

CS 464/564 Spring 2021 Final Exam

Name: _____

Due (upload to Canvas) 11:59 PM Thursday May 13

1. Write relational algebra expressions to answer the following questions. Use the library database we used in class.

a. Find the sid, first name and last name of staff who work at a library with a capacity less than 100.

b. Find the names of customers who have borrowed every book (not necessarily every copy) written by Isaiah Berlin.

c. Find the names of customers who have never borrowed a book written by Alexander Pushkin.

2. Write a Postgres trigger that is fired when an insert happens on the Copy table. The trigger inserts a row into the CopyHistory table. The CopyHistory table consist of the copynum (primary key), booknum (foreign key referencing Book) and date_inserted. The types for copynum, booknum and date_inserted are int, int and Date respectively. A row in the table records the date a copy was inserted into the Copy. The date_inserted is the current_date on the day the insert is done.

3. Write a Postgres stored procedure that is passed a staff id, (an int) and returns the number of employees managed directly or indirectly by the employee whose id was passed to the procedure.

4. a. What does an extendible hash table with a bucket size of 2 look like after the following values are inserted. Redraw the table each time the directory size doubles. Assume the values are inserted in the order shown below and the initial directory size is 2.

Key	Hash Values
12	1010
22	1001
43	1000
4	0111
58	0110
106	0101
77	0100
8	0011
19	0010
10	0001

4b. Show what your answer to problem 4a looks like after the following values are removed. Assume they are removed in the order shown. Redraw the table each time the directory is made smaller.

Key
1
77
106
4
22
58

5. Suppose the Loan table contains 100,000 rows and there are 50 rows per page. Also suppose there are 100 main memory page buffers that can be used for sorting.

a. How many secondary memory accesses (reads and writes) will be done in order to sort the Loan table? Use the external sorting algorithm (i.e. partial sort then k-way merging) described in the lecture on external sorting.

b. How many merge phases are required to complete the sort?

6. Use the following FD set to answer parts a, b, c and d.

FDs = { $AB \rightarrow CD$, $A \rightarrow F$, $CD \rightarrow AG$, $C \rightarrow AF$, $F \rightarrow D$ }

a. Find a key for the universal table.

b. Find a minimal cover set for the FD set.

c. Create a set of 3NF tables from the minimal cover set.

d. If any of the tables in your answer to part c are not in BCNF decompose them into BCNF tables.

7. What are the ACID properties of transactions in relational databases? Define the meaning of each of the four terms that are part of the ACID properties.

8. What does the acronym CAP mean in the context of NoSQL databases?

9. What is the CAP "Theorem"?