15-414: Bug Catching: Automated Program Verification

Lecture Notes on Arrays and Ghosts

Matt Fredrikson

Carnegie Mellon University
Lecture 4
January 23, 2025

1 Introduction

At this point we have experimented with simple imperative programs over integers using loops, recursive functions, and immutable data structures. Today, we start by looking at mutable data structures, specifically arrays. How do we specify and verify simple loops that operate on arrays? The key constructs we have (loop invariants and variants) are sufficient but become more complex because they have to express not only properties of what we *change* in an array, but also about the parts that *do not change*. It is this additional requirement that makes reasoning about ephemeral (mutable) data structures in many cases more difficult than reasoning about persistent (immutable) ones.

After we understand the basics, we consider a particularly important technique to tame more complex structures, namely simple *models* of complex data structures. Consider, for example, a red/black tree implementation of a map. To control its complexity we have introduced the concept of a data structure invariant. In the red/black tree example, this would include the *ordering invariant* (keys in left subtrees are all smaller and keys in the right subtree are all larger than the key in a node), the *color invariant* (there are no two adjacent red nodes in the tree), and the *black height invariant* (the number of black nodes on any path from a leaf to the root is the same). These are all *internal invariants* in the sense that the implementation of the data structure must maintain them but the client should only care that red/black trees provide a correct and efficient implementation of a map from keys to values. In this case, we say the map provides a *model* of the intended behavior of the data structure. We would like the model to be *logical* and as high-level as possible to support reasoning by the client. Maps and sets

are common models. In today's lecture we exemplify sets as a model of an ephemeral (mutable) implementation of bit vectors as arrays.

When implementing a data structure we have to maintain the correspondence between the low-level implementation and the high-level model. The purpose of the model is to *reason* about a data structure, but not to *compute* with it. So we would like to *erase* the code that maintains the model: actually computing it would negate all the advantages of the efficient implementation! This is the primary purpose of *ghosts*. They are pieces of code or data that exist solely for the purpose of verification and do not contribute to the outcome of the computation. This has to be checked by the verification engine. Ghost variables, or ghost fields of records, can only be used in other ghost computations. Otherwise, erasing them before the program is run would lead to incorrect code. This condition is related to the fact that *executable contracts* in C0 can not have any externally observable effects: running the program with or without executable contracts should yield the same answer (as long as all contracts are satisfied, of course).

Learning goals. After this lecture, you should be able to:

- Learn to formulate data structure invariants, and how work with them in Why3
- Verify simple imperative programs computing over arrays
- Model data structures using ghosts

2 Arrays

Because of their efficiency arrays are a common data structure in imperative programs. Reasoning about arrays requires a number of techniques due to their inherent mutability and range requirements for array access.

When we modify an array as we traverse it the loop invariant generally has to be more complicated. This is because with any assignment to an array element inside a loop, we lose all information about what *any* element in the array may be. Therefore, we generally need to specify the entries of the array we change and how, and in addition that the remaining entries do not change.

As an example, we consider a function to negate every element in the set. In order to verify this, we define a predicate negated. As is common (and perhaps we should have done this for the mem predicate) test the property for a slice of the array in the interval $[\mathsf{lower}, \mathsf{upper})$ (inclusive the lower bound and exclusive the upper bound).

```
1 predicate negated (a : array int) (b : array int) (lower : int) (upper
      : int) =
2 a.length = b.length /\ 0 <= lower <= a.length /\ 0 <= upper <= a.
     length
3 /\ forall j:int. lower <= j < upper -> a[j] = -b[j]
```

We also ensure that the arrays a and b have the same length.

The postcondition for our negation function expresses that the state of the array upon the return is equal to the negated initial values of the array all the way up to the last element.

```
let negate (a: array int) : unit =
ensures { negated (old a) a 0 a.length}
let n = a.length in
for i = 0 to n-1 do
  (* no variant, or invariant on i needed *)
  invariant { ... }
  a[i] <- -a[i]
done ; ()
```

This code has two additional new constructs. We use the type `unit` (whose only element is `()`) to express that the function returns no interesting value. This usually implies that it mutates some of its arguments, in this case the array a.

