

Instructor: Matt Fredrikson

TA: Ryan Wagner

Incremental Proof Development in Dafny

In this discussion, we'll see in more detail how to go about proving the total correctness of imperative code in Dafny. We'll use the precondition method discussed in last lecture to develop loop invariants and assertions that are strong enough to establish our specification, while simultaneously developing the proof in Dafny.

The program that we'll consider implements bubble sort, a basic sorting algorithm with worst-case performance $O(n^2)$, where n is the number of elements to be sorted. Although there are more efficient sorting algorithms, bubble sort is a good way of illustrating several techniques and concepts that are useful when reasoning about the correctness of imperative programs.

Partial Correctness

We begin by writing the basic code, absent any specification or annotation. Recall that bubble sort takes a single array as input, and is composed of a nested loop. The outer loop begins indexing at the end of the array, working its way towards the beginning, while the inner loop starts at the beginning of the array and works towards the current index of the outer loop. With each iteration of the outer loop, the largest element from the beginning of the array is moved towards the end by swapping out-of-order pairs in the inner loop. In this way, the loop becomes more sorted with each iteration of the outer loop by ensuring that elements in the range $[i, a.Length - 1]$ are always sorted.

```

1 method BubbleSort(array<int> a)
2 {
3     var i := a.Length - 1;
4     while(i > 0)
5     {
6         var j := 0;
7         while(j < i)
8         {
9             if(a[j] > a[j+1])
10            {
11                a[j], a[j+1] := a[j+1], a[j];
12            }
13            j := j + 1;
14        }
15        i := i - 1;
16    }
17 }
```

Regarding the specification, our basic requirement is that whenever `BubbleSort` terminates, the contents of `a` are in sorted ascending order. We encode this requirement in a predicate. The predicate is *true* whenever for any indices $i \leq j$ within the index bounds, $a[i] \leq a[j]$.

```

1 predicate sorted(a: array<int>, l: int, u: int)
2   reads a
3   requires a != null
4 {
5   forall i, j :: 0 <= l <= i <= j <= u < a.Length ==> a[i] <= a[j]
6 }
```

`sorted` takes three arguments: the array `a`, a lower-bound index `l`, and an upper-bound index `u`. Note that our predicate essentially ignores calls with invalid index bounds, so that if $0 \leq l \leq u < a.Length$ is *false*, then the predicate will always return *true* regardless of the contents of `a`. There are two additional annotations present in the signature of `sorted`:

1. `reads a` tells Dafny that this predicate will read the contents of `a`. Whenever we write a function or predicate that examines the contents of an array, we need to use this annotation to let Dafny know.
2. `requires a != null` is a precondition on our predicate which says that we are not allowed to invoke `sorted` on an array reference that might be set to `null`. This means that we are not defining sortedness on `null` array references, so whenever we want to reason about this predicate, we must prove that the reference we pass takes a proper value.

Having defined what it means for an array to be sorted in ascending order, we can write the postcondition for our specification to state that `sorted` holds for every index in `a` (i.e., between 0 and `a.Length - 1`):

```

1   ensures sorted(0, a, a.Length-1)
```

Because our postcondition stipulates that the reference we're passing to `sorted` isn't `null`, we will also need this precondition in our method. This is also evident by the lack of dynamic checks in `BubbleSort` for `null` values; if we didn't have this precondition, we could potentially dereference a `null` pointer:

```

1   requires a != null
```

Looking for basic facts to add to our annotations, we start by observing that the outer loop index variable `i` will always lie in the range $[-1, a.Length - 1]$. For the inner loop, we can write a stronger invariant for `i`, namely that it will lie in the range $[1, a.Length - 1]$. Additionally, the inner-loop index variable `j` will lie in the range $[0, i]$, giving us our initial set of basic annotations.

```

1 method BubbleSort(a: array<int>)
2   requires a != null
3   ensures sorted(a, 0, a.Length-1)
4 {
5   var i := a.Length - 1;
6   while(i > 0)
7     invariant -1 <= i < a.Length
8   {
9     var j := 0;
10    while(j < i)
11      invariant 0 < i < a.Length && 0 <= j <= i
12    {
13      if(a[j] > a[j+1])
14      {
15        a[j], a[j+1] := a[j+1], a[j];
16      }
17      j := j + 1;
18    }
19    i := i - 1;
20  }
21 }

```

These annotations are clearly not strong enough to prove our desired postcondition, as they say nothing about the contents of **a**. We'll need to strengthen them, and in particular the loop invariants. Starting with the outer loop, we propagate the postcondition backwards to the point where the outer loop terminates. This will tell us what conditions need to hold when the loop exits to support the postcondition, which will hopefully guide our intuition towards better loop invariants.

```

1 invariant -1 <= i < a.Length
2 assume i <= 0
3 ensures sorted(a, 0, a.Length-1)

```

Computing the verification condition,

$$\begin{aligned}
 -1 \leq i < a.Length &\rightarrow \text{wlp}(\text{assume } i \leq 0, \text{sorted}(a, 0, a.Length - 1)) \\
 \Leftrightarrow (-1 \leq i < a.Length \wedge i \leq 0) &\rightarrow \text{sorted}(a, 0, a.Length - 1) \\
 \Leftrightarrow -1 \leq i \leq 0 &\rightarrow \text{sorted}(a, 0, a.Length - 1)
 \end{aligned}$$

