

Call Invariants

Shuvendu K. Lahiri and Shaz Qadeer

Microsoft Research

Abstract. Program verifiers based on first-order theorem provers model the program heap as a collection of mutable maps. In such verifiers, preserving unmodified facts about the heap across procedure calls is difficult because of scoping and modification of possibly unbounded set of heap locations. Existing approaches to deal with this problem are either too imprecise, require introducing untrusted assumptions in the verifier, or resort to unpredictable reasoning using quantifiers. In this work, we propose a new approach to solve this problem. The centerpiece of our approach is the *call invariant*, a new annotation for procedure calls. A *call invariant* allows the user to specify at a call site an assertion that is inductively preserved across an arbitrary update to a heap location modified in the call. Our approach allows us to leverage existing techniques for reasoning about call-free programs to precisely and predictably reason about programs with procedure calls. We have implemented the approach and applied it to the verification of examples containing dynamic memory allocations, linked lists, and arrays. We observe that most call invariants have a fairly simple shape and discuss ways to reduce the annotation overhead.

1 Introduction

Floyd-Hoare logic is a framework for decomposing the partial correctness checking of a program into smaller proof obligations, where a Floyd-Hoare triple $\{P\} s \{Q\}$ is associated with each statement s in the program [10]. Verification condition (VC) generation based on Dijkstra's *weakest liberal precondition* (*wlp*) predicate transformer allows precise reasoning about Floyd-Hoare triples without requiring intermediate assertions for loop-free and call-free statements [7]. The use of automated theorem provers (including *satisfiability modulo theories* (SMT) solvers [20]) for checking the verification conditions provide a scalable and precise approach to program verification, and forms the basis of several tools (e.g. ESC/Java [8], Spec# [5], HAVOC [12]).

However, this framework is not as effective in the presence of the heap and procedure calls. The main issue is to preserve unmodified facts about the part of the heap in the caller's scope that is not in scope of the callee. More formally, a procedure specification comprises of (a) preconditions, (b) postconditions and (c) set of variables modified by the procedure. Since the heap is modeled as a collection of maps, a procedure that modifies a location in a map has to specify that the entire map is potentially modified. The only facts related to such a modified map, known after a procedure call has to come from the postconditions

of the callee procedure. However, the postconditions can only refer to part of the map *in scope*, i.e. the locations reachable from globals and parameters. This means that unmodified facts at caller's scope about a modified map may not be preserved across a procedure call. Matters are further complicated as a procedure call might update an unbounded number of locations in a map. Efficiently decidable SMT-based logics (e.g. linked lists [12], arrays [6]) that deal with a bounded number of heap updates (for a loop-free, call-free program fragment) are rendered ineffective in the presence of the unbounded number of updates.

An existing approach to address this problem has been to introduce *frame axioms* to allow preserving certain unmodified facts [14]. These axioms are not verified in the same spirit as the rest of the user annotations, and may introduce unsoundness in the verifier. Moreover, these frame axioms are encoded using complex quantified facts in the verification condition that severely compromises the predictability of the underlying theorem prover. Verification of such quantified formulas require expert users to be able to guide the theorem provers. These shortcomings make *wp* based Floyd-Hoare reasoning less appealing for reasoning about programs with scoping and the heap.

In this work, we present an alternative approach based on the following insight:

For any statement s in the program, if an assertion R is preserved by an abstract (re-)execution of s in which the heap locations modified by s are updated nondeterministically, then R is preserved across the statement s .

We use an instance of this general rule for procedure calls to deal with the imprecision due to scoping. We also provide a new annotation called *call invariant* for the user to specify an inductive hypothesis when the abstract execution could be unbounded. Given a program with call invariant annotations, we perform a source-to-source transformation to create another program that can be reasoned with any existing technique for call-free programs. In particular, this allows a user to leverage existing *wp*-based verifiers to analyze programs with procedure calls with unbounded heap updates.

One can also view our approach as a strategy to augment the underlying first-order theorem provers with an induction scheme to verify formulas containing unbounded number of heap updates. The call invariant (provided by the user) plays the role of an inductive hypothesis and the underlying theorem prover is used to discharge the proof obligations for establishing the inductive hypothesis. However, there are several advantages of formalizing the inductive hypothesis at the program level instead of at the level of a formula:

1. The call invariants are specified as program annotations independent of the underlying prover. Therefore, the user does not need to interact with the specific syntax of the underlying theorem provers.
2. We formulate the call invariants as loop invariants. This opens the possibility to leverage existing loop invariant synthesis techniques to infer call invariants in many cases.

