

Using the Coq theorem prover to verify complex data structure invariants

Kenneth Roe

Department of Computer Science
The Johns Hopkins University
Email: roe@cs.jhu.edu

Scott Smith

Department of Computer Science
The Johns Hopkins University
Email: scott@jhu.edu

Abstract—While automated static analysis tools can find many useful software bugs, there are many important bugs that are beyond the reach of these tools. Most large software systems have complex data structures with complex invariants, and many bugs can be traced to code that does not maintain these invariants. Additionally, often the invariants are not well documented. However, maintaining them is necessary for correct operation of the software. These invariants cannot be easily inferred by automated tools. One must use an interactive system in which developers first enter these invariants to document their software and then use a theorem prover to verify their correctness.

We describe the PEDANTIC framework for verifying the correctness of C-like programs using the Coq theorem prover. PEDANTIC is designed to prove invariants over complex dynamic data structures such as inter-referencing trees and linked lists. The language for the invariants is based on separation logic. The PEDANTIC tactic library has been constructed to allow program verifications to be done with reasonably compact proofs.

We have completed the verification of a tree traversal program (though the correctness proofs of some of the tactics is incomplete). We are currently working on the verification of a C implementation of the DPLL algorithm in order to demonstrate the utility of the framework on more complex programs. Verifying programs using an interactive theorem prover is quite tedious. We discuss work being done to improve proof development productivity.

I. INTRODUCTION

Many software bugs can be traced to data structure invariant violations. These invariants can be quite complex. They are created as part of the process of writing a program. As an example, a SAT solver algorithm may use both a stack (represented as a linked list) and an array to represent variable assignments. The stack allows the program to efficiently remove assignments in the reverse order that they were created and the array allows the program to quickly find the assignment for a particular variable. Many efforts in the research literature aim at creating fully automated static analysis tools, including [29], [26], [22]. Each tool while being automatic has substantial limitations. [29] only covers a subset of separation logic and not the separation logic augmented with invariants covered in this paper. [26] cannot easily verify invariants with non-linear equations as underlying SMT solvers are limited in this area. [22]’s algorithm only works on a subset of the examples presented in the paper. Finding data structure invariants is a difficult task and hence no program will be able to automatically find and verify all of them. Hence, we

focus on creating a tool that allows the user to first document data structure invariants and then verify them.

This paper presents first steps towards a framework for verifying programs in imperative languages such as C that use complex heap based data structures. Our extensible PEDANTIC (Proof Engine for Deductive Automation using Non-deterministic Traversal of Instruction Code) framework is implemented in Coq and is based on ideas from many other separation logic frameworks such as [8], [4], [20]. The framework implements a logic language based on separation logic for describing program invariants, and a set of tactics to propagate these assertions through the program.

One of the key challenges of designing this system is to find a good division of labor between the user and the tool. The tool obviously cannot automate everything. However, if it automates too little, then the verification process will be too tedious and developers will not use the tool. As part of our effort, we have developed a new IDE for Coq, CoqPIE, aimed at optimizing the effort needed for complex proof development tasks. This project is described in [31]. This paper concentrates on the Framework we developed in Coq for verifying data structure invariants. In our design, data structure invariants need to be entered by the user. Our experience (both from the example presented as well as a DPLL verification presently being done) shows that the size of these invariants tend not to be larger than the size of the program being verified. One key difficulty in many formal systems is inferring loop invariants. Our experience has shown that loop invariants are usually very similar to the pre- and post- condition invariants. Since the task of automatically generating a loop invariant is difficult, we have made the decision to require the user to enter all loop invariants.

Forward propagation of an assertion over a block of statements can be automated. At the end of the block, often there will be a post condition (which often looks very much like the precondition). However, the mechanical steps performed in the forward chaining produce an assertion that does not look like the post-condition. The user will likely have to perform many manual theorem proving steps to prove that the result of the forward propagation implies the post condition.

```

struct list {
  struct list *n;
  struct tree *t;
  int data;
};

struct tree {
  struct tree *l, *r;
  int value;
};

struct list *p;
void build_pre_order(struct tree *r) {
  l1 struct list *i = NULL, *n, *x;
  l2 struct tree *t = r;
  l3 p = NULL;
  l4 while (t) {
  l5   n = p;
  l6   p = malloc(sizeof(struct list));
  l7   p->t = t;
  l8   p->n = n;
  l9   if (t->l==NULL && t->r==NULL) {
  l10    if (i==NULL) {
  l11     t = NULL;
  l12    } else {
  l13     struct list *tmp = i->n;
  l14     t = i->t;
  l15     free(l);
  l16     i = tmp;
  l17    }
  l18  } else if (t->r==NULL) {
  l19    t = t->l;
  l20  } else if (t->l==NULL) {
  l21    t = t->r;
  l22  } else {
  l23    n = i;
  l24    i = malloc(
  l25      sizeof(struct list));
  l26    i->n = n;
  l27    x = t->r;
  l28    i->t = x;
  l29    t = t->l;
  l30  }
  l31 }

```

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

A. Overview of PEDANTIC contributions

In order to create an easy to use yet extensible framework, we took the work of [8], [4], [20], [25], [3] and introduced a number of enhancements. We discuss four areas of development: (1) use of a deep model; (2) use of pointers in the functional representation of a data structure; (3) introduction of a merge tactic to merge together branches at the end of an if-then-else construct; and (4) use of simplification after other operations to put assertions into a canonical form.

1) *Deep model*: A deep model means that rather than representation our invariants directly as propositions in the Coq theorem prover, we created a data structure which is the AST for the invariants. We then created a number of Gallina functions to manipulate these data structures. This gives us greater flexibility in designing tactics. It also separates the design of the algorithm from the verification of the tactic. Once a correctness theorem for a tactic is created in the framework, it never has to be regenerated in its application.

2) *Embedding pointers in a functional representation*: The Btree example in [20] illustrates how embedding pointers in a functional representation can be used to simplify the expression of invariants. We expand on this technique by showing how it can be used to represent cross referencing

relationships from one data structure to another.

