Name________________ Student ID______________________ Department/Year___________

Final Examination

Introduction to Computer Networks
Class#: 901 E31110
Fall 2006

9:20-11:00 Tuesday
January 16, 2007

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 14 pages (including this title page), 7 questions.
2. Write your **name in Chinese**, student ID, and department/year down on top of the first page.
3. There are in total 150 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. (Congestion Control) Suppose the change of the congestion window size of a TCP connection is as depicted below.

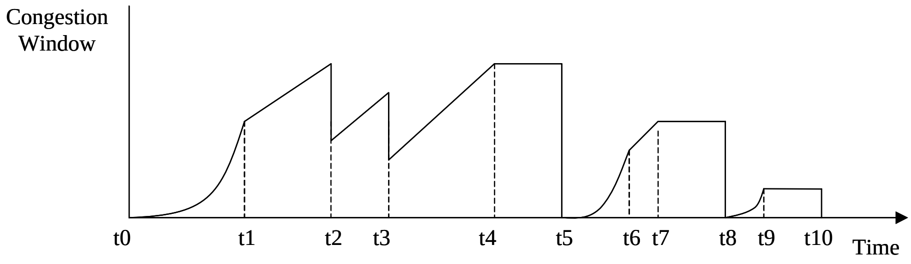

(1) State the time periods that the TCP connection is in the slow start phase. (5%)
(2) State the time periods that the TCP connection is in the congestion avoidance phase. (5%)

Sample solution:
(1) t0-t1, t5-t6, t8-t10
(2) t1-t5, t6-t8

2. (LS Routing) Consider a 5-node string network as follows. The link costs are equal and the value is 1. Follow the link state (LS) routing principle to obtain the routing table.

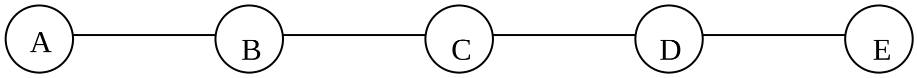

(1) Suppose each node sends its link state (LS) report to all the outgoing links and then the LS reports are further propagated using the algorithm on the right to reach the whole network.

> ***if*** (LS report received on incoming link)
> ***then*** flood LS report onto all but the incoming link

Suppose the LS reports are the same in size, M bytes each. In order for all nodes to receive LS reports from all other nodes, how many bytes of LS reports are transmitted over the network? (5%)

(2) Continue from (1). Suppose the delay to send a LS report over any link is T seconds. If A, B, C, D, and E nodes start sending their LS reports all at the same time, how much time does it take for all nodes to receive all LS reports? (5%)

(3) Continue from (1) and (2). Suppose all LS reports have arrived at all nodes. Compute the shortest paths from node E to every other node using the LS routing principle by filling in the blanks in the tables below. (5%)

Step	Travel Set	D(A),p(A)	D(B),p(B)	D(C),p(C)	D(D), p(D)
0	E	∞	∞	∞	1,E
1	ED	∞	∞	2,D	
2	EDC	∞	3,C		
3	EDCB	4,B			
4	EDCBA				

Sample Solution:

(1) 20M bytes

Each LS report travels 4 links to reach all other nodes in the network. There are 5 reports in total. There are, therefore, 20M bytes consumed over the network to deliver the LS reports.

(2) 4T seconds

All nodes start simultaneously. In the worst case (A to E or E to A), it takes 4 hops to deliver the LS reports. That is, therefore, 4T seconds required to deliver the LS reports.

(3) See the filled-up table above

3. (DV Routing)

(1) Describe how Distance Vector routing works in principle. (5%)

(2) Name one example of DV routing protocols. (5%)

(3) Describe the ‘Count To Infinity’ problem in DV routing. (Hint: easier by an example) (5%)

(4) State the main difference between Path Vector and Distance Vector routing. (5%)

(5) Name one example of PV routing protocols. (5%)

(6) Would the ‘Count To Infinity’ problem exist is PV routing? (5%)

Sample Solution:

(1) Each node on the network keeps a vector of best (next hop, distances) to every other node. Whenever a route report is received, the node updates the distance vector if the route report provides a better route to a particular destination via the neighbor from which the report is received. If this results in changes in the route (next hop or distance) to that destination, a route report is sent which might in turn change the distance vector of the node’s neighbors. In principle, each node will tell the neighbors the best information it’s got.

(2) RIP is a DV routing protocol

(3) Consider the scenario below. A goes to B through link A-B, to C through A-B-C. B goes to A through link A-B, to C through link B-C. C goes to A through C-B-A, to B through link B-C. Suddenly, link A-B breaks down.

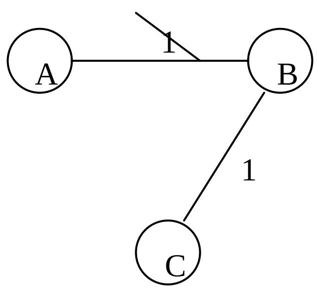

1. In B, the distance to go to A via A is set to infinity. Therefore, B decides going via C to A is a better route (distance of 3, B-C-B-A). B reports to C that its route to A is now via C with distance 3.
2. C updates the distance to A via B to 4. C reports to B that its route to A is still via B but with distance 4.

3. B updates the distance to A via C to 5 and reports to C that its route to A is still via C but with distance 5.
4. C updates the distance to A via B to 6 and reports to B that its route to A is via B with new distance 6.
5. The process continues until B updates the distance to A via C to infinity+1 and reports to C that its route to A is now via A with distance infinity.
6. C updates the distance to A via B to infinity+1 and reports to B that its route to A is with distance infinity+1
7. B updates the distance to A via C to infinity+2 and the routing tables finally converge.

