

***Modern Database Management 13th edition Hoffer
Solutions Manual***

SKU: 9780134773650-SOLUTIONS

Chapter 2 Modeling Data in the Organization

Chapter Overview

The purpose of this chapter is to present a detailed description of the entity-relationship model and the use of this tool within the context of conceptual data modeling. This chapter presents the basic entity-relationship (or E-R) model. Advanced features of conceptual data modeling will follow in Chapter 3.

Chapter Objectives

Specific student learning objectives are included in the beginning of the chapter. *From an instructor's point of view, the objectives of this chapter are to:*

1. Emphasize the importance of understanding organizational data and convince your students that unless they can represent data unambiguously at the conceptual level, they cannot implement a database that will effectively serve the needs of various organizational stakeholder groups.
2. Present the E-R model as a conceptual data model that can be used to capture the structure and much, although not all, of the semantics (or meaning) of data.
3. Apply E-R modeling concepts to several practical examples including the Pine Valley Furniture Company case.

Key Terms

Associative entity	Entity-relationship model (E-R model)	Relationship type
Attribute	Entity type	Required attribute
Binary relationship	Fact	Simple (or atomic) attribute
Business rule	Identifier	Strong entity type
Cardinality constraint	Identifying owner	Term
Composite attribute	Identifying relationship	Ternary relationship
Composite identifier	Maximum cardinality	Time stamp
Degree	Minimum cardinality	Unary relationship
Derived attribute	Multivalued attribute	Weak entity type
Entity	Optional attribute	
Entity instance	Relationship instance	
Entity-relationship diagram (E-R diagram)		

Classroom Ideas

1. Review the major steps in the database development process (Figure 1-8) and highlight the importance of data modeling in determining the overall data requirements of information systems. Lead a discussion regarding the actors within an organization that typically are most heavily involved in each of the steps and how end users may best

- participate in the process.
2. Introduce the concept of drawing models to represent information in a concise manner by having your students participate in a small active exercise in map-making. Divide the students into teams of three or four students each so that you have an even number of teams in the class. Instruct each team to work together to investigate and develop a map to selected campus locations (you develop the list ahead of time; e.g., from this classroom to the library, from this classroom to a colleague's office, etc.). Ask each team to verify the map they draw and then return to the classroom. Pair up each team with a unique location with another team; ask the teams to exchange maps. Instruct each team to then verify the map they received by following it and then returning to the classroom. Conduct a debriefing discussion about how easy/hard it was to follow the maps, how useful were the symbols used, how easily understood were the symbols, etc. Use this discussion to lead into the use of E-R notation used to represent data models and why standardization is useful to systems development activities.
 3. Use the sample E-R diagram shown in Figure 2-1 to introduce the first conceptual model to your students. Ask them to explain the business rules represented in this diagram.
 4. Use Figure 2-2 to summarize the basic E-R notation used in this chapter and throughout the remainder of the text.
 5. Contrast the terms entity type and entity instance (see Figure 2-3). Discuss other examples: STUDENT with each student in the classroom as an instance, etc. Warn the students that the term "entity" is often used to denote either an entity type or an entity instance; the meaning is intended to come from the *context* in which it is used.
 6. Give examples of common errors in E-R diagramming, including inappropriate entities (see Figure 2-4). Ask your students for other examples.
 7. Compare strong versus weak entities using Figure 2-5. Ask your students for other examples.
 8. Discuss the various types of attributes that are commonly encountered (Figures 2-7 through 2-9). Again, ask your students to think of other examples.
 9. Make sure your students understand the difference between relationship types and relationship instances (Figure 2-10).
 10. Introduce the notion of an associative entity by using Figure 2-11. Discuss the four reasons (presented in the text) for converting a relationship to an associative entity.
 11. Discuss unary, binary, and ternary relationships (Figure 2-12). Have the students brainstorm at least two additional examples for each of these relationship degrees.
 12. Discuss the bill-of-materials unary relationship (Figure 2-13). Use a simple and familiar product (such as a toy) to illustrate this essential structure, which is often difficult for students to understand.
 13. Introduce the concept and notation of cardinality constraints in relationships (Figures 2-16, 2-17, and 2-18). Emphasize that these constraints are important expressions of business rules.
 14. Introduce the problem of representing time dependent data. Use Figures 2-19 and 2-20 to illustrate different means of coping with time dependencies.
 15. Discuss examples of multiple relationships between entities (Figure 2-21). Ask your students to suggest other examples.
 16. Use the diagram for Pine Valley Furniture Company (Figure 2-22) to illustrate a more comprehensive E-R diagram. Stress that in real-world situations, E-R diagrams are often much more complex than this example.
 17. As time permits, have your students work in small teams, two or three students each, to solve some of the E-R diagramming tasks in the Problems and Exercises section of the

chapter. We have included a number of new examples for this purpose. Also, you may assign the project case as a homework exercise.

Answers to Review Questions

2-1. *Define each of the following terms:*

- a. *Entity type*. A collection of entities that share common properties or characteristics
- b. *Entity-relationship model*. A logical representation of the data for an organization or for a business area
- c. *Entity instance*. A single occurrence of an entity type
- d. *Attribute*. A property or characteristic of an entity type that is of interest to the organization
- e. *Relationship type*. A meaningful association between (or among) entity types
- f. *Strong entity type*. An entity type that exists independently of other entity types
- g. *Multivalued attribute*. An attribute that may take on more than one value for a given entity instance
- h. *Associative entity*. An entity type that associates the instances of one or more entity types and contains attributes that are peculiar to the relationship between those entity instances
- i. *Cardinality constraint*. Specifies the number of instances of one entity that can (or must) be associated with each instance of another entity
- j. *Weak entity*. An entity type whose existence depends on some other entity type
- k. *Identifying relationship*. The relationship between a weak entity type and its owner
- l. *Derived attribute*. An attribute whose values can be calculated from related attribute values
- m. *Business rule*. A statement that defines or constrains some aspect of the business

2-2. *Match the following terms and definitions:*

- | | |
|---|------------------------|
| i | composite attribute |
| d | associative entity |
| b | unary relationship |
| j | weak entity |
| h | attribute |
| m | entity |
| e | relationship type |
| c | cardinality constraint |
| g | degree |
| a | identifier |
| f | entity type |
| k | ternary |
| l | optional attribute |

2-3. *Contrast the following terms:*

- a. *Stored attribute; derived attribute*. A stored attribute is one whose values are stored in the database, while a derived attribute is one whose values can be calculated or derived from related stored attributes.

- b. *Minimum cardinality; maximum cardinality.* A minimum cardinality specifies the minimum number of instances of one entity associated with an instance of the related entity (typically zero or one), whereas a maximum cardinality specifies the maximum number of such instances.
- c. *Entity type; relationship type.* An entity type is a collection of entity instances that share common properties or characteristics, while a relationship type is a meaningful association between (or among) entity types.
- d. *Strong entity type; weak entity type.* A strong entity type is an entity that exists independently of other entity types, while a weak entity type depends on some other entity type.
- e. *Degree; cardinality.* The degree (of a relationship) is the number of entity types that participate in that relationship, while cardinality is a constraint on the number of instances of one entity that can (or must) be associated with each instance of another entity.
- f. *Required attribute; optional attribute.* A required attribute must have a value for each entity instance, whereas an optional attribute may not have a value for every entity instance.
- g. *Composite attribute; multivalued attribute.* A composite attribute has component parts that give meaning, whereas a multivalued attribute may take one or more values for an entity instance.
- h. *Ternary relationship; three binary relationships.* A ternary relationship is a simultaneous relationship among the instances of three entity types and often includes attributes unique to that simultaneous relationship. Three binary relationships reflect the three two-way relationships between two entity types, and do not depict the same meaning as a ternary relationship.

2-4. *Four reasons underlying the importance of data modeling:*

- a. The characteristics of data captured during data modeling are crucial in the design of databases, programs, and other system components. Facts and rules that are captured during this process are essential in assuring data integrity in an information system.
- b. Data, rather than processes, are the most important aspects of many modern information systems and hence, require a central role in structuring system requirements.
- c. Data tend to be more stable than the business processes that use the data. Thus, an information system that is based on a data orientation should have a longer useful life than one based on a process orientation.
- d. Data modeling facilitates interaction between designers, programmers, and end users.

