Chapter 6: Database Design Using the E-R Model

Outline

- Overview of the Design Process

- The Entity-Relationship Model

- Complex Attributes

- Mapping Cardinalities

- Primary Key

- Extended E-R Features

- Entity-Relationship Design Issues

- Alternative Notations for Modeling Data (UML)

- Other Aspects of Database Design

Design Phases

- Initial phase -- characterize fully the data needs of the prospective database users.

- Second phase -- choosing a data model

  - Applying the concepts of the chosen data model

  - Translating these requirements into a conceptual schema of the database.

  - A fully developed conceptual schema indicates the functional requirements of the enterprise.

    - Describe the kinds of operations (or transactions) that will be performed on the data.

Design Phases (Cont.)

- Final Phase -- Moving from an abstract data model to the implementation of the database

  - Logical Design – Deciding on the database schema. Database design requires that we find a "good" collection of relation schemas.

    - Business decision – What attributes should we record in the database?

    - Computer Science decision – What relation schemas should we have and how should the attributes be distributed among the various relation schemas?

  - Physical Design – Deciding on the physical layout of the database

Design Alternatives

- In designing a database schema, we must ensure that we avoid two major pitfalls:

  - Redundancy: a bad design may result in repeat information.

    - Redundant representation of information may lead to data inconsistency among the various copies of information

  - Incompleteness: a bad design may make certain aspects of the enterprise difficult or impossible to model.

- Avoiding bad designs is not enough. There may be a large number of good designs from which we must choose.

Design Approaches

- Entity Relationship Model (covered in this chapter)

  - Models an enterprise as a collection of *entities* and *relationships*

    - Entity: a "thing" or "object" in the enterprise that is distinguishable from other objects

      - Described by a set of *attributes*

    - Relationship: an association among several entities

  - Represented diagrammatically by an *entity-relationship diagram:*

- Normalization Theory (Chapter 7)

  - Formalize what designs are bad, and test for them

Outline of the ER Model

ER model -- Database Modeling

- The ER data mode was developed to facilitate database design by allowing specification of an enterprise schema that represents the overall logical structure of a database.

- The ER data model employs three basic concepts:

  - entity sets,

  - relationship sets,

  - attributes.

- The ER model also has an associated diagrammatic representation, the **ER diagram**, which can express the overall logical structure of a database graphically.

Entity Sets

- An **entity** is an object that exists and is distinguishable from other objects.

  - Example: specific person, company, event, plant

- An **entity set** is a set of entities of the same type that share the same properties.

  - Example: set of all persons, companies, trees, holidays

- An entity is represented by a set of attributes; i.e., descriptive properties possessed by all members of an entity set.

  - Example:

    instructor = (ID, name, salary)
    course= (course_id, title, credits)

- A subset of the attributes form a **primary key** of the entity set; i.e., uniquely identifying each member of the set.

Entity Sets -- *instructor* and *student*

76766	Crick
45565	Katz
10101	Srinivasan
98345	Kim
76543	Singh
22222	Einstein

instructor

98988	Tanaka
12345	Shankar
00128	Zhang
76543	Brown
76653	Aoi
23121	Chavez
44553	Peltier

student

Representing Entity sets in ER Diagram

- Entity sets can be represented graphically as follows:
    - Rectangles represent entity sets.
    - Attributes listed inside entity rectangle
    - Underline indicates primary key attributes

instructor
ID
name
salary

student
ID
name
tot_cred
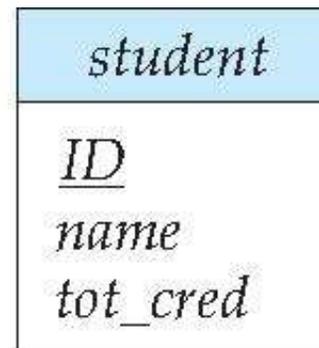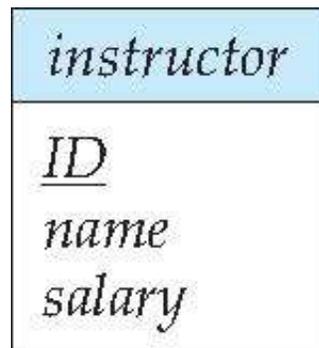
Relationship Sets

- A **relationship** is an association among several entities

  Example:

  44553 (Peltier)   *advisor*   22222 (Einstein)
  student entity   relationship set   *instructor* entity

- A **relationship set** is a mathematical relation among $n \geq 2$ entities, each taken from entity sets

$$\{(e_1, e_2, \dots e_n) \mid e_1 \in E_1, e_2 \in E_2, \dots, e_n \in E_n\}$$

where $(e_1, e_2, \dots, e_n)$ is a relationship

- Example:

$$(44553, 22222) \in advisor$$

Relationship Sets (Cont.)

- Example: we define the relationship set *advisor* to denote the associations between students and the instructors who act as their advisors.

- Pictorially, we draw a line between related entities.

	instructor			student	
76766	Crick		98988	Tanaka	
45565	Katz		12345	Shankar	
10101	Srinivasan		00128	Zhang	
98345	Kim		76543	Brown	
76543	Singh		76653	Aoi	
22222	Einstein		23121	Chavez	
			44553	Peltier	

instructor

student
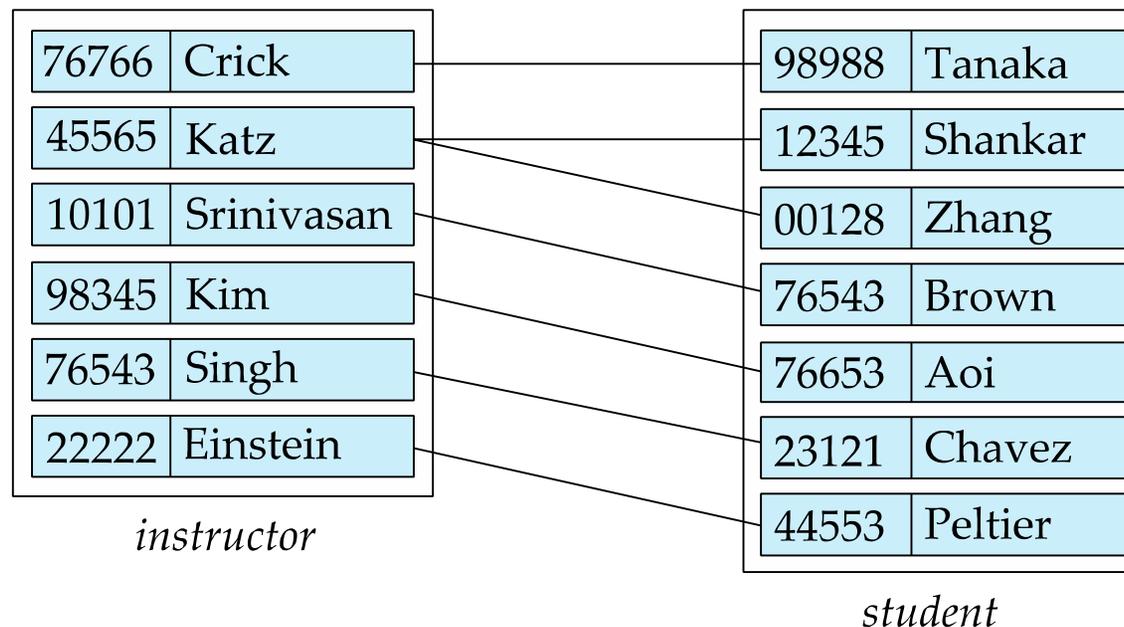
Representing Relationship Sets via ER Diagrams

- Diamonds represent relationship sets.

instructor
ID
name
salary

advisor

student
ID
name
tot_cred
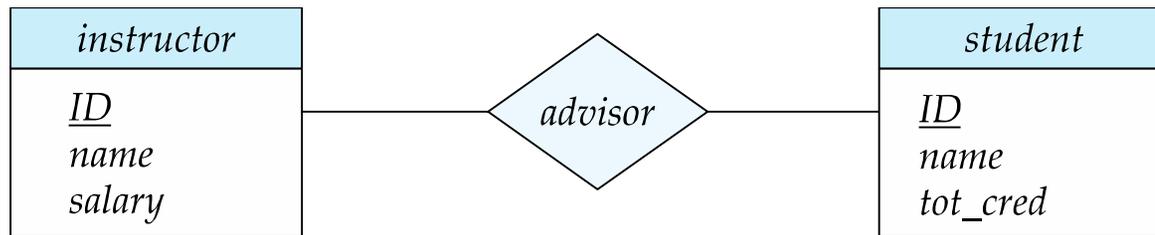
Relationship Sets (Cont.)

- An attribute can also be associated with a relationship set.

- For instance, the *advisor* relationship set between entity sets *instructor* and *student* may have the attribute *date* which tracks when the student started being associated with the advisor

