

Homework 7

Due Thursday, 3/14/2019

Please turn in your homework solutions online at <https://submit.cs.pomona.edu/2019sp/cs131>. There is no programming, so only a pdf is required.

1. (15 points) **Ada Parameter Modes**

The Ada programming language permits parameters to be labeled as `in`, `out`, or `in out`, as in the following procedure definitions, where `T` is some type:

```
procedure test1(in x: T) is begin ... end
procedure test2(out x: T) is begin ... end
procedure test3(in out x: T) is begin ... end
```

The modifiers, or modes, have the following meaning:

- `in`: The value of the parameter `x` may not be changed inside the procedure. (That is, it is a compile-time error if there is code assigning a value to `x` in the body of the procedure.) If we call `test1(y)`, the value of `y` is the same before and after the call.
- `out`: The parameter `x` can be written to, but it cannot be read. If we call `test2(y)`, the value of `y` after the call is the last value written to `x` in the procedure.
- `in out`: The parameter `x` can be both read and written, and the value of `y` after a call to `test3(y)` is the last value written to `x` in the procedure.

The language definition does not specify how each mode should be implemented, and the compiler may use any appropriate parameter passing mechanism to implement them.

- (a) Which parameter passing mechanism could be used by the compiler to implement `test1`, `test2`, and `test3` with the desired semantics of `in`, `out`, or `in-out`? The choices are pass-by-reference, pass-by-value, and pass-by-value-result (as described in problem 7.6 of the text). If more than one is possible, describe the advantages/disadvantages of each.
- (b) Consider the following procedure that takes two parameters. Does the following program print the same value for all strategies you outlined for `in out` parameters above?

```
procedure incTwo(in out x:integer, in out y:integer) is
begin
  x := x + 1;
  y := y + 1;
end

procedure main() is
w : integer = 3;
begin
  incTwo(w,w);
  print w;
end
```

- (c) Discuss the advantages and disadvantages of permitting the compiler such flexibility in how it implements parameter modes. Consider issues of time and space complexity.

2. (15 points) **Type Inference to Detect Race Conditions**

The following problem refers to the use of synchronized blocks in Java to solve concurrency issues when using multiple threads. I've tried to include enough info about synchronized blocks in Java that you won't need to read up on them. However, if you feel you would like to read up more about them, I recommend the following two documents:

- <http://web.mit.edu/6.031/www/fa17/classes/21-locks/> is a reading for an MIT course. It has the advantage of being relatively short.
- <http://www.cs.pomona.edu/classes/cs062-f2017/handouts/CS62sophomoricParallelismAndConcurrency.pdf>, is a link to "A Sophomoric introduction to parallelism and concurrency". This was used in CS 62. Chapter 6 is the chapter to read.

We will cover more on concurrency toward the end of the semester.

The general techniques from our type inference algorithm can be used to examine other program properties as well. In this question, we look at a non-standard type inference algorithm to determine whether a concurrent program contains race conditions. Race conditions occur when two threads access the same variable at the same time. Such situations lead to non-deterministic behavior, and these bugs are very difficult to track down since they may not appear every time the program is executed. For example, consider the following program, which has two threads running in parallel:

<p>Thread 1:</p> <pre>t1 := !hits; hits := !t1 + 1;</pre>	<p>Thread 2:</p> <pre>t2 := !hits; hits := !t2 + 1;</pre>
---	---

In the above code, “:=” is used for assignment, “!” is used to extract the value of a variable. Thus $x := !x + 1$ increases the value of variable x by 1.

Since the threads are running in parallel, the individual statements of Thread 1 and Thread 2 can be interleaved in many different ways, depending on exactly how quickly each thread is allowed to execute. For example, the two statements from Thread 1 could be executed before the two statements from Thread 2, giving us the following execution trace:

$$\text{hits} = 0 \xrightarrow{t1 := !\text{hits}} \text{hits} = 0 \xrightarrow{\text{hits} := !t1 + 1} \text{hits} = 1 \xrightarrow{t2 := !\text{hits}} \text{hits} = 1 \xrightarrow{\text{hits} := !t2 + 1} \text{hits} = 2$$

After all four statements execute, the `hits` counter is updated from zero to 2, as expected. Another possible interleaving is the following:

$$\text{hits} = 0 \xrightarrow{t2 := !\text{hits}} \text{hits} = 0 \xrightarrow{\text{hits} := !t2 + 1} \text{hits} = 1 \xrightarrow{t1 := !\text{hits}} \text{hits} = 1 \xrightarrow{\text{hits} := !t1 + 1} \text{hits} = 2$$

This again adds 2 to `hits` in the end. However, look at the following trace:

$$\text{hits} = 0 \xrightarrow{t1 := !\text{hits}} \text{hits} = 0 \xrightarrow{t2 := !\text{hits}} \text{hits} = 0 \xrightarrow{\text{hits} := !t1 + 1} \text{hits} = 1 \xrightarrow{\text{hits} := !t2 + 1} \text{hits} = 1$$

