

UNIVERSITY OF CALIFORNIA  
College of Engineering  
Department of EECS, Computer Science Division

CS186  
Fall 2003

Eben Haber  
Midterm

**Midterm Exam: Introduction to Database Systems**

This exam has seven problems, worth different amounts of points each. Each problem is made up of multiple questions. You should read through the exam quickly and plan your time-management accordingly. Before beginning to answer a question, be sure to read it carefully and to *answer all parts of every question!*

Please **write your login** at the top of **every page**; you must turn in all the pages of the exam. **Do not** remove pages from the stapled exam! Two pages of extra answer space have been provided at the back in case you run out of space while answering. If you run out of space, be sure to make a “forward reference” to the page number where your answer continues.

Good luck!

<b>I</b>		/24
<b>II</b>		/10
<b>III</b>		/6
<b>IV</b>		/10
<b>V</b>		/10
<b>VI</b>		/20
<b>VII</b>		/20

<b>Total</b>		/100
--------------	--	------

**I. Query Languages (30 points).**

Given the following database tables, choose all queries in SQL, Relational Algebra or Tuple Relational Calculus that return the proper answer to the corresponding question. Note that each question may have more than one correct answers. Circle all that apply.

Primary keys are underlined. Note that the driver involved in a car accident may not always be the owner of the car. Assume that `accident_date` is of type integer, and represents a year (e.g. 1980). Year is also of type integer. We assume that a car cannot get involved in more than one accident at a certain date.

Person(SSN, name, address)

Car(license, year, model)

Accident(license, accident\_date, driver, damage\_amount)

Owns(SSN, license)

1. (4 points) Find the SSN of every person who owns one or more cars, none of which has ever been involved in a car accident.

A. SELECT O.SSN

FROM Owns O  
WHERE O.license NOT IN (SELECT A.license  
FROM Accident A)

B.  $\pi_{SSN}(Owns) - \pi_{SSN}(Owns \triangleright \triangleleft Accident)$

C.  $\{O \mid \exists O1 \in Owns(O.SSN = O1.SSN \wedge \neg \exists A \in Accident(O1.license = A.license))\}$

D. None of the above

2. (4 points) Find the SSN of every person, who owns a TOYOTA or a DODGE

A. SELECT O.SSN

FROM Owns O, Car C  
WHERE O.license=C.license AND  
(C.model='TOYOTA' OR  
C.model='DODGE')

B.  $\pi_{SSN}(Owns \triangleright \triangleleft \sigma_{model='TOYOTA' \vee model='DODGE'}(Car))$

C.  $\{O \mid \exists O1 \in Owns(O1.SSN = O.SSN \wedge \exists C \in Car(O1.license = C.license \wedge (C.model = 'TOYOTA' \vee C.model = 'DODGE')))\}$

D. None of the above

3. (4 points) Find the SSN of all persons who have had all of their cars involved in an accident.

- A. 

```
SELECT O.SSN
FROM Owns O, Accident A
WHERE O.license=A.license
EXCEPT
SELECT O.SSN
FROM Owns O
WHERE O.license NOT IN (SELECT A.license
                        FROM Accident A)
```
- B.  $\pi_{SSN,license}(Owns \bowtie Accident) \div \pi_{license}(Accident)$
- C.  $\{O \mid \exists O1 \in Owns(O1.SSN = O.SSN \wedge \forall O2 \in Owns(O1.SSN = O2.SSN \Rightarrow \exists A \in Accident(O2.license = A.license)))\}$
- D. None of the above.

4. (4 points) Who is the driver who participated in the most costly accident? Return the driver and the amount of damage.

- A. 

```
SELECT driver, damage_amount
FROM Accident
WHERE damage_amount IN (SELECT MAX(damage_amount)
                        FROM Accident)
```
- B. 

```
SELECT driver, damage_amount
FROM Accident
WHERE damage_amount=MAX(damage_amount)
```
- C. 

```
(SELECT driver, damage_amount FROM Accident)
EXCEPT
(SELECT a1.driver, a1.damage_amount
FROM Accident a1, Accident a2
WHERE a1.damage_amount<a2.damage_amount AND a1.driver<>a2.driver)
```
- D. None of the above.

