

DATABASE MANAGEMENT SYSTEMS

Module - III (Part-2)

Database Design Using the E-R Model

Lecture - 2

Mapping from ER to Relational Model



Reduction to Relation Schemas



Reduction to Relation Schemas

- **Entity sets** and **relationship sets** can be expressed uniformly *as relation schemas* that represent the contents of the database.
- A database which conforms to an **E-R diagram** can be represented by a collection of **schemas**.
- For each **entity set** and **relationship set** there is a unique schema that is assigned the name of the corresponding entity set or relationship set.
- Each schema has a number of **columns** (generally corresponding to attributes), which have unique names.



Representing Entity Sets


- A **strong entity set** reduces to a schema with the same attributes

student(ID, name, tot_cred)

- A **weak entity set** becomes a table that includes a column for the primary key of the identifying strong entity set

section (course_id, sec_id, sem, year)

- **Example**





Representation of Entity Sets with Composite Attributes

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

- **Composite attributes** are flattened out by creating a separate attribute for each component attribute
 - **Example:** given entity set *instructor* with composite attribute *name* with component attributes *first_name* and *last_name* the schema corresponding to the entity set has two attributes *name_first_name* and *name_last_name*
 - Prefix omitted if there is no ambiguity (*name_first_name* could be *first_name*)
- Ignoring **multivalued attributes**, extended instructor schema is
 - *instructor*(ID, *first_name*, *middle_initial*, *last_name*, *street_number*, *street_name*, *apt_number*, *city*, *state*, *zip_code*, *date_of_birth*)



Representation of Entity Sets with Multivalued Attributes


- A **multivalued attribute** M of an entity E is represented by a separate schema EM
- Schema EM has attributes corresponding to the **primary key** of E and an attribute corresponding to multivalued attribute M
- **Example:** Multivalued attribute *phone_number* of *instructor* is represented by a schema:
 $inst_phone = (\underline{ID}, \underline{phone_number})$
- Each value of the multivalued attribute maps to a separate tuple of the relation on schema EM
 - **For example**, an *instructor* entity with primary key 22222 and phone numbers 456-7890 and 123-4567 maps to two tuples:
 $(22222, 456-7890)$ and $(22222, 123-4567)$



Representing Relationship Sets

- A **many-to-many** relationship set is represented as a schema with attributes for the primary keys of the two participating entity sets, and any descriptive attributes of the relationship set.
- **Example:** schema for relationship set *advisor*

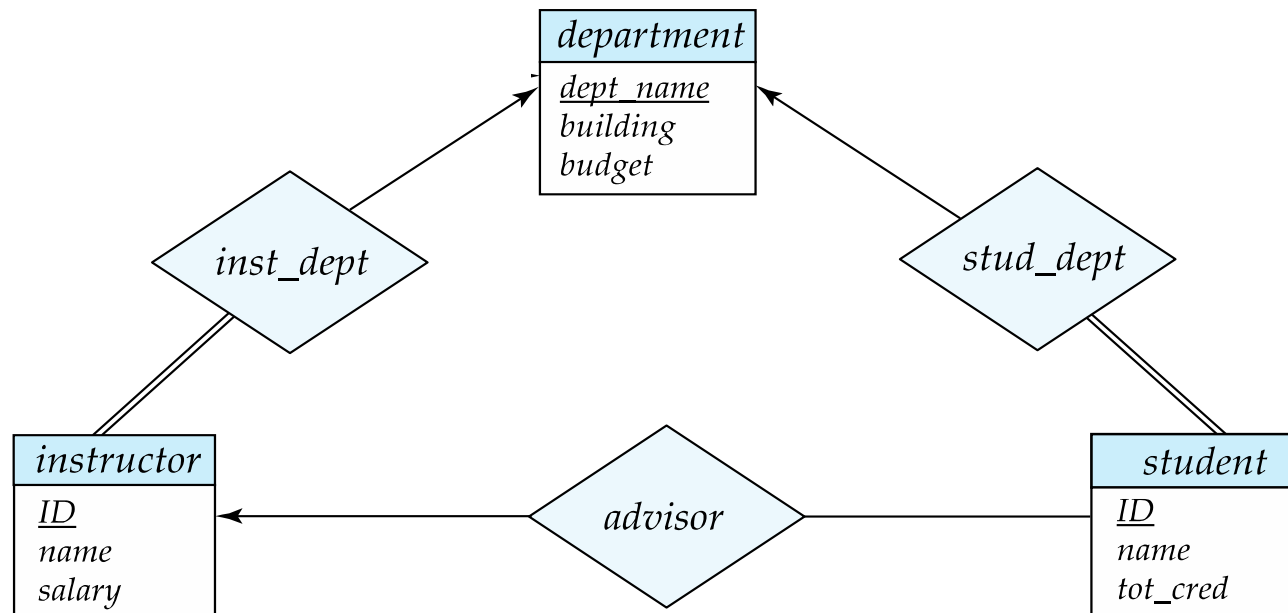
advisor = (*s_id*, *i_id*)





