

COE 431 – Computer Networks

Welcome to Exam I
Thursday April 02, 2015

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 6 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: Comparing terminologies (**10 minutes**) [10 Points]

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

1. Active mode FTP and passive mode FTP

With active mode FTP, the data connection is initiated by the server while in passive mode, the client creates the data connection.

2. Stateless HTTP and stateful HTTP

Stateless HTTP does not store any information about the client while stateful HTTP does.

3. Authoritative DNS server and local DNS server

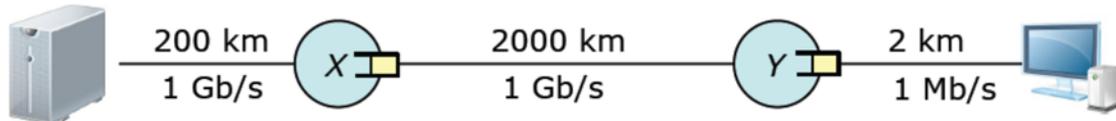
Local DNS acts as a proxy forwarding DNS queries to the DNS hierarchy of servers while authoritative DNS maintains records for a given domain.

4. Canonical name and alias name

Alias name is mnemonic while canonical is not.

5. ServerSocket class and DatagramSocket class

ServerSocket is used to create a welcome sockets while DatagramSocket creates a UDP socket.

Problem II: Delays (10 minutes) [20 Points]

The figure above shows a network path connecting a server to a client.

1. What is the total propagation delay for a packet going from the server to the client (you may assume that the speed of light is 200,000 Km/s)?

$$1 \text{ ms} + 10 \text{ ms} + 0.01 \text{ ms} = 11.01 \text{ ms}$$

2. What is the total transmission delay of a 5000 bit packet on all the links as it travels from the server to the client?

$$5 \text{ }\mu\text{s} + 5 \text{ }\mu\text{s} + 5 \text{ ms} = 5.01 \text{ ms}$$

3. Assume an infinite buffer size at router X and a traffic intensity of 0.9. What is the average queueing delay at X if the average packet size is 5000 bits?

$$1/\mu = 5 \times 10^{-6}; \text{ queueing delay} = 1/\mu \times \rho / (1-\rho) = 5 \times 10^{-6} \times 9 = 45 \text{ }\mu\text{s}$$

4. Assume an infinite buffer size at router Y and a traffic intensity of 0.8. What is the average queueing delay at Y if the average packet size is 5000 bits?

$$1/\mu = 5 \times 10^{-3}; \text{ queueing delay} = 1/\mu \times \rho / (1-\rho) = 5 \times 10^{-3} \times 4 = 20 \text{ ms}$$

5. Building on the above, what is the end-to-end delay experienced by a 5000 bits packet going from the server to the client?

$$11.01 \text{ ms} + 5.01 \text{ ms} + 20.045 \text{ ms} = 36.065 \text{ ms}$$

Problem III: Persistent and non-persistent HTTP (10 minutes) [16 Points]

A user in Beirut is connected to the Internet with a **100 Mb/s** (b=bits) access link. The said user retrieves a **250 KB** (B=bytes) webpage from a server in London, where the page references **3** images of **500 KB** each. Assume that the one way propagation delay is **75 ms** and that the capacity of the user's access link is the bandwidth bottleneck for this connection. Approximately, how long does it take for the page (including images) to appear on the user's screen, assuming:

1. Non-persistent HTTP using a single connection at a time?

8 RTTs + transfer delay for 4 objects = $8 \times 150 \text{ ms} + 140 \text{ ms} = 1340 \text{ ms}$

2. Persistent HTTP with no pipelining?

5 RTTs + transfer delay for 4 objects = 890 ms

3. Non-persistent HTTP with a maximum of 2 parallel connections?

6 RTTs + transfer delay for 4 objects = 1040 ms

4. Persistent HTTP with a maximum of 3 pipelined requests?

3 RTTs + transfer delay for 4 objects = 590 ms

Problem IV: Socket Programming (15 minutes) [16 Points]

Suppose a Server having an IP address of 1.2.3.4 executes the following lines of Java:

```
1. DatagramSocket sSocket = new DatagramSocket(__4321__);
2. byte[] sData = new byte[1024];
3. DatagramPacket sPckt = new DatagramPacket(__sData__, __sData.length__);
4. sSocket.receive(sPckt);
5. System.out.println(sPckt.getPort());
6. System.out.println(new String(sPckt.getData()));
7. System.out.println(sPckt.getAddress());
```

Now suppose a host with an IP of 2.3.4.5 executes the following lines to send a String to the server above:

```
8. DatagramSocket cSocket = new DatagramSocket(1234);
9. InetAddress add = InetAddress.getByName("1.2.3.4");
10. int port = 4321;
11. String str = "Client 1";
12. byte[] cData = str.getBytes();
13. DatagramPacket cPckt =
14. new DatagramPacket(cData, cData.length, port, add);
15. cSocket.send(cPckt);
```

1. Complete below the implementation of lines 1, 3, 9, 10, and 12 of the above presented code fragments?

See above.