5. (4 points) Find the license number of all cars that had an accident in less than 5 years after their production. Return the license number and the number of years

- A. 

```
SELECT c.license
FROM Car c, Accident a
WHERE a.accident_date-c.year<5
```
- B.  $\pi_{license}(\sigma_{accident\_date-year<5}(Car \bowtie Accident))$

C.  $\{C \mid \exists C1 \in Car(C.license = C1.license \wedge \exists A \in Accident(C1.accident\_date - A.year < 5))\}$

D. None of the above

6. (4 points) Find the accident with the smallest cost and return the corresponding owner of the car, his/her address and the amount of damage.

A. SELECT O.SSN, P.address, MIN(A.damage\_amount)  
FROM Accident A, Owns O, Person P  
WHERE O.license=A.license AND O.SSN=P.SSN  
GROUP BY O.SSN, P.address

B. SELECT O.SSN, P.address, (SELECT MIN(damage\_amount) FROM Accident)  
FROM Owns O, Person P, Accident A  
WHERE O.SSN=P.SSN AND A.license=O.license

C. SELECT O.SSN, P.address, A.damage\_amount  
FROM Owns O, Person P, Accident A  
WHERE O.SSN=P.SSN AND A.damage\_amount=(SELECT MIN(damage\_amount)  
FROM Accident)

D. None of the above.

7. (4 points) Find the license number of all cars that have been involved in more than one accident. (DO NOT RETURN DUPLICATES)

A. SELECT A1.license  
FROM Accident A1, Accident A2  
WHERE A1.license=A2.license AND A1.accident\_date<>a2.accident\_date

B. SELECT DISTINCT A1.license  
FROM Accident A1  
WHERE A1.license IN (SELECT A2.license  
FROM Accident A2  
WHERE A1.accident\_date<>A2.accident\_date)

C. SELECT license  
FROM Accident  
GROUP BY license  
HAVING COUNT(accident\_date)>1

D. None of the above

8. (2 points) Explain, in English, what the following Relational Calculus query expresses:

$$\left\{ C \mid \exists C1 \in Car \left( C1.model = C.model \wedge \forall C2 \in Car \left( C1.model = C2.model \Rightarrow \left( \exists A \in Accident (A.license = C2.license) \right) \right) \right) \right\}$$

**II. External Sorting (10 points)**

1. (6 points) Suppose we wish to sort the following values:

84, 22, 19, 11, 60, 68, 31, 29, 58, 23, 45, 93, 48, 31, 7

Assume that:

- you have three pages of memory for sorting
- you will use an external sorting algorithm with a 2-way merge
- a page only holds two values

For each sorting pass, show the contents of all temporary files.

2. (2 points) If you have 100 pages of data, and 10 pages of memory, what is the minimum number of passes required to sort the data.

3. (2 points) If you have 500,000 pages of data, what is the minimum number of pages of memory required to sort the data in 3 passes?

**III. Join Algorithms (6 points)**

Consider a relation R with attributes (x,y) and a relation S with attributes (y, z). Column y in S is a key and the set of values of y in R are the same as the set of values of y in S. Assume that there are no indexes available and that there are 25 pages in the buffer available. Table R is 1,500 pages with 50 tuples per page. Table S is 400 with 100 tuples per page.

Compute the I/O costs for the following joins. Assume the simplest cost model, where pages are read and written one at a time.

A. Block nested loops join with R as the outer relation and S as the inner relation

B. Block nested loops join with S as the outer relation and R as the inner relation

C. Sort merge join with R as the outer relation and S as the inner relation

D. Sort merge join with S as the outer relation and R as the inner relation

E. Hash join with S as the outer relation and R as the inner relation

**IV Query Optimization**

Consider the following schema

Sailors(sid, sname, rating, age)

Reserves(sid, did, day)

Boats(bid, bname, size)

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:

Reserves contains 10,000 records with 40 records per page.

Sailors contains 1000 records with 20 records per page.

Boats contains 100 records with 10 records per page.

There are 50 values for Reserves.bid.

