

Obligatory Exercise 5

Due date: December 16, 2006 at 0800 (8am)

The following table applies to Problems 1 and 2.

Number of tuples in the relation	4×10^{10}
Size of one tuple	512 bytes
Size of pointer to B-tree node	8 bytes
Page size	8192 bytes

1. A B-tree is used as the data structure for a relation. Assume that each B-tree node is stored in one page, and that each tuple is stored as one record. Answer the following.
 - (a) Compute the maximum number of records (tuples) which may be stored in a single page.
 - (b) Compute the maximum height that such a B-tree may have. (The height is defined as the length of a path from the root to a leaf.)
 - (c) Compute the minimum height that such a B-tree may have.
2. This question investigates record distributions in B-trees, using the values computed in Question 1.
 - (a) For a tree of the maximum height computed in 1(b), compute the average number of records in non-root nodes, assuming that the root contains exactly two (2) records. State whether or not this value is consistent with the constraints on the minimum and maximum number of records in a node of a B-tree. If it is consistent, give also the total number of nodes in the tree.
 - (b) For a tree of the maximum height computed in 1(b), compute the average number of records in a node, assuming a uniform distribution of records in all nodes, including the root. State whether or not this value is consistent with the constraints on the minimum and maximum number of records in a node of a B-tree. If it is consistent, give also the total number of nodes in the tree.
 - (c) For a tree of the maximum height computed in 1(b), assume a uniform distribution of records in all nodes but the root. Compute the maximum number of records which is possible for the root node, under the condition that the average number of records in the other nodes is consistent with the constraints on the minimum and maximum number of records in a node of a B-tree. (This number of records in the root node must in this case be a positive integer.) Give also the total number of nodes in the tree.

5DV021, Obligatory Exercise 5, page 2

[Hint: Using the values obtained in parts (a) and (b), use a binary search to find the desired value.]

- (d) For a tree of the minimum height computed in 1(c), compute the average number of records in non-root nodes, assuming that the root contains the maximum number of records possible. State whether or not this value is consistent with the constraints on the minimum and maximum number of records in a node of a B-tree. If it is not consistent, then compute the minimum number of records in the root node for which a uniform distribution of records in all other nodes is consistent. In either case, give also the total number of nodes in the consistent tree.
- (e) For a tree of the minimum height computed in 1(c), assume a uniform distribution of records in all nodes but the root. Compute the minimum number of records which is possible for the root node, under the condition that the average number of records in the other nodes is consistent with the constraints on the minimum and maximum number of records in a node of a B-tree. (This number of records in the root node must in this case be a positive integer.) Give also the total number of nodes in the tree. [Hint: use a binary search, as in part (c).]

The following table applies to Problems 3 and 4.

Number of tuples in the relation	4×10^{10}
Size of one tuple	512 bytes
Size of pointer to a B ⁺ -tree node	8 bytes
Size of internal B ⁺ -tree key	30 bytes
Page size	8192 bytes
Sequential-access pointers in the leaf nodes	32 bytes total

3. Suppose now that the relation will be stored in a B⁺-tree. Assume that each node, interior or leaf, is stored in one page, and that each tuple is stored as one record. Answer the following.
- (a) Compute the maximum number of indices per non-leaf node.
- (b) Compute the maximum number of records per leaf node.
- (c) Compute the maximum height that such a B⁺-tree may have. (The height is defined as the length of a path from the root to a leaf.)
- (d) Compute the minimum height that such a B⁺-tree may have.
4. This question investigates record distributions in B⁺-trees, using the values computed in Question 3.

5DV021, Obligatory Exercise 5, page 3

- (a) For a B^+ -tree of the minimum height computed in 3(d), compute the average number of indices in the interior nodes other than the root, assuming that the root contains exactly ten (10) indices and that the leaf nodes contain the maximum number of records possible. State whether or not this value is consistent with the constraints on the minimum and maximum number of indices in an interior node of a B^+ -tree. In addition, if this value is consistent, compute the number of interior nodes and the number of leaf nodes, as well as the average number of records in each leaf node.
- (b) For a B^+ -tree of the minimum height computed in 3(d), assume a uniform distribution of indices in all interior nodes except the root, and assume further that each leaf node contains the maximum possible number of records. Compute the minimum number of indices which is possible for the root node, under the condition that the average number of indices in the other interior nodes is consistent with the constraints on the minimum and maximum number indices for such a B^+ -tree. (This number of indices in the root node must in this case be a positive integer.) Give also the total number of interior nodes and the total number of leaf nodes. [Hint: Again use binary search.]
- (c) Repeat part (b), this time finding the *maximum* number of indices which is possible for the root node.
- (d) For a B^+ -tree of the minimum height computed in 3(d), assume a uniform distribution of indices in all nodes, including the root, and assume further that each leaf node contains the maximum number of records possible. Compute the average number of indices per interior node, and state whether or not this value is consistent with the constraints on the minimum and maximum number of indices in an interior node of a B^+ -tree. In addition, if this value is consistent, compute the number of interior nodes and the number of leaf nodes.
- (e) Repeat part (d), this time assuming that each leaf node contains the *minimum* number of records possible.

Notes:

- As stipulated in the course syllabus, this exercise may be done either individually, in a group of two, or in a group of three.
- Remember that there are point penalties for late submission. See the course syllabus.
- It is strongly recommended that you use a text-processing tool to display your results. If you write them by hand, they must be very neat.
- It is not allowed to copy the work of others. The submission must be the original work of the individual(s) in the working group.
- The grader reserves the right to interview members of the working group about the solution.
- So that solutions may be discussed in a class meeting, students and/or groups that are very late in preparing a solution may be required to solve an alternate problem to receive credit for this exercise.

5DV021, Obligatory Exercise 5, page 4

- If you have solved this problem for a previous offering of the course, you may re-use your old solution, subject to the following conditions: (a) You may not work with any partners, except possibly those with whom you worked to prepare the solution in the previous course. (b) You must explicitly note any partners from the previous course with whom you submitted a joint solution for that course. Note that grading criteria may not be identical between years, so that a solution which was found to be satisfactory last year may not be evaluated similarly this year. Note also that the problems themselves may be different from those of previous years. You must in any case solve the problems for this year!