Redundancy of Schemas

- **Many-to-one** and **one-to-many** relationship sets that are total on the many-side can be represented by adding an extra attribute to the “many” side, containing the primary key of the “one” side
- **Example:** Instead of creating a schema for relationship set *inst_dept*, add an attribute *dept_name* to the schema arising from entity set *instructor*
- **Example**






Redundancy of Schemas (Cont.)

- For **one-to-one** relationship sets, either side can be chosen to act as the “many” side
 - That is, an extra attribute can be added to either of the tables corresponding to the two entity sets
- If participation is *partial* on the “many” side, replacing a schema by an extra attribute in the schema corresponding to the “many” side could result in **null** values



Redundancy of Schemas (Cont.)

- The schema corresponding to a relationship set linking a **weak entity set** to its identifying **strong entity set** is redundant.
- **Example:** The *section* schema already contains the attributes that would appear in the *sec_course* schema





Extended E-R Features



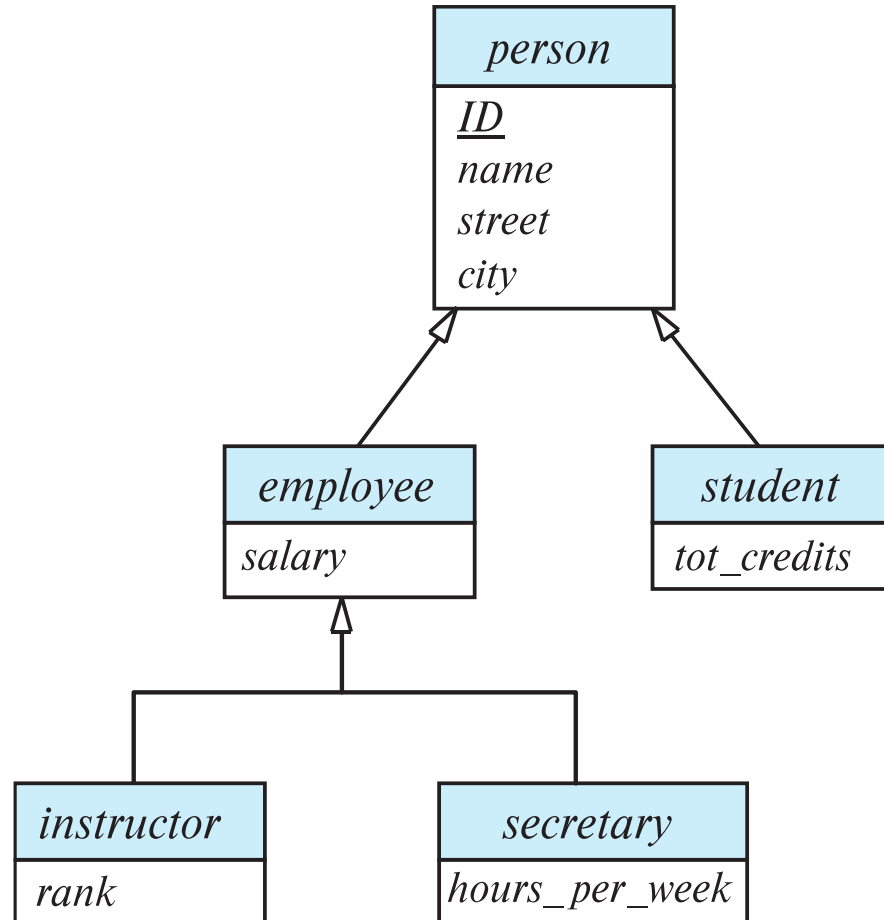
Specialization

- **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





Representing Specialization via Schemas

■ Method 1:

- Form a schema for the higher-level entity
- Form a schema for each lower-level entity set, include primary key of higher-level entity set and local attributes

schema	attributes
person	ID, name, street, city
student	ID, tot_cred
employee	ID, salary

- **Drawback:** getting information about, an *employee* requires accessing two relations, the one corresponding to the low-level schema and the one corresponding to the high-level schema



Representing Specialization as Schemas (Cont.)

■ Method 2:

- Form a schema for each entity set with all local and inherited attributes

