

CSE 344 Section 2 Worksheet

1. Joins Examples

Given tables created with these commands:

```
CREATE TABLE A (a int);
CREATE TABLE B (b int);
INSERT INTO A VALUES (1), (2), (3), (4);
INSERT INTO B VALUES (3), (4), (5), (6);
```

What's the output for each of the following:

```
SELECT *
  FROM A INNER JOIN B
    ON A.a=B.b;
```

```
a|b
3|3
4|4
```

```
SELECT *
  FROM A RIGHT OUTER JOIN B
    ON A.a=B.b;
```

```
a|b
3|3
4|4
  |5
  |6
```

```
SELECT *
  FROM A LEFT OUTER JOIN B
    ON A.a=B.b;
```

```
a|b
1|
2|
3|3
4|4
```

```
SELECT *
  FROM A FULL OUTER JOIN B
    ON A.a=B.b;
```

```
a|b
1|
2|
3|3
4|4
  |5
  |6
```

```
SELECT * FROM A INNER JOIN B; (Challenging question!)
```

```
a|b
1|3
1|4
1|5
1|6
2|3
2|4
2|5
2|6
3|3
3|4
3|5
3|6
4|3
4|4
4|5
```

```
4|6
```

Sidenote:

sqlite3 supports neither RIGHT OUTER nor FULL OUTER.

Right outer can be implemented with

```
SELECT *
  FROM B LEFT OUTER JOIN A
    ON A.a=B.b;
```

Full outer can be implemented with

```
SELECT *
  FROM A LEFT OUTER JOIN B
    ON A.a=B.b
 UNION
SELECT *
```

```
FROM B LEFT OUTER JOIN A ON A.a=B.b;
```

2. Self Join

Consider the following over simplified Employee table:

```
CREATE TABLE Employees (id int, bossId int);
```

Suppose all employees have an id which is not null. How would we find all distinct pairs of employees with the same boss?

```
SELECT E1.id, E2.id
FROM Employees AS E1, Employees AS E2
WHERE E1.id < E2.id
AND E1.bossId = E2.bossId;
```

Sidenote: The predicate "E1.id < E2.id" could also be written as "E1.id > E2.id". We cannot use plain inequality "E1.id <> E2.id" as the predicate condition because this would lead to duplicate pairs.

3. SQL Aggregates Practice

```
CREATE TABLE Movies (id int PRIMARY KEY, name varchar(30),
                     budget int, gross int, rating int, year int);
CREATE TABLE Actors (id int PRIMARY KEY, name varchar(30), age int);
CREATE TABLE ActsIn (mid int REFERENCES Movies(id),
                     aid int REFERENCES Actors(id));
```

What is the minimum age of an actor who has appeared in a movie where the gross of the movie has been over \$1,000,000,000?

```
SELECT MIN(A.age)
FROM Movies as M, ActsIn as AI, Actors as A
WHERE M.id = AI.mid AND A.id = AI.aid
AND M.gross > 1000000000;
```

What is the name and budget of each movie released in 2017 whose oldest actor is less than 30?

```
SELECT M.name, M.budget
FROM Movies as M, ActsIn as AI, Actors as A
WHERE M.id = AI.mid AND A.id = AI.aid AND M.year = 2017
GROUP BY M.id, M.name, M.budget
HAVING MAX(A.age) < 30;
```

More Challenging

```
CREATE TABLE Class (  
  dept VARCHAR(50),  
  number INT,  
  title VARCHAR(50),  
  PRIMARY KEY (dept, number));
```

```
CREATE TABLE Instructor (  
  username VARCHAR(50) PRIMARY KEY,  
  fname VARCHAR(50),  
  lname VARCHAR(50),  
  started_on CHAR(10));
```

```
CREATE TABLE Teaches (  
  username VARCHAR(50), -- alternately, can use the REFERENCES syntax here  
  dept VARCHAR(50),  
  number INT,  
  PRIMARY KEY (username, dept, number),  
  FOREIGN KEY username REFERENCES Instructor,  
  FOREIGN KEY (dept, number) REFERENCES Class);
```

1. For each instructor, return their username and the number of classes they teach.

```
SELECT I.username, COUNT(T.number) AS count  
  FROM Instructor AS I LEFT OUTER JOIN Teaches AS T  
    ON I.username = T.username  
 GROUP BY I.username;
```

2. Return the first and last name of the newest instructor(s) (who started on the latest date). Assume Instructor.started_on uses yyyy-mm-dd format. If there are multiple instructors, list all of them.

- Example of the witnessing problem that doesn't require subqueries.

```
SELECT I.fname, I.lname  
  FROM Instructor AS I, Instructor AS I2  
 GROUP BY I.username, I.fname, I.lname, I.started_on  
HAVING I.started_on = MAX(I2.started_on);
```

Or using a subquery:

```
SELECT fname, lname  
  FROM Instructor  
 WHERE started_on = (  
   SELECT MAX(started_on)  
   FROM Instructor  
 );
```

3. How many classes are currently being taught by at least one instructor?

By the nature of our data, we know that any class that appears in Teaches must be taught by at least 1 teacher. Thus, if we categorize the tuples in Teaches by dept and coursenum (the primary key), we can get our answer by counting the number of groups. The sticking point of this query is how to count the number of groups. The easy solution is to wrap the grouping query in a count(*) query.

```
SELECT COUNT(*)
FROM (SELECT DISTINCT dept, coursenum
      FROM Teaches);
```

4. Return the first name and last name of the instructors who teach the most number of classes. If there are multiple instructors, list all of them.

This is another example of the witnessing problem. There are multiple approaches - see below.

SOLUTION 1:

```
WITH (
  SELECT username, COUNT(*) AS count
  FROM Teaches
  GROUP BY username
) AS ClassCounts
SELECT I.fname, I.lname
FROM ClassCounts AS C1, ClassCounts AS C2, Instructor AS I
WHERE C1.username = I.username
GROUP BY I.username, I.fname, I.lname, C1.count
HAVING C1.count = MAX(C2.count);
```

SOLUTION 2:

```
WITH (
  SELECT username, COUNT(*) AS count
  FROM Teaches
  GROUP BY username
) AS ClassCounts,
(
  SELECT MAX(count) AS max
  FROM ClassCounts
) AS MaxCounts
SELECT I.fname, I.lname
FROM ClassCounts AS C, MaxCounts AS M, Instructor AS I
WHERE C.username = I.username AND C.count = M.max;
```

SOLUTION 3:

```
SELECT I.fname, I.lname
  FROM Instructor AS I, Teaches AS T
 WHERE I.username = T.username
  GROUP BY I.username, I.fname, I.lname
HAVING COUNT(*) >= ALL (
                        SELECT COUNT(*)
                          FROM Teaches
                         GROUP BY username
                        );
```