CSE 473 – *Introduction to Computer Networks* *Jon Turner*

Final Exam Solution

12/18/2013

1. (10 points). What is the queueing delay at a network link with a link rate of 100 Mb/s, an arriving traffic rate of 9,000 packets per second an average packet length of 1250 bytes and a queue length of 500 packets?

   The traffic intensity is 0.9, so the average number of packets in the queue is 9. Since the time needed to send one packet is 100 μs, the average queueing delay is 900 μs.

   What is the delay if the arriving traffic rate increases to 15,000 packets per second?

   In this case, the queue fills, so the average delay is about 50 ms.

   Consider an audio application that sends audio data using RTP and UDP over an Ethernet network. How many overhead bytes does each packet include, assuming that the RTP header is 12 bytes? You may ignore the Ethernet preamble and flag.

   The overhead is 18+20+8+12=58 bytes.

   If audio is encoded at 32 Kb/s and we want or audio packets to be large enough so that the overhead is at most 20% of the total, how many audio bytes must each packet carry?

   We need four times 58 or 232 bytes.

   How long does it take to acquire the audio data needed to "fill" a packet?

   At 32 Kb/s, we get four bytes of audio data per ms, so it takes 58 ms to fill the packet.

2. (15 points) The diagram at right shows two TCP senders at left and the corresponding receivers at right. The first sender uses TCP *Tahoe*, the second uses *Reno*. Assume that the MSS is 1 KB, that the one-way propagation delay for both connections is 50 ms and that the link joining the two routers has a bandwidth of 6 Mb/s. Let $cwnd_1$ and $cwnd_2$ be the values of the senders' congestion windows. What is the smallest value of $cwnd_1 + cwnd_2$ for which the link joining the two routers stays busy all the time?

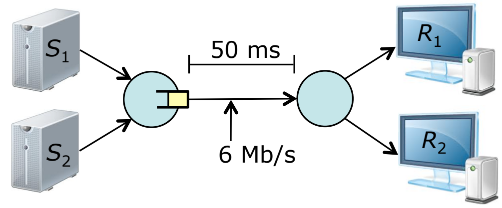

The RTT is 100 ms in this case, so the link rate is equivalent to 600 Kb per RTT or 75 KB.
So, $cwnd_1+cwnd_2= 75$ *KB.*

Assume that the link buffer overflows whenever $cwnd_1+cwnd_2 \geq 150$ KB and that at time 0, $cwnd_1=30$ KB and $cwnd_2=120$ KB. Approximately, what are the values of $cwnd_1$ and $cwnd_2$ one RTT later? Also, what are the values of *ssthresh* for each of the two connections? Assume that all losses are detected by triple duplicate ACKs.

Since the first uses Tahoe and the second uses Reno, $cwnd_1=1$ *KB and* $cwnd_2= 60$ *KB,* $ssthresh_1=15$ *KB and* $sshresh_2= 60$ *KB.*

After 8 more RTTs, approximately what are the values of $cwnd_1$ and $cwnd_2$?

19 KB and 68 KB

Approximately, how many more RTTs before $cwnd_1+cwnd_2 \geq 150$ KB again?
What is $cwnd_2-cwnd_1$ at this point?

(150–87)/2=31.5 RTTs $cwnd_2-cwnd_1=49$ *KB*

Approximately, how many more RTTs pass before $cwnd_1+cwnd_2 \geq 150$ KB and $cwnd_2-cwnd_1 \leq 20$ KB?

The difference goes down by a little less than a factor of 2 after every cycle of the congestion control algorithm. So, 2 more cycles will be required to get the difference below 10 KB.
So it will take about 2(1+8+31.5)=81 RTTs.*

3. (10 points). The figure at right shows two residential networks with routers that implement NAT. Suppose host *A* is connected to the web server at host *E*.

internal (IP,Port)	external port
10.1.1.5,1111	2222
10.1.1.3,7777	3333

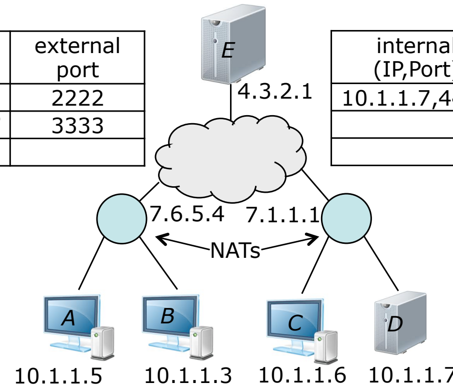

internal (IP,Port)	external port
10.1.1.7,4444	5555

In the left-hand NAT table, add an entry that would allow *A* to communicate with *E*. You may choose any port numbers you like, but the internal port numbers should be different from the external port numbers.