76766	Crick		
45565	Katz	3 May 2008	
10101	Srinivasan	10 June 2007	
98345	Kim	12 June 2006	
76543	Singh	6 June 2009	
22222	Einstein	30 June 2007	
		31 May 2007	
		4 May 2006	

instructor

98988	Tanaka
12345	Shankar
00128	Zhang
76543	Brown
76653	Aoi
23121	Chavez
44553	Peltier

student
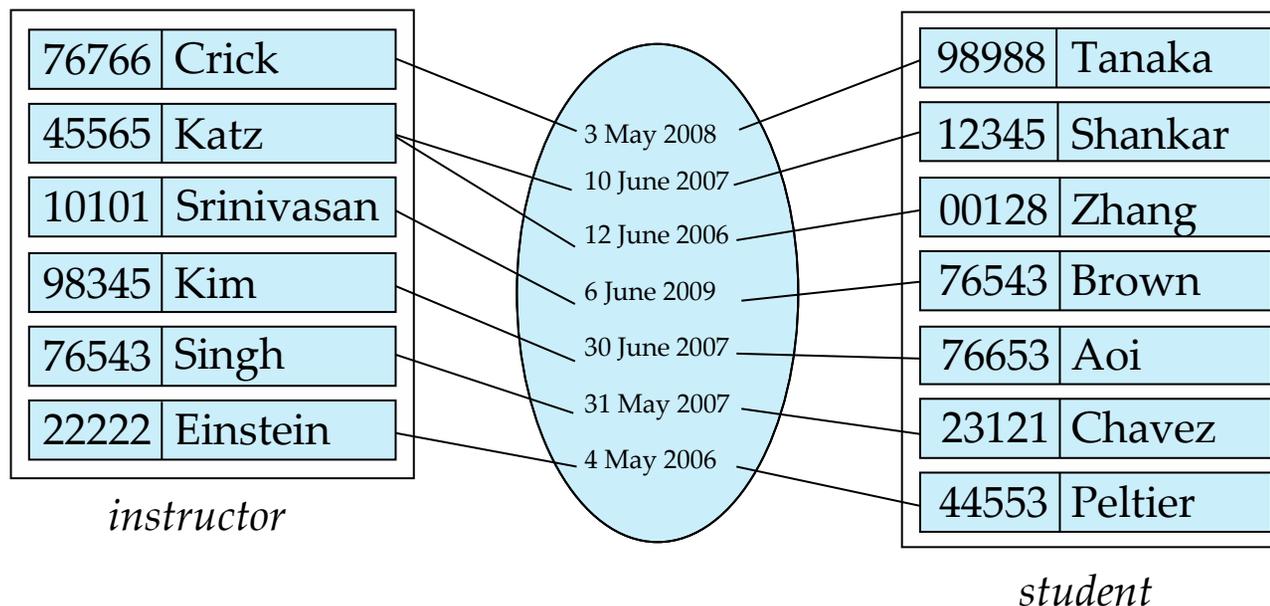
Relationship Sets with Attributes
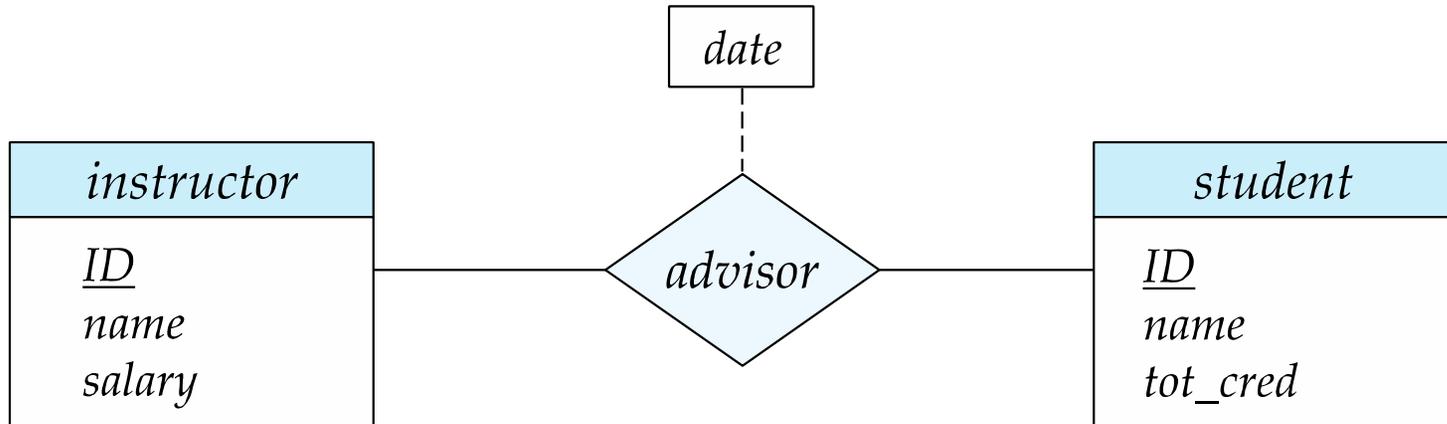
Roles

- Entity sets of a relationship need not be distinct

  - Each occurrence of an entity set plays a "role" in the relationship

- The labels "*course_id*" and "*prereq_id*" are called **roles**.
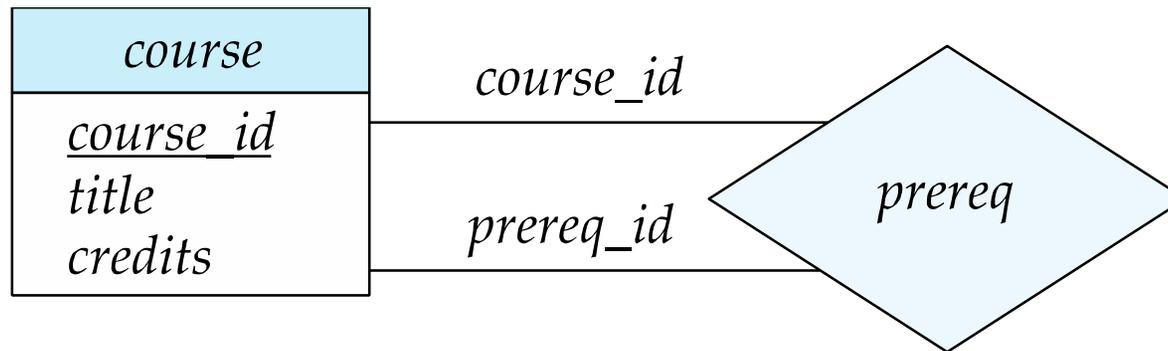
Degree of a Relationship Set

- Binary relationship

  - involve two entity sets (or degree two).

  - most relationship sets in a database system are binary.

- Relationships between more than two entity sets are rare. Most relationships are binary. (More on this later.)

  - Example: *students* work on research *projects* under the guidance of an *instructor*.

  - relationship *proj_guide* is a ternary relationship between *instructor, student,* and *project*

Non-binary Relationship Sets

- Most relationship sets are binary

- There are occasions when it is more convenient to represent relationships as non-binary.

- E-R Diagram with a Ternary Relationship
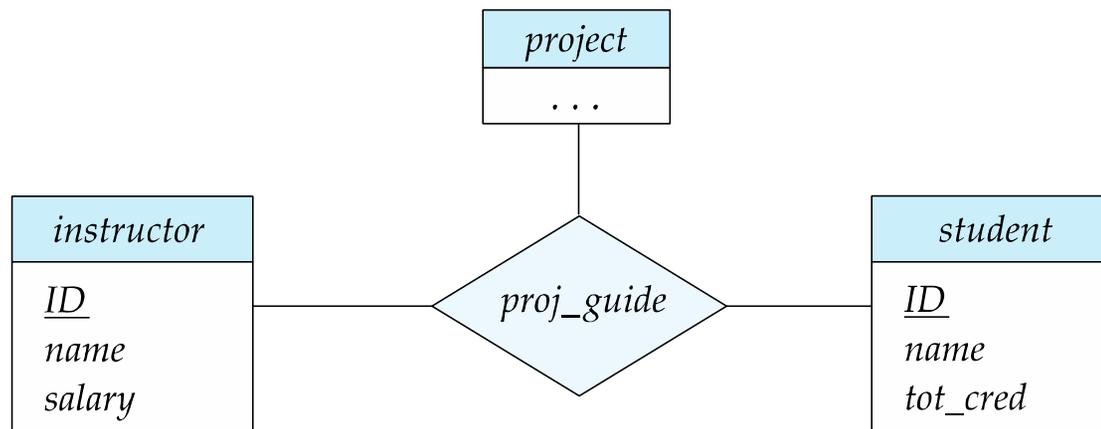
Complex Attributes

- Attribute types:

  - **Simple** and **composite** attributes.

  - **Single-valued** and **multivalued** attributes

    - Example: multivalued attribute: *phone_numbers*

  - **Derived** attributes

    - Can be computed from other attributes

    - Example: age, given date_of_birth

- **Domain** – the set of permitted values for each attribute

Composite Attributes

- Composite attributes allow us to divided attributes into subparts (other attributes).

composite
attributes

name

first_name *middle_initial* *last_name*

component
attributes

address

street *city* *state* *postal_code*

street_number *street_name* *apartment_number*
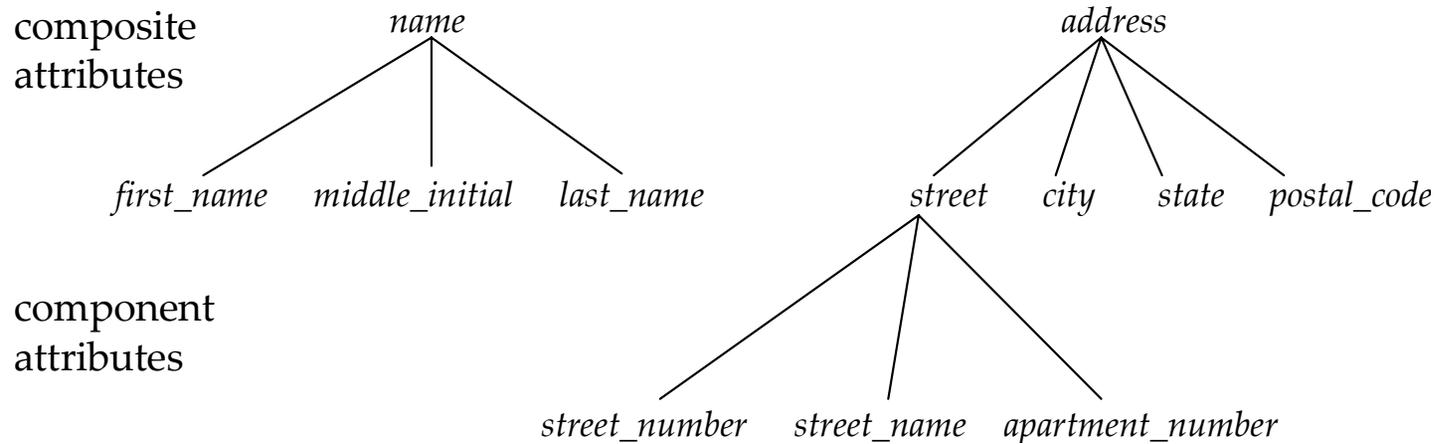
Representing Complex Attributes in ER Diagram

instructor
<u>ID</u>
name
first_name
middle_initial
last_name
address
street
street_number
street_name
apt_number
city
state
zip
{ phone_number }
date_of_birth
age ( )
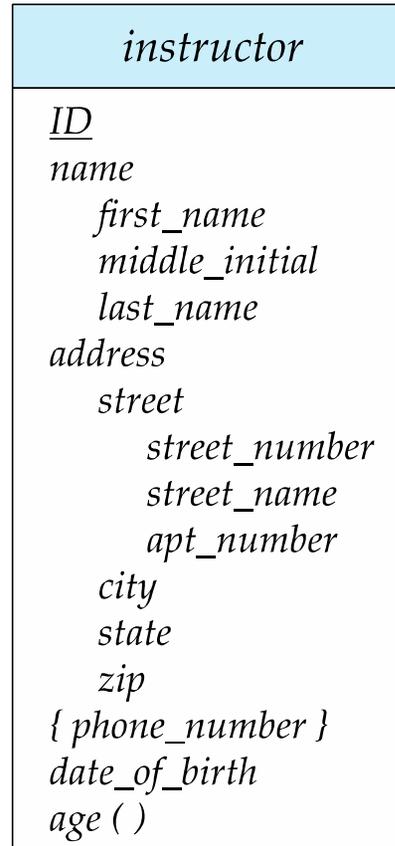
Mapping Cardinality Constraints

- Express the number of entities to which another entity can be associated via a relationship set.

