

Oracle Academy Database Design Instructor Resource Guide

Notes For PowerPoint Slides:

SECTION 7 LESSON 1 – Arcs

Slide 1: Arcs

Some students will be confused as to when to use an arc vs. a supertype/subtype:
Use supertypes/subtypes when you want to represent classifications or types of things.
For example: STUDENT (graduating, nongraduating).
Use arcs when you want to represent mutually exclusive relationships between entities.
For example: STUDENT may attend a UNIVERSITY or a TRADE SCHOOL.
A UNIVERSITY or a TRADE SCHOOL may be attended by one or more STUDENT. It is not a type of STUDENT.
It is also a M:1 relationship, which therefore would not be suited to a supertype/subtype construct.

Slide 2: What Will I Learn?

When discussing arcs, point out that they can sometimes also be modeled as supertype/subtypes. Supertypes/subtypes are best used when you are trying to classify within an entity (e.g, university student: full-time, part-time). Arcs should be used when the exclusive relationships between the entities are more important to the business. The lesson will show examples of these.

Slide 3: Why Learn It?

No instructor notes for this slide

Slide 4: Tell Me /Show Me – What is a constraint?

Constraint- A restriction that applies to data, which is often dictated by the business rules.
Another example already discussed: an EMPLOYEE **must** have a last name (mandatory attribute). In this slide we simply introduce a single generic word (constraint) to cover all of them.

Slide 5: Tell Me /Show Me – Mutually exclusive relationships

Exclusive OR relationships- A logical operator that returns a true value if one, but not both, of its operands is true
Mutually exclusive relationship- A relationship that has choices that are unable to be true at the same time
Point out the arc in the diagram. It represents the exclusive relationship across billboard and movie, product advertisement and public announcement. The arc tells the reader of the diagram that only one of these “features” will have the relationship with

BILLBOARD for each instance. The optional nature of the relationships indicates that each BILLBOARD may be empty -- that is, it may not have any of the “features.”

Slide 6: Tell Me /Show Me – The DJ business has a mutually ...
No instructor notes for this slide

Slide 7: Tell Me /Show Me – Arcs are a way to represent mutually ...

Arc- A curved line used to represent an exclusive relationship in an entity-relationship diagram.

Point out that each entity has its own attributes. Each PUBLIC SPACE may be the venue for one or more EVENTS. Each PRIVATE HOME may be the venue for one or more EVENTS.

Some students may ask: why not just model one entity and make the rental fee optional? Answer: Although that may well be the decision made when the database is built, at this conceptual level, we want to capture all the business rules as closely as possible. The rental fee is mandatory for a PUBLIC SPACE so we want our model to reflect that, as well as the exclusive OR relationship between EVENT and VENUE.

Slide 8: Tell Me /Show Me – Arcs can sometimes be drawn as supertypes ...

If the mutually exclusive entities on the “one” end of the M:1 relationship can be (and should be) combined into a single entity with subtypes, then clearly we don’t need an arc. This would actually be a good choice here, since there are many common attributes between the entities PRIVATE HOME and PUBLIC SPACE, so there is no overriding need to model them as two distinct entities.

Slide 9: Tell Me /Show Me – Supertypes and subtypes can sometimes ...

In this example, EVENT PLANNER, DJ and MANAGER have distinct sets of attributes, so there is a strong case for modeling them as distinct entities. Note that when a supertype/subtype is redrawn as an arc, the relationships in the arc are mandatory 1:1 relationships. Each PARTNER must be an EVENT PLANNER, or a DJ or a PROJECT MANAGER, or an OTHER. Each EVENT PLANNER must be a PARTNER, each DJ must be a PARTNER, each MANAGER must be a PARTNER, each OTHER must be a PARTNER.

