

CMPSCI 645 — Midterm

April 3, 2006

100 Points

90 Minutes

Problem 1 (7 parts, 30 points total) Queries, query languages, and analysis.

Consider the following relational schema. An employee can work in more than one department.

- Emp(eid, ename, age, salary)
- Works(eid, did)
- Dept(did, dname, budget, managerid)

(a) (6 points) Write an SQL to print the names and ages of each employee who works in both the Hardware department and the Software department.

(b) (1 point) Is the query above expressible as a Conjunctive Query? (Yes or No)

(c) (6 points) Write an SQL query to find the managerids of managers who manage **only** departments with budgets greater than 1 million dollars.

(d) (1 point) Is the query above expressible as a Conjunctive Query? (Yes or No)

(e) (3 points) Consider the following two SQL queries:

```
Q1: SELECT DISTINCT E.eid
     FROM Emp E, Works W, Dept D
     WHERE E.eid = W.eid AND W.did = D.did AND D.budget = 100000
```

```
Q2: SELECT DISTINCT E.eid
     FROM Emp E, Works W, Dept D1, Dept D2
     WHERE E.eid = W.eid AND W.did = D1.did AND W.did = D2.did
           D1.budget > 50000 AND D2.budget < 200000
```

State which of the following holds, and **briefly explain** your answer:

- i. $Q_1 \subset Q_2$
- ii. $Q_2 \subset Q_1$
- iii. Q_1, Q_2 are equivalent
- iv. None of the above.

(f) (6 points) Associate each query language class listed below with an equivalent class from: the conjunctive queries (**CQ**), the first order queries (**FO**), or the recursive queries (**Recursive**).

- i. Unions of conjunctive queries with negation
- ii. The relational algebra restricted to the operators \times, σ, Π
- iii. Datalog queries (no negation)
- iv. Queries expressible as single SELECT DISTINCT - FROM - WHERE blocks.
- v. Non-recursive datalog queries (no negation)
- vi. The relational algebra

(g) (7 points) A Boolean tree has leaves labeled with 0 or 1, and internal nodes which are operators AND, OR, **NOT**. The root of a boolean tree returns 0 or 1. We can represent a Boolean tree using the following stored (EDB) predicates:

- Leaf0(x): a unary table containing the identifiers of each leaf labeled 0. (Alternatively, you can think of Leaf0(x) as a logical predicate that is true of node x if x is a leaf labeled 0.)
- Leaf1(x): a unary table containing the identifiers of each leaf labeled 1.
- And(x, y_1, y_2): a ternary table containing tuples (x, y_1, y_2) if x is an AND node whose children are nodes y_1 and y_2 .
- Or(x, y_1, y_2): a ternary table containing tuples (x, y_1, y_2) if x is an OR node whose children are nodes y_1 and y_2 .
- **Not(x, y): a binary table containing tuples (x, y) if x is a NOT node whose child is node y .**
- Root(x): a unary table containing the root node.

Write a datalog query that computes the boolean relation Answer() which is true if and only if the root node returns 1. Hint: your program should compute IDB relations One(x) and Zero(x), in addition to Answer().

Problem 2 (5 parts, 20 points total) Normalization and redundancy

(a) (4 points) Consider the following table:

A	B	C
a11	b22	c55
a11	b22	c66
a11	b33	c77
a22	b44	c66

Indicate which of the functional dependencies hold for this instance. Answer Yes or No for each part:

- $B \rightarrow A$
- $B \rightarrow C$
- $A, B \rightarrow C$
- $B, C \rightarrow A$

(b) (6 points) For $R(A,B,C,D)$ and functional dependencies $AB \rightarrow C$, $AB \rightarrow D$, $C \rightarrow A$, $D \rightarrow B$ state the strongest normal form the relation is in. Briefly explain your answer.

(c) (2 points) Is redundancy (with respect to functional dependencies) possible for this relation? Answer Yes or No.

(d) (6 points) For $R(A,B,C,D,E)$ and functional dependencies $A \rightarrow BC$, $BC \rightarrow E$, and $E \rightarrow DA$, state the strongest normal form the relation is in. Briefly explain your answer.