We have augmented the `Boogie` [3] verifier with call invariant annotations, and have applied it to verify a set of examples containing dynamic memory allocation, linked lists and arrays. These examples were already annotated with preconditions and postconditions — we discuss the additional annotation burden due to call invariants. We introduce useful syntactic sugars and observe the common shape of most call invariants and additional specification required to prove these examples. We also discuss tradeoffs in reducing the additional burden at the cost of slight complication of the assertion logic, without sacrificing soundness.

2 Motivation

Consider two versions of a program in Figure 1 written in a variant of the `Boogie` language [3]. The example is an abstraction of a real-life device driver `kbdcClass` [21] that uses multiple lists of device extensions. The first version (on the left) has single procedure with no procedure calls, and the second version (on the right) has a procedure call.

2.1 Program without Procedure Calls

Let us first look at the example in Figure 1(a). Initially ignore the lines starting with **pre**, **post**, **inv** and **modifies**, which denote annotations. The example contains two map (or array) variables `N` and `D` to model two fields in an object. The procedure *Proc1* takes two pointers `p` and `q` to denote the heads of the two disjoint acyclic lists $\{p, N[p], N[N[p]], \dots, \text{nil}\}$ and $\{q, N[q], N[N[q]], \dots, \text{nil}\}$ respectively. The procedure first initializes the `D` field of all the pointers in the linked list starting at `p` in a while loop, and then non-deterministically deletes some entries from the list starting at `q` — this mimics removing elements from a list that satisfy some criteria. We would like to prove that the `D` field has been correctly initialized for the list from `p`.

The assertions in **pre** and **post** denote preconditions and postconditions of a procedure. The precondition states that the two lists are disjoint and acyclic: we use the set constructor $\text{Btwn}(m, u, v)$, where m is a map value of type $\text{int} \rightarrow \text{int}$ and u and v are values of type int , to denote the set of values $\{u, m[u], m[m[u]], \dots, v\}$ when v lies in the set, or $\{\}$ otherwise [17,12]. The postcondition states that the value of `D` map at all the elements of list from `p` is 1. The “modifies” clause in **modifies** says that the maps `N` and `D` are modified by the procedure, possibly at all locations. Loop invariant assertions are provided using **inv** annotations. The expression $\text{old}(x)$ denotes the value of a variable x at the entry to a procedure (when used in a postcondition), or at the entry of a loop (when used in a loop invariant). The loop invariants on the first loop states that the variable `iter` points to the list from `p`, and all the entries upto `iter` have been initialized to 1. The first loop establishes the postcondition of the procedure on exit from the loop — the problem is to preserve it across the second loop. The first loop invariant for the second loop states that the set of pointers in the list from `p` remains unchanged. The second loop invariant says that the iterator variable `iter` points to elements in the list from `q`.

<pre> var N : int → int; var D : int → int; pre Btwn(N, p, nil) ∩ Btwn(N, q, nil) = {nil} post ∀u ∈ Btwn(N, p, nil). u = nil ∨ D[u] = 1 modifies D, N proc Proc1(p : int, q : int) : void = var iter : int; iter := p; inv iter ∈ Btwn(N, p, nil) inv ∀u : int ∈ Btwn(N, p, nil). u ∈ Btwn(N, iter, nil) ∨ D[u] = 1 while (iter) do D[iter] := 1; iter := N[iter]; iter := q; inv Btwn(N, p, nil) = Btwn(oid(N), p, nil) inv iter ∈ Btwn(oid(N), q, nil) while (iter ∧ N[iter]) do if (*) N[iter] := N[N[iter]]; iter := N[iter]; </pre>	<pre> var N : int → int; var D : int → int; pre Btwn(N, p, nil) ∩ Btwn(N, q, nil) = {nil} post ∀u ∈ Btwn(N, p, nil). u = nil ∨ D[u] = 1 modifies D, N proc Proc1(p : int, q : int) : void = var iter : int; iter := p; //Loop invariants omitted while (iter) do D[iter] := 1; iter := N[iter]; cinv cframe(Btwn(N, p, nil)) call Proc2(q); updates N @ Btwn(N, t, nil) proc Proc2(t : int) : void = var iter : int; iter := t; inv iter ∈ Btwn(oid(N), t, nil) while (iter ∧ N[iter]) do if (*) N[iter] := N[N[iter]]; iter := N[iter]; </pre>
--	--

Fig. 1. Example with (a) no procedure calls and (b) a procedure call

These annotations are sufficient to prove the postcondition, since the map D does not change in the second loop. The annotated program can be encoded precisely in the assertion logic if the logic is closed under weakest (liberal) precondition [7] of statements in the programming language. Such logics with decision procedures have been proposed in [12], thereby providing an algorithm for checking such annotated programs.

