

CPS 216 Spring 2004

Homework #2

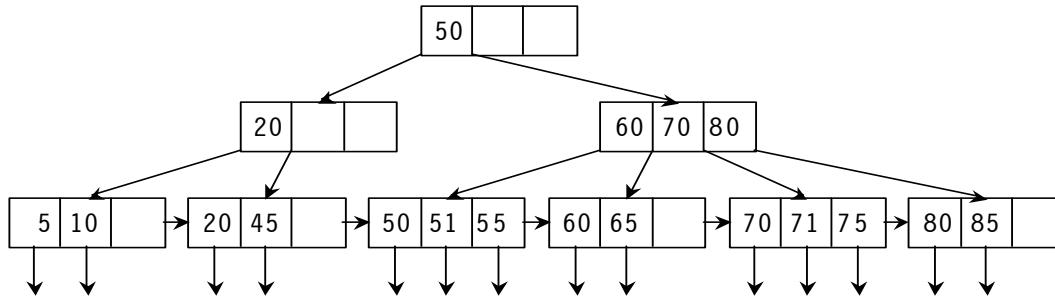
Assigned: Thursday, February 12

Due: Thursday, February 26

Note: This is a long homework. Start early!

Problem 1 (10 points).

Consider the following B⁺-tree with a maximum fan-out of 4.



For all questions below, assume that you always start with the tree shown above—*not* the result tree you get for the previous question.

- (a) Show the result tree after inserting 49.
- (b) Show the result tree after deleting 50.
- (c) Show the result tree after inserting 72.
- (d) Show the result tree after deleting 5.
- (e) What is the minimum number of keys you must delete for this tree to shrink down to two levels? Show the sequence of deletions.

Problem 2 (6 points).

Suppose keys are hashed to 4-bit sequences, and each block can hold three records. We start with a hash table with two empty blocks (corresponding to 0 and 1), and insert 16 records with keys 0000, 0001, 0010, ..., 1111, in order. Show the final state of the index:

- (a) If the index is based on extensible hashing.
- (b) If the index is based on linear hashing, with a capacity threshold of 100%. Here, we define capacity to be (actual number of records indexed) / (maximum number of records that can be held by primary buckets).

Problem 3 (6 points).

Consider a secondary B⁺-tree index. Each internal node contains m child pointers, and each leaf node contains m pointers to records. There are a total of N records indexed by the B⁺-tree. We wish to choose the value of m that will minimize search times on a particular disk drive, given the following information:

- (1) For the disk holding the index, the time to read a given block into memory is given by $(30 + 0.01m)$ milliseconds. The 30 milliseconds represent the seek time and the rotational delay of the read; the $0.01m$ milliseconds is the transfer time.
- (2) Once the block is in memory, a binary search is used to find the right pointer. The time to process a block in main memory is $(a + b \log_2 m)$ milliseconds, for some constants a and b .
- (3) The main memory time constant a is negligible compared to the disk seek time and rotational delay of 30 milliseconds.
- (4) You can assume that the index is completely full and perfectly balanced, and that N is conveniently a power of m .

Questions:

- (a) What value of m minimizes the time to search for a given record? An approximate answer is okay.
- (b) What happens as the disk latency (30 milliseconds) decreases? For instance, if this constant is cut in half, how does the optimal m value change?

Problem 4 (8 points).

To build a hash index for a multi-attribute search key, we can use an approach called partitioned hashing. The partitioned hash function is really a list of hash functions, one for each attribute in the search key. Suppose that we wish to build a partitioned hash index on $R(A, B)$ with 2^n buckets numbered 0 to $2^n - 1$. In this case, the partitioned hash function consists of two hash functions h_A and h_B . Hash function h_A takes a value of A as input and produces a result with n_A bits, and h_B takes a value of B as input and produces a result with $(n - n_A)$ bits. The two results are concatenated together to produce the result of the partitioned hash function, which is then used to index the buckets. To locate records with $A = a$ and $B = b$, we simply go directly to the bucket numbered $h_A(a)h_B(b)$ (in binary).

- (a) Which buckets do we have to examine in order to locate all records with $A = a$?
- (b) Suppose we are given a query mix. Each query in this mix will either ask for records with a given value of A , or it will ask for records with a given value of B (but never both). With probability p , the value of A will be specified. Give a formula in terms of n , n_A , and p for the expected number of buckets that must be examined to answer a random query from the mix.
- (c) Find the value of n_A , as a function of n and p , that minimizes the expected number of buckets examined per query.

Problem 5 (5 points).

The following question is based on the paper “Generalized Search Trees for Database Systems,” by Hellerstein et al. Suppose we want to use GiST to index ranges of the form $[x, y)$, where x and y are integers. These ranges may overlap with each other. The queries are of the form “find all ranges that overlaps with $[a, b)$.” Discuss briefly how you would implement the six basic methods required by GiST, and whether you should set *IsOrdered* to true and define *Compare*. Highlight the differences between your implementation and the B⁺-tree implementation described in the paper.

Problem 6 (12 points).

Consider the join $R \bowtie_{R.A = S.B} S$, given the following information about the tables to be joined. The cost metric is the number of disk I/O's and the cost of writing out the result should be uniformly ignored.

- R contains 10,000 rows and has 10 rows per page.
 - S contains 2,000 rows and also has 10 rows per page.
 - $S.B$ is a key of S .
 - Both tables are stored compactly on disk in no particular order.
 - No indexes are available.
 - 52 memory blocks are available for query processing.
- (a) What is the expected cost of joining R and S using a block-based nested-loop join, with R as the outer table?
- (b) What is the expected cost of joining R and S using a block-based nested-loop join, with S as the outer table?
- (c) What is the expected cost of joining R and S using a sort-merge join? What is the minimum number of memory blocks required for this cost to remain unchanged?
- (d) What is the expected cost of joining R and S using a hash join? What is the minimum number of memory blocks required for this cost to remain unchanged?

Problem 7 (5 points).

Consider a hash join of two tables R and S . R occupies 20000 blocks and S occupies 15000 blocks. Unfortunately, we have an imperfect hash function that consistently assigns twice as many tuples to even-numbered partitions as to odd-numbered partitions (although all even-numbered partitions are of the same size, and so are odd-numbered partitions).

- (a) Using this hash function, what is the minimum amount of memory required for the basic hash join algorithm described in lecture to complete in two passes?
- (b) Can you come up with a strategy to make the required amount of memory nearly as low as the amount required by a perfect hash function?

Problem 8 (8 points).

The *mode* of a list of values is the most common (frequent) value. There can be more than one mode for a list. For example, the list 1, 2, 2, 3, 4, 4, 2, 1, 1 has two modes 1 and 2.

- (a) Assume that no index is available. Design the best algorithm you can come up with to compute one mode (any one) of a column in a table. Analyze the worst-case performance of your algorithm in terms of number of I/O's.
- (b) Consider the following query that compute one mode for a subset of a table:
SELECT mode(r.A) FROM R r WHERE r.B > 0;
Suppose that there is a B⁺-tree index on R(A) and a bit-sliced index on R(B); both are secondary indexes. What would be a good algorithm for processing this query?