

1 Reliable Transport and Congestion Control

Your company is considering possible changes to the networking stack used on their machines, but they have some questions.

1. In the first approach, when a packet drop is detected by duplicate ACKs, CWND becomes 0.25 CWND, and when a full window's worth of packets is ACKed, the window is increased by 2MSS. Is this more aggressive than normal TCP, or less aggressive, or is this TCP-friendly?

Note: a more aggressive protocol will, in the long run, take more bandwidth than its fair share.

2. Another proposal is that the receiver only send an ACK once two packets have been received (these packets can contain the same payload). Is this reliable? Would this interoperate with existing TCP implementations?

2 TCP Handshakes

Your scores for the CS 168 midterm have just been released. You want to check your score on Gradescope and you know that you must connect to Gradescope over TCP. Unfortunately, it seems your Mac does not support TCP because Apple removed it. So you must manually manage your TCP connection.

Consider the following statements:

1. Packet has TCP ACK flag set
2. Packet has TCP SYN flag set
3. Packet has TCP FIN flag set
4. Packet has TCP RST flag set
5. Packet carries HTTP request
6. Packet carries HTTP response

You have to remember the process from establishment of the TCP connection to tearing the TCP connection down (which is initiated by the host) in order to connect to Gradescope. This process takes, in this case, seven ordered messages (without data transmission). For each of these packets, list which of the above properties apply (list all that apply).

- (a) Packet 1
- (b) Packet 2
- (c) Packet 3
- (d) Packet 4
- (e) Packet 5
- (f) Packet 6
- (g) Packet 7

3 Congestion Control

We start during Slow Start, with $CWND=1$ and $SSTHRESH$ is 50. Recall that the $CWND$ is adjusted after every new ACK as follows (in units of MSS):

- Slow-start (new) ACK: $CWND += 1$
- Congestion avoidance (new) ACK: $CWND += \frac{1}{\lfloor CWND \rfloor}$
- Fast-recovery duplicate ACK: $CWND += 1$
- On timeout: $SSTHRESH = \lfloor \frac{CWND}{2} \rfloor$, $CWND = 1$
- On 3 dupACKs: $SSTHRESH = \lfloor \frac{CWND}{2} \rfloor$, $CWND = SSTHRESH + 3$

State what $CWND$ is after each event, in terms of MSS (maximum segment size). Assume that every event maintains the state from prior events.

1. 9 ACKs are received. $CWND =$
2. 8 ACKs are received. $CWND =$
3. A timeout occurs. $CWND =$
4. 45 ACKs are received. $CWND =$
5. 10 ACKs are received. $CWND =$
6. 9 dupACKs are received. $CWND =$
7. 35 (new) ACKs are received. $CWND =$
8. A timeout occurs. $CWND =$
9. 8 ACKs are received. $CWND =$
10. 2 dupACKs are received. $CWND =$
11. 4 ACKs are received. $CWND =$

4 Fair Queueing


Consider four flows traversing a single link. The four flows pass through a router implementing Fair Queueing before they reach the link. At their sources, the four flows are sending at the following rates:

- Flow A: 3 Mbps
- Flow B: 4 Mbps
- Flow C: 5 Mbps
- Flow D: 6 Mbps

There are no other flows on the link.

1. If the link has a capacity of 15 Mbps, how much bandwidth does each flow get after passing through the router?
2. If the link has a capacity of 17 Mbps, how much bandwidth does each flow get after passing through the router?
3. Now continue assume a capacity of 17 Mbps, but now the router drops packets from flow i with probability p_i instead of using fair queueing. What values of p_i should it use to achieve the same (fair) effective rates as from the previous part?

5 ARP



Consider the topology above. ARP caches start off empty. Circles refer to hosts, whereas Cylinders refer to routers. Everyone knows each other's IP addresses from previous DNS lookups. Assume switches learn promiscuously from broadcasts and other packets it receives.

