CSEE 4911: Computer Networks — Tuesday, May 8th, 2012

Final

Laurent Charignon, Avner May, Mathias Lécuyer (TAs) – A. Chaintreau (Instructor)

Please indicate your uni:

- This final should be answered in 180mn. It is graded on a total of 350pt. The exercices are *approximately* sorted from the easiest to the hardest. However, it might depend on what part of the class you feel the most comfortable with, so we recommend you read it entirely first.
- Please start a new page on your book when you start a new exercice, and return the exam sheet along with your blue book.

Exercise 1: General Multiple Choice Review ($12 \times 8 = 96$ pt)

For each question, indicate all the correct answers (there may be one or several). Correct answers will get full point even without justification. However, if you feel unsure, and would like partial credit, you may justify your choice **on the answer sheet!** (so be brief, no more than a couple of sentences).

1. It takes a single bit ten times longer to propagate over a 10Mb/s link than over a 100Mb/s link.
   (a) True
   (b) False
   (c) I need more information: is the link connected to a router or a switch?
   (d) Please stop asking me question, I am so tired I can't even think straight.
2. TCP. The TCP protocol uses a sliding window protocol. The window size varies because:
   (a) Routers along the route advertise a varying window size to prevent congestion.
   (b) The destination advertises a reduced window size when its buffers are congested.
   (c) The destination advertises a reduced window size when packets take a long time to reach it.
   (d) The source reduces its window size when it detects that congestion is occurring.
3. TCP. Which of the following are true statements about TCP:
   (a) A sources retransmission timeout value (RTO) is always set equal to the measured RTT. Setting RTO less than the measured RTT may lead to unnecessary retransmissions.
   (b) A sources retransmission timeout value is set to a value that increases with the variance in measured RTT values.
   (c) RTO is picked randomly when a TCP connection starts, and remains fixed until TCP reset.

4. Link layer. Which of the following are true?
   (a) An Ethernet switch can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
   (b) An Ethernet hub can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
   (c) An Ethernet network cannot detect collisions until it has computed a checksum over the frame.
   (d) The 802.11b wireless protocol incorporates a link-layer ACK not present in regular Ethernet.

5. IP. During normal IP packet forwarding by a router, which of the following packet fields are updated?
   (a) IP header Source address
   (b) IP header Destination address
   (c) IP header TTL
   (d) IP header checksum
   (e) Destination UDP address
   (f) Destination UDP port number

6. Consider a router inside one AS that receives a packet destined for a network inside another AS. To choose the next-hop router:
   (a) An intra-AS routing protocol is used
   (b) An inter-AS routing protocol (BGP) is used
   (c) Both of the above
   (d) None of the above

7. IP. During normal TCP packet forwarding by a NAT box going out from a private network, which of the following packet fields are updated?
   (a) IP header Source address
   (b) IP header Destination address
   (c) IP header TTL
   (d) IP header checksum
   (e) Source TCP port number
   (f) Destination TCP port number

8. Tunneling:
   (a) is a technique to ensure reliability when a part of the network in unreliable.
   (b) is used to allow interoperability between different networking standard.
   (c) is a possible solution used to deploy security.
   (d) is a technique used in Ethernet switch and circuit switched network in which the bits already received from a packet are transferred before the whole packet is received.
   (e) is a technique in which an IP datagram (*e.g.*, IPv6) is carried as the payload of another IP datagam.

9. In an Ethernet network, which of the following are true:

   (a) Ethernet switches (a.k.a bridges) learn addresses by looking at the destination address of packets as they pass by.

   (b) Ethernet hubs and repeaters learn addresses by looking at the addresses of packets as they pass by.

   (c) A correctly operating Ethernet switch never sends a packet to the wrong outgoing port.

   (d) Ethernet switches (a.k.a bridges) learn addresses by looking at the source address of packets as they pass by.

   (e) Collisions occur less on a switched Ethernet network because links run faster.

10. In a CSMA/CD network, let $TRANP$ be the transmission time for the minimum packet size, and $PROP$ be the propagation time between the two farthest nodes. Then, to detect collisions unambiguously, the following is a sufficient condition:

    (a) $TRANP > PROP$

    (b) $TRANP > 2PROP$

    (c) $\frac{TRANP - PROP}{PROP} \leq 1$

    (d) None of the above

11. BGP:

    (a) is a Distance Vector protocol

    (b) always ends up choosing the shortest path to a node

    (c) both of the above

    (d) none of the above

12. To improve the data-rate to your home, you connect your computer using separate links to ten different service providers! To balance the traffic you have a device that will randomly balance transmitted and received packets over all the links. In other words, for each packet that you send or receive there is an equal chance that it is sent over any of the ten links.

    (a) TCP performance is likely to decrease because retranmissions might be triggered even without packet loss.

    (b) Traceroute packets to your computer would always come over the link associated with the smallest IP address.

    (c) Your computer only needs one IP address.

    (d) It is impossible to increase the download speed of this computer for any application using the above setup.

    (e) Having multiple links would increase the risk of packet fragmentation, even if all the links have the same MTU.