- Most useful in describing binary relationship sets.

- For a binary relationship set the mapping cardinality must be one of the following types:
  - One to one
  - One to many
  - Many to one
  - Many to many

Mapping Cardinalities
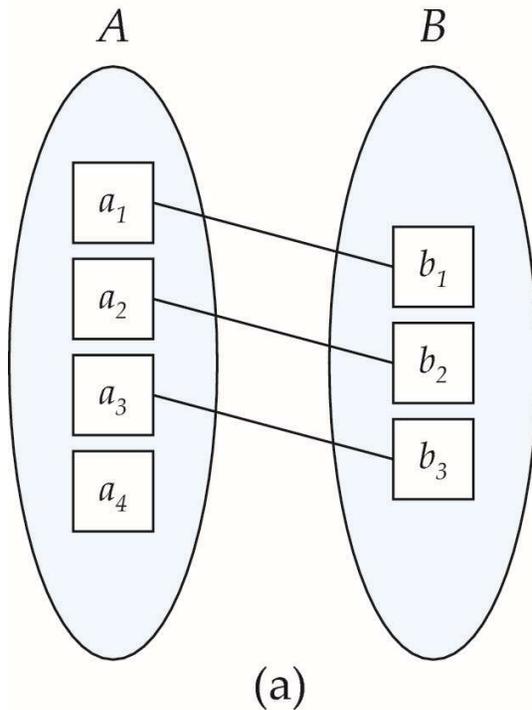
One to one
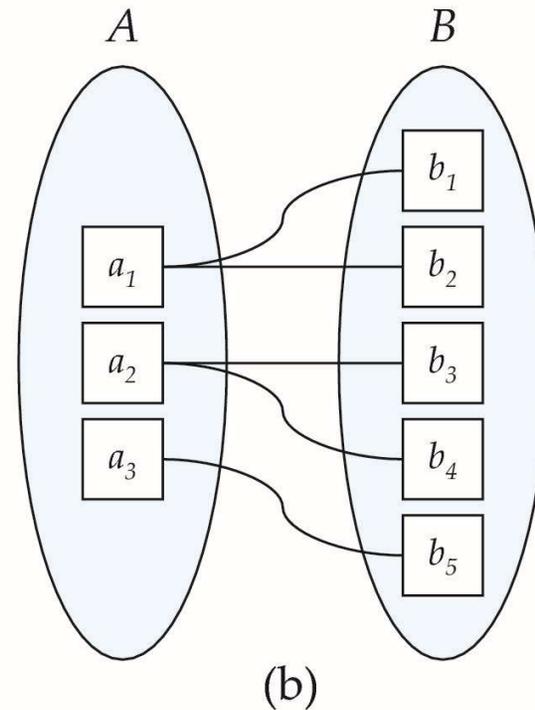
One to many

Note: Some elements in *A* and *B* may not be mapped to any elements in the other set

Mapping Cardinalities
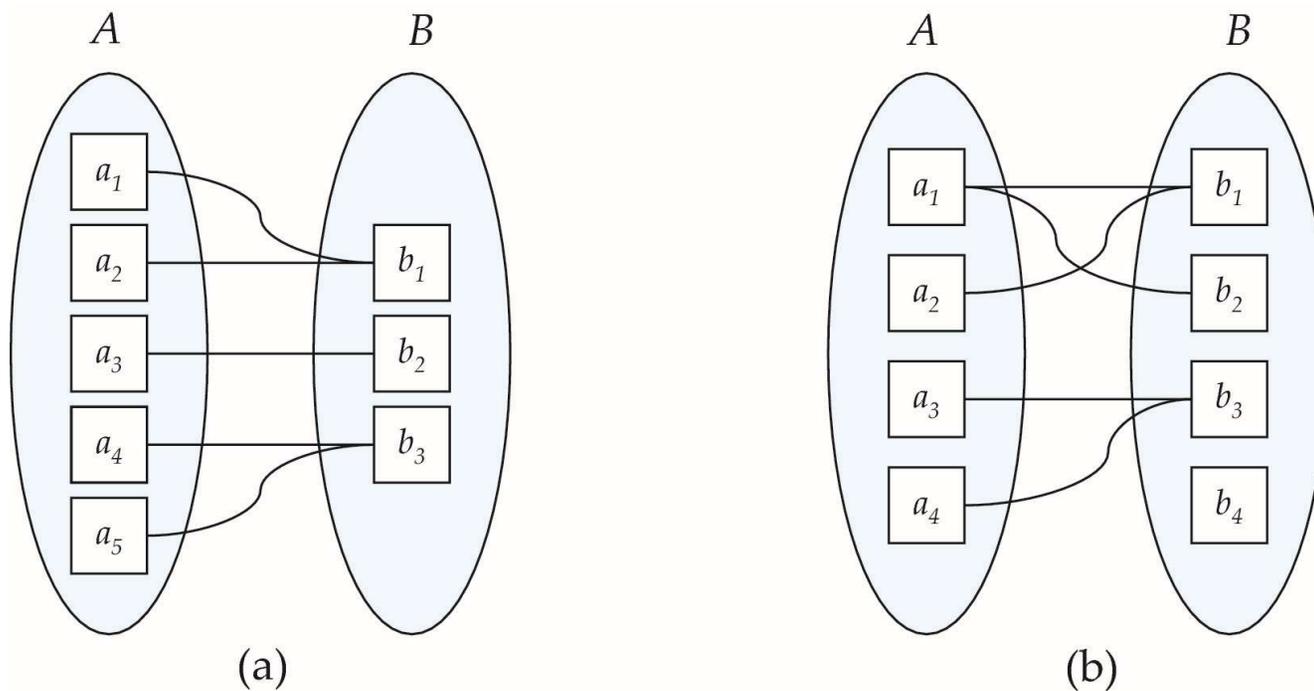
Many to one

Many to many

Note: Some elements in A and B may not be mapped to any elements in the other set

Representing Cardinality Constraints in ER Diagram

- We express cardinality constraints by drawing either a directed line (→), signifying "one," or an undirected line (—), signifying "many," between the relationship set and the entity set.

- One-to-one relationship between an *instructor* and a *student* :
  - A student is associated with at most one *instructor* via the relationship *advisor*
  - A *student* is associated with at most one *department* via *stud_dept*
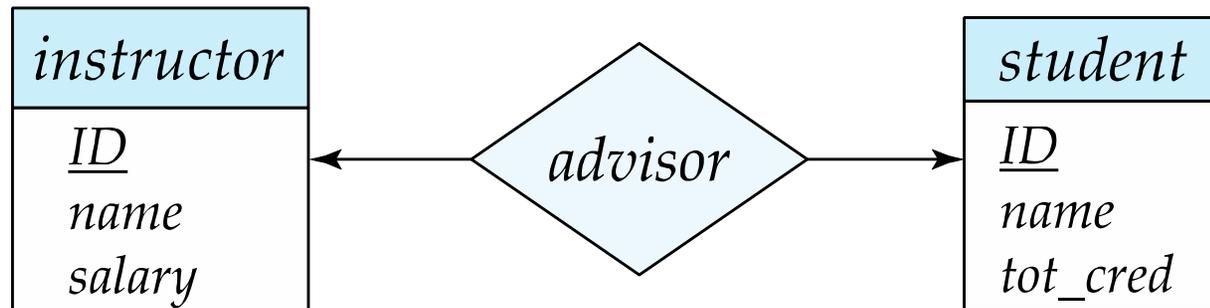
One-to-Many Relationship

- one-to-many relationship between an *instructor* and a *student*

  - an instructor is associated with several (including 0) students via *advisor*

  - a student is associated with at most one instructor via advisor,
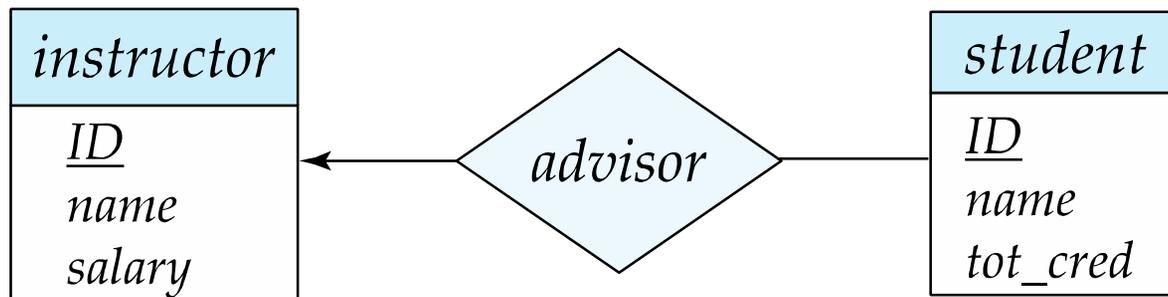
Many-to-One Relationships