We also use a `for` loop, of the form for $i = lower$ to $upper$ do ... done for which we give *inclusive bounds*. It generates suitable loop invariants for the index $lower \leq i \leq upper + 1$ and a variant of $upper + 1 - lower$. There is an analogous form for $i = upper$ downto $lower$ do ... done.

What remains is to state the invariants regarding the array. Intuitively, at iteration i we have negated all the elements up to i while the elements at indices greater or equal to i have remained unchanged. To specify this we find in the useful `array.ArrayEq` module the function

```
array_eq_sub (a : array 'a) (b : array 'a) (lower : int) (upper : int)
```

which is true if for all $lower \leq i < upper$ we have $a[i] = b[i]$. We use this where b is the original version of a.

```
let negate (a: array int) : unit =
ensures { negated (old a) a 0 a.length}
let n = a.length in
for i = 0 to n-1 do
  (* no variant, or invariant on i needed *)
  invariant { negated a (old a) 0 i }
  invariant { array_eq_sub a (old a) i n }
  a[i] <- -a[i]
done ; ()
```

Observe how the invariants in this form express concisely that the values in the interval $[0, i)$ are negated, while those in the interval $[i, n)$ are still the same as in the original array. It is generally easier to understand and express the invariant in this form than using complicated quantified formulas in-line.

It is not necessary to explicitly return the unit element (since the for-loop already returns the unit element), but we write it out for emphasis.

The code now verifies, but true to our methodology we can also try it out.

```
let test () =
let a = Array.make 4 0 in
( a[0] <- 1 ; a[1] <- 27 ; a[2] <- 4 ; a[3] <- 3 ;
```

```
4   (search 1 a , search 2 a , search 27 a ) )
```

We obtain the expected answer.

```
% why3 execute intset.mlw --use="IntSet" 'test ()'
result: (int, int, int) = (0, (-1), 1)
globals:
%
```

Instead of executing it, we can also *prove* that this must be the outcome of the function.

```
1 let test () =
2 ensures { result = (0, (-1), 1) }
3 let a = Array.make 4 0 in
4 ( a[0] <- 1 ; a[1] <- 27 ; a[2] <- 4 ; a[3] <- 3 ;
5   (search 1 a , search 2 a , search 27 a ) )
```

```
% why3 prove -P alt-ergo intset.mlw
File intset-test.mlw:
Verification condition search'vc.
Prover result is: Valid (0.01s, 94 steps).

File intset-test.mlw:
Verification condition negate'vc.
Prover result is: Valid (0.03s, 383 steps).

File intset-test.mlw:
Verification condition test'vc.
Prover result is: Valid (0.03s, 160 steps).
```

3 Data Structure Invariants

As a simple example of data structure invariants we reconsider our implementation of queues. We would like to extend the interface with a function `qsize` that returns the number of elements in a queue. The obvious way to compute this would be to compute the lengths of the front and back and add them, but the complexity would be linear in the size of the queue. To answer this query in constant time we add a `size` field to the queue and maintain it as we enqueue or dequeue elements. The data structure invariant here is that the size field always contains the sum of the lengths of the front and back.

```
1   type queue 'a = { front : list 'a ;
2                     back : list 'a ;
3                     size : int }
4   invariant { size = length front + length back }
```

In the logic, the verifier assumes the invariant when reasoning about queues. This could create a logical inconsistency (*everything is provable!*) if the invariant is unsatisfiable. For example, the trivial invariant *false* means any program in the scope of this

declaration could now be verified. To avoid this, Why3 will try to prove that there exists an instance of the data structure for which the invariant is satisfied. For complex invariants this can be difficult, so WhyML gives you a way to specify an instance of the data structure satisfying the invariant with a 'by' clause. In this particular case it would be unnecessary, but it is good practice to always supply it.

```
type queue 'a = { front : list 'a ;
                  back : list 'a ;
                  size : int }
invariant { size = length front + length back }
by { front = Nil ; back = Nil ; size = 0 }
```