3) *Merging and pairing*: At the end of an if-then-else block, our forward propagation tactics will have produced two distinct state assertions. As these assertions are derived from the same starting point, they will be similar but not the same. We have developed a merge tactic which works by first pairing off components in the states that are the same, and then proceeding to merge other components through more complex operations. In this process it may be necessary to invoke fold or unfold tactics to align recursive data structures to facilitate the merge. We also use this pairing technique for proving assertion entailment properties.

4) *Simplification*: Often the results of forward propagating assertions through a program (or performing other operations) produces fairly complicated state assertions. PEDANTIC introduces a fairly sophisticated simplification algorithm to reduce these assertions to a more canonical form.

5) *Framework extension*: While our framework provides a number of useful constructs and tactics for reasoning about programs, we anticipate that many programs will require customized constructs. Programming styles vary greatly from system to system and trying to design generic constructs for all systems is impossible. Our framework allows for extensibility. Theorems and proofs in the framework are parameterized so that when new constructs are added, only simple proofs of certain properties are needed to verify the integrity of the extended framework. For example, our framework currently provides a Tree construct for linked lists or trees. This construct is parameterized to allow for a few basic variants. While the Tree construct works well for many data structures, there are many cases where it may not work well. For example a linked list with nodes of different types or a graph data structure in which multiple nodes point to the same child. For the former, the linked list may contain records with are either the name of a person or business. Depending on the type, the record may be of a different size. The Tree construct in the framework assumes all records are of the same size. However, the framework is setup so that one could easily add a Tree2 construct (which is a slight variation of Tree) which handles the different sized records.

PEDANTIC is designed to be extensible. We use a deep encoding for defining abstract program states in Coq, which allows use of Galina functions to transform our assertions. The data structure type declarations for abstract states has a generic form for function calls and predicates. This allows us to add new functions and predicates without changing the data structure. We have designed our proofs so that they can be reused as new constructs are added. For example, if the generic tree predicate does not quite fit for a particular verification, a new predicate can be created and added to the system without needing to make major changes to the system's infrastructure.

II. A SAMPLE PROGRAM AND ITS INVARIANT

PEDANTIC excels at reasoning about cross-structure invariants, and we give a small program example here which contains such invariants that PEDANTIC can verify. Figure 1

```

Inductive absExp : Type :=
| AbsConstVal : (@Value unit) -> absExp
| AbsVar : id -> absExp
| AbsQVar : absVar -> absExp
| AbsFun : id -> list absExp -> absExp.

Inductive absState :=
| AbsExists : absExp -> absState -> absState
| AbsAll : absExp ev eq f -> absState -> absState
| AbsCompose : absState -> absState -> absState
| AbsEmpty : absState
| AbsLeaf : id -> (list absExp) -> absState
| AbsMagicWand : absState -> absState -> absState
| AbsUpdateVar : id -> absExp -> absState -> absState
| AbsUpdateLoc : absExp -> absExp -> absState -> absState
| AbsUpdateWithLoc : id -> absExp -> absState -> absState
| AbsUpdState : absState -> absState -> absState -> absState
...

```

Fig. 2. Separation expression and assertion data structures, from file `AbsState.v`.

shows a C-like program which performs a traversal of a tree, given in the parameter `r` to the function `build_pre_order` and places the result in a linked list `p`. This function contains a pointer `t` which walks the tree. As the tree is walked, elements are added to the `p` list.

Our Coq proof verifies a number of important properties of the data structures in this program, including: (1) the program maintains two well formed linked lists, the heads of which are pointed to by `n` and `p`; (2) the program maintains a well formed tree pointed to by `r`; (3) `t` always points to an element in the tree rooted at `r`; (4) the two lists and the tree are disjoint in memory; (5) no other heap memory is allocated; and, (6) the `t` field of every element in both list structures points to an element in the tree. Invariant 6, in particular is a cross-structure invariant which PEDANTIC is highly suited to verify.

A. Separation Invariants

In this section we describe the internal structure of PEDANTIC in more detail.

Figure 2 shows the Coq definitions used for our state assertions. We use the abbreviation $a * b$ for `AbsCompose a b`. We use de Bruijn indices for bound variables, so `AbsAll/AbsExists` do not specify the names of the variable they introduce. We use sugared quantifier notation with variables v_0, v_1, \dots below for readability.

Shallow embeddings such as `Bedrock`[8] allow one to create arbitrary recursive data structure invariants by defining a Coq function that can be used in a separation logic based assertion. With a deep embedding where the logic is defined as a data structure (rather than a Coq expression), it is non-trivial to make the framework extensible. To address this problem, we have generic function (`AbsFun`) and predicate (`AbsLeaf`) constructs in the separation logic grammar. The semantics provide a predefined set of functions and predicates. However, new functions and predicates can easily be added as the separation logic data structure does not need to be modified.