- In a many-to-one relationship between an *instructor* and a *student,*

    - an instructor is associated with at most one student via *advisor*,

    - and a student is associated with several (including 0) instructors via *advisor*

instructor
$\underline{ID}$ *name* *salary*

advisor

student
$\underline{ID}$ *name* *tot_cred*
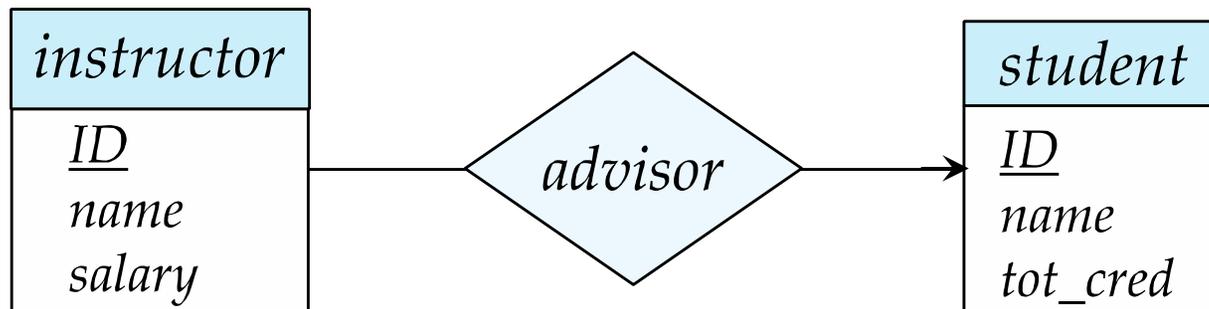
Many-to-Many Relationship

- An instructor is associated with several (possibly 0) students via *advisor*

- A student is associated with several (possibly 0) instructors via *advisor*
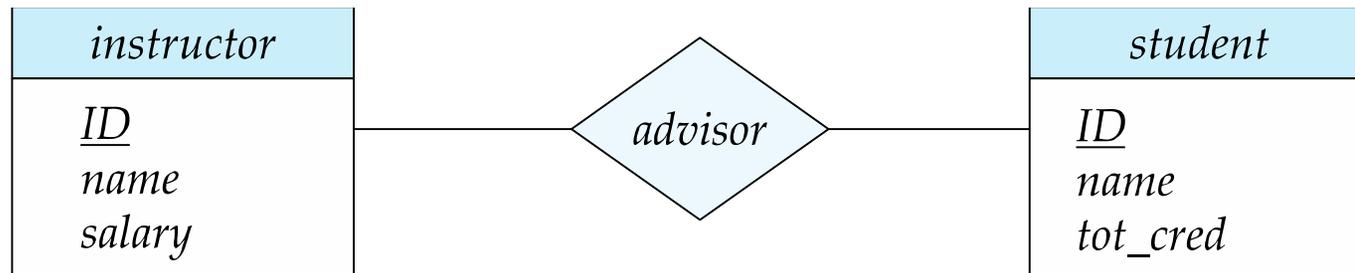
Total and Partial Participation

- **Total participation** (indicated by double line): every entity in the entity set participates in at least one relationship in the relationship set
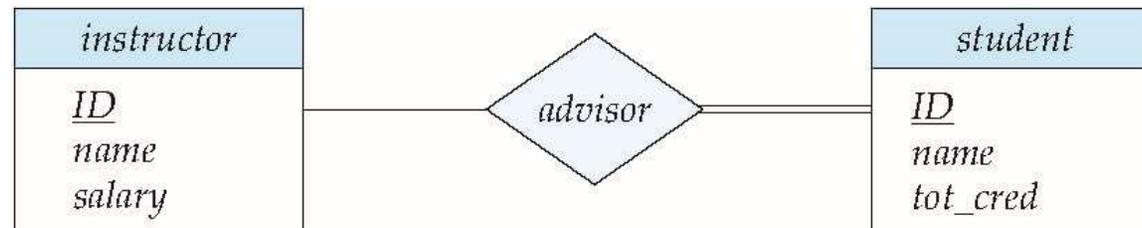
    participation of *student* in *advisor r*elation is total

- every *student* must have an associated instructor

- **Partial participation**: some entities may not participate in any relationship in the relationship set

  - Example: participation of *instructor* in *advisor* is partial

Notation for Expressing More Complex Constraints

- A line may have an associated minimum and maximum cardinality, shown in the form *l..h*, where *l* is the minimum and *h* the maximum cardinality

  - A minimum value of 1 indicates total participation.

  - A maximum value of 1 indicates that the entity participates in at most one relationship

  - A maximum value of * indicates no limit.

instructor		student
ID *name* *salary*	*0..* advisor 1..1*	*ID* *name* *tot_cred*
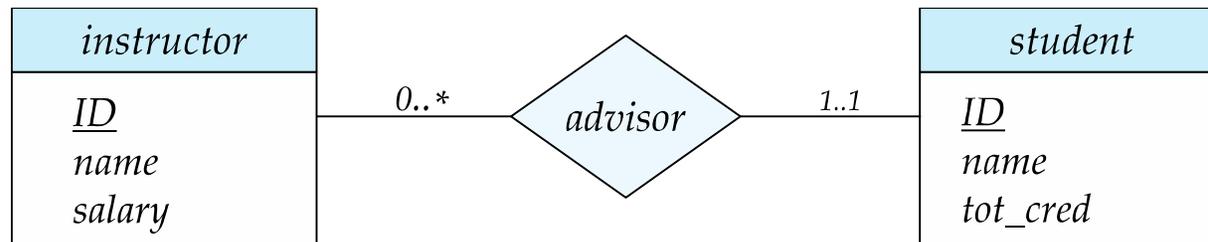
Instructor can advise 0 or more students. A student must have 1 advisor; cannot have multiple advisors

Cardinality Constraints on Ternary Relationship

- We allow at most one arrow out of a ternary (or greater degree) relationship to indicate a cardinality constraint

- For example, an arrow from *proj_guide* to *instructor* indicates each student has at most one guide for a project

- If there is more than one arrow, there are two ways of defining the meaning.

  - For example, a ternary relationship R between A, B and C with arrows to B and C could mean

    1. Each A entity is associated with a unique entity from B and C or

    2. Each pair of entities from (A, B) is associated with a unique C entity, and each pair (A, C) is associated with a unique B

  - Each alternative has been used in different formalisms

  - To avoid confusion we outlaw more than one arrow

Primary Key

- Primary keys provide a way to specify how entities and relations are distinguished. We will consider:

    - Entity sets

    - Relationship sets.

    - Weak entity sets

Primary key for Entity Sets

- By definition, individual entities are distinct.

- From database perspective, the differences among them must be expressed in terms of their attributes.

- The values of the attribute values of an entity must be such that they can uniquely identify the entity.

  - No two entities in an entity set are allowed to have exactly the same value for all attributes.

- A key for an entity is a set of attributes that suffice to distinguish entities from each other

Primary Key for Relationship Sets

- To distinguish among the various relationships of a relationship set we use the individual primary keys of the entities in the relationship set.

  - Let R be a relationship set involving entity sets E1, E2, .. En

  - The primary key for R is consists of the union of the primary keys of entity sets E1, E2, ..En

  - If the relationship set R has attributes a1, a2, .., am associated with it, then the primary key of R also includes the attributes a1, a2, .., am

- Example: relationship set "advisor".

  - The primary key consists of *inrsructor.ID* and *student.ID*

- The choice of the primary key for a relationship set depends on the mapping cardinality of the relationship set.

Choice of Primary key for Binary Relationship

- Many-to-Many relationships. The preceding union of the primary keys is a minimal superkey and is chosen as the primary key.

- One-to-Many relationships . The primary key of the "Many" side is a minimal superkey and is used as the primary key.

- Many-to-one relationships. The primary key of the "Many" side is a minimal superkey and is used as the primary key.

- One-to-one relationships. The primary key of either one of the participating entity sets forms a minimal superkey, and either one can be chosen as the primary key.

Choice of Primary key for Nonbinary Relationship

- If no cardinality constraints are present, the superkey is formed as described earlier. and it is chosen as the primary key.

- If there are cardinality constraints are present:

  - Recall that we permit at most one arrow out of a relationship set.

  - AVI

Weak Entity Sets

- Consider a *section* entity, which is uniquely identified by a *course_id*, *semester, year*, and *sec_id*.

- Clearly, section entities are related to course entities. Suppose we create a relationship set *sec_course* between entity sets *section* and *course*.

- Note that the information in *sec_course* is redundant, since *section* already has an attribute *course_id*, which identifies the course with which the section is related.

- One option to deal with this redundancy is to get rid of the relationship s*ec_course*; however, by doing so the relationship between *section* and *course* becomes implicit in an attribute, which is not desirable.

Weak Entity Sets (Cont.)

- An alternative way to deal with this redundancy is to not store the attribute *course_id* in the *section* entity and to only store the remaining attributes *section_id*, *year*, and *semester.*

  - However, the entity set *section* then does not have enough attributes to identify a particular *section* entity uniquely

- To deal with this problem, we treat the relationship *sec_course* as a special relationship that provides extra information, in this case, the *course_id*, required to identify *section* entities uniquely.

- A **weak entity set** is one whose existence is dependent on another entity, called its **identifying entity**

- Instead of associating a primary key with a weak entity, we use the identifying entity, along with extra attributes called **discriminator** to uniquely identify a weak entity.

Weak Entity Sets (Cont.)

- An entity set that is not a weak entity set is termed a **strong entity set**.

- Every weak entity must be associated with an identifying entity; that is, the weak entity set is said to be **existence dependent** on the identifying entity set.

- The identifying entity set is said to **own** the weak entity set that it identifies.

- The relationship associating the weak entity set with the identifying entity set is called the **identifying relationship**.

- Note that the relational schema we eventually create from the entity set *section* does have the attribute *course_id*, for reasons that will become clear later, even though we have dropped the attribute *course_id* from the entity set *section*.

Expressing Weak Entity Sets

- In E-R diagrams, a weak entity set is depicted via a double rectangle.

- We underline the discriminator of a weak entity set with a dashed line.

- The relationship set connecting the weak entity set to the identifying strong entity set is depicted by a double diamond.

