

Homework 2

ITCS-4010/5010: Cloud Computing for Data Analysis

Due: Wednesday, February 25, 2009

Homeworks are due at the beginning of class on February 25, and are to be done individually. Homeworks will be graded on the basis of clarity and legibility. See course syllabus for late submission policy.

1. The *seek time* is the time needed to position the disk head in a new position. The *transfer time per byte* is the rate of transfer from disk to memory when the head is in the right position. Assume the seek time is 5 milliseconds and the transfer time per byte is 0.02 microseconds. What is the time to transfer 64 MB from disk to memory if it is stored as one chunk, versus if it is stored in 100 noncontiguous chunks (requiring the disk head to move up to 100 times)?
2. If in a non-fault-tolerant system with 1000 nodes, each node has an 99.9% uptime, what is the percentage of time that all nodes in the system are up?
3. You must submit a reading report on the following paper, which describes the original Google search engine.

The Anatomy of a Large-Scale Hypertextual Web Search Engine, Sergey Brin and Lawrence Page. Proceedings of the Seventh World Wide Web Conference (WWW7), pp. 107-117, 1998.

The paper is available off the course web page <http://www.cs.uncc.edu/~sakella/cloud/> or at <http://infolab.stanford.edu/~backrub/google.html>.

For this homework, your reading report may be up to **two** pages in length, and must be written on a computer.

- (a) Your reading report should provide a concise summary of the paper (e.g., what was the problem addressed, what did they do and what was their approach, what were their contributions).
- (b) It should also provide a brief technical critique of the paper (e.g., is the problem important, what are the merits of the approach, what could have been done better or described more clearly, what might the next steps be).