2-5. *Four reasons underlying the preference for the business rules approach:*

- a. Business rules are a core concept in an enterprise since they are an expression of business policy, and they guide individual and aggregate behavior. Well-structured business rules can be stated in a natural language for end users and in a data model for system developers.
- b. Business rules can be expressed in terms that are familiar to end users. Thus, users can define and then maintain their own rules.
- c. Business rules are highly maintainable: they are stored in a central repository and each rule is expressed only once, then shared throughout the organization.

- d. Enforcement of business rules can be automated through the use of software that can interpret the rules and enforce them using the integrity mechanisms of the database management system.

2-6. *Where can you find business rules?*

Business rules appear in descriptions of business functions, events, policies, units, stakeholders, and other objects. These descriptions can be found in interview notes from individual and group information systems requirements collection sessions, organizational documents, and other sources. Rules are identified by asking questions about the who, what, when, where, why, and how of the organization.

2-7. *Six general guidelines:*

- a. Data names should relate to business, not technical characteristics.
- b. Data names should be meaningful, almost to the point of being self-documenting.
- c. Data names should be unique from the name used for every other distinct data object.
- d. Data names should be readable. The names should be structured in a way that is consistent with how the concepts would most naturally be said.
- e. Data names should be composed of words taken from an approved list.
- f. Data names should be repeatable, meaning that different people or the same person at different times should develop exactly or almost the same name.

2-8. *Four criteria:*

- a. Choose an identifier that will not change its value over the life of each instance of the entity type.
- b. Choose an identifier such that for each instance of the entity the attribute is guaranteed to have valid values and not be null (or unknown).
- c. Avoid the use of so-called intelligent identifiers (or keys), whose structure indicates classifications, locations, and so on.
- d. Consider substituting single-attribute surrogate identifiers for large composite identifiers.

2-9. *Why some identifiers must be composite rather than simple?*

An identifier attribute is an attribute (or combination of attributes) whose value distinguishes individual instances of an entity type. Often, a simple attribute will not be unique for all instances of an entity type (e.g., FlightNumber for an instance of an airline flight). Rather, a combination of simple attributes will be needed to uniquely identify the entity instance (e.g., FlightID and FlightDate would make the instance unique).

Please note that you can always create a surrogate key that is guaranteed to have a unique value.


2-10. *Three conditions for an associative entity type:*

- a. All of the relationships for the participating entity types are “many” relationships.
- b. The resulting associative entity type has independent meaning to end users, and it preferably can be identified with a single-attribute identifier.


- c. The associative entity has one or more attributes in addition to the identifier.

2-11. *Four types of cardinality constraints:*

- a. Optional one:



- b. Mandatory one:



- c. Optional many:



- d. Mandatory many:



2-12. *Example of weak entity:*


Phone Call (see below) is an example of a weak entity because a phone call must be placed by a PERSON and thus, an instance of PHONE CALL cannot exist without an instance of PERSON. In this simple example, PHONE CALL is related to only one other entity type. Thus, it is not necessary to show the identifying relationship; however, if this data model were ever expanded so that PHONE CALL related to other entity types, it is good practice to always indicate the identifying relationship.



2-13. *Degree of relationship definition & examples:*

The degree of a relationship is the number of entity types that participate in the relationship.


a) Unary (one entity type):



b) Binary (two entity types):



c) Ternary (three entity types):




2-14. *Attribute examples:*

- a. Derived – Distance (rate x time); both rate and time could be stored, and then when the data is retrieved from the database (e.g., at run-time) the distance could be calculated from the already-stored data elements
- b. Multivalued – spoken language; a person can speak more than one language
- c. Atomic – Social Security Number; this United States National Identification number cannot be broken down into component parts
- d. Composite – Phone Number; a phone number is often broken down into country code, area code, and the rest of the phone number
- e. Composite identifier – Flight ID could consist of Flight Number and Flight Date, together forming a unique identifier for an airline flight.
- f. Optional – Middle Initial; a person's middle initial may be optional for identification

purposes or also because some people may not have a middle name


2-15. *Examples of relationships:*

(a) Ternary



The sale of a property is a simultaneous relationship among the PROPERTY, a BUYER, and an OWNER entity types. This “event” cannot be modeled appropriately with three binary relationships; any one of the three binary relationships (PROPERTY-BUYER; BUYER-OWNER; and PROPERTY-OWNER) is missing an essential element of the sale.


(b) Unary



In an on-campus dormitory/apartment situation, this diagram shows a recursive/unary relationship among instances of the STUDENT entity type. This notation indicates only the current roommate situation between instances of the STUDENT entity type.

2-16. *Effective (or effectivity) dates:*

Effective (or effectivity) dates are used in a data model when the organization wishes to record historical data, rather than just the current instance. A few examples might include the effective date of a product price or service rate. Another example might be the start and end date of an advisor’s assignment to work with a student at a university (see E-R segment below, which includes a multivalued composite attribute Advisor).



2-17. *Rule for moving attribute to another entity type:*

A data modeler should consider extracting an attribute from one entity type and placing it in another entity type linked by a relationship when the attribute is the identifier or some other characteristic of an entity type in the data model, and multiple entity instances need to share these same attributes.

2-18. *Special guidelines for naming relationships:*

- A relationship name should always be a verb phrase and should state the action taken, as opposed to the result of the action taken.
- Use descriptive, powerful verb phrases as opposed to vague names.

2-19. *The relationship definition should also explain the following:*

- a. any optional participation
- b. the reason for any explicit maximum cardinality
- c. any mutually exclusive relationships
- d. any restrictions on participation in the relationship
- e. the extent of history that is kept in the relationship
- f. whether an entity instance involved in a relationship instance can transfer participation to another relationship instance

2-20. *Manages relationship in Figure 2-12a:*

Presently, the cardinality is one-to-many. One possible scenario is an employee who is supervised by more than one manager. This would make the cardinality many-to-many. Another possibility (although quite rare in practice) is that the employee is supervised by one manager, and the manager only supervises one employee. This would result in a one-to-one cardinality. If we take time/history into consideration, the idea of someone being managed currently versus never being managed could affect the cardinality. As we can see here, you cannot always tell what the business rule is by looking at the ERD. These possible scenarios will need to be discussed with the end user to determine the “correct” modeling representation for the business rules at this organization.

2-21. *Entity type vs. Entity instance:*

An entity type can be thought of as a template, defining all of the characteristics of an entity instance. For example, “student” would be an entity type, whereas *you* are an instance of “student.”

2-22. *Conversion of ternary relationship into an associative entity:*

Converting a ternary relationship into an associative entity is recommended for two main reasons: (1) research has shown that participation/cardinality constraints cannot be accurately represented for a ternary relationship with current notation; and (2) most E-R diagramming tools cannot represent ternary relationships. By converting a ternary relationship into an associative entity with three mandatory binary relationships, a data modeler can accurately represent the participation/cardinality constraints although there is a risk that the meaning/semantics of the original ternary relationship will be lost with this solution.

Answers to Problems and Exercises

2-23. *Cellular Operator Database Figure 2-24 questions:*

a. *Can a customer have an unlimited number of plans?*

Yes. A Customer may be responsible for 0, 1, or many Plans.

b. *Can a customer exist without a plan?*

Yes. The minimum cardinality of the Belongs relationship from the Customer to the Plan states that a Customer may exist without a Plan (the minimum cardinality is 0).

c. *Is it possible to create a plan without knowing who the customer is?*

No. The minimum cardinality of both the “responsible for” and “belongs” relationships between Plan and Customer states that at least one Customer must be related to a Plan.

d. *Does the operator want to limit the types of handsets that can be linked to a specific plan type?*

Yes, the cellular operator requires that a Handset (that is a particular type and a particular operating system) is linked to one Plan (that is a particular type of plan). This business rule is to be implemented in this design by indirectly requiring that a Plan Type has 0:M Plans, and each Plan is associated with certain Handsets, and each Handset is of some Handset Type. A given Plan Type is related to Handset Type through the intermediary entity types in this design.

