CPS 196.3 Fall 2002
Homework #1
Assigned: Tuesday, September 3
Due: Thursday, September 12

Problem 1.

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) Give an E/R design for this database. Very briefly explain the intuitive meaning of any entity and relationship sets. Do not forget to indicate any keys, multiplicity of relationships, and weak entity sets in their 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 2.

In this problem you will explore the other side of the dot-com. Your job is to design a database that holds information about startups and venture capital funds that invest in these startups.

- Every venture capital (VC) fund is identified by its name and number, e.g., "Kleiner-Perkins, VII." A VC fund also has a size, which is the amount of money to be invested by the fund, and a closing date, by which the fund must invest the money.
- VC funds invest in two types of startup companies: stealth and private.
- Every startup has a name, an address, and belongs to exactly one industry. While, in general, there may be several companies with the same name, within each industry company names are unique.
- Every industry has a unique name, market size, and consists of several sectors. Sectors have names and projected growths. Sector names are unique within an industry but not across industries.

- Stealth companies operate in a "stealth mode" and have a buzz factor. A stealth company is financed by at most one VC fund since it is difficult to remain in stealth mode when several sets of lawyers get involved.
- A private company has a Web site and a CEO. Every private company belongs to exactly one industry sector. For ethical reasons, VC funds invest in at most one company in any given industry sector.
- Finally, stealth companies may be acquisition targets of private companies.

(a) Give an E/R design for this database. Very briefly explain the intuitive meaning of any entity and relationship sets. Do not forget to indicate any keys, multiplicity of relationships, and weak entity sets in their 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 3.

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

Frequents (*drinker*, *bar*)
Likes (*drinker*, *beer*)
Serves (*bar*, *beer*, *price*)

Write the following queries in *relational algebra*. You may use expression trees to improve readability.

(a) Find all drinkers who frequent James Joyce Pub.
(b) Find all bars that serve both Bud and Miller.
(c) Find all bars that serve at least one of the beers Homer likes for under $5.
(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.

Problem 4.

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.