5

SQL: QUERIES, CONSTRAINTS, TRIGGERS

Online material is available for all exercises in this chapter on the book's webpage at

`http://www.cs.wisc.edu/~dbbook`

This includes scripts to create tables for each exercise for use with Oracle, IBM DB2, Microsoft SQL Server, Microsoft Access and MySQL.

Exercise 5.1 Consider the following relations:

Student(*snum:* `integer`, *sname:* `string`, *major:* `string`, *level:* `string`, *age:* `integer`)
Class(*name:* `string`, *meets_at:* `string`, *room:* `string`, *fid:* `integer`)
Enrolled(*snum:* `integer`, *cname:* `string`)
Faculty(*fid:* `integer`, *fname:* `string`, *deptid:* `integer`)

The meaning of these relations is straightforward; for example, Enrolled has one record per student-class pair such that the student is enrolled in the class.

Write the following queries in SQL. No duplicates should be printed in any of the answers.

1. Find the names of all Juniors (level = JR) who are enrolled in a class taught by I. Teach.
2. Find the age of the oldest student who is either a History major or enrolled in a course taught by I. Teach.
3. Find the names of all classes that either meet in room R128 or have five or more students enrolled.
4. Find the names of all students who are enrolled in two classes that meet at the same time.

5. Find the names of faculty members who teach in every room in which some class is taught.
6. Find the names of faculty members for whom the combined enrollment of the courses that they teach is less than five.
7. For each level, print the level and the average age of students for that level.
8. For all levels except JR, print the level and the average age of students for that level.
9. For each faculty member that has taught classes only in room R128, print the faculty member's name and the total number of classes she or he has taught.
10. Find the names of students enrolled in the maximum number of classes.
11. Find the names of students not enrolled in any class.
12. For each age value that appears in Students, find the level value that appears most often. For example, if there are more FR level students aged 18 than SR, JR, or SO students aged 18, you should print the pair (18, FR).

Answer 5.1 The answers are given below:

1.
```
SELECT DISTINCT S.Sname
FROM   Student S, Class C, Enrolled E, Faculty F
WHERE  S.snum = E.snum AND E.cname = C.name AND C.fid = F.fid AND
       F.fname = 'I.Teach' AND S.level = 'JR'
```

2.
```
SELECT MAX(S.age)
FROM   Student S
WHERE  (S.major = 'History')
       OR S.snum IN (SELECT E.snum
                     FROM   Class C, Enrolled E, Faculty F
                     WHERE  E.cname = C.name AND C.fid = F.fid
                            AND F.fname = 'I.Teach' )
```

3.
```
SELECT   C.name
FROM     Class C
WHERE    C.room = 'R128'
         OR C.name IN (SELECT   E.cname
                       FROM     Enrolled E
                       GROUP BY E.cname
                       HAVING   COUNT (*) >= 5)
```

4.
```
SELECT DISTINCT S.sname
FROM    Student S
WHERE   S.snum IN (SELECT E1.snum
                   FROM    Enrolled E1, Enrolled E2, Class C1, Class C2
                   WHERE   E1.snum = E2.snum AND E1.cname <> E2.cname
                   AND E1.cname = C1.name
                   AND E2.cname = C2.name AND C1.meets_at = C2.meets_at)
```

5.
```
SELECT DISTINCT F.fname
FROM    Faculty F
WHERE   NOT EXISTS (( SELECT *
                      FROM    Class C )
                      EXCEPT
                      (SELECT C1.room
                      FROM    Class C1
                      WHERE   C1.fid = F.fid ))
```

6.
```
SELECT   DISTINCT F.fname
FROM     Faculty F
WHERE    5 > (SELECT COUNT (E.snum)
              FROM    Class C, Enrolled E
              WHERE   C.name = E.cname
              AND     C.fid = F.fid)
```

