

# CSE-3421M Test #2

## “Queries”

**Family Name:** \_\_\_\_\_  
**Given Name:** \_\_\_\_\_  
**Student#:** \_\_\_\_\_  
**CS&E Account:** \_\_\_\_\_

- 
- **Instructor:** Parke Godfrey
  - **Exam Duration:** 75 minutes
  - **Term:** Winter 2013

Answer the following questions to the best of your knowledge. Your answers may be brief, but be precise and be careful. The exam is closed-book and closed-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, each with parts. Points for each question and sub-question are as indicated. In total, the test is out of 50 points.

In schemas, the underlined attributes denote a table's key. Attributes that are in *italics* are not nullable. Foreign keys are indicated by FK.

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

---

MARKING BOX	
<b>1.</b>	/15
<b>2.</b>	/15
<b>3.</b>	/10
<b>4.</b>	/10
<b>Total</b>	/50

1. (15 points) **SQL**. *San Carlos Municipal Airport*.

[Exercise]

Congratulations! You have just made the major leagues by landing a job with **Amazing.com**, the famous Internet bookstore.

<b>Customer</b> ( <u>cno</u> , <i>cname</i> , <i>address</i> )	<b>Stock</b> ( <u>isbn</u> , <u>from</u> , <i>qnty</i> , <i>price</i> )
<b>Author</b> ( <u>ano</u> , <i>aname</i> , <i>birth</i> , <i>country</i> )	FK ( <i>isbn</i> ) refs <b>Book</b>
<b>Book</b> ( <i>isbn</i> , <i>title</i> , <i>year</i> , <i>publisher</i> , <i>language</i> )	<b>Purchase</b> ( <u>cno</u> , <u>isbn</u> , <u>when</u> , <i>from</i> , <i>qnty</i> )
<b>Wrote</b> ( <u>ano</u> , <u>isbn</u> )	FK ( <i>cno</i> ) refs <b>Customer</b>
FK ( <i>ano</i> ) refs <b>Author</b>	FK ( <i>isbn</i> , <i>from</i> ) refs <b>Stock</b>
FK ( <i>isbn</i> ) refs <b>Book</b>	<b>Payment</b> ( <u>cno</u> , <u>when</u> , <i>amount</i> )
	FK ( <i>cno</i> ) refs <b>Customer</b>

Figure 1: **Amazing.com** Schema.

The basic schema of the main database at **Amazing.com** is shown in Figure 1. The underlined attributes indicate a table's primary key (and are hence not nullable). Attributes marked in *italics* are not nullable. Foreign keys are indicated by FK.

The attribute *price* in **Stock** tells how much a customer pays for that book. The attribute *from* in **Purchase** indicates where the stock was pulled from.

a. (6 points) Write an SQL query to report customers by *cno*, name, and address who have ever bought (**Purchase**) a book authored by an author whose nationality (*country*) is *Canadian*.

Do not report duplicate rows.

```
select distinct C.cno, C.cname, C.address
  from Customer C, Purchase P, Wrote W, Author A
 where C.cno = P.cno
    and W.isbn = P.isbn
    and W.ano = A.ano
    and A.country = 'Canada';
```

- 
- 
- b. (6 points) Write an SQL query to report customers by `cno`, name, and address who have ever bought (**Purchase**) a book that has two or more authors for more than \$200. Do not report duplicate rows.

```
with
  Multi (isbn) as (
    select W.isbn
      from Wrote W
     group by W.isbn
    having count(*) > 1
  )
select distinct C.cno, C.cname, C.address
  from Customer C, Purchase P, Multi M, Stock S
 where C.cno = P.cno
    and P.isbn = M.isbn
    and P.isbn = S.isbn and P.from = S.from
    and S.price > 200;
```

---

---

c. (3 points) Consider the SQL query

```
select cname, cno, aname, ano
  from Customer, Author
except
select cname, cno, aname, ano
  from ( select cno, ano, isbn
        from Customer, Wrote
        except
        select P.cno, W.ano, W.isbn
        from Purchase P, Wrote W
        where P.isbn = W.isbn
        ) as X,
      Customer C, Author A
where X.cno = C.cno and X.ano = A.ano;
```

State in *plain, concise English* what the query does.

