

CS 245 Database Systems Principles – Summer 2001

Assignment 3

- Due Date: Monday July 23rd 2001, 5 pm.
 - Submission through
 - Box Placed outside Gates 412.
 - Email solution to masood@cs.stanford.edu (only pdf, ps or plain text files).
 - SITN homework delivery.
 - The deadline is hard, No Late days.
 - Do not forget to write your Leland Ids at the start of your solution.
 - State all assumptions.
 - Email questions to cs245-staff@lists.stanford.edu
-

Problem 1

20 Points

Consider an extensible hash structure where buckets can hold up to three records. Initially, the structure is empty. The hashed key values are the following (in binary):

h(Biochemistry)	= 00001
h(Biology)	= 00011
h(Chemistry)	= 00110
h(Computer Science)	= 01110
h(Economics)	= 01111
h(Electrical Eng)	= 10001
h(History)	= 10101
h(Math)	= 10111
h(Philosophy)	= 11000
h(Physics)	= 11001
h(Psychology)	= 11101
h(Statistics)	= 11111

- a) First, suppose we insert the values in the order given above. Show the structure after all the values have been inserted.
- b) Once again, assume we start with an empty structure. Now suppose we insert all the values in the following order:

Physics, Chemistry, Biology, Math, Electrical Eng, Computer Science, Psychology, Statistics, History, Philosophy, Economics, Biochemistry.

Show the structure after all the values have been inserted.

Problem 2

20 Points


Consider a linear hash structure where buckets can hold up to two records. Initially the structure is empty. Assume that the utilization threshold value is 100% (i.e., when the average number of keys per non-overflow slot is greater than 1, we allocate another bucket). The hashed keys are the same as in Problem 1.

- a) Show the structure after all the values have been inserted in the order they are given (Problem 1 a).

Problem 3

15 Points

Consider the following Multi-key B Tree over keys that are points in two dimensions (i.e. x,y co-ordinates). This tree has one root and four leaf nodes.




Note: The ordering is determined by first comparing the x co-ordinate. If the x co-ordinates are identical, then the order is determined by comparing the y co-ordinates. For example, (2,1) < (12,2) and (12,2) < (12,4).

In this problem, rectangles in two dimensions are represented by a pair of two-dimensional points. A rectangle representation has the form: (x1,y1)-(x2,y2). The two points define the lower-left and upper-right corners of the rectangle, respectively.

- a) How many nodes (blocks) of the index will be read while executing the following lookups:
 - i) X <= 10
 - ii) X <= 9
 - iii) Y <= 10
 - iv) (X,Y) in the rectangle defined by (5,5)-(15,15).
 - v) (X,Y) in the rectangle defined by (2,2)-(8,17).

vi) (X,Y) in the rectangle defined by (4,2)-(19,8).

Now Consider a R-Tree index over the same set of points:



Recall that the non-leaf nodes of an R-Tree have one key per child pointer. Keys in non-leaf nodes are rectangles (represented by a pair of points).

b) How many nodes (blocks) of the R-Tree index will be read while executing the six lookups given in part (a)?

Problem 4

10 Points

Given: $M = \#$ of memory buffer pages = 20 blocks
 Relation R with $B(R) = \#$ of blocks = 520

Lets say we want to sort R using the external merge sort algorithm (without the heap-sort in memory optimization).

- How long will the runs be after the first pass?
- How many runs will there be after the first pass?
- Can the sort be performed in two passes?
- If not, how much memory would we need in order to sort R in two passes?

Problem 5

20 Points

Trace the execution of the merging phase of the merge-join algorithm. Assume that the two relations have already been fully sorted. Show a sequence of values for “i” and “j” as used in the pseudo-code in lecture notes 6. Indicate which (i, j) pairs generate a result (joined) tuple.

The sorted relations R and S have the following keys:

R: 1, 2, 3, 3, 4, 5, 5, 6, 8
 S: 1, 1, 3, 4, 5, 5, 7, 8, 9, 10

Problem 6**15 Points**

Consider a hash-join of two relations R and S having $B(R) = 1000$ and $B(S) = 500$.

The values in R and S are skewed such that the hash function assigns twice as many tuples to even-numbered hash buckets as to odd-numbered buckets.

How much memory would be required to perform the join in two passes?