Name_______________ Student ID____________________ Department/Year__________

Final Examination

Introduction to Computer Networks
Class#: 901 E31110
Fall 2015

16:00-17:40 Wednesday
December 23, 2015

Prohibited

1. You are not allowed to write down the answers using pencils. Use only black- or blue-inked pens.
2. You are not allowed to read books or any references not on the question sheets.
3. You are not allowed to use calculators or electronic devices in any form.
4. You are not allowed to use extra sheets of papers.
5. You are not allowed to have any oral, visual, gesture exchange about the exam questions or answers during the exam.

Cautions

1. Check if you get 18 pages (including this title page), 8 questions.
2. Write your **name in Chinese**, student ID, and department/year down on top of the first page.
3. There are in total 100 points to earn. You have 100 minutes to answer the questions. Skim through all questions and start from the questions you are more confident with.
4. Use only English to answer the questions. Misspelling and grammar errors will be tolerated, but you want to make sure with those errors your answers will still make sense.
5. If you have any extra-exam emergency or problem regarding the exam questions, raise your hand quietly. The exam administrator will approach you and deal with the problem.

1. (Forwarding) A major functionality of a router is to forward a packet out on the right network interface so the packet can continue on and eventually reach its destination. Suppose the packets destined to the following address ranges should go to the corresponding network interfaces in a router.

Destination Address Range	Network Interface
11001000 00010111 00011000 00000000 through 11001000 00010111 00011000 11111111	0
11001000 00010111 00011010 00000000 through 11001000 00010111 00011011 11111111	1
11001000 00010111 00011001 00000000 through 11001000 00010111 00011001 11111111	2
otherwise	3

Based on how longest match works, fill in below the destination prefixes for network interface 0, 1, 2, and 3 in the forwarding table. (8%)

Sample Solution:

Destination Prefix	Network Interface
11001000 00010111 00011000	0
11001000 00010111 0001101	1
11001000 00010111 00011001	2
default	3

2. (Routing) Based on the algorithm of Distance Vector (DV) routing, fill in the blank DV tables below.

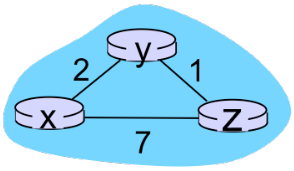

(a) Initialize the DV tables at node x, y and z at t0. (2%)
(b) Develop the DV table at node x at t1, after receiving the DV tables from y and z. (2%)
(c) Develop the DV table at node y at t1, after receiving the DV tables from x and z. (2%)
(d) Develop the DV table at node z at t1, after receiving the DV tables from x and y. (2%)

Sample Solution:

node x table

t_0 (cost to):

from	x	y	z
x	0	2	7
y	∞	∞	∞
z	∞	∞	∞

t_1 (cost to):

from	x	y	z
x	0	2	3
y	2	0	1
z	7	1	0

node y table

t_0 (cost to):

from	x	y	z
x	∞	∞	∞
y	2	0	1
z	∞	∞	∞

t_1 (cost to):

from	x	y	z
x	0	2	7
y	2	0	1
z	7	1	0

node z table

t_0 (cost to):

from	x	y	z
x	∞	∞	∞
y	∞	∞	∞
z	7	1	0

t_1 (cost to):

from	x	y	z
x	0	2	7
y	2	0	1
z	3	1	0

time →

3. (Routing) Consider the following AS-level topology running BGP routing. Suppose x starts by announcing a path "x" for destination x to B. Upon receiving the path announcement, B learns a new path to reach x – "Bx". B updates its routing table and announces the new path "Bx" for destination x to A. Upon receiving the announcement, A learns a new path to reach x – "ABx" and updates its routing table.

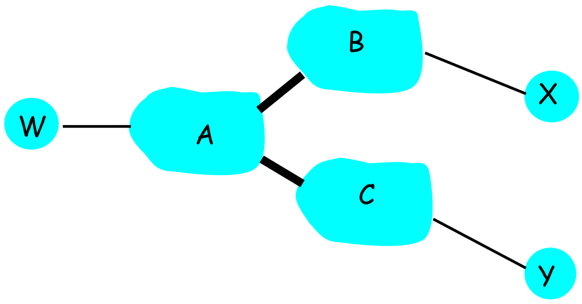

   (a) Suppose A announces the new path for destination x to C. What will C learn as a new path to x? (2%)
   (b) Continue from (a). Suppose C announces the new path for destination x to y. What will y learn as a new path to x? (2%)
   (c) Based on your understanding of Link State routing and Distance Vector routing, which one does BGP routing resemble more? And why? (4%)

Sample Solution:

(a) "CABx"
(b) "yCABx"
(c) Distance Vector. Routing table entries propagated and triggers neighboring routers to update their routing tables. If there's any route update, continue to propagate the newly discovered routes. Alternatively, one may argue BGP resembles Link State routing more if there're sensible arguments.