7.
```
SELECT   S.level, AVG(S.age)
FROM     Student S
GROUP BY S.level
```

8.
```
SELECT   S.level, AVG(S.age)
FROM     Student S
WHERE    S.level <> 'JR'
GROUP BY S.level
```

9.
```
SELECT   F.fname, COUNT(*) AS CourseCount
FROM     Faculty F, Class C
WHERE    F.fid = C.fid
GROUP BY F.fid, F.fname
HAVING   EVERY ( C.room = 'R128' )
```

10.
```
SELECT   DISTINCT S.sname
FROM     Student S
WHERE    S.snum IN (SELECT   E.snum
                    FROM     Enrolled E
                    GROUP BY E.snum
```

```
                                HAVING   COUNT (*) >= ALL (SELECT   COUNT (*)
                                                           FROM     Enrolled E2
                                                           GROUP BY E2.snum ))
```

11.
```
SELECT DISTINCT S.sname
FROM    Student S
WHERE   S.snum NOT IN (SELECT E.snum
                       FROM   Enrolled E )
```

12.
```
SELECT    S.age, S.level
FROM      Student S
GROUP BY  S.age, S.level,
HAVING    S.level IN (SELECT   S1.level
                      FROM     Student S1
                      WHERE    S1.age = S.age
                      GROUP BY S1.level, S1.age
                      HAVING   COUNT (*) >= ALL (SELECT   COUNT (*)
                                                 FROM     Student S2
                                                 WHERE s1.age = S2.age
                                                 GROUP BY S2.level, S2.age))
```

Exercise 5.2 Consider the following schema:

Suppliers(*sid:* `integer`, *sname:* `string`, *address:* `string`)
Parts(*pid:* `integer`, *pname:* `string`, *color:* `string`)
Catalog(*sid:* `integer`, *pid:* `integer`, *cost:* `real`)

The Catalog relation lists the prices charged for parts by Suppliers. Write the following queries in SQL:

1. Find the *pname*s of parts for which there is some supplier.
2. Find the *sname*s of suppliers who supply every part.
3. Find the *sname*s of suppliers who supply every red part.
4. Find the *pname*s of parts supplied by Acme Widget Suppliers and no one else.
5. Find the *sid*s of suppliers who charge more for some part than the average cost of that part (averaged over all the suppliers who supply that part).
6. For each part, find the *sname* of the supplier who charges the most for that part.
7. Find the *sid*s of suppliers who supply only red parts.
8. Find the *sid*s of suppliers who supply a red part and a green part.

9. Find the *sid*s of suppliers who supply a red part or a green part.
10. For every supplier that only supplies green parts, print the name of the supplier and the total number of parts that she supplies.
11. For every supplier that supplies a green part and a red part, print the name and price of the most expensive part that she supplies.

Answer 5.2 The answers are given below:

1.
```
SELECT DISTINCT P.pname
FROM   Parts P, Catalog C
WHERE  P.pid = C.pid
```

2.
```
SELECT S.sname
FROM   Suppliers S
WHERE  NOT EXISTS (( SELECT P.pid
                     FROM   Parts P )
                   EXCEPT
                   ( SELECT C.pid
                     FROM   Catalog C
                     WHERE  C.sid = S.sid ))
```

3.
```
SELECT S.sname
FROM   Suppliers S
WHERE  NOT EXISTS (( SELECT P.pid
                     FROM   Parts P
                     WHERE  P.color = 'Red' )
                   EXCEPT
                   ( SELECT C.pid
                     FROM   Catalog C, Parts P
                     WHERE  C.sid = S.sid AND
                            C.pid = P.pid AND P.color = 'Red' ))
```

4.
```
SELECT P.pname
FROM   Parts P, Catalog C, Suppliers S
WHERE  P.pid = C.pid AND C.sid = S.sid
AND    S.sname = 'Acme Widget Suppliers'
AND    NOT EXISTS ( SELECT *
                    FROM   Catalog C1, Suppliers S1
                    WHERE  P.pid = C1.pid AND C1.sid = S1.sid AND
                           S1.sname <> 'Acme Widget Suppliers' )
```

