

COE 431 – Computer Networks

Welcome to
The Midterm Exam
Tuesday April 10, 2012

Instructor: Wissam F. Fawaz

Name: _____

Student ID: _____

Instructions:

1. This exam is **Closed Book**. Please do not forget to write your name and ID on the first page.
2. You have exactly **85 minutes** to complete the 7 required problems.
3. Read each problem carefully. If something appears ambiguous, please write your assumptions.
4. Do not get bogged-down on any one problem, you will have to work fast to complete this exam.
5. Put your answers in the space provided only. No other spaces will be graded or even looked at.

Good Luck!!

Problem I: Multiple choice questions (20 minutes) [17 Points]

For each question, choose the **single** correct answer.

1. An organization's Web server and mail server:
 - a. Must run on the same machine.
 - b. Can run on the same machine if they share one DNS record.
 - c. **Can have exactly the same alias for a hostname (for example, foo.com).**
 - d. Cannot have the same aliased names.
2. The end to end delay depends on
 - a. The congestion of the network.
 - b. The number of hops between source and destination.
 - c. The application layer protocols.
 - d. **(a) and (b)**
3. IMAP
 - a. is built on top of POP3.
 - b. **allows users to create remote folders.**
 - c. is not compatible with SMTP.
 - d. is stateless across sessions.
4. Packet switching versus circuit switching: which of the following is correct?
 - a. Packet switching is more flexible, circuit switching is more robust to link failures.
 - b. Packet switching is more secure, circuit switching is more efficient.
 - c. **Packet switching allows more users to access the network, circuit switching provides quality of service guarantees.**
 - d. Packet-switching and circuit-switching offer the same bandwidth guarantees.
5. Consider an HTTP client that wants to retrieve a Web document at a given URL. Assuming that the IP address of HTTP server is initially unknown, what **transport** -layer protocol is needed to obtain the IP address of the server?
 - a. DNS
 - b. **TCP**
 - c. UDP
 - d. FTP
6. Which of the following is **true** about encapsulation:
 - a. Packets encapsulate frames.
 - b. Messages encapsulate datagrams.
 - c. **Frames encapsulate segments.**

- d. None of the above.
7. For a type MX DNS resource record, which of the following is **true**:
- a. “Value” is the hostname of the DNS server that is authoritative for the domain name given in the “Name” field.
 - b. **“Value” is the canonical name of the mail server whose alias is stored in the “Name” field.**
 - c. “Value” is the IP address of the mail server that has the alias hostname “Name”.
 - d. None of the above
8. Which of the following is **true**:
- a. Layers 2 and 3 of the Internet protocol stack are implemented in the end systems but not in the routers in the core of the network.
 - b. **With SMTP, it is possible to send multiple messages over the same TCP connection.**
 - c. Before sending a packet into a datagram network, the source must determine all of the links that the packet will traverse between source and destination.
 - d. FTP sends control information “in-band”.
9. For BitTorrent, which of the following is **true**:
- a. In BitTorrent, a common problem is free-riding, in which a peer downloads files from the system without uploading files.
 - b. **Among the chunks that a peer does not have, the peer will request the chunks that are the rarest among his neighbors.**
 - c. None of the above
 - d. All of the above
10. Which of the following is **true** about HTTP
- a. **Two distinct web pages can be sent over the same persistent connection.**
 - b. The `Date:` header in the HTTP response message indicates when the object in the response was last modified.
 - c. HTTP uses persistent UDP connections in its default mode.
 - d. None of the above
11. Transport-layer chunks of data are called:
- a. Messages
 - b. Datagrams
 - c. **Segments**
 - d. Frames
12. The File Transfer Protocol uses
- a. A persistent TCP data connection and a non-persistent TCP control connection.

