CSE-4411(A) Test #2

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

- **Instructor:** Parke Godfrey
- **Exam Duration:** 75 minutes
- **Term:** Fall 2005

The exam is open-book and open-notes. Answer the following questions to the best of your knowledge. Your answers may be brief, but be precise and be careful. Make clear any assumptions that 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 exam is out of 50 points.

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

Regrade Policy

- You should not write anything on the exam paper after the exam period, if it is to be considered for regrading. Write any subsequent notes on separate paper.
- Regrading should only be requested in writing. Write what you would like to be reconsidered. Note, however, that an exam accepted for regrading will be reviewed and regraded in entirety (all questions).

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

1. (10 points) **Join Algorithms.** *A sinful join.* [analysis]

   Your new boss, Dr. Datta Bas, has invented a new join algorithm which is a hybrid of general sort-merge join (MJ)—so not the two-pass version—and index-nested-loop (INL) join. He has nick-named it SIN (sort-index-nested) join. For SIN join, the outer "table" (stream) should be sorted first on the join key. Then the inner table is probed via an index as is done in the INL join.

---

   a. (5 points) Under what circumstances—sizes of the outer and inner, the number of distinct values over the join column(s), the type of probe index, etc.—if any, would Dr. Bas's SIN join be more efficient than regular INL join?

      Either explain *why* SIN join is *not* more efficient than INL join under any circumstances, or explain *how* it is more efficient under certain circumstances.

b. (5 points) Under what circumstances, if any, would Dr. Bas's SIN join be more efficient than general sort-merge join (MJ)?

Either explain *why* SIN join is *not* more efficient than MJ under any circumstances, or explain *how* it is more efficient under certain circumstances.

2. (15 points) **General Optimization.** *Plan your plan.* [multiple choice]

For each of the following, choose *one* best answer.

a. *Pipelining* in query plans
   A. reduces CPU overhead.
   B. eliminates the need to write temporary results to disk.
   C. is a type of access path.
   D. can never reduce the cost of executing the query, so is avoided whenever possible.
   E. is typically used with the inner stream of a join.

b. Which of the following is *true* about the join algorithms?
   A. The larger of the two tables should always be designated the outer table.
   B. The smaller of the two tables should always be designated the outer table.
   C. In any situation (2-pass) hash join (HJ) could be used, it would also be *possible* to use 2-pass sort-merge join (SMJ).
   D. In any situation (2-pass) SMJ could be used, it would also be *possible* to use HJ.
   E. If an appropriate index is available for an index-nested loops join (INL), then INL is guaranteed to be less expensive than SMJ or HJ.

c. The index-nested-loops join
   A. can only use a *clustered* index to probe the inner table.
   B. cannot be pipelined.
   C. probes the inner table for *each* record of the outer table / stream.
   D. could be used for any join in a *bushy-tree* plan.
   E. is generally more efficient when the outer table / stream is large and the inner table is small than vice-versa.

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

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

---

d.
```
select S.sid, R.day
    from Sailor S, Reserves R
    where S.sid = R.sid and
        R.bid = 7;
```

A. 2
B. 5
C. 20
D. 50
E. 200
F. 500

---

e.
```
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 = 47 and S.level <= 7 and S.level > 3;
```

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

---

f.
```
select S.sid, S.sname, B.bid, B.bname
    from Sailor S, Boats B
    where S.fav_colour = B.colour;
```

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

[1] The estimation techniques discussed in Chapter 12 of the textbook are from System R.

g.
```
select part_no from Product
    where (price > 1000 and weight < 20)
        or (price > 5000 and quantity < 200);
```

There is an index on weight and on quantity. There is not an index on price.

A. Evaluation of the query above requires a table scan, so the indexes do not help.
B. An access path that employs just the index on weight is possible.
C. An access path that employs just the index on quantity is possible.
D. No access path with just one index is possible, but an access path that employs both the indexes on weight and quantity is possible.
E. There is not enough information provided to answer this question.

---

h. Restricting focus to left-linear join trees is beneficial for all *except* which of the following reasons?