Exercise 2: Prefix (24 pt)

As an ISP, you have been able to purchase 6 prefixes with 24 bits: 79.59.179.0/24, 79.59.180.0/24, 79.59.181.0/24, 79.59.182.0/24, 79.59.183.0/24, and 79.59.184.0/24.

1. To advertise your addresses more efficiently, you would like to aggregate those prefixes in the smallest number of prefixes. What is the smallest number of prefixes that you can use to do so and what are these prefixes?

Exercise 3: Prefix and forwarding table (30 pt)

You operate a very simple ISP that has two customers and one router with 3 ports: port 1 connects to Columbia, port 2 connects to NYU, and port 3 connects to the rest of the Internet. For historical reasons, NYU is assigned the prefix 79.128.72/25 and Columbia receives all remaining addresses in the prefix 79.128.64/18.

1. What is the range of binary addresses assigned to Columbia? to NYU?

2. What prefix(es) does the router advertise to the rest of the Internet for Columbia and NYU to be reachable?

3. Draw and describe the content of the forwarding table in the router.

Exercise 4: Dijkstra's algorithm (25 pt)

We consider the following network, and, as a reminder, a table showing how to present the steps of Dijkstra's algorithm (on a complete different network):

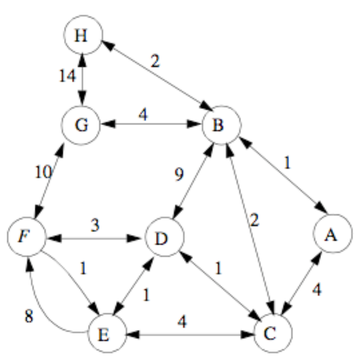

Step	N'	D(v) p(v)	D(w) p(w)	D(x) p(x)	D(y) p(y)	D(z) p(z)
0	u	7,u	3,u	5,u	∞	∞
1	uw	6,w		5,u	11,w	∞
2	uwx	6,w			11,w	14,x
3	uwxv				10,y	14,x
4	uwxvy					12,y
5	uwxvyz					

1. Using Dijkstra's algorithm, compute the shortest-path tree from router F to every other router in the network. Present the result in a table where each line denotes a step of the algorithm, and the current estimation for all nodes that remain to be treated by the algorithm, as shown in the example.

Exercise 5: Hubs and Switches (25 pt)

We consider the following Local Area Network containing 3 local Ethernet Switches (S2,S3,S4) and 1 Central Ethernet Switch (S1), as well as various end-hosts (9 clients and 2 servers). We assume that all links that appear on the Figure are 50Mbps full duplex (they can be used at this rate in both directions).

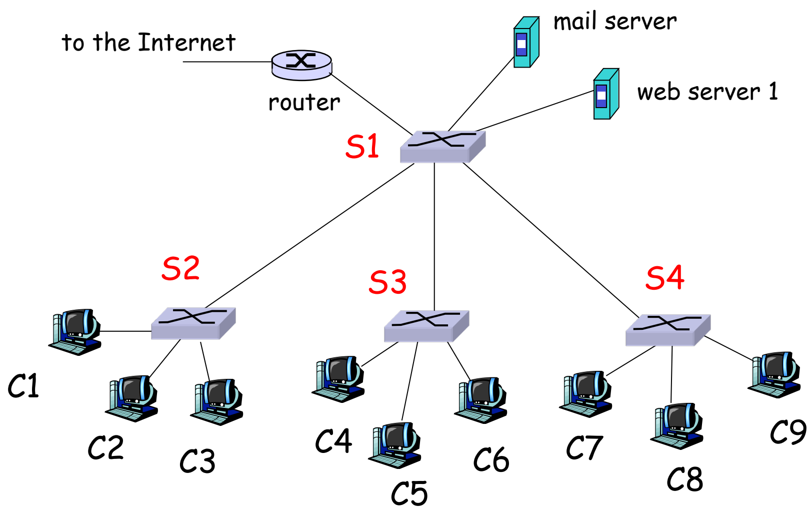

1. Assuming that you can choose arbitrarily to whom each host is sending (either to a server or a client on this LAN or even to another machine on the Internet), what is the maximum amount of data that can simultaneously be sent from all the hosts (clients and servers)?

2. Same question if the local switch S2, S3 and S4 are replaced by a hub.

Exercise 6: Answering Email request to SSOL (40 pt)

You are operating a support service for SSOL through email. Through experiences you made the following observations:

- Students typically ask questions totally independently of each other. In particular the chance to receive an email at anytime during the day is the same.
- You receive on average 120 request during an hour.

You would like to know what is the chance that you get a burst of request arriving in the queue, which would create delay and stress in your staff.

1. Through some observations, you notice that receiving more than 3 requests in a minute is a psychological threshold that essentially makes things worse.

   What is the probability that, in a given minute, there is at least 3 request arriving in the queue to request support? (Hint: It is recommended that you first derive the probability that exactly k requests arrive during this minute.)

