Stimulus

Background on the TCP handshake

These questions are about how a socket is established for the TCP protocol. TCP is a transport layer protocol that provides reliable packet delivery. In assignment 1, you used TCP sockets to communicate between a client and server. Every TCP packet has a TCP header *and* an IP header. The IP header has the usual fields, including source and destination IP addresses and ports.

A TCP connection is established via a three-step handshake between the client and server. The TCP header has two bits allocated for this purpose, the SYN bit and the ACK bit:

1. **SYN (Synchronize)**: The client initiates the connection by sending a TCP packet with the SYN (synchronize) flag set.
2. **SYN-ACK (Synchronize-Acknowledge)**: Upon receiving the SYN packet, the server responds with a SYN-ACK packet. This packet acknowledges the client's request by setting both the SYN and ACK flags.
3. **ACK (Acknowledge)**: The client responds to the SYN-ACK with an ACK packet. Once this third packet is received by the server, the connection is fully established, and data can be exchanged.

These packets also contain other crucial information, but that is not relevant for now.

1 Multiple Choice 1 point

A server is listening on IP address `19.16.1.1` and port `80`. A client at IP address `100.12.0.5` sends a TCP **SYN** packet from port `50000` to the server at port `80`, requesting a connection.

When the server responds to the client's SYN packet with a SYN-ACK packet, what will be the **source IP address and port** in the server's response?

- (●) Source IP: `19.16.1.1`, Source Port: `80`
- ( ) Source IP: `100.12.0.5`, Source Port: `80`
- ( ) Source IP: `19.16.1.1`, Source Port: `50000`
- ( ) Source IP: `100.12.0.5`, Source Port: `50000`

2 Multiple Choice 1 point

What will be the destination IP address and port in the server's response (SYN-ACK packet)?

- Destination IP: `19.16.1.1`, Destination Port: `80`
- **Destination IP: `100.12.0.5`, Destination Port: `50000`**
- Destination IP: `19.16.1.1`, Destination Port: `50000`
- Destination IP: `100.12.0.5`, Destination Port: `80`

3 Multiple Choice 1 point

Does the client need to know the port that the server was listening to (in this case, 80) before initiating the connection?

- (●) Yes, the client must know the exact port number.
- ( ) No, the client can request any port.
- ( ) Yes, but the default port, 80, always works.
- ( ) No, the server will notify the client of the correct port number during the handshake.

4 Essay 1 point

Can the client create another connection to the same server? If yes, what source/destination IP/port should it use on its packet? If no, why not?

5 min / 100 max (word limit)

Must use at least 5 words

Network A:

Stimulus

BGP Convergence

Background (covered in class):

Customer-Provider relationship

- Customer pays provider for access to the Internet
- Provider exports its customer's routes to everybody
- Customer exports provider's routes only to its customers

5 True or False 1 point

Network A

- ○ True
- ◉ False

Peer-peer relationship:

- Peers exchange traffic between their customers
- AS exports *only* customer routes to a peer
- AS exports a peer's routes *only* to its customers

Ranking rule: Everyone prefers customer routes over peer and provider routes

Question:

The diagrams below show 6 topologies. In it, a peer-peer relationship is denoted by an undirected edge. Customer-provider relationships are denoted by a directed edge going from customer to provider. Circles represent autonomous systems (ASes).

In the questions to the right, for each network topology shown, answer true if there exists AS preferences and an order of updates such that the routes oscillate forever (aka route flapping)

