

Homework 2 – Application-Level Protocols

Assigned: Tuesday, February 1, 2011

Due: Tuesday, February 15, 2011 *at the beginning of class*

CS 455 – 100 points

CS 555 – 112 points

Note: All homework assignments must be done on your own, and your answers should be in your own words. The textbook and lecture notes may be used, but you should not copy verbatim from either of them. Consider the written homework as practice for the exam. *Use of previous years' assignments/solutions or the textbook's solutions manual is **not** permitted.*

Review Questions (50 points – 2 pts each)

1. Explain the main difference between the operation of a client-server application and a peer-to-peer application.
2. What are the two transport protocols offered by the Internet and what are their main characteristics?
3. What is the purpose of the TCP 3-way handshake?
4. What two elements are needed to address a communicating process?
5. What is a socket?
6. The UDP example server we looked at needed only one socket, but the TCP server needed two sockets. Why? If the TCP server were to support n simultaneous connections, each from a different client host, how many sockets would the TCP server need? Explain.
7. Does TCP preserve message boundaries? Explain your answer (include an example).
8. Consider the sample TCP and UDP client and server applications provided on Jan 25.
 - a. What happens if the TCP client is executed before the TCP server? Why?
 - b. What happens if the UDP client is executed before the UDP server? Why?
9. Is 2011 Tues 2/1 9:30am EDT in an acceptable date/time format for HTTP according to RFC 2616? If not, write the date/time in an acceptable format.
10. What is the default port used for HTTP servers?
11. Consider the following string of ASCII characters that were captured when a browser sent an HTTP request. The characters `<cr><lf>` are carriage return and line feed characters, respectively.

```
GET /~mweigle/unc/diss/abstract.html HTTP/1.1<cr><lf>
Host: www.cs.odu.edu<cr><lf>
Connection: keep-alive<cr><lf>
Accept:
application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0
.8,image/png,*/*;q=0.5<cr><lf>
User-Agent: Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10_6_6; en-US)
AppleWebKit/534.10 (KHTML, like Gecko) Chrome/8.0.552.237
Safari/534.10<cr><lf>
```

```
Accept-Encoding: gzip, deflate, sdch<cr><lf>
Accept-Language: en-US, en;q=0.8<cr><lf>
Accept-Charset: ISO-8859-1, utf-8;q=0.7, *;q=0.3<cr><lf>
<cr><lf>
```

Answer the following questions:

- a. What is the URL of the document requested by the browser?
 - b. What browser and version is used and on what operating system is it running?
12. The text below shows the HTTP response header sent from the server in response to the HTTP request in Question 11.

```
HTTP/1.1 200 OK<cr><lf>
Date: Thu, 27 Jan 2011 20:16:29 GMT<cr><lf>
Server: Apache/2.2.17 (Unix) PHP/5.3.5 mod_ssl/2.2.17
OpenSSL/0.9.8q<cr><lf>
Last-Modified: Thu, 06 Jan 2011 15:06:32 GMT<cr><lf>
ETag: "9cb-4992ed8d18e02"<cr><lf>
Accept-Ranges: bytes<cr><lf>
Content-Length: 2507<cr><lf>
Connection: close<cr><lf>
Content-Type: text/html<cr><lf>
<cr><lf>
```

Answer the following questions:

- a. When was the document last modified?
 - b. How many bytes are in the document being returned?
13. What is the difference between non-persistent HTTP connections and persistent HTTP connections?
14. What is the difference between persistent HTTP connections and persistent HTTP connections with pipelining?
15. Describe how Web caching can reduce the delay in receiving a requested object.
16. Telnet into the ODU-CS web server located at `www.cs.odu.edu` and request `/~mweigle/files/foo.txt`. Include in the request message an appropriate `If-modified-since:` header option line to force a response message with the 304 Not Modified status code. Submit a printout of the telnet session.
17. What does it mean that FTP uses out-of-band control?
18. What is the difference between active FTP and passive FTP?
19. Is SMTP secure (*i.e.*, by default, is information, including usernames and passwords, encrypted)?
20. SMTP requires that all data be encoded in 7-bit ASCII. What allows us to email non-ASCII attachments such as images or Word documents? How does it work?
21. Suppose Alice uses a webmail client (such as Gmail or Yahoo! Mail) to send an email message to her friend Bob who uses a traditional POP3 email client to read his mail (such as Eudora or Thunderbird). Discuss how the message gets from Alice's host to Bob's host,

including the series of application-layer protocols that are used to move the message between the two hosts. *You may draw a diagram to illustrate this.*

22. What is the difference between a computer's hostname and its IP address?
23. Does DNS name resolution occur before the TCP handshake or after the TCP handshake? Explain.
24. What is the difference between a CNAME DNS record and an A DNS record?
25. What is the purpose of an MX DNS record?

