CS3114 Spring 2011 Homework Assignment 3
Due Wednesday, May 4 at 11pm
60 points

1. Assume that a disk drive is configured as follows. The total storage is approximately 1024GB divided among 16 surfaces. Each surface has 32K tracks, there are 4K sectors/track, 512 bytes/sector, 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 9.0 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.

2. Assume that a computer system has disk blocks of 1024 bytes, and that you are storing records that have 4-byte keys and 4-byte data fields. The records are sorted and packed sequentially into the disk file.

(a) Assume that a linear index uses 4 bytes to store the key and 4 bytes to store the block ID for the associated records. What is the greatest number of records that can be stored in the file if a linear index of size 256KB is used?

(b) What is the greatest number of records that can be stored in the file if the linear index is also stored on disk (and thus its size is limited only by the second-level index) when using a second-level index of 1024 bytes (i.e., 256 key values) as illustrated by Figure 10.2? Each element of the second-level index references the smallest key value for a disk block of the linear index.

3. You are given a series of records whose keys are letters. The records are inserted in the following order: C, S, D, T, A, M, P, I, B, W, N, G, U, R, K, E, H, O, L, J. Show the B^+-tree of order four that results from inserting these records. Assume that the leaf nodes are capable of storing up to three records.

4. Assume that you have a B^+-tree whose internal nodes can store up to 50 children and whose leaf nodes can store up to 50 records. What are the minimum and maximum number of records that can be stored by the B^+-tree for 1, 2, 3, 4, and 5 levels?

5. Show the shortest paths generated by running Dijkstra's shortest-paths algorithm on the graph of Figure 11.25, beginning at Vertex 4. Show the D values as each vertex is processed, as in Figure 11.18.

6. Questions on Minimal Cost Spanning Trees.

(a) List the order in which the edges of the graph in Figure 11.25 are visited when running Prim's MST algorithm starting at Vertex 3. Show the final MST.

(b) List the order in which the edges of the graph in Figure 11.25 are visited when running Kruskal's MST algorithm. Each time an edge is added to the MST, show the result on the equivalence array, (e.g., show the array as in Figure 6.7).