CS 168 Fall 2022 | Introduction to the Internet Sylvia Ratnasamy

Finals Practice Questions (#2)

1 True/False

Please answer the following (on the answer sheet) with T or F as appropriate.

1. The Internet's packet switching approach was borrowed from the design of the telephone network.
2. In the Internet architecture, we think of transport as layer 3.
3. The TCP header is the outermost header on packets that use TCP.
4. AIAD (Additive Increase Additive Decrease) is "unfair" in that it could lead to different average bandwidth allocations to flows with the same RTT.
5. UDP never uses a checksum.
6. HTTP typically runs over UDP.
7. DNS typically runs over UDP.
8. DHCP typically runs over UDP.
9. ARP typically runs over UDP.
10. ARP requests are typically broadcast.
11. ARP responses are typically broadcast.
12. TCP packets with the ACK flag set cannot carry payloads.
13. In DNS, the MX record is the name of the authoritative server for that domain.
14. If we disabled all timers in TCP and relied exclusively on duplicate ACKs for loss detection and retransmissions, the resultant protocol might be inefficient but would still ensure reliable delivery.
15. According to Rob Shakir, a benefit of SDN is that it allows a cloud provider to write their own network management code rather than rely on vendor code.
16. OpenFlow is an abstraction of the forwarding model of switches.
17. For scalability, cellular networks avoid storing per-user state on the data plane.
18. Google runs one unified wide-area network.

2 Short Multiple Choice Questions

(a) Which is the faster way of delivering a 1 TB file from Berkeley to Stanford (assuming the distance between these two locations for both methods is 45 miles): (Pick one.)

1. Sending it over a 1Gbps link
2. Driving at an average speed of 25mph carrying a 1 TB disk.

(b) Which of the following is provided by both UDP and TCP? (List all that apply.)

1. Reliable transfer
2. Congestion control
3. Mux and demux from/to application processes
4. A bytestream abstraction

(c) Which one of the following enables a sender to respond to congestion before router queues are full?

1. ECN
2. Fair-Queueing
3. Faster route convergence
4. AIMD

(d) How does HTTP compensate for the fact that it is a stateless protocol? (Pick one.)

1. Caches
2. Encryption
3. Reverse proxies
4. Cookies

(e) In SDN, what is the interface between the NOS and switches? (Pick one.)

1. Global Network View
2. Abstract Network View
3. OpenFlow

(f) When a host wants to obtain the IP address of www.ebay.com, where does it send its initial DNS request? (Pick one.)

1. The DNS root
2. The .com TLD server
3. The authoritative server for ebay.com
4. The host's local DNS server

(g) A host gets information in many ways. For each question below, choose from this list of possible answers (you can use the same answer for more than one question):

- DNS
- DHCP
- ARP
- The host's own hardware
- Doesn't need to know

When you take your laptop to campus and want to access the www.chocolics.com website, which of the answers above best describe how the laptop determines the following (assuming no information resides in caches, no proxies are being used, etc.)? (List one option for each of the following.)

1. Laptop IP address
2. www.chocolics.com's IP address
3. Laptop's first-hop router IP address
4. Laptop MAC address
5. The MAC address of www.chocolics.com's server, assuming the server is not on your laptop's subnet
6. Laptop's first-hop router MAC address

(h) Consider two TCP connections A and B, with the following properties:

- Connection A: MSS = 1000bytes, RTT = 100msec, drop probability = 1%
- Connection B: MSS = 2000bytes, RTT = 500msec, drop probability = 4%

What is the ratio of A's throughput to B's?

3 HTTP and TCP

A client wants to download an article from a server. This involves two steps:

- Step 1: Download a master index page which precisely fits into one MSS-sized segment (including all HTTP headers). The index page contains links to three images that are stored on the same server.
- Step 2: Download the three images, each of which requires **two MSS-sized segments** to send (including all HTTP headers).

Assumptions:

- **Ignore packet transmission times** and assume no packets are dropped.
- Assume HTTP requests are very small compared to the MSS.
- We are interested in how long it takes until all the relevant data is delivered; do not include the time required for any final acknowledgements or to close the TCP connection.
- TCP uses an initial congestion window of 1 MSS and an initial slow start threshold of 20 MSS. There are no other constraints on the sending behavior other than TCP's congestion control.

We consider two implementation options:

- Concurrent requests with non-persistent TCP connections, where there is no limit on the number of concurrent connections.
- Pipelined requests with persistent TCP connections

(a) Including the TCP connection establishment, how many RTTs are needed for the client to obtain **only** the master index page from the server for each of the above implementation options? (**Give your answer in multiples of the RTT.**)

1. Using concurrent, non-persistent connections
2. Using **pipelined, persistent** connections

(b) After the master index page is received, how many additional **TCP connection establishment exchanges** (considering the sequence SYN-SYNACK-ACK as a single exchange) are necessary for each of the above implementation options to finish the downloading of the article?