schema	attributes
person	ID, name, street, city
student	ID, name, street, city, tot_cred
employee	ID, name, street, city, salary

- **Drawback:** *name, street* and *city* may be stored redundantly for people who are both students and employees



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**.



Completeness constraint

- **Completeness constraint** -- specifies whether or not an entity in the higher-level entity set must belong to at least one of the lower-level entity sets within a generalization.
 - **total**: an entity must belong to one of the lower-level entity sets
 - **partial**: an entity need not belong to one of the lower-level entity sets




Completeness constraint (Cont.)

- **Partial generalization** is the default.
- We can specify total generalization in an **ER diagram** by adding the keyword **total** in the diagram and drawing a dashed line from the keyword to the corresponding hollow arrow-head to which it applies (for a total generalization), or to the set of hollow arrow-heads to which it applies (for an overlapping generalization).
- The **student generalization** is **total**: All student entities must be either graduate or undergraduate. Because the higher-level entity set arrived at through generalization is generally composed of only those entities in the lower-level entity sets, the completeness constraint for a generalized higher-level entity set is usually total



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





Reduction to Relational Schemas

- To represent **aggregation**, create a schema containing
 - Primary key of the aggregated relationship,
 - The primary key of the associated entity set
 - Any descriptive attributes
- In our **example**:
 - The schema *eval_for* is:
$$eval_for (s_ID, project_id, i_ID, evaluation_id)$$
 - The schema *proj_guide* is redundant.




