

15-122: Principles of Imperative Computation

Recitation 4

Josh Zimmerman

Overflow

Overflow can cause potentially nasty bugs, so it's often advisable to check that your operations will not cause overflow before performing them.

For example, let's add $7 + 1$ using two's complement arithmetic:

$$\begin{array}{r} 0111 \\ + 0001 \\ \hline 1000 \end{array}$$

But $1000 = -1 * 2^3 + 0 * 2^2 + 0 * 2^1 + 0 * 2^0 = -8$. So $1 + 7 = -8$ in four-bit two's complement arithmetic, which is not what we wanted to happen. (Similarly, $-8 - 1$ does strange things—you should play with that and see what happens.)

There's some more weird behavior when we negate the minimum int (In a 4-bit system, that's 1000.) Recall that we define the negation of an integer x as $-x = \sim x + 1$. We then have the following:

$$-x = (\sim x + 1) = (\sim 1000 + 1) = (0111 + 1) = (1000) = x$$

This is an important oddity to be aware of: **if you negate the minimum integer, you get the minimum integer back.**

Overflow can occur with unsigned integers as well. Let's add $15+1$ assuming that both are unsigned integers:

$$\begin{array}{r} 1111 \\ + 0001 \\ \hline 0000 \end{array}$$

Overflow is not a theoretical problem. Here are some real-world examples where integer overflow has led to bugs:

- (a) In the original Pac-Man game, the level counter was an unsigned 8-bit integer and could only store values from 0-255 (in this case, the integers were unsigned types, not signed types that we used in the examples above). Reaching level 256 would cause the counter to overflow, leading to a 'kill screen'
- (b) In the original Donkey Kong game, the amount of time to complete the level was directly dependent on the level value and was calculated via a formula $100 \times (10 \times (\text{level} + 4))$. Reaching level 22 would yield the calculation (100×260) ; because of integer overflow, this would be evaluated as (100×4) , leading to an impossibly short amount of time to complete the level.

So, clearly, overflow is bad. However, it is also a non-trivial problem to catch. Let's discuss the process of implementing a function `int safe_add(int a, int b)` that will use a precondition to ensure that the calculation $a + b$ does not overflow. Let's discuss the possible cases:

- (a) `a > 0 && b > 0`: While it tempting to calculate the sum $a+b$ and check for overflow via the expression $a+b > 0$, this solution does not generalize well (i.e. actually perform the computation that we want to check for overflow) although it will work in C0. A better method is to check if $a \leq \text{int_max}() - b$, as $\text{int_max}() - b$ will not overflow (as $b > 0$).
- (b) `a < 0 && b < 0`: Likewise, the expression $a + b < 0$ is not a good approach. We instead check that $a \geq \text{int_min}() - b$, as $\text{int_min}() - b$ will not overflow (as $b < 0$).
- (c) `a <= 0 && b >= 0`: In this case, $a+b$ will not overflow (it will be between `int_min()` and `int_max()`).
- (d) `a >= 0 && b <= 0`: In this case, $a+b$ will not overflow (it will be between `int_min()` and `int_max()`).

Checkpoint 0

Based on the discussion above, what is the pre-condition for the function `safe_add`?

Checkpoint 1

Try to develop a precondition for the function `safe_mult`. For the sake of simplicity, assume that $a > 0$ and $b > 0$ (you can try the other cases as an exercise later).

Here's a hint: if `a > 0 && b > 0`, is it true that if $a * b > 0$ that overflow did not happen? This should also make it apparent why we adopted the strategy used in `safe_add`.

Arrays

Arrays are stored as *addresses* in C0. We'll talk more about pointers later in the course, but for now it's sufficient to know that a variable representing an array stores the address of the array (i.e. where it is located in memory) and that indexing into the array allows us to access the value of the element at that index.

Like other variables, array variables can be assigned to other array variables. However, assigning array variables just changes the *reference* that one array variable stores.

```
1 void addOne(int[] arr, int size)
2 //@requires size == \length(arr);
3 {
4   // It's important to make sure that all array accesses are in bounds.
5   // Ideally, we should be able to prove that they are.
6   for (int i = 0; i < size; i++)
7     //@loop_invariant 0 <= i && i <= size;
8     {
9       //arr[i] gets the value at index i in array arr; we can assign
10      //arr[i] to another value or we can assign a variable to store
11      //the value at arr[i]
12      arr[i] += 1;
13    }
14   // When we take (loop invariant) and not(loop exit condition),
15   // we can easily conclude that i == size, which is often useful
16 }
17
18 int main()
19 {
20   int array_size = 2;
21
22   //allocate an int array of array_size elements, with all
23   //elements initialized to 0. Only way to create array in c0
24   int[] a = alloc_array(int, array_size);
25
26   //create an alias to a; any changes made to elements in a
27   // will be reflected in b
28   int[] b = a;
29
30   //@assert a[0] == 0;
31   //@assert b[0] == 0;
32
33   // We must pass the size of the array to the function because otherwise
34   // the function has no way of verifying that all array accesses are
35   // in bounds.
36   addOne(a, array_size);
37
38   // The function call changes the area in memory that a and b
39   // refer to, so it changes both a and b when we access them.
40   assert (a[0] == 1 && a[1] == 1);
41   assert (b[0] == 1 && b[1] == 1);
42
43   return 0;
44 }
```

Let's try a variation of `addOne`, but with a twist. We will use the same main function as given above.

```
1 void addOne(int[] arr, int size)
2 //@requires size == \length(arr);
3 {
4   arr = alloc_array(int, size);
5   for (int i = 0; i < size; i++)
6     //@loop_invariant 0 <= i && i <= size;
7     {
8       arr[i] += 1;
9     }
10 }
```

Although we assigned `arr` to a new array and added one to each element in the new array, the changes are not reflected back in the calling function! The reason for this is that C0 is a *call by value* language. This means that when we pass arguments to a function, *copies* of the arguments are passed. Changes made to function's arguments will *not* be reflected back in the calling function.

Checkpoint 2

Fill in the blanks in the code to show that there are no out of bounds array accesses. You are given the function `slow_fib` that calculates Fibonacci recursively (so it can be used as a reference function)

```
1 int slow_fib(int n)
2 //@requires n >= 0;
3 ;
4
5 :
6
7 1 int fib(int n)
8 2 //@requires _____;
9 3 //@ensures \result == slow_fib(n);
10 4 {
11 5   int[] F = alloc_array(int, n);
12 6   if (n > __) {
13 7     F[0] = 0;
14 8   }
15 9   else {
16 10    return 0;
17 11   }
18 12   if (n > __) {
19 13     F[1] = 1;
20 14   }
21 15   else {
22 16    return 1;
23 17   }
24 18   for (int i = 2; i < n; i++)
25 19     //@loop_invariant _____;
26 20     /*@loop_invariant F[i - 1] == slow_fib(i - 1)
27 21        && F[i - 2] == slow_fib(i - 2); @*/
28 22     {
29 23       F[i] = F[i - 1] + F[i - 2];
30 24     }
31 25   return F[n - 1] + F[n - 2];
32 26 }
```

Checkpoint 3

Are the invariants strong enough to prove the postcondition? If so, then use the loop invariant to prove the postcondition. Use the outline below to think out your solution

Array access

Showing that the postcondition holds.

Termination