



Written Assignment #1 CAS CS 460

Due: 09/25 in gradescope.

Problem 1: ER-Model. [25pts]

Assume that you have been asked to create a database for a medical clinic. In this database you want to store information about Doctors, Patients, Tests, and Visits. For each doctor, we need to store a doctor id, name, specialty, and medical school that got her/his degree. For a patient, we need to store name, patient id, insurance, and a doctor that is the primary physician for this patient. For a test, we store the test name, test id, date of test, and the result. For a visit, we need to store the visit id, the check in and check out days, the patient associated with the visit and the doctor, who can be any doctor, not necessarily the primary physician. We assume that the doctor ids, patient ids, visit ids, and test ids are unique. Make sure that your design allows the same patient to have many visits on different dates and to have the same test many times.

Complete the following tasks:

1. Create the E-R diagram.
2. List the primary keys, candidate keys, weak entities (if any), partial keys (if any), total participation and any key constraints.
3. Turn the E-R diagram into tables (relational model schema). Provide the SQL CREATE TABLE statement for each table in the relational schema.

Problem 2: ER-Model and Relational Model. [30pts]

Suppose that you need to design a database for an airport. The relevant information that must be stored is:

- Every airplane has a registration number, and each airplane is of specific model.
- The airport accommodates a number of airplane models, and each model

is identified by a model number (e.g. A-320, B-767) and has a capacity and a weight.

- A number of technicians work at the airport. You need to store the name, SSN, address, phone number, and salary of each technician.
 - Each technician is an expert on one or more plane model(s).
 - Traffic controllers work also at the airport. We need to store name, SSN, address, and phone number.
 - Traffic controllers must have an annual medical examination. For each traffic controller, you must store the date of the most recent exam.
 - The airport has a number of tests that are used periodically to ensure that airplanes are still airworthy. Each test has a Federal Aviation Administration (FAA) test number, a name, and a maximum possible score.
 - The FAA requires the airport to keep track of each time a given airplane is tested by a given technician using a given test. For each testing event, the information needed is the date, the number of hours the technician spent doing the test, and the score the airplane received on the test.
1. Give an E/R diagram for this database. List the primary keys, candidate keys, weak entities (if any), partial keys (if any), total participation and any key constraints.
 2. Turn the E-R diagram into tables (relational model schema). Indicate primary and candidate keys and foreign keys. Give the SQL statements to create the corresponding relations in a relational DBMS.

**Problem 3: Relational Algebra I. [25pts]**

Consider the following database schema that is part of a larger database and stores information about electronic products (that in our case can be either PC, Laptop or Printer.)

Product (maker, model, type)

PC (model, speed, ram, hdd, price)

Laptop (model, speed, ram, hdd, screen, price)

Printer (model, color, type, price)

The Product stores the maker, the model and the type that can be either PC, Laptop, or Printer. Each of the other three relations store specific information about each product. Maker and model, or model and type are unique.

Give an expression in relational algebra for each of the queries below:

1. What PC models have speed at least 3?
2. Which manufacturers make laptops with hard disks at least 100GB?
3. Find the manufacturers that make Laptops but not PCs.
4. Find the hard disk sizes that occur in two or more PCs.
5. Find the manufacturer(s) of the PC(s) with the highest available speed.

**Problem 4: Relational Algebra II. [20pts]**

Consider a database with the following schema:

Employee (SSN, name, salary, DNo)

Department (DNo, DeptName, MgrSSN)

Project (PNo, location, ProjName)

HourLog (SSN, PNo, hours)

The Employee relation provides a list of employees with their SSN, name, salary, and department number (DNo). The SSN is unique for each employee. Each employee belongs to only one department. The Department relation contains a list of the departments for the company. Its schema includes a unique department number called DNo. It also includes the name of the department (DeptName) and the social security number of the department's manager (MgrSSN). Each department has only one manager. The Project relation includes a unique project number (PNo), location and the project name (ProjName). An employee can be assigned to any number (including zero) projects. Each project has at least one person assigned to it. Finally, the HourLog relation lists for each project the number of hours of work for each employee who is assigned to that project. The key of this relation is SSN and PNo.

Write the following queries in Relational Algebra. You may use assignment of intermediate results for long queries.

1. Find the name and the SSN of everyone who works more than 100 hours on a project located in Boston.
2. Find the name and SSN of everyone who works for department number 1 and also works on project number 2.
3. Find the name and the SSN of everyone who works on at least two projects.
4. Find the name and the SSN of everyone who works on all projects.