Show the values of the address and port fields in the diagram below, for a typical packet sent by host *A*.

src adr	dest adr	src port	dest port
10.1.1.5	4.3.2.1	7777	80

Show the fields in the packet as it might appear when it reaches *E*.

src adr	dest adr	src port	dest port
7.6.5.4	4.3.2.1	2222	80

Suppose the user in the right-hand network runs a game server on host *D* and invites her friends to join her game sessions. Add an entry to the right-hand table that would allow remote connections to the game server. Again, you may pick your own port numbers, but the internal and external port numbers should be different. Assume host *B* connects to the game server at *D*. Add an entry to the left-hand NAT table for this connection. Show the address and port fields for a typical packet leaving host *B*, the fields in the same packet as it passes through the public internet, and the fields in the packet that is delivered to *D*.

src adr	dest adr	src port	dest port
10.1.1.3	7.1.1.1	7777	5555

src adr	dest adr	src port	dest port
7.6.5.4	7.1.1.1	3333	5555

src adr	dest adr	src port	dest port
7.6.5.4	10.1.1.7	3333	4444

4. (15 points) Does SSL use public key encryption? If so, what does it use it for and how does it use it, in a typical situation?

Yes. Typically a server's public key is signed using a certificate authority's private key, so that a client can use its copy of the CA's public key to verify that that the key in the certificate really belongs to the server. The client then uses the server's public key to encrypt some secret information, which it sends to the server. This information is then used to generate additional keys.

Does SSL use symmetric encryption? If so, what does it use it for and how does it use it?

Yes. After the initial handshake, the client and server encrypt data using symmetric encryption.

How many different keys are created for a typical SSL session? What are they used for?

Four keys are created, two encryption keys (one for each direction) and two MAC keys (one for each direction.

Is it necessary for a single SSL record to be sent entirely within one IP packet? Why or why not?

It is not. The TCP layer does not know about SSL records, it just sends bytes. So the IP packet boundaries may come in the middle of SSL records. This does no harm, as SSL is unaware of IP packet boundaries and simply operates on bytes.

Suppose an attacker learned the encryption key used by one endpoint of an SSL connection. Could the attacker use this to modify one of the records in the connection, without this being detected by the receiver? If so how, if not, why not?

No. While the attacker could encrypt the packet, without also knowing the MAC key, it could not compute a new MAC to go with the modified record. Consequently, the receiver would be able to detect the modification to the record.

5. (15 points) The diagram at right shows a corporate network with three sites connected by the internet. The company's n etwork administrator has configured the gateway routers at the three sites to use IPsec to encrypt all traffic going between the three sites. How many Security Associations are needed to allow the sites to communicate securely with each other?

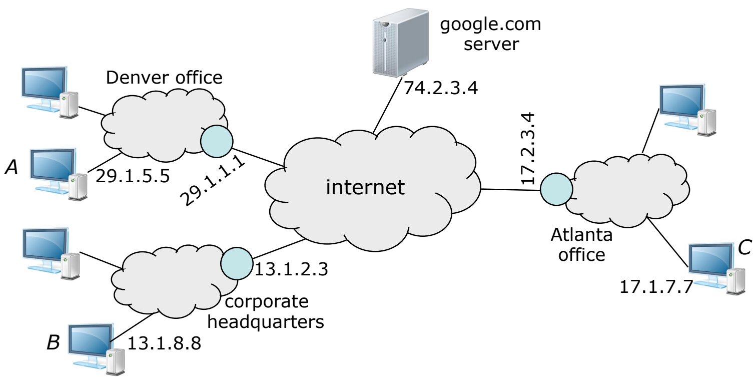

Six SAs are needed, two for each pair of offices (one in each direction).