4. (Routing) Consider a slightly different AS-level topology running BGP routing below. Suppose x starts by announcing a path "x" for destination x to B. Upon receiving the path announcement, B learns a new path to reach x – "Bx". B updates its routing table and announces the new path "Bx" for destination x to A and C. Upon receiving the announcement, A learns a new path to reach x – "ABx" and updates its routing table.

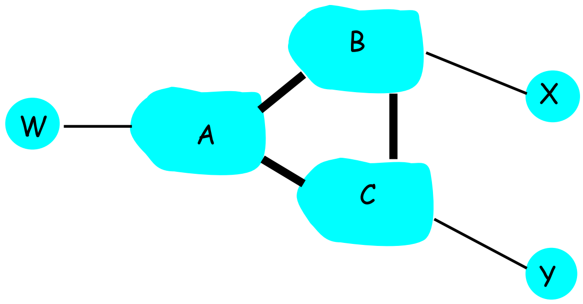

(a) What will C learn as a new path to x upon receiving the announcement from B? (2%)

(b) Suppose A announces the new path for destination x to C. What will C learn as a new path to x upon receiving the announcement from A? (2%)

(c) Continue from (b). If y is C's customer and forwarding packets for y with shorter delay will satisfy y better, which path for destination x would C prefer to announce to y? And why? (4%)

(d) Continue from (b). y is C's customer and C cares less of the customer's satisfaction. If A charges less per unit traffic to forward for C and B charges more, which path would C prefer to announce to y? And why? (4%)

Sample Solution:

(a) "CBx"

(b) "CABx"

(c) "CBx". There's likely a significant amount of delay for traffic to pass through A in addition to passing through C, B, and x.

(d) "CABx". y knows only to send traffic to x through the path "yCABx". In that, C will forward traffic to A instead of B where C pays less for the same job.

5. (EDC) Given a short data sequence 101010 to transmit over a link using CRC.
   (a) What are the EDC bits with G=101? (2%)
   (b) What are the EDC bits with G=1001? (2%)
   (c) If the link flips the first 3 bits of the whole sequence transmitted, will CRC with G=101 be able to detect the error? And what about CRC with G=1001? (4%)
   (d) If the link flips the first 4 bits of the whole sequence transmitted, will CRC with G=101 be able to detect the error? And what about CRC with G=1001? (4%)
   (e) Compare CRC with G of different length. (4%)

Sample Solution:

(a) 10
(b) 111
(c) yes. yes
(d) no, yes
(e) Longer G requires more EDC bits, but longer G is more reliable at detecting errors.

6. (MAC) We have come across three Carrier Sense Multiple Access (CSMA) based protocols. Namely, the original CSMA, CSMA/CD, and CSMA/CA. Based on your knowledge to these variants of CSMAs, address the following questions.
   (a) How does CSMA work in principle? (4%)
   (b) Can frames collide in CSMA and how? (4%)
   (c) How does CSMA/CD work in principle? (4%)
   (d) Will CSMA/CD be able to detect collisions in a wireless network? (4%)
   (e) How does CSMA/CA work in principle? (4%)

Sample Solution:

(a) Listen before transmit. Send when the channel is idle. Hold when the channel is busy.
(b) Yes. Multiple nodes may transmit about the same time. The channel may appear idle while the packet propagation from each node has begun but hasn't yet reached the other nodes.
(c) CSMA/CD stops the frame transmission as soon as the collision is detected so to reduce the channel wastage. In CSMA, the entire frame transmission time will be wasted as the collision occurs.
(d) No. Due to the hidden terminal problem, in which transmissions from certain nodes might not be visible by other nodes on the same wireless LAN.
(e) CSMA/CA avoids the undetectable collisions due to the hidden terminal problem by exchanging CTS and RTS frames, which in a sense alerts all the reachable nodes from the data sender and receiver. These nodes will reframe from sending any frames during the DATA-ACK exchange between the sender and receiver.

7. (Mobility) There are two ways of handling node mobility in general, indirect routing and direct routing. Consider the following topology. Network A, B, C, D, E, and F equally divide the circle of radius R. The mobile node lives in network A. Sitting at the center is the network where the correspondent resides. Let's assume the inquiries and replies exchanged between the correspondent and the mobile node are so small that the transmission and the processing delay are negligible. Let's assume also the propagation delay of sending messages within a network is so small that it is also negligible. The delay for the correspondent to send an inquiry and get a reply back from the mobile node is therefore 2R.