This isn't all that surprising, and just points out that when the loop terminates, if **i** is either -1 or 0 , then the contents of **a** should be sorted. Informally, we know that **i** will be either -1 or 0 when the loop terminates, and if **i** = -1 , then **sorted** will necessarily be true, so this fact doesn't seem to tell us anything new. We want to generalize this fact to something that reflects the progress made at each iteration of the outer loop, such that when it is conjoined with the negated loop guard $\neg(i > 0)$, the postcondition will hold. After a bit of thought, we arrive at:

$$\text{sorted}(a, i, a.Length - 1)$$

Adding this to our outer loop invariant makes sense, as it corresponds to our observation that as the method proceeds, it maintains the fact that the array remains sorted in the range $[i, a.Length - 1]$. Additionally, we see that:

$$\neg(i > 0) \wedge \text{sorted}(a, i, a.Length - 1) \rightarrow \text{sorted}(a, 0, a.Length - 1)$$

Continuing on, we need to find an invariant for the inner loop that will be sufficient to preserve our outer invariant. Applying the precondition method from the exit point of the inner loop to the next iteration of the outer loop, we have:

```

1 invariant 0 < i < a.Length && 0 <= j <= i
2 assume i <= j;
3 i := i - 1
4 invariant -1 <= i < a.Length && sorted(a, i, a.Length-1)

```

We can effectively ignore the portion of the postcondition constraining the bounds of i , as a straightforward calculation will verify that they are preserved by the loop. Then, computing the corresponding verification condition:

$$\begin{aligned}
 & 0 < i < a.Length \wedge 0 \leq j \leq i \rightarrow \text{wlp}(\text{assume } i \leq j; i := i - 1, \text{sorted}(a, i, a.Length - 1)) \\
 \Leftrightarrow & (0 < i < a.Length \wedge 0 \leq j \leq i \wedge i \leq j) \rightarrow \text{sorted}(a, i - 1, a.Length - 1) \\
 \Leftrightarrow & (0 < i < a.Length \wedge 0 \leq j \wedge i = j) \rightarrow \text{sorted}(a, i - 1, a.Length - 1)
 \end{aligned}$$

We might be tempted to generalize this fact by taking the consequent and adding it to our inner loop invariant, giving us:

```

1 method BubbleSort(a: array<int>)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5 {
6   var i := a.Length - 1;
7   while(i > 0)
8     invariant -1 <= i < a.Length
9     invariant sorted(a, i, a.Length-1)
10 {
11   var j := 0;
12   while(j < i)
13     invariant 0 < i < a.Length && 0 <= j <= i
14     invariant sorted(a, i-1, a.Length-1)
15   {
16     if(a[j] > a[j+1])
17     {
18       a[j], a[j+1] := a[j+1], a[j];
19     }
20     j := j + 1;
21   }
22   i := i - 1;
23 }
24 }

```

This would make the basic path from the beginning of the outer loop to the beginning of the inner loop:

```

1 invariant -1 <= i < a.Length && sorted(a, i, a.Length-1)
2 assume 0 < i;
3 j := 0
4 invariant 0 < i < a.Length && 0 <= j <= i && sorted(a, i-1, a.Length-1)

```

Which would lead to the verification condition:

$$0 < i < a.Length \wedge \text{sorted}(a, i, a.Length - 1) \rightarrow \text{sorted}(a, i - 1, a.Length - 1)$$

This is not valid, as there is nothing in the antecedent to imply that a larger portion of the list is sorted (i.e., $[i - 1, a.Length - 1]$) than what was assumed (i.e., $[i, a.Length - 1]$). We can observe the following counterexample as concrete evidence that this condition is invalid:

$$\begin{array}{ccccc} 1 & 2 & 3 & 5 & 4 \\ & & & (j) & (i) \end{array}$$

This array satisfies the antecedent, because it is sorted starting at i . It does not satisfy the consequent, because it is not sorted starting at $i - 1$. To find a suitable invariant, we will need to give this more thought. In the meantime, we will add the main invariant on the outer loop,

$$\text{sorted}(a, i, a.Length - 1)$$

to the inner loop, as we know that at the very least, it must be preserved.

It oftentimes helps to step through the execution of the code in question, examining they key invariants that arise in actual traces. Let's look at a small trace of bubble sort on the array $[4, 1, 3, 2]$.

4 (j)	1	3	2 (i)	Swap 4, 1
1	4 (j)	3	2 (i)	Swap 4, 3
1	3 (j)	4	2 (i)	Swap 4, 2
1	3	2	4 (j) (i)	
1 (j)	3	2 (i)	4	
1 (j)	3	2 (i)	4	Swap 3, 2
1 (j)	2	3	4 (j) (i)	
1 (j)	2 (i)	3	4	
1 (j) (i)	2	3	4	
1 (j) (i)	2	3	4	

Considering the outer invariant we're currently trying to support, one fact stands out immediately: throughout execution, all elements of a indexed by a value greater than i are greater than those that are at most i . Formally, we observe the invariant:

$$\forall k, k'. 0 \leq k \leq i \wedge i < k' < a.Length \rightarrow a[k] \leq a[k']$$

We'll rename this invariant `partitioned(a, i)`, and add the predicate to our code.

```

1 predicate partitioned(a: array<int>, i: int)
2   reads a
3   requires a != null
4 {
5   forall k, k' :: 0 <= k <= i < k' < a.Length ==> a[k] <= a[k']
6 }
```

