Name_______________ Student ID____________________ Department/Year__________

Midterm Examination

Introduction to Computer Networks
Class#: EE 4020, Session#: 02
Fall 2016

10:20-12:10 Thursday
November 10, 2016

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 16 pages (including this title page), 6 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. (Overview) Consider a hierarchical home network such as this one below.

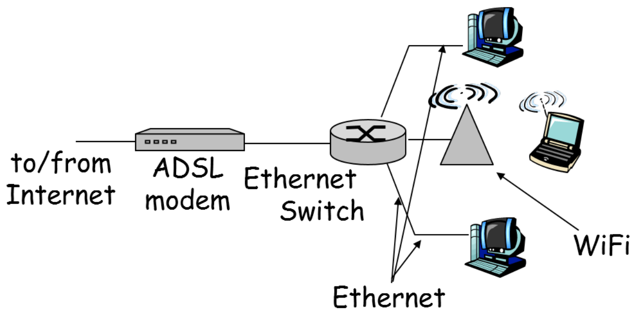

(1) Circle out (right on the figure) the devices that are end systems. (4%)
(2) Which is the access network the home owner pays monthly/quarterly/yearly to the ISP to connect to the Internet? (4%)
(3) Which is(are) the access network(s) that facilitates connectivity to the end systems? (4%)

Sample Solution:

(1)

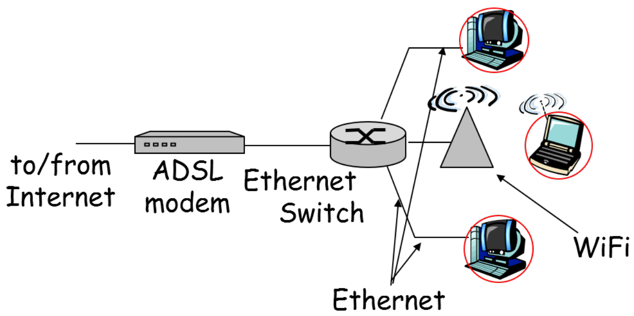

(2) ADSL
(3) Ethernet and WiFi

2. (Overview) Compare and contrast the two public transportation networks in Taipei city: the bus network and the Metropolitan Rapid Transit (MRT) network, based on your understanding of circuit switching and packet switching principle.
   (1) Is the bus network more circuit switching or packet switching? Why do you think so? (4%)
   (2) Is the MRT network more circuit switching or packet switching? Why do you think so? (4%)

Sample Solution:

   (1) Packet switching with virtual circuit. There is a fixed route but the route may not be dedicated/reserved solely for busses only. The buses can well be stuck in traffic during the rush hour and therefore resembles more the packet switched network. It is very much like the virtual circuit form of delivery in packet switched network.
   (2) Circuit switching. The MRT tracks are scheduled and to be switched at a designated time to ensure the people in train follow a pre-determined path to the destination.

3. (Application) We learn in Chapter 2 several application-layer protocols. In particular, HTTP for Web transfer, SMTP for email delivery, IMAP for mailbox access, FTP for file transfer, and DNS for name-IP lookup. Each of these application-layer protocol uses a transport-layer service to deliver the messages.
   (1) Which of these 5 protocols uses TCP as the transport-layer service and Why? (4%)
   (2) Which of these 5 protocols uses UDP as the transport-layer service and Why? (4%)

Sample Solution:

(1) HTTP, SMTP, IMAP, FTP. These application are not loss-tolerant and do not require timely delivery.
(2) DNS. The name-IP lookup is the 1^{st} thing general Internet services will do before the actual transmission of control/data messages. It is therefore essential to ensure timely DNS query/reply delivery. UDP in this regard serves better. Also, DNS service allows configuration of multiple local DNS servers and re-tries in the absence of timely replies.

4. (Application) Below is a network of DNS servers and hosts. To simplify the analysis, assume the delay in between any of the 2 nodes is d and there's no packet loss.