(e) (2 points) Is redundancy (with respect to functional dependencies) possible for this relation? Answer Yes or No.

Problem 3 (10 points total) State if the following statements are TRUE or FALSE.

- (a) Suppose that a randomly chosen block is being fetched from disk. The overall access delay per byte decreases as the block size increases.
- (b) The LRU (Least Recently Used) buffer replacement strategy is a good strategy for repeated sequential scans of the same file.
- (c) Insertion order into a B+ Tree will effect the tree's end structure.
- (d) One way to handle duplicate keys in an unclustered index is to maintain a pointer to only the first occurrence of the key in the file.
- (e) B+ trees are typically more efficient than linear or extensible hashing for all types of queries.
- (f) A B+Tree on $\langle \text{age, salary, department} \rangle$ can be used to answer a selection condition "age=20 AND department = 'CS'".
- (g) A hash index on $\langle \text{age, salary} \rangle$ can be used to answer a selection condition "age=20 AND salary > 5000".
- (h) Since there is no downside to having indexes, we may as well build an index on every column of every relation to speed up as many queries as possible.
- (i) $\sigma_p(E_1 - E_2) = \sigma_p(E_1) - E_2$
- (j) $\pi_s(E_1 - E_2) = \pi_s(E_1) - \pi_s(E_2)$

Problem 4 (4 parts, 25 points total): Indexes and File Organization

- (a) (6 points) **Circle** the basic file organization (heap, sorted, or hash) that is best for a large file where the most frequent operations are as follows (answer each separately – no explanation needed):
- i. Search for records based on a range of field values.
HEAP SORTED HASH
 - ii. Perform inserts and scans where the order of records does not matter.
HEAP SORTED HASH
 - iii. Search for a record based on a particular field value.
HEAP SORTED HASH
- (b) (8 points) Create a B+tree where each node can hold at most 4 pointers and 3 keys when the following keys are inserted in the following order:
1, 10, 2, 11, 3, 4, 8, 5, 7, 6

- (c) (6 points) Assume that a B+tree can hold at most n keys in each node and the height of the tree is h . Also assume there is a 50% minimum occupancy requirement for each node, except the root. What are the **minimum** and **maximum** number of keys can this B+tree hold?
- (d) (5 points) Suppose that you have a file that is already sorted in key order and you want to construct a clustered B+ tree index on this file using (key, RID) pairs for data entries. A simple way to accomplish this is to create an empty B+tree, and then sequentially scan the file, inserting an index entry for each record using the normal B+tree insertion routine. What **storage utilization** problem does this approach have? **Briefly** describe a solution that would solve the problem.

Problem 5 (3 parts, 15 points total) Query Optimization

Consider the following schema:

- Sailors(sid, sname, rating, age)
- Boats(bid, bname, size)
- Reserves(sid, bid, day)

Reserves.sid is a foreign key to Sailors and Reserves.bid is a foreign key to Boats.bid.

We are given the following information about the database:

- Sailors contains 1000 records with 20 records per page.
- Boats contains 100 records with 10 records per page.
- Reserves contains 10,000 records with 40 records per page.
- There are 100 values for Reserves.sid.
- There are 50 values for Reserves.bid.
- There are 1000 values for Reserves.day

In the following queries, assume that a System R style optimizer is used.

Consider Query 1:

```
SELECT S.sid, S.sname, B.bname
FROM   Sailors S, Reserves R, Boats B
WHERE  S.sid = R.sid AND R.bid = B.bid AND R.day = 'July 4, 2003';
```

- a) (6 points) Assuming uniform distribution of values and column independence, estimate the number of tuples returned by this query.

Consider Query 2:

```
SELECT S.sid, S.sname, B.bname
FROM Sailors S, Reserves R, Boats B
WHERE S.sid = R.sid AND R.bid = B.bid
```

b) (4 points) Draw all possible left-deep query plans for this query:

c) (5 points) Assume that the first join in a query plan (the one at the bottom of the tree) is $R \bowtie S$. Assume that you have 51 pages of memory. There are no indexes, so indexed nested loops join is not an option. What join algorithm would work best? Briefly explain your answer.