

◆ Review your homework and quizzes. Review the “Review Questions” at the end of each chapter. **READ THE BOOK.**

Know and understand the design steps,

- ▶ products of design steps
- ▶ purpose
- ▶ audience

What in the output of each design step?

Differences between Semantic/Relational models

Define Key.

Key Constraint vs Participation Constraint.

Be able to explain/use Entity/Relationship/Primary Key/Candidate Key/Class Hierarchy/Role

Relate these to the Real World

Be able to determine when to use an Entity/Relationship/Descriptive Attribute

Key constraint/Participation constraint

- ▶ identify/draw/explain

▶ A student takes exactly one course. What is the key constraint? What is the participation constraint?

You will need to draw a small ER diagram.

Given a small ER diagram, provide example data (a set of tables) that shows how the tables would be built and how the constraints would be implemented.

Be able to describe all the pieces of an ER diagram.

Given a small E-R diagram, be able to answer questions about what data it can store/queries it can answer.


- ▶ how would you change the ER diagram to satisfy these requests?

Given a set of CREATE TABLE statements, build some simple SQL queries to satisfy users' requests.

Select, join, sort, group by.  
no views, subqueries

Work through all the queries in the SQL notes. Make sure you understand how to build all the queries.

Questions:



Build a small example database (draw the tables as below and fill the tables with data) that correctly implements the above ER diagram. The Garages Table is laid out for you. Note, the Garages/ParksIn arrow is NOT bold. The Cars/ParksIn arrow IS bold.

GarageID	GarageNumber	Rent

Build CREATE TABLE commands for each table above.

Write SQL queries to determine:

1. All the Garage Numbers.
2. All the Garage Numbers that contain a car.
3. The Garage Number of all the Red cars.
4. The average rent for Blue cars
5. How many cars of each color are in the database?

Should Color be an Entity or Attribute?

What type of relationship is ParksIn? Many to many, one to many, one to one?

Alter the diagram: (Answer each of the following questions assuming the previous questions have been answered and the appropriate changes have been made to the diagram.)

Add an owner for each car. An owner may have multiple cars but each car has an owner. Other people are permitted to drive a car, however. Any time a car leaves or enters a garage the date and time is logged as well as the driver. Only people listed as a driver for a car may drive the car out of the garage.

Record the license plate for each car.

Record when each car started parking in its current garage.

What are all the candidate keys for ParksIn? What are all of the super keys for ParksIn?

Why is garage number not the primary key above?

Why does the car have a date of birth?

Log all the previous garages that a car has parked in and when the car started and stopped parking there.

Bonus: Detect any instances where a car left a garage with one driver and returned, from the same trip, with another driver

Detect any instances where a car returns to the garage with a non-permitted driver.

Describe each of the advantages of using a database over a text file.

What is a descriptive attribute? Give an example of a descriptive attribute, or add one if none exists, to the Car/Garage example. Where is the descriptive attribute located? Why?

Why is the Client/SoftwareProduct an aggregation connected to Service Plan rather than having a single relationship directly connecting all three (why does the latter solution allow you to produce bad data?).

Why is using a Foreign Key constraint necessary in ParksIn?

What is the command to build a ssh tunnel from your VM to the database machine? Why does this need to go through zeus?

What is a key constraint? How is a key constraint represented in the ER diagram and in SQL?

What columns or set of columns might need to be marked as an INDEX in the Garage/Car ER diagram. Why?

How are the access patterns to a social media site like Facebook or LinkedIn different than the access patterns to a site like YourBank.com or WeatherAlerts.com? How do those differences impact the database hosting each of those sites?