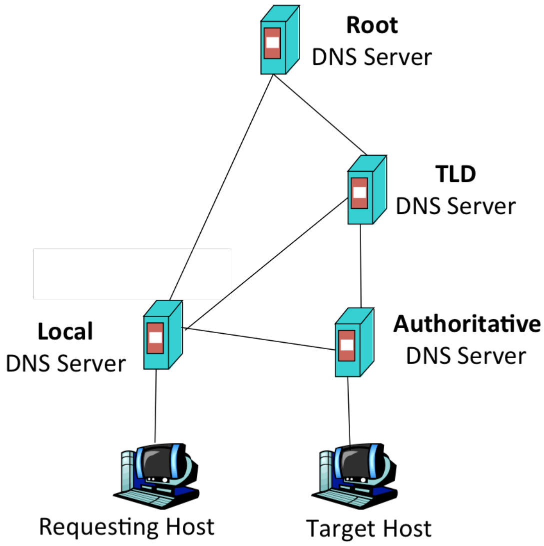

Let's start from a very primitive form of DNS where DNS queries are still recursive and the local DNS server does not maintain the <name, IP address> cache. Now quantify the delay and root server load and compare when 'iterative query' technique is applied.

(1) How long does it take for the DNS reply to come back with a recursive query? And how long with an iterative query? (4%)

(2) How many messages (including sent and received) does the root DNS server need to handle with a recursive query? And how many with an iterative query? (4%)

Now, assume all queries are iterative. Analyze next the benefit of applying the 'caching' technique.

(3) How long does it take for the DNS reply to come back if the <name, IP address> mapping being requested is already cached at the local DNS server? And how long if the mapping is not cached? (4%)

(4) How many messages (including sent and received) does the root DNS server need to handle with the requested <name, IP address> mapping cached? And how many if not cached? (4%)

Lastly, compare and contrast quantitatively the 'iterative query' and 'caching' techniques based on the scenario provided above.

(5) Which of the 2 techniques help(s) reducing the delay? (4%)

(6) Which of the 2 techniques help(s) reducing the root DNS server load? (4%)

Sample Solution:

(1) 8d, 8d
(2) 4, 2
(3) 2d, 8d
(4) 0, 2
(5) caching
(6) both

5. (Application) Consider a circular DHT storing the <file-ID, torrent-IPs> mappings for a BitTorrent tracker – PiratesAreHumansToo.tw. The DHT has 8 nodes, whose IDs are 0, 1, 2, … 7 respectively. All 8 nodes are up and running reliably. The file ID ranges from 0 to arbitrarily large. The hash function is defined as simple as: key = (file-ID mod 8), meaning a <file-ID, torrent-IPs> mapping is stored in the node whose ID = key. Upon receiving a search query for a file (file-ID), each node operates based on a simple algorithm: check if node-ID == key. If so, done and search the local <file-ID, torrent-IPs> entries for the matching file. Otherwise, relay the search query to the successor or the shortcut node that is the closest but not exceeding the key.

   The DHTs below depicts (a) the version without any shortcuts, (b) the version with 1 shortcut each to the node straight across (4 nodes apart), and (c) the version with 1 shortcut each to the node that is 2-node away. (d) the version with 2 shortcuts each to the nodes that are 2- and 4-node away