[Alternative interpretation: No, there is nothing in the current model that creates a condition that would limit – in advance – the handset types that can be related to a specific plan type.]

e. *Is it possible to maintain data regarding a handset without connecting it to a plan?*

Yes. The minimum cardinality of the Includes relationship between Plan and Handset states that a Handset may be included in 0 or 1 plan. The 0 minimum cardinality means that we can track data about the handset even if it is not connected to a plan; the Handset has optional participation in the Includes relationship with Plan.

f. *Can a handset be associated with multiple plans?*

No. The minimum cardinality of the Includes relationship between Plan and Handset states that a Handset may be included in 0 or 1 plan, not multiple plans.

g. *Assume a handset type exists that can utilize multiple operating systems. Could this situation be accommodated within the model included in Figure 2-24?*

No. The current model shows that a handset type is associated with one and only one operating system.

h. *Is the company able to track a manufacturer without maintaining information about its handsets?*

Yes. The minimum cardinality of the relationship between Manufacturer and Handset Type indicates that we can track data about a Manufacturer even if we have no (or zero) Handset Types in our database.

i. *Can the same operating system be used on multiple handset types?*

Yes. The maximum cardinality on the relationship between Operating System and Handset Type indicates that an Operating System may be used on 0, 1, or many Handset types.

j. *There are two relationships between Customer and Plan. Explain how they differ.*

The *Responsible For* relationship is an overall 1:M relationship between Customer and Plan. A Customer can be responsible for 0, 1, or many Plans yet any one Plan will be linked to only 1 Customer for responsibility purposes. The *Belongs* relationship is an overall M:M relationship that permits the linking of multiple customers to a single plan, as in the case of family members being part of a particular plan or different plans.

k. *Characterize the degree and the cardinalities of the relationship that connects Customer to itself. Explain its meaning.*

The “Family Member” relationship that connects Customer to itself has a degree of 1 (unary). It permits the tracking of each family member as a Customer. Any Customer may be a Family Member of 0, 1, or many Customer(s); as a Family Member Customer, the Customer may be linked to 0 or 1 Customer.

l. *Is it possible to link a handset to a specific customer in a plan with multiple customers?*

No, this is not possible according to the current model. However, the current model could be adjusted to create an Associative Entity to track the particular Customer instance with a particular Plan instance that is then associated with a particular Handset. This suggested extension to the current model also permits a design that will easily extend the database’s ability to track additional data about the particular Customer instance with a particular Plan instance.

m. *Can the company track a handset without identifying its operating system?*


No. The minimum cardinality of the relationship between Handset Type and Operating System is 1 and only 1; the minimum of 1 is a mandatory participation for the Handset Type with the Operating System.

2-24. For each of the descriptions below, perform the following tasks:

- i) Identify the degree and cardinalities of the relationship.
- ii) Express the relationships in each description graphically with an E-R diagram.

a. A book is identified by its ISBN, and it has a title, a price, and a date of publication. It is published by a publisher, each of which has its own ID number and a name. Each book has exactly one publisher, but one publisher typically publishes multiple books over time.


- (i) This relationship is a degree of 2 (binary). This relationship is one-to-many from Publisher to Book.
- (ii)



Note: This solution assumes that we have a reason to track a Publisher even if it does not yet have a Book published.


b. A book (see above in (a)) is written by one or multiple authors. Each author is identified by an author number and has a name and date of birth. Each author has either one or multiple books; in addition, occasionally data are needed also regarding prospective authors who have not yet published any books.

- (i) This relationship is a degree of 2 (binary). This relationship is many-to-many from Author to Book.
- (ii)




c. In the context specified above in (a) and (b), better information is needed regarding the relationship between a book and its authors. Specifically, it is important to record the percentage of the royalties that belong to a specific author, whether or not a specific author is a lead author of the book, and each author's position in the sequence of the book's authors.

- (i) This relationship is a degree of 2 (binary). This relationship is many-to-many from Author to Book.
- (ii)




- d. A book (see (a) above) can be part of a series, which is also identified as a book and has its own ISBN number. One book can belong to several sets and a set consists of at least one but potentially many books.

- (i) This relationship is a degree of 1 (unary). This relationship is many-to-many.
- (ii) This solution assumes that “series” and “sets” are synonymous terms. The question does not require that a series have any special attributes or distinguishing features, so it can be represented in the data model like any other Book instance and identified by ISBN.



- e. Ebony and Ivory, a piano manufacturer, wants to keep track of all the pianos it makes individually. Each piano has an identifying serial number and a manufacturing completion date. Each instrument represents exactly one piano model, all of which have an identification number and a name. In addition, the company wants to maintain information about the designer of the model. Over time, the company often manufactures thousands of pianos of a certain model, and the model design is specified before any single piano exists.


- (i) These relationships have a degree of 2 (binary). These relationships are one-to-many.
- (ii)



f. *Ebony and Ivory* (see (e) above) employs piano technicians who are responsible for inspecting the instruments before they are shipped to the customers. Each piano is inspected by at least two technicians (identified by their employee number). For each separate inspection, the company needs to record its date and a quality evaluation grade.

(i) This relationship is a degree of 2 (binary). This relationship is many-to-many.


(ii)



g. *The piano technicians* (see (f) above) have a hierarchy of reporting relationships: some of them have supervisory responsibilities in addition to their inspection role and have multiple other technicians report to them. The supervisors themselves report to the chief technician of the company.

(i) This relationship is a degree of 1 (unary). This relationship is one-to-many.


(ii) Because the chief technician is not represented as a separate entity type, that person does not have a supervisor. This, in turn, leads to the 0 minimum cardinality on the 1 side of the unary relationship.



h. *Chiclets Electronics (CE)* builds multiple types of tablet computers. Each type has a type identification number and a name. The key specifications for each type include amount of storage space and display type. The company uses multiple processor types, exactly one of which is used for a specific tablet computer type; obviously, the same


processor can be used in multiple types of tablets. Each processor has a manufacturer and a manufacturer's unique code that identifies it.

- (i) This relationship is a degree of 2 (binary). This relationship is many-to-many.
- (ii)




- i. *Each individual tablet computer manufactured by CE (see (h) above) is identified by the type identification number and a serial number that is unique within the type identification. The vendor wants to maintain information about when each tablet is shipped to a customer.*

- (i) These relationships are a degree of 2 (binary). These relationships are one-to-many. If, over time, shipment of a tablet computer to multiple customers (e.g., as in a refurbished unit) is possible, the Tablet Computer – Customer relationship would become many-to-many and the Shipping Date attribute would become an attribute of that M:N relationship.
- (ii)



- j. *Each of the tablet computer types (see (h) above) has a specific operating system. Each technician the company employs is certified to assemble a specific tablet type – operating system combination. The validity of a certification starts on the day the employee passes a certification examination for the combination, and the certification is valid for a specific period of time that varies depending on tablet type – operating system combination.*

- (i) This relationship is a degree of 2 (binary). This relationship is many-to-many.
- (ii) Based on the limited situation description, it appears that there is no need to model a separate entity type for Operating System. If the situation required additional data about the Operating System and the Technician's certification for this element, the diagram would need to be revised accordingly.



2-25. Each answer refers to Figure 2-22 found in the chapter text.

- a. Where is a unary relationship, what does it mean, and for what reasons might the cardinalities on it be different in other organizations?

A unary relationship is shown with the EMPLOYEE entity; An EMPLOYEE Supervises 0:M EMPLOYEEs, An EMPLOYEE Is Supervised By 0:1 EMPLOYEE. This relationship tells us that we can determine which employees are supervised by another employee, as well as determine which employees are supervisors in this company.

In other organizations, there may be different policies regarding employee supervision that could cause the data relationships among EMPLOYEE instances to be different. For instance, another company might allow an employee to have multiple supervisors (e.g., in an organization with a matrix structure).