Slide 10: Tell Me /Show Me - Terminology
No instructor notes for this slide

Slide 11: Summary – In this lesson you have learned how to:
No instructor notes for this slide

Slide 12: Summary - Practice Guide
No instructor notes for this slide

SECTION 7 LESSON 2 – Hierarchies and Recursive Relationships

Slide 1: Hierarchies and Recursive Relationships

What to Watch For

Be sure the class understands that hierarchy and recursive relationships essentially model the same structure. Be prepared to answer questions about pros and cons of each one.

PRO Hierarchical: Hierarchical structures are more explicit and are easier for most people to understand because they are very similar to an organizational chart. Each entity can have its own mandatory attributes and relationships, if the business requires this (instead of all optional attributes and relationships, as you would have in a recursive). In this way, your data model truly reflects the business rules.

PRO Recursive: Recursive relationships tend to be simpler because you are using only one entity. Your diagram will be less “busy.” However, they are less specific -- you cannot have mandatory attributes or relationships unless they are mandatory in all instances of the entity.

Another thing to consider when modeling these types of relationships is how often the structure changes. If it changes often, then a recursive relationship is easier to maintain. If it is fairly fixed, then you may consider the more explicit hierarchy.

A structure that doesn’t change very much would be a building with suites and floors and rooms. This is a good case for using an explicit hierarchy. On the other hand, a company that reorganizes frequently (going from a fairly flat organizational structure to one with many levels and vice versa) would probably be better modeled with a recursive relationship.

Connections

Go back to the ERD with ROOM, SUITE, FLOOR, and BUILDING. Ask about the transferability of the relationships. Can a floor be moved from one building to another? How should the ERD be modified?

Answer: These are fairly permanent physical structures. Although walls could be torn down or built up to reorganize rooms into different suites, a suite is not likely to move up or down a floor. Similarly, a floor cannot physically relocate to another building. The relationships between BUILDING and FLOOR and between FLOOR and SUITE are nontransferable and should be represented with a diamond. The business rules will determine whether the relationship between ROOM and SUITE is transferable or not.

Slide 2: What Will I Learn?

No instructor notes for this slide

Slide 3: Why Learn It?

Ask students to name other hierarchical structures.

Answers could include: government, military, biology (species, genus).

Slide 4: Tell Me /Show Me – An organizational chart can be represented ...

Explain that one way to model a hierarchy would be to create an entity for each level with a relationship to the next level.

The UIDs would probably be an employee ID or number for each entity.

Ask students: how would the ERD change if the business was reorganized to add or remove a level ?

Slide 5: Tell Me /Show Me – What about this ERD?

Here you have a case of the cascading UIDs: the UID of FLOOR is the combination of FLOOR number and the BUILDING id, the UID of SUITE is the combination of SUITE number and the FLOOR number and the BUILDING id, and the UID of ROOM is the combination of ROOM id and SUITE number and FLOOR number and the BUILDING id. It is important to have the students realize this.

This may be convenient in the sense that the UID of a ROOM will also tell you the SUITE and FLOOR and BUILDING that it's in. However, this makes for a rather long UID. Unique independent, artificial codes may be more practical. Also, if the hierarchical structure changes often, use artificial identifiers. (If we were to add a level called APARTMENT in between FLOOR and SUITE, think of how that would affect the UIDs of all SUITES and ROOMs!)

Slide 6: Tell Me /Show Me – Both of these models represent all employees

Hierarchical relationship- A series of relationships that reflect entities organized into successive levels.

Recursive relationship- A relationship between an entity and itself.

Explain that the mandatory attributes specific to an entity (bonus plan, car plan, budget) become optional in the recursive model.

Why is the recursive relationship fully optional?

Answer: Otherwise, the hierarchy does not have a top or bottom. Ask students who manages the president? No one! That is why it must be fully optional. One instance of the entity does not have a manager. Hence, it is optional in the recursive model.

Important: point out that these two models are not identical. In the hierarchical model, it is clear that every employee except the president **must** report to someone. In the recursive model, the fully optional relationship implies that it is possible that **no** employees report to someone. A programmatic rule would also be needed. Or, the different levels of employee could be represented as subtypes within EMPLOYEE.