The `absFun` and `absLeaf` datatypes are given semantics by the Coq functions `basicEval` and `basicState`. The `basicState` function defines three types of leaf predicates, `AbsPredicate`, `AbsCell` and `AbsTree`.

```

(e, h) ⊢ AbsCompose s1 s2      iff ∃h1 h2, dom(h1) ∩ dom(h2) = ∅ ∧
                               (e, h1) ⊢ s1 ∧ (e, h2) ⊢ s2 ∧ h = h1 ∪ h2
(e, h) ⊢ AbsMagicWand s1 s2   iff ∃h1 h2, dom(h) ∩ dom(h2) = ∅ ∧
                               (e, h1) ⊢ s1 ∧ (e, h2) ⊢ s2 ∧ h1 = h ∪ h2
(e, h) ⊢ AbsEmpty             iff dom(h) = ∅
(e, h) ⊢ AbsExists v.s        iff ∃x ∈ V.(e, h) ⊢ s[x/v]
(e, h) ⊢ AbsAll v ∈ R(p̄).s    iff ∀x ∈ R(⟨⟨ p̄ ⟩⟩ e)
(e, h) ⊢ AbsUpdateVar i v s   iff (e[i] = ⟨⟨ v ⟩⟩ e ∧ ∃x ∈ V.(e[x/i], h) ⊢ s
(e, h) ⊢ AbsUpdateWithLoc i v s iff (e[i] = h(⟨⟨ v ⟩⟩ e) ∧ ∃x ∈ V.(e[x/i], h) ⊢ s
(e, h) ⊢ AbsUpdateLoc l v s   iff (⟨⟨ v ⟩⟩ e = h(⟨⟨ l ⟩⟩ e) ∧
                               ∃x ∈ V.(e, h[x/⟨⟨ l ⟩⟩ e]) ⊢ s
(e, h) ⊢ AbsUpdState s1 s2 s3 iff (e, h) ⊢ AbsCompose (AbMagicWand s1 s2) s3
(e, h) ⊢ [P]                  iff ⟨⟨ P ⟩⟩ e ≠ 0 and dom(h) = ∅
(e, h) ⊢ l ↦ v                 iff h(⟨⟨ l ⟩⟩ e) = ⟨⟨ v ⟩⟩ e ∧
                               ∀x ∈ N, x ≠ ⟨⟨ l ⟩⟩ e → x ∉ dom(h)
(e, h) ⊢ TREE(r, v, n, f̄)     iff (⟨⟨ r ⟩⟩ e = 0 ∧ h = ∅) ∨
∃h0, …, hn+m, x0, …, xn-1, v0, …, vm. ⟨⟨ r ⟩⟩ e ≠ 0 ∧
(e, h0) ⊢ (⟨⟨ r ⟩⟩ e + 0) ↦ x0 ∧ … ∧ (e, hn-1) ⊢ (⟨⟨ r ⟩⟩ e + n - 1) ↦
xn-1 ∧
(e, hn) ⊢ TREE(h(⟨⟨ r ⟩⟩ e + f0), vf0, s, f̄) ∧ … ∧
(e, hn+m-1) ⊢ TREE(h(⟨⟨ r ⟩⟩ e + fm-1), vfm-1, s, f̄) ∧
combine(h, h0, …, hn+m-1) ∧
v = [⟨⟨ r ⟩⟩ e, v0, …, vn-1] ∧ ∀i < n. i ∉ f̄ implies vi = xi
where combine(h, h0, …, hn) iff ∀i, j ≤ n. dom(hi) ∩ dom(hj) = ∅ ∧
∀i ∈ dom(h). ∃j. i ∈ dom(hj) ∧ h(i) = hj(i)

```

Fig. 3. Semantics for the separation predicates of the three constructs defined by `basicState` for `absLeaf`. (e, h) represents a concrete state. $(e, h) \vdash s$ means that (e, h) satisfies the assertion s . $e \in \text{Id} \rightarrow \text{Nat}$ represents assignments for program variables and $h = \text{Nat} \rightarrow \text{option Nat}$ represents the values on the heap. The notation $\ll exp \gg e$ represents evaluation of the expression `exp` with the set of variables in the environment e .

`AbsLeaf AbsPredicate P`, denoted $[P]$, indicates an arbitrary predicate on the heap, `AbsLeafAbsCell l v`, abbreviated $l \mapsto v$, asserts that cell l has contents v , and `AbsLeaf AbsTree`, abbreviated `TREE`, is used for characterizing recursive data structures on the heap, either lists or trees. By defining a fixed grammar for tree structure assertions we can in tandem define generalized fold and unfold tactics which will work over different recursive data structures in different derivations that are defined with `AbsTree`. The semantics for these three predicates and some of the core separation logic are given in Figure 3.

B. Expressions and values

The values in our logic can be either integers or lists of values as represented by the following Coq data type:

```

Inductive Value :=
| NatValue : Int -> Value
| ListValue : (list Value) -> Value.

```

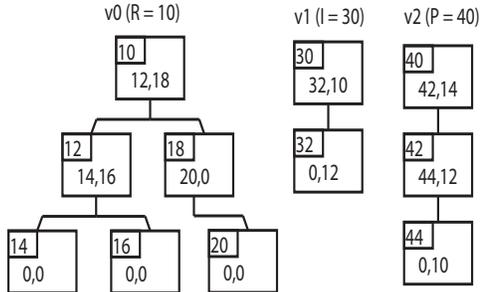
Expressions in our system consist of basic operators on the value structure as well as specialized operations to facilitate our logic. The semantics of these operators are given in figure 4. A more detailed explanation of the latter is given in the next section where we outline the verification of our tree traversal invariant.

C. An example state assertion

To better understand state assertions we present a sample heap snapshot of our example program in Figure 5. This example shows the heap memory after two loop iterations have

$\ll c \gg e b$	$= c$	constant
$\ll v \gg e b$	$= e[v]$	environment variable
$\ll v \gg e b$	$= b[v]$	quantifier variable
$\ll l_1 + l_2 \gg e b$	$= \text{NatValue } v_1 + v_2$	iff $\text{NatValue } v_1 = \ll l_1 \gg e b \wedge$ $\text{NatValue } v_2 = \ll l_2 \gg e b$
	NoValue	otherwise
$\ll l_1 * l_2 \gg e b$	$= \text{NatValue } v_1 * v_2$	iff $\text{NatValue } v_1 = \ll l_1 \gg e b \wedge$ $\text{NatValue } v_2 = \ll l_2 \gg e b$
	NoValue	otherwise
$\ll l_1 = l_2 \gg e b$	$= \text{NatValue } 1$	iff $\ll l_1 \gg e b = \ll l_2 \gg e b$
	$\text{NatValue } 0$	otherwise
$\ll l_1 < l_2 \gg e b$	$= \text{NatValue } 1$	iff $\text{NatValue } v_1 = \ll l_1 \gg e b \wedge$ $\text{NatValue } v_2 = \ll l_2 \gg e b \wedge v_1 < v_2$
	NoValue	otherwise
	$\text{NatValue } 0$	iff $\text{NatValue } v_1 = \ll l_1 \gg e b \wedge$ $\text{NatValue } v_2 = \ll l_2 \gg e b \wedge v_1 \not< v_2$
	NoValue	otherwise
...		
$\ll \text{find}(l, v) \gg e b$	$= F(\ll l \gg e b, \ll v \gg e b)$	iff $v = \text{ListValue}[l, \dots]$
where	$F(l, v) = v$	iff $v = \text{ListValue}[l', v_1, \dots, v_n] \wedge$ $v_i = \text{ListValue}[l, \dots] \wedge$ $i \in [1, \dots, n]$
	$F(l, v) = v_i$	otherwise
	$F(l, v) = \text{NoValue}$	
$\ll \text{rangeSet}(f) \gg e b$	$= \text{RS}(\ll f \gg e b)$	iff $v = \text{ListValue}[l, \dots]$
where	$\text{RS}(v) = l + +\text{RS}(v_1) + + \dots + +\text{RS}(v_n)$	otherwise
	$\text{RS}(v) = []$	iff $\ll s \gg e = \text{NatValue } st \wedge$ $\ll f \gg e b = \text{NatValue } en \wedge$
$\ll \text{rangeNumeric}(s, f) \gg e b$	$= \text{RN}(st, en)$	iff $v = \text{ListValue}[l, \dots]$
where	$\text{RS}(v) = [st, \dots, en]$	