2. What output is printed to the screen when the code fragment corresponding to the server executes?

1234

Client 1

2.3.4.5

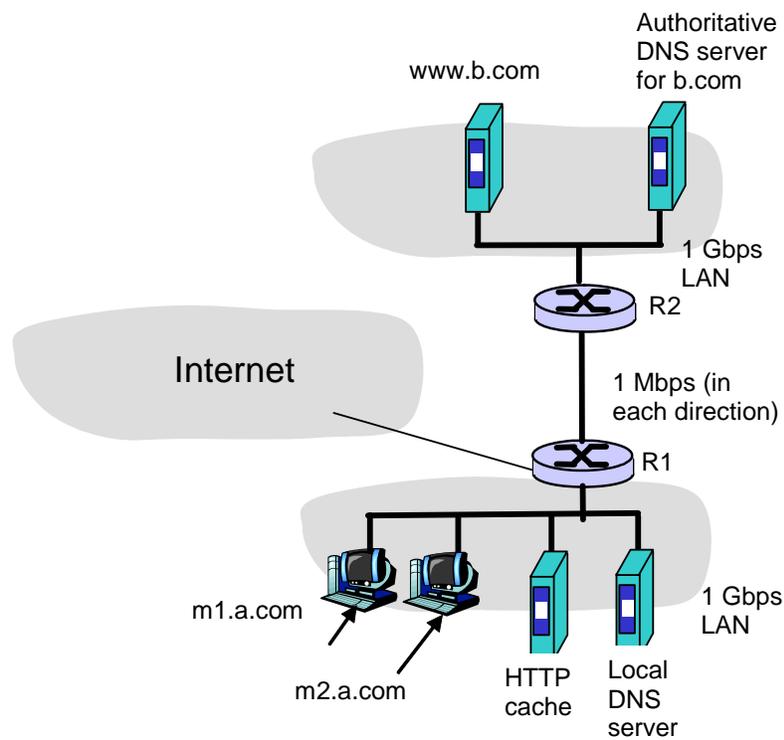
3. How many sockets are created at the server? How long the connection between the client and the server remain open?

Only one. There is no connection in this case.

4. Can the provided server implementation be used to handle multiple clients? Justify your answer.

No since no infinite loop is used.

Problem V: Putting it all together (20 minutes) [20 Points]



Consider the scenario depicted in the Figure above, in which two user machines called respectively `m1.a.com` and `m2.a.com` belong to a 1 Gbps high speed LAN. The figure shows also another 1 Gbps LAN in which resides a webservice called `www.b.com` and an authoritative DNS server for the `b.com` domain called `dns.b.com`. The two LANs are interconnected through a 1 Mbps access link. In what follows, assume that: a) every HTTP request in the bottom LAN is first directed to the HTTP cache, b) the HTTP cache in the bottom LAN is initially empty, and c) the cache of the local DNS server present in the bottom LAN is initially empty.

- Suppose the user at `m1.a.com` types in the URL `www.b.com/bigFile.htm` into his browser to retrieve a 1 Gbit file from `www.b.com`. Describe in details the step-by-step process of obtaining the file from the point that the URL is entered into the browser until the requested file is completely received by `m1.a.com`. In particular, list the sequence of DNS and HTTP messages that will flow through the network.
 - `m1.a.com` needs to resolve the name `www.b.com` to an IP address so it sends a DNS REQUEST message to its local DNS resolver (this takes no time given the assumptions below)
 - Local DNS server does not have any information so it contacts a root DNS server with a REQUEST message (this take 500 ms given the assumptions below)