The ER diagramming convention that shows a recursive relationship is known as a "pig's ear."

The loop can appear on any side of the entity's softbox.

Slide 7: Tell Me /Show Me – In the DJ model, the project manager has ...

Explain to students that, according to the relationship, only managers can supervise other employees. The use of subtypes makes this clear. However, the ERD still makes it possible for managers to manage other managers !

Slide 8: Tell Me /Show Me - Terminology
No instructor notes for this slide

Slide 9: Summary – In this lesson you have learned how to:
No instructor notes for this slide

Slide 10: Summary - Practice Guide
No instructor notes for this slide

SECTION 7 LESSON 3 – Modeling Historical Data

Slide 1: Modeling Historical Data

Lesson Preparation

Modeling data that changes over time can be a tricky subject. In this course, we will cover this subject twice -- here and again in Section 8. There are several implications to a model that incorporates the time element, so we are spreading out the material.

In this lesson, we will discuss the factors that influence the decision to model historical data, create a new entity to track it, and define a UID for such an entity.

Slide 2: What Will I Learn?

No instructor notes for this slide

Slide 3: Why Learn It?

Some possible answers to the questions posed:

A pharmaceutical company may want to track complaints filed about their drugs. This could help them detect problems early and issue recalls if necessary. Or they may want to conduct a study that follows patients on a new drug for an extended period of time. This could help scientific research and lead to further improvements in their products.

A commercial dairy may want to keep a record of rejected units per batch, which machine it came from or the farm that was the source of the cows or products. If there is a pattern over time, it could suggest ways to improve the processing methods to minimize waste and reduce costs.

Slide 4: Tell Me /Show Me – When is it necessary to model data over time?

Historical data- Data that records changes over time.

Audit trail- A step-by-step record of data changes.

Point out that in the real world, most data changes over time, but the business may not need to track all of it.

They need to validate requirements for storing historical data with the user. Storing unnecessary historical data can be costly. It takes up space -- whether it's in a filing cabinet or in the database. It takes up employee resources -- someone has to be responsible for updating the historical records with the latest data, for making sure that there is enough space as the data grows, and for allocating more space or moving data around if necessary.

Slide 5: Tell Me /Show Me – An organization needs to keep data about...

Ask the class why the model shown will not work based on the new requirements.

Answer: an employee may have many salaries at different times. Since each employee is only one instance of EMPLOYEE (note the UID), multiple salary amounts and start dates would be a repeating group of multivalued attributes, thus violating First Normal Form.

Slide 6: Tell Me /Show Me – To model salary changes over time, add a SALARY...

Point out that this relationship is mandatory at both ends, because obviously every salary belongs to an employee, but also every employee must have at least one salary. The organization does not expect people to work for no money!

Why is the salary end date optional? Answer: because if it is the current salary, there is no end date (yet).

Which salary is the current one? We could add a status attribute to SALARY HISTORY with possible values Current or Previous, but this is not needed. The current salary is the one whose end date is null.

There would also need to be a programmatic constraint to ensure that an employee has exactly one salary at a time: the start date of a new salary cannot be earlier than the end date of the previous salary.

Slide 7: Tell Me /Show Me – A jewelry store rents pieces ...

Ask students how the store would know when the jewelry piece was rented? How would they know when it was returned? This will help the students realize that to track the rental history, we really need a M:M relationship between JEWELRY PIECE and MOVIE STAR. They may try to put "date rented on" in JEWELRY PIECE, but this will be problematic because it can be rented out on different dates by different MOVIE STARS. Neither does the date rented belong in MOVIE STAR because different JEWELRY PIECES may be rented out by the same MOVIE STAR but on different dates.

Slide 8: Tell Me /Show Me – The relationship between JEWELRY ...

Ask the class if drawing barred relationships makes an adequate UID for the RENTAL HISTORY entity.