1. Using concurrent, non-persistent connections
2. Using **pipelined, persistent** connections

(c) Including the TCP connection establishment, what is the total number of RTTs needed to download the entire article using each of the implementation options? (**Give your answer in multiples of the RTT.**)

1. Using concurrent, non-persistent connections
2. Using **pipelined, persistent** connections

4 Detailed Sequence of Packets

Consider subnet 1 and subnet 2 connected by a router R. Each subnet has a local DNS server and DHCP server, and there are two existing servers: server B on subnet 1 and server C on subnet 2. At the start of this scenario, host A arrives on subnet 1, and then uses HTTP to fetch an object from server B, and then an object from server C. Host A only knows the DNS names of B and C at the beginning of this scenario.

Host A tears down the TCP connection to B before starting the download from C (where we use the two one-sided teardowns that involves a FIN then an ACK in response, and then from the other side a FIN then the returning ACK). We assume that all objects fit within a single packet. There is no piggy-backing of data on ACKs sent in response to an arriving packet (i.e. an ACK is sent immediately after each TCP packet arrives, without waiting for the HTTP process to respond with data). Similarly, the initial FIN in the teardown is not piggy-backed on a TCP ACK of newly arriving data. The ARP and DNS caches are initially empty, but they retain all data during the rest of the scenario.

(Part a) During this process, A will send and receive a set of packets. Use the following options to answer the questions:

- DHCP packet
- ARP request
- ARP response
- DNS request
- DNS response
- TCP packet carrying no payload
- HTTP request
- HTTP response

1. What kind of packet is the third packet **received** by A? (This is counting only the packets received by A.)
2. What kind of packet is the eleventh packet **received** by A? (This is counting only the packets received by A.)
3. What kind of packet is the eighth packet **sent** by A? (This is counting only the packets sent by A.)
4. What kind of packet is the fifteenth packet **sent** by A? (This is counting only the packets sent by A.)

(Part b) Assume that for the scenario described above, subnet 1 uses an L2 network comprised of learning switches. Further assume that at the beginning of the scenario they have no cached routing state, but that all learned entries are retained for the rest of the scenario. We can categorize all packets leaving A as:

- Delivered directly to the intended destination in subnet 1 (i.e., not broadcast, and not flooded at any switches)
- Flooded by one or more switches in subnet 1
- Broadcast to all destinations in subnet 1

In part (a) you had to analyze the sequence of messages sent and received by A. In this portion of the question, we are asking you to identify to which category (from the three choices above) do the following packets belong. For each packet, choose only one option.

1. DHCP Discovery
2. DNS request for server B's address
3. DNS request for server C's address
4. SYN packet to server B
5. SYN packet to server C

5 SDN

[Questions on SDN in the final will be limited to T/F, simple multiple-choice. This question isn't quite that, but use it for practice. – Sylvia]

As presented in Scott Shenker's lecture, an SDN network has several components which are listed below in alphabetical order.

a) Abstract Network View

b) Control Program

c) Global Network View

d) Network Operating System

e) Switch Interface (e.g., OpenFlow)

f) Virtualization Layer

In the following Faulknerian paragraph, fill in each of the blanks using the above phrases (list one letter from the list above in each blank below): (hint: start with the first and last blanks, and work inwards)

In a Software-Defined Network, the ____ expresses the operator's intentions for network control by configuring the switches in the ____ , which is part of the API provided by the ____ that is based on the ____ it receives from the ____ , which in turn uses the ____ to control the physical switches.

6 Reliability

Consider a flaky link where the initial transmission of a data packet is dropped if its number is prime (in other words, the initial transmissions of D2, D3, D5, D7, D11, D13 are dropped, but subsequent retransmissions are OK). Note that the ACKs are cumulative and numbered according to the next expected packet (hence, A4 indicates the receipt of D1, D2, and D3).Hosts x and y are using a transport protocol with sliding window flow control with a constant window size of 5 packets; three duplicate ACKs trigger a retransmission. Assume that the latency of the link is significantly longer than the transmission time of 5 packets. Indicate what packets are sent from host x (you don't need to indicate what ACKs are generated, though it may be helpful and we have entered the first few entries below). Mark which packets are retransmits due to timeouts and which are retransmits due to duplicate acknowledgements.

1. D1 A2
2. D2 (dropped) -
3. D3 (dropped) -
4. D4 A2
5.
6.
7.
8.
9.
10.
11.
12.
13.
14.
15.
16.
17.
18.
19.