2.2 Program with Procedure Calls

Now let us look the second version in Figure 1(b), where the second loop has been moved to a procedure $Proc2$. Let us initially ignore the annotation in **cinv**, and the **updates N @** annotation on $Proc2$. Instead, let us pretend that we only have a **modifies N** annotation for $Proc2$. Since $Proc2$ modifies the map N , any fact involving N will be invalidated after the call to $Proc2$. Therefore, the postcondition of $Proc1$ will not be provable. It is not hard to see that we cannot write any specification about the heap in the scope of $Proc2$ (namely the pointers in the list reachable from t) that would allow us to prove the postcondition for $Proc1$.

One approach to address the imprecision has been to use *frame axioms* that allow the user to specify how to preserve certain unmodified facts [14]. However, the use of frame axioms can lead to unsoundness as they are not verified. Besides,

these frame axioms have complex quantified structure that may destroy the predictability of the underlying theorem provers. In the rest of this section, we show how our approach helps retain precision in the presence of procedure calls, without requiring the use of frame axioms.

First, let us look at the new annotation on *Proc2*. The annotation `updates X @ S` denotes that the map X could have been modified only at locations in S by the procedure., where the set-expression S is interpreted at entry to a procedure. This annotation is actually a syntactic sugar for a particular postcondition that we explain in Section 3.2, and does not introduce any new annotation construct. In this example, the annotation is used to specify that the map N is only modified in the locations present in the list from t at the start of the procedure.

Second, we introduce a new annotation construct called *call invariants* (using `cinv`) that allows the user to annotate a call site of a procedure. A user can specify an assertion R inside `cinv` at a call site of a procedure — with the intention that R is preserved across modifications to the maps in the callee. We use the syntactic sugar `cframe(e)` to denote the assertion that the value of the expression e is preserved across the call. In this example, the assertion in `cinv` states that the set of pointers in the list from p is preserved across the call to *Proc2*.

These annotations suffice to prove the postcondition of *Proc1*. Indeed, we can prove the specifications of this example (including the new proof obligations for showing that assertions in `cinv` are really preserved across a call). Not only that, the proof obligations can be encoded using the same logic that was used to prove the example without a procedure call in Figure 1(a).

3 Call Invariants

3.1 Source and Assertion Language

Figure 2 shows a simple programming language. The language supports scalar and map variables (*Scalars* and *Maps* respectively) and various operations on them. Let $Vars = Scalars \cup Maps$. The type of any variable $x \in Scalars$ is integer (`int`), and the type of any variable $X \in Maps$ is a map from integers to integers (`int \rightarrow int`). The standard assignment statement for scalars is extended with assignment statements for maps. The statement for variable introduction `var x in s endvar` introduces a variable x with an arbitrary value in s (the variable introduction rule for a map variable is similar). The statement `assert ϕ` behaves as a `skip` when the formula ϕ evaluates to `true` in the current state; else the execution of the program *fails*. The statement `assume ϕ` behaves as a `skip` when the formula ϕ evaluates to `true` in the current state; else the execution of the program is *blocked*. Expression terms of the statements and formulas are denoted by *Expr*, and include scalar variables, constants, arithmetic expressions and map lookups. The language also supports sequential composition, procedure calls, conditional statements and while loops. Allocation and deallocation can be modeled by introducing a special map `Alloc` to track the allocation status of objects, but is not built into the language. We show an example of dynamic allocation in the next section 4.1.

x, y	$\in Scalars$
X, Y	$\in Maps$
e	$\in Expr ::= x \mid c \mid e \pm e \mid X[e]$
s, t	$\in Stmt ::= skip \mid \mathbf{assert} \phi \mid \mathbf{assume} \phi \mid x := e \mid X := Y \mid X[e] := e \mid$ $\quad \mathbf{var} \ x \ \mathbf{in} \ s \ \mathbf{endvar} \mid \mathbf{var} \ X \ \mathbf{in} \ s \ \mathbf{endvar} \mid s; s \mid x := \mathbf{call} \ f(e) \mid$ $\quad \mathbf{if} \ (e) \ \mathbf{then} \ s \ \mathbf{else} \ t \mid \mathbf{while} \ (e) \ \mathbf{do} \ s$
P, Q, R, ϕ, ψ	$\in Formula ::= e \leq e \mid \phi \wedge \phi \mid \neg \phi \mid e \in \mathcal{S} \mid \mathcal{S} \subseteq \mathcal{S} \mid \dots$
\mathcal{S}	$\in SetExpr ::= \mathbf{Btwn}(X, e, e) \mid \mathbf{Inverse}(X, e) \mid [e, e] \mid \dots \mid \mathcal{S} \cup \mathcal{S} \mid \mathcal{S} \setminus \mathcal{S} \mid \dots$

