Stanford University Computer Science Department
2006 Comprehensive Exam in Databases
SAMPLE SOLUTION

1. *Data Models*

   (a) Why does the E/R model require a notion of "weak entity set"? Be as specific as you can. (For example, "Because some entity sets don't have a key" is not an acceptable answer.)

   There are entity sets whose members cannot be distinguished by their own attributes, but require the help of entities to which they are linked by relationships. These are weak entity sets.

   (b) Why do object-oriented models like ODL *not* have a notion equivalent to weak entity sets? Again, be as specific as you can.

   Since object-oriented models have a notion of object-identity, all members of a class are distinguishable from one another, even if they have the same attribute-values.

2. *Relational Algebra*

   Suppose there are relations $R(A, B)$ and $S(B, C)$. Write the following queries in *relational algebra* using standard operators only. You may write either single expressions or sequences of assignments to temporary variables, the last of which is the desired result. If you want your answers to be considered for partial credit you may provide some explanation of your reasoning, but it is sufficient just to provide the correct relational algebra queries.

   (a) The tuples of R such that their B-component is neither a B-component of a tuple of S nor a C-component of a tuple of S.

   $RA := \pi_A(R)$
   $Svals := \pi_B(S) \cup \pi_C(S)$
   $Ans := R - (RA \times Svals)$

   (b) The A-values that appear in at least two distinct tuples of R.

   $\pi_A(\sigma_{A=D \wedge B \neq E}(R \times \rho_{T(D,E)}(R)))$

(c) The A-values a such that:

i. a appears in at least one tuple of R, and

ii. Every B-value b, for which (a, b) is a tuple of R, appears as a B-value in at least one tuple of S.

RA := $\pi_A(R)$
$BadB$:= $\pi_B(R) - \pi_B(S)$
$BadA$:= $\pi_A(R \cap (RA \times BadB)$
Ans := $RA - BadA$

3. *Relational Design*

Suppose we are given relation $R(A, B, C, D)$ with functional dependencies $AB \rightarrow C$, $CD \rightarrow A$, and $B \rightarrow D$. Answer each of the following four questions. In each part, you may choose to justify your answer for partial credit, but it is sufficient just to provide an answer.

(a) Find all the (minimal) keys of R.

AB and BC

(b) How many superkeys does R have?

6 (all sets that contain, not necessarily properly, either AB or BC or both)

(c) Of the given FD's, which violate Boyce-Codd normal form?

$CD \rightarrow A$ and $B \rightarrow D$

(d) Of the given FD's, which violate third normal form?

Only $B \rightarrow D$

4. *XML*

Consider the following four DTDs for XML documents containing data about books and authors. You may assume that `isbn` and `ssn` are "keys" for books and authors, in the sense that no two actual books have the same `isbn`, and no two actual authors have the same `ssn`.

```
DTD1: <!DOCTYPE BA [
        <!ELEMENT BA (Book*)>
        <!ELEMENT Book (Title, Author+)>
        <!ELEMENT Title (#PCDATA)>
        <!ELEMENT Author>
        <!ATTLIST Author ssn ID name CDATA> ]>

DTD2: <!DOCTYPE BA [
        <!ELEMENT BA (Author*)>
        <!ELEMENT Author (Book)>
        <!ATTLIST Author ssn ID name CDATA>
        <!ELEMENT Book (Title)>
        <!ELEMENT Title (#PCDATA)> ]>

DTD3: <!DOCTYPE BA [
        <!ELEMENT BA (Book*, Author*)>
        <!ELEMENT Book (Title)>
        <!ATTLIST Book isbn ID author IDREF>
        <!ELEMENT Title (#PCDATA)>
        <!ELEMENT Author>
        <!ATTLIST Author ssn ID name CDATA> ]>

DTD4: <!DOCTYPE BA [
        <!ELEMENT BA (Author*, Book*)>
        <!ELEMENT Author>
        <!ATTLIST Author ssn ID book IDREFS>
        <!ELEMENT Book (Title)>
        <!ATTLIST Book isbn ID>
        <!ELEMENT Title (#PCDATA)> ]>
```

(a) Which of the DTDs can be used for XML documents that encode *many-many* relationships between books and authors? (circle those that can) `DTD4`

(b) Which of the DTDs can be used for XML documents that encode *many-one* relationships from books to authors? (circle those that can) `DTD3 DTD4`

(c) Which of the DTDs can be used for XML documents that encode *one-many* relationships from books to authors? (circle those that can) `DTD1 DTD4`

(d) Which of the DTDs can be used for XML documents that encode *one-one* relationships between books and authors? (circle those that can) `DTD1 DTD2 DTD3 DTD4`

5. *SQL*

Consider the following SQL schema for a database about schools in a university and departments in those schools. Data types are omitted for brevity.

```
create table School(schoolName primary key, dean)
create table Dept(deptName primary key,
                  schoolName references School on delete cascade)
```

Write a *single SQL statement* that deletes all schools with more than 10 departments and also deletes all departments in those schools.

```
delete from School
where 10 < (select count(*) from Dept
            where Dept.schoolName = School.schoolName)
```

6. *Transactions*

Consider a relation `Temperatures(time,temp)` and the following transaction `T` with SQL isolation level `REPEATABLE READ`:

```
T: (Q1) Select Avg(temp) From Temperatures;
   (Q2) Select Avg(temp) From Temperatures;
```

Is it possible for a concurrently-running transaction to cause queries `Q1` and `Q2` to return different values? If so, show the simplest such transaction. If not, explain why not.

Yes. This transaction could run between `Q1` *and* `Q2`:

```
insert into Temperatures values (12:00:00, 1000)
```