Final Review Sheet
ENEE 457
Fall 2020

Static Analysis:

Consider the following code snippet on which we would like to perform a taint analysis. Type qualifiers are represented by capital letters: A, B, C, D.

```
1   int printf(A char *fmt, ..);
2
3   int main(B int argc, C char *argv[]) {
4           if (argc < 2 || argc > 2){
5                   printf("enter 1 string only");
6                   return 0; }
7           D char *mystring;
8           if (!strcmp(argv[1], "Hello")){
9                   mystring = argv[1];
10          }else{
11                  mystring = "Goodbye";
12                  printf(mystring);
13          }
14          return 0;
15  }
```

1. Identify all the sources and sinks in the code snippet and determine the corresponding settings for the type qualifiers.

2. List all of the constraints on the type qualifiers.

3. Is there a vulnerability in the above code? Is there a solution for the undetermined type qualifiers that satisfies all the constraints? If there is no vulnerability and no solution, it means that our taint analysis has produced a false positive. How can the taint analysis be modified so that the false positive is removed?

Malware:

1. What is the difference between a virus and a worm?

2. What is the difference between a polymorphic and metamorphic virus?

3. What is a virus signature?

4. What is a crypting service?

Bitcoin

Assume the current Blockchain looks like the following and that the current difficulty level is n.

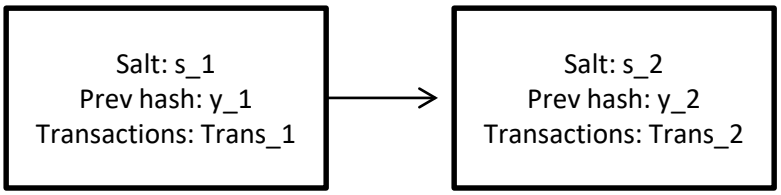

In order to successfully mine the next block (Block 2), a miner needs to find a salt s_3 such that h(s_3, x) = y_3, for x,y_3 of a specific form. What is x in the above example? What is the form of y_3?

Differential Privacy

Show that mechanism M defined below is not differentially private, using the definition of differential privacy and the databases D and D' defined below.

M chooses a value v uniformly at random from $\{-1, 0, 1\}$ and returns the number of UMD students in the database plus v.

Name	UMD Student
Alice	0
Bob	0
Charlie	0
Daniel	1
Edgar	0

Name	UMD Student
Alice	0
Bob	0
Charlie	0
Edgar	0

Dining Cryptographers/MixNets

Using pseudocode, specify how the Dining Cryptographers protocol would work for 4 parties. What happens if two parties collude? Can they combine forces to learn which of the other two parties is broadcasting in a given round? Why or why not?

Password Hashing

Briefly explain what a Rainbow Table is.

In class we saw that adding a “salt" to the hash can prevent attacks via Rainbow Tables. Is this countermeasure still effective if a 256-bit “salt" is chosen, but the same “salt" value is used for each entry in the password table? Explain your answer.

Briefly list some properties of a good password hash and explain why they are desirable for a password hash, but may not be desirable for other settings.

Consider the following graph (Figure 1) showing the memory usage of a hash function H over time when evaluated on a single input x. Let H’ be a hash function that on input $(x_1||x_2)$, outputs $H(x_1)||H(x_2)$. How can we minimize the space-time complexity of the computation of H'? (Recall that the space-time complexity is the maximum amount of space used in any time step multiplied by the number of time steps.) Draw a graph on the back of the page showing the memory usage of H' over time.

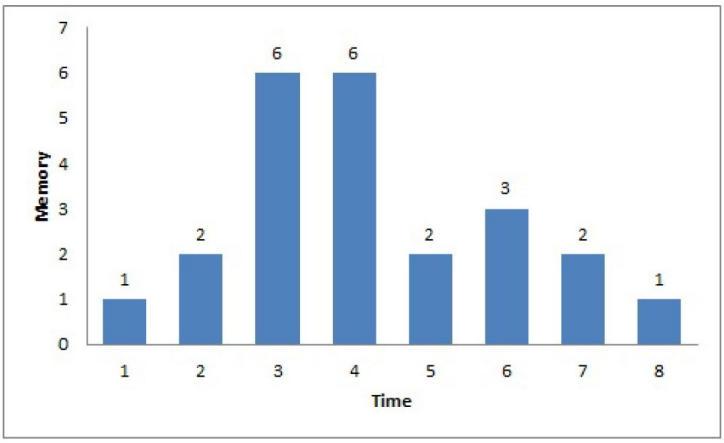

Figure 1: Memory usage of H over time

Network Security

What is TCP/IP?

What is Packet Sniffing/Spoofing and what tools can be used to perform these actions?

What is the TCP 3-way handshake?

Name 3 TCP Header “code bits” (also called flags) and what they do.

What is the significance of the sequence number and acknowledgement number in a TCP header?

What is a SYN Flooding attack?

What is a TCP RST attack?

What is a session hijacking attack?

What is a reverse shell?

Adversarial machine learning

Inputs in the training set are 0/1 vectors of dimension n.

Assume the target function is $x_1 \vee x_2 \vee x_3 \ldots \vee x_n$

Examples occurring in the training set all satisfy $x_2 \neq x_3$.

Give an example of a model that might be output by a machine learning algorithm that would correctly classify all the examples in the training set.

Give an example of an adversarial input that is misclassified by your model given above.