# CSci 5271: Introduction to Computer Security

**Exercise Set 4** **due: Thursday, November 20th, 2014**

**Ground Rules.** You may choose to complete these exercises in a group of up to three students. Each group should turn in **one** copy with the names of all group members on it. You may use any source you can find to help with this assignment but you **must** explicitly reference any source you use besides the lecture notes or textbook. An electronic (plain text or PDF) copy of your solution should be submitted on the course Moodle by 11:55pm on Thursday, November 20th.

**1. Hashing and Signing.** (30 pts) Nearly every digital signature scheme works by first hashing a message to be signed (with a cryptographic hash function) and then performing some operation on the hash—so in essence, we are "signing the hash" and not the message. In particular, if Eve sees Alice's signature on the message $M$ and can find a message $M' \neq M$ so that $H(M) = H(M')$, she can convince people that Alice signed $M'$. This is OK, since a good crypto hash function $H$ will resist finding targeted collisions (second pre-images) like this.

Suppose our signature scheme uses a hash function $H$ with an output length $\ell$ that is sufficient to resist second pre-images but NOT resistant to free collisions (e.g. the hash length is around 100-120 bits). Then it is possible that if Eve can get Alice to sign one of a pair of colliding messages, she can later claim that Alice signed the other.

(a) It might be tempting to think that the risk of such an attack is minimal, since the birthday attack works by hashing random messages until two have the same hash; why would Alice want to sign a random message, and even if she did, why would we care that Eve could claim she signed a different random message? Give a simple explanation why even this attack could be troublesome.

(b) Let's show that in fact, it is worse than that. Suppose that a message is "favorable" if it is something that Alice would sign, for example "I will pay $5 to McDonald's for my lunch." Suppose that a message is "undesirable" if it is something that Alice would not sign, like "I will pay $500,000 to Eve for her lunch." Notice that we can generate 256 different "favorable" messages from the example above, for instance by varying the number of space characters between words between 1 and 2. Extend this idea to show how to generate a pair of messages, one favorable and one undesirable, with the same hash. Your attack should compute about as many hashes as the birthday attack.

(c) Complete the attack: how would Eve use the pair she generates in part (b) to her advantage?

**2. Random numbers with limited entropy.** (30 pts) Alice, Bob, and Carol are employees of a company (in a small island nation) setting up an online casino website based on card games like blackjack. They realize that if users could predict the sequence of pseudorandom numbers used to deal cards, they could win reliably and hurt the company's bottom line. They've found a good cryptographically-strong pseudorandom number generation algorithm to use in the shuffling process, but they're having trouble deciding what to use as the seed when they initialize the generator at the start of each user's session.

(Following the usual good security design principles, they don't want the security of the games to depend on the choice of the pseudorandom generator or the shuffling algorithm being secret;

they might also want to franchise their casino out in the future. But practically speaking, reverse-engineering those algorithms would be a significant effort, so attacks that worked without the attacker needing to do so would be particularly damaging.)

(a) Alice suggests seeding the PRNG with the time: specifically the date and time as returned by the Unix `time` system call, equal to the number of seconds since midnight, January 1st 1970 UTC. Explain why this is a bad idea by describing an easy attack.

(b) Bob suggests seeding the PRNG with the process ID of the login CGI script. Assuming this script runs once for each login, and process ID numbers are assigned sequentially in the range of 2 to 65535, describe an attack against this scheme.

(c) Carol suggests combining Alice and Bob's ideas by taking the time and the PID and XORing them together. But Alice points out a problem with this scheme that involves a user logging in once every second. Explain the details of her attack and why it's a problem.

(d) After the problems with their previous schemes, Alice, Bob, and Carol have called you in as a consultant. Suppose that because of the architecture of the system, the seed is required to be a deterministic function of the time in seconds and the PID. Propose a better combining function that takes these two pieces of information as input and produces a bit string (of any length) than can be used as a seed. Would it help if the function could also take another input that was like a key, fixed per-site but secret? Evaluate the security of your approach.

**3. Firewall Schmirewall.** (20 pts) Sarah is installing a network firewall for her company. Being familiar with the principle of fail-safe defaults, she has configured the firewall to DENY all packets by default. Now she needs to identify the minimal access rules that will allow her organization to use its Internet connection. For example, her organization will need to be able to send and receive email through the firewall, and uses a central mail server at IP address 10.1.100.100. So she has added rules to the firewall that look like this:

| SRC ADDR | DEST ADDR | SRC PORT | DST PORT | PROTOCOL | ACTION |
|---|---|---|---|---|---|
| 10.1.100.100 | * | * | sendmail | TCP | ALLOW |
| * | 10.1.100.100 | * | sendmail | TCP | ALLOW |

The organization has determined that it will also require the following kinds of Internet access:

- Incoming SSH access to a VPN server, at 10.1.100.200.
- Access to the web, through a proxy that whitelists approved sites. The proxy's address is 10.1.200.200.
- Outgoing SSH access to three client sites: 0.1.2.3, 42.42.42.42, and 3.14.15.9.

List the minimal set of firewall rules necessary to allow these connections. List some potential vulnerabilities associated with this ruleset. Can the firewall and proxy servers defend against these vulnerabilities?

**4. False Positive Answer.** (20 pts) Anderson's chapter 11 details several ways to defeat physical intrusion detection systems (a.k.a. "burglar alarms"). One of the common ones is to artificially create "false" alarms so that the true alarm is ignored. Let's investigate this idea with respect to computer intrusion detection systems.

(a) An old Snort rule says that any HTTP packet that includes "`/..%c0%af../`" should trigger an alarm, as an attempted IIS exploit. Explain why in "normal" usage this rule would have a low false positive rate.

(b) Suppose Eve discovers a web server, `vulnerable.org`, that is vulnerable to the IIS Unicode exploit and she wants to exploit the hole without having it noticed. What are a few ways Eve can temporarily increase the false positive rate at `vulnerable.org` for the rule, without getting her IP address noticed?

(c) What can you conclude about "advertised" false positive and false negative rates?