

Comprehensive Exam
Databases
Fall 2000

1) (20 points) Match up the functionality/statement on the right with the DBMS component on the left that is primarily responsible for implementing the stated functionality.

DML compiler	_____	a) Provides fast access to data items that hold particular values.
Embedded DML precompiler	_____	b) Responsible for fetching data from disk storage into main memory, and deciding what data to cache in (and out of) memory.
Authorization and integrity manager	_____	c) Manages the allocation of space on disk storage and the data structures used to represent information stored on disk.
Transaction manager	_____	d) Stores the actual data that is in a database.
File manager	_____	e) Stores metadata about the structure of the attributes, tables, users, and other data items that are managed by the database management system.
Buffer manager	_____	f) Stores information about the data in the database that is used by the query processor to select efficient ways to execute a query.
Data files	_____	g) Converts DML statements in an application program to normal procedure calls in the host language.
Statistical data	_____	h) Tests for the satisfaction of constraints and checks the access privileges of users trying to access data.
Data dictionary	_____	i) Ensures that the database remains in a consistent (correct) state despite system failures, and that concurrent operations proceed without conflicting.
Indices	_____	j) Translates DML statements in a query language into low-level instructions that the query evaluation engine understands.

2) Suppose that a logical model is to be developed for the information associated with a medical clinic. In this clinic patients are allowed to make appointments with doctors in advance, or they are allowed to walk-in without an appointment. In either case, clinic requirements dictate that the date and time they are admitted, as well as the time they are checked out be tracked, in addition to the total amount billed for the visit (cost). Note that there are no overnight stays, but there can be multiple visits in the same day by the same patient. Also of great concern to the clinic is the policy number for each patient's medical insurance. It is assumed that patients will have only one medical insurance policy, members of the same family may be listed on the same policy, and patients without insurance are *not* allowed to visit the clinic. Each patient is assigned to exactly one doctor the first time they are admitted to the clinic, and this doctor is the one who sees the patient on all future visits.

Finally, each doctor has a unique office number.

Note that the following relations could be identified from the above description. Also note that attributes forming the primary key for each relation have been underlined.

- patient(ss#, name, policy#)
- visitation(ss#, date-admitted, time-admitted, time-checked-out, cost)
- doctor(ss#, name, office#)

(a) (7 points) Give a collection of DDL statements for creating database tables for the above relations. Note that your solution should include declaration of all appropriate integrity constraints.

(b) (5 points) Give a relational algebra statement for listing the dates on which *Al Gore* was admitted to the clinic.

(c) (5 points) Give a tuple calculus expression for the query from part (b).

(d) (5 points) Give a domain calculus expression for the query in part (b).

(e) (5 points) Give an SQL statement for the query in part (b).

(f) (7 points) Give an SQL statement that, for each patient, will list their social security number and the total cost of all their visits to the clinic.

3) (36 points) For each of the following, consider the relation scheme and associated functional dependencies.

(a) customer (customer-name, customer-street, customer-city)

customer-name \Rightarrow customer-street
customer-name \Rightarrow customer-city

Is the above relation scheme in 1NF (yes or no)?

Is the above relation scheme in 2NF (yes or no)?

Is the above relation scheme in 3NF (yes or no)?

Is the above relation scheme in BCNF (yes or no)?

(b) Loan-Info(branch-name, loan-number, customer-name, amount)

loan-number \rightarrow amount
loan-number \rightarrow branch-name

Is the above relation scheme in 1NF (yes or no)?

Is the above relation scheme in 2NF (yes or no)?

Is the above relation scheme in 3NF (yes or no)?

Is the above relation scheme in BCNF (yes or no)?

(c) depositor (customer-name, account-number)

*No functional dependencies (except trivial ones)

Is the above relation scheme in 1NF (yes or no)?

Is the above relation scheme in 2NF (yes or no)?

Is the above relation scheme in 3NF (yes or no)?

Is the above relation scheme in BCNF (yes or no)?

4) (10 points) Define a multivalued dependency. For this question you may provide a formal definition, an example, or both.