Problems (455 – 50 points, 555 – 62 points)

1. You are in charge of developing your company's website. Your boss is concerned about the download times for the main company webpage, `www.imsprt.com`. She wants to ensure that a typical user will be able to download the entire main webpage (including all objects) in less than 300 ms.

The current main webpage has a size of 1,000 bytes and includes the following embedded objects:

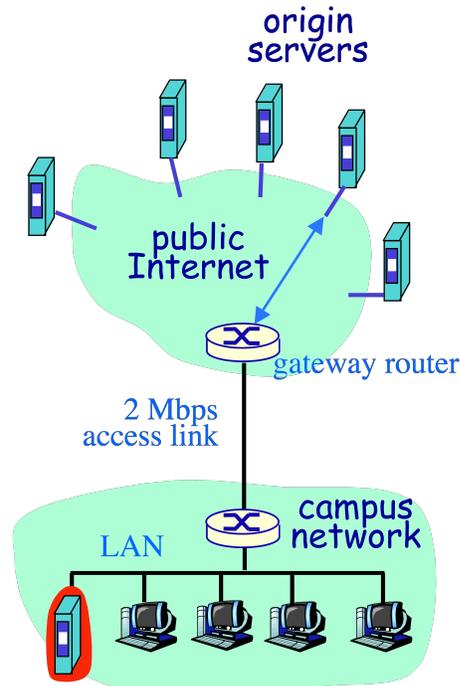
<code>http://www.imsprt.com/ads/banner.jpg</code>	500 bytes
<code>http://www.imsprt.com/pics/boss.jpg</code>	3,000 bytes
<code>http://www.fontmaster.com/img/sprt.jpg</code>	2,000 bytes

Your boss has specified the following parameters for a typical user:

- The round-trip time (RTT) from a typical user's machine to `www.imsprt.com` is 50 ms (propagation delay only).
 - The RTT from typical a user's machine to `www.fontmaster.com` is 30 ms (propagation delay only).
 - The typical user has an access link of 2 Mbps.
 - The access link to the servers at both `www.imsprt.com` and `www.fontmaster.com` is 10 Mbps.
 - The links in between the user's access link and the servers' access link are multiple Gbps, so transmission delays are negligible.
 - For parallel connections, as soon as one of the parallel connections finishes, a new connection can be started.
- a. [2 pts] Draw a diagram of the network described above, including link speeds and link propagation delays.
 - b. [4 pts] How many *round-trip times* (RTTs) are required to download the entire page using non-persistent connections? *Remember that RTTs involve propagation delay only.*
 - c. [6 pts] What is the *total download time* for the entire page (including transmission times) using non-persistent connections? Is your boss' requirement met?
 - d. [4 pts] How many *RTTs* are required to download the entire page using persistent connections (no pipelining)?
 - e. [6 pts] What is the *total download time* for the entire page (including transmission times) using persistent connections (no pipelining)? Is your boss' requirement met?

- f. [4 pts] How many *RTTs* are required to download the entire page using 2 parallel connections?
- g. [555-only, 6 pts] What is the *total download time* for the entire page (including transmission times) using 2 parallel non-persistent connections? Is your boss' requirement met?

2. [4 pts] Consider the following network:



- On average there are 35 objects downloaded per second.
- The average object size is 50,000 bits.

What is the traffic intensity on the 2 Mbps access link?

3. Refer back to the network used in Problem 2. Add the following parameters:

- A proxy cache with a 45% hit ratio has been installed on the campus network.
- It takes on average 15 ms to send a request to the proxy cache, check the cache, and receive the response.
- On average, the response time between the gateway router and any web server is 800 ms, which includes transmission time for web objects.

- [4 pts] What is the traffic intensity on the 2 Mbps access link?
- [555-only, 6 pts] What is the average response time for an object?

4. Use the Unix commands `nslookup` or `dig` to answer the following questions. *In addition to providing the answer, indicate what you typed on the command line to determine the answer.*

- [2 pts] What is the IP address of `www.lsu.edu`?
- [2 pts] What are the nameservers for the `lsu.edu` domain?
- [2 pts] What are the mail servers for the `lsu.edu` domain?

5. [8 pts] Use the Unix command `dig` and act as the local authoritative server for a client trying to resolve the IP address of `www.cs.odu.edu`.

- Use the `+norecurse` option to specify an iterative query.
- Use the `@nameserver` option to query a specific nameserver.
- Do not use the `+trace` option.

For each query, indicate what you typed on the command line.

Once the query has been resolved, indicate the cache entries that would have been stored at the local authoritative server.

Assume the cache contains the following entries upon startup:

d.root-servers.net	A	128.8.10.90
a.root-servers.net	A	198.41.0.4

6. [2 pts] Using the cache produced in Question 5, what nameserver would be the first to be contacted to resolve a query for `www.wm.edu`? Why?