

CS3114 Fall 2012 Homework Assignment 3
Due Thursday, December 13 at 11pm
50 points

Submissions for this assignment must obey the course General Assignment Guidelines. You can find a link to this at the top of the Assignments page on the course website. You may draw the trees in Question 3 and the graph in Question 4. Everything else must be typeset.

1. Suppose that a record is 32 bytes, a block is 1024 bytes (thus, there are 32 records per block), and that working memory is 1MB (there is also additional space available for I/O buffers, program variables, etc.). What is the *expected* size for the largest file that can be merged using replacement selection followed by a *single* pass of multiway merge? Explain how you got your answer.

2. A typical disk drive from 2004 has the following specifications. The total storage is approximately 120GB on 6 platter surfaces or 20GB/platter. Each platter has 16K tracks with 2560 sectors/track (a sector holds 512 bytes) and 8 sectors/cluster. The disk turns at 7200 rpm. The track-to-track seek time is 2.0 ms, and the average seek time is 10.0 ms. Calculate the time required to read a 16MB file assuming

- (a) The file is stored on a series of contiguous tracks, as few tracks as possible.
- (b) The file is spread randomly across the disk in 4KB clusters.

Show your calculations.

3. B⁺-tree Questions

- (a) Show the result of inserting the values 1, 2, 3, 4, 5, and 6 (in that order) into the B⁺-tree of Figure 10.18 in the textbook.
- (b) Show the result of deleting first the value 31 and then the value 52 from the B⁺-tree of Figure 10.18 in the textbook.

4. Show the shortest paths generated by running Dijkstra's shortest-paths algorithm on the graph of Figure 11.26, beginning at Vertex 5. Show the D values as each vertex is processed, as in Figure 11.19.

5. Write an algorithm to determine whether a directed graph of $|V|$ vertices contains a cycle. Your algorithm should run in $\Theta(|V| + |E|)$ time.