Fig. 2. A simple programming language and assertion logic

The formulas in *Formula* constitute the *assertion logic* for specifying contracts for programs in this language. The language of formulas in *Formula* is extensible, and includes relational, Boolean operations and set operations. *SetExpr* represent set-valued expressions and can be constructed from various set constructors such as *Btwn* and $[e, e]$ and other operations on sets. The weakest (liberal) precondition of an assertion $\phi \in Formula$ with respect to a statement $s \in Stmt$ is denoted as $wp(s, \phi)$. Intuitively, $wp(s, \phi)$ is a formula that represents the set of states for which executing s does not fail any assertions in s and moreover if the execution terminates, it does so in a state satisfying ϕ . The weakest precondition for the simple statements in our language are described in the appendix of the detailed report [13], and are fairly standard [4]. The assertion logic in *Formula* is *closed under wp*, when for any $\phi \in Formula$ and $s \in Stmt$, $wp(s, \phi) \in Formula$. The following proposition relates checking partial correctness of statements using Floyd-Hoare triples [10] and provability in a logic. Let us refer to loop-free and call-free statements as *simple* statements.

Proposition 1. *If the assertion logic in Formula is closed under wp for simple statements in Stmt, then for any simple statement s, (i) the logical formula $(P \implies wp(s, Q))$ is in Formula, and (ii) is valid if and only if the Floyd-Hoare triple $\{P\} s \{Q\}$ holds.*

In such a case, an automated theorem prover for checking assertions in *Formula* provides a method (an algorithm when the theorem prover is complete and terminating) to check Floyd-Hoare triples expressed in the logic. In the presence of loops that may have unbounded updates, the user decomposes the problem by specifying a loop invariant.

However, the presence of procedure calls and heap makes reasoning in Floyd-Hoare logic imprecise because it introduces the challenge of preserving unmodified facts about the heap in the callers scope (and not in the callee's scope) across (a possibly unbounded) update to the heap in the callee.

3.2 Call Invariant and Program Instrumentation

In this section, we instrument the source program with additional (ghost) variables to track the set of locations modified by a procedure and then introduce a

new annotation called *call invariants* at the call site of a procedure. The purpose of the call invariants at a call site is to specify the assertions that are preserved across the procedure call.

First, for each global map $X \in Maps$, we introduce a state variable MS_X , whose interpretation is a set of locations. Intuitively, the value of MS_X at exit from a procedure captures the set of locations where X was modified between the entry and exit to the procedure. The source program is automatically instrumented to update the MS_X variables as follows:

- For any explicit update to the map X , $X[e_1] := e_2$; we insert $MS_X := MS_X \cup \{e_1\}$; before the update to X .
- Every procedure has a precondition **pre** $MS_X = \{\}$.
- Before any procedure call that has a **modifies** X annotation, we save the current value of MS_X into a caller local variable, set MS_X to $\{\}$. Upon return from the procedure, we union the saved set with the value of the set after the procedure call.

It is not hard to see that the value of MS_X at exit from a procedure captures the set of locations where X was modified between the entry and exit to the procedure. For this instrumented program, the user can specify preconditions, postconditions and loop invariants in terms of MS_X variable just as any other state variable.

Next, we introduce an annotation construct called *call invariant* at the call site of a procedure specified using **cin**v R , where $R \in Formula$. For a call site that may (transitively) modify locations in the map X , we provide the following instrumentation in addition to the updates for MS_X :

```

var  $X_{pre}, X_{post}$  in
   $X_{pre} := X$ ;
  call  $Foo(e)$ ; //procedure call
   $X_{post} := X$ ;
   $X := X_{pre}$ ;

  inv  $R[X_{pre}/old(X)]$ 
  while (*) do
    var  $u, v$  in assume  $u \in MS_X; X[u] := v$ ; endvar

  assume  $X = X_{post}$ ;
endvar

```

First, it copies the value of X before and after the call into local variables X_{pre} and X_{post} respectively. It restores X to the value before the call, and introduces a non-deterministic loop that updates X at one of the locations in MS_X . Finally, the assume relates the value of X after the loop with the value at the end of the procedure call X_{post} . The purpose of the loop is to model the abstract re-execution of the procedure call that nondeterministically modifies the heap locations modified by the callee, starting from the state before the call. The call

invariant R (specified using `cinv` R) is checked as a loop invariant for this loop; any occurrence of `old(X)` is replaced with copy X_{pre} , the value of X just before the procedure call. Although we have described the instrumentation for a single map variable, our implementation allows for multiple maps and is presented in the detailed report [13].