This phenomenon that the network needs to wait until the routes are counted to infinity before the routing tables stablize is referred to as the 'Count to Infinity' problem. In the process of the routes counting to infinity, there could be a substantial amount of data looping in between without realizing that the destination is no longer reachable.

(4) Path Vector routing protocols propagate not only the distance, but also the entire path.
(5) BGP is a PV routing protocol.
(6) No

4. (MAC) We have come across three generations of Carrier Sense Multiple Access protocols -- the original CSMA, CSMA/CD, and CSMA/CA. Based on your knowledge to these variants of CSMA MAC protocols, address the following questions.
   (1) How does CSMA work in principle? (5%)
   (2) Can frames collide in CSMA and how? (5%)
   (3) How does CSMA/CD work in principle? (5%)
   (4) Can frames collide in CSMA/CD in a wireless network? (5%)
   (5) How does CSMA/CA work in principle? (5%)

Sample Solution:

(1) Listen before transmit. Send when the channel is sensed idle. Hold when the channel is sensed busy.
(2) Yes. Multiple CSMA transmissions might start about the same time when the channel is sensed idle. They could collide during the propagation delay. In CSMA, the entire frame transmission time will be wasted as the collision occurs.
(3) CSMA/CD tries to stop the frame transmission as soon as the collision is detected so to reduce the channel wastage. (Re-send after a random exponential backoff.)
(4) Due to the hidden terminal problem, in which transmissions from certain nodes might not be visible by other nodes on the same wireless LAN.
(5) CSMA/CA avoids the potential collisions due to the hidden terminal problem by the sending of CTS and RTS frames which in a sense alerts all the visible nodes from the data sender and receiver of the data-ack exchange coming up next.

5. (Mobility) Given your knowledge about (a) indirect routing and (b) direct routing for handling mobility, try to address the following questions.
   (1) Which one results in less efficient routes and why? (5%)
   (2) Which one offers nice IP address transparency to the correspondent and why? (5%)
   (3) Which one has better connection continuity and why? (5%)
   (4) Which one places more burden on the home agent and why? (5%)
   (5) Which one suffers from the packet encapsulation (a packet within a packet) overhead and why? (5%)

Sample Solution:

(1) (a) indirect routing.

Indirect routing results in triangle routing. It is not efficient in the route length. Using direct routing, the correspondent communicates directly to the destination. Therefore, it is more efficient.

(2) (a) indirect routing.

Indirect routing provides transparency to the correspondent. I.e., the correspondent does not need to know the current address of the destination. Direct routing does not provide transparency to the correspondent. The correspondent needs to learn a new address every time.

(3) (a) indirect routing.

Indirect routing could keep the connections ongoing to the destination's permanent address when the destination mobile node moves from one foreign network to another. Direct routing will receive a new address each time when the destination node moves to a new foreign network. It will not be able to keep the ongoing connections running continuously without re-establishing the connection to the new address.

(4) (a) indirect routing.

The home agent for indirect routing needs to track where the destination is currently and encapsulates the original packets to be sent to the care-of address of the mobile destination. The home agent of the direct routing only needs to track the designation's new address and let the correspondent node know when requested.

(5) (a) indirect routing.

In indirect routing, every original data packet is encapsulated within another packet. The extra header overhead (the size of IP header in TCP/IP) accumulates. As opposed to indirect routing, direct routing does not have the encapsulation overhead.

6. (QoS) Suppose you are designing the communication part of the following three Internet applications. The table below shows the QoS (quality of service) the general users demand for the three types of network applications.

	Loss	Time Sensitive
Web	no loss	no
Skype	loss tolerant	yes
Online Game	no loss	yes

(1) Which transport layer services, TCP or UDP, will you choose to transfer Web data and why? (10%)
(2) Which transport layer services, TCP or UDP, will you choose to transfer Skype calls and why? (10%).
(3) Which transport layer services, TCP or UDP, will you choose to transfer game data and why? (10%)

Sample Solution:

(1) TCP. No-loss is the only requirement.
(2) UDP. Delay is important. Reliability isn't. This way, there is no need to wait for the retransmission that might be too late already
(3) Your pick as long as you can justify your answer. For example, one may choose UDP and implement forward error correction (FEC) mechanisms on the game data transmission so that the lost data can be recovered at the receiver without retransmission. Alternatively, one may choose TCP and shorten the retransmission timeout delay so that the retransmission can be as timely as possible.

7. (EWMA) In the adaptive playout delay calculation for Internet phone multimedia applications, we estimate the average network delay and the average delay deviation using the following formula. The formula basically computes a new average (A_n) by taking the weighted (α) sum of current average (A_{n-1}) and new sample (S_n), where $0 < \alpha < 1$. When the value of α is close to 1, the new average adapts faster to the latest sample value. And vice versa, when the value is close to 0.

$$A_{n+1} = (1-\alpha)A_n+(\alpha)S_n$$

(1) Let $A_0 = 0$. Expand A_{n+1} in terms of S_i, i=0…n. (10%)
(2) Explain why A_{n+1} is called the 'exponentially weighted moving average'. (5%)

Sample Solution:

(1) $A_{n+1} = (1-\alpha)^n (\alpha)S_0 + (1-\alpha)^{n-1}(\alpha)S_1 + (1-\alpha)^{n-2}(\alpha)S_2 + \ldots + (1-\alpha)^2(\alpha)S_{n-2} + (1-\alpha)(\alpha)S_{n-1} + (\alpha)S_n$

(2) If the value of α is close to 1/0, the average is exponentially less/more relevant (i.e., $(1-\alpha)^n$) to earlier samples (some n time periods before).