- b. Why is Includes a one-to-many relationship and why might this ever be different in some other organization?

Includes is a one-to-many (1:M) relationship because of the business rules that PVFC has in place: “a product line may group any number of products but *must group at least one product*; and each product *must belong to exactly one product line*.” Another organization may have other business rules that could permit a product being assigned to more than one product line (changing Includes to a M:N relationship). Alternatively, another organization might also show Includes as a (1:M) overall relationship but might permit the establishment of a PRODUCT LINE without identifying PRODUCTS that belong to this group (e.g., thus permitting an optional minimum cardinality on the PRODUCT side of the Includes relationship).

- c. Does Includes allow for a product to be represented in the database before it is assigned to a product line (e.g., while the product is in research and development)?

No, Figure 2-22 shows that the PRODUCT must be Included in at least 1 PRODUCT LINE by the mandatory 1 and only 1 cardinality notation near the PRODUCT LINE portion of the Includes relationship line. The cardinality notation would have to be changed to show optional 1 cardinality in order to represent the research and development situation.

- d. Suppose there is a rating of the competency for each skill an employee possesses, where in the data model would we place this rating?

The Has Skill associative entity, which associates a single instance of a SKILL with a single instance of an EMPLOYEE, would permit the tracking of a competency rating for each skill in which an employee has competence.

- e. What is the meaning of the DOES BUSINESS IN associative entity and why does each DOES BUSINESS IN instance have to be associated with exactly one TERRITORY and CUSTOMER?

The DOES BUSINESS IN associative entity associates a single instance of a TERRITORY with a single instance of a CUSTOMER for the overriding M:N DOES BUSINESS IN relationship between TERRITORY and CUSTOMER. Each DOES BUSINESS IN instance must be related to exactly one TERRITORY and one CUSTOMER because the business rules of PVFC indicate that sales territories have been established for its customers. In particular, the rules are: *a TERRITORY has one-to-many CUSTOMERs; and a CUSTOMER may do business in 0:M TERRITORIES*. When converting this M:N relationship on the ERD, the cardinalities near the originating entities will always be mandatory one, indicating the exactly one relationship with each entity's instances and the associative entity's instance.

- f. In what way might Pine Valley change the way it does business that would cause the Supplies associative entity to be eliminated and the relationships around it to change?

According to current business practice at PVFC, each RAW MATERIAL is provided by 1 or more VENDORS and a VENDOR supplies 0, 1, or many RAW MATERIALs and this is represented by the Supplies associative entity. The PVFC could consider entering into exclusive supplier arrangements with particular vendors such that an instance of RAW MATERIAL is supplied by only 1 VENDOR. If that situation should occur, then the overall relationship between RAW MATERIAL and VENDOR would change to 1:M (instead of M:N) and the Supply Unit Price attribute could become part of the RAW MATERIAL entity instance; the Supplies associative entity would no longer need to be on the ERD.

2-26. *Analysis of Figure 2-22:*

- a. Entities PRODUCT, PRODUCT LINE; relationship Includes
- b. Entities CUSTOMER, ORDER; relationship Submits
- c. Entities ORDER, PRODUCT; associative entity ORDER LINE
- d. Entities CUSTOMER, TERRITORY; associative entity DOES BUSINESS IN
- e. Entities SALESPERSON, TERRITORY; relationship Serves
- f. Entities PRODUCT, RAW MATERIAL; relationship Uses
- g. Entities RAW MATERIAL, VENDOR; relationship Supplies
- h. Entities WORK CENTER, PRODUCT; associative entity PRODUCED IN
- i. Entities EMPLOYEE, WORK CENTER; associative entity WORKS IN
- j. Entity EMPLOYEE; relationship Supervises, Is Supervised By

2-27. *Use of CASE or drawing tool:*

Student answers will vary based on the CASE or drawing tool that is used and their personal experiences. The answers should describe their experiences with the CASE or drawing tool in terms of the requirements of the E-R notation used in the chapter. Expect


to see students make reference to noting identifiers, using associative entities, using cardinality constraints properly, indicating required vs. optional attributes, and noting derived/composite/multivalued attributes.

2-28. *ER diagrams in Figure 2-25:*

- a. The ERD for City B does not (nor does any ERD) tell us why the cardinality is 1:M. The more restrictive cardinality for City B could be due to a business rule that they want to maintain only current volunteers but it could also be due to only tracking the agency for which the volunteer works the most hours of assistance. More detailed discussions would need to be held with the end users to properly document this business rule; notes should be added to the diagram to depict the appropriate business rule.
- b. The ERD for City A shows that a volunteer may assist one, none, or several agencies.
- c. The native notation used in ERDs does not show whether membership in a relationship can change (i.e., whether a volunteer can change agencies or whether an agency can change its volunteers). Some DBMSs can be told whether membership can change or not, and special notation or textual notes can be added to an ERD to state such business rules. The minimum cardinality next to Agency does address whether a Volunteer must always be associated with an Agency to exist in the database, but none of the cardinalities control whether linkages between specific agencies and volunteers can change. More detailed discussions would need to be held with the end users to properly document this business rule; notes should be added to the diagram to depict the appropriate business rule.

	City A	City B	Can't Tell
a. Which city maintains data about only those volunteers who currently assist agencies?			X
b. In which city would it be possible for a volunteer to assist more than one agency?	X		
c. In which city would it be possible for a volunteer to change which agency or agencies she assists?			X

2-29. ERD for ShinyShoesForAll:



2-30. Associative entities vs. Weak entities?

A weak entity requires the presence of another entity type; the weak entity does not exist independently from the other entity type and has no business meaning in the ERD without the other entity type. A weak entity will not have its own identifier, but will have a partial identifier attribute that will later be combined with the identifier of its strong entity owner to create a full identifier.

An associative entity is an entity type that associates the instances of one or more entity types and contains attributes specific to the relationship between those entity instances. An associative entity generally has an independent business meaning to end users and can be identified with a single-attribute identifier. If an associative entity meets these conditions, then it would not be considered a weak entity.

2-31. Figure 2-22 associative entities:

DOES BUSINESS IN: between TERRITORY and CUSTOMER

Although this entity has no attributes and no independent meaning, it is the only way that Visio can represent the M:N relationship between TERRITORY and CUSTOMER.

ORDER LINE: between PRODUCT and ORDER

This relationship has an attribute: Ordered Quantity that reflects the amount of product on each line of the order by the customer. It has independent meaning on the Customer's Order.

USES: between PRODUCT and RAW MATERIAL

This relationship has one attribute, Goes Into Quantity. It also may have independent meaning, although there is no obvious independent identifier.

SUPPLIES: between RAW MATERIAL and VENDOR

Since there is an attribute on this entity and it can have independent meaning, it might be a good candidate to convert to an associative entity.

PRODUCED IN: between WORK CENTER and PRODUCT

Although this entity has no attributes and no independent meaning, it is the only way that Visio can represent the M:N relationship between WORK CENTER and PRODUCT.

WORKS IN: between WORK CENTER and EMPLOYEE

Although this entity has no attributes and no independent meaning, it is the only way that Visio can represent the M:N relationship between WORK CENTER and EMPLOYEE.

HAS SKILL: between EMPLOYEE and SKILL


Although this entity has no attributes and no independent meaning, it is the only way that Visio can represent the M:N relationship between SKILL and EMPLOYEE.

There are so many associative entities because there are many M:N relationships that have independent meaning and because Visio's templates cannot represent M:N relationships.

2-32. *ERD for Figure 2-26 Grade Report:*

Student ID was chosen as the identifier for the STUDENT entity type as it is likely unique. Course ID was chosen as the identifier for the COURSE entity type as it is likely unique. Instructor Name was chosen as the identifier for the INSTRUCTOR entity type and it is assumed to be unique—should discussions during analysis work prove otherwise, it may be wise to create either (a) a composite identifier comprised of Instructor Name and Location, or (b) a new attribute Instructor ID that will be a unique number which can serve as an identifier (latter option would, in practice, be the most likely one).

Note: The addition of the Semester and Year attributes on the Registers For relationship allows this diagram (and resulting database) to reflect multiple semesters of data.