A. The inner "relation" for every join is a base table—or the evaluation of a single relation sub-query—so a more accurate size estimation of the output can be obtained.
B. The inner "relation" for every join is a base table—or the evaluation of a single relation sub-query—so an index-nested-loops join remains possible.
C. This helps prune the search search space of all possible join trees dramatically.
D. For any query plan based on a join tree that is not left linear, there is guaranteed to be a query plan based on left-linear join tree that is less expensive.
E. Left-linear join trees typically enable pipelining along the outer "relations".

---

i. *Double buffering* is useful in the external sort algorithm because it

A. increases fan-in during merges so that fewer passes may be necessary.
B. results in initial runs that are twice as long (than if double buffering were not used).
C. likely ensures that merge processing does not need to wait while pages are being read and written.
D. allows for sequential reads and writes.
E. speeds up I/O operations.

Consider table **R** with attributes A and B, table **S** with attributes B and C, and table **T** with attributes C and D. All joins below are natural joins; that is, the join columns are those columns named the same between the two tables.

j. Consider the following relational-algebra expressions:

I. $(\mathbf{R} \bowtie \mathbf{S}) \bowtie \mathbf{T}$
II. $(\mathbf{R} \bowtie \mathbf{T}) \bowtie \mathbf{S}$
III. $\mathbf{R} \bowtie (\mathbf{S} \bowtie \mathbf{T})$

Which of the above relational algebra expressions necessarily evaluate to the same result?

A. All three must evaluate to the same result.
B. They each may evaluate to a different result.
C. **I** & **II**
D. **I** & **III**
E. **II** & **III**
F. Not enough information is provided to determine this.

k. Consider the following relational-algebra expressions:

I. $\pi_{\mathsf{A,D}}((\mathbf{R} \bowtie \mathbf{S}) \bowtie (\mathbf{S} \bowtie \mathbf{T}))$
II. $\pi_{\mathsf{A,D}}(\pi_{\mathsf{A,B,D}}(\mathbf{R} \bowtie \mathbf{S}) \bowtie \pi_{\mathsf{A,B,D}}(\mathbf{S} \bowtie \mathbf{T}))$
III. $\pi_{\mathsf{A,D}}(\pi_{\mathsf{A,C,D}}(\mathbf{R} \bowtie \mathbf{S}) \bowtie \pi_{\mathsf{A,C,D}}(\mathbf{S} \bowtie \mathbf{T}))$

Which of the above relational algebra expressions necessarily evaluate to the same result?

A. All three must evaluate to the same result.
B. They each may evaluate to a different result.
C. **I** & **II**
D. **I** & **III**
E. **II** & **III**
F. Not enough information is provided to determine this.

l. Consider the following relational-algebra expressions:

I. $\pi_{\mathsf{B}}(\mathbf{R} \bowtie (\mathbf{S} \bowtie \mathbf{T}))$
II. $\pi_{\mathsf{B}}((\pi_{\mathsf{B}}(\mathbf{R}) \bowtie \mathbf{S}) \bowtie \mathbf{T})$
III. $\pi_{\mathsf{B}}((\pi_{\mathsf{B}}(\mathbf{R} \bowtie \mathbf{S})) \bowtie \mathbf{T})$

Which of the above relational algebra expressions necessarily evaluate to the same result?

A. All three must evaluate to the same result.
B. They each may evaluate to a different result.
C. **I** & **II**
D. **I** & **III**
E. **II** & **III**
F. Not enough information is provided to determine this.

The following questions are about System R, and how it works.

Consider a query that joins tables **R**, **S**, **T**, **U**, **V**, **X**, and **Y**. Assume any pair of tables (e.g., **S** & **V**) have a join condition between them.

---

m. How many left-join trees are possible?
   A. 1
   B. 7
   C. 49
   D. 128
   E. 5,040

---

n. In phase 4 of join enumeration, (when query plans of four tables are generated), *at least* how many four-table plans are retained?
   A. 1
   B. 16
   C. 24
   D. 35
   E. 2,401

---

o. In each phase of join enumeration, additional plans may be retained than just the plans with the best estimated cost. These additional plans are ones that
   A. are estimated to return *fewer* output records than the best plan is estimated to return.
   B. generate some order on the output stream.
   C. use hash join, which is nice.
   D. each make it possible to eliminate one, or more, of the remaining joins.
   E. are randomly chosen.

(Scratch space.)

