

CS 474, Spring 2019
Written Exam #1

Name: _____ Email: _____@dukes.jmu.edu

*By writing your name, you acknowledge the following honor code statement:
“I have neither given nor received unauthorized assistance on this examination.”*

Directions

1. This is a closed-book, closed-notes, no-calculator exam. Do not refer to any materials other than the exam itself.
2. Do not look at anyone else’s exam. Do not talk to anyone but the instructor during the exam. Turn off all cell phones, etc.
3. If you have a question, please come to the front of the room. If you need to leave the room, you must first ask for permission.
4. You will have up to 50 minutes to complete the exam. Please write clearly and show all your work.

Question	Maximum	Score
MC	8	
MC	8	
MC	4	
SA	10	
EER	10	
Map	10	
Total	50	

Multiple Choice

(2pts each) Select the best answer for each question. For partial credit, you may write a brief explanation for the answer you selected.

_____ Which statement is **not correct**?

- A. In a database approach, applications don't have their own files, but all applications access the same version of the data by interfacing with the DBMS.
- B. In a database approach, the data definitions or metadata are stored in the applications accessing the data.
- C. In a database approach, one generally avoids storing multiple copies of the same data.
- D. In a database approach, maintenance of data and metadata is easier.

_____ Which statement is **not correct**?

- A. In a file-based approach, every application has its own query and access procedures, even if they want to access the same data.
- B. SQL is a database language to manage DBMSs without having to write a substantial amount of programming code.
- C. SQL is a database language that focuses on how to access and retrieve the data from the file system.
- D. SQL is a database language that allows different applications to access different subsets of the data necessary for each application.

_____ Complete this sentence: In the three-layer architecture, between the external layer and the conceptual/logical layer, there is ...

- A. logical data independence.
- B. physical data independence.
- C. no independence; they are basically the same thing.
- D. the internal layer.

_____ Fill in the gaps in the following sentences:

1. When, during crash recovery, aborted transactions need to be undone, that is the task of the ... **1** ...
 2. The part of the storage manager that guarantees the ACID properties is the ... **2** ...
- A. 1: recovery manager 2: buffer manager
 - B. 1: lock manager 2: recovery manager
 - C. 1: recovery manager 2: transaction manager
 - D. 1: lock manager 2: lock manager

_____ Which of the following components is part of the query processor?

- A. DDL compiler
- B. DML compiler
- C. Security manager
- D. Transaction manager

_____ Evaluate the following statements:

1. Record-at-a-time DML means that the query gets recorded from the user at the time the user inputs the query and then gets processed.
 2. Record-at-a-time DML means that navigating the database starts with positioning on one specific record and going from there onwards to other records.
 3. Set-at-a-time DML means that the query gets set beforehand and then gets processed by the DBMS.
 4. Set-at-a-time DML means that many records can be retrieved in one DML statement.
- A. 1 and 3 are correct
 - B. 2 and 3 are correct
 - C. 1 and 4 are correct
 - D. 2 and 4 are correct

_____ Which of the following statements is **correct**?

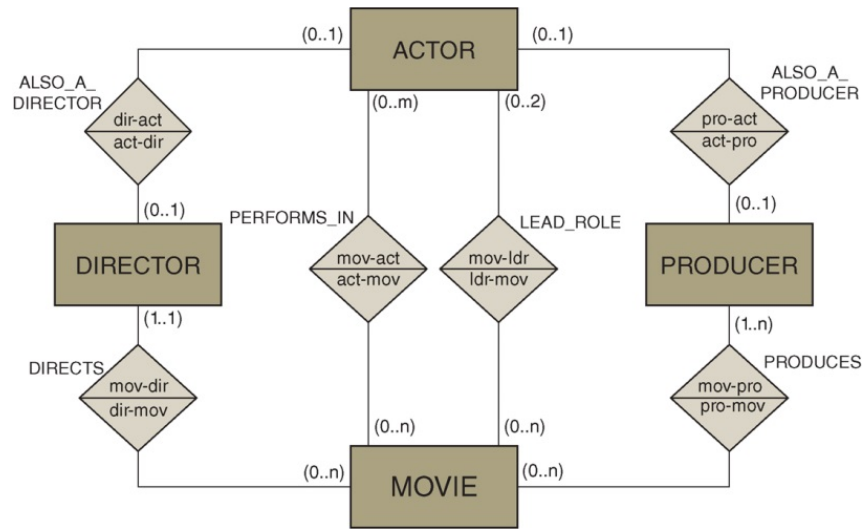
- A. A foreign key of a relation T cannot refer to the primary key of the same relation T .
- B. A relation cannot have more than one foreign key.
- C. Every relation must have a foreign key.
- D. A foreign key can be NULL.