This is pertinent to our main invariant, because bubble sort works by maintaining a partition on `a`: the first half is unsorted, whereas the second half is sorted, and per our new invariant, contains values that are at least as large as those in the unsorted portion. We add this invariant to both loops, as is a property of the outer loop that is preserved by the inner loop.

```

1 method BubbleSort(a: array<int>)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5 {
6   var i := a.Length - 1;
7   while(i > 0)
8     invariant -1 <= i < a.Length
9     invariant sorted(a, i, a.Length-1)
10    invariant partitioned(a, i)
11    {
12      var j := 0;
13      while(j < i)
14        invariant 0 < i < a.Length && 0 <= j <= i
15        invariant sorted(a, i, a.Length-1)
16        invariant partitioned(a, i)
17        {
18          if(a[j] > a[j+1])
19          {
20            a[j], a[j+1] := a[j+1], a[j];
21          }
22          j := j + 1;
23        }
24        i := i - 1;
25      }
26 }
```

After attempting to verify the above code, it appears that our invariants are still not strong enough. The verifier's error message tells us that the two non-basic outer invariants might not be preserved by the loop. Let's examine the relevant basic path, and its corresponding verification condition to gain further clarity.

```

1 invariant 0 < i < a.Length && 0 <= j <= i &&
2             sorted(a, i, a.Length-1) &&
3             partitioned(a, i)
4 assume i <= j;
5 i := i - 1;
6 invariant -1 <= i < a.Length &&
7             sorted(a, i, a.Length-1) &&
8             partitioned(a, i)

```

The corresponding VC, after a bit of simplification, is:

$$\begin{aligned}
 0 < i < a.Length \wedge i = j \wedge \text{sorted}(a, i, a.Length - 1) \wedge \text{partitioned}(a, i) \\
 \rightarrow \text{sorted}(a, i - 1, a.Length - 1) \wedge \text{partitioned}(a, i - 1)
 \end{aligned}$$

This invariant is not valid, and the problem is similar to what we encountered earlier when we tried to generalize by accepting an invariant that widened the sorted region. Consider the following counterexample, which satisfies the antecedent but not the consequent.

$$\begin{array}{ccccc}
 2 & 1 & 4 & 3 & 5 \\
 (j) & (i)
 \end{array}$$

However, a bit of thought leads us to conclude that bubble sort would never produce this counterexample state. We know this because bubble sort always propagates the *largest* value from the unsorted region to the bottom of the sorted region. Indeed, throughout the inner loop $a[j]$ contains the largest value processed so far. Formalized, the corresponding invariant is relevant to the inner loop:

$$\forall k. 0 \leq k \leq j \rightarrow a[k] \leq a[j]$$

After adding this, we have an implementation and proof that Dafny “almost” accepts.

```

1 method BubbleSort(a: array<int>)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5 {
6   var i := a.Length - 1;
7   while(i > 0)
8     invariant -1 <= i < a.Length
9     invariant sorted(a, i, a.Length-1)
10    invariant partitioned(a, i)
11  {
12    var j := 0;
13    while(j < i)
14      invariant 0 < i < a.Length && 0 <= j <= i
15      invariant sorted(a, i, a.Length-1)
16      invariant partitioned(a, i)
17      invariant forall k :: 0 <= k <= j ==> a[k] <= a[j]
18    {
19      if(a[j] > a[j+1])
20      {
21        a[j], a[j+1] := a[j+1], a[j];
22      }
23      j := j + 1;
24    }
25    i := i - 1;
26  }
27 }

```

At this point, the verifier no longer complains about the loop invariants being preserved. Rather, it notifies us that the postcondition might not hold on a path through our method, and points to the head of the loop as the possible location at which our code may be incorrect (or our proof incomplete). What is going on here? The loop exit condition conjoined with our key loop invariant appears sufficient to establish the postcondition:

$$\text{sorted}(\mathbf{a}, \mathbf{i}, \mathbf{a.Length} - 1) \wedge \mathbf{i} \leq 0 \rightarrow \text{sorted}(\mathbf{a}, 0, \mathbf{a.Length} - 1)$$

Informally, we know that when the loop exits, \mathbf{i} is either 0 or -1 . When $\mathbf{i} = 0$, the relevant VC is certain to hold:

$$\text{sorted}(\mathbf{a}, \mathbf{i}, \mathbf{a.Length} - 1) \wedge \mathbf{i} = 0 \rightarrow \text{sorted}(\mathbf{a}, 0, \mathbf{a.Length} - 1)$$

When $\mathbf{i} = -1$, then we know that $\mathbf{a.Length} = 0$, and our **sorted** predicate becomes:

$$\begin{aligned} \forall i, j. 0 \leq 1 \leq i \leq j \leq u < -1 \rightarrow a[i] \leq a[j] \\ \Leftrightarrow \text{false} \rightarrow a[i] \leq a[j] \end{aligned}$$

So it would seem that the postcondition holds in either case. It seems that the verifier is not picking up on the relationship we need,

$$\mathbf{i} < 0 \rightarrow \mathbf{a.Length} = 0$$

This is clearly true before the outer loop executes (i.e., immediately after \mathbf{i}) is assigned, and cursory inspection of the loop body suggests that it is maintained until the loop terminates. We can add it as an invariant to our outer loop, and the verifier accepts our proof.

```
1 method BubbleSort(array<int> a)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5 {
6   var i := a.Length - 1;
7   while(i > 0)
8     invariant i < 0 ==> a.Length == 0
9     invariant -1 <= i < a.Length
10    invariant sorted(a, i, a.Length-1)
11    invariant partitioned(a, i)
12  {
13    var j := 0;
14    while(j < i)
15      invariant 0 < i < a.Length && 0 <= j <= i
16      invariant sorted(a, i, a.Length-1)
17      invariant partitioned(a, i)
18      invariant forall k :: 0 <= k <= j ==> a[k] <= a[j]
19    {
20      if(a[j] > a[j+1])
21      {
22        a[j], a[j+1] := a[j+1], a[j];
23      }
24      j := j + 1;
25    }
26    i := i - 1;
27  }
28 }
```

