

Homework 3 – Application-Level Protocols

Assigned: Tuesday, October 5, 2010

Due: Thursday, October 14, 2010 *at the beginning of class*

Submit either a hard-copy or via Blackboard

100 points

Note: All homework assignments must be done on your own, and your answers should be in your own words. The lecture notes may be used, but you should not copy verbatim from either of them. *Use of previous years' assignments/solutions is **not** permitted.*

Review Questions (70 points – 5 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 two elements are needed to address a communicating process?
4. What is the purpose of the TCP 3-way handshake?
5. What is the default port used for HTTP servers?
6. What is the difference between non-persistent HTTP connections and persistent HTTP connections?
7. What is the difference between persistent HTTP connections and persistent HTTP connections with pipelining?
8. What is the minimum number of round-trip times (RTTs) needed to download a HTML-only webpage? Explain.
9. 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 screen capture of the telnet session.
10. 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.
11. 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?
12. What is the difference between a computer's hostname and its IP address?
13. Does DNS name resolution occur before the TCP handshake or after the TCP handshake? Explain.
14. What is the purpose of having a single hostname resolve to multiple IP addresses?

Problems (30 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 base page 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.
 - You must consider the transmission delays over the user's access link and the server's access link.
- a. **[3 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. **[3 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?*
2. For each object given below, indicate the DNS servers that the local authoritative server contacts in order to resolve the hostnames of each web server to an IP address. Also, describe the IP address that the local authoritative server obtains from the DNS server it contacts (*i.e.*, is the IP address obtained for the desired host or for a particular nameserver?). Assume that the hostnames are resolved in the order given and that caching is used, but that the cache starts out empty.

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

- a. **[3 pts]** `http://www.cnn.com/`
- b. **[3 pts]** `http://i.cdn.turner.com/cnn/.element/css/3.0/common.css`
- c. **[3 pts]** `http://i2.cdn.turner.com/cnn/2010/images/10/04/tzwide.alex.jpg`
- d. **[3 pts]** `http://marquee.blogs.cnn.com/2010/10/04/wonder-woman-back-to-tv/`