1. Host A_1 sends a message to host A_3 . Which entities see the ARP request?
2. A_3 responds. Who knows its MAC address?
3. Host B_1 sends a message to A_3 . Whose MAC address does it need first?
4. When R_A gets the packet from the last part, does it need to send an ARP request?
5. Host A_3 sends a message to host B_2 . Whose MAC does A_3 need first, and does it have it?
6. After B_2 gets the packet, which end hosts does R_B have in its ARP cache?
7. At this point, which end hosts does R_A have in its ARP cache?

6 DHCP


Suppose we have an L2 network with D DHCP servers and H hosts (none of whom have an IP address yet). Each host then goes through the DHCP protocol to be assigned an IP address.

Assume no packets are sent besides DHCP messages and everything works correctly (no failures).

Consider the format of each DHCP packet sent. How many packets will the network see with the following formats?

1. Packets with source IP 0.0.0.0 and destination IP 255.255.255.255
2. Packets with source IP X and destination IP Y , where X is an offered IP address and Y is any DHCP server's IP.
3. Packets with source IP 0.0.0.0 and destination IP Y , where Y is any DHCP server's IP.
4. Packets with source IP X and destination IP 255.255.255.255 where X is the IP of any DHCP server.
5. Packets with source MAC X and destination MAC FF:FF:FF:FF:FF:FF, where X is the MAC of any host.
6. How many ARP requests were sent?

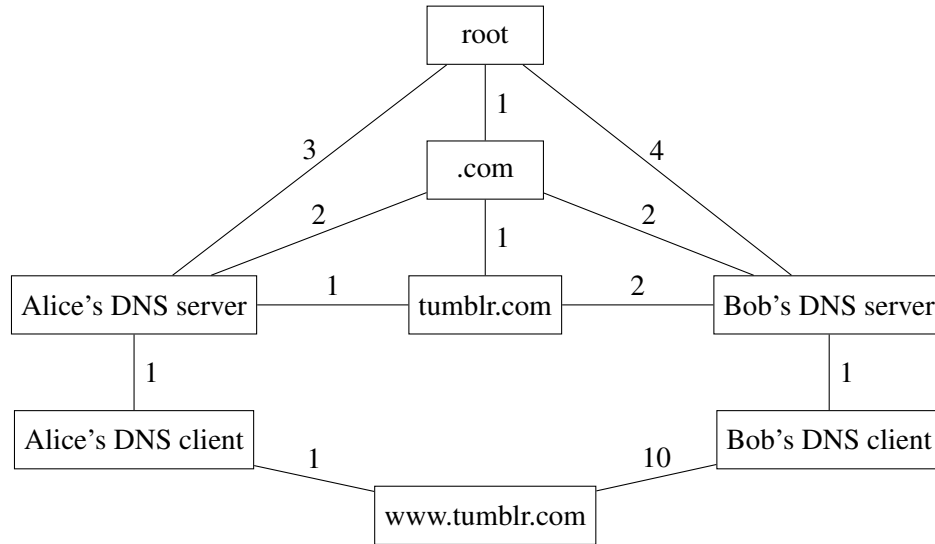
Suppose the network is connected with L2 switches as shown below:



Assume all H hosts are attached to switch A and all DHCP servers are attached to switch D. If there is not enough information to answer the question, please indicate.

1. How many packets would Switch B see?
2. How many packets would Switch G see?

7 DNS



Consider two siblings, Alice and Bob, who want to access `www.tumblr.com` from their web browser. Alice is in the US, while Bob is in China. The latencies between nodes are the edge weights on the topology above. Neither siblings' resolving DNS server caches any data.

“.com” is the authoritative server for the TLD .com, and the “tumblr.com” is the authoritative server for tumblr.com.

1. What kind of DNS record does Alice and Bob need to ask for?
2. What kind of DNS record(s) will the root server return?
3. If iterative DNS resolution is being used, what is the total latency of a DNS request for each sibling?
4. Suppose there is only one Tumblr server. Based on the latencies given in the topology, are Alice and Bob being directed to a server in China or the United States?
5. Seeing as the latency to `www.tumblr.com`'s server hurts Bob's performance, what could be done to improve his user experience? Assume that DNS servers can determine the geographical region of a client, using IP geolocation.