Fig. 4. Semantics for expressions. The first three cases are for the AbsExp core and the rest are for some of the functions embedded in absFun.



```

AbsExists v3. AbsExists v2. AbsExists v1.
  TREE(R, v3, 2, [0, 1])*
  TREE(I, v2, 2, [0])*
  TREE(P, v1, 2, [0])*
  AbsAll v0 ∈ TreeRecords(v2).
    [nth(find(v2, v0), 2) inTree v3]*
  AbsAll v0 ∈ TreeRecords(v1).
    [nth(find(v1, v0), 2) inTree v3]*
    [T = 0 ∨ T inTree v3]

```

Fig. 5. A snapshot of a possible heap state of the program from Figure 1, and the general PEDANTIC assertion describing the state invariant for the loop over all executions.

been completed. To the right of the trees we give the general state invariant for the loop as we formalize it in Coq. The three TREE assertions characterize the tree and the two lists of the heap. The four parameters of TREE are: (1) the root of the tree; (2) a variable holding an equivalent functional representation (discussed below); (3) the word size of a record; (4) and, a list of offsets for all the child node pointers. Here the tree needs two such offsets, $[0, 1]$ and the lists only need one, $[0]$.

The functional representations $v_0/v_1/v_2$ defined by TREE characterize both the recursive list/tree structure and additionally embed information on the pointer structure, similar to

methods used in [20]. For the snapshot from the Figure, v_0 would for example be the list

$[10, ([12, [14, [0], [0]], [16, [0], [0]]), ([18, [20, [0], [0]], [0]])]$

Getting back to the invariant assertion, following the TREE assertions, the two AbsAll clauses assert for each list that the pointers in the second field of each node in the list point into the tree. TreeRecords is a function that returns the list of record pointers given a functional representation. For example, for the representation above, TreeRecords returns $[10, 12, 14, 16, 18, 20]$. x inTree y is shorthand for $x \in \text{TreeRecords}(y)$, and find takes an address in a tree and a functional representation of the tree (as shown above), and returns the subtree rooted at that address. For example, $\text{find}([30, [32, [0], 12], 10], 32)$ will return $[32, [0], 12]$. Finally, the last line states that either T is 0 or that T points to an element in R.

III. VERIFYING THE ASSERTIONS

We now show how to verify that the invariant of figure 5 holds throughout our program. Figure 9 shows the top level Coq proof that the invariant holds throughout the program in figure 1. Notice that the proof follows the structure of the program. The comments show the program that has been converted to Coq notation. One starts by forward propagating the invariant (called afterInitAssigns here). When we hit the while loop, we need to supply a loop invariant (called loopInv here). In this case, the invariant turns out to be exactly the same as the one given in figure 5.

This main proof generates 12 lemmas that need to be verified. treeRef1 and treeRef2 verify that a heap location being read belongs to an allocated block. storeCheck1

$$\begin{array}{c}
\frac{}{\{s\} x := e \{ \text{absUpdateVar } x \ e \ s \} d} \text{assignment} \\
\\
\frac{\text{canSave}(e + f, s)}{\{s\} x := e \rightarrow f \{ \text{AbsUpdateWithLoc } x \ (e + f) \ s \} \text{load}} \\
\text{where } \text{canSave}(l, s) = \forall eh. (e, h) \vdash s \rightarrow h(l) \neq \text{None} \\
\\
\frac{\text{canSave}(x + f, s)}{\{s\} x \rightarrow f := e \{ \text{AbsUpdateLoc } l \ x \ s \} \text{store}} \\
\\
\frac{l = e \quad \text{validExp}(e, s) \quad n : \text{nat} \quad \exists st.st \vdash s * - \quad (\exists v_0, \dots, v_{n-1}. l \mapsto v_0, \dots, l + n - 1 \mapsto v_{n-1})}{\{s\} \text{DELETE } e, n \{ s * - (\exists v_0, \dots, v_{n-1}. l \mapsto v_0, \dots, l + n - 1 \mapsto v_{n-1}) \} \text{del}} \\
\\
\frac{n : \text{nat}}{\{s\} v := \text{NEW } n \{ s * v \mapsto c_0 * \dots * v + n - 1 \mapsto c_{n-1} \} \text{new}}
\end{array}$$

Fig. 6. Forward propagation rules. The inference rules above are used to propagate over the statements of a program.

$$\begin{array}{c}
\frac{\text{absUpdateVar } v \ e \ s}{\text{absExists } x.s[x/v] * [v = e]} \text{absUpdateVar elim} \\
\\
\frac{\text{absUpdateWithLoc } v \ e \ s}{\text{AsExists } x.(s[x/v] * [v = q[x/v]]) \text{absUpdateWithLoc elim1}} \\
\text{where } s = \dots * e \mapsto q * \dots \\
\\
\frac{\text{absUpdateWithLoc } v \ e \ s}{\text{absUpdateWithLoc } v \ e \ s} \text{elim2} \\
\frac{[\text{absExists } x.ss[x/v] / \text{absUpdateWithLoc } v \ e' \ ss]}{\text{absUpdateWithLoc } v \ e \ s} \\
\text{where } s = \dots * \text{absUpdateWithLoc } v \ e' \ ss * \dots \\
\\
\frac{\text{absUpdateWithLoc } v \ e \ s}{s'} \text{absUpdateWithLoc elim3} \\
\text{where } s' \text{ is } s \text{ with all predicates involving } e \text{ removed.} \\
\text{This rule requires that } v \text{ only be used inside of predicates of the form } [\dots] \text{ within } s. \\
\\
\frac{\text{absUpdateLoc } l \ e \ s}{s[l \mapsto e/l \mapsto q]} \text{absUpdateLoc elim} \\
\text{where } s = \dots * l \mapsto q * \dots
\end{array}$$