Design Issues

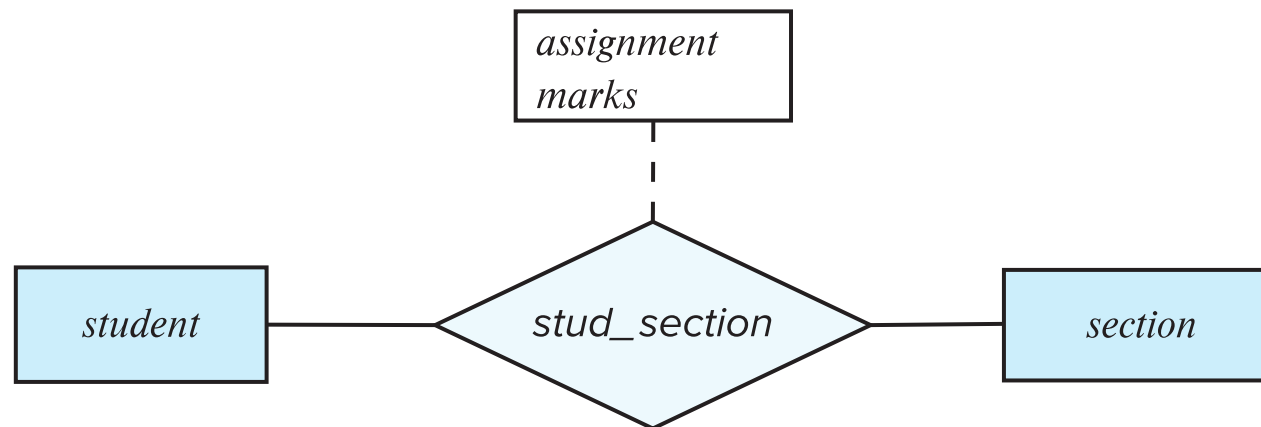


Common Mistakes in E-R Diagrams

- Example of erroneous E-R diagrams



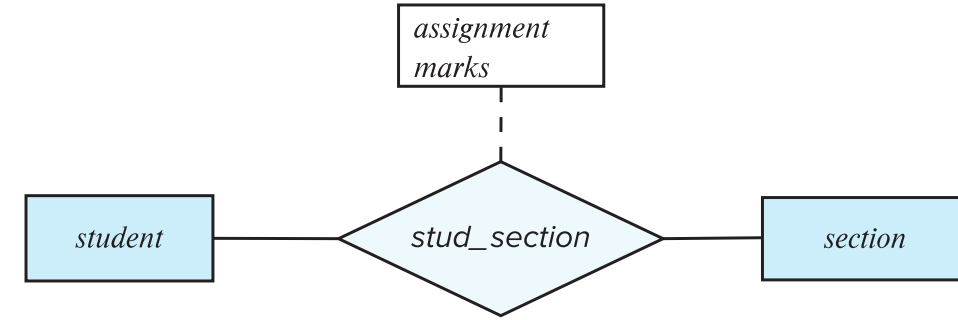
(a) Incorrect use of attribute



(b) Erroneous use of relationship attributes

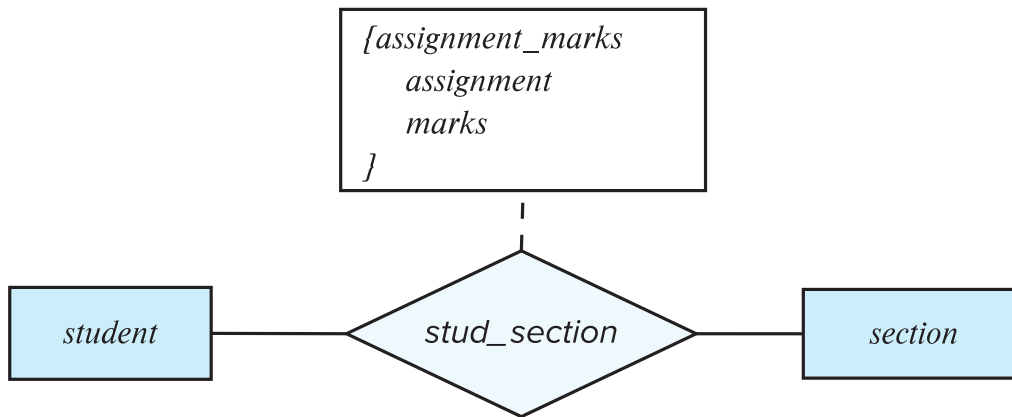


Common Mistakes in E-R Diagrams (Cont.)



