Lab 2: Relational Algebra

You may work with one other person on this lab. To submit your assignment, place a PDF in your `~/cs44/labs/2/` *directory and use* `handin44` *to electronically submit the lab. Be sure both names are on the document. Your assignment should be submitted by 11:59pm on* **Friday, February 20, 2015**.

1. True or False: if a query language is *relationally complete*, we are able express every desired query in that language. Explain your answer.

2. Assume you have two relations R and S, where R contains N tuples and S contains M tuples, and $N > M$ (i.e., R has more rows). For each expression below, give the (1) *minimum* and (2) *maximum* number of tuples possible for the resulting relation. Additionally, describe if there are any (3) requirements for the schemas to ensure the expression is legal:

   (a) $R \cup S$
   (b) $R \cap S$
   (c) $R - S$
   (d) R/S
   (e) $\sigma_{x=10}R$
   (f) $R \times S$

3. You are given the following schema:

   Suppliers(sid:`integer`, sname: `string`, address: `string`)

   Parts(pid:`integer`, pname: `string`, color: `string`)

   Catalog(sid:`integer`, pid: `integer`, cost: `real`)

   Underlined fields form the primary key for the relation. Write each of the following queries as a relational algebra expression:

   (a) Find the names of suppliers who supply some red part.
   (b) Find the sids of suppliers who supply some red part or are at 500 College Avenue.
   (c) Find the sids of suppliers who supply some red part and some green part.
   (d) Find the sids of suppliers who supply every part.
   (e) Find the sids of suppliers who supply every red part.
   (f) Finds sids of suppliers who supply every red part or supply every green part.
   (g) Find the pids of parts supplied by at least two different suppliers.
   (h) Find the pids of the most expensive parts supplied by suppliers named "Parts R' Us"

Note: one benefit to the renaming operator is that it makes a copy of a relation. So, you could use it to store an intermediate result if you want to break up an expression into pieces. For example, to simplify $\pi_{a,b,c}((R \bowtie S) \cup (X \bowtie Y))$ we could produce:

$\rho(Temp, R \bowtie S)$
$\rho(Temp2, X \bowtie Y)$
$\pi_{a,b,c}(Temp \cup Temp2)$

4. Using the same schema as above, state the query that the following expressions compute. If the query is illegal, please state why:

   (a) $\pi_{sname}((\sigma_{color=`red'}Parts) \bowtie (\sigma_{cost<100}Catalog) \bowtie Suppliers)$

   (b) $\pi_{sname}(\pi_{sid}((\sigma_{color=`red'}Parts) \bowtie (\sigma_{cost<100}Catalog) \bowtie Suppliers))$

   (c) $(\pi_{sname}((\sigma_{color=`red'}Parts) \bowtie (\sigma_{cost<100}Catalog) \bowtie Suppliers)) \cap (\pi_{sname}((\sigma_{color=`green'}Parts) \bowtie (\sigma_{cost<100}Catalog) \bowtie Suppliers))$