Syntactic sugars. We introduce two syntactic sugars to make the specifications concise and readable:

1. We introduce updates $X @ \mathcal{S}$ for a set-valued expression \mathcal{S} , as a syntactic sugar for the following annotations:

$$\begin{array}{l} \text{modifies } X, MS_X; \\ \text{post } MS_X \subseteq \text{old}(\mathcal{S}) \end{array}$$

2. The use of `old(X)` in a call invariant R refers to the value of X prior to the procedure call (same as X_{pre} at the time of the call). We provide a sugar `cframe(.)` to denote that an expression is preserved by the procedure. For an expression e , `cframe(e)` expands to $e = \text{old}(e)$. This is most useful when specifying that the value of a scalar expression or a set-valued expression is preserved.

Bounded updates. Recall that the scoping problem exists even when a callee modifies a bounded number of heap locations. Our program instrumentation is still essential to preserve facts at the callers scope. However, if callee has a postcondition that bounds MS_X to a bounded number of locations (say n), then one does not require the user to specify a call invariant R . In such cases, it suffices to unroll the loop that modifies locations in MS_X n times, and eliminate the need for the call invariant.

4 Evaluation

We have built a prototype implementation of call invariant annotations over the `Boogie` program verifier [3]. In addition to specifying procedure preconditions, postconditions and loop invariants, the user can specify call invariants at call sites using the `cinv` annotation. We highlight the annotations required related to call invariants in the program in addition the already provided preconditions and postconditions. We discuss more about the `cinv` call invariants (the call invariants that are both highlighted and underlined) later in Section 4.4.

We perform the program instrumentation to create the transformed program where the call invariants are desugared as loop invariants as described earlier. The resultant annotated program is verified by `Boogie` by generating a logical formula (verification condition) and checking the formula with *satisfiability modulo theories* (SMT) solvers. The assertion logics used in these programs use sophisticated set constructors in addition to the usual theories of uninterpreted functions, select-update arrays and arithmetic supported by most SMT solvers. In spite of the complexity, the assertion logics used in these examples are closed


```

var Alloc : int → bool;
var D : int → int;

updates D @ Inverse(Alloc, false)
proc Proc4() : void =
  var y : int;
  while (*) do
    y := new; D[y] := 0;
    /* Add y to a list */

proc Proc3() : void =
  var x : int;
  x := new;
  D[x] := 5;

  cinv cframe(D[x])
  call Proc4();

  assert D[x] = 5;

```

Fig. 3. Example with dynamic memory allocation and the `Inverse` set constructor

under the *wp* predicate transformer with respect to call-free and loop-free statements in the language. For a few cases, we even have decision procedures (i.e. a sound, complete and terminating procedure) for deciding formulas in the assertion logic [6,12].

4.1 Dynamic Allocation

In the example in Figure 3, we consider a map `Alloc` whose range contains two Boolean values `true` and `false`. The map is used to track the set of *allocated* elements of the domain; `Alloc[u] = true` if and only if *u* is an allocated element. The statement `x := new` is a sugar for the following statements: `{var u in x := u endvar; assume Alloc[x] = false; Alloc[x] := true;}`. In this example, the map `D` is mutated at an unbounded set of freshly allocated locations in procedure `Proc4` — namely at the locations *u* for which `Alloc[u] = false` at entry to `Proc4`. In this case, this modified set excludes `x` in `Proc3`.

To specify the modified set, we use the set constructor `Inverse : (int → int) * int → 2int`, which takes a map and returns all elements of the domain that map to a given value; i.e. `Inverse(X, v) ≐ {u | X[u] = v}`. For this set constructor, `wp(X[x] := y, u ∈ Inverse(X, v))` is given by [12]:

$$(y = v \wedge u \in \text{Inverse}(X, v) \cup \{x\}) \vee (y \neq v \wedge u \in \text{Inverse}(X, v) \setminus \{x\})$$

4.2 Linked Lists

Reverse. Consider the recursive implementation of list reversal in Figure 4(a). This implementation performs an in-place reversal of an input list. The precondition requires the argument `h` to point to a nonempty acyclic list. The procedure may modify the `N` map only at the pointers in this list. The first postcondition asserts that the set of elements in the output list is the same as the set of elements in the input list. The second postcondition strengthens this assertion to ensure that the ordering in the output list is the reverse of the ordering in the input list.

The recursive call to *reverse* requires a call invariant stating that the value of `N[h]` remains unchanged by the call. This assertion is crucial for ensuring that subsequent updates to `N` first, do not trash the list reversal performed by the recursive call itself and second, successfully reverse the link from `h` to `N[h]`.