Answer: No, because although RENTAL HISTORY tracks the date of the rental as well, a UID consisting only of MOVIE STAR id plus JEWELRY PIECE code will not allow a MOVIE STAR to rent the same JEWELRY piece more than once (on different dates). To explore this further, you can play the "moving UID" game. Draw the model on the board, but this time, just have a UID bar between RENTAL HISTORY and MOVIE STAR, but keep date as part of the UID. Ask the class: What does this model say?

Answer: It says that a MOVIE STAR can rent only one JEWELRY PIECE on a given day. What if she/he wanted to rent more? Data example: If (give a movie star or celebrity that the class will recognize) wanted to rent both the Princess necklace and the Maharajah bracelet on 14-Jun, she could not, because MOVIE STAR identifier plus date rented must be unique.

(continued on next slide)

Slide 9: Tell Me /Show Me – The UID of RENTAL HISTORY is ...

Try It / Solve It

Now, revise the model. This time, draw barred relationships from RENTAL HISTORY to JEWELRY PIECE **and** to MOVIE STAR. Ask the class: Now what does the model say?

Answer: Two different MOVIE STARS could rent the same JEWELRY PIECE on the same DAY! That would not work at all.

Data example: (Movie Star 1 and Movie Star 2 -- name people that the class will recognize) want to rent the Winston diamond earrings on 26-Mar. The combination of MOVIE STAR identifier, JEWELRY PIECE code, and date rented are unique for each instance, so the model would allow this. However, we know that, in reality, this cannot happen.

Revise the model again. This time, just have a barred relationship between RENTAL HISTORY and JEWELRY PIECE, but keep the date as part of the UID – as shown on the slide.

Ask the class: Now what does the model say? Answer: It says that a JEWELRY PIECE can be rented only once on the same date. This is what we want.

Data example: Movie Star 1 wants to rent the Princess necklace on 05-May. Movie Star 2 can rent it on 12-May. The combination of JEWELRY PIECE code and date rented is unique.

Tell the class that this stresses the importance of defining the right UID -- it has implications for how the business can operate!

Slide 10: Tell Me /Show Me - Terminology
No instructor notes for this slide

Slide 11: Summary – In this lesson you have learned how to:
No instructor notes for this slide

Slide 12: Summary - Practice Guide
No instructor notes for this slide

Notes For Practice Activities:

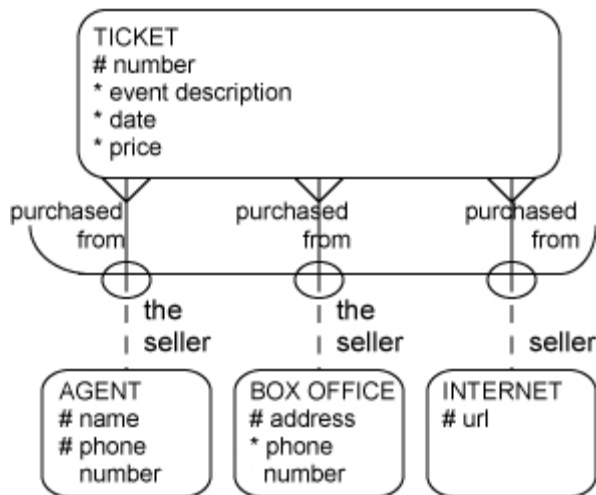
Arcs S07-L01

Vocabulary

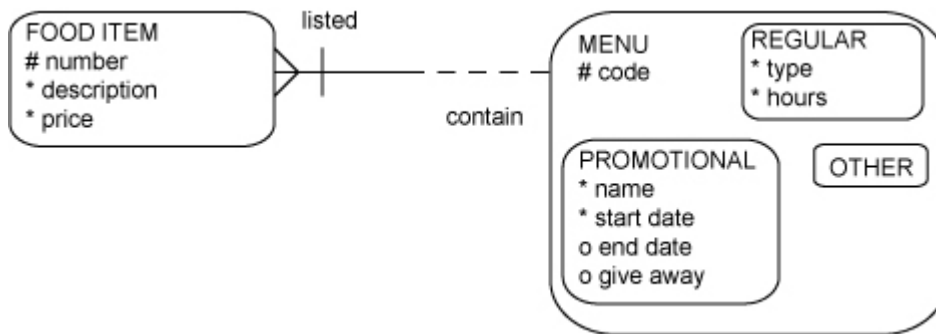
<u>Arc</u>	A curved line used to represent an exclusive relationship in an entity-relationship diagram
<u>Exclusive OR relationship</u>	A logical operator that returns a true value if one, but not both, of its operands is true.
<u>Mutually exclusive relationship</u>	A relationship that has choices that are unable to be true at the same time.
<u>Constraint</u>	A restriction that applies to data, which is often dictated by the business rules.

Try It / Solve It

1.



2. One acceptable solution would be to have four subtypes of MENU (BREAKFAST, LUNCH, PROMOTIONAL, OTHER). The type attribute would not be needed in the BREAKFAST and LUNCH subtypes.



Hierarchies and Recursive Relationships S07-L02

Vocabulary

Recursive relationship A relationship between an entity and itself.

Hierarchical relationship A series of relationships that reflect entities organized into successive levels.

Try It / Solve It

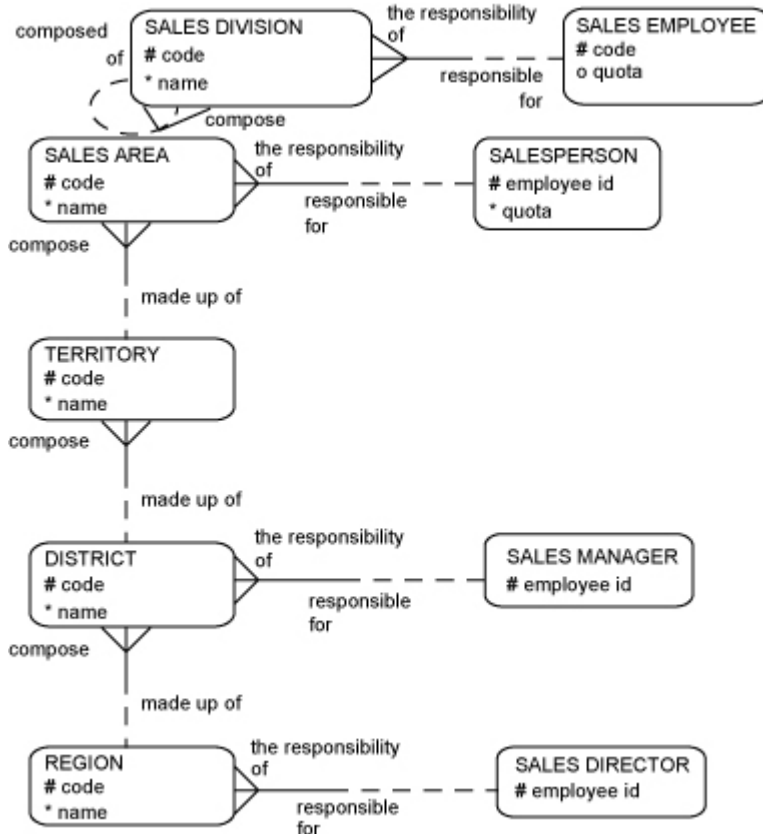
1. You may want to explain "quota" as an amount of sales that the salesperson must achieve.

A lot of students will also include hierarchical relationships from SALESPERSON, SALES MANAGER, and SALES DIRECTOR (Each SALESPERSON must be reporting

to one and only one SALES MANAGER, each SALES MANAGER may be the manager of one or more SALESPERSONs, etc.).

They may also model a recursive relationship between SALES EMPLOYEE (Each SALES EMPLOYEE may be the manager of one or more SALES EMPLOYEEs. Each SALES EMPLOYEE may be managed by one and only one SALES EMPLOYEE.)