Strengthening the Specification. We have managed to prove that when `BubbleSort` terminates, its parameter contains a sorted list whenever it begins non-null. However, our specification is not as strong as we may like it to be for some purposes. For example, consider the following program, which satisfies the specifications we've given to `BubbleSort`.

```

1 method BubbleSort(a: array<int>)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5 {
6   var i := a.Length - 1;
7   while(i >= 0)
8     invariant -1 <= i < a.Length
9     invariant i >= 0 ==> forall j :: i < j < a.Length ==> a[j] == a[a.Length-1]
10    invariant sorted(a, i+1, a.Length-1)
11  {
12    a[i] := a[a.Length-1];
13    i := i - 1;
14  }
15
16 }
```

While a well-intentioned programmer is unlikely to deviate in spirit so far from the contract, the point remains that our specification does not require that the contents of **a** differ in the post-state only in their order from the pre-state; this is crucial to our expectations of what a sorting procedure does.

We can strengthen our specification by adding to the postcondition that the final contents of **a** be a *permutation* of their initial contents. Dafny has a syntactic facility for referring to the initial contents of a reference type in specifications: the `old(a)` notation.

```

1 method BubbleSort(a: array<int>)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1)
5   ensures permutation(a[...], old(a[...]))
```

This specification says that when `BubbleSort` terminates, the new contents of **a** must be a permutation of their initial contents. Notice that when we refer to **a** in the arguments of `permutation`, we suffix with `[...]`. This is Dafny's notation for computing an *array slice*, which returns the contents of an array between specified indices as a `seq`. Because we do not provide any indices, this returns the entire contents of **a** as a `seq<int>`, and we do this so that the `permutation` predicate we write can take `seq<int>` arguments rather than `array<int>`; sequences are nicer to work with in specifications than arrays. We could have computed a sequence corresponding to any range of **a** by writing `a[low..high]` had we needed to refer to specific portions of **a**.

Now we must write a predicate for `permutation`. Intuitively, we know that an array **a** is a permutation of **b** if and only if for each possible value **v**, the number of occurrences of **v** in **a** is the same as the number of occurrences in **b**. So, to encode the right predicate we will need a function that counts the number of occurrences of a particular value in a sequence.

```

1 function count(a: seq<int>, v: int): nat
2 {
3     if ( $|a| > 0$ ) then
4         if (a[0] == v) then 1 + count(a[1..], v)
5         else count(a[1..], v)
6     else 0
7 }
```

With this function in hand, we can write our predicate for `permutation` in a straightforward manner.

```

1 predicate permutation(a: seq<int>, b: seq<int>)
2 {
3     forall v :: count(a, v) == count(b, v)
4 }
```

When we attempt to verify our method against the new specification, Dafny tells us that the postcondition might not hold. Of course, it is referring to the postcondition mentioning `permutation`, which means that our current annotations are not sufficient to establish the newly-strengthened specification. We must obtain new annotations sufficient to prove this property.

Unlike in the case of `sorted`, our implementation does not gradually establish that the current `a` is a permutation of `old(a)`. So, adding an invariant that incrementally asserts that a larger portion of the array is a permutation of the old may not be necessary, and in fact may be more difficult to prove than the straightforward choice of:

```
1 invariant perm(old(a[..]), a[..])
```

Intuitively, the only way in which our code modifies the contents of `a` is by swapping adjacent pairs of elements. If we place this invariant, which is clearly strong enough to prove the postcondition in question, on both loops, then we can reasonably expect to prove it by reasoning about localized portions of `a` near the point of our modifications inside the inner loop. However, Dafny does not accept our annotations yet.

Although it may seem obvious to us that no elements are added or removed from `a` in the line:

```
1 a[j], a[j+1] := a[j+1], a[j];
```

Dafny does not know how to connect these manipulations of the array to the properties relevant for `permutation`. We need to demonstrate that the contents of `a` immediately after the update are a permutation of the contents before it. To do this, it will be helpful to have some way of referring to both versions of `a` simultaneously. Dafny supports *ghost variables* for exactly this purpose. Ghost variables are declared with the modifier `ghost`, and are used only at verification time; they are not compiled into variables that are allocated and updated when the program is run, so we do not need to worry about the impact on performance that manipulating these variables would otherwise have.

We declare and initialize a new ghost variable to track the contents of `a` before the update:

```

1 ghost var a' := a[..];
2 a[j], a[j+1] := a[j+1], a[j];

```

We now need to demonstrate that `permutation(a[..], a')` holds after the update:

```

1 ghost var a' := a[..];
2 a[j], a[j+1] := a[j+1], a[j];
3 assert permutation(a[..], a');

```

The verifier still is not convinced that we've proven the claim, although we can check to make sure that this assertion is sufficient to establish our loop invariants by temporarily switching the `assert` to an `assume`:

```

1 ghost var a' := a[..];
2 a[j], a[j+1] := a[j+1], a[j];
3 assume permutation(a[..], a');

```

When we do this, the loop invariants and postconditions successfully verify, so we know that we're on the right track. This is an effective proof debugging technique that should be used frequently to ensure that we aren't wasting time proving facts that are irrelevant to our specifications.