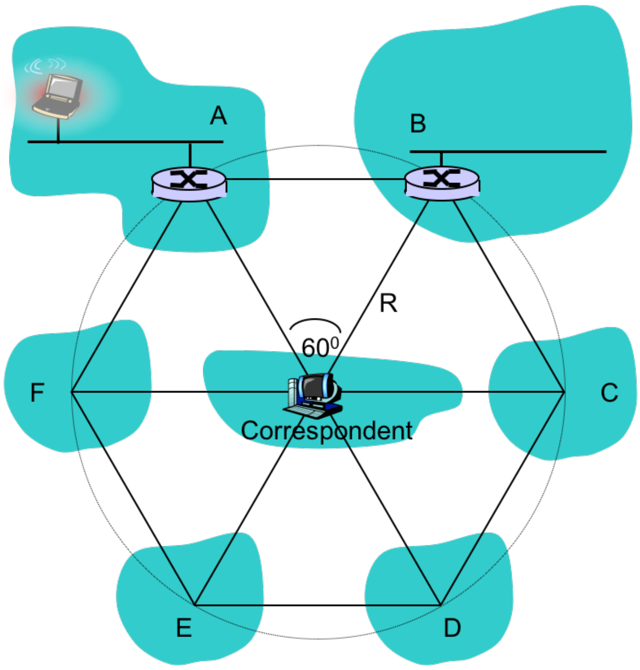

(a) Suppose now the mobile node moves to network B. What will be the delay for the correspondent to get a reply back from the mobile node using indirect routing? And what will be the delay using direct routing? (2%)

(b) Suppose now the mobile node moves to network B and stays in B for a duration in which the correspondent inquires the mobile node 5 times. How much delay in total the 5 correspondences will take using indirect routing? And how much delay in total using direct routing? (4%)

(c) Compare the delay in (a) and (b). Would you prefer to use indirect routing or direct routing if the mobile node will stay in a foreign network for a long while (more than 3 correspondences)? And why? (4%)

(d) Suppose now the mobile node moves to network B and stays in B for one correspondence and moves next to C. Similarly, the mobile node stays in C, D, E, F for one correspondence and then moves on to D, E, F, A. Suppose also that messages always travel the shortest delay path. How much delay in total the 5 correspondences (1 in B, 1 in C, 1 in D, 1 in E, and 1 in F) will take using indirect routing? And how much for direct routing (without the anchor foreign agent redirect routing)? (4%)

(e) Compare the delay in (a) and (d). Would you prefer to use indirect routing or direct routing if the mobile node continues to move from a foreign network to another with short stays (only 1 correspondence)? And why? (4%)

Sample Solution:

(a) Indirect routing: 3R (3-hop triangle route)
direct routing: 4R (2R to get the new address from A + 2R to exchange messages)

(b) Indirect routing: 15R (3R*5)
direct routing: 12R (2R to get the new address, 2R*5 to exchange messages 5 times)

(c) Direct routing. In direct routing, the overhead of getting the new address of the mobile node occurs only once. For the rest of the correspondences, the delay is minimum. The more correspondences to make within a foreign network, the more efficient using direct routing.

(d) Indirect routing: 18R (3-hop route for B and F; 4-hop route for C, D, E).
Direct routing: 20R (2R+2R, 5 times)

(e) Indirect routing. Using direct routing, the overhead of getting the new address of the mobile node will reoccur. In a shortest delay path network, the delay directly between the home and foreign network is less than or equal to going indirectly through a 3rd node (the correspondent).

8. (Stored Media) Three mechanisms that are common in the most popular stored video streaming services, such as YouTube and Netflix today are (a) client buffering, (b) adaptive streaming, and (c) CDN. Describe the purpose for each of the mechanisms:
   (a) client buffering (2%)
   (b) adaptive streaming (4%)
   (c) CDN (4%)

Sample Solution:

(a) Network end-to-end delay is not constant. To avoid having no video frames to play back when network delay suddenly increases, the client typically buffers video frames to some degree before the client starts playing back the video.
(b) Network bandwidth is not a constant either. To avoid dropping bits when downloading a high bitrate version of the video while the network bandwidth suddenly decreases, the client typically monitors (by accounting how much bits received per second) the network bandwidth and re-selects the version of the video with right bitrate (not exceeding the network bandwidth usually) to download.
(c) Even when the client and server are doing the best they can, the network delay is inevitably long if the two are physically far apart and the bandwidth is inevitably low if there are many more other middle points where cross traffic may interfere. Preloading the videos to the CDN and therefore having a local duplicate (of the server) close to the client reduces the delay and avoids cross traffic, which potentially allows transmitting the higher bitrate versions.