

CS 4375, Operating Systems, Final Examination

Fall 2024, December 12

name _____

165 minutes 10:00 – 12:45

Three pages of handwritten notes are allowed. There are no room re-entry privileges. Please answer concisely, but if necessary use the last page to continue any answer.

/80

1. [22 points] True/False

- T F** Every IP packet contains the IP address of the destination host.
 - T F** Every TCP packet contains the IP address of the destination host.
 - T F** Every TCP *header* contains the IP address of the destination host.
 - T F** Every TCP packet contains the domain name of the destination host.
 - T F** A routing table is a complete listing of all reachable nodes in the internet.

 - T F** Every packet has a payload, that is, no packet contains only headers.
 - T F** IP is used primarily for short-distance communication, and TCP for longer distances.
 - T F** TCP is used primarily for short-distance communication, and IP for longer distances.
 - T F** Real-time protocols, such as streaming video or Voice over IP, run on top of TCP.
 - T F** To guarantee that some lines of code execute atomically, you can surround them with a lock.

 - T F** In bash, ending a command with & makes the command execute uninterruptedly.
 - T F** `threading.join()`, or, equivalently in C, `pthread_join()`, joins the stacks of two threads.
 - T F** A process can be associated with more than one port.
 - T F** Cached DNS-to-IP mappings are authoritative.
 - T F** The Posix API for files on disk is more complex than the API for files on a USB memory stick.

 - T F** Each process has its own memory, which other processes cannot access (except if using `mmap()`).
 - T F** Each process has its own file system, which other processes cannot access, ever.
 - T F** It is more important to keep your encryption key secret than the encryption algorithm.
 - T F** It is more important to keep your public key(s) secret than your private key(s).
 - T F** If a non-masked interrupt occurs while process X is running, the OS may immediately suspend X so that it can take control and handle the interrupt.

 - T F** If an interrupt occurs while a process X is running, the OS may then decide to run it again, or to run a different process, depending on the scheduling algorithm.
 - T F** A mutex is a lock.
2. [1] Upgrading to “high-speed internet” usually means
- a) The transmission time will be significantly shorter
 - b) The propagation time will be significantly shorter
 - c) The ISP will handle your packets using more advanced protocols
 - d) none of the above

/23

3. [1 point] Would it be possible, in principle, to deliver IP packets across the French Optical Telegraph?
- No, because the latency (delay) is too high.
 - No, because the bandwidth is too low.
 - No, because there is no way to transmit packet headers.
 - No, for some other reason.
 - Yes, in principle.
4. [1] The domain name earth.cs.utep.edu corresponds to the IP address 129.108.156.28. This tells us that
- 28 is the code for .edu and 129 is the code for earth
 - 129 is the code for .edu and 28 is the code for earth
 - 129 is the code for .edu and 156.28 is the code for earth
 - None of the above
5. [1] In an IP address, like 129.108.111.214, the dots serve to
- Mark the boundary between network (backbone), subnet, subsubnet, and LAN
 - Indicate that the numbers are decimal or hex, not binary
 - Separate the hierarchical part of the address from the flat part
 - none of the above
6. [1] Every running process starts with 3 file descriptors open. Which of the following is *not* one of these?
- stderr
 - stdio
 - stdin
 - stdout
7. [1] If I view a website, and all packets from the server to my client machine follow the same path through the internet, this is because
- Packets always follow the best path, and this is guaranteed by the Internet Protocol (IP)
 - The path is fixed for the duration of the connection, and this is guaranteed by the Transmission Control Protocol (TCP)
 - Consistency of web packet paths is guaranteed by the Hypertext Transfer Protocol (HTTP)
 - all of the above
 - none of the above
8. [1] If a socket does a `recv()` with a timeout (versus without), this will potentially cause more network traffic if
- It's a UDP socket
 - It's a TCP socket
 - both
 - neither
9. [1] Steve implements a producer-consumer problem with a circular buffer, two locks, and two semaphores. However his program never finishes. What could explain this?
- Deadlock due to a circular dependency between the semaphores
 - Deadlock due to a circular dependency between the locks
 - Deadlock due to a circular dependency involving one semaphore and one lock
 - Something else
 - All of the above
10. [2] Imagine that a process on computer X wants to send a packet to a specific process on computer Y. In order to do this, it needs of course the address of Y, but also some way to specify the process it wants to talk to. How can do this?

11. [1] Name one thing that will reliably cause increased propagation time. (Hint: perhaps consider increasing the distance between two nodes, increasing the amount of data to send, or switching from IPv4 to IPv6.)

12. [5] In an unlucky case, a simple memory access can involve many steps, including
- a) the high-order bits of the virtual address are used as an index into the page table
 - b) a page for the process that was swapped out to disk is copied into physical memory
 - c) some page of physical memory is copied onto disk
 - d) the physical address is computed by adding the low-order bits (the offset) to the page start address
 - e) the specified value is loaded from or stored to the physical address

In what order do these happen? first ___, second ___, third ___, fourth ___, fifth ___

13. [1] For the above example, what would differ if the memory access is not a data read/write, but an instruction fetch?