- b. A non-persistent TCP data connection and a non-persistent TCP control connection.
 - c. Persistent TCP data and control connections.
 - d. **None of the above**
13. Which of the following is **false**:
- a. It is possible for a mail client to send e-mail messages over HTTP rather than SMTP.
 - b. **ADSL bandwidth is shared.**
 - c. All of the above
 - d. None of the above
14. Which of the following is **true** about a web cache:
- a. A web cache makes use of the cookie technology to reduce the response time for a client request.
 - b. **It responds to HTTP requests on the behalf of a web server.**
 - c. It substantially increases the amount of traffic flowing through the access link connecting an institution to the Internet.
 - d. None of the above
15. Statistical multiplexing is
- a. A method to optimally shared files in a peer-to-peer system
 - b. An algorithm that divides up the frequency spectrum of a link
 - c. Pre-allocates use of transmission link regardless of demand
 - d. **An on-demand form of resource sharing**
16. A user requests a web page that consists of some text and 3 images. The browser's cache is empty. For this page, the client's browser:
- a. sends 1 HTTP request message and receives 1 HTTP response message.
 - b. sends 1 HTTP request message and receives 4 HTTP response messages.
 - c. sends 3 HTTP request messages and receives 3 HTTP response messages.
 - d. **Sends 4 HTTP request messages and receives 4 HTTP response messages.**
17. HTTP with persistent connections
- a. Requires 2 RTTs per object.
 - b. Requires the server to open a new TCP connection to send a response message.
 - c. All of the above.
 - d. **None of the above.**

Problem II: Comparing terminologies (5 minutes) [8 Points]

What is the difference between each of the following pairs of concepts?

1. HTTP and FTP

HTTP and FTP they have many common characteristics. However, they differ in many fundamental ways. The most striking difference is that FTP uses two parallel TCP connections.

2. Persistent HTTP and non-persistent HTTP

With non-persistent HTTP, each request/response pair is sent over a separate TCP connection. With persistent HTTP, all pairs are sent over the same TCP connection.

3. FDM and TDM

With FDM, the frequency spectrum is divided up among the connections established across a link. TDM on the other hand is a frame based solution where the time domain is segmented into frames.

4. Recursive DNS queries and iterative DNS queries

A recursive DNS query is sent by one entity to another asking it to obtain the mapping on its behalf. If an iterative query is used the reply is directly returned to the requesting entity.

Problem III: Socket Programming (15 minutes) [15 Points]

1. Consider the following Java code fragment:

```
try {
    ServerSocket ss = new ServerSocket(2011);
    Socket s1 = ss.accept();
    Socket s2 = ss.accept();
    BufferedReader is = new BufferedReader(
        new InputStreamReader(
            s1.getInputStream()));
    DataOutputStream os = new DataOutputStream(
        s2.getOutputStream());

    while(true)
        os.writeBytes(is.readLine() + '\n');
} catch(Exception e) {
    // crash!
}
```

What does this code do?

- a. It does not compile because the ServerSocket constructor takes no parameters.
- b. It always crashes in line 4, because you can only accept() once.
- c. It creates an “echo” program that responds to the client with the exact same lines he sent.
- d. **It creates a “relay” program that sends the bytes received from one client to another client.**

2. Consider the following Java program:

```
1: public class TCPServer {
2:     public static void main(String[] args) throws Exception {
3:         String clientSentence;
4:         String modSentence;
5:         ServerSocket welcomeSocket = new ServerSocket(6789);
6:         while(true) {
7:             Socket connectionSocket = welcomeSocket.accept();
8:             BufferedReader inFromClient = new BufferedReader(new
9:                 InputStreamReader(connectionSocket.getInputStream()));
10:            DataOutputStream outToClient = new DataOutputStream(
11:                connectionSocket.getOutputStream());
12:            clientSentence = inFromClient.readLine();
13:            modSentence = clientSentence.toUpperCase() + '\n';
14:            outToClient.writeBytes(modSentence);
15:        }
16:    }
17: }
```