2-33.

Note: attributes are omitted from the ERD solutions for this Problem and Exercise in order to save space in the Instructor's Manual.


a. Figure 2-5




b. Figure 2-10a




c. Figure 2-12 (all parts)



d. Figure 2-13c




e. Figure 2-14




2-34. *Is Married To* relationship (Figure 2-12a) with time variations:


Problem & Exercise (a)



Problem & Exercise (b)



Problem & Exercise (c)



Problem & Exercise (d)




Diagram Notes for (d):

This solution presumes that Marriage Date is a partial identifier of the MARRIAGE entity; a full composite identifier will include Marriage Date and the two Person IDs involved in


the marriage. The solution also assumes that the same two people do not get married, divorced, and re-married on the same date. Adding a Marriage Time attribute (also a part of the identifier) would permit this situation to be covered by this model.

An alternate solution would be to use a surrogate identifier of License No instead of the suggested composite identifier of Marriage Date and the two Person IDs for the MARRIAGE entity.

(e):

The solution in 2-34d does not place any restrictions on the number of persons to whom any one person is simultaneously married, thus the 2-34d solution is sufficient in representing the lack of legal restrictions regarding the number of marriage partners.

2-35. Figure 2-27 Student, Club, School situation:



- A STUDENT Works For 0:1 SCHOOL; A SCHOOL Employs 0:M STUDENTS
- A STUDENT may belong to a CLUB only when located in the SCHOOL s/he Attends
- Student answers may vary. Alternative solutions include:
 - Since the STUDENT may not Work For a SCHOOL (the employment is optional), the Works For relationship is needed in the diagram in order to properly represent this business rule. This solution makes it harder for the database to enforce the business rule that a STUDENT works for the SCHOOL that s/he attends, but opens up the possibility that a STUDENT could Work For a SCHOOL that s/he is not currently attending.
 - An alternative design would be to remove the Works For relationship, and add an attribute to STUDENT named Works that would have a binary (Y/N) value to represent whether or not the STUDENT instance is working for the SCHOOL s/he Attends. The advantage of this design is that it would enforce the business rule that a STUDENT can only Work For a SCHOOL that s/he is currently attending.


2-36. Figure 2-28 diagrams showing stock price history:

Note: Student answers may vary. The crux of the answer relies upon what the purpose of the E-R diagram is for the modeling situation and how end users in the organization “see” the situation. In particular, do people in the organization have a term for stock price and refer to it as its own concept?



If so, solution B may be the “better” way to model this situation. Instructors may also use solution B to demonstrate an issue related to view integration (topic in Chapter 4) where transitive dependencies emerge; solution B makes the model easy to expand so that stock prices may have relationships that do not directly involve the STOCK entity.

Solution A indicates that each STOCK has multiple prices and is well-suited to early discussions with end users about the data needs of a system. Solution B adds the precision of multiple STOCK PRICE entity instances occurring for each STOCK entity instance. Solution B indicates that STOCK PRICE is a weak entity whose instances do not exist independently in the database without a corresponding STOCK entity instance. Solution B presents more precise detail of the data relationships that will likely be developed in the logical design of the database; this model may more closely resemble the relational model implementation of this design. Solution B also makes it easy to expand the model so that stock prices may have relationships with other entities that do not directly involve the STOCK entity.

2-37. Figure 2-11b (Modified):




- b) There could be more than one product finish for a product, which could affect the price.



c) Yes, this would be possible. For example, a customer could have more than one address.


2-39. ERD for Employee & Project situation:

a)



Yes, the attribute names do generally follow the guidelines for naming attributes.

b) ERD for Chemist, Project, Equipment situation:



Assignment: All three entities participate in the Chemical Use relationship that is modeled as an associative entity, since the Volume for each CHEMIST's use of a chemical in a particular project and equipment item must be tracked. However, CHEMICAL and PROJECT do not need to have any USE instances. All entities can have multiple assignments.

c) ERD for Course, Section situation:




Diagram Notes for 2-39c: SECTION is modeled as a weak entity. It could have been modeled as a multivalued attribute; however, using a weak entity is better, since SECTION may have a relationship with another entity. A multivalued attribute could not be used to show this relationship.

d) ERD for Hospital situation:






Diagram Notes for 2-39d: Both Admits and Treats relationships were created since the patient could be treated by other PHYSICIANS in addition to the admitting PHYSICIAN. Hospital was not included as an entity in this case as there was insufficient information in the scenario write-up to indicate that the data model needed to allow for multiple hospitals (e.g., in the case of a large health-care organization). The current ERD does not allow for the tracking of multiple admissions over time by different physicians. The ERD would need a M:N relationship between PHYSICIAN and PATIENT in order to track that kind of data. If the date of admission needs to be tracked, under the circumstances of tracking multiple admissions over time, the ERD could be revised to show Date Admitted as an attribute of the M:N Admits relationship, just as Treatment Detail is an attribute of the Treats relationship. The ERD could also be revised to show ADMISSION and

TREATMENT DETAIL associative entities (with corresponding attributes) instead of the M:N relationships currently discussed.


- e) **First situation:** credit check can be used by more than one request.



Second Situation: CREDIT CHECK can only be used by 1 CREDIT REQUEST (2 entities)




Second Situation: CREDIT CHECK can only be used by 1 request (1 entity)




Using one entity type seems much simpler since the credit check and rating only apply to this credit request. However, Credit Check Date and Credit Rating will have blank values (null) until the credit check is received.


f) Starting point diagram:



Situation 1 – Adding Hourly Rate attribute. This could be added to the CONSULTANT entity if the business rule is that a CONSULTANT Works for only one COMPANY at a time.



Situation 2 – Tracking a CONSULTANT's contract. Notice that CONTRACT is added as another entity that participates in a binary relationship with COMPANY and a binary relationship with CONSULTANT. We have moved the Hourly Rate attribute to the CONTRACT entity, which permits a CONSULTANT to vary his/her Hourly Rate as a function of the particular CONTRACT for a COMPANY. As only current CONTRACTs are tracked, an alternative solution would be to move the CONSULTANT attributes into the CONTRACT entity and eliminate the CONSULTANT entity from the model. The downside to this alternative solution is that Consultant Name and Consultant Specialty would occur redundantly in the CONTRACT entity instances.



Situation 3 – Tracking historical CONTRACT information. We can create an associative entity for CONTRACT. We’ve also added Contract ID as a surrogate identifier that is a unique serial number (not a composite identifier, as shown in Situation 2 above).



g) *Parking Garage ERD*





Diagram notes for 2-39g:

Please note that the solution assumes one customer instance that is assumed to be the unidentifiable hourly customer. For example, the customer instance with Cust ID of 99999999 could be reserved for this purpose.

h) *Law Firm ERD*

Note: This problem and exercise is a good lead-in for Chapter 3 modeling notation for the Enhanced Entity Relationship Diagram (EERD). The P&E offers several chances to provide better representation in the EERD (with subtyping) than the ERD notation that is provided in Chapter 2. Using EERD notation, a single LEGAL ENTITY can be shown as a supertype, with subtypes of DEFENDANT and PLAINTIFF. The ‘type’ (person or Organization) characteristic of both DEFENDANT and PLAINTIFF may also be considered for further subtyping. The solution presented here is a valid answer to the P&E, given the limitations of basic ERD notation and what is currently known about the situation.

This P&E also provides the instructor with an opportunity to discuss how history might be modeled if the business assumption regarding the tracking of Net Worth for both Plaintiff and Defendant was changed from only being concerned with Net Worth at the time of the CASE, to wanting to track the Net Worth over time of each party to the CASE. Refer to the chapter section on “Modeling Time-Dependent Data” and Figure 2-19 for more information on how this ERD could be revised.

