

# CSE 444 Midterm --- Winter 2005

Name \_\_\_\_\_

**Total number of points: 100 points**

**Time: 50 minutes.**

1. [40 points] Consider the following relational schema for database with information about TV Series. For example, “Friends Season 1” and “Friends Season 2” would be two different TV Series.

- *TVSERIES*(seriesname, season, year)
- *ACTORS*(stagename, lname, fname, gender, dob, dod)
- *TVSERIESCASTS*(seriesname, season, stagename)

(*seriesname*, *season*) forms the primary key for TVSERIES. For the table ACTORS, the primary key is *stagename*. TVSERIESCASTS has a foreign key to TVSERIES and one to ACTORS.

a. [30 points] Write SQL queries for the following

- [15 points] Find the names of all TV Series that have at least one male actor and at least one female actor. The male and female actors do not necessarily have to have cast in the same season of that series.

- ii.** [15 points] For each TV Series list the first name and last name of all the actors who acted in all the seasons of that series. Order series alphabetically; order actors by the last name.

**b.** [10 points] Write SQL statements that express the following constraints. You do not have to write the entire CREATE TABLE statement. You should write only the constraint in SQL.

**i.** Specify all the foreign keys that you can identify in the above schema.

**ii.** No actor can die before he was born

2. [20 points] A publishing company hires staff writers to write electronic documents and web pages. Draw an E/R diagram to model the following entity sets:

- writers: name(key)
- editors: name(key), phone
- document: filename(key), title
- webpage: filename(key), title, URL

and the following relationships

- writers write documents
- every writer has exactly one homepage (which is a webpage)
- every webpage is a document
- every editor is a writer
- every document is checked by exactly one editor

(a) [10 points] Draw the corresponding E/R diagrams.

(b) [10 points] Then design the relational schema for your E/R diagram. You have to turn in: the relation names, their attributes, their keys, and their foreign keys. You do not write SQL CREATE TABLE statements. Use a format like in the following example:

Product (Name, Price, MfgName) Key: Name, Price,  
Foreign Key MfgName to Company  
Company(name, address) Key: name

3. [20 points]

a. [10 points] Consider the relation  $R(A,B,C,D,E)$  and the functional dependencies:

$A \rightarrow B$   
 $D \rightarrow AC$   
 $DE \rightarrow C$

Compute the Boyce Codd Normal form for  $R$ . Your answer should consist of (1) a list of table names and attributes, and (2) an indication of the keys in each table. (like  $R1(A,\underline{B})$ ,  $R2(\underline{B},C)$ ,  $R3(\underline{B},\underline{D},E)$  -- but note that this is NOT the correct answer.)

b. [10 points] Consider two relations  $R(A,B,C)$ ,  $S(D,E,F)$  and the following functional dependencies:

In  $R$ :  $A \rightarrow C$   
In  $S$ :  $D,E \rightarrow F$

Consider the following view  $T(A,B,C,D,E,F)$  defined by:

```
CREATE VIEW T AS
  SELECT R.A, R.B, R.C, S.D, S.E, S.F
  FROM R, S
  WHERE R.B = 'xyz' and R.C = S.D
```

Find the keys in  $T$ .

4. [20 points] Consider the log below, where the <END CKPT> record has been omitted.
- c. [10 points] Assume that the log is UNDO log:

```
<START S>
<S, A, 60>
<COMMIT S>
<START T>
<T, A, 10>
<START CKPT (T)>
<START U>
<U, B, 20>
<T, C, 30>
<START V>
<U, D, 40>
<V, F, 70>
<COMMIT U>
<T, E, 50>
<COMMIT T>
<V, B, 80>
<COMMIT V>
```

- i. [5 points] Indicate where should the <END CKPT> record be? If there is no unique place, indicate the range where it might go.
- ii. [5 points] Assuming that <END CKPT> in a correct position, indicate which portion of the log the recovery manager must read in the case of a crash.

d. [10 points] Assume now that the log is a redo log:

```
<START S>
<S, A, 60>
<COMMIT S>
<START T>
<T, A, 10>
<START CKPT (T)>
<START U>
<U, B, 20>
<T, C, 30>
<START V>
<U, D, 40>
<V, F, 70>
<COMMIT U>
<T, E, 50>
<COMMIT T>
<V, B, 80>
<COMMIT V>
```

iii. [5 points] Indicate where should the <END CKPT> record be? If there is no unique place, indicate the range where it might go.

iv. [5 points] Assuming that <END CKPT> in a correct position, indicate which portion of the log the recovery manager must read in the case of a crash.