descriptions.

$$\Sigma * \text{emp} = \Sigma \qquad \mathcal{R}_R(\text{nil}, \bar{f}) = \text{emp}$$

For the  $\Pi$  portion, there are all the usual propositional logic rules (for example,  $a \wedge \text{true} \rightarrow a$ ). Axioms for  $\rightsquigarrow$  and  $\exists$  for example include the following:

$$\neg(\text{nil} \overset{f}{\rightsquigarrow} x) \qquad E \in \Sigma \rightarrow \exists x \in \Sigma. x = E$$

## 5 Outline of Coq proof of the example

To give a feel for how the correctness proof for the example in Section 2 was performed in Coq, we highlight a few portions of the proof here. We start by outlining the first few steps in which Hoare axioms are used to decompose the proof obligation for the entire program into obligations that implement back chaining on the individual statements. We then show how an unfold is done to allow an `as3` axiom to be applied.

### 5.1 Using Hoare axioms to back chain

We walk through the first few Coq steps in the the correctness proof. We start with a single proof obligation for the following Hoare triple:

$$\begin{aligned} & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) | \\ & \quad \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \} \\ & \text{while}(t)\{ \dots \} \\ & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) | t = \text{nil} \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \} \end{aligned}$$

We first apply our while loop axiom `absWhile` with the invariant:

$$\begin{aligned} & \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) | \\ & \quad \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \end{aligned}$$

This generates three proof obligations corresponding to the three preconditions of the rule. The first is the most interesting—it is the obligation to prove the correctness of the body:

$$\begin{aligned} & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) | \\ & \quad \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \wedge t \neq \text{nil} \} \\ & \mathbf{n} = \mathbf{p}; \\ & \dots \\ & \text{if}(t \rightarrow \mathbf{1} == \text{NULL} \ \&\& \ t \rightarrow \mathbf{r} == \text{NULL})\{ \dots \} \text{else}\{ \dots \} \\ & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) | \\ & \quad \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \overset{p}{\rightsquigarrow} m \} \end{aligned}$$

The other two obligations verify that the pre-condition implies the invariant and that the invariant implies the post condition.

We then use the `absSeq` rule to break up the sequence of statements into individual proof obligations. We use Coq's `eapply` tactic which leaves the pre-conditions existentially quantified. We first get an obligation for the if block which appears as follows:

$$\begin{aligned} & \{?1|?2\} \\ & \text{if } (\mathbf{t} \rightarrow \mathbf{1} == \text{NULL} \ \&\& \ \mathbf{t} \rightarrow \mathbf{r} == \text{NULL}) \{ \dots \} \text{else} \{ \dots \} \\ & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) \} \\ & \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \} \end{aligned}$$

These existential obligations `?1|?2` get filled in as we back trace individual statements. As an example, one of the obligations we will eventually work down to is an obligation for an assignment that looks like this:

$$\begin{aligned} & \{?1|?2\} \\ & \mathbf{t} = \mathbf{t} \rightarrow \mathbf{1} \\ & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) \} \\ & \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \} \end{aligned}$$

Here we apply `as2`, and then apply a number of of auxiliary axioms. This produces the following pre-condition.

$$\begin{aligned} & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) \} \\ & t \xrightarrow{l} v' \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \} \end{aligned}$$

which replaces the existential variables.

## 5.2 Unfolding prior to applying `as3`

Often we need to do an unfolding prior to performing a back chaining operation. Consider the following proof obligation:

$$\begin{aligned} & \{?1|?2\} \\ & \mathbf{i} \rightarrow \mathbf{t} = \mathbf{x}; \\ & \{ \mathcal{R}_{lr}(r) * \mathcal{R}_n(i) * \mathcal{R}_n(p) \} \\ & t \in \mathcal{R}_{lr}(r) \wedge t \xrightarrow{l} v' \wedge \forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \\ & \forall n \in \mathcal{R}_n(p). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \} \end{aligned}$$

The problem here is that the assignment requires us to reason about an individual cell within the  $\mathcal{R}_n(i)$  construct. Ideally, we would like to use the `as3` rule to generate the precondition. However, the rule will not match with the abstract state in its current form. We need to apply our `unfold` rule from section 4.3. This yields the following:

$$\begin{aligned} & \{?1|?2\} \\ & \mathbf{i} \rightarrow \mathbf{t} = \mathbf{x}; \\ & \{ i \mapsto \{ n : n, ?3 \} * \mathcal{R}_{lr}(r) * \mathcal{R}_n(n) * \mathcal{R}_n(p) \} \\ & t \in \mathcal{R}_{lr}(r) \wedge t \xrightarrow{l} v' \wedge \forall n \in \mathcal{R}_n(p). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \\ & (\forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m) \wedge \exists m \in \mathcal{R}_{lr}(r). i \xrightarrow{p} m \} \end{aligned}$$

Now we can apply our **as3** rule to fill in the precondition yielding the following:

$$\begin{aligned}
& \{i \mapsto \{n : n\} * \mathcal{R}_{lr}(r) * \mathcal{R}_n(n) * \mathcal{R}_n(p) | \\
& \quad t \in \mathcal{R}_{lr}(r) \wedge t \xrightarrow{l} v' \wedge \forall n \in \mathcal{R}_n(p). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \\
& \quad (\forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m) \wedge \exists m \in \mathcal{R}_{lr}(r). i \xrightarrow{p} m\} \\
& \quad \mathbf{i} \rightarrow \mathbf{t} = \mathbf{x}; \\
& \{i \mapsto \{n : n, t : x\} * \mathcal{R}_{lr}(r) * \mathcal{R}_n(n) * \mathcal{R}_n(p) | \\
& \quad t \in \mathcal{R}_{lr}(r) \wedge t \xrightarrow{l} v' \wedge \forall n \in \mathcal{R}_n(p). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m \wedge \\
& \quad (\forall n \in \mathcal{R}_n(i). \exists m \in \mathcal{R}_{lr}(r). n \xrightarrow{p} m) \wedge \exists m \in \mathcal{R}_{lr}(r). i \xrightarrow{p} m\}
\end{aligned}$$

We have shown a few illustrative steps in this section to give a flavor of how Coq verifications using our framework are constructed.

