

# **CSE344 Section 3 Worksheet**

## **1. SQL and Physical Tuning (Midterm Autumn 2011)**

You are in charge of managing the program committee for an important conference. The following database stores information about papers submitted to the conference (table Paper), reviewers on the program committee (table Reviewer), and the assignment of reviewers to papers (table Reviews). Each reviewer on the program committee will have to review a set of papers. Each paper will be reviewed by some subset of reviewers.

**Paper(pid, title)**

**Reviewer(rid, name)**

**Reviews(rid, pid)**

- pid is a unique paper identifier and the primary key of the Paper table.
- rid is a unique reviewer identifier and the primary key of the Reviewer table.
- Reviews.rid is a foreign key that references Reviewer.rid.
- Reviews.pid is a foreign key that references Paper.pid.
- A reviewer is assigned zero or more papers.
- A paper is assigned zero or more reviewers.

(a) Write a SQL query that has all papers with fewer than three reviewers assigned to them. The output of the query should be a list of paper titles. The result should include papers without any reviewers assigned to them.

Answer (write a SQL query):

**Paper(pid, title)**

**Reviewer(rid, name)**

**Reviews(rid, pid)**

(b) (15 points) Write a SQL query that has the reviewers with the most papers assigned to them. There can be more than one such reviewer. The output of the query should be a list of reviewer names. A reviewer should be listed if no other reviewer has strictly more papers to review.

Answer (write a SQL query):

(c) (10 points) Suggest 2 indexes that would speed-up your queries from the previous questions. Explain why you are selecting these indexes.

(solution available at

<http://www.cs.washington.edu/education/courses/cse344/13wi/exams/midterm-sol-11au.pdf>)