CMPT 354 — Database Systems & Structures: Solution Assignment No. 2

Total scores: 40 Spring 1998
Due: June 18 Instructor: Osmar Zaiane

Question 1: (16 points) Consider the following relational database, answer each of the queries in SQL and QBE:

person(pname, street, city)
works_for(pname, cname, salary)
company(cname, city)
manages(pname, mname)

1. Find the street and city of all employees who work for the SuperStore, live in Coquitlam, and earn more than $35,000. (2 marks each)

   SQL:

```
SELECT P.street, P.city
FROM person P, works_for W
WHERE P.city = "Coquitlam" AND
      W.cname="SuperStore" AND
      W.salary>35000 AND
      W.pname=P.pname
```

   QBE:

person	*pname*	*street*	*city*
	_x	P.	P. Coquitlam

works_for	*pname*	*cname*	*salary*
	_x	SuperStore	>35000

2. Find the names and the companies they work for, for all people who have a higher salary than their manager. (2 marks each)

   SQL:

```
SELECT W1.pname, W1.cname
FROM works_for W1, works_for W2, manages M
WHERE W1.pname = M.pname AND
      W2.pname = M.mname AND
      W1.salary > W2.salary
```

QBE:

manages	*pname*	*mname*
	_p	_m

worksfor	*pname*	*cname*	*salary*
	P._p	P.	_s1
	_m		_s2

Conditions
_s1 > _s2

3. Assume that companies may be located in several cities. Find all companies located in every city in which the SuperStore is located. (2 marks each)

   SQL:

```
SELECT C.cname
FROM company C
WHERE not exists ((SELECT city
                   FROM company
                   WHERE cname = "SuperStore")
                  except
                   (SELECT city
                   FROM company N
                   WHERE C.cname = N.cname))
```

   QBE:

company	*cname*	*city*
	SuperStore	_x
	SuperStore	_y
	P.	_y

Conditions
CNT.UNQ.ALL._x = CNT.UNQ.ALL._y

4. Find the names of managers who manage more than 5 employees living in Richmond. (2 marks each)

   SQL:

```
SELECT M.mname
FROM manages M, person P
WHERE P.city = "Richmond" AND
        P.pname = M.pname AND
GROUP BY M.mname
HAVING COUNT M.pname >5
```

QBE:

person	*pname*	*street*	*city*
	_p		Richmond

manages	*pname*	*mname*
	_p	P.G._m

Conditions
CNT.ALL._p > 5

Question 2: (12 points) Give expressions in QBE, Quel, and Datalog equivalent to each of the following two queries:

1. $\Pi_{A,F}(\sigma_{C=D}(r \times s)), where R = (A, B, C) and S = (D, E, F)$.

   SQL: (not required)

```
SELECT DISTINCT A,F
FROM r, s
WHERE C=D
```

   QBE:

r	*A*	*B*	*C*
	_a		_c

s	*D*	*E*	*F*
	_c		_f

result	*A*	*F*
	_a	_f

   Quel:

```
RANGE OF a is r
RANGE OF b is s
RETRIEVE (a.A, b.F)
WHERE a.C = b.D
```

   Datalog:

```
query(X,Y) :- r(X,V,W),s(W,Z,Y)
```

2. $\{< a, b, c > | (< a, b > \in r \wedge < a, c > \in s\}, where R = (A, B) and S = (A, C)$

   SQL: (not required)

```
SELECT DISTINCT r.A, r.B, s.C
FROM r, s
WHERE r.A=s.A
```

   QBE:

r	A	B
	_a	_b

s	A	C
	_a	_c

result	A	B	C
P.	_a	_b	_c

   Quel:
   RANGE OF a is r
   RANGE OF b is s
   RETRIEVE (a.A, a.B, b.C)
   WHERE a.A = b.A

   Datalog:
   query(X,Y,Z) :- r(X,Y),s(X,Z)

Question 3: (6 points) Consider rhe relational database of question 1. Using SQL, define a view consisting of the manager name and the average salary of all employees who work for that manager. Explain with an example why the database system should not allow updates to be expressed in terms of this view.

```
CREATE VIEW salaryavg AS
    SELECT mname, AVG(salary)
    FROM manages M, works_for W
    WHERE M.pname = W.pname
    GROUP BY mname
```

Updates should not be allowed in this view because there is no way to determine how to change the underlying data. For example, suppose the request is "change the average salary of employees working for Smith to $200". Should everybody who works for Smith have their salary changed to $200? or shoud the first (or mor if necessary) employee found who works for Smith have their salary adjusted so that the average is $200? Neither approach really makes sense.

Question 4: (6 points) Give an SQL schema definition for the relational database of question1. Choose an appropriate domain for each attribute and an appropriate primary key for each relation schema.

```
CREATE DOMAIN companyname CHAR(20)
CREATE DOMAIN cityname    CHAR(30)
CREATE DOMAIN personname  CHAR(30)

CREATE TABLE person
(pname personname,
street CHAR(30),
city cityname,
PRIMARY KEY (pname))

CREATE TABLE works_for
(pname personname,
cname companyname,
salary NUMERIC(8,2),
PRIMARY KEY (pname))

CREATE TABLE company
(cname companyname,
city cityname,
PRIMARY KEY (cname))

CREATE TABLE manages
(pname personname,
mname personname,
PRIMARY KEY (pname))
```