- Root DNS server returns name of DNS Top Level Domain server for .com (this takes 500 ms given the assumptions below)
 - Local DNS server contacts .com TLD (this take 500 ms given the assumptions below)
 - TLD .com server returns authoritative name server for b.com (this takes 500 ms given the assumptions below)
 - Local DNS server contacts authoritative name server for b.com (this takes 100 ms given the assumptions below)
 - Authoritative name server for b.com returns IP address of `www.b1.com`. (this takes 100 ms given the assumptions below)
 - HTTP client sends HTTP GET message to `www.b1.com`, which it sends to the HTTP cache in the a.com network (this takes no time given the assumptions).
 - The HTTP cache does not find the requested document in its cache, so it sends the GET request to `www.b.com`. (this takes 100 ms given the assumptions below)
 - `www.b.com` receives the GE request. There is a 1 sec transmission delay to send the 1Gbps file from `www.b.com` to R2. If we assume that as soon as the first few bits of the file arrive at R1, that they are forwarded on the 1Mbps R2-to-R1 link, then this delay can be ignored.
 - The 1 Gbit file (in smaller packets or in a big chunk, that's not important here) is transmitted over the 1 Mbps link between R2 and R1. This takes 1000 seconds. There is an additional 100 ms propagation delay.
 - There is a 1 sec delay to send the 1Gbps file from R1 to the HTTP cache. If we assume that as soon as the first few bits of the file arrive at the cache, that they are forwarded to the cache, then this delay can be ignored.
 - There is a 1 sec delay to send the 1Gbps file from the HTTP cache to `m1.a.com`. If we assume that as soon as the first few bits of the file arrive at the cache, that they are forwarded to the cache, then this delay can be ignored.
 - The total delay is thus: $.5 + .5 + .5 + .5 + .1 + .1 + 1 + 1000 + 1 + 1 = 1105.2$ secs (1002.2 is also an OK answer).
2. How much time does it take to accomplish the steps you outlined in your answer to the first question? Show your work. In answering this question, you can make the following assumptions:
- a. DNS query/reply messages and HTTP REQUEST messages are very small compared to the size of the file, and thus their transmission times can be neglected (but not **their propagation times**).
 - b. Propagation delays within the LANs are small enough to be ignored.
 - c. The propagation time from router R1 to router R2 is **100 ms**.
 - d. The propagation delay from anywhere in **a.com** to any other site in the Internet (except **b.com**) is **500 ms**.

(See above for answer. Note that we have neglected to account for TCP hand-shaking delays for the HTTP exchanges!)

3. Now assume that machine **m2.a.com** makes a request to exactly the same URL that was requested by **m1.a.com**. List the sequence of DNS and HTTP messages underlying the step-by-step process of obtaining the file from the point that the URL is entered until the browser of **m2.a.com** receives the file in its entirety.
- **m2.a.com** needs to resolve the name **www.b.com** to an IP address so it sends a **DNS REQUEST** message to its local DNS resolver (this takes no time given the assumptions above)
 - The local DNS server looks in its cache and finds the IP address for **www.b.com**, since **m1.a.com** had just requested that that name be resolved, and returns the IP address to **m2.b.com**. (this takes no time given the assumptions above)
 - HTTP client at **m2.a.com** sends HTTP GET message to **www.b1.com**, which it sends to the HTTP cache in the **a.com** network (this takes no time given the assumptions).
 - The HTTP cache finds the requested document in its cache, so it sends a GET request with an If-Modified-Since to **www.b.com**. (this takes 100 ms given the assumptions)
 - **www.b.com** receives the GET request. The document has not changed, so **www.b.com** sends a short HTTP REPLY message to the HTTP cache in **a.com** indicating that the cached copy is valid. (this takes 100 ms given the assumptions)
 - There is a 1 sec delay to send the 1Gbps file from the HTTP cache to **m2.a.com**.

The total delay is thus: $.1 + .1 + 1 = 1.2$ secs

4. How much time does it take to accomplish the steps you outlined in your previous answer? Use the assumptions that were given in the second question.

(Answer: see above)

Problem VI: Web Caches (20 minutes) [18 Points]

Consider a scenario similar to the one given in Problem V. Assume that requests for objects from the **www.b.com** server are generated by the users of the bottom LAN at the rate of 3 requests/s. Suppose that the average object size is 500 Kbits and that the HTTP request messages are negligibly small and thus create no traffic in the network or in the access link.

- a) For this question, assume that there is no web cache in the bottom LAN. Find the total response time observed by a user of the bottom LAN requesting objects from **www.b.com** when the capacity of the access link connecting the bottom and upper LANs increases from 1 Mbps to 10 Mbps. For delay computation use $\frac{1}{\lambda} \frac{\rho}{1-\rho}$, where λ is the arrival rate and ρ is the traffic intensity. Under what condition would such a solution be unnecessary?

Traffic intensity on LAN1 = (3 requests/sec) × (0.5 Mbits/request)/1 Gbps = 0.0015

=> LAN1 delay = 0.5 ms

Whereas the traffic intensity on the access link is: 1.5/10 = 0.15

=> Access delay = 58.8 ms

Traffic intensity on LAN2 = 0.0015 => LAN2 delay = 0.5 ms

Total response time = 59.8 ms

- b) Find the total response time when a cache with a hit rate of 0.5 is installed in the bottom LAN. What is the minimum cache hit value that is needed to reduce response time?

The traffic intensity on the access link is reduced by 50% since 50% of the requests are satisfied within the institutional network. Thus, the average access delay is: (1/3*0.5) (0.75)/(1-0.75) = 2 sec

Traffic intensity on LAN2 = 0.0015/2 => LAN2 delay = 0.5 ms

=> Total response time = 0.5 × 0.5 ms + 0.5 × 2001 = 1000.75 ms