Complete the code of the TCP client given below. This client takes the first command line argument, sends it to the server, and then displays the modified version of the line received from server on the screen. Assume that the name of server is “compnet.lau.edu.lb”. For simplicity, ignore exception handling.

```
public class TCPServer {
    public static void main(String[] args) throws Exception {
        // Get the first command line argument
        String clientSentence = args[0] + '\n';
        Socket connectionSocket =
            new Socket("compnet.lau.edu.lb", 6789);
        BufferedReader inFromServer = new BufferedReader(
            new InputStreamReader(
                connectionSocket.getInputStream()));
        DataOutputStream outToServer =
            new DataOutputStream(
                connectionSocket.getOutputStream());
        outToServer.writeBytes(clientSentence);
        String capitalizedSentence = inFromServer.readLine();
        System.out.println(capitalizedSentence);
        connectionSocket.close();
    }
}
```

Problem IV: Domain Name System (15 minutes) [20 Points]

1. Suppose you decided to create a startup company called “netnet”☺. Your network has the following servers:

- A primary DNS server which has the canonical name “dns1.netnet.com” and an IP address of “128.119.12.40”.
- A secondary DNS server that has the canonical name “dns2.netnet.com” and an IP address of “128.119.12.41”.
- A Web server with the canonical name “www.netnet.com” and the alias name “netnet.com”. Moreover, the web server has two IP addresses associated with it: “128.119.12.55” and “128.119.12.56”.
- A Mail server assigned a canonical name of “galaxy.netnet.com” and an IP address of “128.119.12.60”. Note that your employees’ email addresses will take the form firstName@netnet.com.

What actions should be taken in order to allow people from outside of your company to visit your website and send email to the employees at your company?

The following resource records will be provided to the upper layer “.com” Registrar: (netnet.com, dns1.netnet.com, NS) and (dns1.netnet.com, 128.119.12.40, A) along with (netnet.com, dns2.netnet.com, NS) and (dns2.netnet.com, 128.119.12.41, A). In addition, the following resource records should be put in your DNS server

(www.netnet.com, 128.119.12.55, A)

(www.netnet.com, 128.119.12.56, A)

(www.netnet.com, netnet.com, CNAME)

(galaxy.netnet.com, 128.119.12.60, A)

(netnet.com, galaxy.netnet.com, MX)

If the web user sends an e-mail to boss@netnet.com, to which IP address does the e-mail get delivered?

“128.119.12.60”

2. Assume that the following lines are inserted into a DNS server:
- (midterm1.edu, service1.edu, CNAME)
 - (midterm2.edu, service1.edu, MX)
 - (service1.edu, 8.11.11.101, A)
 - (service2.edu, 8.11.11.104, A)

As a consequence, when you open midterm.edu in your browser:

- a. **The request is served by the server with IP 8.11.11.101**
 - b. The request is served by the server with IP 8.11.11.104
 - c. The request is served by the server with IP 8.11.11.101 or 8.11.11.104, depending on the current server load.
 - d. The DNS server cannot hold two records for midterm.edu
3. Suppose within your web browser you click on a link to obtain a Web page. The IP address for the associated URL is not cached in your local host, so a DNS lookup is necessary to obtain the IP address. Suppose that n DNS servers are visited before your host received the IP address from DNS; the successive visits incur an RTT of $RTT_1, RTT_2, \dots, RTT_n$. Further suppose that the Web page associated with the link contains exactly one object, consisting of a small amount of HTML text. Let RTT_0 denote the RTT between the local host and the server containing the object. Assuming zero transmission time of the object, how much time elapses between when the client clicks on the link until it receives the object?
- a. $RTT_1 + RTT_2 + \dots + RTT_n + RTT_0$
 - b. $RTT_1 + RTT_2 + \dots + RTT_n + 2 RTT_0$
 - c. $2 (RTT_1 + RTT_2 + \dots + RTT_n) + RTT_0$
 - d. $2 (RTT_1 + RTT_2 + \dots + RTT_n + RTT_0)$

Problem V: Persistent/non-persistent (10 minutes) [15 Points]

Assume that you want to retrieve a webpage that references 9 very small objects.

1. When using non-persistent HTTP with no parallel TCP connections, how much time elapses from when the client requests the webpage until the entire webpage comes in? How many messages will the client and the server exchange before the entire webpage is received by the client's host?

