

CSE-4411A Test #2

Sur / Last Name:
Given / First Name:
Student ID:

- **Instructor:** Parke Godfrey
- **Exam Duration:** 80 minutes
- **Term:** Fall 2009

Answer the following questions to the best of your knowledge. Your answers may be brief, but be precise and be careful. The exam is open-book and open-notes. Calculators, etc., are fine to use. Write any assumptions you need to make along with your answers, whenever necessary.

There are four major questions. Points for each question and sub-question are as indicated. In total, the test is out of 50 points.

If you need additional space for an answer, just indicate clearly where you are continuing.

Grading Box	
1.	/15
2.	/10
3.	/15
4.	/10
Total	/50

-
-
1. (15 points) **Indexes and Access Paths.** *Forest for the trees.* [SHORT ANSWER]

Table **R** has a *clustered* tree index of type alternative #2 on A, B, C, D (in that order, all ascending). The index pages contain 99 *index records* (encompassing 99 pointers, plus a 100th pointer), on average;¹ data-entry pages contain 50 data entries each, on average; and data-record pages contain 20 data records each, on average.

A's values range over 1..10,000; B's over 1..1,000; C's over 1..100; and D's over 1..10.

- a. (3 points) An index record contains effectively the information A, B, C, D, and a pointer (an address to another page). A data entry contains effectively A, B, C, D, and an RID (which is an address to another page *and* a slot number). A data entry is *slightly* larger than an index record—by a slot number—but only slightly. So explain how it is possible there are 99 index records per page, on average, but *only* 50 data entries per page, on average.

-
- b. (3 points) Say that table **R** has 1,000,000 records. How deep is the index tree?

¹This accounts for the fill factor. A page could hold more than 99 index records.

Consider the query

```
select * from R where A = 3141 and B > 500 and D > 2;
```

- c. (3 points) What is the cardinality estimation of the query?
What assumption, often faulty, are you having to make to estimate the cardinality?

-
- d. (3 points) Estimate the I/O cost of the query using the index as the access path.

-
- e. (3 points) Name a *functional dependency* that, if true for the table **R**, would guarantee that use of the index would satisfy *order by* A, B, D.

2. (10 points) **Query Optimization.** *One is not like the others.* [MULTIPLE CHOICE]

Choose *one* best answer for each of the following. Each is worth one point. There is no negative penalty for a wrong answer.

- a. An iterator, pull-based architecture for the query processor
- A. eliminates the need for a buffer pool manager.
 - B. eliminates the need for a query optimizer.
 - C. is good for object-oriented database systems, but not for relational systems.
 - D. implements pipelining.
 - E. does not allow for on-the-fly operations.
-
- b. A System R style optimizer—the type of optimizer the textbook presents—preserves sub-plans that generate *interesting orders* while enumerating possible query plans because
- A. such plans are *always* less expensive than plans that have no interesting order.
 - B. pipelining is impossible without them.
 - C. such plans may allow certain subsequent operations to be done on-the-fly.
 - D. hash joins are impossible without them.
 - E. this prunes the search space of possible query plans.
-
- c. All are access path options *except*
- A. block-nested-loop join
 - B. file scan
 - C. index-only scan
 - D. index intersection
 - E. index probe
-
- d. Consider the two-pass hash join and the two-pass sort-merge join. Which of the following is not true?
- A. The hash join can be used whenever the sort-merge join can be used.
 - B. The sort-merge join can be used whenever the hash join can be used.
 - C. Both are blocking operations that break pipelining.
 - D. Sort-merge join is often preferred as it is immune to skew.
 - E. Sort-merge join is often preferred as it produces an interesting order.
-
- e. Which of the following is true?
- A. Anywhere a (2-pass) hash join can be used, a (2-pass) sort-merge join could be used instead, at the same cost.
 - B. A block-nested-loops join is never better than a hash join or a sort-merge join. It is part of the repertoire simply for the cases when hash join and sort-merge join cannot be applied.
 - C. Select conditions can only be evaluated by use of an index.
 - D. Index-nested-loops join will rarely be the best choice unless there are other highly selective—returning few results—conditions in the query.
 - E. Hash joins and sort-merge joins cannot be used together in the same query plan.




Figure 1: Join Trees.

f. Which of the trees in Figure 1 is a left linear (left-deep) join tree?

- A. I.
- B. II.
- C. III.
- D. IV.
- E. V.
- F. They *all* are left linear join trees.

g. Which of the trees in Figure 1 is not a linear join tree?

- A. I.
- B. II.
- C. III.
- D. IV.
- E. V.
- F. They *all* are linear join trees.

The following information is available on tables **Sailors** and **Reserves**.

- **Reserves:** 10,000 records
 - bid: 50 values (1..50)
- **Boats:** 50 records
 - colour: 5 values
- **Sailors:** 1000 records
 - level: 10 values (1..10)
 - fav_colour: 5 values

The primary key of **Sailors** is sid, of **Reserves** is sid + bid + day, and of **Boats** is bid. Table **Reserves** reserves has a foreign key on sid referencing **Sailors** (on sid), and a foreign key on bid referencing **Boats** (on bid). All columns are *not null*.

