M. Nascimento

NAME: _____ STUDENT ID: _____

CMPUT 291-B1 Final Exam - April 17, 2000

General Guidelines: *Exam duration: 120 min. Closed book, no collaboration. No questions during the exam, if you are unsure, state your assumptions clearly. All questions need to be answered. Write your name and student id number in all pages. Use the spaces after the questions for your answers (additional pages can be used, but should not be necessary). Marked exams, for this section, will be available at 2pm on April 21 at the instructor's office, GSB 733. Absolute final deadline date for appeals is April 21 at 5pm. It was nice being your instructor. Best of luck in your future (and in this final exam !).*

Question 1 (15 points). Map the following ER diagram into relations. Make sure all attributes attributes and primary keys are defined. For each resulting relation indicate, if any, the foreign keys and their original relations.

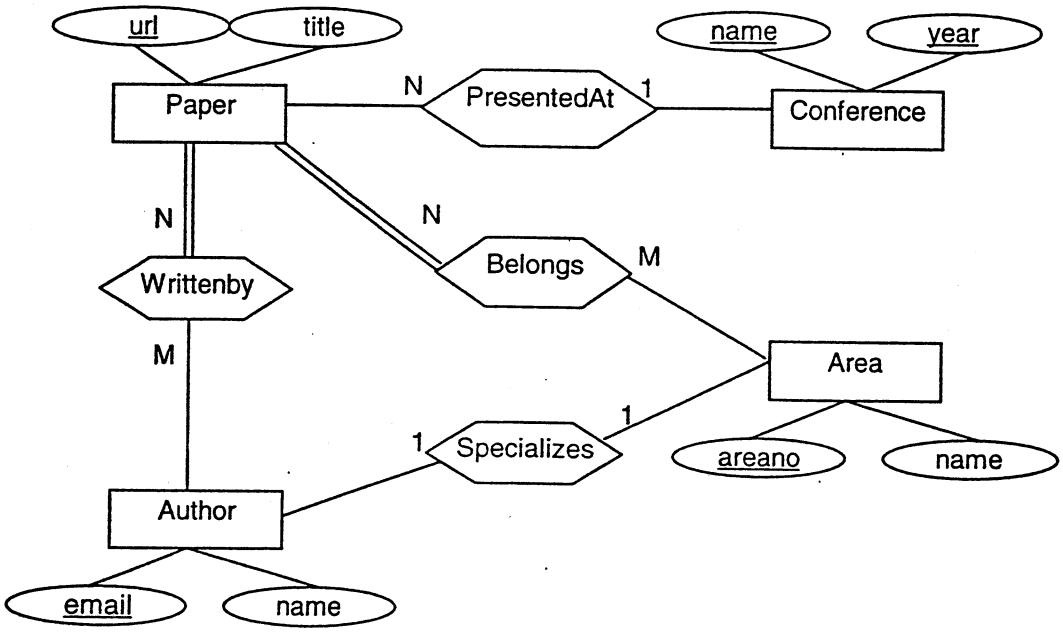

Question 2. Consider the following relations regarding a library setting. Assume that attributes with same name across relations represent a primary key-foreign key relationship.

Books (bookId, bookTitle, publisherId)
Publishers (publisherId, cityName)
Authors (bookId, authorName)
Users (userId, userName, userPhone)
Loans (bookId, userId, dateOut, dueDate, dateIn)

Write the following queries using SQL:

A (5 points). What is (are) the names(s) of the user(s) who has (have) borrowed at most 5 books during the library's lifetime ? (It doesn't matter whether the books have been already returned, are currently borrowed or are late).

B (5 points). What is the average number of books published by all publishers ?

C (5 points). What are the userIds of the users who, during the library's lifetime, have borrowed at least once all books ?

Question 3 (5 points). Using the same relations given in Question 2, write the following query using Relational Algebra: What are the names of the users who have at least once, returned a book exactly one day after borrowing it ?

Question 4 (5 points). Using the relations given in Question 2, give an example of a query which you can express in SQL but you cannot express using Relational Algebra. Explain why is that.

Question 4 (10 points). Using Relational Algebra show that the natural join operator can be obtained using only two of the following operators: select, division, project, cross-product and set-intersection.

Question 5 (10 points). When reusing freed disk space (blocks) in a file with variable length records one may use the best-fit, worst-fit and first-fit policies. Which one should be a best compromise between external fragmentation and speed necessary to find a suitable block ? Explain why.

Question 6. Consider the following following B-tree (of order 4). The lower-case letter above each node denotes its address (or id).

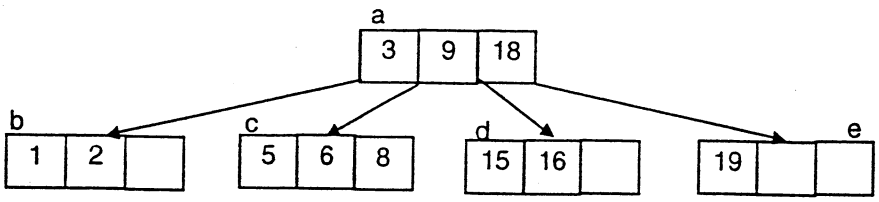

A (5 points). Insert key value 4. Indicate the steps performed during the insertion and show the resulting tree. Assume the use of left promotion.

B (5 points). Considering the original tree (i.e., before the insertion in item A above), delete key value 1. Indicate the steps performed during the deletion and show the resulting tree.

C (5 points). Considering again the original tree (i.e., before the insertion and deletion in items A and B above), assume you have a three page buffer pool to improve I/O performance of the B-tree. Assuming a First-In-First-Out policy for reclaiming space in the buffers, what is the number of I/Os necessary to obtain the data pointer (not the data record) associated with each of the following key values (in this order): 5, 3 and 1. Explain your answer. Assume no other accesses are made to the B-tree at this time.

Question 7 (10 points). Given the following relation R and the set of functional dependencies FD1, FD2 and FD3, decompose R so that all resulting relations are in 3rd normal form. Indicate clearly what are the final relations. (Hint: using the FDs to determine the candidate keys of R and assume the shortest one to be the primary key).

R (p, c, l, a, r, t)
FD1: p → c, l, a, r, t
FD2: c, l → p, a, r, t
FD3: a → r
FD4: c → t

NAME: _____ STUDENT ID: _____

Question 8. Consider a disk with the following characteristics:

Sector size in bytes: 512
Number of tracks per surface: 2000
Number of sectors per track: 50
Number of double-sided platters: 5
Average seek time in msecs: 10
Disk speed in RPMs: 6000

A (5 points). What is the capacity, in bytes, of the whole disk ?

C (10 points). If one track of data can be transferred per disk revolution, what is the data transfer rate of a track in bytes/sec ? Is this the actual transfer rate of this disk ? Justify your answer.