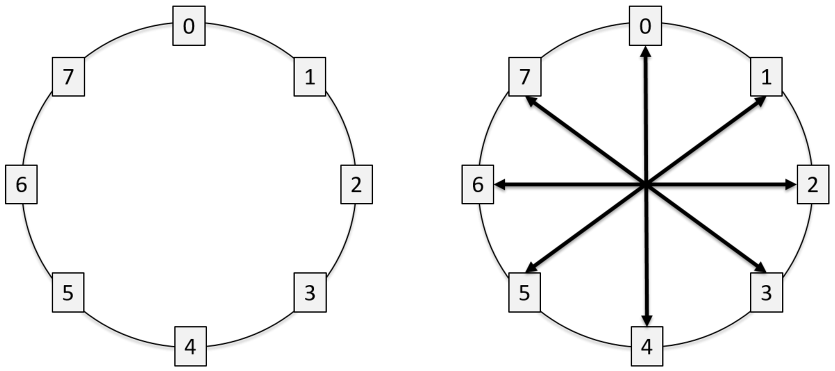

(a) No shortcut (b) 1 shortcut 4-node away

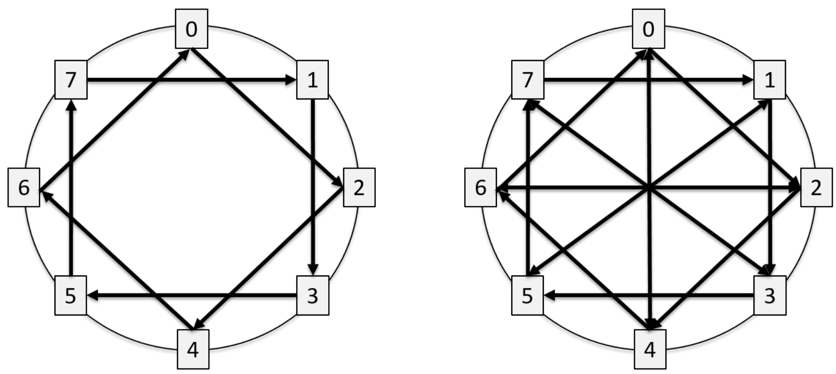

(c) 1 shortcut 2-node away (d) 2 shortcuts 2- and 4-node away

(1) On DHT (a), if the peer node ID=0 is searching for file ID=16, 17, 18, 19, 20, 21, 22, and 23, how many times the search query needs to be relayed to find the corresponding <file-ID, torrent-IPs > mapping? (4%)
(2) On DHT (b), if the peer node ID=0 is searching for file ID=16, 17, 18, 19, 20, 21, 22, and 23, how many times the search query needs to be relayed to find the node storing the <file-ID, torrent-IPs > mapping? (4%)
(3) On DHT (c), if the peer node ID=0 is searching for file ID=16, 17, 18, 19, 20, 21, 22, and 23, how many times the search query needs to be relayed to find the node storing the <file-ID, torrent-IPs > mapping? (4%)
(4) On DHT (d), if the peer node ID=0 is searching for file ID=16, 17, 18, 19, 20, 21, 22, and 23, how many times the search query needs to be relayed to find the node storing the <file-ID, torrent-IPs > mapping? (4%)
(5) Suppose the file IDs of the queries are uniformly random (equal chance of each file being requested). What is the average number of relays for DHT (a) and (b)? Which one relays more the search queries and why? (4%)
(6) Suppose the file IDs of the queries are uniformly random (equal chance of each file being requested). What is the average number of relays for DHT (b) and (c)? Which one relays more the search queries and why? (4%)
(7) Suppose the file IDs of the queries are uniformly random (equal chance of each file being requested). What is the average number of relays for DHT (c) and (d)? Which one relays more the search queries and why? (4%)

Sample Solution:

(1) 0, 1, 2, 3, 4, 5, 6, 7
(2) 0, 1, 2, 3, 1, 2, 3, 4
(3) 0, 1, 1, 2, 2, 3, 3, 4
(4) 0, 1, 1, 2, 1, 2, 2, 3
(5) The average number of relays for DHT (a) and (b) are 3.5 and 2.0 respectively. DHT (a) relays more, because there is no shortcut to allow jumping to the closer node when possible.
(6) The average number of relays for DHT (b) and (c) are the same, both 2.0. Both adds 1 short cut per node and evenly spread out (therefore both maximizing the opportunities to jump when possible).
(7) The average number of relays for DHT (c) and (d) are 2.0 and 1.5 respectively. DHT (c) relays more because there are less shortcuts allowing jumping to the closer node.

6. (Transport) Provided below are the FSMs of rdt 2.1 sender and receiver. Indicate the order of the transitions (in terms of t1, t2, …, t12) taking place for the following scenarios.

rdt 2.1 sender:

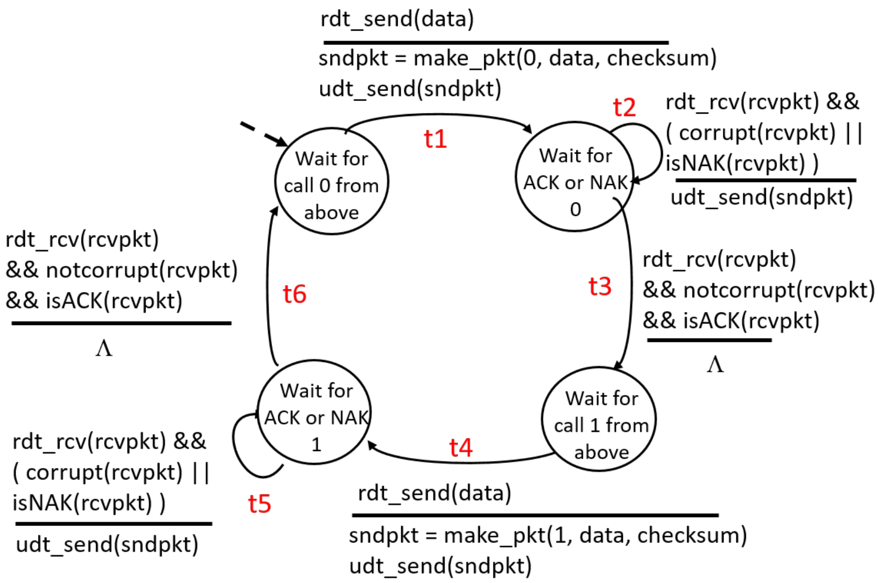

rdt 2.1 receiver:

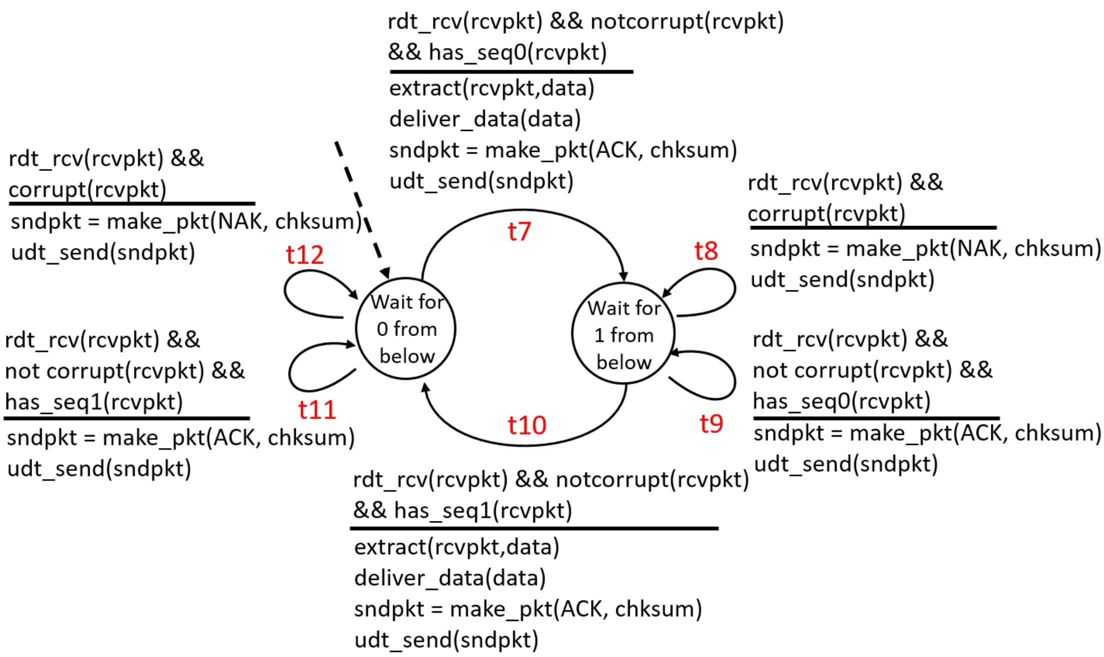

(1) Both sender and receiver start from the initial state. The sender gets a call from above and there is no bit error. (4%)
(2) Both sender and receiver start from the initial state. The sender gets two calls from above and there is no bit error. (4%)
(3) Both sender and receiver start from the initial state. The sender gets a call from above and there is a bit error in the data packet. (4%).
(4) Both sender and receiver start from the initial state. The sender gets a call from above and there is a bit error in the ACK packet coming back. (4%)
(5) Both sender and receiver start from the initial state. The sender gets a call from above and the ACK packet does not come back at all. (4%)

Sample Solution:

(1) t1, t7, t3
(2) t1, t7, t3, t4, t10, t6
(3) t1, t12, t2, t7, t3
(4) t1, t7, t2, t9, t3
(5) t1, t7, and then the sender hangs. None of the t2 and t3 defines such an event. rdt 2.1 does not handle cases where there is any packet loss.