<pre> var N : int → int; pre h ≠ nil pre nil ∈ Btwn(N, h, nil) updates N @ Btwn(N, h, nil) post Btwn(old(N), h, nil) = Btwn(N, r, nil) post ∀u ∈ Btwn(old(N), h, nil). u = nil ∨ Btwn(old(N), h, u) = Btwn(N, u, h) proc reverse(h : int) : (r : int) = if (N[h] = nil) { r := h; } else { <u>cinv cframe(N[h])</u> r := reverse(N[h]); N[N[h]] := h; N[h] := nil; } </pre>	<pre> var N : int → int; var D : int → int; pre r ∈ Btwn(N, l, r) updates N @ ROS(N, l, r); post Sorted(N, D, hd, r) post r ∈ Btwn(N, hd, r) post ROS(old(N), l, r) = ROS(N, hd, r) proc quick_sort(l : int, r : int) returns (hd : int) = var ret : int; if (l = r ∨ N[l] = r) { hd := l; } else { hd := partition(l, r); <u>cinv cframe(ROS(N, hd, N[l]))</u>; <u>cinv cframe(N[l])</u>; ret := quick_sort(N[l], r); N[l] := ret; <u>cinv cframe(ROS(N, N[l], r))</u>; <u>cinv Sorted(N, D, N[l], r)</u>; ret := quick_sort(hd, N[l]); hd := ret; } } </pre>
--	--

Fig. 4. List examples (a) reverse, (b)list-based quick sort

List sort. We have also verified an implementation of quick sort for lists. This example (present in Figure 4(b)), required nontrivial call invariants. We have used the following helper predicates to define the annotations for this example:

$$\begin{aligned}
 ROS(X, u, v) &\doteq \text{Btwn}(X, u, v) \setminus \{v\} \\
 UpperBound(X, Y, l, r, d) &\doteq \forall u \in ROS(X, l, r) : Y[u] \leq d \\
 LowerBound(X, Y, l, r, d) &\doteq \forall u \in ROS(X, l, r) : Y[u] \geq d \\
 Sorted(X, Y, l, r) &\doteq \forall u \in ROS(X, l, r) : \forall v \in ROS(X, u, r) : Y[u] \leq Y[v]
 \end{aligned}$$

The example requires reasoning about shapes of lists, properties of a collection of pointers, and arithmetic relationship on the data elements of the list. The recursive nature of the procedure makes the proof highly non-trivial — which explains the complexity of the call invariants. The proof illustrates the benefits of combination frameworks present in first-order theorem provers for precise reasoning of such examples.

Merge and append. In addition to these programs, we have also successfully verified recursive implementations for appending and merging two lists. These implementations and their specifications are described in the appendix of the detailed report [13]. Interestingly, in spite of the presence of recursive calls

```

var A : int → int;

post idx ∈ [l, r)
post Deref(A, [l, r)) = Deref(olds(A), [l, r))
updates A @ [l, r)
proc Partition(l : int, r : int, pivot : int) : (idx : int) =
  /* Partitions A and returns the final index of pivot */

pre 0 < l ≤ r
post Deref(A, [l, r)) = Deref(olds(A), [l, r))
updates A @ [l, r)
proc QuickSort(l : int, r : int) : void =
  var pivot : int, idx : int;
  if (l = r) return;
  pivot := A[l];
  idx := Partition(l, r, pivot);
  cinv cframe(Deref(A, [idx, r)));
  QuickSort(l, idx);
  cinv cframe(Deref(A, [l, idx + 1)));
  QuickSort(idx + 1, r);

```

Fig. 5. Example of quicksort over an array A

and unbounded number of updates, no call invariants were required for proving the correctness of these examples.

4.3 Arrays

In this example, we illustrate the use of two new set constructors, (i) the range set constructor $[i, j)$, and (ii) a set constructor **Deref** to collect the content of a set of locations. Consider the quicksort algorithm in Figure 5 where the array map **A** is being sorted with recursive invocations to *QuickSort*. The procedure *QuickSort* sorts the indices of the array **A** in the range $[l, r)$. The procedure *Partition* (we omit the procedure body) takes a value **pivot** and returns an index **idx** such that $\text{idx} \in [l, r)$.

Let us check a simple property that the algorithm preserves the contents of **A**, assuming distinct elements in the array. Since the postcondition of *Partition* establishes this constraint, the main challenge is in establishing this fact across the two recursive calls to *QuickSort*. The two call invariants serve to preserve facts at the call-site required to establish the postcondition. For example, the first call invariant states that the contents of **A** in the range $[\text{idx}, r)$ is preserved by the call to *QuickSort*(*l*, *idx*) since the procedure does not modify these locations.