Diagram Notes for 2-39h:

- Def Type and Plaintiff Type are used to denote Person or Organization type of legal entity.
- Net Worth of both Plaintiff and Defendant is relevant only at the time of the CASE, thus are modeled as attributes of the M:N relationships between CASE and PLAINTIFF, DEFENDANT.

i) Professional society (CAM)




Diagram notes for 2-39i:

- In this solution, it is assumed that dietary restrictions can be captured sufficiently as instances of DIETARY RESTRICTION TYPE. If further specification is needed a supporting column could be added to PARTICIPANT and/or REGISTRATION.

2-40. Star Hoist




Diagram Note for Problem & Exercise 2-40:

Note the difference between Part Number in PART and Part Instance Number in PART INSTANCE. In this case, PART represents a part type and PART INSTANCE the individual parts that all have their own individual identifying number.



2-41. *Emerging Electric ERD*


Diagram Notes for Problem & Exercise 2-41:

- A RATE may be for one, none, or many LOCATIONs.
- A LOCATION may have multiple CUSTOMERs.
- A CUSTOMER may own multiple LOCATIONs.

2-42. *STUDENT and ADVISER ERD*




2-43. Figure 2-4a Revised for Sarbanes-Oxley compliance purposes



2-44. Virtual Campus

a. ERD, First Phase




Note: It is assumed that there is no data about a PERSON until that person registers for an ACCOUNT on some SITE. It is also assumed that a SITE always has at least one ACCOUNT. POSTING is a weak entity because it cannot exist unless associated with some ACCOUNT.

b. One person, multiple accounts on the same site

Yes, one person can have multiple accounts on the same site. There is nothing in the ERD that would prevent this from happening: a person can have multiple accounts and multiple accounts can be associated with one site. There is no simple way to enforce this requirement in an ERD. Such a rule would have to be shown on the ERD as a note and then enforced in application software.

c. Phase 2, additional requirements



Note: The responses to a POSTING are other postings, represented by the Respond relationship. Each “response” posting will be associated with one ACCOUNT. READ is an associative entity that records when an ACCOUNT reads a particular POSTING (whether it be an initial posting or a response posting).

d. Ability to answer questions


It is possible to determine the number of postings each person has created for a specific site. Each ACCOUNT is associated with both a person and a site and thus, as long as we remember to take into account that each person can have multiple ACCOUNTs on the same site, we can determine which accounts need to be included in the query and then count the number of POSTINGs on each of these accounts.

The current structure associated each posting with only one account and, consequently, only one site. The business rules specified for Part c make this query invalid.

It is possible to answer this, but the logic is not simple. The process would start with a POSTING and identify the PERSON making that posting. Then search through any response postings for that POSTING to find the first response, and then follow the path from POSTING to ACCOUNT to PERSON to see who made that posting. If it is the same PERSON who made the initial posting, then this is a person who satisfies the query.


It would be straight forward to count the number of POSTINGs per SITE (through the associated ACCOUNT associative entity) and identify any with a count of zero.

e. Phase three: Complaints added




Note: A COMPLAINT is associated with an ACCOUNT, and then to a PERSON, and a COMPLAINT is associated with a specific POSTING. The Removed Status attribute has been added to the POSTING entity to store hold a logical value for yes/no to remove. The site administrator is an actor using the social network site, but that person is indirectly associated with a COMPLAINT by the attribute in the SITE for the ACCOUNT that make the objectionable POSTING.

f. Final phase: Private sites



Note: GROUP is a weak entity because an instance cannot exist without some PERSON creating it. It is not an associative entity, but rather has a separate M:N relationship with the PERSONs who are in the group, and the groups a person is a member of. A GROUP is associated with exactly one private SITE, but since not all sites are private, some SITES are not associated with a group. Application software would make sure that the value for Site Administrator for a private site matches the PERSON who created the site. And, application software would enforce the business rule that only group members can post to their private site.

2-45. Preliminary ERD for Symphony Orchestra




Business Rule: A concert includes the performance of one or more compositions; a composition may be performed at one or more concerts or may not be performed. This business rule is modeled in the ERD above through the use of the COMPOSITION and CONCERT entities, together with the PERFORMANCE Associative Entity.


Note: The use of the Associative Entity PERFORMANCE also permits the independent binary relationship between SOLOIST and PERFORMANCE, which permits the model to support the tracking of derived data, Date Last Performed. Although the diagram *appears* to represent a ternary relationship among COMPOSITION, CONCERT, SOLOIST and PERFORMANCE, such a ternary relationship would not accurately reflect the requirements of the problem. Rather, the needs of the problem state that there is an overall M:N binary relationship between SOLOIST and PERFORMANCE, which permits the

tracking of multiple soloists performing any given composition as well as a given soloist performing multiple compositions.


2-46. *Miami-Dade County court system*




2-47. *The Sensing Building Company ERD*



2-48. Stillwater Antiques ERD




2-49. A.M. Honka School of Business ERD



Note: Contact Type refers to mail, email, telephone, fax, or personal discussion.

2-50. Wally's Wonderful World of Wallcoverings ERD:

a. Part I




Note: So far, this is a fairly standard order processing database structure. Any reasonable customer demographic attributes and product characteristic attributes, that might be used to identify customer-product affinity, would be fine (e.g., customer address might be used to identify climate or décor related products).

b. Explanations for queries


- It is possible to tell which other customers bought the same product(s) as a specific customer had bought: we can first identify the products based on the ORDER, LINE ITEM and PRODUCT entities and after that the customers who had ordered the same product(s) based on these same entities by identifying first all ORDER LINES in which the product(s) are included and then the ORDERs that include the relevant ORDER LINES. Each ORDER is, in turn, associated with exactly one PERSON.
- Yes, it is also possible to continue the process above and identify other products that the relevant customer set had purchased (using the process described above).
- Yes, this is possible: the database includes complete information regarding the customer interests, and it is possible to identify those customers whose interest set has at least three overlapping items with the customer of interest.

c. Part II




Note: The Similar to M:N relationship represents the subjective assessments of product similarity, if there are any, for each product. Each VIEW entity instance holds the associative entity attributes associated with each product a person views.

d. Part III



Note: ADDRESS and REVIEW are weak entities. A PERSON will have one or more ADDRESSES, and each add ADDRESS may be used for multiple ORDERS. Each LINE ITEM on an ORDER may have a review submitted for it, and READ is a typical associative entity linking the PERSONS with the REVIEWS they read.


e. Part IV



- Note: The Address attribute in ADDRESS can be used as a qualifier to find other PERSONS with the same address attribute value (which ever part of Address they want to use).
- The attributes added to PRODUCT will make this query possible.
- This requirement simply demands a way to sort the results of a query, so this will be possible because the possible sorting attributes are available.


2-51. Doctors Information Technology (DocIT):

a. Part I




Note: APPOINTMENT is a classic associative entity. A patient may not have any allergies, but a staff member must have some credentials. The multi-valued attribute Reason may come from a standard list of reason codes, or may include custom text provided by the patient.

b. Part II



Note: The concept of a relationship between PATIENT/CONTACTs is an associate entity, which each relationship pair from one patient/contact to another described by the multi-valued attribute Nature of Relationship. This structure is similar to the bill-of-materials structure shown in the chapter.

c. Phase III




Note: Each PATIENT (but not CONTACT) may have many POLICY(ies), kept in Sequence, and each POLICY is from one company, so POLICY is an associative entity. CLAIM is an associative entity between an APPOINTMENT and a POLICY.

d. Responses to queries


- This can be calculated by counting the number of CLAIM entity instances with Amount Paid < Amount Claimed.
- The count in the first query can be grouped by Company ID following the relationship from CLAIM to POLICY then to INSURANCE COMPANY.
- Similar to the prior query, the Amount Claimed can be summed for each group defined by each Staff ID, by following the relationship from CLAIM to APPOINTMENT then from APPOINTMENT to STAFF.
- Given the database design in Part c, there is no easy way to determine what are likely different Staff ID and Patient ID values. This could be handled in a different database design in which all persons are in one super entity type (this concept is introduced in Chapter 3).

e. Part IV

Note: CLAIM PROCESSING STEP is an associative entity relating STAFF and CLAIM; each CLAIM may be the result of a new submission from a claim processing step, and a claim processing step may cause a new claim to be submitted. See the solution on the next page.