Perhaps the verifier is having difficulty reasoning monolithically about the entirety of `a` and `a'`. To help it along, let's break both values down into simpler parts that highlight their differences. We know that the only positions in which the two differ are `j, j + 1`. So, it should be the case that

$$a[..j] = a'[..j] \wedge a[j + 2..] = a'[j + 2..]$$

A quick check with an `assert` verifies that this is the case, and is recognized by Dafny. However, we don't need this assertion by itself in our proof, so we remove it before moving on.

Decomposing `a` and `a'` in this manner makes it obvious that the two are permutations of each other: `a[..j]` and `a'[..j]` are permutations of each other, `a[j + 2..]` and `a'[j + 2..]` are as well, and finally `a[j..j + 2]` and `a'[j..j + 2]` are as well. However, Dafny does not accept this final fact, so we make it very clear by being explicit about the contents of `a[j..j + 2], a'[j..j + 2]`:

```

1 ghost var a' := a[..];
2 a[j], a[j+1] := a[j+1], a[j];
3 ghost var v1, v2 := a[j], a[j+1];
4 assert a[..] == a[..j] + [v1, v2] + a[j+2..];
5 assert a' == a[..j] + [v2, v1] + a[j+2..];
6 assert permutation([v1, v2], [v2, v1]);
7 assume permutation(a[..], a');

```

The only line that the verifier will not accept is the final. It seems that we have done enough to prove the desired property, but let's take a closer look to understand where Dafny might be getting tripped up. `permutation` is a predicate over `count`, and we expect that `a` is a permutation of `a'` because for every value, `count` is the same in `a[..j]` and `a'[..j]`, `a[j..j + 2]` and `a'[j..j + 2]`, and `[v1, v2], [v2, v1]`. Because `a` and `a'` are composed of only these parts, we expect that `count` on the full contents will match up as well. We can check this intuition with an assertion:

```

1 assert forall v :: count(a[..], v) == count(a[..j], v) + count([v1, v2], v) +
2   count(a[j+2..], v);

```

The verifier complains, so presumably this is our culprit. Dafny does not realize that `count` is distributive across sequence concatenation, i.e., that:

$$\forall v. \text{count}(a + b + c, v) = \text{count}(a, v) + \text{count}(b, v) + \text{count}(c, v)$$

Now that we know this is the issue, it is not particularly surprising: `count` is non-trivial function over sequences, and we require quantified reasoning over its behavior. We will need to prove it, and because it is a general property that does not depend on the specifics of `BubbleSort`, we will do so by creating a lemma that can potentially be reused later.

```

1 lemma count_distrib3(a: seq<int>, b: seq<int>, c: seq<int>)
2   ensures forall v :: count(a + b + c, v) ==
3     count(a, v) + count(b, v) + count(c, v)

```

Before we begin proving the main argument that `count` is distributive across concatenation, we observe that the signature of this lemma may be more complicated than we'd like. The context in which we'd like to use the lemma needs to relate this property to the concatenation of three sequences, but intuition should tell us that we can achieve this by sequencing assertions of distributivity with respect to pairs of sequences. Formally, this corresponds to the claim that:

$$\begin{aligned} \forall a, b, c, v. \quad & (\text{count}(a + b, v) = \text{count}(a, v) + \text{count}(b, v) \wedge \\ & \text{count}((a + b) + c, v) = \text{count}(a + b, v) + \text{count}(c, v)) \\ \rightarrow & \text{count}(a + b + c, v) = \text{count}(a, v) + \text{count}(b, v) + \text{count}(c, v) \end{aligned}$$

To reason in this way, we need a lemma `count_distrib2` with the signature:

```

1 lemma count_distrib2(a: seq<int>, b: seq<int>)
2   ensures forall v :: count(a + b, v) == count(a, v) + count(b, v)

```

This allows us to complete `count_distrib3` by invoking `count_distrib2` twice:

```

1 lemma count_distrib3(a: seq<int>, b: seq<int>, c: seq<int>)
2   ensures forall v :: count(a + b + c, v) ==
3     count(a, v) + count(b, v) + count(c, v)
4 {
5   count_distrib2(a, b);
6   count_distrib2(a + b, c);
7 }

```

Note that we could have used `count_distrib2` directly in the body of `BubbleSort` to prove the claim we're after, but it would have cluttered the proof with more calls to the lemma.

Now we need to prove the postcondition for `count_distrib2`. This is a universally-quantified claim over all sequence element values `v`, so we'll start off with the `forall` statement in the body, which allows us to write a postcondition that will hold for every value in the iteration range of the statement:

```

1 lemma count_distrib2(a: seq<int>, b: seq<int>)
2   ensures forall v :: count(a + b, v) == count(a, v) + count(b, v)
3 {
4   forall(v: int)
5     ensures count(a + b, v) == count(a, v) + count(b, v)
6   {
7   }
8 }
9 }
```

Copying the formula within the scope of our postcondition's quantifier is appropriate here, and sufficient to prove the claim. Although it's fairly straightforward to see in this example, we can test it by inserting an `assume false` statement in the body of the loop.

