INST346 HW04

Congestion Control, Routers, and the Internet

Cody Buntain

October 30, 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?

   Solution Since we are already in congestion avoidance, we scale up linearly. So, after every RTT, we increase the congestion window by 1 MSS. Therefore, it will take 6 RTTs for `cwnd` to equal 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?

   Solution We sum the congestion window size from t_0 to t_6, so $\frac{1}{6}\sum_{j=6}^{11} j = 8.5$ MSS per 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?

   Solution In both implementations of TCP (Tahoe or Reno), TCP responds to a *timeout* by resetting the `cwnd` to 1 MSS.

   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?

   Solution TCP returns to congestion avoidance mode when it reaches or exceeds one half of the congestion window it had when it experienced some loss. Since we were at `cwnd` = 12 MSS when this timeout occurred, we transition from slow start to congestion avoidance at `cwnd` = 8 MSS (note we would expect this to be 6 MSS, but our slow-start increments by powers of 2, and 8 is the first power of 2 to exceed 6).

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?

Solution TCP only increments its congestion window after *all bytes* of the current window have been acknowledged, so after all 4 1,000-byte packets have been acknowledged, TCP will increment its congestion window one time. Since we are in congestion avoidance mode, this increment will be a linear of 1 MSS (or 1,000 bytes), so after 4 ACKs, the new congestion window will be 5,000 bytes.

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?

Solution As before, TCP only increments its congestion window after *all bytes* of the current window have been acknowledged. Since we are in slow start mode, this increment will be multiplicative and increases by a factor of 2, so after 4 ACKs, the new congestion window will be 8,000 bytes.

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.

   Solution One difference has to do with complexity. Because your router is a home device with likely only a single Internet-facing link, much of the routing algorithm complexity is removed. Instead, your home router likely has an isolated network behind it and routes all traffic from that internal network through the external, Internet-facing link. A network-core router instead has peer networks on other side (and likely more than two), and devices on either side of a network-core router can reach other as peers. Internal networks for at-home routers instead isolate their internal networks from the rest of the Internet.

   Another difference is with how devices on other side of the at-home routers are addressable. For your consumer-grade at-home router, it likely uses network address translation (NAT) to share one forward-, Internet-facing IP address with all the devices behind it. Network core routers do not do this NAT and instead present all devices equally, with their own IP addresses addressable from either side of the router.

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

   Solution One source of delay is the time needed to transmit packets from the router's output queue onto the transmission medium. Another source of delay is the time a packet may be sitting in a queue before it is forwarded across the router's switching fabric if the arrival rate is higher than the switching fabric's transfer rate. In a related fashion, a packet may also sit in the input queue waiting for packets ahead of it to be transferred to a popular output port (called head-of-line blocking).

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?

   Solution Having a much faster switching fabric is important to reduce congestion in the router's input queues. We want the switching fabric to be fast enough such that, even if a new packet arrived on each and every input port simultaneously, the fabric would be able to transmit all packets across itself before the next set of packets arrived.

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.

   Solution One possible solution is for all of a router's to be augmented with an additional set of **priority** queues, and any packet destined to or sent from the single server's IP address could be pushed to this high-priority queue instead of the regular input/output queues. Switching fabric and transmission should then pull from this high-priority queue before pulling packets from the regular queues. In this way, packets to/from this single server can effectively jump the line in transmission.

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

Solution The above answer in which an ISP gives preference to a particular IP address or block of addresses may violate net neutrality. If Netflix were to lease one of these high-priority IP address blocks for a fee, an upstart competitor may not be able to compete with Netflix because the ISP is offering Netflix customers a better experience.

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?

   Solution Since NAT tables use input ports as the mechanism to differentiate between devices inside the NAT, if a single host was behind your NAT, that host could maintain up to $2^{16} = 65,536$ simultaneous connections. As the number of devices behind the NAT increases, however, these 65,536 possible simultaneous connections must be shared among them.

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

   Solution Without port forwarding set up beforehand, a device behind a NAT cannot be contacted by a device outside the NAT without the internal device first opening a connection to the external device (otherwise, the NAT device would not know to which internal node an unsolicited packet should be sent). NAT traversal techniques are used to allow "direct" connectivity with these internally NAT devices.

   With regard to P2P applications, NAT traversal is a significant issue since an external client would not be able to contact an internal client directly. For example, if your internal machine was running BitTorrent or a similar P2P application, external peers would not be able to notify you that they want to download data from you. Similarly, if you were running a Skype-like client within your NAT, external clients could not call you directly. Instead, some external, publicly available server would be needed to coordinate the beginning of the communication.

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?

   Solution This IP block contains 16 bits of address space, so this block contains $2^{16} = 65,536$ addresses.

   b) How is this subnet described using CIDR?

   Solution The CIDR version is 128.8.0.0/16.

   c) What would the netmask of this subnet be?

   Solution The netmask can be derived from the /16, so the first 16 bits (or two octets) are all 1s. Therefore, the netmask is 255.255.0.0.

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.

   Solution 10000001 000000010 01100101 00010111 = 2,164,417,815

   b) List **seven** different CIDR-ized subnets to which this IP address could belong.

   Solution 129.2.101.23/32, 129.2.101.20/30, 129.2.101.16/28, 129.2.101.0/26, 129.2.101.0/24, 129.2.100.0/22, 129.2.96.0/20, and many others.

   c) For each of the **seven** different CIDR-ized subnets you listed, provide their netmasks in dotted-decimal notation.

   Solution
   - 129.2.101.23/32 = 255.255.255.255
   - 129.2.101.20/30 = 255.255.255.252
   - 129.2.101.16/28 = 255.255.255.240
   - 129.2.101.0/26 = 255.255.255.192
   - 129.2.101.0/24 = 255.255.255.0
   - 129.2.100.0/22 = 255.255.252.0
   - 129.2.96.0/20 = 255.255.240.0

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?

   Solution It will match both.

   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?

   Solution It will go out on port 1 because of the longest-prefix-match rule.

   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?

   Solution Since the router is configured to forward broadcast packets, this packet will be sent to both output links.

4. Describe two methods we covered to assign an IP address to a computer's network card.

   Solution The first method was static assignment, where the administrator hardcodes a specific IP address into the network interface. The second method uses DHCP, which allows the computer to request an IP address from the network's DHCP server when the computer boots up.

5. If your machine's DHCP client fails to obtain an IP address from a DHCP server, describe a failover measure might it employ.

   Solution Modern systems, when they are unable to obtain an IP address from a DHCP server and are not statically configured, fall back to a "self-assigned IP address" generally in the range of 169.254.0.0/16.

6. If a router directly connects n different subnets, how many IP addresses does that router require?

   Solution The router needs an IP address on each subnet, so it must have at least n IP addresses.

7. Why does the IP protocol header require a length field?

   Solution In IPv4, the protocol header **might** include additional option headers. As a result, the header must specify a length, so devices know where these optional headers end. In contrast, IPv6 has a pointer to the "next header" and no optional headers, so the IPv6 header is always the same length.

8. In IPv4, why do routers need to recalculate a packet's checksum whenever a packet is forwarded through the router?

   Solution Regenerating the checksum is necessary since each router changes the TTL value in the IP header. As such, the packet's content changes and its checksum must be regenerated.

9. What is the **maximum** number of routers an IP packet can traverse?

   Solution Since the TTL field is 8 bits, and no packet can pass through more than the max TTL routers, this maximum must be $2^8 = 256$ routers.