

CS 450 - Database Concepts

Spring 2012

Instructor: Dr. Jessica Lin

Homework 1 – (extended deadline) Due Feb 15 at 4:35pm (both electronically and in-class)

Problem 1.

We may think of relationships in the ER model as having keys, just as entity sets do. Let R be a relationship among the entity sets E_1, E_2, \dots, E_n . Then a key for R is a set K of attributes chosen from the attributes of E_1, E_2, \dots, E_n such that if (a_1, a_2, \dots, a_n) and (b_1, b_2, \dots, b_n) are two different tuples in the relationship set for R , then it is not possible that these tuples agree in all the attributes of K . Now, suppose $n = 2$; that is, R is a binary relationship. Also, for each i , let K_i be a set of attributes that is a key for entity set E_i . In terms of E_1 and E_2 (or K_1 and K_2), give a smallest possible key for R under the assumption that:

- a) R is many-many
- b) R is many-one from E_1 to E_2 .
- c) R is many-one from E_2 to E_1 .
- d) R is one-one.

Problem 2.

By now you should be familiar with Piazza, the free online discussion forum for college classes. Suppose you are hired by Piazza to design their database, which contains information on the users, schools, courses, and posts.

Each user has the following information: first name, last name, and email address. The email address is also used as the login ID, so it's unique. A user can enroll in multiple courses, but can have only one role in each of the courses that he/she enrolls in (i.e. as an Instructor, Teaching Assistant, or Student). For example, a user can enroll in a course as an Instructor, and in another course as a Student.

Each course in the database is offered by a school. A school has a school ID (unique), school name, city, state, and zip code. A course has a subject code (e.g. "CS"), course number (e.g. "450"), section number (e.g. "001"), course title (e.g. "Database Concepts"), and course status ("Active" or "Inactive"). A course can be offered in different semesters. Note there is no unique course ID available for a course (do not generate an artificial ID either), as different schools may have the same course numbers/titles.

Each user can specify his/her own email notification preference: real-time notification, daily email digest, or no email. Note a user can set different preference for each course

that he/she enrolls in.

Each course has a list of posts. Only members of a course (i.e. those who enroll in the course) can post something for the course. A post could be either an original/new post, or a follow-up post to an existing post. In addition to the author information (the poster), each post should also have: (unique) post ID, date and time the post was made, the actual post, and the post ID of the original post, if any, that the current post follows up on.

Draw the ER diagram for the given database, identifying the following:

- (i) all the entity sets and their attributes;
- (ii) all the relationship sets;
- (iii) all the required key constraints and participation constraints;
- (iv) the primary key for each entity set (and weak entity set, if any).

In some cases, you may have to make your own constraint assumptions. Clearly state, and justify your assumptions. Does your diagram capture all the given requirements? If not, specify which requirement(s) is (are) not captured. **You will be graded both on the correctness of the design, and the quality of the design.**

While there are many variations of ER diagrams, please follow the same notations used in the textbook. If you use different notations, you may get points taken off.

Problem 3.

The following questions concern the ER diagram resulting from Problem 2:

- a) Convert the entity-relationship design to a schema for a relational database. List all relational schema. For each relational schema, state
 - (i) the name of the relation,
 - (ii) the names of its attributes,
 - (iii) the domain (or data type) of each attribute,
 - (iv) the primary key, and
 - (v) the foreign key(s).

Use the following format:

relation_name(attr₁: domain₁, attr₂: domain₂, ... attr_n: domain_n)

Primary Key: attr₁

Foreign Key: X references Y

- b) Use Oracle on ITE LAB machine to create the tables from a) above and insert at least two tuples to each table. Implement all these in one script (text) file.

What/when/how to submit: Submit a hard copy AND an electronic copy of ALL problems (no loose paper please). Make sure you turn in a separate script file for item b) of Problem 3 (see below). Use PDF for the other problems. Please follow the guidelines below. Any violations might cause you to lose points.

1. Prepare your submission:

Create a script file by putting all the commands (SQL statements, no sqlplus prompts and outputs) in **ONE** file, then run it in sql*plus. Send the script file in text format. Note that text file means **only ASCII code is allowed**. If you use a word processor, make sure you save the file in ASCII format. No other formats are acceptable.

If you are using Windows, make sure to use Notepad, or save as text file.

About how to create a script file and run it in sql*plus, please refer to the Oracle user guide. Before you submit, please run your script file to make sure it is executable. **Un-executable script files will receive zero for the part of script file! No exceptions.**

Please name your script file in this way: Initial letter of your first name + your last name + ".sql". For example, if your name is John Smith, your script file should be JSmith.sql.

Again, the file you are submitting is **an executable .sql file**, no .txt file or other types.

2. Format of your .sql file:

As the first few lines of the text file, put your **name**, the **assignment number**, and your **email address** in these few lines.

Be sure to add comment symbol before and after them.

Here is an example:

```
// Student: John Smith  
// Assignment #1  
// Instructor – Dr. Jessica Lin  
// Email: jsmith@gmu.edu
```

3. Where to submit:

Please note that the only place for submission is "Assignments" module in Blackboard - click on the "Assignments" icon on the homepage of Blackboard class site, then click on the "HW1" icon and the homework submission page comes out. If you'd like, you can also add your comments. **Submission via email or Blackboard mail is not acceptable.**

Please note that you are not allowed to take back and correct after submission, so please make sure everything is correct before you click on the "Submit" button.

4. Always submit a paper copy in class for all your electronic submissions. The due

time is the same as the electronic submission and in-class submission.

The deadline for submitting both paper version and the electronic version is **September 15, 2011 by 1:35PM. The hardcopy needs to be turned in class, within five minutes of the beginning of the lecture.**

Note: Please type your solution if possible. Otherwise, **write clearly**. Always assume that the TA will **INCORRECTLY** decipher difficult handwriting.

Some Hints:

Free tools for drawing ER diagrams:

You can download Microsoft Visio Professional 2010 for free from GMU's msdn academic alliance (http://msdn05.e-academy.com/elms/Storefront/Home.aspx?campus=gmu_bsit). You need to register with the software center then download the program for free (or you can have the Visio 2010 CD mailed to you for \$19.95). If you have problem you can contact the IT&E support @703-993-1515 (more on the support page: <http://labs.ite.gmu.edu/index.php/FAQ/MSDNAA#a2>). Please be aware that the notation employed by Visio is not exactly the same as the one used in the textbook.

Also, you can download a trail version of SmartDraw (<http://www.smartdraw.com/>). The good news is the notation is similar to the textbook, and the bad news is you have to pay if you want to get a full version.

Oracle DBMS:

1. If you have troubles creating a table, try to clean up your database. Proceed as follows:

```
select table_name
from all_tables
where owner='YOURACCOUNTNAMEINUPPERCASE';
to see what tables you have in your database. Then use
drop table table_name;
to delete them in sqlplus.
```

2. If you still have troubles creating a particular table, try a different name. A name (like order, group, user, etc) may be a reserved word.

3. When creating tables that contain foreign keys, make sure the tables that are referenced have already been created.