We must now decide how to proceed with the main proof of the distributive property. Because the objects that we refer to in the lemma postcondition are sequences, it is worth taking a moment to consider the set of possible values that these objects can take. Sequences can be defined inductively, with the empty sequence `[]` serving as the base case. The definition is as follows:

$$\text{seq } < T > ::= [] | [T] + \text{seq } < T >$$

This suggests that we can use structural induction to prove our claim. We will induce on the structure of `a`, starting with the base case of `a = []`:

```

1 lemma count_distrib2(a: seq<int>, b: seq<int>)
2   ensures forall v :: count(a + b, v) == count(a, v) + count(b, v)
3 {
4   forall(v: int)
5     ensures count(a + b, v) == count(a, v) + count(b, v)
6   {
7     if(a == []) {
8       calc == {
9         count(a + b, v);
10        { assert a + b == b; }
11        count(b, v);
12        0 + count(b, v);
13        { assert count(a, v) == 0; }
14        count(a, v) + count(b, v);
15      }
16    }
17  }
18 }
```

We use an equality calculation to demonstrate the claim, and proceed by using the fact that `a`'s emptiness makes concatenation with `b` the identity. The last few steps of this calculation are not necessary for Dafny to accept the claim in this case, but are provided here to make the proof explicit.

Continuing on with the inductive case, we need to separate `a` into its constituent parts according. Having done so, we can recursively invoke the lemma on the smaller constituents to prove the claim.

```

1 lemma count_distrib2(a: seq<int>, b: seq<int>)
2   ensures forall v :: count(a + b, v) == count(a, v) + count(b, v)
3 {
4   forall(v: int)
5     ensures count(a + b, v) == count(a, v) + count(b, v)
6   {
7     if(a == []) {
8       calc == {
9         count(a + b, v);
10        { assert a + b == b; }
11        count(b, v);
12        0 + count(b, v);
13        { assert count(a, v) == 0; }
14        count(a, v) + count(b, v);
15      }
16    } else {
17      calc == {
18        count(a + b, v);
19        { assert a + b == [a[0]] + (a[1..] + b); }
20        count([a[0]] + (a[1..] + b), v);
21        (if a[0] == v then 1 + count(a[1..] + b, v) else count(a[1..] + b, v));
22        { count_dist(a[1..], b); }
23        (if a[0] == v then 1 + count(a[1..], v) + count(b, v)
24        else count(a[1..], v) + count(b, v));
25        count(a, v) + count(b, v);
26      }
27    }
28  }
29 }

```

Notice that the third line comes directly from the definition of `count`. As before, several of the final steps are not necessary for Dafny to accept our proof, but are provided here for documentation.

This completes the proof that `count` distributes across concatenation, and thus the proof of `BubbleSort`. The complete listing of `BubbleSort`'s code is given in the next figure.

```

1 method BubbleSort(array<int> a)
2   modifies a
3   requires a != null
4   ensures sorted(a, 0, a.Length-1) && perm(old(a[..]), a[..])
5 {
6   var i := a.Length - 1;
7   while(i > 0)
8     invariant i < 0 ==> a.Length == 0
9     invariant sorted(a, i, a.Length-1)
10    invariant part(a, i)
11    invariant perm(old(a[..]), a[..])
12  {
13    var j := 0;
14    while(j < i)
15      invariant 0 <= j <= i
16      invariant sorted(a, i, a.Length-1)
17      invariant part(a, i)
18      invariant forall k :: 0 <= k <= j ==> a[k] <= a[j]
19      invariant perm(old(a[..]), a[..])
20    {
21      if(a[j] > a[j+1]) {
22        ghost var a' := a[..];
23        a[j], a[j+1] := a[j+1], a[j];
24        ghost var v1, v2 := a[j], a[j+1];
25        assert a[..] == a[..j] + [v1, v2] + a[j+2..];
26        assert a' == a[..j] + [v2, v1] + a[j+2..];
27        count_distrib3(a[..j], [v1, v2], a[j+2..]);
28        count_distrib3(a[..j], [v2, v1], a[j+2..]);
29        assert forall v :: count(a[..], v) ==
30          count(a[..j], v) + count([v1, v2], v) + count(a[j+2..], v);
31        assert perm([v1, v2], [v2, v1]);
32        assert perm(a[..], a');
33      }
34      j := j + 1;
35    }
36    i := i - 1;
37  }
38 }
```

Appendix: Using Z3 to check verification conditions

When reasoning about the validity of verification conditions on basic paths, it can be tedious to look for counterexamples or attempt validity proofs by hand. Here, we'll see how to use Z3, the decision procedure utilized by Dafny, directly to assist our reasoning.

Running Z3 You will find the binaries for Z3 in the Dafny distribution you downloaded earlier (assuming that you did not opt for the Visual Studio extension). You can also use Z3 via the web interface at <http://rise4fun.com/z3>.

To run Z3 from the command line, invoke its binary (`z3`) with a single parameter corresponding to a file containing the SMT problem that you would like to solve. Z3 accepts several input formats, and the examples that we'll see in this section are in the SMT2 format that you can read more about at <http://smtlib.cs.uiowa.edu/>.

Encoding VCs To demonstrate the use of Z3 to encode verification conditions, we'll look at all of the partial correctness conditions in our completed `BubbleSort` method. However, to make the examples relatively simpler and easier to follow, we'll look at the version of `BubbleSort` that doesn't prove `permutation(old(a[..]), a[..])`.