Fig. 7. Rules for eliminating update constructs

and `storeCheck2` similarly verify that a heap location being written to is in an allocated block. `deleteExists1` verifies that a block being deallocated is actually an allocated block on the heap. `mergeTheorem1`, `mergeTheorem2`, `mergeTheorem3` and `mergeTheorem4` are used to merge the resulting states from forward propagating the then and else branches of an if-then-else. The theorems `implication1`, `implication2`, `implication3` verify that the state generated by forward propagation matches the post-condition at the end of a block.

$$\frac{s * - (\text{absExists } v.l_1 \mapsto v) * \dots * (\text{absExists } v.l_n \mapsto v)}{s[\text{AbsEmpty} / (l_1 \mapsto e_1, \dots, l_n \mapsto e_n)]} \\
\text{where } s = \dots * l_1 \mapsto e_1, \dots, l_n \mapsto e_n \dots$$

Fig. 8. Rule for eliminating magic wand

```

Theorem loopInvariant :
  {(afterInitAssigns)}loop{(afterWhile return nil with AbsNone)}.
Proof.
  (* Break up the while portion of the loop *)
  unfold loop. unfold afterWhile. unfold afterInitAssigns.

  (* WHILE ALnot (!T == A0) DO *)
  eapply strengthenPost.
  eapply whileThm with (invariant := loopInv). unfold loopInv.
  eapply strengthenPost.

  (* N := !P; *)
  eapply compose. pcrunch.
  (* NEW P,ANum(Size_1);*)
  eapply compose. pcrunch. simp. simp. simp.
  (* CStore (!P)++(ANum F_P) (!T) *)
  eapply compose. pcrunch.
  apply treeRef1. apply H. apply H0.
  (* CStore (!P)++(ANum F_n) (!N) *)
  eapply compose. pcrunch.
  apply treeRef2. apply H. apply H0. simp.
  (* CLoad Tmp_1 (!T)++ANum(F_1) *)
  eapply compose. pcrunch.
  (* CLoad Tmp_r (!T) ++ A1 *)
  eapply compose. pcrunch.

  (* IF (ALand (!Tmp_1 == A0) (!Tmp_r == A0)) *)
  eapply if_statement. simpl.

  (* IF (!I == A0) *)
  eapply if_statement. simpl.

  (* T := A0 *)
  pcrunch.

  (* ELSE *)

  (* CLoad T (!I)++A1 *)
  eapply compose. pcrunch.
  (* CLoad Tmp_1 (!I)++A0 *)
  eapply compose. pcrunch.
  (* DELETE !I, A2 *)
  eapply compose. pcrunch.
  eapply deleteExists1. apply H0.
  (* I := !Tmp_1 *)
  pcrunch. pcrunch. pcrunch. pcrunch. pcrunch.

  (* FI *)
  apply mergeTheorem1.

  (* (CIF (!Tmp_1 == A0) *)
  simpl.
  eapply if_statement.

  (* CLoad T (!T ++ A1) *)
  simpl. pcrunch.

  (* ELSE *)

  (* (CIF (!Tmp_r == A0) *)
  simpl. eapply if_statement.

  (* CLoad T (!T ++ A0) *)
  simpl. pcrunch.

  (* ELSE *)

  (* N := !I *)
  simpl. eapply compose. pcrunch.
  (* NEW I, A2 *)
  eapply compose. pcrunch.
  (* CStore (I +++ A0) (!N) *)
  eapply compose. pcrunch.
  apply storeCheck1. apply H. apply H0.
  (* CLoad Tmp_1 (! T +++ A1) *)
  eapply compose. pcrunch.
  (* CStore (! I +++ A1) (! Tmp_1) *)
  eapply compose. pcrunch.
  apply storeCheck2. apply H. apply H0.
  (* (CLoad T (! T +++ A0) *)
  pcrunch.

  (* FI *)
  pcrunch. pcrunch. pcrunch. pcrunch. pcrunch. pcrunch.
  apply mergeTheorem2.

  (* FI *)
  pcrunch.
  apply mergeTheorem3.

  (* FI *)
  pcrunch.
  apply mergeTheorem4.
  pcrunch. pcrunch. pcrunch. pcrunch. pcrunch. pcrunch.

  apply implication1.
  intros. inversion H. intros. inversion H.
  apply implication2. apply implication3.
  intros. apply H. intros. inversion H.
Qed.

```

Fig. 9. Top level Coq proof of our program. Notice how the structure of the proof matches the structure of the program

A. Forward propagation

Figure 6 shows the rules for forward propagating over a state. We have set up our framework so that these rules can be applied mechanically. Notice that the assertion before the statement in each rule does not need to be broken apart. We use `absUpdateVar`, `absUpdateLoc` and `absUpdateWithLoc` to represent updates. To show how this propagation works, below is shown the state before the assignment $N ::= !P$ in figure 9:

```
[ T = 0]*
AbsExists v3.AbsExists v2.AbsExists v1.
  TREE(R, v3, 2, [0, 1])*
  TREE(I, v2, 2, [0])*
  TREE(P, v1, 2, [0])*
AbsAll v0 ∈ TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v3]*
AbsAll v0 ∈ TreeRecords(v1).
  [nth(find(v1, v0), 2) inTree v3]*
[T = 0 ∨ T inTree v3]
```

After the assignment, we get the state shown below:

```
AbsUpdateVar N P
([ T = 0]*
AbsExists v3.AbsExists v2.AbsExists v1.
  TREE(R, v3, 2, [0, 1])*
  TREE(I, v2, 2, [0])*
  TREE(P, v1, 2, [0])*
AbsAll v0 ∈ TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v3]*
AbsAll v0 ∈ TreeRecords(v1).
  [nth(find(v1, v0), 2) inTree v3]*
[T = 0 ∨ T inTree v3])
```