(b) Erroneous use of relationship attributes

(c) Correct alternative to erroneous E-R diagram (b)

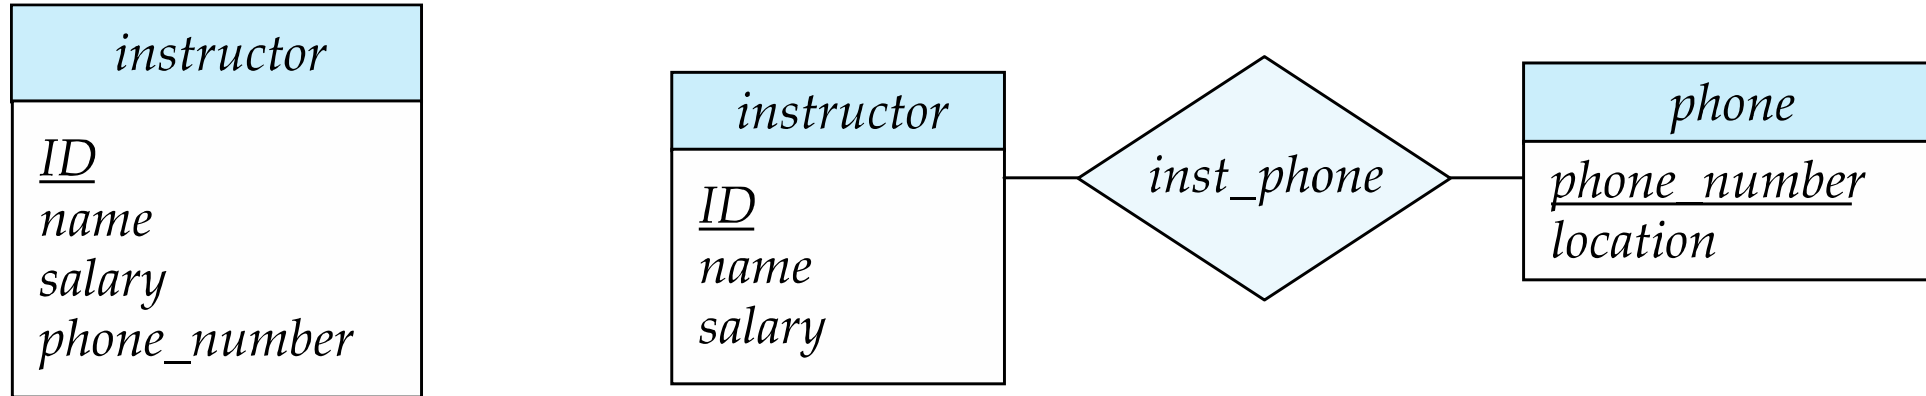


(d) Correct alternative to erroneous E-R diagram (b)



