

Database Technology

Exercise 6: Indexing and Hashing

6.1. Indices

- a. Indices speed query processing, but it is usually a bad idea to create indices on every attribute, and every combinations of attributes, that is a potential search keys. Explain why?
- b. Is it possible in general to have two clustering indices on the same relation for different search keys? Explain your answer.

6.2. B+ Trees

- a. Construct a B+ tree for the following set of key values: (2,3,5,7,11,17,19,23,29,31)
Assume that the tree is initially empty and values are added in ascending order.
Construct B+ trees for the cases where the number of pointers that will fit in one node is:
 - i. Four
 - ii. Six
 - iii. Eight
- b. For each B+ tree of the previous task, show the form of the tree after each of the following series of operations:
 1. Insert 9
 2. Insert 10
 3. Insert 8
 4. Delete 23
 5. Delete 19
- c. What would the occupancy of each leaf node of a B+ tree be, if index entries are inserted in sorted order? Explain why.
- d. Create a formula for a b-order B+ tree with height h (levels of index) which computes
 - i. The maximum number of keys
 - ii. The maximum number of records

6.3. Hashing

- a. Suppose that we are using extendable hashing on a file that contains records with the following search-key values: 2,3,5,7,11,17,19,23,29,31. Show the extendable hash structure for this file if the hash function is $h(x)=x \bmod 8$ and buckets can hold three records.
- b. Show how the extendable hash structure of the previous task changes as the result of each of the following steps:
 - i. Delete 11
 - ii. Delete 31
 - iii. Insert 1
 - iv. Insert 15

6.4. Bitmaps

Consider the following relation on the right:

- a. Construct a bitmap index on the attribute salary, dividing values into 4 ranges:
 - below 50000
 - 50000 to 60000
 - 60000 to 70000
 - 70000 and above
- b. Consider a query that requests all instructors in the Finance department with salary of 80000 or more. Outline the steps in answering the query, and show the final and intermediate bitmaps constructed to answer the query.

ID	name	dept_name	salary
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000