INST346 HW04

Congestion Control, Routers, and the Internet

Cody Buntain

October 18, 2017

1 TCP and Congestion Control

1. Consider sending a large file from host A to host B over a TCP connection that has no loss.

   a) Suppose these hosts use additive-increase, multiplicative decrease (AIMD) for their congestion control. Host A has already transitioned to congestion avoidance mode. Assume the RTT values are approximately constant. At time $t = 0$, Host A's `cwnd = 6` MSS. In terms of RTT, how long does it take for Host A to increase its `cwnd` to 12 MSS?

   b) What is the average throughput (in terms of MSS and RTT) for this connection from $t = 0$ up through $t = 6$ RTT?

   c) After sending all its data after $t = 6$ RTT, host A's timer expires while waiting for the next acknowledgement. To what multiple of the MSS will host A set its congestion window after this expiration?

   d) This timer expiration transitioned host A back to slow start mode. At what size of congestion window will host A transition back to congestion avoidance mode given its congestion window when it timed out?

2. At time t, a TCP connection has a congestion window of 4,000 bytes. The MSS used by the connection is 1000 bytes. Suppose there is one ACK per packet.

   a) If the connection is currently in **congestion avoidance** mode, how large will the congestion window be after it sends out 4 packets and receives ACKs for all of them?

   b) If the connection is currently in **slow-start** mode, how large will the congestion window be after it sends out 4 packets and receives ACKs for all of them?

2 Routers in the Network Layer

1. Describe two differences between consumer routers you may find in a residential home and a router in the network core.

2. List and describe **three** hardware sources of a queuing delay in routers.

3. In a router with many ports, why is it desirable for the router's switching fabric to be *much faster* than the bandwidth of the input ports?

4. In a network with with many users consuming a significant portion of the bandwidth, describe a possible packet scheduling mechanism that could be used to give priority to a single server with a known IP address when forwarding across the router.

5. Describe a scenario in which packet scheduling mechanisms and/or router queue management might be used to violate net neutrality.

6. Assume your router at home is using networking address translation to provide Internet access to several machines.

   a) How many simultaneous connections can **one host** behind your NAT router maintain?

   b) Describe the NAT traversal problem and how it might affect a peer-to-peer application running on your machine behind the NAT router.

3 The Internet Protocol

1. According to ARIN, the University of Maryland owns the IP block from 128.8.0.0 - 128.8.255.255.
   a) How many different IP addresses does this block contain?
   b) How is this subnet described using CIDR?
   c) What would the netmask of this subnet be?
2. A possible IP address in the University of Maryland is 129.2.101.23.
   a) Convert this IP address in dotted-decimal notation to a 32-bit number in base-10.
   b) List **seven** different CIDR-ized subnets to which this IP address could belong.
   c) For each of the **seven** different CIDR-ized subnets you listed, provide their netmasks in dotted-decimal notation.
3. Consider a router with a forwarding table that contains both (prefix, port) pairs (129.2.0.0/16, 9) and (129.2.100.0/23, 1). Assume this router is configured to forward broadcast packets.
   a) Which prefixes will the IP address 129.2.101.23 match?
   b) If a packet comes into this router with a destination address of 129.2.101.23, to which port will the packet be forwarded?
   c) If a packet comes into this router with a destination address of 255.255.255.255, to which port will the packet be forwarded?
4. Describe two methods we covered to assign an IP address to a computer's network card.
5. If your machine's DHCP client fails to obtain an IP address from a DHCP server, describe a failover measure might it employ.
6. If a router directly connects n different subnets, how many IP addresses does that router require?
7. Why does the IP protocol header require a length field?
8. In IPv4, why do routers need to recalculate a packet's checksum whenever a packet is forwarded through the router?
9. What is the **maximum** number of routers an IP packet can traverse?