Assignment 2: Safe to the Last Command
15-316 Software Foundations of Security and Privacy

Due: **11:59pm**, Thursday 2/15/18
Total Points: 50

1. **Verification conditions (15 points).** Recall that a verification condition for a program path $\alpha_1; \ldots; \alpha_n$ is a formula of arithmetic whose validity implies the validity of the DL formula $[\alpha_1; \ldots; \alpha_n]P$, for some P. Consider the following program that we will denote α:

```
i := 0;
while(i < k) {
  if(i < 0) { i := k; }
  if(0 <= Mem(i)) { x := x + Mem(i); }
  i := i + 1;
}
```

   Part 1 (5 points). Assume that k is a constant that we fix in advance. How many verification conditions are required to check $[\alpha]0 \leqslant x$, as a function of k? Explain the rationale behind your answer.

   Part 2 (10 points). List the verification conditions for $k = 2$. Note that it is not necessary to conduct a sequent calculus proof for this problem, and you will receive full credit for listing out the paths and their corresponding verification conditions. It is also not necessary to say whether the VCs are valid or not.

2. **Unfinished business (10 points).** In lecture 7, we discussed two cases of the structural induction used to prove the security of SFI. Complete the inductive case for conditional commands. That is, assuming that Equation 1 is valid for α and β whenever $0 \leqslant s_l \leqslant (x \,\&\, s_h) \mid s_l \leqslant b_h < U$:

$$\forall i. \neg(s_l \leqslant i \leqslant s_h) \wedge \texttt{Mem}(i) = v_i \rightarrow [\alpha]\texttt{Mem}(i) = v_i \tag{1}$$

   Prove that it is also valid for $\texttt{if}(Q)\,\alpha\,\texttt{else}\,\beta$.

3. **Leaky sandbox (20 points).** Consider the following language, which resembles a simplified assembly language.

$\texttt{and}(x, y)$	Take the bitwise-and of variables x and y, store the result in x
$\texttt{or}(x, y)$	Take the bitwise-or of variables x and y, store the result in x
$x := y$	Copy the value stored in y to x
$x := \texttt{Mem}(y)$	Read the memory at address stored in variable y, save result in x
$\texttt{Mem}(x) := y$	Store the value in y at the address pointed to by x
$\texttt{if}(Q)\,\texttt{jump}\,x$	If Q is true in the current state, jump to the instruction pointed to by x

   Programs in this language are sequences of instructions indexed on integers 0 to n, and we refer to the instruction at index i of program Π with the notation Π_i. Note that there are no expressions in this program. Results of operations are stored in variables, and can be moved into memory when necessary. Think of variables as acting like registers, so to implement the computation $w := (x \,\&\, y) \mid z$ from our language in lecture we would write the program:

$$\begin{array}{ll} 1: & \texttt{and}(x, y) \\ 2: & \texttt{or}(x, z) \\ 3: & w := x \end{array}$$

   It is *not* valid to write $w := \texttt{or}(\texttt{and}(x, y), z)$ because neither $\texttt{or}(\texttt{and}(x, y), z)$ or $\texttt{and}(x, y)$ is a variable.

Part 1 (10 points). We want to implement a sandboxing policy for this language using software fault isolation. So the proposal is to replace all memory read and write operations as follows. Assume that s_l = `0x15316000` and s_h = `0x15316fff`, so the memory sandbox is contained in the range of addresses `0x15316000` − `0x15316fff`.

$x := \texttt{Mem}(y)$	becomes	$\texttt{and}(y, \texttt{0x15316fff})$ $\texttt{or}(y, \texttt{0x15316000})$ $x := \texttt{Mem}(y)$
$\texttt{Mem}(x) := y$	becomes	$\texttt{and}(x, \texttt{0x15316fff})$ $\texttt{or}(x, \texttt{0x15316000})$ $\texttt{Mem}(x) := y$

Additionally, we want to prevent indirect jumps from leaving a code sandbox restricted to the range of instruction addresses `0x00000a00` − `0x00000aff`. So each indirect jump is rewritten as follows.

$\texttt{if}(Q)\,\texttt{jump}\,x$	becomes	$\texttt{and}(x, \texttt{0x00000aff})$ $\texttt{or}(x, \texttt{0x00000a00})$ $\texttt{if}(Q)\,\texttt{jump}\,x$

Any untrusted code is rewritten using these rules prior to being executed. Unfortunately, we were on a tight deadline and didn't have time to prove that this implementation of SFI is secure. Explain why this instrumentation still allows untrusted code to read and write outside the memory sandbox, and provide an example program in the language that violates the policy.

Part 2 (15 points). Propose an alternative implementation in this language for the policy in Part 1 that is secure. You may assume that the untrusted code is not allowed to modify some variables that you select, but be sure to state any assumptions about what invariants must hold of those variables for your implementation to be secure.

4. ***(Extra Credit)*** **Tough conditions (5 points).** As discussed in lecture 5, symbolic execution can be used to find inputs that drive a program down a particular path. It does this by generating the corresponding path condition, and checking it for satisfiability. If the path condition is satisfiable, then it generates a *model*, or satisfying assignment to the variables. When this assignment is used as the input to the program, it will necessarily end up taking the path used to derive the condition.

   However, this is all contingent on being able to first determine the satisfiability of the path condition and then subsequently generating a satisfying assignment. The decision procedures used to do this are subject to the same laws of computability as any other algorithm, and so there is no guarantee that they will be able to provide answers for every path condition.

   Write a short program for which it is unlikely that a decision procedure will be able to produce satisfying assignments to drive execution down at least one path. Your program is allowed to call outside functions, e.g. $\texttt{Fib}(n)$ to return the nth Fibonacci number, but be sure to describe precisely what any such external function computes, and why it is unlikely that a decision procedure will be able to solve the resulting path conditions.