

Submission Instructions:

Due Date: Tuesday, October 9, 2007 1:00 p.m.

- Answer both of the questions provided in the assignment.
- The assignment must be **typed**. Handwritten assignments will **not** be marked.
- If you are working in a team of two, submit only **one** assignment. (If two assignments are submitted, the assignment with the lower mark will be considered.)
- Include a cover page. A template for the cover page is available on the Assignment One web page.
- All pages must be **numbered** and **stapled** together. Do not put them in an envelope. Attach all pages well. Only those pages actually stapled to the front page will be marked – the markers cannot be responsible for pages that go astray. For this reason – don't use a paperclip, fold your pages together at the corner, or attach them by any other creative means.
- Late assignments must be emailed to your instructor (either jm@cs.toronto.edu, Professor Mylopoulos; or faye@cs.toronto.edu, Professor Baron). Acceptable electronic formats for submission are PDF and Postscript. Any other format will not be marked.
- Assignments cannot be more than three days late – assignments received after 1:00 p.m. Friday, October 12, 2007, will not be considered. A 10% penalty (five marks) will be applied for each late day.

Question 1 – The Relational Model and DDL [20 marks]

Le Challow Stores is building a database to manage purchasing, inventory control and sales. They must capture the following kinds of information:

- Le Challow Stores has company locations. Each location has a location ID, address, city, province, telephone number, and location type.
- Location types may be one of office, warehouse or store.
- There are employees. Each employees has a base location, last name, first name, job type (sales, purchasing, receiving, office), address, city, province, telephone number, social insurance number, pay type (salary which is annual rate, and hourly which is paid for each hour worked), and pay rate.
- There are clothing items that will be sold in the stores. Each item has a unique ID, a short description, long description, a clothing category (shirt, dress, sweater, scarf, etc.), intended sex of wearer (women, men), a vendor ID, cost price, sale price, and a discontinued flag to indicate the product may no longer be purchased.
- Clothing items also have sizes and colours, where applicable. They require a record for each size in which the clothing item is manufactured, and a record for each colour in which it is manufactured.
- Purchase orders are used to order new items from vendors. Each order indicates the location to which the items must be shipped, a unique purchase order number for reference, the vendor, the order date, the employee who placed the purchase. If a purchase order is cancelled, it is deleted from the system. All records dependent on the purchase order are also deleted.
- Each purchase order has order item details. For each item ordered we must specify the purchase order ID, the clothing item ID, size, colour, quantity, and status (on order, received). Note that there must be a separate order item detail for each size and colour combination ordered. If an order item is deleted, all items dependent on this item must be deleted.
- There is a distribution record for each order item detail. The purchaser predetermines the stores to which stores the ordered items will be sent, the status of the distribution (sent, received, and none) and how many of each size and colour. Distribution details are determined by the buyer at the time of purchase. A single purchase order item received at the warehouse may be distributed to ten different locations.
- Inventory is managed for each location. Inventory information is kept for each different clothing item by item ID, size, and colour. The quantity of each item received, sold, and on-hand (in the store at the time a manual inventory is taken) is maintained.
- There are Vendors. For each Vendor, we maintain information about their name, address, city, province or state, country, primary contact name, primary contact telephone number, relationship rating (e.g., reliable, usually late, excellent, closed).
- There are Sales in all of the stores. For each sale we record the ID, location, date and time, form of payment (cash, debit, credit card), credit card type and number (if applicable), and approval or authority number received when debits and credit cards are approved for the purchase amount.
- For each sale tax information is recorded (in a separate table?) which contains the sale ID, tax type, (GST, PST), the amount charged, the province code (if it is a PST record), and when the tax has been remitted to the appropriate government.
- Finally, there are Sales Items. These are the details about individual clothing items that are sold to the customer. They include the Clothing item ID, size, colour, quantity, sale price, and Sale ID.

Define a relational schema that can accommodate the above information **in DDL**. Remember: Enforce all key and integrity constraints.

