

# CS110 - Principles of Computer Systems

## Final Exam

**(Total time = 180 minutes, Total Points = 174)**

Name: (please print) \_\_\_\_\_

**In recognition of and in the spirit of the Stanford University Honor Code, I certify that I will neither give nor receive unpermitted aid on this exam.**

Signature: \_\_\_\_\_

**This examination is closed book and closed notes. You may not collaborate in any manner on this exam. You have 180 minutes to complete the exam. Please write your answers on the exam. Note there is one problem per page so the amount of space provided does not necessarily provide an indication of the expected length of the answer. In other words, do not feel compelled to fill every nanoacre of the exam with writing. Before starting, please check to make sure that you have all 13 pages.**

Question	Points	Score
1	16	
2	12	
3	12	
4	16	
5	12	
6	12	
7	18	
8	12	
9	18	
10	12	
11	18	
12	16	
<b>Total</b>	<b>174</b>	

- 1 . (16 points) For each of the pairs of performance optimization techniques listed below, describe any synergistic benefits<sup>1</sup> that can be achieved when using the pair together. If no synergistic benefits can be achieved state why.

(a) *Speculation and batching*

(b) *Hiding latency using concurrency and caching*

---

<sup>1</sup> Two optimization techniques are synergistic if together they can produce a speedup that is greater than the sum of the individual techniques.

2 . (12 points) Assume you have a computer system with a modern CPU scheduler (such as Linux) and you modify the scheduler to assign all jobs the same priority, removing the code that dynamically updates the priorities. Describe how you expect your changes to the scheduler would affect the system under the following workloads. Be sure to explain your answers.

- (a) All processes are I/O-bound on the disk.
- (b) All processes are CPU-bound and do little I/O.
- (c) Half the processes are I/O-bound and half are CPU-bound.

3. (12 points) Assume you are remotely logged into one of the class Linux machines (say corn12.stanford.edu) using a TCP connection. Each character you type is echoed back over the TCP connection by the software running on the Linux machine.

(a) Although TCP has special acknowledgement packets, if you were to peer at the packets exchanged between the two endpoints, you wouldn't see any special acknowledgement packets. Explain how TCP can provide its guarantees without using special acknowledgement packets.

(b) In response to packet loss, the TCP protocol reduces the size of the window for the connection. Explain why it does this.

- 4 . (16 points) When the IP header was designed in the early 1980s, it used a 16-bit integer to describe the packet's length, thus limiting packet sizes to 64KB. This limit has remained in place to this day, even though since then the underlying networking technologies have changed and even though machines have increased in speed by several orders of magnitude.
- (a) Give three reasons why the maximum packet size hasn't increased.
  - (b) Provide a reason why it might be desirable to increase packet sizes.

5. (12 points) Ethernet, as well as several other systems we talked about in class, uses *exponential backoff*. Describe what the exponent is and how it is computed in these systems.

6. (12 points) Describe the relationship between protocol headers and networking layers.

7. (18 points) Assume you have been given the task of implementing a high performance RPC system for a 1000-node cluster of machines. The RPC runs on top of the UDP protocol. Your design can limit the system to one outstanding RPC call per machine. Answer the following three questions about your design.
- (a) Given the limit of one outstanding call, would your design need to worry about flow control?
  - (b) How about congestion control?
  - (c) If you had to choose between making the RPC semantics “at least once”, “at most once”, or “exactly once” which would you think would be the easier to implement.

Be sure to justify your answers.

8. (12 points) In class we talked about the encoding of the access matrix using either access control lists (ACLs) or capabilities. Answer the following questions:

(a) Browsers support the notion of “cookies” by which a website can have the browser store data on its behalf. Would cookies be more useful for implementing access control lists or capabilities? Justify your answer.

(b) Although both access control lists and capabilities are different encodings of the same access matrix, explain how one of them typically subsumes the naming functionality and why the other doesn't.

9. (18 points) Recently WikiLeaks<sup>2</sup> published a bunch of secret US diplomatic cables. Answer the following security-related questions:

(a) Shortly after posting the documents, WikiLeaks posted a Twitter message saying its website was “under a mass distributed denial of service attack.” Describe what is meant by that.

(b) Given that WikiLeaks functions as a document distribution center, describe whether it should be worried about each of the following attacks. If no worries, justify your answer. If worries provide a description of the attack WikiLeaks should worry about.

- (i) Timing Attack
- (ii) Eavesdropper Attack
- (iii) Trojan Horse Attack

(c) The founder of WikiLeaks appears to be feeling a little paranoid these days. He released a statement say that he possessed a “poison pill” in the form of a super sensitive document that would be released if anything bad happened to him. Using cryptography primitives discussed in class, describe how he could publish evidence that would convince the US government he had a particular document but wouldn't disclose the contents of the document to anyone.

---

<sup>2</sup> WikiLeaks is a website that publishes documents supposed leaked from non-public sources. Its main functionality is the public distribution of documents.

10. (12 points) Public key cryptography was invented much more recently than shared key cryptography. Its invention was greeted with much excitement since it appears to solve one of the big challenges with using shared key cryptography for confidentiality. Describe this challenge and how public key cryptography helped bring about the widespread use of cryptography for confidentiality.

11. (18 points) Traditional Unix systems supported two different access paths to disks. One of the paths was the Unix file system abstraction we studied in Assignment #1. The other path was called raw disk devices allowed an application to directly read and write sectors from the disk. Since the system is constructed in layers, it is possible to view the raw disk device access as an example of layer bypass. Answer the following questions:
- (a) Although it was possible to use both paths (raw disk device and file system) simultaneously, it was strongly discouraged. Explain why this wasn't a safe operation to do.
  - (b) Even if you imposed mutual exclusion so that a process accessing the raw disk device never ran when file system calls were active, it still was unsafe. Explain why.
  - (c) Database systems frequently used the raw disk device path rather than the file system to access disk storage because they didn't like some of the functionality in the file system. Describe how this relates to the end-to-end argument in networking.

12. (16 points) In the first part of the course we talked about *enforced modularity* as being useful for building robust systems. Describe the relationship between enforced modularity and the following concepts from the security portion of the class:
- (a) Complete mediation.
  - (b) Trusted Computing Base.
  - (c) Covert channels.