

## CS 168 Discussion 12 - Final Review

### Question 1: Short Answer

#### A) Internet Principles

We are faced with three basic design decisions:

- a) How to break the overall systems 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 us clarify? (Choose either a, b, or c for each principle below)

- i. Fate Sharing: **c**
- ii. End-to-End Principle: **b**
- iii. Layering: **a**

#### B) Routing

Which of the following three choices form a necessary and sufficient condition for successful routing (choose all that may apply):

- a) Forward all packets along the shortest path, and never send a packet out along the port it came from
- 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 (as a reminder, TCP Reno continues AIMD after a fast retransmit, and TCP Tahoe enters slow-start).

**There is no scenario; TCP Reno can never obtain the three duplicate ACKs required to do fast recovery, so the performance will be the same.**

#### D) Wireless

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

**No. Counterexample: asymmetric links and the sender does not hear a CTS from the receiver, but that receiver can hear what the sender is sending.**

## E) Ethernet and IP

a) Suppose Ethernet were 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 necessary 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).

1. Flooding doesn't scale for broadcast transmission (when address hasn't been learned by the switch). Even if switched Ethernet, flooding doesn't scale for ARP or spanning tree.

2. MAC addresses aren't hierarchical → large routing tables.

b) Why don't we simply assign IP addresses to network adapters, instead of dealing with both IP addresses and MAC (Ethernet) addresses? List one reason.

1. Gives up topological addressing → large routing tables (since no aggregation).

2. End hosts would have to participate in routing protocol and announce addresses when they migrate, like a global ARP, which doesn't scale.

## Question 2. True/False

- a) End-to-end packet delay of cut-through routers is smaller than that of store & forward routers. **T**
- b) According to the end-to-end argument, reliability should be implemented at the network layer. **F**
- c) Pipelined requests can reduce the response time in HTTP. **T**
- d) CIDR allocates IP addresses less efficiently than Classful Addressing. **F**
- e) In the OSI model, the transport layer can directly use the services of the data link layer. **F**

- f) One key goal of Software Defined Networking (SDN) is to define abstractions for the network data plane. **F**

### Question 3: Multiple Choice

*Unless otherwise stated, circle all that apply.*

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 (circle all that apply):
  - 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 the receiver.
  - c. **Collisions are minimized when RTS/CTS mechanisms are used.**
  - d. **TCP congestion control mechanisms work poorly in wireless environments if link-layer retransmission is not performed**
  - e. **Wireless networks generally have higher loss rates than 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 drop packets.**
  - c. Routers cannot change the IP packets that they forward.
  - 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 20ms to 40ms.
  - 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  $1/N$  of the link capacity.
  - b. Fair queuing algorithms conceptually track the number of bytes each flow consumes, rather than the number of packets
  - c. Fair queuing algorithms can be based on unique flows or on unique classes of traffic (e.g. the IP Type-of-service field).
  - d. 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 anyone else.
  - b. Pass a ticket around and only talk if you are holding the ticket.
  - c. Raise your hand and wait until a moderator gives you permission to talk.
  - d. Every person is given a scheduled time to talk.
7. Which of the following is/are true about ARP and learning switches?
- a. A learning switch maintains state that maps IP addresses to 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 addresses.
  - d. A learning switch maintains stat that maps MAC addresses to a port.
8. Which of the following is/are true about a communications channel that uses

time-division multiplexing?

- 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 completely be denied access to the channel.

9. Which of the following is/are true about DNS?

- 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 name servers.
- d. A short TTL on an A record reply runs the risk of increasing traffic at the root or GTLD name servers

10. Which of the following describe a reason why the Spanning Tree Protocol is not used in datacenter networks?

- a. The spanning tree protocol does not scale to a large number of hosts
- b. The spanning tree protocol prevents the use of multiple paths between two hosts.
- c. It is difficult to implement the spanning tree protocol in software switches, which are common in datacenter networks.
- d. Many end-hosts in datacenters are virtual machines that use "virtual" network interface cards, which the spanning tree protocol doesn't support.

11. Which of the following is/are true about persistent HTTP connections?

- 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 webpage, 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.