The first basic path proceeds from the method entry point to the beginning of the outer loop.

```

1 requires a != null
2 i := a.Length - 1
3 invariant -1 <= i < a.Length
4 invariant sorted(a, i, a.Length-1)
5 invariant partitioned(a, i)

```

To encode this in Z3, we'll first declare symbols corresponding to our syntactic entities. Starting with the variables `a`, `i`, we have the following.

```

1 (declare-const a (Array Int Int))
2 (declare-const i Int)

```

We declare variables using `declare-const`, which accepts a type argument. `a` is defined to be an element from the domain of Z3's theory of arrays taking integer indices to integer-valued elements, and `i` an integer. The `Array` types used by Z3 have the same axioms as the theory of arrays that we studied earlier in the semester.

```

1 (declare-const a (Array Int Int))
2 (declare-const i Int)

```

Now we need to define the predicates `sorted` and `partitioned`. Recall that predicates are functions ranging over Booleans, so we use `declare-fun` to declare them, and then assert their semantics as axioms (i.e., universally-quantified logical sentences). We also need some way of referring to the length of our arrays, as the theory of arrays does not account for it, but it is mentioned in our formulas. To do so, we'll introduce an uninterpreted function `length` ranging over arrays, and assert an axiom that requires the length of any array to be at least zero.

```

1 (declare-const a (Array Int Int))
2 (declare-const i Int)
3
4 (declare-fun sorted ((Array Int Int) Int Int) Bool)
5 (declare-fun partitioned ((Array Int Int) (Int)) Bool)
6 (declare-fun length ((Array Int Int)) Int)
7
8 (assert (forall ((a (Array Int Int)))
9   (<= 0 (length a))))
10
11 (assert (forall ((a (Array Int Int)) (l Int) (u Int)))
12   (iff
13     (sorted a l u)
14     (forall ((i Int) (j Int))
15       (=>
16         (and (<= 0 i) (<= l i) (<= i j) (<= j u) (< u (length a)))
17         (<= (select a i) (select a j))))))
18
19 (assert (forall ((a (Array Int Int)) (i Int)))
20   (iff
21     (partitioned a i)
22     (forall ((k1 Int) (k2 Int))
23       (=>
24         (and (<= 0 k1) (<= k1 i) (< i k2) (< k2 (length a)))
25         (<= (select a k1) (select a k2)))))))

```

Before moving on, notice a few things about this input format. Z3 is an SMT solver that is fundamentally based on the “lazy” propositional encoding algorithm that we studied earlier in the semester. Having algorithmic roots in DPLL, you should expect that it accepts CNF formulas over some first-order background theory as input, and returns `sat`, `unsat`, or `unknown` as its result (`unknown` is necessary as the background theory may be undecidable, or the user may want to place a hard limit on the amount of time the solver is given). The lines beginning with `assert` correspond to the CNF clauses that we’re giving as our input, and as such, they take a Boolean-valued term (i.e., formula) as their only argument. When we ask Z3 to solve our instance, it will attempt to decide whether each asserted formula can be satisfied conjunctively. However, it is worth pointing out that Z3 is somewhat flexible with respect to the contents of each clause. It does not strictly require the input to be a CNF, so the clauses don’t necessarily need to be disjunctive; they can contain arbitrary logical connectives and theory symbols, as long as the final result is a Boolean type.

The second thing to notice is the fact that formulas and terms are given in **prefix notation**. That is, rather than writing:

$F_1 \text{ and } F_2 \text{ and } \dots \text{ and } F_n$

Z3 expects its input to be given as:

$(\text{and } F_1 \ F_2 \ \dots \ F_n)$

If you haven’t seen this notation before, it may take a little getting used to. However, the basic idea is straightforward: rather than placing the operator between its arguments, it is placed at the beginning of a parenthetical list, and the arguments follow in-order. In general, when a binary operator is associative, Z3 will allow you to list an arbitrary number of arguments rather than nesting successive instances of the operator within the arguments. This is why we were able to use

one instance of `and` in the example above, rather than:

$$(\text{and } F_1 \text{ (and } F_2 \text{ (and } \dots \text{ } F_n))$$

Getting back to our example, recall that we'd like to check the verification condition for the basic path:

```

1 requires a != null
2 i := a.Length - 1
3 invariant -1 <= i < a.Length
4 invariant sorted(a, i, a.Length-1)
5 invariant partitioned(a, i)

```

We've encoded our variables and all of the necessary predicates by making universally-quantified assertions that reflect their semantics. For example, recalling the predicate that we wrote in Dafny for `sorted`:

```

1 predicate sorted(a: array<int>, l: int, u: int)
2   reads a
3   requires a != null
4 {
5   forall i, j :: 0 <= l <= i <= j <= u < a.Length ==> a[i] <= a[j]
6 }

```

We encoded this in Z3 using the following assertion:

```

1 (assert (forall ((a (Array Int Int)) (l Int) (u Int))
2   (iff
3     (sorted a l u)
4     (forall ((i Int) (j Int))
5       (=>
6         (and (<= 0 l) (<= l i) (<= i j) (<= j u) (< u (length a)))
7         (<= (select a i) (select a j)))))))

```

This translation is fairly straightforward, with the only noteworthy artifact being the fact that we universally-quantified the arguments to the predicate when encoding it in Z3. This reflects the fact that we want this definition to hold for any arguments that we might apply the predicate to. We did not account for the precondition in the Dafny predicate, as we have not encoded the possibility that the array reference might be null into our SMT satisfiability problem. When Dafny produces verification conditions, it certainly does account for this, but because we only reason about sortedness in situations where the array reference is non-null (recall the precondition on `BubbleSort`), we don't need to address this additional complexity at the moment.