Entities vs. Attributes

- Use of entity sets vs. attributes



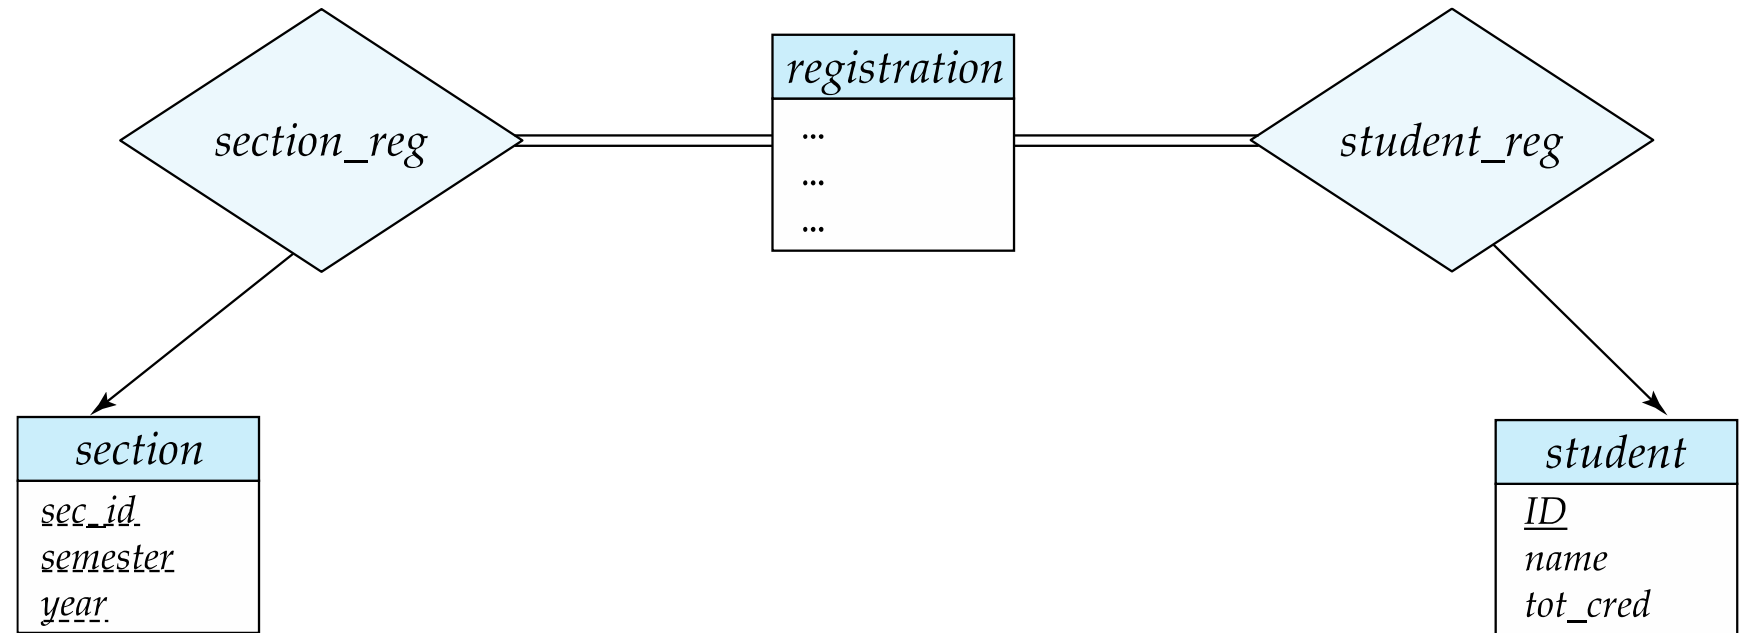
- Use of phone as an entity allows extra information about phone numbers (plus multiple phone numbers)



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*




Converting Non-Binary Relationships to Binary Form

- In general, any non-binary relationship can be represented using binary relationships by creating an artificial entity set.
 - Replace R between entity sets A , B and C by an entity set E , and three relationship sets:
 1. R_A , relating E and A
 2. R_B , relating E and B
 3. R_C , relating E and C
 - Create an identifying attribute for E and add any attributes of R to E
 - For each relationship (a_i, b_i, c_i) in R , create
 1. a new entity e_i in the entity set E
 2. add (e_i, a_i) to R_A
 3. add (e_i, b_i) to R_B
 4. add (e_i, c_i) to R_C



(a)



(b)



Converting Non-Binary Relationships (Cont.)

- Also need to translate **constraints**
 - Translating all constraints may not be possible
 - There may be instances in the translated schema that cannot correspond to any instance of R
 - **Exercise:** *add constraints to the relationships R_A , R_B and R_C to ensure that a newly created entity corresponds to exactly one entity in each of entity sets A , B and C*
 - We can avoid creating an identifying attribute by making E a weak entity set (described shortly) identified by the three relationship sets