These relationships are not part of the solution, because it is not explicitly stated in the business scenario. However, it is not a bad assumption to make. Just remind them that if they make assumptions, they have to eventually check them with the customer.

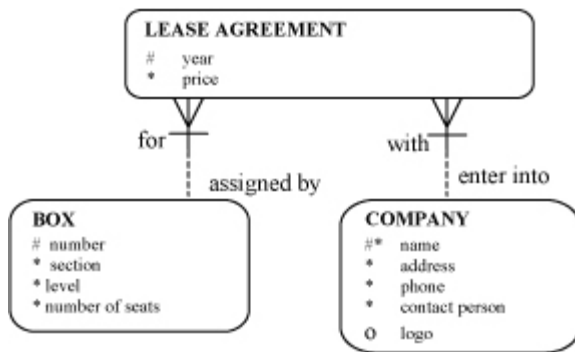


Modeling Historical Data S07-L03

Try It / Solve It

1. Actors at the local theater sign contracts that specify their pay and working conditions. These contracts are negotiated **annually**. Each year the contracts may **change** through new language and pay rates. In order to resolve labor disputes, the theater must sometimes **go back several years** in order to see the contract provisions in place **on a particular date**.

2.

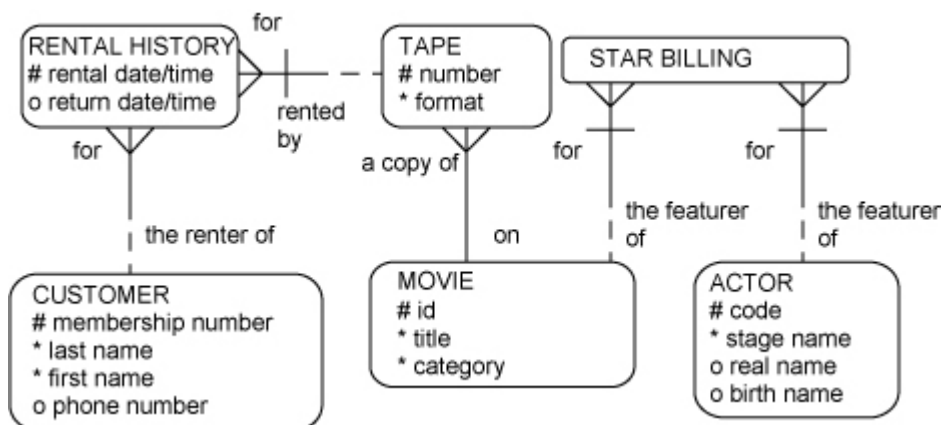


3. If students need guidance, hint that the relationship between TAPE and CUSTOMER has now changed (to a M:M, which then needs to be resolved).

A lot of students will create barred relationships from RENTAL HISTORY to TAPE and from RENTAL HISTORY to CUSTOMER, and also make rental date/time part of the UID. This is incorrect, because all you need to establish uniqueness is the relationship to TAPE and the rental date/time. You may want to play the “moving UID” game again:

- If the UID included rental date/time and the barred relationship to CUSTOMER, it would mean that a CUSTOMER could not rent more than one tape at a time.
- If the UID included rental date/time and the barred relationship to TAPE, it would mean that a tape could be rented out only once on a given date/time. **THIS IS WHAT WE WANT.**

-Adding the second barred relationship is wrong, because it would mean that two different CUSTOMERs could rent the same TAPE on the same date/time.



4. Provide examples and/or formats for the written documentation, such as:

Statement of the problem:

“The IOUG would like to ...” or “The IOUG is having trouble with...” or “The IOUG needs help with...”

Statement of the proposed solution:

“We recommend that the IOUG ...”

Information requirements of the business clearly stated:

“The IOUG would like to keep track of the following information:” (list them here)

The business rules as they apply to the information requirements of the company:

Example of business rules:

Each member is required to pay dues in January.

The assumptions that had to be considered:

Example of assumption:

It is possible to have an event that nobody attends.