If host *A* sends a packet to host *C*, what are the source and destination address fields in the packet header as it passes through the public internet?

Source address is 29.1.1.1, destination address is 17.2.3.4.

If host *B* sends a packet to the Google server shown, what are the source and destination address fields in the packet header as it passes through the public internet?

Source address is 13.1.8.8, destination address is 74.2.3.4.

How does a gateway router decide if an outgoing packet should be encrypted using IPsec? How does it decide which Security Association to use, if it does require encryption?

It looks for a matching entry in its Security Policy Database (SPD), using the source and destination address fields in the packet as part of the lookup key. If there is a matching entry, the packet should be encrypted and the entry will include the Security Parameter Index (SPI) of the appropriate SA.

If the gateway router at Atlanta receives an IPsec packet from the public Internet, how does it determine which encryption key to use when decrypting the packet?

The packet will include the SPI of the Security Association in cleartext. The receiving router will lookup the entry for that SPI in its Security Association Database. This contains the encryption key for the SA.

6. (10 points) In the diagram at right, the nodes represent IP routers and the numbers on the links are OSPF link weights. If the network uses PIM with reverse-path f orwarding, which links are never used for multicast packets going to a host connected to router *J*?

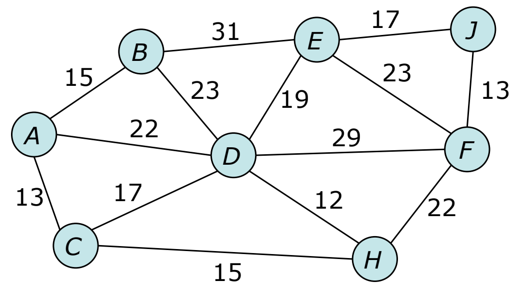

The shortest path tree rooted at J includes the links EJ, FJ, BE, DE, CH, AD and FH. So the links that would never be used for multicast packets going to a host at J are: EF, DF, DH, BD, AB, AC, CD.

Suppose that router *C* is the PIM DR for some layer 2 subnet, and that a host *X* in that subnet sends an IGMP report for address 229.1.2.3. Assume that *J* is the RP for 229.1.2.3 and that no other routers are currently participating in 229.1.2.3. What does router *C* do at this point? Which routers eventually add forwarding state for 229.1.2.3 as a result of *C*'s action?

C would send a PIM Join packet to H. Eventually, routers C, H, F and J would add forwarding state for 229.1.2.3.

If a host connected to router *B* sent a packet to 229.1.2.3, which routers would the packet pass through in order to reach the host connected to the subnet at *C*? List them in the order in which the packet passes through them.

It would pass through routers B, E, J, F, H and C.

How would the previous answer change if router *C* issued a source-specific join on 229.1.2.3 for the host at router *B*.

This would cause packets from B to pass through routers B. A and C.

If the host connected to the subnet at *C* sends a packet to 229.1.2.3, which routers receive a copy? Assume that there has been no other activity on this multicast address, except for what has been described in the earlier parts.

Routers C, H, F and J would all receive copies. J would simply discard its copy.

7. (10 points) Consider an audio/video teleconferencing session with ten participants in the same switched layer 2 network. The application limits the number of simultaneous senders to two. If the RTP session bandwidth is 4 Mb/s, how much bandwidth can be used for RTCP packets, assuming the default behavior?

   5% of the session bandwidth or 200 Kb/s.

   Suppose that all participants send RTCP receiver reports and that at some point in time, two hosts are sending RTCP sender reports. How much bandwidth does each receiver have for its reports? How much does each sender have for its reports?

   The RTCP bandwidth is split between senders and receivers with senders getting 25% of 200 Kb/s or 50 Kb/s. So each sender gets 25 Kb/s. Each receiver gets 15 Kb/s.

   Suppose the number of participants in the conference increases to 20. In this case, how much bandwidth is available to each sender for its reports? How many for each receiver?

   The senders still get 25 Kb/s each. The receivers get 7.5 Kb/s each.