The specifications for this example refer to a dependent set constructor **Deref** : $(\text{int} \rightarrow \text{int}) * 2^{\text{int}} \rightarrow 2^{\text{int}}$ that takes a map and a set and constructs a set with the union of values of the map at elements in the set. Formally, $\text{Deref}(X, S) \doteq \{X[u] \mid u \in S\}$. Interestingly, this assertion logic is still closed under *wp* with respect to statements of our language. For this set constructor, $\text{wp}(X[x] := y, u \in \text{Deref}(X, S))$ is defined as follows:

$$\begin{aligned}
& \bigvee x \notin S \wedge u \in \text{Deref}(X, S) \\
& \bigvee x \in S \wedge X[x] \in \text{Deref}(X, S \setminus \{x\}) \wedge u \in \text{Deref}(X, S) \cup \{y\} \\
& \bigvee x \in S \wedge X[x] \notin \text{Deref}(X, S \setminus \{x\}) \wedge u \in (\text{Deref}(X, S) \setminus \{X[x]\}) \cup \{y\}
\end{aligned}$$

The first disjunct corresponds to the case when $\text{Deref}(X, S)$ remains unchanged; the second (and the third) disjunct corresponds to the cases when the value $X[x]$ is contained (respectively not contained) in X at an index other than x .

4.4 Discussion

We now discuss the call invariant annotations highlighted as **cinv**. These call invariants are interesting because their specification can be obviated by adding the following postcondition automatically for each procedure:

$$\text{post } \forall u : \text{int} :: u \in \text{MS}_X \vee X[u] = \text{old}(X)[u]$$

This postcondition ensures that the map X is preserved at a location disjoint from MS_X . However, the postcondition introduces a quantifier, which may compromise the termination of a theorem prover. For example, the decision procedure for reasoning about linked lists with sets [12] may not terminate when reasoning with quantified facts of the above form. The invariants marked with **cinv** are the ones that do not have to be specified, when the theorem prover can prove the program with this quantified postcondition. The reader may observe that these are precisely those call invariants that do not contain a set constructor or a predicate that depends on the map being modified by the callee. This explains why most call invariants using **Btwn** are not removed.

The example of the list implementation of quick sort uses a call invariant which is not specified using the **cframe**(.) syntactic sugar. We needed to specify a single-state predicate to specify the sortedness of the list, instead of a set or a scalar expression inside **cframe**(.) — one may view this as a way to preserve a relation (in this case sortedness). Finally, several examples did not need a call invariant (e.g. list append and merge) even in the presence of unbounded updates in callees. This is due to the fact that these recursive procedures are actually tail-recursive, where one does not need to carry facts before the procedure call.

5 Related Work

In this work, we provide a simple approach for leveraging precise verifiers for call-free programs to reason about programs with procedure calls. An important benefit of our approach is that it does not require adding additional unchecked or complex frame axioms [14] to the verification conditions. However, the benefit comes at the cost of additional user-specified annotations. Our contribution is complementary to the large body of work on modular verification in the presence of data hiding [5,11,2]. The implementations of these approaches invariably use complex quantifiers to encode variants of frame axioms that make verification unpredictable [1]; our work could potentially be used to eliminate or simplify these frame axioms.

Separation logic [18] is a specialization of Floyd-Hoare logic that requires a specialized assertion logic. The assertion logic contains formulas to describe heaps, where the formula $\phi * \psi$ denotes a heap with two disjoint subheaps for which ϕ and ψ hold respectively. The following *frame rule* [16] allows preservation of the fact R whose domain does not intersect with the modified locations in the statement s .

$$\frac{[\text{FRAME RULE}] \quad \{P\} s \{Q\}}{\{P * R\} s \{Q * R\}}$$

However, the specialized extension prevents leveraging existing tools for doing precise verification of call-free programs. Many of the inference rules for Floyd-Hoare logic do not apply in a straightforward manner; for example, the rule of constancy is no longer sound [18] in this extension. In addition, the specifications of s has to precisely describe the locations in the heap that s reads or write from, which is not required in Floyd-Hoare logic. Besides, one cannot use the standard *wp*-based methods to generate verification conditions. Although scalable and automatic shape analysis engines have been recently developed based on separation logic [22], proof obligations for expressive separation logic properties are often checked with higher order theorem provers [15].

The frame problem due to procedure scoping has also been explored in the context of automatic shape analysis [19,9], where the caller's heap is separated from the callee's heap by identifying a set of *cutpoints* which are dominators for any location in the heap of the callee. These cutpoints are treated as ghost parameters, but precision is lost when the set of cutpoints can be unbounded. We believe that call invariants may alleviate the need to introduce cutpoints to pass a local heap to a callee.

6 Conclusion

This paper makes two important contributions. First, we provide an automatic program instrumentation to address the imprecision due to procedure scoping. The call invariant annotation allows a user to specify an inductive hypothesis while dealing with unbounded updates in a callee. Such an instrumented and annotated program can be verified using any off-the-shelf verifier without any need to interact with the lower-level theorem prover to specify the inductive hypothesis. This has allowed us to leverage existing precise verifiers for call-free programs to verify non-trivial examples with a small annotation burden. Second, we have separated the problem of specifying the frame from inferring the frame automatically for procedure calls. We imagine exploiting various loop invariant inference algorithms to synthesize most call invariants, given their restricted shape in practice. However, it still allows the user to explicitly specify the frame when the inference algorithm fails to discover the necessary frame.