You get *zero* credit if you use database terms in your answer! (E.g., “Well, the query first *joins* two tables, taking the *projection* of...” does not count!)

*Which customers have bought every book of an author?*

---

---

2. (15 points) **General.** *All the beautiful choices!*

[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. The SQL “ORDER BY” clause<sup>1</sup>

- A. can be used to order the results of any sub-query.
  - B. can only be used as the last clause of a query to order the tuples of the answer.
  - C. may affect the results of a join.
  - D. may only be used when “GROUP BY” is used.
  - E. is required in any SQL query.
- 

b. The SQL statement “DELETE FROM R;”

- A. will remove all the tuples from **R**.
  - B. will remove no tuples from **R**.
  - C. may remove just some tuples from **R**.
  - D. is syntactically incorrect because it is missing a WHERE clause.
  - E. *Not enough information to determine.*
- 

c. In a *real* relational database system, if you try to join (natural join) tables **R** and **S** and **R** is empty (that is, it has no tuples),

- A. the system reports an error.
  - B. the answer set is an empty table.
  - C. the answer set is the same as table **S**.
  - D. the answer set consists of just one row.
  - E. an answer set is returned; however, the results are system dependent.
- 

For Questions 2d and 2e, consider a table **T** with attributes A and B.

---

d. Which of the following SQL queries is illegal?

- A. `select * from T;`
  - B. `select count(*) from T;`
  - C. `select count(*) from T group by A;`
  - D. `select count(*), max (B) from T group by A;`
  - E. `select max(count(*)) from T group by A;`
- 

e. Which of the following SQL queries is illegal?

- A. `select A from T;`
  - B. `select A, count(*) from T;`
  - C. `select A, count(*) from T group by A;`
  - D. `select A, count(*) from T group by A, B;`
  - E. `select A, B, count(*) from T group by A, B;`
- 

<sup>1</sup>Accepted ‘A’ too for this. By the SQL standards, only ‘B’ is correct; but most database engines allow “order by” in sub-queries now, regardless.

f. In relational algebra, the intersection operator ( $\cap$ ) is *not* logically redundant if we only have additionally

- A. join ( $\bowtie$ )
- B. crossproduct ( $\times$ ), select ( $\sigma$ ), and project ( $\pi$ )
- C. difference ( $-$ ) and union ( $\cup$ )
- D. crossproduct ( $\times$ ) and difference ( $-$ )
- E. crossproduct ( $\times$ ) and union ( $\cup$ )

g.  $\mathbf{R} \cap \mathbf{S}$  is equivalent to

- A.  $\mathbf{R} - (\mathbf{R} - \mathbf{S})$
- B.  $\mathbf{R} - (\mathbf{S} - \mathbf{R})$
- C.  $(\mathbf{R} - \mathbf{S}) - \mathbf{R}$
- D.  $-((-\mathbf{R}) \cup (-\mathbf{S}))$
- E. *There is not enough information to answer this.*

h. Consider the schema

$\mathbf{R}(\underline{\mathbf{A}}, \mathbf{B})$  FK (B) refs  $\mathbf{S}$

$\mathbf{S}(\mathbf{A}, \underline{\mathbf{B}})$  FK (A) refs  $\mathbf{R}$

*Note that none of the attributes are nullable.*

Which of the following is guaranteed to produce as many as, or more, tuples than each of the others?

