CS4604 Final Exam

May 7, 2001

Please enter the following information:

- **Name:**

- **ID:**

GOOD LUCK and have a nice summer!!
Do not write below this line

Problem	Max Score	Score
1	30	
2	10	
3	10	
4	20	
5	10	
6	10	
7	10	
8	15 (XC)	
Total	100	

1. (30 points) Short answer questions:

   (a) If relation R has m tuples and S has n tuples, what is the maximum number of tuples that $R \cup S$ can have? Assume set-theoretic semantics.

   (b) What was covered in Dr. Chris North's guest lecture? Write enough to indicate that you know what you are talking about.

   (c) How many nontrivial MDs are possible in a two-column relation?

   (d) Which forms of VIEWs are not updatable?

   (e) Why does the SQL SELECT operator not remove duplicates by default?

   (f) Explain what a referential integrity constraint means.

   (g) (True/False) In $R(A, B, C)$, if $AB \rightarrow\rightarrow C$, then $A \rightarrow C$.

   (h) Five popular techniques for database tuning were presented in class. Mention any four of them.

   (i) Express the natural join operator $\bowtie$ in terms of other basic operations: $\{\cup, \cap, -, \times, \sigma, \pi\}$.

   (j) What command would you use to open a database of English language puns?

2. (10 points) Consider the following E/R diagram that models highways and the cities they connect. A highway (e.g., I-81 North) has a starting city, an ending city, and can go through any number of intermediate cities (including zero). Notice that a two-way highway will be modeled as two separate highways (thus, I-81 North is different from I-81 South). Does the diagram model this information accurately? If yes, explain how. If not, explain why not and redraw the diagram so that it is consistent with the above specifications.

3. (10 points) Suppose you are given a relation $R(A, B, C, D, E)$ with

$$R \bowtie_{(R.B=S.B \ \ AND \ \ (R.C<>S.C \ \ OR \ \ R.D.<>S.D))} S = \emptyset$$

where relation S is the same as R and $\emptyset$ is the empty set. Decompose R into a collection of relation schemas that are in BCNF. Explain how you performed the decomposition.

4. (20 points) Consider the relation `HasTaken(student-name,course-number)`, which indicates the courses taken by various students. Write a query in relational algebra to find the students who have taken the same number of courses as Tommy Hanks.

5. (10 points) Consider the relational schema:

```
HasTaken(student-ssn,course-number)
RequiredForGraduation(course-number)
```

with their obvious interpretations *i.e.,* `HasTaken` identifies the courses taken by each student and `RequiredForGraduation` lists all the courses required to graduate. For example, if `RequiredForGraduation` contains 14 courses, this means that all those 14 courses are required to graduate (you may assume that all the `course-number`s in `RequiredForGraduation` will be different). Write a Datalog query to find the students who have satisfied **some, but not all** requirements for graduation.

6. (10 points) Consider the relational schema:

```
Father(son-name,parent-name)
```

which is a list of sons and their parents. Example data for this relation is:

```
Father(Jimmy,Tommy)
Father(Marc,Tommy)
Father(Linus,Jimmy)
Father(Bruce,Jimmy)
Father(Goodman,Marc)
Father(John,Marc)
```

Assume that we are interested in finding the earliest common ancestor (ECA) of a given pair of people. For example, the earliest common ancestor of Linus and Bruce is their father, Jimmy. The earliest common ancestor of Goodman and Bruce is their grandfather, Tommy. The earliest common ancestor of Bruce and Marc is also Tommy, and so on. Give the definition of a Datalog predicate

```
ECA(person1-name,person2-name,ancestor-name)
```

to find the earliest common ancestor of `person1-name` and `person2-name`. For example, the above three examples of ECA can be expressed as:

```
ECA(Linus,Bruce,Jimmy).
ECA(Bruce,Linus,Jimmy).
ECA(Goodman,Bruce,Tommy).
ECA(Bruce,Goodman,Tommy).
ECA(Bruce,Marc,Tommy).
ECA(Marc,Bruce,Tommy).
```

7. (10 points) Consider the relation schema

```
Ships(name,class,launched)
Battles(name,date)
Outcomes(ship,battle,result)
```

where relation `Ships` records the name of a ship, the name of its class, and the year in which the ship was launched. Relation `Battles` gives the name and date of battles involving ships, and relation `Outcomes` gives the result ('sunk', 'damaged', or 'ok') for each ship in each battle. Write a query in SQL to find those ships that 'lived to fight another day'; i.e., they were damaged in one battle, but later fought in another.

8. (Extra credit question: 15 points) Consider the following cryptarithmetic puzzle (on the left):

```
  SEND                   9567
+ MORE                 + 1085
------                 ------
 MONEY                  10652
------                 ------
```

The goal of the puzzle is to substitute numbers for each of the letters so that the addition works out. So, if the number 5 is substituted for E, then the same value should be used for all occurences of E. The answer to the puzzle is given on the right (in case you are interested). Explain how you can solve such cryptarithmetic puzzles by running an SQL database query. For full credit, you have to explain clearly (i) how the above problem is modeled in database tables, (ii) what the tables look like, and what type of information they contain, and (iii) what the query looks like. Again, write enough to indicate that you know what you are talking about.