References

1. Banerjee, A., Barnett, M., Naumann, D.A.: Boogie meets regions: A verification experience report. In: Shankar, N., Woodcock, J. (eds.) VSTTE 2008. LNCS, vol. 5295, pp. 177–191. Springer, Heidelberg (2008)
2. Banerjee, A., Naumann, D.A., Rosenberg, S.: Regional logic for local reasoning about global invariants. In: Ryan, M. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 387–411. Springer, Heidelberg (2008)
3. Barnett, M., Chang, B.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A modular reusable verifier for object-oriented programs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) FMCO 2005. LNCS, vol. 4111, pp. 364–387. Springer, Heidelberg (2006)
4. Barnett, M., Leino, K.R.M.: Weakest-precondition of unstructured programs. In: Program Analysis For Software Tools and Engineering (PASTE 2005), pp. 82–87 (2005)
5. Barnett, M., Leino, K.R.M., Schulte, W.: The Spec# programming system: An overview. In: Barthe, G., Burdy, L., Huisman, M., Lanet, J.-L., Muntean, T. (eds.) CASSIS 2004. LNCS, vol. 3362, pp. 49–69. Springer, Heidelberg (2005)
6. Bradley, A.R., Manna, Z., Sipma, H.B.: What's decidable about arrays? In: Emerson, E.A., Namjoshi, K.S. (eds.) VMCAI 2006. LNCS, vol. 3855, pp. 427–442. Springer, Heidelberg (2006)
7. Dijkstra, E.W.: Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM* 18, 453–457 (1975)
8. Flanagan, C., Leino, K.R.M., Lillibridge, M., Nelson, G., Saxe, J.B., Stata, R.: Extended static checking for Java. In: Programming Language Design and Implementation (PLDI 2002), pp. 234–245 (2002)
9. Gotsman, A., Berdine, J., Cook, B.: Interprocedural shape analysis with separated heap abstractions. In: Yi, K. (ed.) SAS 2006. LNCS, vol. 4134, pp. 240–260. Springer, Heidelberg (2006)
10. Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* 12(10), 576–580 (1969)
11. Kassios, I.T.: Dynamic frames: Support for framing, dependencies and sharing without restrictions. In: Misra, J., Nipkow, T., Karakostas, G. (eds.) FM 2006. LNCS, vol. 4085, pp. 268–283. Springer, Heidelberg (2006)
12. Lahiri, S.K., Qadeer, S.: Back to the Future: Revisiting Precise Program Verification using SMT Solvers. In: Principles of Programming Languages (POPL 2008), pp. 171–182 (2008)
13. Lahiri, S.K., Qadeer, S.: Call invariants. Technical Report MSR-TR-2009-13, Microsoft Research (2009)
14. Leino, K.R.M., Nelson, G.: Data abstraction and information hiding. *ACM Trans. Program. Lang. Syst.* 24(5), 491–553 (2002)
15. McCreight, A., Shao, Z., Lin, C., Li, L.: A general framework for certifying garbage collectors and their mutators. In: Programming Language Design and Implementation (PLDI 2007), pp. 468–479 (2007)
16. O’Hearn, P.W., Reynolds, J.C., Yang, H.: Local reasoning about programs that alter data structures. In: Fribourg, L. (ed.) CSL 2001 and EACSL 2001. LNCS, vol. 2142, pp. 1–19. Springer, Heidelberg (2001)
17. Rakamarić, Z., Bingham, J., Hu, A.J.: An inference-rule-based decision procedure for verification of heap-manipulating programs with mutable data and cyclic data structures. In: Cook, B., Podolski, A. (eds.) VMCAI 2007. LNCS, vol. 4349, pp. 106–121. Springer, Heidelberg (2007)

18. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Logic in Computer Science (LICS 2002), pp. 55–74 (2002)
19. Rinetzky, N., Bauer, J., Reps, T.W., Sagiv, S., Wilhelm, R.: A semantics for procedure local heaps and its abstractions. In: Symposium on Principles of Programming Languages (POPL 2005), pp. 296–309. ACM, New York (2005)
20. Satisfiability Modulo Theories Library (SMT-LIB),
<http://goedel.cs.uiowa.edu/smtlib/>
21. Windows Driver Kit,
<http://www.microsoft.com/whdc/devtools/wdk/default.aspx>
22. Yang, H., Lee, O., Berdine, J., Calcagno, C., Cook, B., Distefano, D., O’Hearn, P.W.: Scalable shape analysis for systems code. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 385–398. Springer, Heidelberg (2008)