In this example it is relatively easy to update the code to maintain the size field and Why3 will prove that it is always correct. In constructing the verification condition we may think of the invariant as being *assumed* at the beginning of a function (like a precondition) and *proved* at the end of a function (like a postcondition). In between, the invariant may be violated, although this possibility does not come into play here. It is necessary because you may build or modify an element of the data structure incrementally and only at the end does the invariant hold.

The `qsize` function just returns the size field. In addition, the postcondition certifies that it is indeed the length of the queue (when viewed as a single sequence of element, which represents the client's perspective.

```
let qsize (q : queue 'a) : int =
ensures { result = length (sequence q) }
q.size
```

4 Ghosts and Models

As an example of a model we use the standard set library to model a bit vector implementation of bounded finite sets. Here is an excerpt of the finite set module.

```
module Fset
  type fset 'a
  predicate mem (x: 'a) (s: fset 'a)
  predicate is_empty (s: fset 'a) = forall x: 'a. not (mem x s)
  constant empty: fset 'a
  function add (x: 'a) (s: fset 'a) : fset 'a
  axiom add_def: forall x: 'a, s: fset 'a, y: 'a.
      mem y (add x s) <-> (mem y s \/ y = x)
  function remove (x: 'a) (s: fset 'a) : fset 'a
  axiom remove_def: forall x: 'a, s: fset 'a, y: 'a.
      mem y (remove x s) <-> (mem y s /\ y <> x)
  ...
end
```

If you look at the full listing in Why3's standard library, you will quickly notice that all of the predicates and functions are defined using axioms rather than being given computational definitions. From a practical standpoint, this means that nothing defined in `Fset` can be used to perform computations—it can be used *only* in specifications.

Our goal will be to implement a computational finite set that matches the theory given in `Fset`.

4.1 A computational finite set

We start by defining the `Bitset` module by defining a `bset` as a record consisting of an array `a`, a `bound` and a ghost field called `model` containing a finite set of integers. Because bitsets are mutable (for example, we actually *change* a bset by adding an element to it), the `model` field must also be mutable.

```
type bset = { a : array bool ;
              bound : int ;
              mutable ghost model : Fset.fset int }
invariant { 0 <= bound <= a.length
            /\ forall j. 0 <= j < a.length -> a[j] <-> Fset.mem j
                model }
by { a = Array.make 0 false ; bound = 0 ; model = Fset.empty }
```

The invariant states that element $a[i]$ of the array is true if and only if the number i is in the model set. We witness the existence of such a bset with the empty array and empty model. Note that the fields `a` and `bound` are immutable, although the elements in the array are mutable.

To create an empty bset we need a bound on the elements we may add to the set, which will be the length of the array.

```
let empty_bset (bound : int) : bset =
requires { bound >= 0 }
ensures { Fset.is_empty result.model /\ result.bound = bound }
{ a = Array.make bound false ; bound = bound ; model = Fset.empty }
```

The model is just the empty set. Note that the postcondition states that `empty_bset` models the empty finite set.

To add an element i to a bset we just set the corresponding array element to `true` (whether it was already true or not). This requires the precondition that the i is in the permissible range. Because this operation is destructive, modifying the given bset, the postcondition needs to state the model *after* the update is equal to the model *before* the update, plus the element i. For this purpose we use again the `old` keyword to refer to the state of the model at the time the function is called.

```
let add_bset (x : int) (s : bset) : unit =
requires { 0 <= x < s.bound }
ensures { s.model = Fset.add x (old s).model }
s.a[x] <- true ;
ghost (s.model <- Fset.add x s.model) ;
()
```

The assignment to `s.model` is enclosed in `ghost (...)` to be explicit that this update of the model will not be carried out if the program is executed. Instead it is there to maintain the data structure invariant which must be restored before the function `add_bset` exits. Just before this line, by the way, the data structure invariant is false

because we have updated the array but not the model. We can see the significance of checking the data structure invariance exactly at function boundaries.