There are 10 values for Sailors.rating (1...10)

There are 10 values for Boat.size

There are 500 values for Reserves.day

Consider the following query

```
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
      B.size > 5 AND R.day = 'July 4, 2003';
```

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

Consider the following query

```
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) Draw all possible left-deep query plans for this query:

(c) For the first join in each query plan (the one at the bottom of the tree, what join algorithm would work best? Assume that you have 50 pages of memory. There are no indexes, so indexed nested loops is not an option.



**V Buffer Management (10 points total)**

In this exercise, you will compare the LRU, MRU, and CLOCK buffer management algorithms for a given workload (just as you did in the first part of homework 1). Here are the assumptions you must make:

- Assume there are four page slots your buffer manager must manage: P1, P2, P3, and P4.
- All four slots are empty to start.
- When the buffer pool has unused slots (such as at the beginning, when all four slots are empty), it will put newly read data in the leftmost empty slot (e.g., if slots 2 and 3 are free, it will use slot 2).
- The pages to be read from disk are labelled A through G.
- For each access the page is pinned, and then immediately unpinned.

Below are three tables for describing the contents of the buffer pool at each time step. A page is read at the beginning of each time step. You should record, in the table, the contents of each Buffer Page after the new page has been read in.

LRU (3 points)					
Time	Page Read	Buffer Page			
		P1	P2	P3	P4
1	G				
2	F				
3	E				
4	D				
5	G				
6	C				
7	F				
8	E				
9	D				
10	C				
11	D				
12	B				
13	A				
14	D				
15	F				
Total number of pages "evicted":					

MRU (3 points)					
Time	Page Read	Buffer Page			
		P1	P2	P3	P4
1	G				
2	F				
3	E				
4	D				
5	G				
6	C				
7	F				
8	E				
9	D				
10	C				
11	D				
12	B				
13	A				
14	D				
15	F				
Total number of pages "evicted":					

CLOCK (4 points)					
Time	Page Read	Buffer Page			
		P1	P2	P3	P4
1	G				
