CS5614 Midterm Exam

October 29, 1999

Please enter the following information:

- **Name:**

- **ID:**

GOOD LUCK!
Do not write below this line

Problem	Max Score	Score
1	30	
2	10	
3	10	
4	20	
5	20	
6	10	
Total	100	

1. (30 points) Short answer questions:

   (a) What is magic?

   (b) Assume that a disk block is 4096 bytes and that keys and pointers take up 4 and 8 bytes respectively. What is the maximum fanout of a disk block (in an index structure), assuming blocks only have pointers to child nodes?

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

   (d) Give example data for a relation $R(A, B, C)$ for which the following hold: $\{A \rightarrow\rightarrow B, A \rightarrow C, CB \rightarrow\rightarrow A\}$.

   (e) Calculate the number of bushy joins of 4 relations. Include all possible permutations of the relations.

   (f) Assume that the FD $A \rightarrow B$ holds in a relation $R(A, B)$. Express this FD as a constraint in relational algebra.

2. (10 points) An E/R diagram when converted to relations (using the method that we now know so well), gives rise to the following schema (without normalization): `R(a,b,c)`, `S(a, d)` and `T(a,d,f,g)`. You may assume that the same symbols refer to the same attribute and that different symbols refer to different attributes (e.g. the `a`, in `R`, `S`, and `T` are the same). Your task is to give at least *two* different E/R diagrams that could have produced these relations.

3. (10 points) Prove that every two-column relation is in 3NF.

4. (20 points) Consider the relation $R(A, B, C, D, E)$ with MDs $\{A \twoheadrightarrow B, AB \twoheadrightarrow C\}$ and FDs $\{A \rightarrow D, AB \rightarrow E\}$. Decompose the relation into a collection of relation schemas in 4NF.

5. (20 points) Consider the relational schema modeling Web pages:

```
Webpage(URL, author, date)
Link(URL1, URL2)
```

where there is a `Link` between two `URL`s if there is a physical hyperlink on the Internet from `URL1` to `URL2` (There may or may not be a `Link` from an `URL` to itself). The 'master relation' `Webpage` models the details of each `URL`, like its creator and the date it was created. You may assume that `URL` is the key for `Webpage`. A web page is called an *authority* if *all other* web pages link to it. Write a query in Datalog to find all the authorities. You are not allowed to use relational algebra or SQL.

6. (10 points) Consider the relational 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 relational algebra or SQL to find those ships that 'lived to fight another day'; i.e., they were damaged in one battle, but later fought in another. Datalog notation is not permissible.