CSci551 Spring 2006 Thursday Section Homework 4

Assigned: April 17, 2006. Due: noon, April 28, 2006.

You are welcome to discuss your homework with other students, but each student is expected write his or her final answer independently.

Answers should be *short* and *to the point*. Complete sentences are not required. It is possible answer the homework in one page of text, and if your answers are much more than that you should ask yourself if you're answering clearly and succinctly. We reserve the right to deduct points on answers that are too long or that miss the point.

This homework must be submitted electronically. Choice of formats is:

- Simple text (ASCII). This format is *strongly* preferred since it can be most easily marked up and returned to you.
- PDF. Please use PDF *only* if you have figures in your text.
- Postscript or HTML. Use these only if you have figures and cannot generate PDF.

Note that MS-Word or other proprietary formats are *not* accepted. See the course web page for details about how to generate postscript or PDF if you start with these formats. (See the course web site for information about how to generate PDF.)

To submit your homework, upload it to aludra.usc.edu or nunki.usc.edu, then use the submit comment as you did for the project:

```
% submit -user csci551 -tag hw4 file1 [ file2 ... ]
```

You should have one file for the body of your homework. You may have additional files if you generate html. If you have multiple files, please list them all on the command line separately (do *not* combine them ahead of time in a tar or zip file). Just a reminder: your name and student id should appear at the top of the first page of your response.

Because of the limited time left in the semester, this homework is a short homework, worth half the credit of the other homework assignments, and due in less time.

1: The Domain Name System is the distributed database that maps DNS names to IP addresses.

"Scalability" of the DNS is important. a) and b) and c) (1pt each part) What are three *different* dimensions in which the DNS must scale? (Make sure you label the parts of your answer (a), (b), and (c).)

The DNS system is designed to support caching. Traditionally Unix systems cached requests at the user's computer, while Windows and Macintosh clients did not cache requests. d) (1pt) Would it be advantageous to cache DNS requests at the user's computer or not? (Just answer "yes, useful" or "no, not useful"). e) (2pts) Why? (Justify your answer to part (d).)

One promise of a future Internet architectures is that each user has access to a huge amount of bandwidth, maybe 10Gb/s to the desk (compared to 100Mb/s common today). f) (1pt) Would the addition of 100 times more bandwidth than common today change your answer to (d)? (Just answer, "yes, my answer changes" or "no, my answer doesn't change".) g) (2pts) Why or why not? (Justify your answer to part (g).)

DNS servers are typically assigned as part of a DHCP request. Today one is typically assigned a local DNS server. At USC you may get a campus-wide server, or if you connect inside a lab, you might get a DNS server that's just for a department or a school. A DNS server handles interacting with remote DNS servers to retrieve a name, then it caches the response and uses it for later requests for the same name. Suppose you have a choice between using the School of Engineering server (shared by all the staff and students) or the campus-wide server. Based on what you learned from other papers in class: h) (1pt) Should you expect a much better cache hit rate using the campus-wide server? (Just say "yes, campus has higher hit rate" or "no, campus does not have higher hit rate".) i) (2pts) Why? (Justify your answer for part (f).)

Answer: (1pt/each) a/b/c) any 3 of: number of names, number of requests, number of servers, update rate, round trip time. Other answers are possible.

(1pt) d) YES, it's helpful. (2pts) e) Because it would reduce the load on the central servers. OR because clients usually make multiple DNS requests to the same host.

(1pt) f) NO, the answer to (d) should not change with more bandwidth. (2pts) g) DNS requests are all small, so they are limited by *latency*, not bandwidth. For the most part, all the bandwidth in the world cannot make them go faster.

(1pt) h) NO, cache hit rate would be about the same. (2pts) i) Because, based on our experience with web hit rates in Wolman et al, we know that hit rates top out at 2000 people for web-like queries.

2: For this semester: (a) (1pts) What are the two papers you learned the most from (or, if you prefer, the two you found the most interesting)? (b) (2pts) What was it that you learned (or that was most interesting)? (c) (1pts) What were the two papers you thought were the weakest or had the least to offer? (d) (2pts) For one of those papers, what was weak about it and how would you change it to improve it? (Just suggest a fix in a sentence or two, you don't need to actually make the change!) (e) (2pt) If there was one subject you wanted to spend more or less time on, what was it (if not, you don't need to answer this part of the question)? Please identify papers by their names in the syllabus like [Clark88a].

(You're welcome to consider any of the papers in the syllabus for this question, not just the ones we've read so far.)

Answer: Answers will vary.