This time, something bad happened. Although both threads updated `hits`, the final value is only 1. This is a race condition: the exact interleaving of statements from the two threads affected the

final result. Clearly, race conditions should be prevented since it makes ensuring the correctness of programs very difficult. One way to avoid many race conditions is to protect shared variables with mutual exclusion locks. A lock is an entity that can be held by only one thread at a time. If a thread tries to acquire a lock while another thread is holding it, the thread will block and wait until the other thread has released the lock. The blocked thread may acquire it and continue at that point. The program above can be written to use lock `l` as follows:

```

Thread 1:
  synchronized(l) {
    t1 := !hits;
    hits := !t1 + 1;
  }

Thread 2:
  synchronized(l) {
    t2 := !hits;
    hits := !t2 + 1;
  }

```

The statement “`synchronized(l) { s }`” acquires lock `l`, executes `s`, and then releases lock `l`. There are only two possible interleavings for the program now:

$$\text{hits} = 0 \xrightarrow{t1 := !\text{hits}} \text{hits} = 0 \xrightarrow{\text{hits} := !t1 + 1} \text{hits} = 1 \xrightarrow{t2 := !\text{hits}} \text{hits} = 1 \xrightarrow{\text{hits} := !t2 + 1} \text{hits} = 2$$

and

$$\text{hits} = 0 \xrightarrow{t2 := !\text{hits}} \text{hits} = 0 \xrightarrow{\text{hits} := !t2 + 1} \text{hits} = 1 \xrightarrow{t1 := !\text{hits}} \text{hits} = 1 \xrightarrow{\text{hits} := !t1 + 1} \text{hits} = 2$$

All others are ruled out because only one thread can hold lock `l` at a time. Note that while we use assignable variables inside the synchronized blocks, the names we use for locks are constant. For example, the name `l` in the example program above always refers to the same lock.

Our analysis will check to make sure that locks are used to guard shared variables correctly. In particular, our analysis checks the following property for a program `P`:

For any variable `y` used in `P`, there exists some lock `l` that is held by the current thread every time `y` is accessed.

In other words, our analysis will verify that every access to a variable `y` will occur inside the synchronized statement for some lock `l`. Checking this property usually uncovers many race conditions.

Let’s start with a simple program containing only one thread:

```

Thread 1:
  synchronized (m) {
    a := 3;
  }

```


For this program, our analysis should infer that lock `m` protects variable `a`.

As with standard type inference, we proceed by labeling nodes in the parse tree, generating constraints, and solving them.

Step 1: Label each node in the parse tree for the program with a variable. This variable represents the set of locks held by the thread every time execution reaches the statement represented by that


node of the tree. Note that these variables keep track of sets of locks names, and NOT types, in this analysis.

Here is the labeled parse tree for the example:



Step 2: Generate the constraints using the following four rules:

- (a) If S is the variable on the root of the tree, then $S = \emptyset$.
- (b) For any subtree matching the form




we add two constraints:

$$T = R$$

$$S = R \cup \{1\}$$

- (c) For any subtree matching the form



where ANY matches any node other than a sync node, we add two constraints:

$$T = R$$

$$S = R$$

(d) To determine lock_y , the lock guarding variable y , add the constraint

$$\text{lock}_y \in S$$

for each node $y : S$ or $!y : S$ in the tree. In other words, require that lock_y be in the set of locks held at each location y is accessed.

Here are the constraints generated for the example program:

$$\begin{aligned} S1 &= \emptyset && \text{(rule 2a)} \\ S2 &= S1 && \text{(rule 2b)} \\ S3 &= S1 \cup \{m\} && \text{(rule 2b)} \\ S4 &= S3 && \text{(rule 2c)} \\ S5 &= S3 && \text{(rule 2c)} \\ \text{lock}_a &\in S4 && \text{(rule 2d)} \end{aligned}$$

Step 3: Solve the constraints to determine the set of locks held at each program point and which locks guard the variables:


$$\begin{aligned} S2 = S1 &= \emptyset \\ S3 = S4 = S5 &= \{m\} \\ \text{lock}_a &\in \{m\} \end{aligned}$$

Clearly, lock_a is m in this case, exactly as we expected.

You will now explore some aspects of this analysis:

(a) Here is another program and corresponding parse tree:

```
Thread 1:
  synchronized (l) {
    synchronized (m) {
      a := 4;
      b := !a;
    }
    b := 33;
  }
```



Compute $lock_a$ and $lock_b$ using the algorithm above. Explain why the result of your algorithm makes sense.

- (b) Let's go back to the original example, but change Thread 2 to use a different lock:

```

Thread 1:
synchronized(l) {
    t1 := !hits;
    hits := !t1 + 1;
}

Thread 2:
synchronized(m) {
    t2 := !hits;
    hits := !t2 + 1;
}
    
```

Compute $lock_{t1}$, $lock_{t2}$, and $lock_{hits}$ using the algorithm above. Since there are two threads in the program, you should create two parse trees, one for each thread. Explain the result of your algorithm.

- (c) Suppose that we allow assignments to lock variables. For example, in the following program, l and m are references to locks, and we can change the locks to which those names refer with an assignment statement:

```

Thread 1:
synchronized(!l) {
    a := !a + 1;
}
m := !l;
synchronized(!m) {
    a := !b + 1;
    b := !a;
}

Thread 2:
synchronized(!m) {
    x := !b + 3;
    b := 11 + x;
}
    
```

Describe any problems that arise due to assignments to lock variables, and what the implications for the analysis are. You do not have to show the constraints from this example or change the analysis to handle mutable lock variables. A coherent discussion of the issues is sufficient. Thinking about what the algorithm would compute for `locka`, `lockb`, and `lockx` may be useful, however.