


Final Exam Review Questions

12/2013

1. The figure below shows a network path connecting a server to a client.




What is the propagation delay for a packet going from the server to the client (you may assume that the speed of light is 200,000 km/s)?

What is the total transmission delay of a 10,000 bit packet on all of the links?

What is the average queueing delay at router X , assuming that the traffic intensity is 1.3, and the buffer can hold 10,000 packets and that the average packet size is 5,000 bits?

What is the average queueing delay at router Y , assuming that the traffic intensity is 0.8, and the buffer can hold 100 packets and the average packet size is 5,000 bits?

2. The diagram at right shows two TCP senders that are transferring large files to a pair of receivers. Both senders use TCP *Tahoe*. Assume that the MSS is 1 KB, that the one-way propagation delay for both connections is 100 ms and that the link joining the two routers has a bandwidth of 4 Mb/s. Let $cwnd_1$ and $cwnd_2$ be the values of the senders' congestion windows and assume that $cwnd_1 = cwnd_2$. What is the smallest value of $cwnd_i$ for which the link joining the two routers stays busy all the time?




Assume that the link buffer overflows whenever $cwnd_1 + cwnd_2 \geq 200$ KB and that at time 0, $cwnd_1 = 40$ KB and $cwnd_2 = 160$ 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.

After 7 more RTTs, approximately what are the values of $cwnd_1$ and $cwnd_2$?

Approximately, how many more RTTs before $cwnd_1 + cwnd_2 \geq 200$ KB again? What is $cwnd_2 - cwnd_1$ at this point?

Approximately, how many more RTTs pass before $cwnd_1 + cwnd_2 \geq 200$ KB and $cwnd_2 - cwnd_1 \leq 10$ KB?

3. The figure below shows a portion of the internet with four autonomous systems, one of which is shown with four routers, each with its own /16 subnet. Note that each of the other ASs has a /8 subnet. The blank routing table at right is for router C. Complete the table, providing an entry for each of the seven subnets. Assume that AS1 uses OSPF as its intra-AS routing algorithm, and BGP as its inter-AS routing algorithm. Use the standard tie- breaking rules and assume there are no policy constraints that must be satisfied.



prefix	output link

Suppose the link joining router B with AS3 fails, which entries would change, and how would they change?

4. Consider a simple block cipher that uses 4 bit blocks, where each block is encrypted by adding 3 to it and discarding any "overflow bits". So for example,

0010 0101 1110 becomes 0101 1000 0001

To make this more secure, we add cipher block chaining. Suppose the initial vector is 1011. What is the cipher text corresponding to the following clear text?

0101 1011 0011

5. Recall that an RSA encryption key is a pair of numbers (n,e) and the corresponding decryption key is another pair (n,d) . Which of the following could be used as RSA key pairs (ignoring the fact that they are too small)? For those that are valid key pairs, show that they meet all the requirements for a key pair, for those that are not, explain why they are not.

- $(31,5), (31,11)$

- $(77,7), (77,43)$

- $(55,7), (55,41)$

Consider the pair $(91,5), (91,29)$. Assume the first is the encryption key and the second is the decryption key. What is the encrypted value of the number 10? (*Hint: $10^2 \bmod 91=9$*)

6. The diagram below represents a portion of the payload of an IP datagram carrying data for an SSL session. Each line represents a distinct SSL record.

type	version	length	data	MAC
0	3.2	35	kasfjla;kja;ljf	2343
0	3.2	41	aiusrooqirgn/nc;p	9783
0	3.2	25	qp98hj;vn	8520
1	3.2	36	z2jt3;r'wjf2h5	2985

Which of the fields in each record are encrypted by the sender, and which are sent as cleartext.

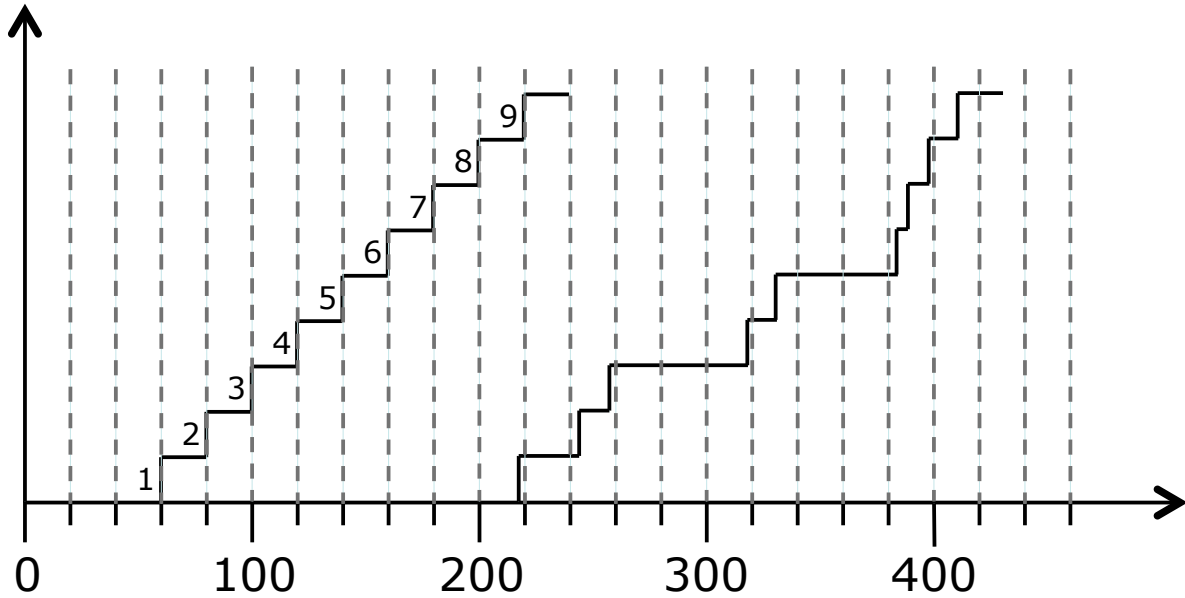
Suppose an attacker was able to intercept and modify packets as they pass through the internet. If the attacker were to modify several bytes in the data field of the second record, what changes would the attacker have to make to the IP and TCP headers, in order to ensure that these records were actually delivered to the SSL software.

