Problem 1 . For each of the following statements, indicate whether they are true (**T**) or false (**F**).

(a)______ Projection in SQL removes duplicates.

(b)______ There is a procedure that tests if any two relational algebra queries are equivalent.

(c)______ $\pi_A(R - S) = \pi_A(R) - \pi_A(S)$ for any relations R and S with attribute A.

(d)______ Assertions in SQL follow the *event-condition-action* scheme.

(e)______ Every schema that is in BCNF is also in 3NF.

(f)______ The size of $R \bowtie S$ is always strictly smaller than the size of $R \times S$.

(g)______ Every relational schema admits a lossless dependency preserving BCNF decomposition.

(h)______ For $R(A, B)$ and $S(B, C)$, $\sigma_{B=1}(R \bowtie S) = \sigma_{B=1}(R) \times \sigma_{B=1}(S)$.

(i)______ Every schema that is in BCNF is also in 1NF.

(j)______ Intersection can be expressed in the $\{\sigma, \pi, \times\}$ fragment of relational algebra.

Problem 2 . For each of the following statements, indicate whether they are true (**T**) or false (**F**).

(a)______ SQL cannot be invoked from C.

(b)______ Cursors can be declared for both tables and results of queries.

(c)______ There can be an arbitrary number of `primary key` declarations in one `create table` statement.

(d)______ There can be an arbitrary number of `unique` declarations in one `create table` statement.

(e)______ Atomicity means that either all operations of a transaction are properly reflected in a database, or none are.

(f)______ The value of any arithemtic operation involving a null value (e.g., `'5-Null'`) is null.

(g)______ The complexity of the sort-merge join is $O(n \log n)$.

(h)______ If a query Q evaluates to the empty table, then the condition $0 >$ `ALL` (Q) is true.

(i)______ Pushing selections often leads to a faster query execution.

(j)______ There are three different types of null values in SQL.

Problem 3 . For problems 3–7, we use the following relational schema:

- `Product(maker, model, category)`
- `PC(model, speed, ram, hd, price)`
- `Laptop(model, speed, ram, hd, screen, price)`
- `Printer(model, color, type, price)`

The constraints are as follows:

- `model` is the key for all relations.
- The only possible values of `category` are “PC”, “laptop”, and “printer”.
- The only possible values of `type` are “laser” and “ink-jet”.
- Every PC, laptop or printer must be referenced in relation `Product`.

part a
Write `CREATE TABLE` statements for this schema.

part b Write an assertion statement for the following constraint: *Products of different types cannot have the same model number.*

part c Write an assertion statement for the following constraint: *For every PC, there is a laser printer, and a color ink-jet printer, by the same maker.*

Problem 4 .
Write the following query in SQL.
Find manufacturers from whom a combination (PC, laptop, laser printer) can be bought for less than \$2,000.
You may not assume that the assertions from Problem 3 are valid.

Problem 5
Write the following query in both SQL and Relational Algebra (for the latter, you may use $\pi, \sigma, \bowtie, \times, \cup, -$ and renaming of attributes. Renaming of relations, and the linear notation, are **not** allowed). As always, you will be graded on both correctness and simplicity of your queries so do not include unnecessary operations.
Find PCs that are faster and have more ram than all the laptops by the same manufacturer.
Again, you should not assume that the assertions of Problem 3 are valid.
SQL

Relational algebra

Problem 6 . Write a query that finds manufacturers whose laptops ought to be avoided. That is, for every laptop they make, there is another one, by a different manufacturer, that has a faster processor, more ram, a larger screen, and costs less.
Write this query in both SQL and relational algebra (using only operations listed in Problem 5). Again, you should not assume that the assertions of Problem 3 are valid.

SQL

Relational algebra

Problem 7
Write the following query in SQL:
For each manufacturer, find the minimum, maximum, and average price of a (laptop, color ink-jet printer) combination.

Problem 8
Part a Consider a schema with attributes A, B, C and FDs $A \to B, BC \to A$. Does it admit a dependency-preserving BCNF decomposition?

Part a Produce a lossless BCNF decomposition of the following schema: $(ABCDE, A \to BC, B \to CE, E \to D)$.