EE122 Final Review Section

Many thanks to previous years finals/midterms at Berkeley and Princeton.

Question 1: Short Answers

A) Internet Principles

We are faced with three basic design decisions
a) How to break the overall system into modules
b) Where these modules are implemented:
c) Where the associated state is stored:

We have three general principles to guide us. Which of the above design decisions do the following help clarify: (choose either a, b, or c for principle below)

i. Fate Sharing:
ii. End-to-End Principle:
iii. Layering:

B) Routing

Which all of the following three choices form a necessary and sufficient condition (independently) for successful routing (assuming routing is deterministic):
a) Forward all packets along the shortest path, and never send a packet out along the port it came in on.
b) Avoid loops (a packet returning to a node it has already visited) and dead-ends (where a packet is dropped before it reaches the destination).
c) Always forward packets to a node closer to the destination.

C) TCP

Assume the advertised window of a TCP receiver does not exceed 3 segments. Give a "scenario" in which TCP Reno achieves better performance than TCP Tahoe, if any. Justify your answer.
(A "scenario" can be specified by the number of segments sent by the sender, and by indicating which segments are lost, if any.)

There is no scenario because TCP Reno can never obtain the three duplicate ACKs necessary to do fast retransmissions and fast recovery when the receiver's advertized window does not exceed 3 segments.

D) Wireless

Is RTS/CTS guaranteed to eliminate hidden terminals? If yes, argue why; otherwise, present a counterexample.

E) Ethernet and IP

a) Suppose Ethernet was the only existing LAN technology, so every host in the Internet was part of a local Ethernet and thus had a globally-unique Ethernet address. Your friend suggests that IP isn't necessarily anymore, and the entire Internet could just be one large, switched Ethernet instead. Give two reasons why using existing Ethernet protocols for this is a bad idea from a networking perspective (i.e., don't consider security or privacy).

(b) What about the other way around, why do we not simply assign IP addresses to network adaptors, instead of dealing with both MAC (Ethernet) and IP addresses? Give one reason.

Question 2. True / False

For each of the following statements, indicate whether the statement is True or False, and provide a very short explanation of your selection (2 points each).

a. End-to-end packet delay of cut-through routers is smaller than of store & forward routers.

b. According to end-to-end argument reliability must be implemented at networking layer.

c. Pipelined requests reduce the response time in HTTP.

d. CIDR allocates IP addresses less efficiently than Classful Addressing.

e. In the OSI model, the transport layer can directly invoke (use) the data link layer.

Question 3: Multiple choice

1. When a TCP packet arrives at a host, in order to direct the segment to the appropriate socket, the operating system's network stack uses the following fields:
a. transport protocol number
b. destination IP address
c. source port number
d. destination port number
e. destination MAC address

2. Which of the following is/are true about wireless networks?
a. All wireless networks must use access points.
b. The sender can always detect a collision without feedback from receiver.
c. Collisions are minimized when RTS/CTS mechanisms are used.
d. TCP congestion control mechanisms work poorly in wireless environments if they do not perform any type of link-layer retransmission.
e. Wireless networks generally have higher loss rates than that in wired networks.

3. Which of the following is/are true about routers?
a. Routers reassemble IP fragments if the next link can handle the full datagram.
b. Routers can arbitrarily drop packets if they want.
c. Routers can not change the IP packets they forward at all
d. On a router with many 1 Gbps ports, the router backplane can only handle 1 Gbps on the shared bus, leading to potential congestion.
e. In their line cards, routers lookup forwarding tables in the incoming direction and queue packets in the outgoing direction.

4. Otto Pilot built a home-brew network with 20 computers. The RTT between each computer is 10 ms. Communication between computers uses a simple UDP query and response protocol. If no response is received within 20 ms, a computer retransmits the request. Soon, Otto notices

congestion collapse in his network. Which of the following techniques is/are guaranteed to prevent congestion collapse?
a. Double the timeout value from 20 ms to 40 ms.
b. Increase the size of the queue in each router from 4 packets to 8 packets.
c. Use exponential backoff in the timeout mechanism while retrying queries.
d. If a query is not answered within a timeout interval, multiplicatively reduce the maximum rate at which the client application sends query packets.
e. Use a flow control window at each receiver to prevent buffer overruns.

5. Which of the following statements is/are true about fair queuing algorithms?
a. For a router serving n flows, fair queuing ensures that no flow can transmit at more than 1/nth the link's capacity.
b. Fair queuing is often used in the core of the Internet to minimize denial-of-service attacks by individual senders.
c. Fair queuing algorithms conceptually track the number of bytes each flow consumes, rather than the number of packets.
d. Fair queuing algorithms can be based both on unique flows (5-tuples) or on unique classes of traffic (specified in the IP TOS field).
e. Fair queuing with drop-tail policies experiences synchronized losses between multiple senders.

6. Which best describes the Ethernet protocol? (Circle ONE)
a. Talk only if you hear no one else talking, but stop as soon as you hear anybody else.
b. Pass a ticket around and only talk if you are holding the ticket.
c. Raise your hand and wait till a moderator gives you permission to talk.
d. Every person is scheduled a time to talk.

7. Which of the following is/are true about Address Resolution Protocol (ARP) and learning switches? (Circle ALL that are correct)
a. A learning switch maintains state that maps IP addresses to hardware (MAC) addresses.

b. A learning switch maintains state that maps MAC addresses to IP addresses.
c. A host's ARP table maintains state that maps IP addresses to hardware (MAC) addresses.
d. A host's ARP table maintains state that maps hardware addresses to IP addresses.

8. Which of the following is/are true about a communications channel that uses time- division multiplexing? (Circle ALL that are correct)
a. There may be times when the channel is idle, even if a sender has data to send on the channel.
b. The channel requires the sender's and receiver's clocks to be closely synchronized.
c. Data in the channel could experience variable delays due to queuing.
d. In times of high utilization, a sender could be completely denied access to the channel.

9. Which of the following is/are true about increase/decrease policies for fairness and efficiency in congestion control? (Circle ALL that are correct)
a. Additive increase reduces fairness.
b. Additive increase improves efficiency.
c. Multiplicative increase improves fairness.
d. Multiplicative decrease improves fairness.

10. Which of the following is/are true about DNS? (Circle ALL that are correct)
a. A query for an A record may return multiple IP addresses in the response.
b. A query for an NS record may return multiple IP addresses in the response.
c. A short TTL on an NS record reply runs the risk of increasing traffic at the root or GTLD nameservers.
d. A short TTL on an A record reply runs the risk of increasing traffic at the root or GTLD nameservers.

11. Which of the following is/are true about web caches? (Circle ALL that are correct)

a. A web cache, or proxy server, is a network entity that satisfies HTTP requests from clients.
b. A web cache is both a client and a server at the same time.
c. HTTP does not explicitly support caching or cache consistency.
d. All HTTP objects are cacheable.

12. Which of the following is/are true about persistent HTTP connections? (Circle ALL that are correct.)

a. Persistent HTTP allows a server to track the client's requests through a persistent session.
b. Only one TCP connection must be opened for downloading a "page", if that page does not include any embedded objects served by other servers.
c. Persistent HTTP shows the greatest performance advantage over non-persistent HTTP when downloading a page with large objects.
d. When the server has finished sending data for all objects referenced by the initially requested page, the server closes the connection.