Our postcondition will allow the client to reason about the effects of it add operations. Note that the pre- and post-conditions *do not reference the array*, only properties of the model and the bound. The client can reason about the behavior of these functions without knowing the representation, using only the model.

The remove operation is entirely analogous.

```
let remove_bset (i : int) (s : bset) : unit =
requires { 0 <= i < s.bound }
ensures { s.model = Fset.remove i (old s).model }
s.a[i] <- false ;
ghost (s.model <- Fset.remove i s.model) ;
()
```

We did not implement any more complex operations such as union or intersection, even though this would certainly be possible. You can find the live-code `Bitset` module in the file bitset.mlw.

5 Appendix: Diagnosis of Failing Goals

We provide here a brief illustration how we can use the Why3 IDE to isolate failures of verification. We change the source so that the middle invariant

```
  invariant { forall j. 0 <= j < i -> a[j] <> x }
```

has an off-by-one error

```
  invariant { forall j. 0 <= j <= i -> a[j] <> x }
```

We start the Why3 IDE with

```
why3 ide intset.mlw
```

select the outermost goal and hit '2' for a relatively advanced strategy. It cannot prove the verification condition but splits it into several parts. We can see that many subgoals are checked as green, but two of them still have question marks. Goals depending on

them higher up in the goal/subgoal tree will then be similarly marked.