## 6 Conclusions

We have constructed extensions to traditional separation logic that allow it to reason about more complex data structures than previous work. The example tree traversal program we presented contains data structures with fairly complex pointer relationships. Notably, the traversal is represented as a linked list of pointers into a tree; existing separation logics cannot handle such a relationship. The invariant proof of this program that was formalized in Coq guarantees that the program maintains the integrity of this data structure and there are no errors such as memory leaks or dangling pointer references. These more complex data structures present a much greater verification challenge, and here we have shown the feasibility of using Coq to fully verify programs with these complex data structure invariants. To accomplish our goals we needed to develop novel axioms to fold and unfold our new recursive separation logic constructs which are in a form amenable to formal verification.

Currently there are many manual steps needed to complete the proof. We anticipate that much of this can be automated through the development of tactics, and that is a subject of future work. There already has been progress in developing tactics for separation logic proofs in Coq [12, 1] that we can build on. Our long-term goal is to get to the point where the only user input required will be the pre- and post- conditions in the form of abstract states as well as loop invariants. We anticipate that programs will be annotated with these conditions in a manner similar to [4]. Because the language for the abstraction is much richer than many of the previous works, automatic abstraction from a concrete domain such as [9] is likely to be difficult.

Another future goal is to formally verify our new axioms in Coq, following the program of [2] where the core axioms of separation logic are verified in Coq. It should be pointed out that there are many related projects pursuing formal theorem proving of imperative program properties in separation logic, including [11, 17] which are more ambitious in terms of the examples verified; our contribution here is to extend the potential for future verifications to more complex heap shapes.

**Acknowledgments** The authors would like to thank Mike Hicks, Henny Sipma and Steven Magill for reviewing drafts of this paper. The first author would also like to thank Dino Distefano, Honseok Yang and other members of the East London Massive for early feedback and pointing out key research.

## References

1. A. W. Appel. Tactics for separation logic, 2006. Early draft.
2. Andrew Appel and Sandrine Blazy. Separation logic for small-step cminor. In *Theorem Proving in Higher Order Logics*, volume 4732 of *Lecture Notes in Computer Science*, pages 5–21. 2007.
3. Yves Bertot and Pierre Casteran. *Interactive Theorem Proving and Program Development*.
4. Bor-Yuh Evan Chang, Xavier Rival, and George C. Necula. Shape analysis with structural invariant checkers.
5. Peter W. O’Hearn Dino Distefano and Hongseok Yang. A local shape analysis based on separation logic. In *TACAS*, 2006.
6. S. P. Rahul Westley Weimer George C. Necula, Scott McPeak. Cil: Intermediate language and tools for analysis and transformation of c programs. In *Proceedings of Conference on Compiler Construction*, March 2002.
7. Bolei Guo, Neil Vachharajani, and David I. August. Shape analysis with inductive recursion synthesis. In *PLDI*, 2007.
8. B Cook D Distefano PW O’Hearn T Wies J Berdine, C Calcagno and H Yang. Shape analysis for composite data structures. In *CAV’07*. Springer, 2007.
9. Stephen Magill, Aleksandar Nanevski, Edmund Clarke, and Peter Lee. Inferring invariants in separation logic for imperative list-processing programs. In *SPACE*, 2006.
10. Nicolas Marti and Reynald Affeldt. A certified verifier for a fragment of separation logic. *Information and Media Technologies*, 4(2):304–316, 2009.
11. Nicolas Marti, Reynald Affeldt, and Akinori Yonezawa. Formal verification of the heap manager of an operating system using separation logic. In Zhiming Liu and Jifeng He, editors, *Formal Methods and Software Engineering*, volume 4260 of *Lecture Notes in Computer Science*, pages 400–419. 2006.
12. Andrew McCreight. Practical tactics for separation logic. In *Proceedings of the 22nd International Conference on Theorem Proving in Higher Order Logics*, TPHOLs ’09, pages 343–358, 2009.
13. Magnus Myreen. Separation logic adapted for proofs by rewriting. In *Interactive Theorem Proving Conference*. Springer, 2010.
14. Hongseok Yang Peter O’Hearn, John Reynolds. Local reasoning about programs that alter data structures. In *Proceedings of CSL’01*, volume LNCS 2142, pages 1–19, 2001.
15. B. Cook G. Ramalingam M. Sagiv R. Manevich, J. Berdine. Shape analysis by graph decomposition.
16. Reinhard Wilhel Shmuel Sagiv, Thomas W. Reps. Solving shape-analysis problems in languages with destructive updating. 20(1):1–50, 1998.
17. Harvey Tuch, Gerwin Klein, and Michael Norrish. Types, bytes, and separation logic. In *Proceedings of the 34th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL ’07, pages 97–108, 2007.