5.
```
SELECT DISTINCT C.sid
FROM   Catalog C
```

```
WHERE  C.cost > ( SELECT AVG (C1.cost)
                  FROM   Catalog C1
                  WHERE  C1.pid = C.pid )
```

6.
```
SELECT P.pid, S.sname
FROM   Parts P, Suppliers S, Catalog C
WHERE  C.pid = P.pid
AND    C.sid = S.sid
AND    C.cost = (SELECT MAX (C1.cost)
                 FROM   Catalog C1
                 WHERE  C1.pid = P.pid)
```

7.
```
SELECT DISTINCT C.sid
FROM   Catalog C
WHERE  NOT EXISTS ( SELECT *
                    FROM   Parts P
                    WHERE  P.pid = C.pid AND P.color <> 'Red' )
```

8.
```
SELECT DISTINCT C.sid
FROM   Catalog C, Parts P
WHERE  C.pid = P.pid AND P.color = 'Red'
INTERSECT
SELECT DISTINCT C1.sid
FROM   Catalog C1, Parts P1
WHERE  C1.pid = P1.pid AND P1.color = 'Green'
```

9.
```
SELECT DISTINCT C.sid
FROM   Catalog C, Parts P
WHERE  C.pid = P.pid AND P.color = 'Red'
UNION
SELECT DISTINCT C1.sid
FROM   Catalog C1, Parts P1
WHERE  C1.pid = P1.pid AND P1.color = 'Green'
```

10.
```
SELECT    S.sname, COUNT(*) as PartCount
FROM      Suppliers S, Parts P, Catalog C
WHERE     P.pid = C.pid AND C.sid = S.sid
GROUP BY S.sname, S.sid
HAVING    EVERY (P.color='Green')
```

11.
```
SELECT    S.sname, MAX(C.cost) as MaxCost
FROM      Suppliers S, Parts P, Catalog C
WHERE     P.pid = C.pid AND C.sid = S.sid
```

```
GROUP BY S.sname, S.sid
HAVING   ANY ( P.color='green' ) AND ANY ( P.color = 'red' )
```

Exercise 5.3 The following relations keep track of airline flight information:

Flights(*flno:* `integer`, *from:* `string`, *to:* `string`, *distance:* `integer`, *departs:* `time`, *arrives:* `time`, *price:* `real`)
Aircraft(*aid:* `integer`, *aname:* `string`, *cruisingrange:* `integer`)
Certified(*eid:* `integer`, *aid:* `integer`)
Employees(*eid:* `integer`, *ename:* `string`, *salary:* `integer`)

Note that the Employees relation describes pilots and other kinds of employees as well; every pilot is certified for some aircraft, and only pilots are certified to fly. Write each of the following queries in SQL. (*Additional queries using the same schema are listed in the exercises for Chapter 4.*)

1. Find the names of aircraft such that all pilots certified to operate them have salaries more than $80,000.
2. For each pilot who is certified for more than three aircraft, find the *eid* and the maximum *cruisingrange* of the aircraft for which she or he is certified.
3. Find the names of pilots whose *salary* is less than the price of the cheapest route from Los Angeles to Honolulu.
4. For all aircraft with *cruisingrange* over 1000 miles, find the name of the aircraft and the average salary of all pilots certified for this aircraft.
5. Find the names of pilots certified for some Boeing aircraft.
6. Find the *aid*s of all aircraft that can be used on routes from Los Angeles to Chicago.
7. Identify the routes that can be piloted by every pilot who makes more than $100,000.
8. Print the *ename*s of pilots who can operate planes with *cruisingrange* greater than 3000 miles but are not certified on any Boeing aircraft.
9. A customer wants to travel from Madison to New York with no more than two changes of flight. List the choice of departure times from Madison if the customer wants to arrive in New York by 6 p.m.
10. Compute the difference between the average salary of a pilot and the average salary of all employees (including pilots).