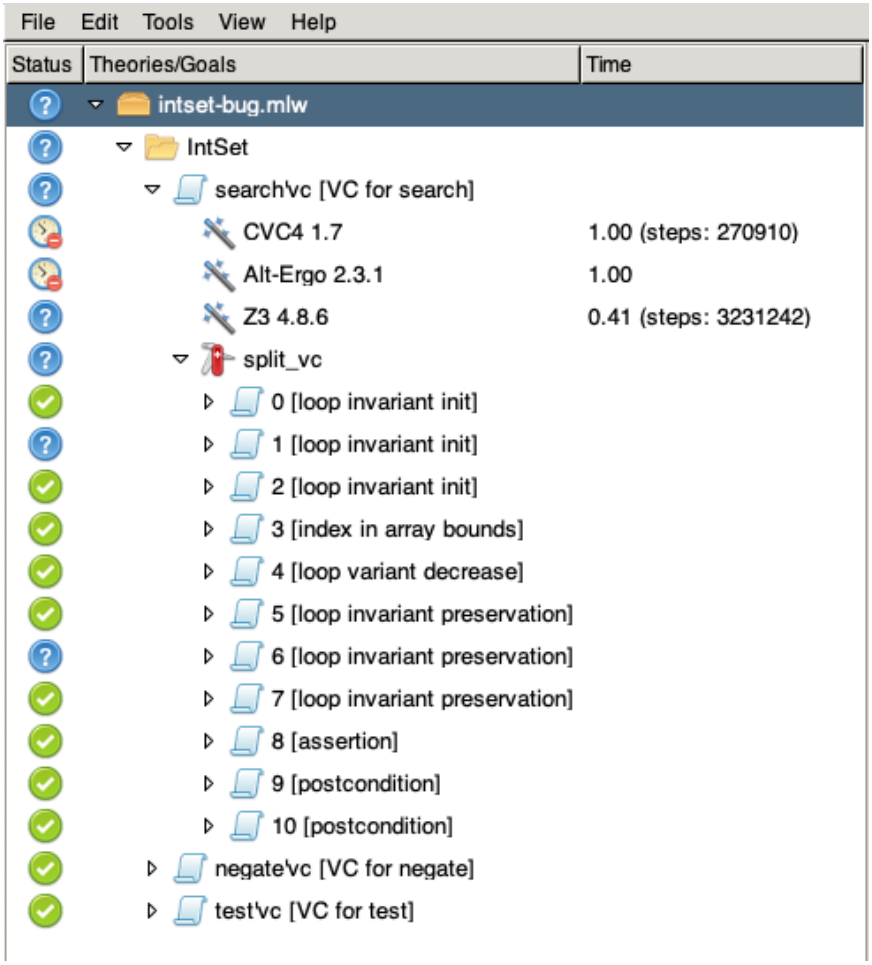

The unproven subgoals are labeled `[loop invariant init]` and `[loop invariant preservation]` which indicates that there is a loop invariant that cannot be established initially, nor can it be shown to be preserved. To see which one we select the second one in the pane on

the left and examine the program in the pane on the right.

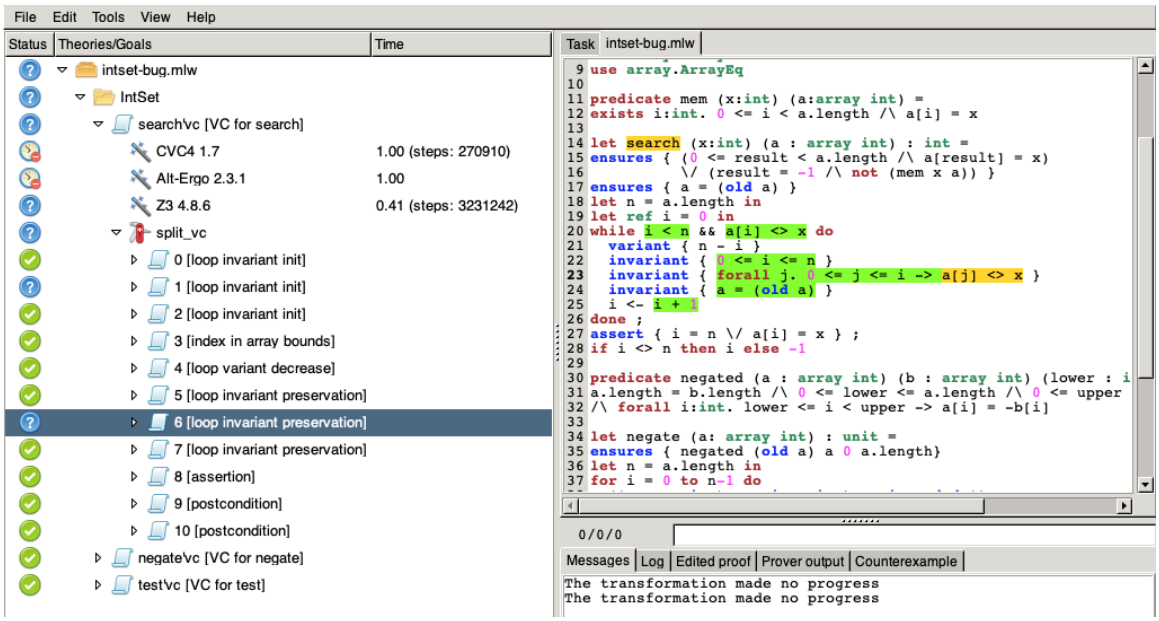

The IDE will highlight in green the assumptions it uses and in yellow the proof goal it is trying to prove. Here we see it is `a[j] <> x` in the middle invariant. Even though it does not occur here, the IDE will highlight formulas in your program as red if it uses its negation in the proof attempt. This occurs for the else-branch of conditionals and loop guards when reasoning about the state after the loop is exited.

Highlighting the other unprovable subgoal (labeled `[loop invariant init]`) leads to the same culprit. Some further forensics and thought about this will hopefully reveal the bug at this point. Fortunately, it is not in the program but in the invariant.