- Primary key for *section* – (*course_id, sec_id, semester, year*)
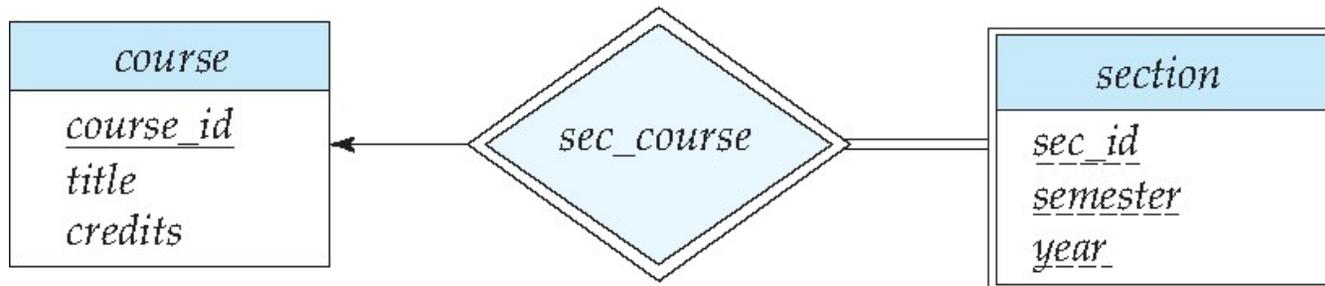
Redundant Attributes

- Suppose we have entity sets:
  - *instructor*, with attributes: *ID*, *name*, *dept_name, salary*
  - *department,* with attributes: *dept_name, building, budget*
- We model the fact that each instructor has an associated department using a relationship set *inst_dept*
- The attribute *dept_name* in *instructor* replicates information present in the relationship and is therefore redundant
  - and needs to be removed.
- BUT: when converting back to tables, in some cases the attribute gets reintroduced, as we will see later.

instructor	inst_dept	department
ID name dept_name salary		dept_name building budget
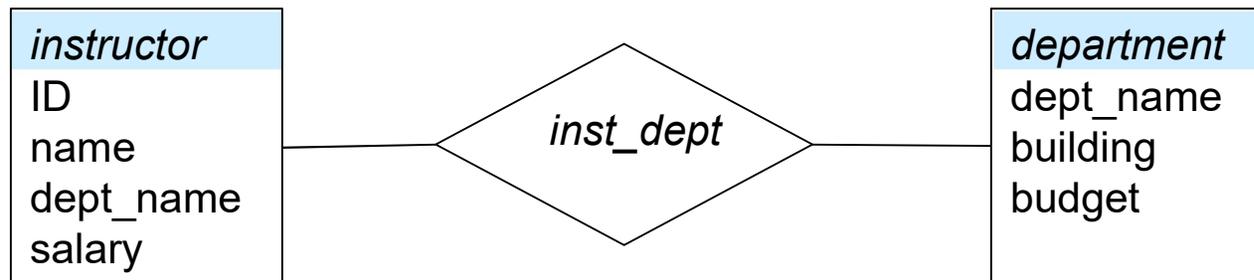
E-R Diagram for a University Enterprise
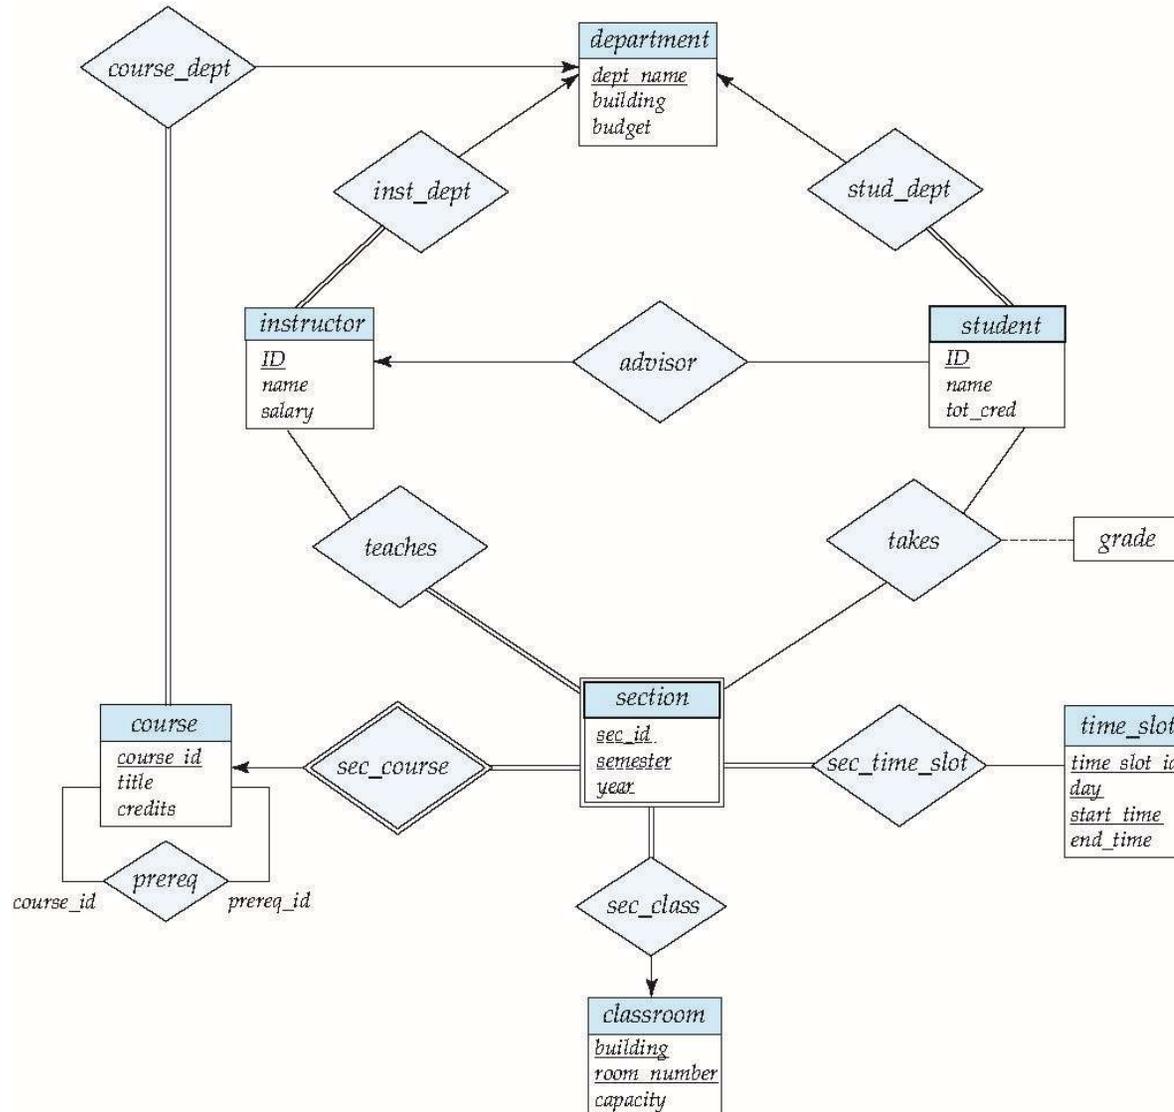
Extended E-R Features

Specialization

- Top-down design process; we designate sub-groupings within an entity set that are distinctive from other entities in the set.

- These sub-groupings become lower-level entity sets that have attributes or participate in relationships that do not apply to the higher-level entity set.

- Depicted by a *triangle* component labeled ISA (e.g., *instructor* "is a" *person*).

- **Attribute inheritance** – a lower-level entity set inherits all the attributes and relationship participation of the higher-level entity set to which it is linked.

Specialization Example

- **Overlapping** – *employee* and *student*
- **Disjoint** – *instructor* and *secretary*
- Total and partial
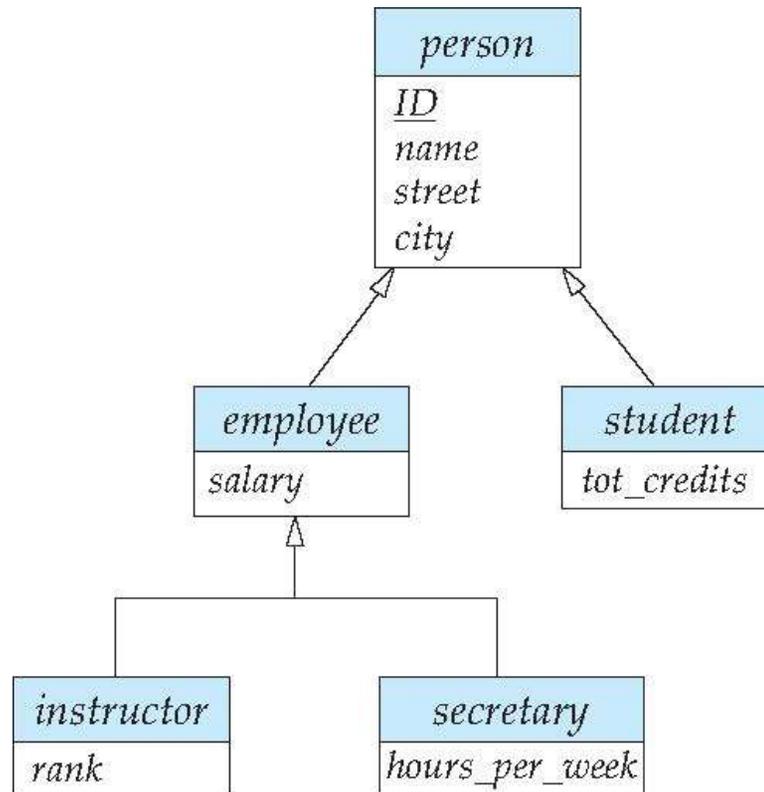
Generalization

- **A bottom-up design process** – combine a number of entity sets that share the same features into a higher-level entity set.

- Specialization and generalization are simple inversions of each other; they are represented in an E-R diagram in the same way.

- The terms specialization and generalization are used interchangeably.

