

EP2120 Internetworking/Internetteknik IK2218 Internets Protokoll och Principer

Homework Assignment 3

Solutions due 23:00, September 30, 2016

Review due 23:00, October 4, 2016

Problems

1. Socket Interface (15 p)

The pseudo-code sample below (with most details omitted) describes an application that uses the socket interface (API) for communication.

```
s = socket(...);
bind(s, ...);
listen(s, ...);
while true {
    t = accept(s, ...);
    recv(t, ...);
    ProcessRequest(...);
    send(t, ...);
    close(t);
}
```

- a) Is the sample code for a client or a server? Does it use TCP or UDP? Explain your answer. (5 p)
- b) The textbook gives two examples of communication using the socket interface: 1) connection-oriented, concurrent communication and 2) connectionless, iterative communication. Characterize the communication in the sample code using this terminology. (5 p)
- c) In practice, this kind of communication is not frequently used. What is the main limitation? (5 p)

2. Web (35 p)

Suppose that you click on a link on a web page, which causes the following HTTP request to be sent.

```
GET /swordfish HTTP/1.1
Host: www.feathers.org
Connection: keep-alive
Cache-Control: max-age=0
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_8_5)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/29.0.1547.76
Safari/537.36
Accept-Encoding: gzip,deflate,sdch
Accept-Language: en-US,en;q=0.8
```

- a) Which web document does the browser request? Answer by giving the URL. (3 p)

- b) Which version of HTTP is the browser using? What kind of connection is requested, persistent or non-persistent? (Note that the textbook does not cover header fields related to persistence, so you may want to consult other resources, on the Internet for example.) (3 p)

The server gives the following response:


```
HTTP/1.1 302 Found
Date: Fri, 23 Sep 2016 12:34:56 GMT
Server: Apache/2.2.3 (Red Hat)
Location: https://www.feathers.org/swordfish
Content-Length: 286
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

more data...

- c) What kind of connection does the server accept, persistent or non-persistent? (3 p)
- d) The server seems to give a positive response (“302 Found”), but this is not the normal response for a successful request. Answer the following three questions: Why does the server give this response? What is the browser supposed to do? Why do you think the server has been configured to respond in this way? (6 p)
- e) Assume that you instead receive a “200 OK” response with an HTML object that contains three other objects (it could be a web page with three images, for instance). Suppose that the client you are using is strictly sequential, so that it fetches one object at a time. How long time does it take from that you click on the link until the entire document has been received? Explain your solution.
The round-trip time between client and server is RTT . All objects are very small and the connection is fast, so transmission time is negligible, and so is processing time on the server. Consider the following two cases:
1. Non-persistent HTTP is used.
 2. Persistent HTTP is used. (10 p)
- f) Suppose instead that the client is configured to use as much parallelism as possible. How long time would it then take for the two cases?
1. Non-persistent HTTP is used.
 2. Persistent HTTP is used. (Hint: *HTTP pipelining*) (10 p)

3. Mail (25 p)

Consider the scenario when Alice (left) sends an email to Bob (right). There are two intermediate systems: *outgoing mail server* and *incoming mail server*.



Consider the following questions:

- a) What is the purpose of the outgoing mail server? Discuss what would happen if it were removed, so that Alice's user agent directly connects to the incoming mail server. What would be the limitations (if any)? (10 p)

- b) What is the purpose of the incoming mail server? Discuss what would happen if it were removed, so that the outgoing mail server connects directly to Bob's user agent. What would be the limitations (if any)? (10 p)
- c) For each of the connections 1–3, explain what application-layer protocol(s) are used for the mail transfer. (5 p)

4. DNS (25 p)

You use the “dig” lookup tool to get the IP address of KTH's web server “www.kth.se”. You want to try different name servers, so you specify the DNS server as an argument to dig (the server's IP address prepended with ‘@’). You run the following three commands:

```
dig @192.36.135.107 www.kth.se
dig @192.36.148.17 www.kth.se
dig @130.237.72.250 www.kth.se
```

- a) Explain the results: Describe the responses from the three DNS servers. What do the responses say? We distinguish between four kinds of name servers: root, TLD, authoritative, and local. You can deduce just from studying the responses what kind of DNS server it is (the flags are useful here, among other things). For each of the three cases, describe the kind of name server that is responding to your query. Explain what the response contains. Also, for each name server, explain whether or not the answer contains the IP address you are looking for. *You only need to discuss the responses at a general level – you should not discuss or describe the details of the messages, the different fields, their contents, etc. However, you probably need to study the responses carefully in order to be able to explain them.* (15 p)
- b) Organizations often have multiple name servers for their domains. What is the purpose of having multiple name servers? How are they kept synchronized? (5 p)
- c) It is sometimes useful to have more than one name in DNS for the same computer. Is it possible for a computer to have two names under different top-level domains? For example, one name in the “com.” TLD and one in “se.”. If so, give an example when this might be useful. (5 p)