Name:

CSci 126 Midterm 2 Fall 2002
Show your work
20 points (10 problems)

Part I

The following Company relational schema applies to problems 1 and 2. An employee can work in more than one department. An employee who is a manager has his/her eid as managerid. Every department is guaranteed to have a manager. If an employee record is deleted, the tuple that is related to this employee in the Work table has to be deleted as well.

Emp(eid: integer, ename: string, age: integer, salary: real)
Works(eid: integer, did: integer, pct_time: integer)
Dept(did: integer, dname: string, budget: real, managerid: integer)

No. 1 SQL DDL (1.5 points)
Define the Dept relation in SQL.

No. 2 SQL DDL (2 points)
Define the Works relation in SQL.

Part II

The following database applies to problem 3.
Suppliers(sid: integer, sname: string, address: string)
Parts(pid: integer, pname: string, color: string)
Catalog(sid: integer, pid: integer, cost: real)

No. 3 Write the following query in relational algebra, QBE and SQL: (4.5 points)
Find the names of suppliers who supply **some** green part.

Part III

The following database applies to problem 4.
Emp(eid: integer, ename: string, age: integer, salary: real)
Works(eid: integer, did: integer, pct_time: integer)
Dept(did: integer, budget: real, managerid: integer)

No. 4 Write a SQL query to find the managerids of managers who manage only departments with budgets greater than $2 million. (1.5 points)

Part IV

The following database applies to problems 5 & 6.
Sailors(sid: integer, sname: string, rating: integer, age: real)
Boats(bid: integer, bname: string, color: string)
Reserves(sid: integer, bid: integer, day: dates, rname: string)

No. 5 Use SQL **aggregate** operators to answer the following query: (2 points)
Find the average age of the sailors for each rating level that has at least two such sailors.

No. 6 Write the following query in relational algebra (1.5 pts) and SQL (2 pts):
Find the names of sailors who have reserved **all** boats.

Part V Query Processing. Nest Loop Join. R and S are relations.

No. 7 State the **block** nested loop join algorithm for relations R and S (1.5 pts)

No. 8 State the **page-at-a time** nested loop join algorithm for relations R and S (1 pt).

Problems 9 and 10:

Show your work. (Zero points for the answer if not accompanied by explanations or detail work)

Consider the join $R \mid\!><\!\mid_{R.a = S.b} S$, given the following information about the relations to be joined. The cost metric is the number of page I/Os unless otherwise noted, and the cost of writing out the result should be uniformly ignored.

> Relation R contains 200,000 tuples and has 20 tuples per page.
> Relation S contains 4,000,000 tuples and has 20 tuples per page.
> Attribute b of relation S is the primary key for S. (Applicable to no. 9 only)
> Attribute a of relation R is the primary key for R. (Applicable to no. 10 only)
> Both relations are stored as simple heap files.
> Neither relation has any indexes built on it.
> 1002 buffer pages are available.

No. 9 What is the cost of joining R and S using a **block** nested loops join with R as the outer relation? (1 point)

No. 10 What is the cost of joining R and S using a **tuple-at-a-time** nested loop join with S as the outer relation? (1 point)