We can then propagate over the next statement, `NEW P, ANum(Size_1)` and we get the state shown below:

```
AbsExists v7.AbsExists v6.AbsExists v5.AbsExists v4.
  v4 + 1 ↦ v6 * v + 4 ↦ v5 * P = v4*
AbsUpdateVar N v7
([ T = 0]*
AbsExists v3.AbsExists v2.AbsExists v1.
  TREE(R, v3, 2, [0, 1])*
  TREE(I, v2, 2, [0])*
  TREE(P, v1, 2, [0])*
AbsAll v0 ∈ TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v3]*
AbsAll v0 ∈ TreeRecords(v1).
  [nth(find(v1, v0), 2) inTree v3]*
[T = 0 ∨ T inTree v3])
```

B. Merging states (+fold/unfold and simplification)

Shown below are the two states from the ends of then- and else- blocks of the if statement when `mergeTheorem1` of figure 9 is applied. These are the states generated by the forward propagation of the two branches. They are both similar as they both were derived from the same root invariant. The merge process involves several steps. Hence, this merge example also illustrates simplification, unfolding and folding.

```
AbsUpdateVar T 0 ([I = 0] * [tmp_l = 0 ∧ tmp_r = 0])*
(AbsUpdateWithLoc tmp_r (T + 1)
(AbsUpdateWithLoc tmp_l (T + 0)
(AbsUpdateLoc P N (AbsUpdateLoc (P + 1) T
(AbsExists v6 (AbsExists v5 (AbsExists v4
(AbsUpdateVar N v2 (
P + 1 ↦ v6 * P ↦ v5 * [P = v4] * [0 < T]*
(AbsExists v3 (AbsExists v2 (AbsExists v1
TREE(R, v1, 2, [0, 1])*
TREE(I, v2, 2, [0])*
TREE(P, v3, 2, [0])*
AbsAll v0 ∈ TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v1]*
AbsAll v0 ∈ TreeRecords(v3).
  [nth(find(v3, v0), 2) inTree v1]*
[T inTree v1])))])))
```

```
AbsUpdateVar I tmp_l
(AbsUpdateWithLoc tmp_l T + 0
(AbsUpdateWithLoc T I + 1
([I ≠ 0] * [tmp_l = 0 ∧ tmp_r = 0])*
(AbsUpdateWithLoc tmp_r (T + 1)
(AbsUpdateWithLoc tmp_l (T + 0)
(AbsUpdateLoc P N (AbsUpdateLoc (P + 1) T
(AbsExists v6 (AbsExists v5 (AbsExists v4
(AbsUpdateVar N v2
P + 1 ↦ v6 * P ↦ v5 * [P = v4] * [0 < T]*
(AbsExists v3 (AbsExists v2 (AbsExists v1
TREE(R, v1, 2, [0, 1])*
TREE(I, v2, 2, [0])*
TREE(N, v3, 2, [0])*
AbsAll v0 ∈ TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v1]*
AbsAll v0 ∈ TreeRecords(v3).
  [nth(find(v3, v0), 2) inTree v1]*
[T inTree v1])))]))) * -
(AbsExists v2 (AbsExists v1 (AbsExists v0
v0 + 1 ↦ v2 * v0 + 0 ↦ v1 * [I = v0])))
```

We start by doing simplifications and transformations. The results are shown below. Rules for eliminating `updateVar`, `updateLoc` and `updateWithLoc` are shown in figure 7. There are many other simplification rules that correspond to simple tautologies. For example, the expression $x + 0$ will be replaced with x . We do not enumerate these rules.

```

(AbsExists v6 (AbsExists v5 (AbsExists v4
  (AbsExists v3 (AbsExists v2 (AbsExists v1
    [T = 0] * [I = 0] * (AbsEmpty * AbsEmpty)*
    [N = v5] * (P + 1)  $\mapsto$  v6 * P  $\mapsto$  N * [P = v4] * [0 < v5]*
    TREE(R, v1, 2, [0, 1])*
    TREE(I, v2, 2, [0])*
    TREE(N, v3, 2, [0])*
    AbsAll v0  $\in$  TreeRecords(v2).
    [nth(find(v2, v0), 2) inTree v1]*
    AbsAll v0  $\in$  TreeRecords(v3).
    [nth(find(v3, v0), 2) inTree v1]*
    [v6 inTree v1])))

```

```

(AbsExists v6 (AbsExists v5 (AbsExists v4
  (AbsExists v3 (AbsExists v2 (AbsExists v1
    [I = tmp_l]*
    (AbsUpdateWithLoc tmp_l v6
      (AbsUpdateWithLoc T (v6 + 1)
        ([0 < v5] * [tmp_l = 0  $\wedge$  tmp_r = 0]*
          (AbsUpdateWithLoc tmp_r (T + 1)
            (AbsUpdateWithLoc tmp_l T
              P + 1  $\mapsto$  T * P  $\mapsto$  N * [P = v4] * [0 < T]*
              TREE(R, v1, 2, [0, 1])*
              TREE(I, v2, 2, [0])*
              TREE(N, v3, 2, [0])*
              AbsAll v0  $\in$  TreeRecords(v2).
              [nth(find(v2, v0), 2) inTree v1]*
              AbsAll v0  $\in$  TreeRecords(v3).
              [nth(find(v3, v0), 2) inTree v1]*
              [T inTree v1]))) * -
            (AbsExists v1 (AbsExists v0
              v6 + 1  $\mapsto$  v2 * v6  $\mapsto$  v1))))))

