```
<head><lf>  <meta http-equiv="Content-Type"
content="text/html; charset=iso-8859-1"><lf> <meta
name="GENERATOR" content="Mozilla/4.79 [en] (Windows NT
5.0; U) Netscape]"><lf>  <title>CMPSCI 453 / 591 /
NTU-ST550A Spring 2005 homepage</title><lf></head><lf>
<much more document text following here (not shown)>
```

a. Was the server able to successfully find the document or not? What time was the document reply provided?

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? Did the server agree to a persistent connection?

P6. Obtain the HTTP/1.1 specification (RFC 2616). Answer the following questions:

a. Explain the mechanism used for signaling between the client and server to indicate that a persistent connection is being closed. Can the client, the server, or both signal the close of a connection?

b. What encryption services are provided by HTTP?

P7. 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 receives the IP address from DNS; the successive visits incur an RTT of $RTT_1, \ldots, RTT_n$. 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 from when the client clicks on the link until the client receives the object?

P8. Referring to Problem P7, suppose the HTML file references three very small objects on the same server. Neglecting transmission times, how much time elapses with

a. Non-persistent HTTP with no parallel TCP connections?

b. Non-persistent HTTP with parallel connections?

c. Persistent HTTP?

P9. Consider Figure 2.12, for which there is an institutional network connected to the Internet. Suppose that the average object size is 900,000 bits and that the average request rate from the institution's browsers to the origin servers is 15 requests per second. Also suppose that the amount of time it takes from when the router on the Internet side of the access link forwards an HTTP request until it receives the response is two seconds on average (see Section 2.2.5).

b. Now suppose that Peer X abruptly disconnects from the Internet without notifying its five neighbors that it is closing the TCP connections. What would happen?

P21. In this problem we explore the reverse-path routing of the query hit messages in query flooding. Suppose that Alice issues a query message. Further suppose that Bob receives the query messages (which may have been forwarded by several intermediate peers) and has a file that matches the query.

a. Recall that when a peer has a matching file, it sends a query hit message along the reverse path of the corresponding query message. An alternative design would be for Bob to establish a direct TCP connection with Alice and send the query hit message over this connection. What are the advantages and disadvantages of such an alternative design?

b. When the peer Alice generates a query message, it inserts a unique ID in the message's MessageID field. When the peer Bob has a match, it generates a query hit message using the same MessageID as the query message. Describe how peers can use the MessageID field and local routing tables to accomplish reverse-path routing.

c. An alternative approach, which does not use message identifiers, is as follows. When a query message reaches a peer, before forwarding the message, the peer augments the query message with its IP address. Describe how peers can use this mechanism to accomplish reverse-path routing.

P22. In this problem we explore designing a hierarchical overlay that has ordinary peers, super peers, and super-duper peers.

a. Suppose each super-duper peer is roughly responsible for 200 super peers, and each super peer is roughly responsible for 200 ordinary peers. How many super-duper peers would be necessary for a network of four million peers?

b. What information might each super peer store? What information might each super-duper peer store? How might searches be performed in such a three-tier design?

P23. Consider query flooding, as discussed in Section 2.6. Suppose that each peer is connected to at most N neighbors in the overlay network. Also suppose that the node-count field is initially set to K. Suppose Alice makes a query. Find an upper bound on the number of query messages that are sent into the overlay network.

P24. Install and compile the Java programs TCPClient and UDPClient on one host and TCPServer and UDPServer on another host.

a. Suppose you run TCPClient before you run TCPServer. What happens? Why?