Aggregation

- Consider the ternary relationship *proj_guide*, which we saw earlier

- Suppose we want to record evaluations of a student by a guide on a project
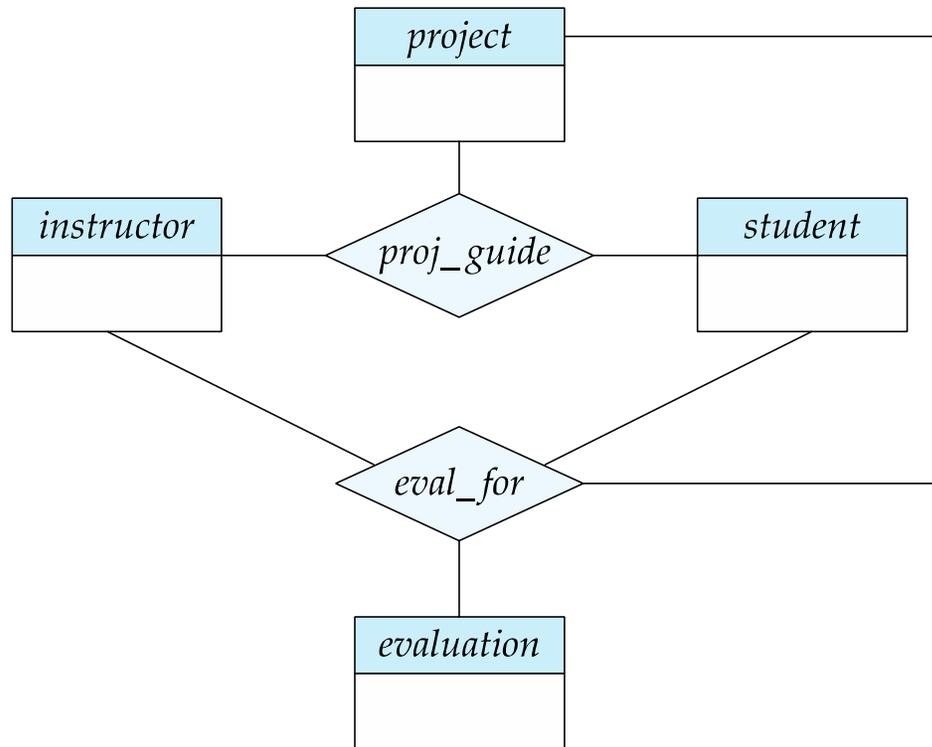
Aggregation (Cont.)

- Relationship sets *eval_for* and *proj_guide* represent overlapping information

  - Every *eval_for* relationship corresponds to a *proj_guide* relationship

  - However, some *proj_guide* relationships may not correspond to any *eval_for* relationships

    - So we can't discard the *proj_guide* relationship

- Eliminate this redundancy via *aggregation*

  - Treat relationship as an abstract entity

  - Allows relationships between relationships

  - Abstraction of relationship into new entity

Aggregation (Cont.)

- Eliminate this redundancy via *aggregation* without introducing redundancy, the following diagram represents:

  - A student is guided by a particular instructor on a particular project

  - A student, instructor, project combination may have an associated evaluation
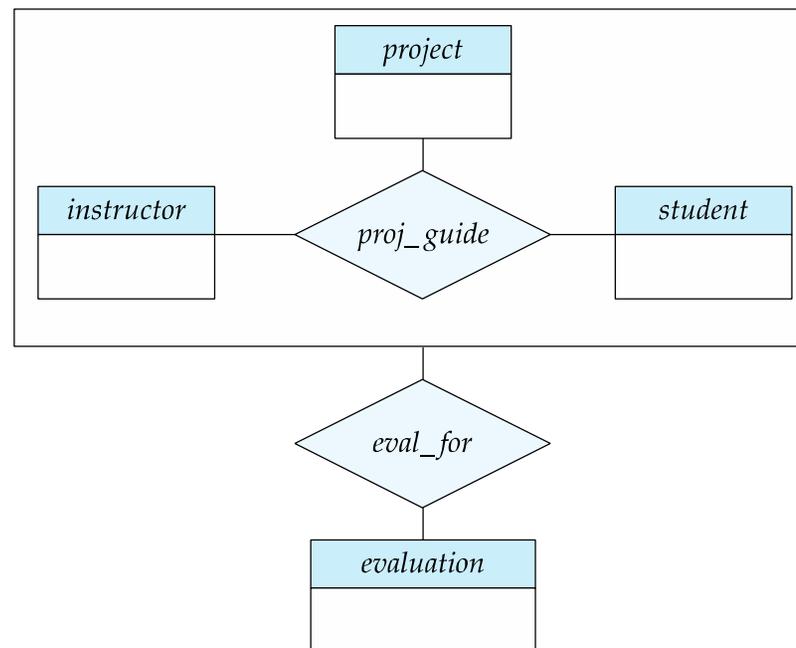
Design Issues

Common Mistakes in E-R Diagrams

- Example of erroneous E-R diagrams
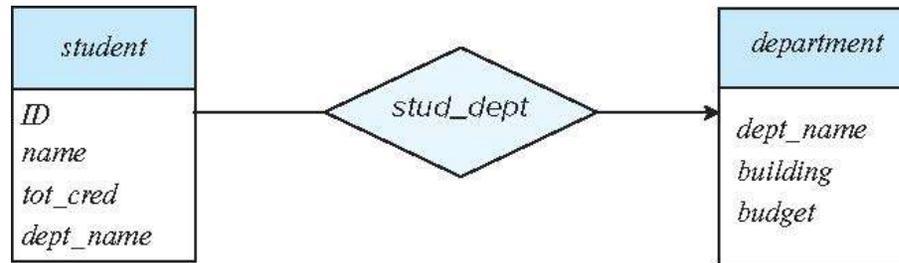
(a) Incorrect use of attribute
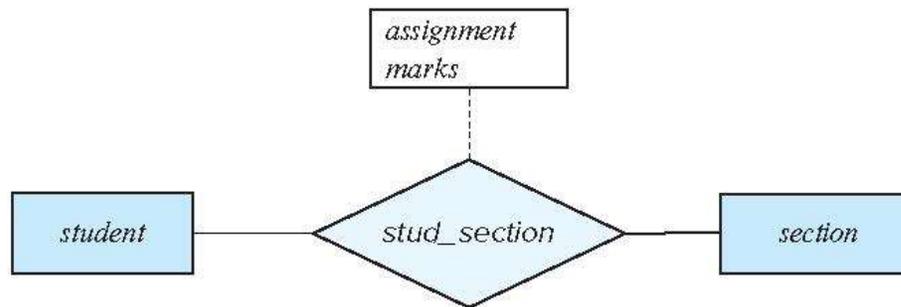
(b) Erroneous use of relationship attributes

Common Mistakes in E-R Diagrams (Cont.)

- Correct versions of the E-R diagram of previous slide
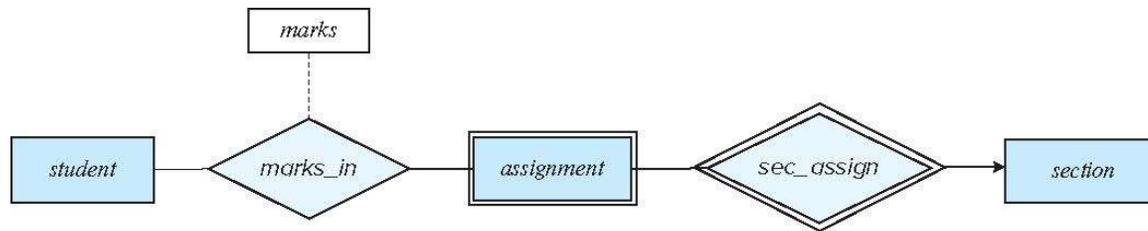
(c) Correct alternative to erroneous E-R diagram (b)
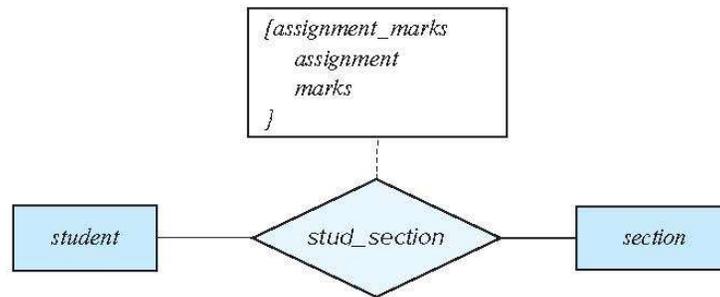
(d) Correct alternative to erroneous E-R diagram (b)

Entities vs. Attributes

- Use of entity sets vs. attributes

instructor
ID
name
salary
phone_number

instructor
ID
name
salary

inst_phone

phone
phone_number
location
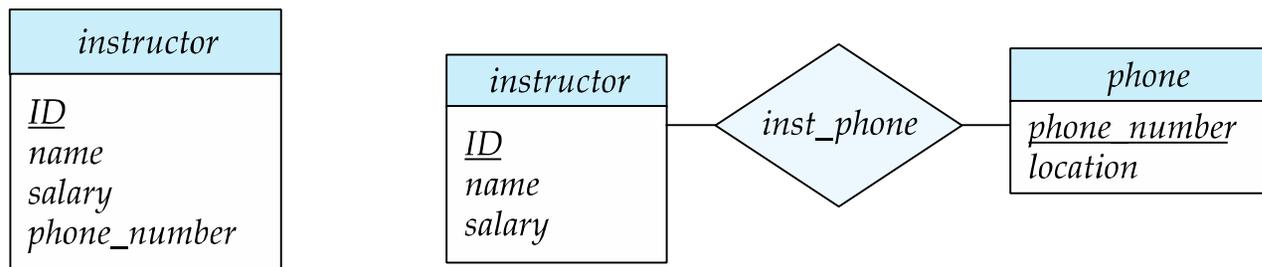
- Use of phone as an entity allows extra information about phone numbers (plus multiple phone numbers)

