

Homework 2 – Application-Level Protocols

Assigned: Thursday, September 17, 2009

Due: Tuesday, September 29, 2009 *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. *Use of previous years' assignments/solutions or the textbook's solutions manual is **not** permitted.*

Review Questions (56 points – 2 pts each part)

1. What is the difference between a computer's hostname and its IP address?
2. What two elements are needed to address a communicating process?
3. What is a socket?
4. What is the TCP 3-way handshake?
5. 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.
6. Does TCP preserve message boundaries? Explain your answer (include an example).
7. Consider the sample TCP and UDP client and server applications provided on **Sep 4**.
 - 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?
8. Is 2009 Tues 9/15 2:09pm EDT in an acceptable date/time format for HTTP according to RFC 2616? If not, write the date/time in an acceptable format.
9. What is the default port used for HTTP servers?
10. Consider the following string of ASCII characters that were captured when a browser sent an HTTP GET message. The characters `<cr>` and `<lf>` are carriage return and line feed characters, respectively. Answer the following questions and *indicate where in the HTTP GET message below you find the answer*.

```
GET /~mweigle/cs312/index.html HTTP/1.1<cr><lf>
Host: www.cs.odu.edu<cr><lf>
User-Agent: Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10.5; en-US;
rv:1.9.0.4) Gecko/2008102920 Firefox/3.0.4<cr><lf>
Accept: text/html, application/xhtml+xml,
application/xml;q=0.9,*/*;q=0.8<cr><lf>
Accept-Language: en-us,en;q=0.5<cr><lf>
Accept-Encoding: gzip,deflate<cr><lf>
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7<cr><lf>
Keep-Alive: 300<cr><lf>
Connection: keep-alive<cr><lf><cr><lf>
```

- a. What is the URL of the document requested by the browser?
- b. What version of HTTP is the browser running?
- c. What browser and version is used and on what operating system is it running?
- d. Does the browser request a non-persistent or a persistent connection?

11. The text below shows the reply sent from the server in response to the HTTP GET message in Question 10. Answer the following questions and *indicate where in the message below you find the answer*.

```
HTTP/1.1 200 OK<cr><lf>
Date: Tue, 02 Dec 2008 16:07:40 GMT<cr><lf>
Server: Apache/2.2.9 (Unix) PHP/5.2.6 mod_ssl/2.2.9
OpenSSL/0.9.8g<cr><lf>
Last-Modified: Fri, 14 Nov 2008 22:58:23 GMT<cr><lf>
ETag: "41b224-275-45bae2a8ff5ad"<cr><lf>
Accept-Ranges: bytes<cr><lf>
Content-Length: 629<cr><lf>
Connection: close<cr><lf>
Content-Type: text/html<cr><lf><cr><lf>
<!doctype html public "-//w3c//dtd html 4.0 transitional//en">
[much more document text following here (not shown)]
```

- a. Was the server able to successfully find the document?
 - b. When was the document last modified?
 - c. How many bytes are there in the document being returned?
 - d. What are the first 5 bytes of the document being returned?
12. Explain the difference between non-persistent HTTP connections, persistent HTTP connections, parallel HTTP connections, and persistent HTTP connections with pipelining.
13. Describe how Web caching can reduce the delay in receiving a requested object. Will Web caching reduce the delay for all objects requested by a user or for only some of the objects? Why?
14. 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.
15. What does it mean that FTP uses out-of-band control?
16. What is the difference between active FTP and passive FTP?
17. Is FTP secure (*i.e.*, by default, is information, including usernames and passwords, encrypted)?
18. 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.

19. Does DNS name resolution occur before the TCP handshake or after the TCP handshake? Why or why not?
20. What is the purpose of having a single hostname resolve to multiple IP addresses?
21. Explain the main difference between the operation of a client-server application and a peer-to-peer application.

Problems (455 – 44 points, 555 – 56 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. [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.*
 - b. [6 pts] What is the total download time for the entire page (including transmission times) using non-persistent connections? Is your boss' requirement met?
 - c. [4 pts] How many RTTs are required to download the entire page using persistent connections (no pipelining)?
 - d. [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?
 - e. [4 pts] How many RTTs are required to download the entire page using 2 parallel connections?
 - f. [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 RTT between the gateway router and any web server is 800 ms.

a. [4 pts] **What is the traffic intensity on the 2 Mbps access link?**

b. [555-only, 6 pts] **What is the average response time for an object?**

4. [12 pts] For each object given below, indicate the DNS servers that the local authoritative server contacts in order to resolve the hostnames to IP addresses. Also, describe the IP address that the local authoritative server obtains from the DNS server it contacts (*i.e.*, is it for the desired host or for a particular nameserver). Assume that the hostnames are resolved in the order given, and that caching is used.

You do not need to provide actual IP addresses, so there is no need to use dig or nslookup

- `http://sportsillustrated.cnn.com/football/ncaa`
- `http://i.a.cnn.net/si/element/js/3.0/main.js`
- `http://ad.doubleclick.net/ad/3475.si`
- `http://t.cnn.net/si/images/1.gif`