

CPS 116 Fall 2004

Homework #1 (8.75% of course grade: 100 points)

Assigned: Tuesday, August 31

Due: Thursday, September 9

This homework should be done in pieces as soon as relevant topics are covered in lectures. If you wait until the last minute, you might be overwhelmed.

Problems 1, 4, and 6 need to be completed on online testing system Gradiance. Please read <http://www.cs.duke.edu/courses/fall04/cps116/faqs/gradiance.html> (accessible from the “Programming Notes” section of the course Web page) for instructions on how to use the system.

Problem 1 (16 points).

Complete the Gradiance homework titled “Homework 1.1 (Relational Algebra Basics).” Note that in the online exercises we use English names of relational algebra operators instead of symbols, because some fonts are unavailable. Also note that “theta-join” means a regular join with an explicit join condition (\bowtie_p).

Problem 2 (36 points).

Consider a database containing information about bars, beers, and bar-goers.

Drinker (name, address)

Bar (name, address)

Frequents (drinker, bar, times_a_week)

Likes (drinker, beer)

Serves (bar, beer, price)

Write the following queries in *relational algebra*. You may use expression trees (as seen in lecture slides) to improve readability.

- (a) Find all drinkers who frequent James Joyce Pub.
- (b) Find all bars that serve both Amstel and Corona.
- (c) Find all bars that serve at least one of the beers Amy likes for no more than \$2.50.
- (d) For each bar, find all beers served at this bar that are liked by none of the drinkers who frequent that bar.
- (e) Find all drinkers who frequent *only* those bars that serve some beers they like.
- (f) Find all drinkers who frequent *every* bar that serves some beers they like.

Optional: You can test your queries on a sample database on rack40. Follow the instructions at <http://www.cs.duke.edu/courses/fall04/cps116/faqs/login.html> to log into rack40 and set up the environment for DB2. Then, run `/home/dbcourse/examples/db-beers/setup.sh` to setup a database with some sample data. Use the command `ra` to test your relational algebra queries. For instructions on using `ra`, please refer to <http://www.cs.duke.edu/courses/fall04/cps116/faqs/ra.html>. Submit a script of running `ra`, showing all the queries and answers.

Problem 3 (14 points).

As discussed in class, the core operators in relational algebra are selection (σ_p), projection (π_L), cross product (\times), union (\cup), and difference ($-$).

- (a) Show that the projection operator is necessary; that is, some queries that use the projection operator cannot be expressed using any combination of the other operators.
- (b) Show that the selection operator is necessary; that is, some queries that use the selection operator cannot be expressed using any combination of the other operators.

Problem 4 (10 points).

Complete the Gradiance homework titled “Homework 1.4 (E/R Design).”

Problem 5 (14 points).

An online used-car trading dot-com hires you to design a database for its Web site. The database will store information about used automobiles for sale.

- Each automobile has a VIN (vehicle identification number), a model (e.g., Camero), a make (e.g., Chevrolet), a year (e.g., 1999), a color (e.g., red), a mileage (e.g., 50,000 miles), and a body style (e.g., coupe).
 - Each automobile has a seller, which may be either a dealer or an individual. For each dealer, the database stores name, address, phone number. For each individual, only phone number and email address are recorded.
 - In addition, the Web site maintains reviews about automobiles. Each review is about one particular model, make, and year. Each review has an author. The same author may write several reviews about the same model, make, and year. The database should connect each car with the reviews about the make, model, and year of that car, although it need not make a *direct* connection.
- (a) Design an E/R diagram for this database. Very briefly explain the intuitive meaning of any entity and relationship sets. Do not forget to indicate keys, multiplicity of relationships, and weak entity sets in appropriate ways.
 - (b) Design a relational schema for this database. (You can start by translating the E/R design.) You may ignore attribute types, and you do not need to show any sample data. Indicate all keys and non-trivial functional dependencies in the schema. Check if the schema is in BCNF. If not, decompose the schema into BCNF.

Problem 6 (10 points).

Complete the Gradiance homework titled “Homework 1.6 (Relational Design Theory: FD).”