- A.  $\mathbf{R} \bowtie \mathbf{S}$
- B.  $\pi_{\mathbf{A}}(\mathbf{R}) \bowtie \mathbf{S}$
- C.  $\mathbf{R} \bowtie \pi_{\mathbf{B}}(\mathbf{S})$
- D.  $\pi_{\mathbf{A}}(\mathbf{R}) \bowtie \pi_{\mathbf{B}}(\mathbf{S})$
- E. *There is not enough information to answer this.*

i. Consider the schema from Question 2h again.

Which of the following is guaranteed to produce as many as, or more, tuples than each of the others?

- A.  $\mathbf{R} \bowtie \mathbf{S}$
- B.  $\mathbf{R} \times \mathbf{S}$
- C.  $\mathbf{R} \cap \mathbf{S}$
- D.  $\mathbf{R} \cup \mathbf{S}$
- E. *There is not enough information to answer this.*

j. Say that A is the primary key of table  $\mathbf{R}$ , and that B is another, numeric attribute of  $\mathbf{R}$ . Let  $r$  be the number of tuples in table  $\mathbf{R}$ .

$$(\pi_{\mathbf{A}}(\sigma_{\mathbf{B} < 5}(\mathbf{R}))) \bowtie (\pi_{\mathbf{A}}(\sigma_{\mathbf{B} \geq 5}(\mathbf{R})))$$

How many tuples will result from the above RA expression?

- A. 0
- B. 1
- C.  $\frac{1}{2}r$
- D.  $r$
- E.  $2r$

---

---

k. Consider the following schema.

```
create table T (  
  c integer primary key,  
  d integer);  
create table S (  
  b integer primary key,  
  c integer references T(c) on delete cascade);  
create table R (  
  a integer primary key,  
  b integer references S(b) on delete set null);
```

Let tables **R**, **S**, and **T** contain  $r$ ,  $s$ , and  $t$  tuples, respectively. Let  $n$  be the number of tuples in the result of the following query:

```
select *  
  from R, S, T  
 where R.B = S.B and S.C = T.C;
```

What is the lowest bound on  $n$  that you can say with certainty?

- A.**  $n \leq r$
  - B.**  $n \leq t$
  - C.**  $n \leq r * s * t$
  - D.**  $n \leq \min(r, s, t)$
  - E.** *It is impossible to establish a bound with the information provided.*
-

---

For Questions 2l to 2n, one of the choices is not like the others. That is, one of the choices could evaluate to a different answer than the others do. Choose the one that may evaluate differently.

---

l. Consider the relations  $\mathbf{R}(A, B)$ ,  $\mathbf{S}(B, C)$ , and  $\mathbf{T}(C, A)$ .

- A.  $\pi_{A,B}((\mathbf{R} \bowtie \mathbf{S}) \bowtie \mathbf{T})$
  - B.  $\pi_{A,B}(\mathbf{R} \bowtie (\mathbf{T} \bowtie \mathbf{S}))$
  - C.  $\mathbf{R} \bowtie \pi_{A,B}(\mathbf{S} \bowtie \mathbf{T})$
  - D.  $\pi_{A,B}(\mathbf{R} \bowtie \mathbf{S}) \bowtie \pi_{A,B}(\mathbf{S} \bowtie \mathbf{T})$
  - E.  $\pi_{A,B}(\mathbf{R} \bowtie \mathbf{T}) \bowtie \pi_{A,B}(\mathbf{R} \bowtie \mathbf{S})$
- 

m. Consider the schema  $\mathbf{R}(\underline{A}, B, C)$ ,  $\mathbf{S}(\underline{A}, \underline{D})$ , and  $\mathbf{T}(\underline{D}, B, E)$ .