2-52. Revision to prior ERD in 2-49



Note: Contact Type refers to mail, email, telephone, fax, or personal discussion.

Changes made to prior 2-49 ERD

- Added qualifiers to attribute names in the STUDENT entity

Entities:

- Student: A person who attended and graduated from the A.M. Honka School of Business.
- Event: School events held around the world.
- Contact: The School's records of any contact made with a former student and graduate of the School.

Attributes on Student:

Student No: A unique identifier for a student when they attended the School. This attribute must be unique and is required.

Student Name: The name of the student when attending the School. This attribute is required.

Student Current Name: Current name of the former student.

Student Current Address: Current address of the former student.

Student Country of Birth: Country where the former student was born.

Student Country of Citizenship: Country where the former student holds citizenship.

Student Major: The name of the academic major completed by the former student. A student may have one or two majors. This attribute is required.

Attributes on Event:

- Event ID: A unique identifier for the event. This attribute is required.
- Event Location: The physical location of the event. This data may be made up of the street address, city, state, postal code, and country.
- Event Date: The year, month, and day of the event.
- Event Type: The type of event (e.g., reception, dinner, or seminar).
- Event Title: The title of the event, as used in the press release and communications with the former students. This attribute is required.

Attributes on Contact:

- Contact Type: The category of contact type made with the former student. Possible values include mail, email, telephone, and fax. This attribute is required.
- Contact Date: The year, month, and day of the contact event with the former student. This attribute is required.
- Contact Info: The updated information about the former student that was learned during the contact event. This attribute is required.

Relationship:

- Attends: A student attends zero, one, or many events. An event has one or more students in attendance.
- Makes: The school makes contact with a specific former student. For each contact with a specific student, the School tracks zero, one, or many instances of contact information.

Attributes on Relationship:

- Comment: Information that school officials learn from a graduate at a specific event.

Suggestions for Field Exercises

- 2-53. The intent of this exercise is to have your students gain some exposure to standards in the business world. This is a good opportunity for your students to learn the benefits of enforcing naming standards, whether for E-R models or for programming code. If standards do not exist in the organization, have your students come up with some guidelines for naming standards. If standards do exist, your students should ask the database or systems analyst for an opportunity to review these standards to see if they are consistent and uniform.
- 2-54. You may choose to use the same organizations for this field exercise that were used in Chapter 1 Field Exercises, or instead choose different organizations. It is likely that some of your students may have contacts in suitable organizations. The main difference that students are likely to find in a manufacturing company (compared to a service company) is the complexity encountered in modeling a product structure (or bill of materials). This often results in a recursive unary relationship, which is described in this chapter.
- 2-55. This field exercise can be performed in conjunction with Exercise 2-54 above. Most

organizations will probably have examples of each of these types of relationships. Be on the alert to discover ternary relationships that are mistakenly modeled as multiple binary relationships.

- 2-56. This field exercise can be combined with Exercise 2-55 above. It is quite likely the organization will be using E-R notations that are different from the text, but students should be able to accommodate different notations with some explanation.
- 2-57. We suggest you combine this with Exercise 2-56 (and perhaps Exercise 2-55) above. If time-dependent data is apparent in the models, you might ask, for example, how the organization tracks customer sales over time.
- 2-58. Students should build a table to compare features of all products.

Project Questions


2-59. *Revised Enterprise Data Model*

Please see page 55 for the revised Enterprise Data model.


2-60. *E-R diagram for FAME*

Note: Although the question asks you to base the E-R diagram on Question 1-52, students should also consider the work they did for 2-59.

Please see Page 56 for the E-R model.



Question 2-59 Figure 1



2-61. *Identify outputs*


Relevant outputs include, for example, the following items:

- A tentative contract for a specific prospective artist
- A final contract for a specific prospective artist
- Artist schedule for a specific timeframe
- List of revenues by manager within a timeframe
- List of FAME fees by manager within a timeframe
- List of revenues by artist within a timeframe
- List of revenues minus FAME fee by artist within a timeframe
- List of artist expenses by artist within a timeframe
- An account statement per artist (including previous balance, expenses since previous statement close, revenues (minus FAME fee) since previous statement close, payments to/from an artist, new balance)
- List of future performances (including anticipated revenue and FAME fee) within a specific timeframe
 - For all artists
 - For a specific artist
 - For a specific manager
- Invoices to customers for a specific time period
- Accounts receivable reports (list of unpaid invoices categorized by due date)
- List of manager expenses by manager within a timeframe
- An account statement per manager (including previous balance, expenses since previous statement close, manager's share of FAME fees since previous statement close, payments to/from a manager, new balance)
- Data regarding event cancellations
- Artist contact information
- List of relevant news items per artist
- Availability of a specific artist within a time frame
- List of all available artists for a specific time frame filtered by specific criteria (instrument, experience level, fee level, etc.)
- Artist contracts to be renewed within a specific future timeframe
- Expired artist contracts

See page 58 for the revised E-R model.


2-62. *Questions based on the E-R modeling efforts*

The questions will vary depending on the students.



Modern Database Management

Thirteenth Edition



Chapter 2

Modeling Data in the
Organization

Learning Objectives

2.1 Define terms

2.2 Understand importance of data modeling

2.3 Write good names and definitions for entities, relationships, and attributes

2.4 Distinguish unary, binary, and ternary relationships

2.5 Model different types of attributes, entities, relationships, and cardinalities

2.6 Convert many-to-many relationships to associative entities

2.7 Model time-dependent data using time stamps

E-R Model Constructs

- Entities:
 - Entity instance – person, place, object, event, concept (often corresponds to a row in a table)
 - Entity Type – collection of entities (often corresponds to a table)
- Relationships:
 - Relationship instance – link between entities (corresponds to primary key–foreign key equivalencies in related tables)
 - Relationship type – category of relationship; link between entity types
- Attributes:
 - Properties or characteristics of an entity or relationship type (often corresponds to a field in a table)

Figure 2-1 Sample E-R Diagram





Figure 2-2 Basic E-R Notation



Business Rules

- Are statements that define or constrain some aspect of the business
- Are derived from policies, procedures, events, functions
- Assert business structure
- Control/influence business behavior
- Are expressed in terms familiar to end users
- Are automated through DBMS software

A Good Business Rule Is:

- Declarative – what, not how
- Precise – clear, agreed-upon meaning
- Atomic – one statement
- Consistent – internally and externally
- Expressible – structured, natural language
- Distinct – non-redundant
- Business-oriented – understood by business people

A Good Data Name is:

- Related to business, not technical, characteristics
- Meaningful and self-documenting
- Unique
- Readable
- Composed of words from an approved list
- Repeatable
- Written in standard syntax

Data Definitions

- Explanation of a term or fact
 - Term – word or phrase with specific meaning
 - Fact – association between two or more terms
- Guidelines for good data definition
 - A concise description of essential data meaning
 - Gathered in conjunction with systems requirements
 - Accompanied by diagrams
 - Achieved by consensus and iteratively refined

Entities

- Entity
 - a person, a place, an object, an event, or a concept in the user environment about which the organization wishes to maintain data
- Entity type
 - a collection of entities that share common properties or characteristics
- Entity instance
 - a single occurrence of an entity type

Entity Type and Entity Instances

FIGURE 2-3 Entity type EMPLOYEE with two instances

Entity type: EMPLOYEE


Attributes	Attribute Data Type	Example Instance	Example Instance
Employee Number	CHAR (10)	64217836	53410197
Name	CHAR (25)	Michelle Brady	David Johnson
Address	CHAR (30)	100 Pacific Avenue	450 Redwood Drive
City	CHAR (20)	San Francisco	Redwood City
State	CHAR (2)	CA	CA
Zip Code	CHAR (9)	98173	97142
Date Hired	DATE	03-21-1992	08-16-1994
Birth Date	DATE	06-19-1968	09-04-1975

An Entity...