2	F				
3	E				
4	D				
5	G				
6	C				
7	F				
8	E				
9	D				
10	C				
11	D				
12	B				
13	A				
14	D				
15	F				
Total number of pages "evicted":					

**VI. Functional Dependencies and Normalization (20 points total)**

Consider the attributes A B C D E F G which have the following functional dependencies:

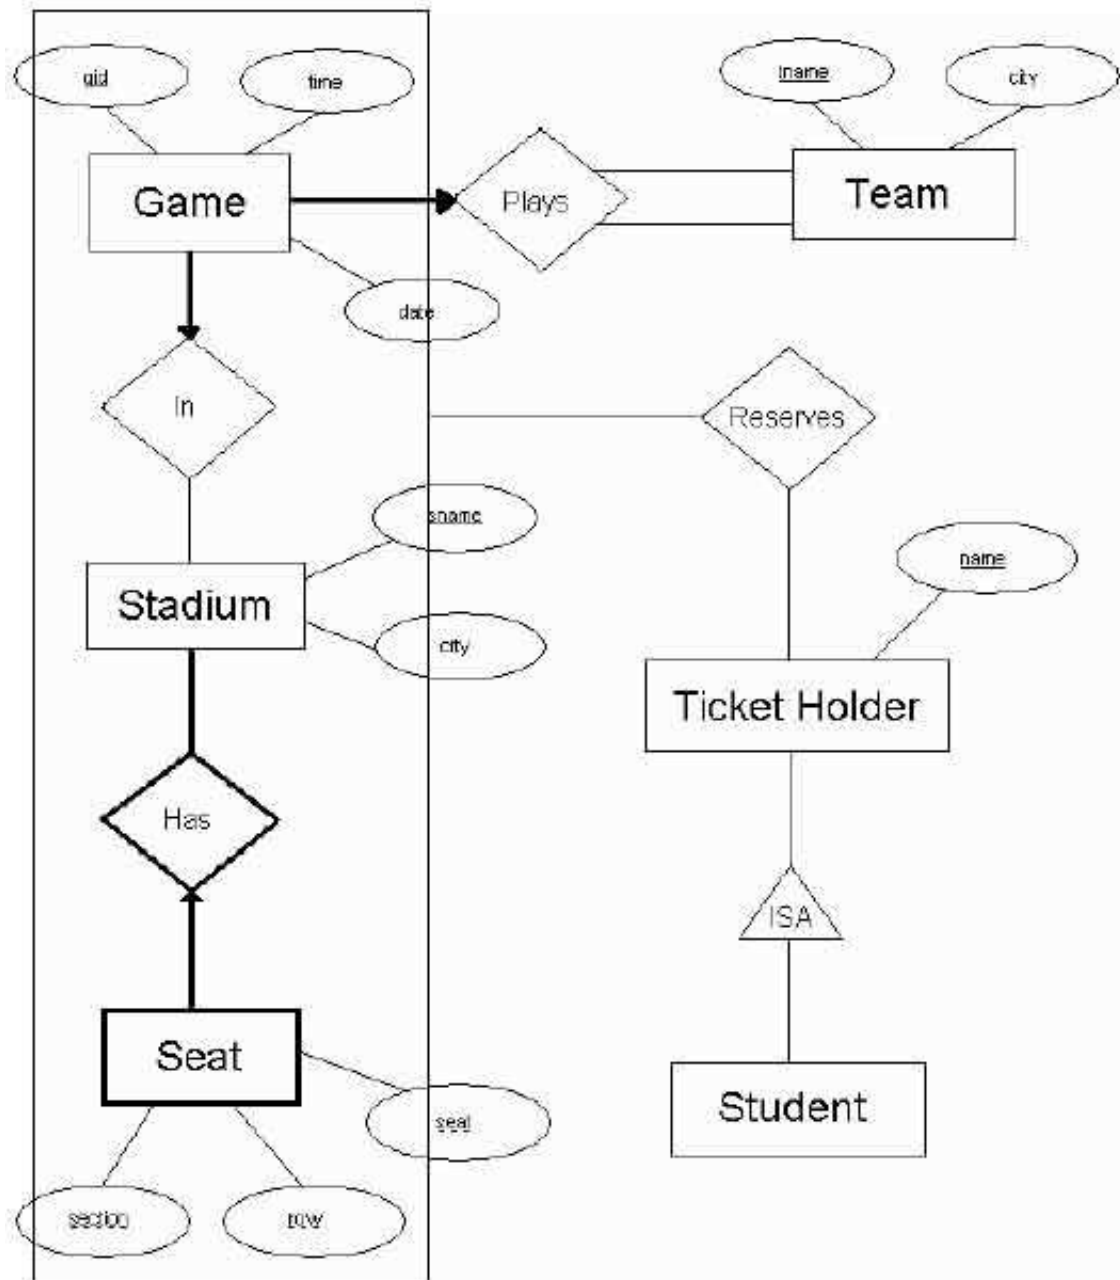
AD → F  
 AE → G  
 DF → BC  
 E → C  
 G → E

1. List all candidate keys (not superkeys): (5 points)
  
2. Which of the following dependencies are implied by those above: (5 points)
  - a. (IS) (IS NOT) G → C
  - b. (IS) (IS NOT) A → G
  - c. (IS) (IS NOT) ADF → E
  - d. (IS) (IS NOT) ADC → B
  - e. (IS) (IS NOT) AGF → E
  
3. Consider the decomposition into 4 relations: (ADF) (CE) (EG) (ABDG).  
 This decomposition: (5 points)
  - a. (IS) (IS NOT) in BCNF
  - b. (IS) (IS NOT) in 3NF
  - c. (IS) (IS NOT) in 1NF
  - d. (IS) (IS NOT) Dependency Preserving
  - e. (IS) (IS NOT) Lossless
  
4. Consider the decomposition into 3 relations: (ADF) (EC) (ABDEG).  
 This decomposition: (5 points)
  - a. (IS) (IS NOT) in BCNF
  - b. (IS) (IS NOT) in 3NF
  - c. (IS) (IS NOT) 1NF
  - d. (IS) (IS NOT) Dependency Preserving
  - e. (IS) (IS NOT) Lossless

**VII. Entity-Relational Model (20 Points)**

a) (10 points) Draw an E-R diagram for the following situation:

- This is a simplified model for reserving baseball tickets.
- There are teams, which are identified by the team name. Teams are also located in a city.
- Teams play each other in games, which occur on a particular date at a particular time. Games are identified by a game ID, and each game has exactly two teams that play in it.
- A game is played in exactly one stadium.
- A stadium is identified by its name, and is also located in a city.
- Stadiums have seats, which have a section number, a row number, and a seat number.
- Ticket holders reserve seats for a game. Ticket holders are identified by their name.
- Some ticket holders are students (students get discounts, but we are not including that in the model).



b) (10 points) Using SQL, convert the following ER diagram to the relational model.  
 Hint: You do not need CHECK constraints.