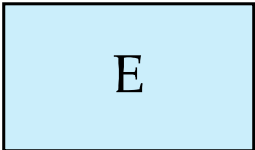
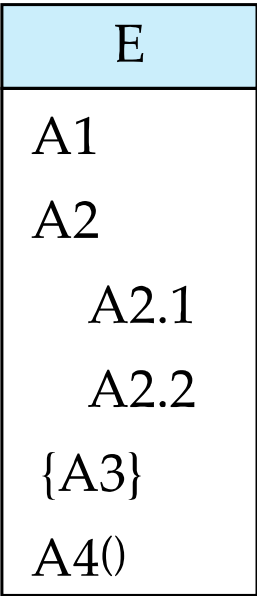
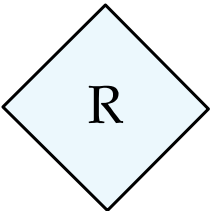
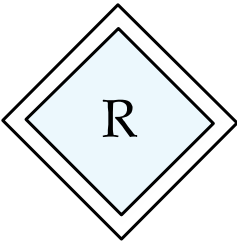
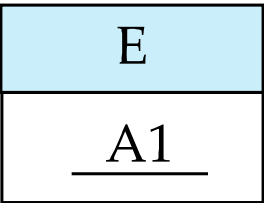
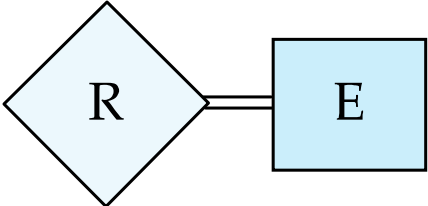
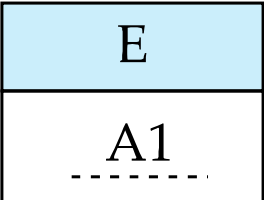


E-R Design Decisions

- 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


	entity set		attributes: simple (A1), composite (A2) and multivalued (A3) derived (A4)
	relationship set		
	identifying relationship set for weak entity set		primary key
	total participation of entity set in relationship		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



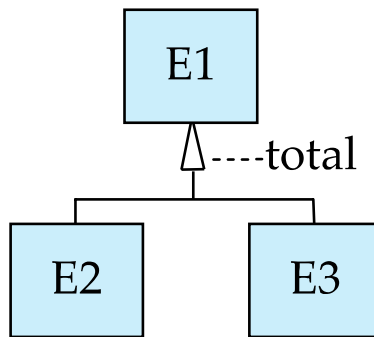
cardinality limits




role indicator



ISA: generalization or specialization



total (disjoint) generalization



disjoint generalization






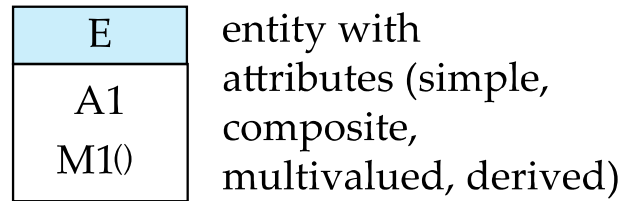
UML

- **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.




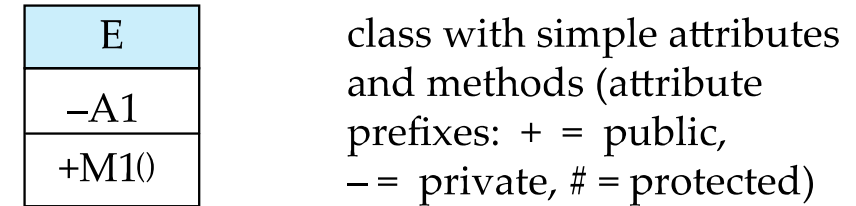


ER vs. UML Class Diagrams

ER Diagram Notation



Equivalent in UML




* Note reversal of position in cardinality constraint depiction








ER vs. UML Class Diagrams

ER Diagram Notation



Equivalent in UML



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



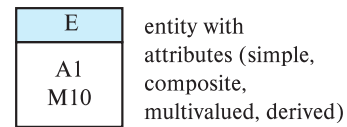
UML Class Diagrams (Cont.)

- 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.













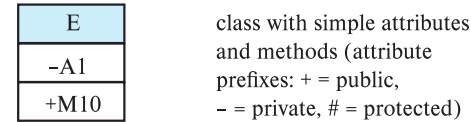


ER vs. UML Class Diagrams

ER Diagram Notation



Equivalent in UML





Other Aspects of Database Design

- Functional Requirements
- Data Flow, Workflow
- Schema Evolution