Entities vs. Relationship sets

- **Use of entity sets vs. relationship sets**

  Possible guideline is to designate a relationship set to describe an action that occurs between entities
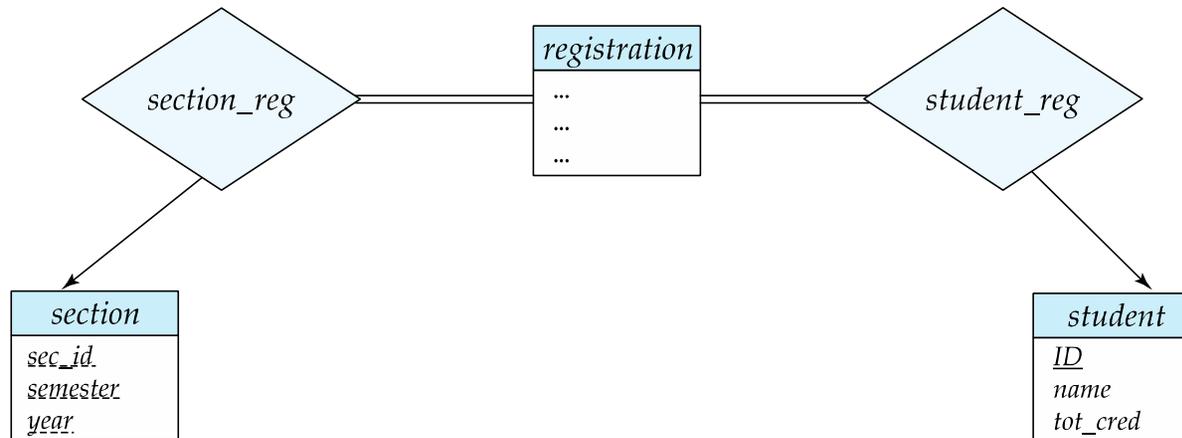
- **Placement of relationship attributes**

  For example, attribute date as attribute of advisor or as attribute of student

Binary Vs. Non-Binary Relationships

- Although it is possible to replace any non-binary (*n*-ary, for *n* > 2) relationship set by a number of distinct binary relationship sets, a *n*-ary relationship set shows more clearly that several entities participate in a single relationship.

- Some relationships that appear to be non-binary may be better represented using binary relationships

  - For example, a ternary relationship *parents*, relating a child to his/her father and mother, is best replaced by two binary relationships, *father* and *mother*

    - Using two binary relationships allows partial information (e.g., only mother being known)

  - But there are some relationships that are naturally non-binary

    - Example: *proj_guide*

E-R Design Decisions

- The use of an attribute or entity set to represent an object.

- Whether a real-world concept is best expressed by an entity set or a relationship set.

- The use of a ternary relationship versus a pair of binary relationships.

- The use of a strong or weak entity set.

- The use of specialization/generalization – contributes to modularity in the design.

- The use of aggregation – can treat the aggregate entity set as a single unit without concern for the details of its internal structure.

Summary of Symbols Used in E-R Notation

E   entity set

R   relationship set

R   identifying relationship set for weak entity set

R — E   total participation of entity set in relationship

E
A1
A2
A2.1
A2.2
{A3}
A4()

attributes:
simple (A1),
composite (A2) and
multivalued (A3)
derived (A4)

E
<u>A1</u>

primary key

E
A1

discriminating attribute of weak entity set
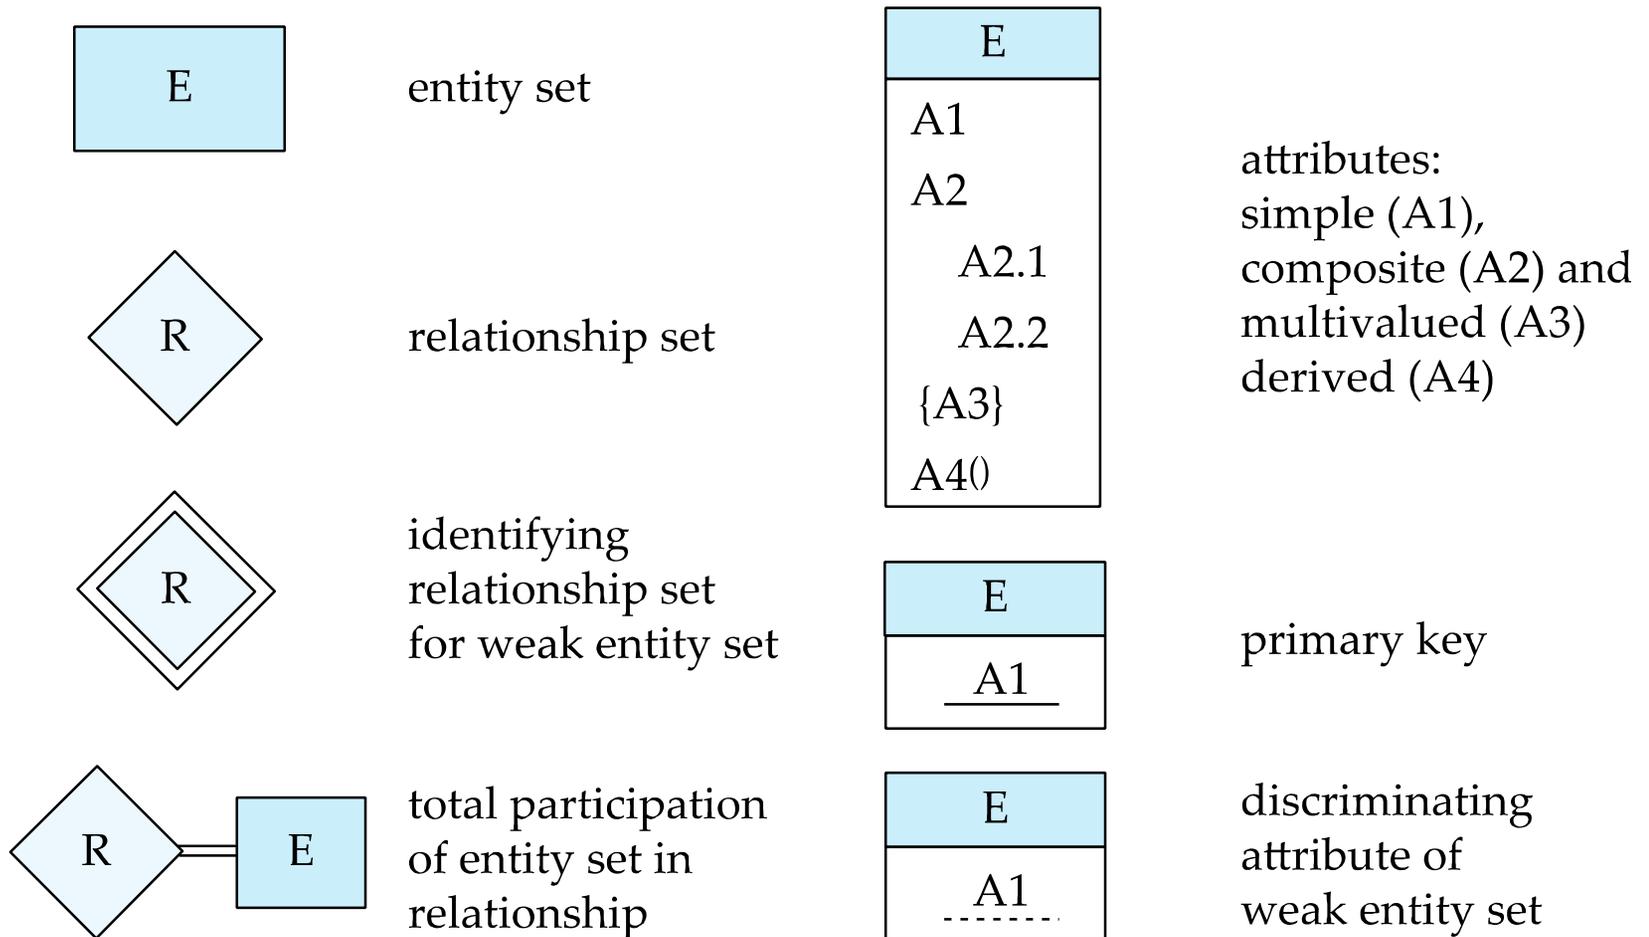
Symbols Used in E-R Notation (Cont.)
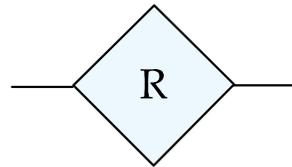
many-to-many relationship
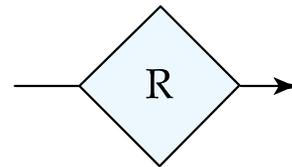
many-to-one relationship
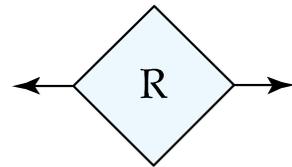
one-to-one relationship
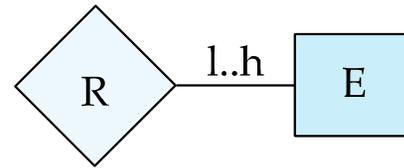
cardinality limits
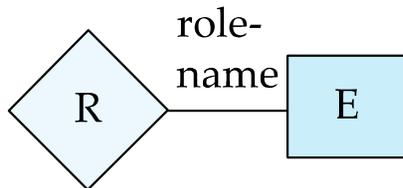
role indicator
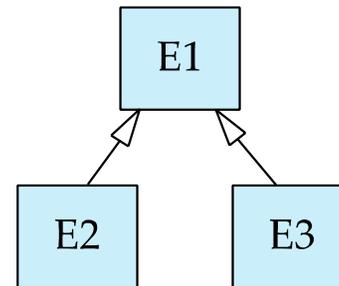
ISA: generalization or specialization
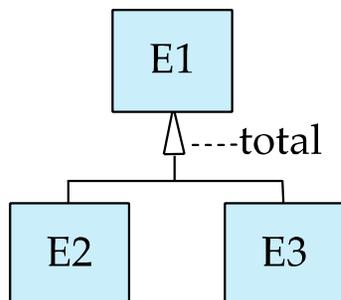
total (disjoint) generalization
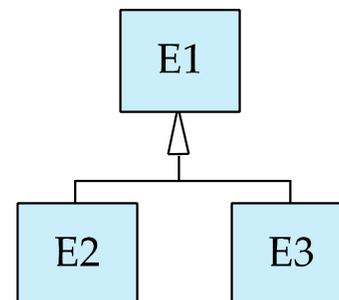
disjoint generalization

