Midterm Review Questions

CS 174A Fundamentals of Database Systems

Luis Ángel Larios Cárdenas · *Section A*

October 31, 2019

A. Fundamental Concepts

1. Enumerate at least 4 advantages of a DBMS.
2. What is a *data model*?
3. What is the *physical* or internal schema?
4. What is the *logical* or conceptual schema?
5. What is an *external schema*?
6. Define *logical data independence*.
7. Define *physical data independence*.
8. What is a DB *transaction*?
9. What is a *key constraint* in terms of the ER model?
10. What is a *participation constraint* in the ER model?
11. What is a *weak entity* set?
12. What is *class hierarchy* in the ER model?
13. Define *aggregation* in terms of the ER model.
14. What is a *superkey*?
15. What is a (candidate) *key*?
16. What is a *primary key*?
17. What is a *partial key*?
18. What is the difference between *overlap* and *covering constraints*?
19. What are *integrity constraints*?
20. What are *domain constraints*?
21. What is a *key constraint* in terms of the Relational model?
22. What are *foreign key constraints*?

B. Entity-Relationship Diagram

You have been hired to design the DB for an automotive dealer. The dealer sells **vehicles** for a given *price*, and cars are uniquely identified by a vehicle ID number or *vin*. Vehicles are of two distinct types: **electric** and **gas**. An electric car requires to store its *battery life*, while a gas vehicle needs to record its *cylinder capacity*.

Each individual vehicle is of a particular **model** which belongs to a brand or **manufacturing company** which we can recognize by a unique *name* (e.g., the Cambri is a model name of the car brand Toyota). We assume that a model name is unique among the models produced by its company. When a brand company is removed from the system, we need **not** keep information about its models in the DB. Additionally, a car company produces the same model in different *years*, but the dealer is interested on selling only the latest version. Further, a model comes with one or more **options**. We can distinguish options by looking at their unique *name*.

1. Provide the ER diagram for the problem described above and make sure to indicate any participation and/or key constraints.
2. Translate your ER diagram to SQL create statements. Indicate if there are any constraints that you can't capture in your create statements.

C. Relational Algebra

Consider the (simplified) relation schemas:

$$\mathbf{R}(a, b) \qquad \mathbf{S}(a, b) \qquad \mathbf{T}(b)$$

Express the following relational algebra operations using only the elementary operators: $\pi, \sigma, \times, \cup$ and $-$.

1. $R \cap S$
2. S/T

D. Relational Algebra + Domain Relational Calculus

Consider these schemas, where the *Flights* relation contains **only non-stop** flight information between two points, and an *Aircraft*'s *cruisingrange* is equivalent to the *distance* that an airplane can travel without recharging fuel.

Flights(*flno*: int, *from*: string, *to*: string, *distance*: int, *departs*: time, *arrives*: time)
Aircraft(*aid*: int, *aname*: string, *cruisingrange*: int)
Certified(*eid*: int, *aid*: int)
Employees(*eid*: int, *ename*: string, *salary*: int)

Express the following queries in relational algebra and domain relational calculus.

1. Find the *eid*s of pilots certified for some 'Boeing' aircraft.
2. Find the *names* of pilots certified for some 'Boeing' aircraft.
3. Find the *aid*s of all aircrafts that can be used on non-stop flights *from* 'LAX' *to* 'GDL'.
4. Identify the *flno*s of the non-stop flights that can be piloted by *every* pilot whose *salary* is more than $100,000.
5. Find the *eid*s of employees who are certified to fly *exactly* 2 aircrafts.
6. Find the *names* of pilots who can operate planes with a range greater than 3000Km but are not certified on any 'Boeing' aircraft.