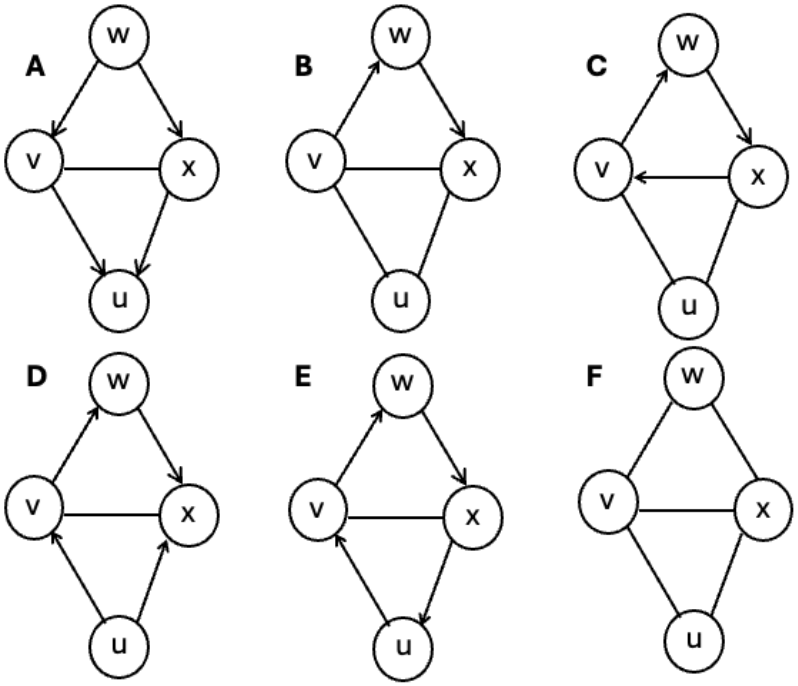

6 True or False 1 point

Network B

- ○ True
- ● False

7 True or False 1 point

Network C

- ● True
- ○ False

Here is the content of the above diagrams in textual form:

All topologies have the same structure and the same vertices: u, v, w, x. Only the directedness is different. Directed edges are denoted as "x > y" denoting an edge going from x to y. "x < y" denotes an edge going the other way. Undirected edges are denoted as "x – y". The topologies are given below

Network A:

w > v

w > x

v – x

v > u

x > u

Network B:

w < v

w > x

v - x

v - u

x - u

Network C :

w < v

w > x

v < x

v - u

8 True or False 1 point

Network D

- True
- False

9 True or False 1 point

Network E

- True
- False

x - u

Network D:

w < v

w > x

v - x

v < u

x < u

Network E

w < v

w > x

v - x

v < u

x > u

Network F

w - v

w - x

v - x

v - u

x - u

10 True or False 1 point

Network F

- ○ True
- ● False

11 Multiple Choice 1 point

Suppose an AS W has a route X to a provider. Typically, ASes will not export these routes to their peers. Why?

- ○ Doing so can cause routing loops
- ● W will not want to do it
- ○ W's peer will not trust W to send its packets
- ○ Doing so will cause route oscillation/flapping

12 Multiple Choice 1 point

Consider a router in the UT Austin campus that routes using IP addresses. Who fills its routing table entries that corresponds to IP addresses outside the campus?

- ○ A spanning tree protocol
- ● An IGP protocol such as OSPF (link state) or RIP (distance vector) shortest path, combined with BGP
- ○ It will learn the routes by observing where the packets are coming from
- ○ BGP first determines where to send it and adds the entries into all internal routers

13 Multiple Choice 1 point

Can the link state routing protocol create routing loops? A routing loop occurs when the routing tables of multiple routers are set up in a way that causes packets to go in a loop at least once

- ○ No, it can never happen
- ● Yes, but it will eventually settle into a loop free topology once links stop forming/breaking
- ○ Yes, and the loops will remain forever, but it only occurs in rare cases
- ○ No, unless there is a pair of endpoints between which there are two different paths of the same length

14 Essay 2 points

Write one reason why link state routing is better than distance vector routing, and another reason why distance vector routing is better

5 min / 100 max (word limit)

Must use at least 5 words

15 Essay 2 points

Is the spanning tree protocol optimal? Argue why this is the case by proposing a *reasonable* metric and either prove/argue convincingly that it is optimal or show a counter-example where it is not.

Any reasonable metric will work. An example of an *unreasonable* metric is one that says A is better than B even though B is better on all measures of throughput and latency.

5 min / 300 max (word limit)

Must use at least 5 words

16 Numeric 1 point

Two end points are communicating over a wired link. Say that the transmitted signal is $f(t)$ and the received signal is $f(t) + \eta(t)$, where $\eta(t)$ is noise.

We are constrained to maintain $0V \leq f(t) \leq 1V$ at all times and we know that $0 < \eta(t) < 0.125$ at all times. Further, the receiver cannot sense any voltages outside the range 0 to 1.

Suppose that we transmit information by encoding them into symbols as discussed in class. What is the maximum number of bits that we can pack into a symbol?

3