How would the SSL software detect the modification to the second record?

Suppose the attacker were to make a copy of the second record and insert it after the third record (so the modified packet had five records). Explain how the SSL software at the receiver would detect this modification?

Suppose the attacker were able to discover the key used for the MAC, but not the encryption key. Could the attacker use this to change the type field of the third record, without it being detected by the SSL software at the receiver. If so, explain how this is done. If not, explain why not.

7. The chart below shows packets sent by an audio source and received at a sink. Assume that the time between packet transmissions is 20 ms and that packets are timestamped by the sender and that the receiver playout buffer uses a fixed playout delay of 200 ms. Also, assume that the sender's and receiver's clocks are exactly synchronized.



List the packet numbers for those packets that arrive too late to be played out.

What is the smallest playout delay required to ensure that packets never arrive too late for playout (for the example)? You may round up to a multiple of 20 ms.

8. Consider a VoIP system that uses a fixed playout delay of 200 ms. If a receiver receives a sequence of four packets with timestamps 17.02, 17.04, 17.06 and 17.08 (where the timestamp is expressed in seconds relative to some arbitrary starting time), when will the receiver play out these packets (you may assume that the receiver's clock is synchronized with the sender's).

Now, suppose that we are using an adaptive playout delay, instead of a fixed playout delay. After a period of silence, the receiver starts receiving a sequence of voice packets with timestamps 30.01, 30.03, 30.05 and 30.07. Let $d_i=50$ ms be the current value for the average measured delay across the network, let $v_i=5$ ms be the current value of the measured delay variation. When computing the playout time, multiply the delay variation by the parameter $K=5$. At what times are the four packets played out?

Now suppose the VoIP system does loss concealment by separating "odd" voice samples from "even" voice samples and sending them in different packets. In this scheme, samples from missing packets can be interpolated using the samples from the previous and following packets. How does this affect the required playout delay if voice packets are still sent at the rate of 50 packets per second, during a talkspurt?

9. Consider a token bucket rate controller, used to control a reserved rate flow. Assume that the token bucket has a capacity of 10 tokens and a token fill rate of 100 tokens per second, and that every packet consumes one token. If no token is available for an arriving packet, it is marked for possible discarding.

Suppose that at time 0, the token bucket is empty and the next token “arrives” at time 10 ms. If packets arrive at times 11, 13, 17, 19, 23 and 29, which packets (if any) are marked. (You may identify the marked packets by their arrival times.)

How many tokens are in the token bucket at time 55 ms if no additional packets arrive?

How many tokens are in the token bucket at time 195 ms?

What is the largest number of packets that can be sent between time 201 ms and time 299 ms without any of the packets being marked?

If this largest possible number is sent between times 201 and 299, what is the largest number of packets that can be sent between times 301 ms and 399 without any of the packets being marked?

10. Let A and B be two nodes in a wireless network with access point X . Assume that A and B cannot hear each others' transmissions, but can hear X . At time 0, X is sending a packet to some other node and it completes sending this packet at time $100 \mu\text{s}$. At time 30, a packet becomes available for transmission at A and at time 70, a packet becomes available for transmission at B . Assume that A needs $200 \mu\text{s}$ to send its packet and B needs $100 \mu\text{s}$ to send its packet. Let t_A be the initial value of the backoff timer for A and let t_B be the initial value of the backoff timer for B .

Is it possible for A and B to transmit their packets without colliding if $|t_A - t_B| = 60$? Explain. You may ignore the inter-frame spacing and the time needed for acknowledgments.

Is it possible for A and B to transmit their packets without colliding if $|t_A - t_B| = 130$? Explain.

11. Consider a user who is downloading a large file from a remote web server to her laptop in a coffee shop with WIFI access. The shop's DSL link has a download rate of 5 Mb/s and the wireless network supports a peak rate of 50 Mb/s. Assume there are no other limits on the user's download rate (no one else in the coffee shop is online and the server and internet path are lightly loaded).


What is the maximum throughput the user can expect in this case?

Approximately how often does TCP have to retransmit a packet in this scenario? Assume that the maximum segment size is 10,000 bits and that the round-trip time is 122 ms.

Now, assume that the wireless network experiences interference whenever the coffee grinder is running. During this time, it fails to deliver approximately 4% of the packets sent to the user. Approximately what throughput can you expect in this situation?

Suppose the throughput went down by a factor of 2, relative to the previous situation. Approximately what would the packet loss rate have to be, in order to explain this?

12. 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 $100 \mu\text{s}$. Also,

- A gets a packet to send at time 50 that takes $100 \mu\text{s}$ to send and is assigned a backoff timer of 70.
- B gets a packet at time 70 that takes $200 \mu\text{s}$ and is assigned a backoff timer of 200.
- C gets a packet at time 120 that takes $150 \mu\text{s}$ 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.

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

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?

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.