

CIS 6930 Midterm, Spring 2008, 02/27/2008

Name: _____

I. (60pts) Given the following database design,

LibraryBranch (lid, bname, address);
BookLoan(bid, lid, rid, dateOut, dueDate);
Reader(rid, name, address, age, phone);
Book (bid, title, publisher).

Write the following queries in SQL:

(a) Retrieve the names of all readers who never checked out any books.

(b) For each book that is loaned out from the branch named "Sharpstown" and whose due date is today, retrieve the book title, the borrower's name, and the borrower's address.

(c) For each library branch, retrieve the branch name and the total number of books loaned out from that branch.

(d) Retrieve the names, addresses, and number of books checked out for all borrowers who have more than five books checked out.

(e) Find the names of readers who checked out any book from Sharpstown branch but none from Central branch;

(f) Find the name(s) of the oldest reader(s) registered in the database.

(g) Find the titles of books that are checked out (i.e., dueDate \geq today) in ALL library branches.

Write the following queries in relational algebra:

(h) For each book that is loaned out from the branch named "Sharpstown" and whose due date is today, retrieve the book title, the borrower's name, and the borrower's address.

(i) Find the names of readers who checked out any book from Sharpstown branch but none from Central branch;

(j) Find the titles of all books that are checked out ($\text{dueDate} \geq \text{today}$) in all library branches.

II. (10pts) True or false. Please clearly mark each statement as either True (T) or False (F). Ambiguous answers will be awarded zero points. You do not need to explain your answers.

(a) In a B+ tree index, if we store the data records directly in the leaf nodes (alternative #1 for storing records), this index is a clustered index. T () F ()

(b) In a clustered B+ tree index, the data records must be stored in the leaf nodes (alternative #1 for storing records). T () F ()

(c) Hash-based indexes are generally useful for both equality search/selection queries and range queries. T () F ()

(d) A hash-based index with a well-designed hash function (i.e., random and uniform) can definitely be a clustered index. T () F ()

(e) In the dynamic hashing strategy introduced in class, two or more entries in the hash directory can point to the same bucket. T () F ()


III. Indexing (30pts)

1. (20pts) Construct a B⁺-tree for the following set of key values:

(38, 27, 52, 48, 26, 64, 42, 20, 57, 18, 69)

Assume the tree is initially empty and the values are added in the order specified by the above list. The number of pointers that will fit in one node is FOUR (i.e, THREE key values). Draw intermediate trees for partial credits.

2. Given the following B⁺ tree



Please finish the following tasks:

2a. (5pts) Draw the tree after deleting key value 9 from the above tree;

2b. (5pts) Draw the tree after deleting key value 31 from the tree resulted in 2a.