```

On the right hand side, we need remove the magic wand. To do this, we first need to unfold the tree rooted at I. After unfolding, we simplify. The result of unfolding is shown below.

```

(AbsExists v8 (AbsExists v7
  (AbsExists v6 (AbsExists v5 (AbsExists v4
    (AbsExists v3 (AbsExists v2 (AbsExists v1
      (AbsUpdateWithLoc tmp_l v6
        (AbsUpdateWithLoc T (v6 + 1)
          (AbsUpdateWithLoc tmp_r (T + 1)
            (AbsUpdateWithLoc tmp_l T
              (P + 1  $\mapsto$  T * P  $\mapsto$  N * [P = v4] * [0 < T]*
              TREE(R, v1, 2, [0, 1])*
              v6 + 1  $\mapsto$  v8 * v6 + 0  $\mapsto$  nth(v7, 0)*
              TREE(nth(v7, 0), nth(v6 :: v7 :: v8 :: nil, 1),
                2, [0])*
              TREE(N, v3, 2, [0])*
              AbsAll v0  $\in$ 
                TreeRecords(v6 :: v7 :: v8 :: nil).
                [nth(find(v6 :: v7 :: v8 :: nil, v0), 2)
                  inTree v1]*
              AbsAll v0  $\in$  TreeRecords(v3).
              [nth(find(v3, v0), 2) inTree v1]*
              [T inTree v1]) * [N = v5])) * -
            [tmp_l = 0] * [tmp_r = 0] * [0 > v6] * -
            (AbsExists v1 (AbsExists v0
              v6 + 1  $\mapsto$  v2 * v6  $\mapsto$  v1))))))

```

We now use the rule shown in figure 8 to eliminate the magic wand construct.

```

(AbsExists v5 (AbsExists v4
  (AbsExists v3 (AbsExists v2 (AbsExists v1
    [I = nth(v5, 0)] * (P + 1)  $\mapsto$  v1 * P  $\mapsto$  N * [0 < v1]*
    TREE(R, v2, 2, [0, 1])*
    TREE(I, v5, 2, [0])*
    TREE(N, v3, 2, [0])*
    [TinTree v2]*
    AbsAll v0  $\in$  TreeRecords(v5).
    [nth(find(v5, v0), 2) inTree v2]*
    AbsAll v0  $\in$  TreeRecords(v3).
    [nth(find(v3, v0), 2) inTree v2]*
    [v1 inTree v2])))

```

We then merge the two states. The result is shown below. Most of the work can be done automatically by pairing off identical components in the two branches. There is one term, $[T = 0 \vee T \text{ inTree } v_1]$, that the user needs to introduce. It is derived from the fact that that $[T \text{ inTree } v_1]$ is in the left state and $[T = 0]$ is in the right state. They can be put together in a disjunction.

```

(AbsExists v6 (AbsExists v5 (AbsExists v4
  (AbsExists v3 (AbsExists v2 (AbsExists v1
    [I = I] * (P + 1)  $\mapsto$  v6 * P  $\mapsto$  N * [0 < v5]*
    TREE(R, v1, 2, [0, 1])*
    TREE(I, v2, 2, [0])*
    TREE(N, v3, 2, [0])*
    AbsAll v0  $\in$  TreeRecords(v5).
    [nth(find(v5, v0), 2) inTree v2]*
    AbsAll v0  $\in$  TreeRecords(v3).
    [nth(find(v3, v0), 2) inTree v2]*
    [v6 inTree v2]*
    [T = 0  $\vee$  T inTree v1])))

```

As a final step, we want to fold back the tree rooted at N with the cells at locations P and T + 1. The result is shown below.

```

(AbsExists v3 (AbsExists v2 (AbsExists v4
  AbsAll v0  $\in$  TreeRecords(v3).
  [nth(find(v3, v0), 2) inTree v1]*
  TREE(P, v3, 2, [0])*
  [nth(nth(v3, 1), 0) = N] * [0 < nth(v3), 2]*
  TREE(R, v1, 2, [0, 1])*
  TREE(I, v2, 2, [0])*
  AbsAll v0  $\in$  TreeRecords(v2).
  [nth(find(v2, v0), 2) inTree v1]*
  AbsAll v0  $\in$  TreeRecords(v3).
  [nth(find(v3, v0), 2) inTree v1]*
  [T = 0  $\vee$  T inTree v1]))

```

C. Entailment Proofs

There are three entailment proofs in our tree traversal example. The general methodology in PEDANTIC is to pair off identical components in the same manner as is done for merge and then to prove the remaining components separately.

IV. COQ SOURCE FILES

The main proof described in this paper with its 12 lemmas can be found in `TreeTraversal.v`. We also included the file `SatSolverMain.v` which shows the beginnings of our DPLL solver verification. We have also included `SatSolver.c` and a DIMACS file `test` which gives the original C implementation of our Sat solver (before we hand translated it to our Coq syntax). The files for both CoqPIE and PEDANTIC

are found at <https://github.com/kendroe/CoqPIE>. They compile with Coq 8.5pl3.

V. RELATED WORK

Separation logic was originally developed by John Reynolds [30]. Tools based on separation logic were developed and found to be effective in finding bugs in Microsoft device drivers [13], [15], [27].

The basic concepts were later introduced in the Coq theorem prover. McCreight developed a set of axioms and tactics to reason about separation logic [25]. Systems such as Charge! [4], Bedrock [1] and VST [2], [3] were then developed to provide frameworks for verifying imperative programs with recursive data structures. Coq's Ltac language for creating tactics is limited and hence limits the ability to do automation in any of these systems. PEDANTIC uses a deep model which gives greater flexibility in designing tactics.

In recent years, the VST system was used to verify the correctness of the HMAC algorithm in OpenSSL [5]. This is an approximately 200-line C program. This verification task was quite tedious and emphasizes the need to study proof development productivity issues. VST is built on top of CompCert-C [17] and hence the C language semantics in VST are fully verified.

VI. CONCLUSION

This paper has covered the basic concepts of the PEDANTIC framework, a tool for verifying the correctness of C programs. We have demonstrated how dynamic data structure invariants including cross referenced dependencies can be expressed and reasoned about. There are many other features for which there is not enough space in the paper to discuss. We currently have an implementation of the framework that contains all of the basic data types and tactics. The proof of the example program invariant of Figure 5 is complete, but correctness proofs of most auxiliary Lemmas still need to be completed and are for now admitted.

We are also working on the verification of the DPLL [11] based SAT solver invariant. After this work is done, the next project will be dealing with function calls. The idea is to use our invariant language to specify pre- and post- conditions which are an abstraction of the function's behavior and then to use those conditions in the verification of code that calls those functions. Separation logic provides some nice framing properties which we anticipate will be useful to generalize pre- and post- conditions after a function verification is done.