- A.  $\pi_B((\sigma_{C=5}(\mathbf{R}) \bowtie \sigma_{E=7}(\mathbf{T})) \bowtie \mathbf{S})$
  - B.  $\pi_B(\pi_{A,B,D}(\sigma_{(C=5) \wedge (E=7)}(\mathbf{R} \bowtie \mathbf{T})) \cap (\mathbf{S} \times \pi_B(\mathbf{T})))$
  - C.  $\pi_B((\sigma_{C=5}(\mathbf{R} \bowtie \mathbf{S})) \bowtie (\sigma_{E=7}(\mathbf{S} \bowtie \mathbf{T})))$
  - D.  $\pi_B(\sigma_{C=5}(\mathbf{R} \bowtie \mathbf{S})) \cap \pi_B(\sigma_{E=7}(\mathbf{S} \bowtie \mathbf{T}))$
  - E.  $\pi_B(\pi_{A,B}(\sigma_{C=5}(\mathbf{R})) \bowtie (\mathbf{S} \bowtie \pi_{D,B}(\sigma_{E=7}(\mathbf{T}))))$
- 

n. Consider the schema  $\mathbf{R}(\underline{A}, \underline{B})$ , and  $\mathbf{S}(\underline{B}, \underline{C})$ . (There are no FKs!)<sup>2</sup>

- A.  $\pi_{A,C}(\mathbf{R} \bowtie \mathbf{S})$
  - B.  $\{\langle A, C \rangle \mid \exists B(\langle A, B \rangle \in \mathbf{R} \rightarrow \langle B, C \rangle \in \mathbf{S})\}$
  - C.  $\{\langle A, C \rangle \mid \neg \forall B(\langle A, B \rangle \in \mathbf{R} \rightarrow \langle B, C \rangle \notin \mathbf{S})\}$
  - D. `select distinct R.A, S.C from R, S where R.B = S.B;`
  - E. `select distinct R.A, S.C from R, S s1  
where R.B in (select s2.B from S s2 where s1.C = s2.C);`
- 

o. Consider the schema  $\mathbf{R}(A, B)$ ,  $\mathbf{S}(B, C)$ , and  $\mathbf{T}(C, A)$ . (There are no FKs!)

- A.  $\{\langle A \rangle \mid \neg \exists B, C(\langle A, B \rangle \in \mathbf{R} \wedge \langle B, C \rangle \in \mathbf{S} \wedge \langle C, A \rangle \in \mathbf{T})\}$
  - B.  $\{\langle A \rangle \mid \neg \exists B(\langle A, B \rangle \in \mathbf{R} \wedge \exists C(\langle B, C \rangle \in \mathbf{S} \wedge \langle C, A \rangle \in \mathbf{T}))\}$
  - C.  $\{\langle A \rangle \mid \forall B(\langle A, B \rangle \in \mathbf{R} \rightarrow \forall C(\langle C, A \rangle \notin \mathbf{T} \vee \langle B, C \rangle \notin \mathbf{S}))\}$
  - D.  $\{\langle A \rangle \mid \forall B(\langle A, B \rangle \in \mathbf{R} \wedge \forall C(\langle B, C \rangle \in \mathbf{S} \wedge \forall A2(\langle C, A2 \rangle \in \mathbf{T} \rightarrow A \neq A2))\}$
  - E.  $\{\langle A \rangle \mid \forall C(\langle C, A \rangle \in \mathbf{T} \wedge \forall B(\langle B, C \rangle \in \mathbf{S} \wedge \forall A2(\langle A2, B \rangle \in \mathbf{R} \rightarrow A \neq A2))\}$
- 

<sup>2</sup>Gave everyone credit for Questions 2n and 2o, regardless, because of typos (fixed here) in the test as given. They are good questions, though!



3. (10 points) **Relational Algebra & Calculus.** *The area under what?!*

[Short Answer]

R	
A	B
1	2
3	2
5	6
7	8
9	8

S	
B	C
6	2
2	4
8	1
8	3
2	5

T	
A	C
7	1
1	2
9	3
5	4
3	5

Three tables: **R**, **S**, & **T**.

a. (3 points) Show the results of  $\mathbf{R} \bowtie \mathbf{S}$ .