_____ Consider the following relational model (primary keys are underlined, foreign keys are in italics). Which statement is **correct**?

STUDENT(student number, student name, street name, street number, city, zip code)
COURSE(course number, course name)
ENROLLED(*student number*, *course number*)
PROFESSOR(professor number, professor name)
TEACHES(*course number*, *professor number*)

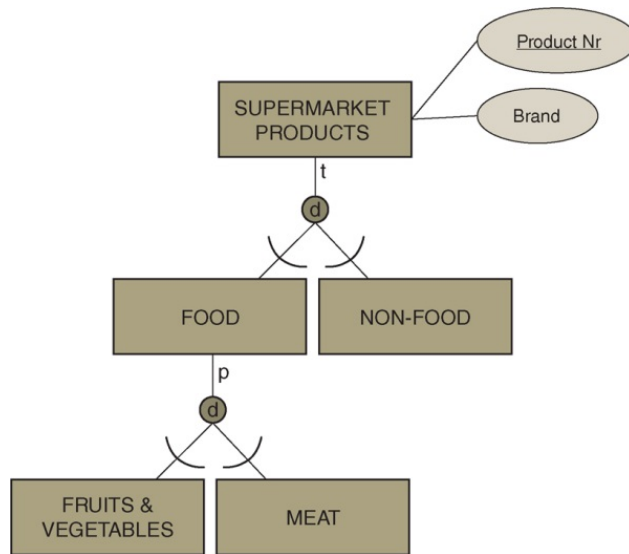
- A. The model can be further normalized.
- B. The model does not allow a professor to teach multiple courses.
- C. The model does not allow a course to be taught by multiple professors.
- D. The model does not allow a course to be enrolled by multiple students.

_____ Given the ER model, which of the following statements is **correct**?



- A. A movie can have as many lead actors as there are actors in the movie.
- B. PRODUCER is an existence-dependent entity type.
- C. A director of a movie can also act in the same movie.
- D. A movie can have multiple actors, producers, and directors.

_____ Given the EER specialization, which of the following is **correct**?



- A. A supermarket product can be a food and non-food product at the same time.
- B. There are certain supermarket products that are not fruits & vegetables, not meat, and not non-food.
- C. All food products are either fruits & vegetables or meat.
- D. A meat product does not have any attribute types.

Short Answer

(5 pts) Explain the purpose of the query optimizer, and give two specific examples of the kinds of decisions it makes.

(5 pts) Normalize the following relation into Third Normal Form. Underline primary keys, and circle foreign keys.

An art gallery wishes to maintain data on their customers, artists, and paintings. They generally have several paintings by each artist in the gallery at one time. Paintings may be bought and sold multiple times, at different prices (i.e, the gallery may sell a painting, buy it back at a later date, and then sell it to another customer). Customers have one address, one phone number, and multiple purchases.

R(custId, custName, address(street, city, state, zip), phone, purchases(artistId, artistName, paintId, title, dateSold, price))

EER Diagram

(10 pts) Draw an EER conceptual data model for the following scenario:

Recently, a social network site called Facepage was founded. When new users want to join Facepage, they must fill out a form with their personal information: name, email, and date of birth. Each user must have a unique email address.

Upon completion of the form, a new account is created with a unique account number generated by the system. The user needs to specify which type of account he/she prefers: a business account, or a personal account:

- A business account is specifically designed to support companies in their marketing campaigns. When a user opens a business account, he/she has to specify the name of the company. Users with a business account pay a monthly fee.
- When a user chooses a personal account, he/she can keep in touch with other Facepage users by forming friendships. Only personal accounts can send or receive friend requests. Each request has a status (e.g., pending, accepted).

Maintaining multiple accounts, regardless of the purpose, violates Facepage's Terms of Use. If a user already has an account of either type, the system will not allow the user to create another one. But an account may be transferred to a user not associated with an account.

Each account can create multiple pages on the site, and each page is owned by the account that created it. For each page, the number of visits is logged. For each account, no two pages can exist with the same title.

Relational Mapping

(10 pts) Map the following EER model for a train company to a relational model. Underline primary keys, and circle foreign keys. You may abbreviate the attribute names.