Alternative ER Notations

- Chen, IDE1FX, …

entity set E with
simple attribute A1,
composite attribute A2,
multivalued attribute A3,
derived attribute A4,
and primary key A1
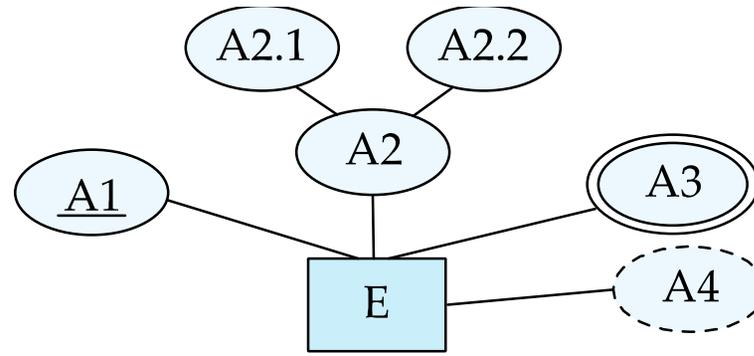
weak entity set   generalization   ISA   total generalization   ISA
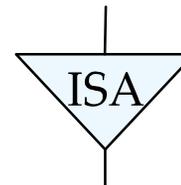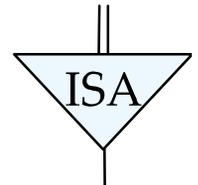
Alternative ER Notations

Chen **IDE1FX (Crows feet notation)**

many-to-many relationship
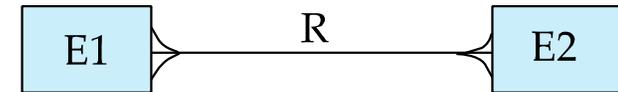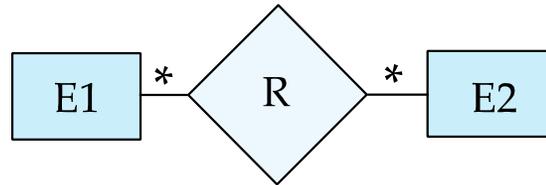
one-to-one relationship
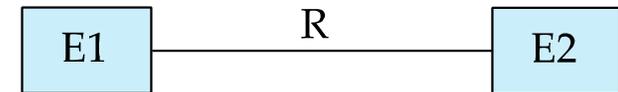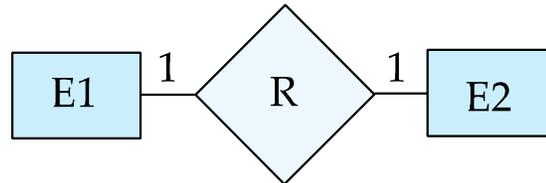
many-to-one relationship
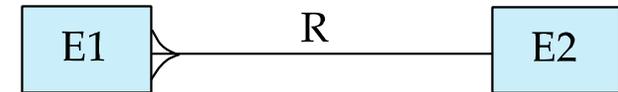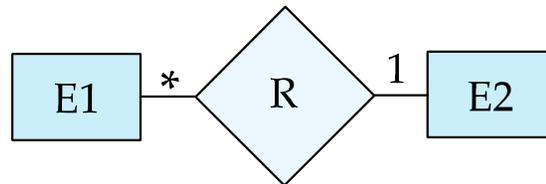
participation in R: total (E1) and partial (E2)
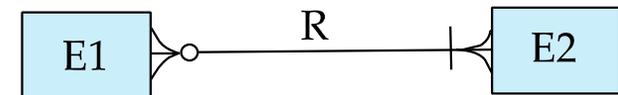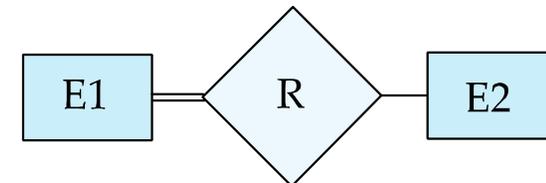
UML

- **UML**: Unified Modeling Language

- UML has many components to graphically model different aspects of an entire software system

- UML Class Diagrams correspond to E-R Diagram, but several differences.

ER vs. UML Class Diagrams

ER Diagram Notation

E
A1
M1()

entity with attributes (simple, composite, multivalued, derived)

Equivalent in UML

E
–A1
+M1()

class with simple attributes and methods (attribute prefixes: + = public, – = private, # = protected)

E1 —role1— < R > —role2— E2   binary relationship

E1 —role1 R role2— E2

A1

E1 —role1— < R > —role2— E2   relationship attributes

R
A1

E1 —role1 : role2— E2

E1 —0..*— < R > —0..1— E2   cardinality constraints

E1 —0..1 R 0..*— E2
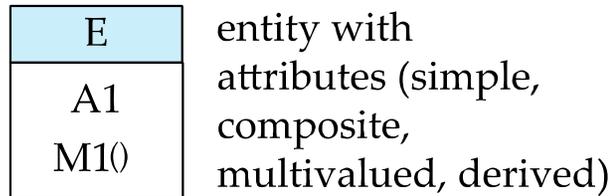
*Note reversal of position in cardinality constraint depiction

ER vs. UML Class Diagrams

ER Diagram Notation

Equivalent in UML

n-ary relationships

overlapping generalization

overlapping
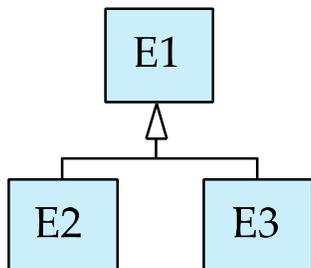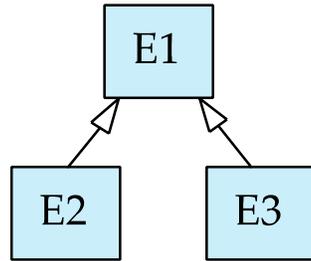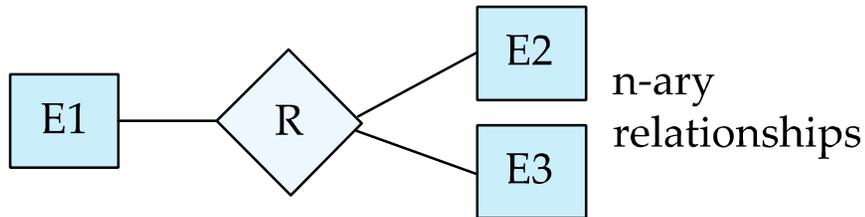
disjoint generalization
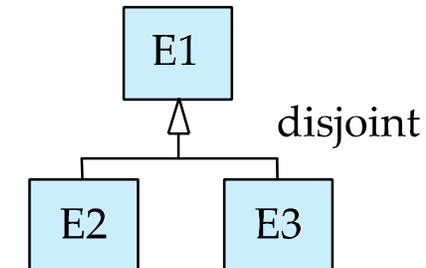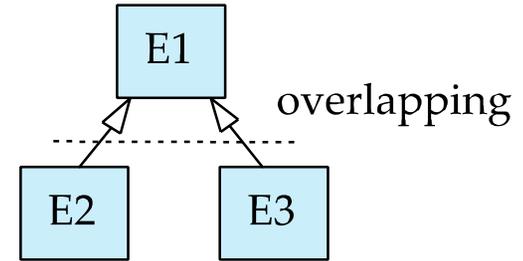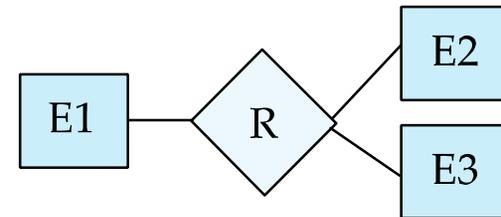
disjoint

*Generalization can use merged or separate arrows independent of disjoint/overlapping

UML Class Diagrams (Cont.)

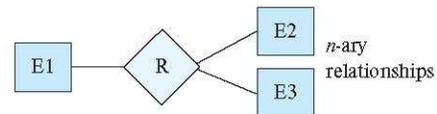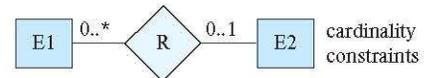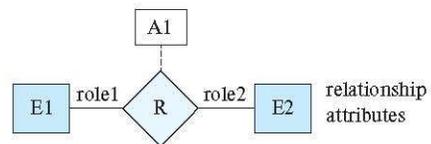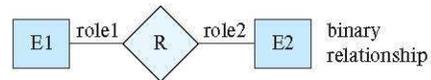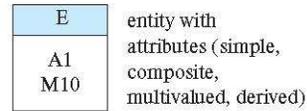- Binary relationship sets are represented in UML by just drawing a line connecting the entity sets. The relationship set name is written adjacent to the line.

- The role played by an entity set in a relationship set may also be specified by writing the role name on the line, adjacent to the entity set.

- The relationship set name may alternatively be written in a box, along with attributes of the relationship set, and the box is connected, using a dotted line, to the line depicting the relationship set.
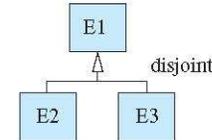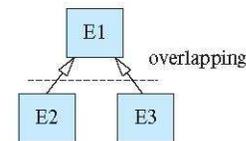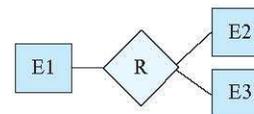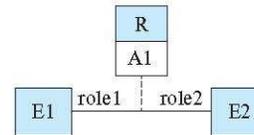
ER vs. UML Class Diagrams

ER Diagram Notation

E	entity with
A1	attributes (simple,
M10	composite, multivalued, derived)

E1 —role1— R —role2— E2   binary relationship

A1
E1 —role1— R —role2— E2   relationship attributes

E1 —0..*— R —0..1— E2   cardinality constraints

E1 — R — E2 / E3   *n*-ary relationships

E1 → E2, E3   overlapping generalization

E1 → E2, E3   disjoint generalization

E1 ← R = E2   weak-entity composition

Equivalent in UML

E	class with simple attributes
-A1	and methods (attribute
+M10	prefixes: + = public, - = private, # = protected)

E1 —role1 R role2— E2

R
A1
E1 —role1   role2— E2

E1 —0_1 R 0_*— E2

E1 — R — E2, E3

E1 → E2, E3   overlapping

E1 → E2, E3   disjoint

E1 ◆— E2   composition

Other Aspects of Database Design

- Functional Requirements

- Data Flow, Workflow

- Schema Evolution

End of Chapter 6