A total of 10 objects (including the base html file) must be retrieved, each requiring a new TCP connection. Thus 40 messages will be exchanged, requiring 20 RTTs.

2. How long would the client wait before the entire webpage is received under persistent HTTP with no pipelining? How many messages must be sent in this case?

Only one connection should be opened in this case, so 22 messages will be sent requiring 11 RTTs.

3. How long does the client wait until it obtains the entire webpage when using non-persistent HTTP with a maximum of 5 parallel connections?

3 groups of objects will be retrieved from the server with each requiring 2 RTTs. As such, total number of RTTs = 6 RTTs.

4. How long does the client wait until it obtains the entire webpage under persistent HTTP with a maximum of 5 pipelined requests?

The base html file, a first set of 5 objects followed by a second set of 4 objects will require each one $RTT + RTT$ to establish a TCP connection, therefore, total number of RTTs = 4 RTT.

Problem VI: Peer to peer (10 minutes) [10 Points]

1. Consider distributing a file of size $F = 700$ Mbytes to $N = 10000$ peers. The server has an upload rate of $u_s = 100$ Mbps, and each peer has a download rate of $d_i = 5$ Mbps and an upload rate of $u = 300$ Kbps. Compute the minimum distribution time for both client-server and P2P architectures. Assume that $1 \text{ Mbps} = 1000 \text{ Kbps} = 1000 \text{ 000 bps}$ and that $1 \text{ Mbyte} = 1024 \text{ Kbytes} = (1024)^2 \text{ bytes}$.

For calculating the minimum distribution time for client-server distribution, we use the following formula: $D_{cs} = \max\{NF/u_s, F/d_{min}\} = 163 \text{ h } 6 \text{ min } 43 \text{ s}$. Similarly, for computing the minimum distribution time for P2P distribution, we use the following formula: $D_{P2P} = \max\{F/u_s, F/d_{min}, NF/(u_s + \sum(u_i))\} = 5 \text{ h } 15 \text{ min } 42 \text{ s}$.

Problem VII: (10 minutes) [15 Points]

Consider two hosts, A and B, separated by three links. These three links are connected by two packet switches. Let d_i , s_i , and R_i denote the length, propagation speed, and the transmission rate of link i , for $i = 1, 2, 3$. Each packet switch introduces a constant processing delay of d_{proc} for each packet. Consider sending a packet of length L from A to B.

- a. Assuming no queueing delays, what is the total end-to-end delay for the packet, in terms of d_i , s_i , and R_i ($i = 1, 2, 3$), L and d_{proc} ?

$$\text{End to end delay} = L/R1 + L/R2 + L/R3 + d1/s1 + d2/s2 + d3/s3 + d_{\text{proc}} + d_{\text{proc}}.$$

- b. Suppose now the packet consists of 1500 bytes, the propagation speed on all three links is 2.5×10^8 m/s, the transmission rate of all three links are 2 Mbps, the packet switch processing delay is 3 msec, the length of the first link is 5000 km, the length of the second link is 4000 km and the length of the last link is 1000 km. For these values, what is the end to end delay? Assume that 1 Mbps = 1000 Kbps = 1000000 bps.

Answer: 64 msec

- c. Consider a router buffer preceding an outbound link. In this part, you will use Little's formula, a famous formula from queueing theory. Let N denote the average number of packets in the buffer plus the packet being transmitted. Let a denote the rate of packets arriving at the link. Let d denote the average total delay (i.e., the queueing delay plus the transmission delay) experienced by a packet. Little's formula is $N = a \times d$. Suppose that on average, the buffer contains 10 packets, and the average packet queueing delay is 10 ms. The link's transmission rate is 100 packets/sec. Using Little's formula, compute the average packet arrival rate, assuming there is no packet loss.

The number of packets in the system includes those in buffer and the packet in service $\Rightarrow N = 10 + 1 = a \times d \Rightarrow a = 550$ packets/sec.