We are working towards greater automation of the proof development process. Much of the work done for the tree traversal proof presented can be automated by a tool such as CoqPIE [31]. The top level `loopInvariant` theorem can be generated automatically for the most part. The two areas that need user input are 1) providing loop invariants and 2) at merge points, some specification of the output state is required. Of the 12 lemmas, many turn out to be straight forward and can be automated. For example, `treeRef1`, `treeRef2`, `storeCheck1` and `storeCheck2` turn out to be fairly simple

proofs that can likely be automated (though this is not always the case). The areas of complexity are in doing implication proofs and merges.

The biggest challenge in making interactive theorem proving tools practical for software development is to better understand proof development productivity issues. Right now, based on the author's experience using Coq, every hour of software development likely requires 100 hours of proof development time. This ratio needs to come down for the tool to be practical. However, once these issues are addressed, proof development will become an important part of software development methodologies as it can find bugs for which there is otherwise no effective methodology.

REFERENCES

- [1] Greg Morrisett Avraham Shinnar Ryan Wisnesky Adam Chlipala, Gregory Malecha. Effective interactive proofs for higher-order imperative programs. 2009.
- [2] A. W. Appel. Tactics for separation logic, 2006. Early draft.
- [3] 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.
- [4] Jesper Bengtson, Jonas Braband Jensen, and Lars Birkedal. Charge! - a framework for higher-order separation logic in Coq. In *Third International Conference, ITP*, 2012.
- [5] Lennart Beringer, Adam Petcher, Q Ye Katherine, and Andrew W Appel. Verified correctness and security of openssl hmac. In *USENIX Security*, volume 15, pages 207–221, 2015.
- [6] Yves Bertot and Pierre Casteran. *Interactive Theorem Proving and Program Development*. Springer, 2004.
- [7] Hao Chen, Xiongnan (Newman) Wu, Zhong Shao, Joshua Lockerman, and Ronghui Gu. Toward compositional verification of interruptible os kernels and device drivers. In *PLDI*, pages 431–447, 2016.
- [8] Adam Chlipala. Mostly-automated verification of low-level programs in computational separation logic. In *32nd Programming Language Design and Implementation (PLDI)*, 2011.
- [9] David Costanzo, Zhong Shao, and Ronghui. End-to-end verification of information-flow security for c and assembly programs.
- [10] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.
- [11] D. Ranjan D. Tang, Y. Yu and S. Malik. Analysis of search based algorithms for satisfiability of quantified boolean formulas arising from circuit state space diameter problems. *Proceedings of the Seventh International Conference on Theory and Applications of Satisfiability Testing (SAT2004)*, May, 2004.
- [12] Jyotirmoy V. Deshmukh, E. Allen Emerson, and Prateek Gupta. Automatic verification of parameterized data structures. In *Proc. of Tools and Algorithms for the Construction and Analysis of Systems*, pages 27–41, 2006.
- [13] Peter W. O'Hearn Dino Distefano and Hongseok Yang. A local shape analysis based on separation logic. In *TACAS*, 2006.
- [14] 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.
- [15] 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.
- [16] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. seL4: Formal verification of an OS kernel. In *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles, SOSP '09*, pages 207–220, New York, NY, USA, 2009. ACM.
- [17] Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, July 2009.
- [18] Huisong Li, Francois Brenger, Bor-Yuh Evan Chang, and Xavier Rival. Semantic-directed clumping of disjunctive abstract states. page 13 pages, 2017.

- [19] Stephen Magill, Josh Berdine, Edmund Clarke, and Byron Cook.
- [20] Gregory Malecha and Greg Morrisett. Mechanized verification with sharing. In *ICTAC*, 2010.
- [21] Gregory Malecha, Greg Morrisett, Avraham Shinnar, and Ryan Wisnesky. Toward a verified relational database management system. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '10, pages 237–248, New York, NY, USA, 2010. ACM.
- [22] Roman Manevich, Boris Dogadov2, and Noam Rinetzkyy. From shape analysis to termination analysis in linear time. In *CAV*, 2016.
- [23] Nicolas Marti and Reynald Affeldt. A certified verifier for a fragment of separation logic. *Information and Media Technologies*, 4(2):304–316, 2009.
- [24] 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.
- [25] 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.
- [26] ThanhVu Nguyen. *Automating Program Verification and Repair Using Invariant Analysis and Test-input Generation*. PhD thesis, University of New Mexico, August 2014.
- [27] 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.
- [28] Benjamin C. Pierce, Chris Casinghino, Michael Greenberg, Vilhelm Sjberg, and Brent Yorgey. *Software Foundations*. 2011.
- [29] Piskac R., Wies T., and Zufferey D. Automating separation logic using smt. In Sharygina N. and Veith H., editors, *Computer Aided Verification. CAV. Lecture Notes in Computer Science*, volume 8044. Springer, Berlin, Heidelberg, 2013.
- [30] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, pages 55–74, 2002.
- [31] Kenneth Roe and Scott Smith. Coqpie: An ide aimed at improving proof development productivity (rough diamond). In *ITP*, 2016.
- [32] Tahina Ramananandro Zhong Shao Xiongnan (Newman) Wu Shu-Chun Weng Haozhong Zhang Yu Guo Ronghui Gu, Jrmie Koenig. Deep specifications and certified abstraction layers. In *POPL*, pages 595–608, 2015.
- [33] Reinhard Wilhel Shmuel Sagiv, Thomas W. Reps. Solving shape-analysis problems in languages with destructive updating. 20(1):1–50, 1998.
- [34] Yavuz-Kahveci T. and Bultan T. Automated verification of concurrent linked lists with counters. In Hermenegildo M.V. and Puebla G., editors, *Static Analysis. SAS. Lecture Notes in Computer Science*, volume 2477, 2002.
- [35] Ting Zhang Zohar Manna, Henny B. Sipma. Verifying balanced trees. In *Proceedings of the Symposium on Logical Foundations of Computer Science (LFCS 2007)*, 2007.