For each of the following queries, estimate the selectivity of the query as the number of tuples it likely returns, as System R would. (The estimation techniques discussed in Chapter 12 of the textbook are from System R.)

h. select S.sid, B.bid
 from Sailor S, Boats B
 where B.colour = S.fav_colour;

- A. 1
 - B. 50
 - C. 100
 - D. 5,000
 - E. 10,000
 - F. 500,000
-

i. select S.sid, R.day
 from Sailor S, Reserves R
 where S.sid = R.sid and
 R.bid = 37;

- A. 1
 - B. 20
 - C. 50
 - D. 100
 - E. 200
 - F. 500
-

j. select S.sid, S.sname, R.bid, B.bname, R.day
 from Sailor S, Reserves R, Boats B
 where S.sid = R.sid and R.bid = B.bid and
 R.bid = 37 and S.level = 7;

- A. 1
- B. 5
- C. 20
- D. 80
- E. 100
- F. 10,000

3. (15 points) **Operator Algorithms.** *Try a left-handed integral join...* [EXERCISE]

For the following, the joins considered are equi-joins; that is, the join predicates are equalities.

- a. (3 points) Consider that table **R** has a foreign key referencing table **T**, on **T**'s primary key **A**. The foreign key is mandatory. **T** has a clustered tree index on **A**. Consider a query with an equi-join on **A** between **R** and **T**. The query has no other predicates, selections or otherwise.

In what scenario—sizes of tables, whether the join columns are *distinct* for the outer or the inner, and so forth—would an index nested loop join be significantly less expensive than a two-pass sort-merge join or hash join?

Explain why.

-
- b. (3 points) Name a scenario—sizes of tables, etc.—when having the outer sorted *or* partitioned with respect to the join columns would benefit an index-nested-loop join.

Explain why.

Schema:

Student(sid, sname, startdate, major, advisor)
 FK (advisor) refs **Prof** (pid)
Class(cid, dept, number, section, term, year, room, time, pid, ta)
 FK (pid) refs **Prof**
 FK (ta) refs **Student** (sid)
Enrol(sid, cid, date, grade)
 FK (sid) refs **Student**
 FK (cid) refs **Class**
Prof(pid, pname, pdept, office)

Assume no attribute is nullable. The attribute **pid** in **Class** refers to the the professor / instructor for the class. The attribute **ta** in **Class** refers to the teaching assistant for the class. The attribute **advisor** in **Student** refers to the student's academic advisor.

Statistics:

- **Student:** 50,000 records on 1,000 pages
 - advisor: 2,500 distinct values
- **Enrol:** 2,000,000 records on 20,000 pages
 - sid: 50,000 distinct values
 - cid: 80,000 distinct values
- **Class:** 80,000 records on 1,600 pages
 - pid: 4,000 distinct values
 - ta: 5,000 distinct values
- **Prof:** 4,000 records on 40 pages

Indexes:

- **Student:**
 - clustered tree index on sid (200 data entries per page)
- **Enrol:**
 - clustered tree index on cid, sid (167 data entries per page)
 - unclustered tree index on sid, cid (167 data entries per page)
- **Class:**
 - clustered tree index on cid (200 data entries per page)
- **Prof:**
 - clustered tree index on pid (200 data entries per page)

All indexes are of alternative #2. For each tree index, the index pages are 3 deep, except for the index on **Prof.pid** which is 2 deep.

Query:

```
select sid, sname, cid, dept, number, section, term, year
  from Student S, Class C
 where S.sid = C.ta;
```

c. (3 points) What is the cardinality of the query?

d. (6 points) Devise a least-cost query plan for the query. Assume you have a buffer-pool allocation of 50 frames. Show your plan's cost.

4. (10 points) **Join Mechanics.** *Maybe it's the carburetor?*

[ANALYSIS]

- a. (3 points) Dr. Datta Bas has an idea for how to share partitions with two-pass hash joins. Say we have a query joining tables **R**, **S**, and **T** with $\mathbf{R.A} = \mathbf{S.A}$ and $\mathbf{S.A} = \mathbf{T.A}$. He claims we can reuse partitions of one hash join in the next hash join. Is he right? Will this save I/O's? If so, explain. If not, explain why he is wrong.

-
- b. (2 points) Say you have a fixed buffer pool allocation of B frames for the query in Question 4a, and will use Dr. Bas's shared-partition idea with two two-pass hash joins. What would have to be true so that it could be used?

c. (3 points) We have a query that joins tables **R**, **S**, and **T**, with join predicates $\mathbf{R.A} = \mathbf{S.A}$, $\mathbf{R.B} = \mathbf{S.B}$, and $\mathbf{S.A} = \mathbf{T.A}$.

- **R** is 50,000 tuples on 1,000 pages.
- **S** is 1,000,000 tuples on 20,000 pages.
- **T** is 200,000 tuples on 2,500 pages.
- $\mathbf{R} \bowtie_{\mathbf{A,B}} \mathbf{S}$ is 5,000,000 tuples on 143,000 pages.
- $\mathbf{R} \bowtie_{\mathbf{A}} \mathbf{T}$ is 300,000 tuples on 10,000 pages.
- $\mathbf{S} \bowtie_{\mathbf{A}} \mathbf{T}$ is 10,000,000 tuples on 335,000 pages.
- The result of $(\mathbf{R} \bowtie_{\mathbf{A,B}} \mathbf{S}) \bowtie_{\mathbf{A}} \mathbf{T}$ is 10,000 tuples on 500 pages.

Using two-pass sort-merge joins, show an efficient plan to evaluate the query. There are no indexes. You have a buffer pool allocation of 150 frames.

What is its I/O cost, in total?

d. (2 points) Dr. Dogfurry thinks that an even better plan can be found for the query in Question 4c using two-pass sort-merge joins if *over sorting* is allowed; that is, the first sort-merge join sorts the stream *more* than it needs to for the benefit of future operations.

Is he right? If so, what efficient plan can you find using over-sorting, and what is its I/O cost, in total?

(Scratch space.)