- The order of steps would differ
- One step would never occur
- An additional step would occur
- none of the above

14. [3] Senator Smith introduces the Wonderfully Improve National Security Act (WINS). This will
- Prohibit the manufacture or sale of any computer or device with less than 512MB of memory.
 - Prohibit any computer or device from having more than 8 processes at any one time.
 - Prohibit any process from using more than 64MB of memory.
- in order to remove the need for virtual memory. Evaluate this proposal.

15. [4 pts] Match

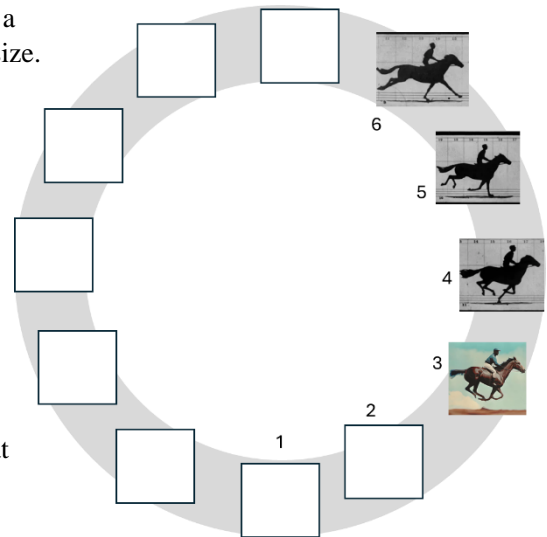
- | | |
|-----------------------|---|
| a. applications layer | ___ supports some specific user-visible functionality |
| b. physical layer | ___ handles addressing, such that a packet reaches the intended host |
| c. transport layer | ___ defines the hardware of a network connection |
| d. network layer | ___ may handle error detection, flow control, packet sequencing, etc. |

16. [3] The system call to rename a file, namely `rename()`, is usually atomic. What problems could arise if this were not the case? Hints: Recall that a directory contains a list of filenames, each mapped to some sectors(s) on disk. Consider the case of renaming `file09.csv` to `file10.csv`.

17. [6] A certain video display tool has three threads, all sharing a circular buffer with space for 12 frames, each about 5MB in size.

- Thread A repeatedly reads video frames from a file and copies them into the buffer
- Thread B crops and adds color to frames
- Thread C shows each colorized frame to the user, and then frees up the space.

The figure (smithsonianmag.com) shows the situation after A has read in six frames (1-6), B has colorized three (1-3), and C has already shown two to the user and freed them up (1-2).



- To avoid threads doing busywait, we'd like to use condition variables (or semaphores). How many would we need? What name(s) would you give to each condition variable (or semaphore)?
- To keep track of things, we would need various pointers into the buffer, such as `nextFrameToColorize`, and we code to modify them, such as `nextFrameToColorize++`. In addition to condition variables (or semaphores) would we also need locks around the code for updating these pointers? Why or why not?

18. [3] Jason proposes a new Virtual Machine system, in which every *thread* runs in its own virtual machine, for the sake of isolation. Evaluate this proposal.

19. [3] (Chris Dees) System calls might seem like regular function calls to programmers, but they are fundamentally different. Briefly explain.
20. [1] (Martin Mendez, Billy Baird) Which of the following statements about virtual memory is true?
- a. Virtual memory increases the size of physical memory in the computer
 - b. Virtual memory eliminates the need for physical RAM
 - c. Virtual memory can use a portion of the secondary storage (disk, etc.) as if it were additional RAM
 - d. Virtual memory is only used in realtime operating systems
 - e. Page faults occur when the CPU cache misses a memory reference.
 - f. none of the above
21. [7] Imagine a pancake-printing device, whose hardware includes a data register, a status register, and a command register. Further imagine a pancake printing app that interacts with the device driver in order to enable the user to 1) send a jpg file to the printer, 2) start printing, and 3) show the status of the printer (ready, printing, or done). Explain how the device driver would support this, with a rough execution trace. Assume that efficiency considerations rule out polling.

22. [4] Samantha in a job interview states that “Every network protocol is independent of all other protocols: that is, it does not depend on higher-level protocols or on lower-level protocols”. What is correct about her statement and what is incorrect.

23. [7] Trying to speed up some code, Janine discovers that half the time is spent in one while loop:

```
while not otherThreadIsDone:
    pass    // does nothing, just syntax for an empty loop
// now the other thread is finally done
complicatedFinalFunction()
```

Her boss suggests trying four things:

- a. Running the code on a machine with a faster processor
 - b. Running the code on a machine with more memory
 - c. Modifying the code for the thread that sets `otherThreadIsDone`
 - d. Giving higher priority to the other thread
- i) Depending on the situation, any of these could be a good solution or completely pointless. Pick one and explain why it could be good but why it could also be pointless.
- ii) What else could she try to reduce the wasted time?