- **Should Be:**
 - An object that will have many instances in the database
 - An object that will be composed of multiple attributes
 - An object that we are trying to model
- **Should Not Be:**
 - A user of the database system
 - An output of the database system (e.g., a report)

Figure 2-4 Example of Inappropriate Entities

(a) System user (Treasurer) and output (Expense Report) shown as entities



b) E-R diagram with only the necessary entities




Strong vs. Weak Entities, and Identifying Relationships

- Strong entity
 - exists independently of other types of entities
 - has its own unique identifier
 - identifier underlined with single line
- Weak entity
 - dependent on a strong entity (identifying owner); cannot exist on its own
 - does not have a unique identifier (only a partial identifier)
 - entity box and partial identifier have double lines
- Identifying relationship
 - links strong entities to weak entities

Figure 2-5 Example of a Weak Identity and Its Identifying Relationship

Weak entity with identifying relationship.



Guidelines for Naming Entities

- Singular noun
- Specific to organization
- Concise, or abbreviation
- For event entities, the result not the process
- Name consistent for all diagrams

Guidelines for Defining Entities

- “An X is...”
- Describe unique characteristics of each instance
- Explicit about what is and is not the entity
- When an instance is created or destroyed
- Changes to other entity types
- History that should be kept

Attributes

- Attribute – property or characteristic of an entity or relationship type
- Classifications of attributes:
 - Required versus Optional
 - Simple versus Composite
 - Single-Valued versus Multivalued
 - Stored versus Derived
 - Identifier

Figure 2-6 Required vs. Optional Attributes

Entity type: STUDENT

Attributes	Attribute Data Type	Required or Optional	Example Instance	Example Instance
Student ID	CHAR (10)	Required	28-618411	26-844576
Student Name	CHAR (40)	Required	Michael Grant	Melissa Kraft
Home Address	CHAR (30)	Required	314 Baker St.	1422 Heft Ave
Home City	CHAR (20)	Required	Centerville	Miami
Home State	CHAR (2)	Required	OH	FL
Home Zip Code	CHAR (9)	Required	45459	33321
Major	CHAR (3)	Optional	MIS	

Required – must have a value for every entity (or relationship) instance with which it is associated

Optional – may not have a value for every entity (or relationship) instance with which it is associated

Figure 2-7 A Composite Attribute

- **Composite attribute** – An attribute that has meaningful component parts (sub-attributes)





Figure 2-8 Multivalued and Derived Attributes

- **Multivalued**

- May take on more than one value for a given entity (or relationship) instance
- An employee can have more than one skill

- **Derived**

- Values can be calculated from related attribute values (not physically stored in the database)
- Years employed calculated from date employed and current date



Identifiers (Keys)


- Identifier (Key) – an attribute (or combination of attributes) that uniquely identifies individual instances of an entity type
- Simple versus Composite Identifier
- Candidate Identifier – an attribute that could be an identifier; it satisfies the requirements for being an identifier

Criteria for Identifiers


- Choose Identifiers that
 - Will not change in value
 - Will not be null
- Avoid intelligent identifiers (e.g., containing locations or people that might change)
- Substitute new, simple keys for long, composite keys

Figure 2-9 Simple and Composite Identifier Attributes

a) Simple identifier attribute



b) Composite identifier attribute



Naming Attributes

- Name should be a singular noun or noun phrase
- Name should be unique
- Name should follow a standard format
- Similar attributes of different entity types should use the same qualifiers and classes

Defining Attributes (1 of 2)

- State what the attribute is and possibly why it is important
- Make it clear what is and is not included in the attribute's value
- Include aliases in documentation

Defining Attributes (2 of 2)

- State source of values
- State whether attribute value can change once set
- Specify whether required or optional
- State min and max number of occurrences allowed
- Indicate relationships with other attributes

Modeling Relationships


- Relationship Types vs. Relationship Instances
 - The relationship type is modeled as lines between entity types. The relationship instance is between specific entity instances
- Relationships can have attributes
 - These describe features pertaining to the association between the entities in the relationship
- Two entities can have more than one type of relationship between them (multiple relationships)
- Associative Entity – combination of relationship and entity

Figure 2-10 Relationship Type and Instances

a) Relational type
(Completes)



b) Relationship
instances



Degree of Relationships


- Degree of a relationship is the number of entity types that participate in it
 - Unary Relationship
 - Binary Relationship
 - Ternary Relationship

From Figure 2-2 Degrees of Relationships

Unary – entities of the same entity type related to each other

Binary – entities of one type related to entities of another

Ternary – entities of three different types involved in the same relationship



Cardinality of Relationships

- One-to-One
 - Each entity in the relationship will have exactly one related entity
- One-to-Many
 - An entity on one side of the relationship can have many related entities, but an entity on the other side will have a maximum of one related entity
- Many-to-Many
 - Entities on both sides of the relationship can have many related entities on the other side

Figure 2-12 Examples of Relationships of Different Degrees (1 of 3)

a) Unary relationships




Figure 2-12 Examples of Relationships of Different Degrees (2 of 3)

b) Binary relationships




Figure 2-12 Examples of Relationships of Different Degrees (3 of 3)

c) Ternary relationships






Figure 2-15 Using Relationships and Entities to Link Related Attributes

Multivalued attributes can be represented as relationships

(a) Multivalued attribute versus relationships via bill-of-materials structure



(b) Composite, multivalued attribute versus relationship



Cardinality Constraints

- Cardinality Constraints — the number of instances of one entity that can or must be associated with each instance of another entity
- Minimum Cardinality
 - If zero, then optional
 - If one or more, then mandatory
- Maximum Cardinality
 - The maximum number

Figure 2-17 Examples of Cardinality Constraints (1 of 3)

a) Mandatory cardinalities




Figure 2-17 Examples of Cardinality Constraints (2 of 3)

b) One mandatory, one optional cardinality




Figure 2-17 Examples of Cardinality Constraints (3 of 3)

c) Optional cardinalities



Figure 2-21 Example of Multiple Relationships (1 of 2)

a) Employees and departments





Figure 2-21 Example of Multiple Relationships (2 of 2)

b) Professors and courses (fixed lower limit constraint)




Associative Entities

- When should a **relationship with attributes** instead be an **associative entity**?
 - All relationships for the associative entity should be many
 - The associative entity could have meaning independent of the other entities
 - The associative entity preferably has a unique identifier, and should also have other attributes
 - The associative entity may participate in other relationships other than the entities of the associated relationship
 - Convert ternary relationships to associative entities

Figure 2-11 Associative Entities

An associative entity is an entity. It has attributes. It is also a relationship. It serves to link other entities together in a many-to-many relationship.

(a) Attribute on a relationship



(b) An associative entity (CERTIFICATE)



Figure 2-13 Representing a Bill-of-Materials Structure (1 of 2)

Two ways to represent a bill-of-materials structure

(a) Many-to-many relationship



(c) Associative entity




Figure 2-13 Representing a Bill-of-Materials Structure (2 of 2)

A bill-of-materials is a hierarchy of items. Here, bikes are composed of handle bars, transmissions, wheels, etc. Transmissions include their own sub-parts, brakes and derailleurs.

(b) Two ITEM bill-of-materials structure instances




Figure 2-18 Cardinality Constraints in a Ternary Relationship

The three relationships are marked by their mandatory or optional constraints. #1 states that a vendor may or may not supply parts to warehouses. #2 states that each part must come from a vendor and go to a warehouse. #3 says a warehouse must have at least one part from a vendor.




Figure 2-19 Simple Example of Time-Stamping

Time stamp – a time value that is associated with a data value, often indicating when some event occurred that affected the data value.

The Price History attribute is both multivalued and composite.





Figure 2-22 Data Model for Pine Valley Furniture Company in Microsoft Visio Notation

Different modeling software tools may have different notation for the same constructs.



Copyright



This work is protected by United States copyright laws and is provided solely for the use of instructors in teaching their courses and assessing student learning. Dissemination or sale of any part of this work (including on the World Wide Web) will destroy the integrity of the work and is not permitted. The work and materials from it should never be made available to students except by instructors using the accompanying text in their classes. All recipients of this work are expected to abide by these restrictions and to honor the intended pedagogical purposes and the needs of other instructors who rely on these materials.