

Comprehensive Exam

Databases

Fall 2011

1) (20 points) Consider the following relational scheme for storing information associated with courses and faculty members at some university, which is not necessarily Florida Tech (Note that some of the rules may have changed from what you have seen before).

University_Info= (FACULTY#, DEPARTMENT, COURSE#, SECTION#, BOOK)

FACULTY#	A faculty identification number.
DEPARTMENT	A department name.
COURSE#	A course identification number.
SECTION#	The number of a course section, e.g., 01, 02, 03, etc.
BOOK	The ISBN number of a book.

Suppose that the following rules apply in the university.

1. Department names are math, computer science, biology, etc.
2. Course identification numbers are unique.
3. Every faculty member is assigned to one or more departments, i.e., joint appointments are allowed.
4. Each course that is taught in a given semester, has one or more sections.
5. Every faculty member teaches one or more course sections.
6. Different faculty members can teach different sections of the same course.
7. No more than one faculty member is assigned the same section of the same course.
8. Every faculty member selects one or more books for each course that they teach.
9. Different faculty teaching the same course can use the same or different books for their sections.
10. A faculty member may use different books for each section of a specific course that they teach.
11. No other rules apply.

Based on the above rules, circle each of the following functional dependencies that hold.

FACULTY# => DEPARTMENT

SECTION# => COURSE#

BOOK => COURSE#,SECTION#

COURSE#,SECTION# => FACULTY#

FACULTY#,COURSE#,SECTION# => BOOK

FACULTY#,BOOK => COURSE#

COURSE#,SECTION# => BOOK

SECTION#,BOOK => COURSE#

FACULTY#,COURSE# => SECTION#

DEPARTMENT,COURSE#,SECTION#,BOOK => FACULTY#

2) (20 points)

a) Define the term *transaction*.

b) Describe one example of a (non-trivial) transaction.

c) Explain why transactions are important to a DBMS.

d) Is there any special SQL syntax for specifying transactions (yes or no)? If so, then describe the syntax.

3) (20 points) Give a formal definition for each of the following normal forms.

(a) First normal form

(b) Second normal form

(c) Third normal form

(d) BCNF

4) (10 points) Define (formally) what is meant by a loss-less join.

5) (30 points) Consider the following collection of relation schemes.

employee(employee-name, street, city)

works(employee-name, company-name, salary)

company(company-name, city)

manages(employee-name, company-name, manager-name)

Note that an employee can have multiple employers (and managers) in the above database.

Give an SQL statement for each of the following.

(a) A list of employee names and for each employee a count of the number of companies they work for plus a total income, i.e., the sum of all the salaries for that employee.

(b) A list of all employees who do not work for either CVS or for IBM.

(c) A list of company names, and for each company the largest salary for that company, plus the name of the employee(s) who have the largest salary for that company. Note that, as shown below, if more than one employee has the maximum salary for a particular company, then that company will be listed multiple times, once with each employee who makes the maximum.

<u>company</u>	<u>max-salary</u>	<u>employee</u>
IBM	55000	Smith
CVS	42500	Jones
Harris	65000	Brown
Harris	65000	Carson
Wendys	35000	Pierce
	:	
	:	

(d) Give a relational algebra or tuple calculus expression (your choice) for a list of the names of those employees who work for every company that *Smith* works for, but none of the companies that *Jones* works for.