- Provide default values where appropriate.
- Avoid duplicating information (except for foreign keys used to link relations.)
- Indent your DDL statements consistently for readability.
- You may make text fields any (reasonable) size you want.
- If you are using a descriptive field that occurs in more than one tuple, you may want to encode it in a separate table.

Question 2 – Relational Algebra [30 marks]

For this problem we use the schema of the Ontario Hospital Health Network database. It includes the following relations. Keys are underlined and each tuple is described:

- Hospital(hospitalID, hName, hAddress, hCity) A hospital including the hospital ID, name, address and city of that hospital.
- Room(roomID, hospitalID, rNumber, rSize) A hospital room including the hospital ID, room number, and room capacity – the number of patients it accommodates.
- Employee(employeeID, eFirstName, eLastName, eRole, eSalary, eAddress, eCity, eSIN) An employee including their employee ID number, first and last names, role in the hospital (or occupation type, such as doctor, technician, etc.), salary, address, city and Social Insurance Number.
- Department(departmentID, hospitalID, dName, dHead) a department in a hospital including the hospital ID, the department ID, the name of the department, and the employee ID of the head of the department.
- EmployeeDepartment(employeeID, departmentID, edStartDate, edEndDate) the departments to which an employee belongs. An employee may work in more than one department. This tuple contains the employee ID, the department ID, and the dates that they started and finished employment in that department.
- Patient(patientID, pFirstName, pLastName, pSex, pOHIP, pOther, pDOB, pAddress, pCity, pProvince, pCountry) A patient in the hospital system including the patient ID, first name, last name, sex (M or F), OHIP number, other form of payment if not on OHIP, date of birth, residence address including the address, city, province, and country.
- PatientAppointment(patientID, departmentID, employeeID, paDate, paTime, paReason) An appointment for a patient, including the department ID of the hospital department, the employee ID of the person whom they are seeing (could be technician), the date, the time, and the purpose of the visit.
- PatientRoomAssignment(patientID, roomID, prInDate, prOutDate, prDoctor) A room that is assigned to a patient who needs to stay in the hospital - including the ID of the patient and room, the date that they were admitted, the date that they left the hospital, and the employee ID of their attending physician during their stay.

The following inclusion dependencies (foreign keys) hold:

Room(hospitalID) \subseteq Hospital(hospitalID)

Department(hospitalID) \subseteq Hospital(hospitalID)

Department(dHead) \subseteq Employee(employeeID)

EmployeeDepartment(departmentID) \subseteq Department(departmentID)

PatientAppointment(patientID) \subseteq Patient(patientID)

PatientAppointment(departmentID) \subseteq Department(departmentID)

PatientAppointment(employeeID) \subseteq Employee(employeeID)

PatientRoomAssignment (patientID) \subseteq Patient(patientID)

PatientRoomAssignment (roomID) \subseteq Room(roomID)

Write the following queries in Relational Algebra. You are not allowed to use the renaming or relations as described in the text though you may create interim tables as shown in tutorial, nor can you use the extended relational algebra. You may use renaming of attributes, as well as numerical and date comparisons (e.g., $R.A > 5$, or $R.A \geq YYYY-MM-DD$). You may specify a date as YYYY-MM-DD and time as HH:MM (using a 24 hour clock). Also assume that you can subtract one date from another to get the number of days. So if a first date minus a second date equals 5, then the first date is five days after the second date. Use CurrentDate to represent today's date and be sure to clearly state any assumptions that you make.

1. Find the first name, last name, hospital, department name, date, and time of all patients who had an appointment with Dr. Joe Smith. (assume there is only one Dr. Joe Smith)
2. Find the hospital name and department name for all departments whose head is also head of another department.
3. Find the first name and last name of all Doctors whose patients are never in the hospital for less than 7 days. (They must have had patients admitted in the hospital at some time.)
4. Find the first and last name of all patients who have exactly two doctors in the hospital system. (Note that doctors are specified in both appointments and room assignment.)
5. Find the hospital and department name for the department at which there were no appointments from June 3, 2007 to June 8, 2007.
6. Find the hospital, room number, and patient first and last names of all semi-private (size = 2) rooms which were occupied by both a male and female patient both of whom were over 50 on August 13, 2007.