3. (15 points) **Joins.** *Join up free, today only!* [exercise / short answer]

Consider the query

```
select *
    from R, S
    where R.A = S.A;
      and S.D is between x and y
      and R.E = z;
```

The schema are $\mathbf{R}(\underline{\mathsf{A}}, \mathsf{B}, \mathsf{E})$ and $\mathbf{S}(\underline{\mathsf{C}}, \mathsf{A}, \mathsf{D})$. The underlined attributes designate the primary key. **S** has a foreign key cast on **R** through A, and **S**.A is not nullable.

The σ_{D} has a reduction factor of $\frac{1}{10}$. Likewise, the σ_{E} has a reduction factor of $\frac{1}{10}$. There are no indexes on **R**.E or **S**.D, however.

There is a clustered tree index on **R**.A and an unclustered tree index on **S**.A. Each index is of alternative #2 and has two layers of index pages. So the third layer in each case consists of the data-entry pages.

Let $N_{\mathbf{T}}$ generically denote the number of records in table **T**. Let $P_{\mathbf{T}}$ denote the number of pages of **T**. And let $V_{\mathbf{T}.\mathsf{C}}$ denote the number of distinct values found in column C of table **T**.

$N_{\mathbf{R}} \ll N_{\mathbf{S}}$ and $P_{\mathbf{R}} \ll P_{\mathbf{S}}$. That is, the number of records in table **R** is *much* less than the number of records in table **S**; and likewise, the number of pages of table **R** is *much* less than the number of pages of table **S**.

For the following questions, push selections as far down in the tree as possible.

---

a. (4 points) Show the annotated query plans (trees) both for using an index nested loops join with **R** as the outer and for using an index nested loops join with **S** as the outer. Be complete. Do not leave any operators out.

b. (6 points) Which would be the better plan? The index nested loops (INL) join with **R** as the outer and using the unclustered index on **S**.A to probe, or the INL join with **S** as the outer and using the clustered index on **R**.A to probe?

Justify your claim. Use N_{R}, N_{S}, P_{R}, P_{R}, $V_{\mathsf{R.A}}$, etc. in your argument.

c. (5 points) Is the better of the two INL plans likely to be better than a (two-pass) sort-merge join plan or a hash-join plan for this query? Explain.

4. (10 points) **Query Plans.** *Okay, here's our plan...* [exercise]

The following information is available on tables **Purchase** and **Books**.

- **Purchase**(cust#, book#, when, qnty)
  - 20,000,000 records on 200,000 pages
  - foreign keys:
    - on (book#) referencing **Books**
  - column cardinalities:
    - book#: 10,000 values
  - indexes:
    - unique clustered tree index on cust# + book# + when#
      depth of 3 index pages
      120 data entries per page
    - unclustered tree index on cust# + when + qnty
      depth of 3 index pages
      120 data entries per page
    - unclustered tree index on book# + when + qnty
      depth of 3 index pages
      120 data entries per page
- **Books**(book#, title, author, year, publisher, price)
  - 10,000 records on 200 pages
  - column cardinalities:
    - author: 1,000 values
  - indexes:
    - unique clustered tree index on book#
      depth of 3 index pages
      200 data entries per page
    - unclustered tree index on title + author + year
      depth of 3 index pages
      100 data entries per page

Underlined attributes in the schema (e.g., book#) indicate components of the primary key.

Consider the following query.

```
select B.title, B.year, sum(P.qnty) * B.price as total
    from Purchase P, Books B
    where P.book# = B.book# and
          B.author = 'Agatha Christie'
    group by B.book#, B.title, B.year;
```

---

a. (2 points) Estimate the cardinality—the number of records returned—of the query *without* the group by and sum steps.

Also estimate the cardinality of the query itself (*with* the group by and sum steps).

---

b. (6 points) Assume that you have a quota of 20 buffer frames. Design a query plan for the query. Show the annotated query tree. Estimate the cost of your plan.

Part of the grading is based on how inexpensive your plan is. For full credit, your plan should cost less than 500 I/O's.

c. (2 points) Agatha Christie has written *many more* books than the average author (in **Books**). Given this statistic (say, that Agatha Christie has 66 books in **Books**), would the plan you generated for Question 4b have been different? Why or why not?

(Scratch space.)

(Scratch space.)