2. You introduce an automatic response that is received by all request before they arrive to the queue. This automatic response answers a few common questions, and asks for confirmation to request additional support. You notice that for every request, there is a 2/3 probability that it is handled through this automatic response and hence never actually arrives in the queue. This event, that occurs with probability 2/3, is independent from everything else. If the request is not handled this way, the student immediately click on the confirmation and the request arrives in the queue.

   What is the probability that, in a given minute, there is at least 3 requests that are confirmed and arrive in the queue? (Hint: There are at least two ways to answer this question; one of them is very short and the other is very long. It means that if you end up with very complicated formulae you should probably come back and consider a simpler way to derive the result).

Exercise 7: Collision with partial deployment of carrier sensing (50 pt)

Consider 4 nodes pictured above, where A, B and C each wish to transmit to D. All solid edges are the same length, hence nodes separated by a solid edge are the same distance from one another, and the propagation delay between them is τ. The only exception is the pair of nodes A and C, who are separated by distance $c\tau$ where $c = \sqrt{3}$.

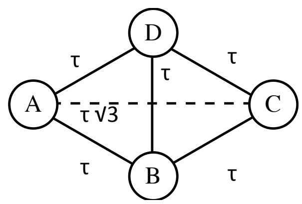

All nodes transmit frames whose transmission time L on the medium satisfies $L > \sqrt{3}\tau$. Prior to transmitting, each device runs a backoff timer that is exponentially distributed, with devices A and C using rate λ and device B using rate μ.

1. Suppose device B uses carrier sensing whereas devices A and C do not. When device B's timer expires, if it senses another device transmitting, it resets its timer and backs off again. A and C always transmit when their timers expire. When device B transmits a frame, what is the probability that its transmission is successful (*i.e.*, the probability that A and C's transmissions do not collide with B's transmission attempt)?

2. In the above setting where B is still the only device utilizing carrier sensing, what is the probability that A's transmission is successful (*i.e.*, the probability that B and C's transmissions do not collide with A's transmission attempt)?

3. Suppose in this part that all three devices use carrier sensing, $\lambda = \mu$, and C is transmitting a frame. Because of the different distance B will sense C terminating its transmission $(\sqrt{3} - 1)\tau$ time units before A. Does this give B an advantage (*i.e.*, does it mean that B has a larger probability to use the medium than A)? Explain your answer.

Exercise 8: 3 Properties of Bellman Ford Algorithm (60 pt)

In this exercise, we consider a graph G with weighted edges in which shortest paths are computed using the Distance Vector algorithm.

Optimality rule. Shortest paths are composed of shortest paths.

1. More precisely, if you consider a shortest path (*i.e.*, a path with minimum cost) $P1$ from A to B. If you assume that P1 goes through an intermediate node C, then the path from A to C used in $P1$ forms a shortest path from A to C. Prove this assertion.

Speed of convergence.

2. Assume the algorithm runs in a synchronous round-by-round manner: a node whose shortest path information changes in a round i notifies all of its neighbors of the change in round $i+1$. Consider a node A whose shortest path to a node B goes through k hops, and no path from A to B with fewer than k hops is as short. Prove that if all nodes initially overestimate their shortest path distance to B in the 0_th round, then A learns its shortest path distance to B during the $(k-1)$_th round. (Hint: you may start with small values of k).

Distance vector with multiple metrics. Consider a network where each edge e is annotated with two weights, w_e and v_e, that represents each a possible "cost" metric. We assume a priori no correlation between the value of w_e and the v_e, it is possible that the edge with a large w_e has a small value v_e.

The Distance Vector protocol is implemented to compute routes to each destination, with a small enhancement to account for the two costs metrics. When a node sends to a neighbor its distance vector, the vector contains two distances per destination: the distance of the path using weights w along the edges and the distance of the path using weights v. Upon receiving this vector, the node updates its table accordingly, with two entries for each neighbor/destination, as shown in the example below.

	A	C	D
A	5,2	4,6	7,10
B	4,9	8,3	9,4
C	2,6	5,2	8,8
D	7,10	6,3	5,1

For instance, in this example, the entry (4,9) in column A, row B indicates that according to weight w, the distance to B through A will be 4 and according to weight v, this distance will be 9.

Suppose the node is required to choose a single next hop for each destination. There are many possible approaches for choosing the next hop. Below, three are considered.

RULE A All nodes select the next hop that minimizes the distance along the path using the v weight.

RULE B Some nodes select the next hop that minimizes the distance along the path using the v weight, while others use the w weight.

RULE C All nodes select the next hop that minimizes the distance along the sum of the v weights and the w weights.

For instance, given the above table, if this node N were to choose its next hop on the path to B using weight function w, it would choose node A (with distance 4). Using weight function v, it would choose node C as the next hop, and using v + w, it would also choose C (total weight 11).

3. As a general motivation for this problem, can you propose at least four metrics that can be considered in reality to measure the cost of a link in a path?

4. For each of the rule A-B-C above, justify either with a counter-example that forwarding loops may occur, or explain using a couple of sentences why all destinations remain reachable?