   When a participant in the conference receives an audio packet, how does it determine the actual time when the packet was sent, so that it can determine the appropriate playout delay?

   The audio packet contains a timestamp that is based on the audio sample clock. This is essentially just just the value of an audio sample counter maintained by the source. To relate this to real-time, the receiver uses information it received from the last sender report issued by the audio source. This sender report relates the audio sample clock to the actual real-time.

8. (10 points) Consider a residential network that connects to the internet with a DSL link that has a download rate of 4 Mb/s. Assume that there are three UDP flows sharing the link and the remote hosts are sending at rates of 1 Mb/s, 2 Mb/s and 3 Mb/s. Assume that the ISP router has a link buffer that can hold 300 packets (assume all packets have the same length). For each flow, what fraction of the packets it sends are discarded?

   They each lose about 33% of the packets they send.

   For each flow, about how many packets does it have in the queue.

   The first flow has about 50 packets in the queue, the second flow has about 100 and the third has 150.

   Now, suppose the queue at the ISP router is replaced by three queues that can each hold 100 packets and that the queues are scheduled using weighted-fair queueing, where the weights are all 0.33. In this case, what fraction of packets are discarded from each flow?

   In this case, the first flow loses 0, the second loses 25% and the third loses 50%.

   How many packets does each flow have in the queue?

   The first flow's queue is empty or close to empty. The other two each have about 100 packets in their queues.

   Now, suppose the weights are 0.2 for the first flow, 0.6 for the second and 0.2 for the third. In this case, what fraction of packets are discarded from each flow?

   In this case, the first two flows each lose 0, while the third loses 67%.

   How many packets does each flow have in the queue?

   The second flow's queue is empty or close to empty. The third has about 100 packets in its queue.

   The first flow is a little tricky, as its allocated portion of the link bandwidth exactly matches its data rate. If the packets arrive with uniform spacing, the queue will remain close to empty, but if the times between packets are more random, significant queueing can occur. Indeed, there can even be some packet loss in this case. For a single flow however, it's more likely that the queue will never accumulate a large backlog.

9. (10 points) The diagram at right shows a WIFI network with an access point, *X* and three hosts, *A*, *B* and *C*. The large circles indicate the *coverage areas* of the three hosts. The coverage area for *X* is not shown, but you may assume that it includes all three hosts. Assume RTS/CTS are not used.

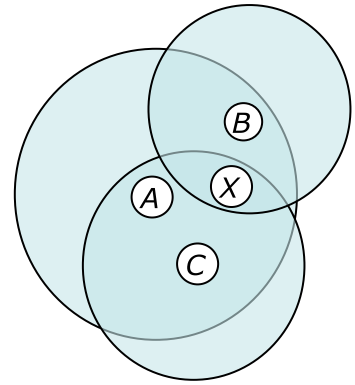

Suppose *X* is transmitting a packet at time 0 and finishes sending it at time 100 µs. Also,

- *A* gets a packet to send at time 50 that takes 100 µs to send and is assigned a backoff timer of 100.
- *B* gets a packet at time 70 that takes 200 µs and is assigned a backoff timer of 50.
- *C* gets a packet at time 90 that takes 150 µs and is assigned a backoff timer of 150.

For each of the three hosts, what time do they start sending their packets? You may ignore the inter-frame spacing and the time required for acks.

A starts sending at time 200 and finishes at 300
B starts sending at 150 and finishes at 350
C starts sending at 350 and finishes at 500

Of the three packets sent, which are successfully delivered on the first attempt?

Only the one from C is successfully delivered.

For each packet that is not successfully delivered on the first attempt, approximately when does the sending host learn that the packet was lost and must be sent again?

Hosts learn of lost packets from the absence of ACKs. Here, A would expect an ACK at 300 and so would learn of the lost packet when the ACK fails to arrive at time 300. Similarly, B would learn of its lost packet at time 350.

Now, suppose RTS/CTS is enabled. In this case, approximately when does each host send its data packet? You may assume that the time needed to send RTS, CTS and ACK packets is negligible.

A sends at 400.
B sends at 150.
C sends at 550.