A	B	C
1	2	4
1	2	5
3	2	4
3	2	5
5	6	2
7	8	1
7	8	3
9	8	1
9	8	3

b. (3 points) Show the results of  $\mathbf{R} \bowtie (\mathbf{S} \bowtie \mathbf{T})$ .

A	B	C
3	2	5
7	8	1
9	8	3

- 
- 
- c. (2 points) Consider any three tables **J**, **K**, and **L**. Is it possible that  $\mathbf{J} \bowtie (\mathbf{K} \bowtie \mathbf{L})$  and  $(\mathbf{J} \bowtie \mathbf{K}) \bowtie \mathbf{L}$  might evaluate to different tables?

If so, construct a *small* example that demonstrates this.

If not, *briefly* explain why.

*No. Join is associative.*

- 
- d. (2 points) Which is SQL based upon, the *domain* relational calculus or the *tuple* relational calculus?

Briefly justify your answer.

*Tuple-relational calculus. Variables (in SQL, range variables) are over a domain of tuples, not over domains of individual attributes as in the domain-relational calculus (which is standard in predicate-calculus notation).*

---

---

EXTRA SPACE.

---

---

4. (10 points) **Aggregation.** *Groups for a thousand, Alex.*

[Exercise]

**Person**(sin, name, gender, job, studio, title, year)  
FK (studio) references **Studio**,  
FK (title, year) references **Movie**  
**Category**(cat)  
**Movie**(title, year, length, *director*, *studio*, *cat*)  
FK (director) references **Person** (sin),  
FK (studio) references **Studio**,  
FK (cat) references **Category**  
**Cast**(sin, title, year, character)  
FK (sin) references **Person**,  
FK (title, year) references **Movie**  
**Studio**(studio, city)

Figure 2: Movies and Casts.

Write SQL queries for the following with respect to the schema in Figure 2.

- 
- a. (3 points) For each *director* (someone who has *directed* at least one movie), report the person's *sin*, name, gender, and the number of movies that he or she has directed. Call this last column *#movies*.

```
select P.sin, P.name, P.gender, count(*) as #movies
  from Person P, Movie M
 where P.sin = M.director
 group by P.sin, P.name, P.gender
```

- 
- b. (4 points) For each actor (someone who has been *cast* in a movie) and movie category (*cat*), report the *sin*, name, gender, and title of the actor and the number of movies of *that* category in which he or she has been cast. Call this last column *#movies*.  
Only count a given actor or actress *once* per movie; that a person may have played several *characters* in the movie does *not* count multiple times.

```
with
  InMovie (cid, title, year) as (
    select distinct cid, title, year
      from Cast
  )
select P.sin, P.name, P.gender, P.title, M.cat, count(*) as #movies
  from Person P, InMovie C, Movie M
 where P.sin = C.sin
    and C.title = M.title and C.year = M.year
 group by P.sin, P.name, P.gender, P.title, M.cat;
```

---

---

EXTRA SPACE.

- c. (3 points) For each movie, report how many male actors and how many female actresses—gender = 'M' for *male* and gender = 'F' for *female*—were *cast* in the movie in columns *male* and *female*, respectively.

Again, only count a given actor or actress *once* per movie; that a person may have played several *characters* in the movie does *not* count multiple times.

The column *male* or *female* should report 0 (zero) if there were no actors or actresses, respectively, in the movie.

```
with
Counts (title, year, male, female) as (
  select title, year, 0, 0
    from Movie
  union
  select C.title, C.year, count(distinct C.sin), 0
    from Cast C, Person P
   where C.sin = P.sin
      and P.gender = 'M'
  union
  select C.title, C.year, 0, count(distinct C.sin)
    from Cast C, Person P
   where C.sin = P.sin
      and P.gender = 'F'
)
select C.title, C.year, max(C.male) as male, max(C.female) as female
  from Counts C
 group by C.title, C.year;
```

EXTRA SPACE.

HAPPY PI DAY!

RELAX. BREATHE. TURN IN YOUR TEST. GO HOME.