Inspecting the above basic path, we'll replace the precondition `a != null` with `true`. This gives us the verification condition:

$$\begin{aligned} & \text{true} \rightarrow \text{wlp}(i := |a| - 1, -1 \leq i < |a|) \wedge \text{sorted}(a, i, |a| - 1) \wedge \text{partitioned}(a, i)) \\ \Leftrightarrow & -1 \leq |a| - 1 < |a| \wedge \text{sorted}(a, |a| - 1, |a| - 1) \wedge \text{partitioned}(a, |a| - 1) \end{aligned}$$

We'd like to use Z3, a satisfiability solver, to check the validity of this formula. Recall that we can exploit the connection between satisfiability and validity to accomplish this. Namely, we will check

to determine whether the negation of our formula is satisfiable:

$$\neg(-1 \leq |a| - 1 < |a| \wedge \text{sorted}(a, |a| - 1, |a| - 1) \wedge \text{partitioned}(a, |a| - 1))$$

Continuing from the Z3 instance we left off with before, we simply add another assertion corresponding to this condition:

```

1 (declare-const a (Array Int Int))
2 (declare-const i Int)
3
4 (declare-fun sorted ((Array Int Int) Int Int) Bool)
5 (declare-fun partitioned ((Array Int Int) (Int)) Bool)
6 (declare-fun length ((Array Int Int)) Int)
7
8 (assert (forall ((a (Array Int Int)))
9   (<= 0 (length a))))
10
11 (assert (forall ((a (Array Int Int)) (l Int) (u Int))
12   (iff
13     (sorted a l u)
14     (forall ((i Int) (j Int))
15       (=>
16         (and (<= 0 l) (<= l i) (<= i j) (<= j u) (< u (length a)))
17         (<= (select a i) (select a j))))))
18
19 (assert (forall ((a (Array Int Int)) (i Int))
20   (iff
21     (partitioned a i)
22     (forall ((k1 Int) (k2 Int))
23       (=>
24         (and (<= 0 k1) (<= k1 i) (< i k2) (< k2 (length a)))
25         (<= (select a k1) (select a k2))))))
26
27 (assert (not (and
28   (<= -1 (- (length a) 1))
29   (< (- (length a) 1) (length a))
30   (sorted a (- (length a) 1) (- (length a) 1))
31   (partitioned a (- (length a) 1))))))
32
33 (check-sat)

```

The last line instructs Z3 to decide the satisfiability of our assertions. If we save our instance to `vc.smt2`, then we invoke Z3 as follows:

```
1 $ z3 vc.smt2
```

It should return promptly with:

```
1 $ z3 vc.smt2
2 unsat
```

Which tells us that our verification condition is valid: its negation is unsatisfiable.

Now let's take a look at an invalid VC. Recall from earlier that the basic path:

```

1 invariant 0 < i < a.Length && 0 <= j <= i &&
2           sorted(a, i, a.Length-1) &&
3           partitioned(a, i)
4 assume i <= j;
5 i := i - 1;
6 invariant -1 <= i < a.Length &&
7           sorted(a, i, a.Length-1) &&
8           partitioned(a, i)

```

With corresponding VC:

$$\begin{aligned}
0 < i < a.Length \wedge i = j \wedge \text{sorted}(a, i, a.Length - 1) \wedge \text{partitioned}(a, i) \\
\rightarrow \text{sorted}(a, i - 1, a.Length - 1) \wedge \text{partitioned}(a, i - 1)
\end{aligned}$$

is invalid. We can use Z3 to check this by deciding the satisfiability of:

$$\begin{aligned}
0 < i < a.Length \wedge i = j \wedge \text{sorted}(a, i, a.Length - 1) \wedge \text{partitioned}(a, i) \\
\wedge \neg(\text{sorted}(a, i - 1, a.Length - 1) \wedge \text{partitioned}(a, i - 1))
\end{aligned}$$

Translating the problem into Z3's input format, assuming that we've already defined the relevant variables and axioms:

```

1 (assert (< 0 i))
2 (assert (< i (length a)))
3 (assert (= i j))
4 (assert (sorted a i (- (length a) 1)))
5 (assert (partitioned a i))
6 (assert (not (and (sorted a (- i 1) (- (length a) 1)) (partitioned a (- i 1))))))
7
8 (check-sat)

```

It appears that Z3 hangs on this instance. Adding a 10-second timeout to the solver's configuration at the top of the input file:

```

1 (set-option :timeout 10000)

```

We see that the solver returns `unknown`. Because program verification is one of the primary applications that Z3 was developed to enable, it has been engineered to identify `unsat` instances very quickly, as these correspond to valid verification conditions. However, for invalid conditions, which correspond to `sat` instances, it is oftentimes not as effective. When Dafny's verifier sees an `unknown` answer from Z3, it assumes that the VC is invalid and returns an error message. The program locations that it identifies in its error message are generated by obtaining a “reason” for the unknown answer, which Z3 is usually able to return in the form of a subset of assertions on which it had difficulty reasoning.

In our case, the difficulty is due to the quantifiers that we used to define `sorted` and `partitioned`; as we've discussed previously, quantified formulas are especially difficult for automated reasoning techniques to handle. Nonetheless, Dafny and Z3 do remarkably well handling such instances in a large number of cases. A significant part of the challenge in developing a verifier lies in finding a good SMT encoding for the generated verification conditions. The exact encoding that Dafny

uses differs from the one we used here in numerous ways, which is why the verifier did not respond with a timeout error message when we used the wrong loop invariant corresponding to this basic path. You can examine the Z3 verification conditions generated by Dafny by invoking it with the `proverLog` argument:

```
1 $ mono Dafny.exe prog.dfy /proverLog:@PROC@-vc.smt2
```

This will cause the compiler to generate a file for each verification condition, where `@PROC@` expands to the procedure that the condition corresponds to. As an exercise, examine the verification conditions generated for `BubbleSort`, and try to make the Z3 verification conditions we wrote above correspond more closely to those generated by Dafny. Understanding, at least in some level of detail, how the verifier works behind the scenes will further remove the mystery from some of the error messages that you encounter when proving your code.