

These slides are for use with

Database Systems

Concepts, Languages and Architectures

Paolo Atzeni • Stefano Ceri • Stefano Paraboschi • Riccardo Torlone
© McGraw-Hill 1999

Concepts,
Languages
and
Architectures


Paolo Atzeni,
Stefano Ceri,
Stefano Paraboschi
and Riccardo Torlone

Mc
Graw
Hill

To view these slides on-screen or with a projector use the arrow keys to move to the next or previous slide. The return or enter key will also take you to the next slide. Note you can press the 'escape' key to reveal the menu bar and then use the standard Acrobat controls — including the magnifying glass to zoom in on details.

To print these slides on acetates for projection use the escape key to reveal the menu and choose 'print' from the 'file' menu. If the slides are too large for your printer then select 'shrink to fit' in the print dialogue box.

Press the 'return' or 'enter' key to continue . . .



Chapter 7

Logical design

Logical design

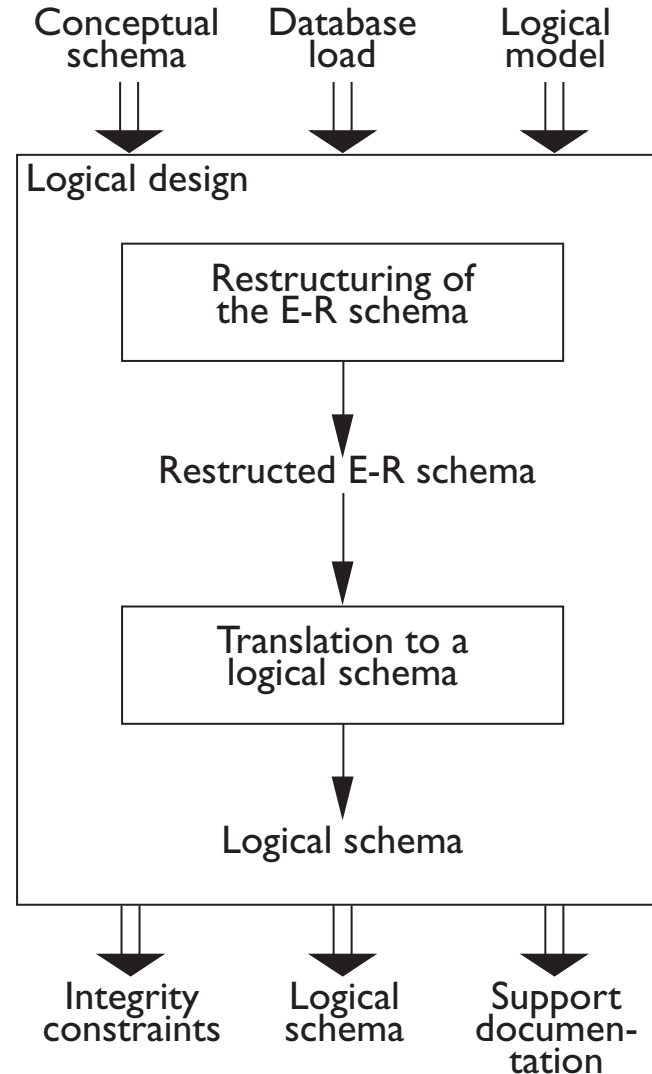
- The aim of logical design is to construct a relational schema that correctly and efficiently represents all of the information described by an Entity-Relationship schema produced during the conceptual design phase.
- This is not just a simple translation from one model to another for two main reasons:
 - not all the constructs of the Entity-Relationship model can be translated naturally into the relational model;
 - the schema must be restructured in such a way as to make the execution of the projected operations as efficient as possible.

Logical design steps

It is usually helpful to divide the logical design into two steps:

- **restructuring of the Entity-Relationship schema**, based on criteria for the optimization of the schema and the simplification of the following step;
- **translation into the logical model**, based on the features of the logical model (in our case, the relational model).

Logical database design



Performance analysis on E-R schemas

- An E-R schema can be restructured to optimize two parameters:
 - **cost of an operation** (evaluated in terms of the number of occurrences of entities and relationships that are visited to execute an operation on the database);
 - **storage requirement** (evaluated in terms of number of bytes necessary to store the data described by the schema).
- In order to study these parameters, we need to know:
 - the volume of data;
 - the operation characteristics.

An E-R schema on the personnel of a company




Table of volumes and table of operations

The volume of data and the general characteristics of the operations can be summed up in special tables.

Table of volumes

Concept	Type	Volume
Branch	E	10
Department	E	80
Employee	E	2000
Project	E	500
Composition	R	80
Membership	R	1900
Management	R	80
Participation	R	6000

Table of operations

Operation	Type	Frequency
Operation 1	I	50 per day
Operation 2	I	100 per day
Operation 3	I	10 per day
Operation 4	B	2 per day

Example of a navigation schema

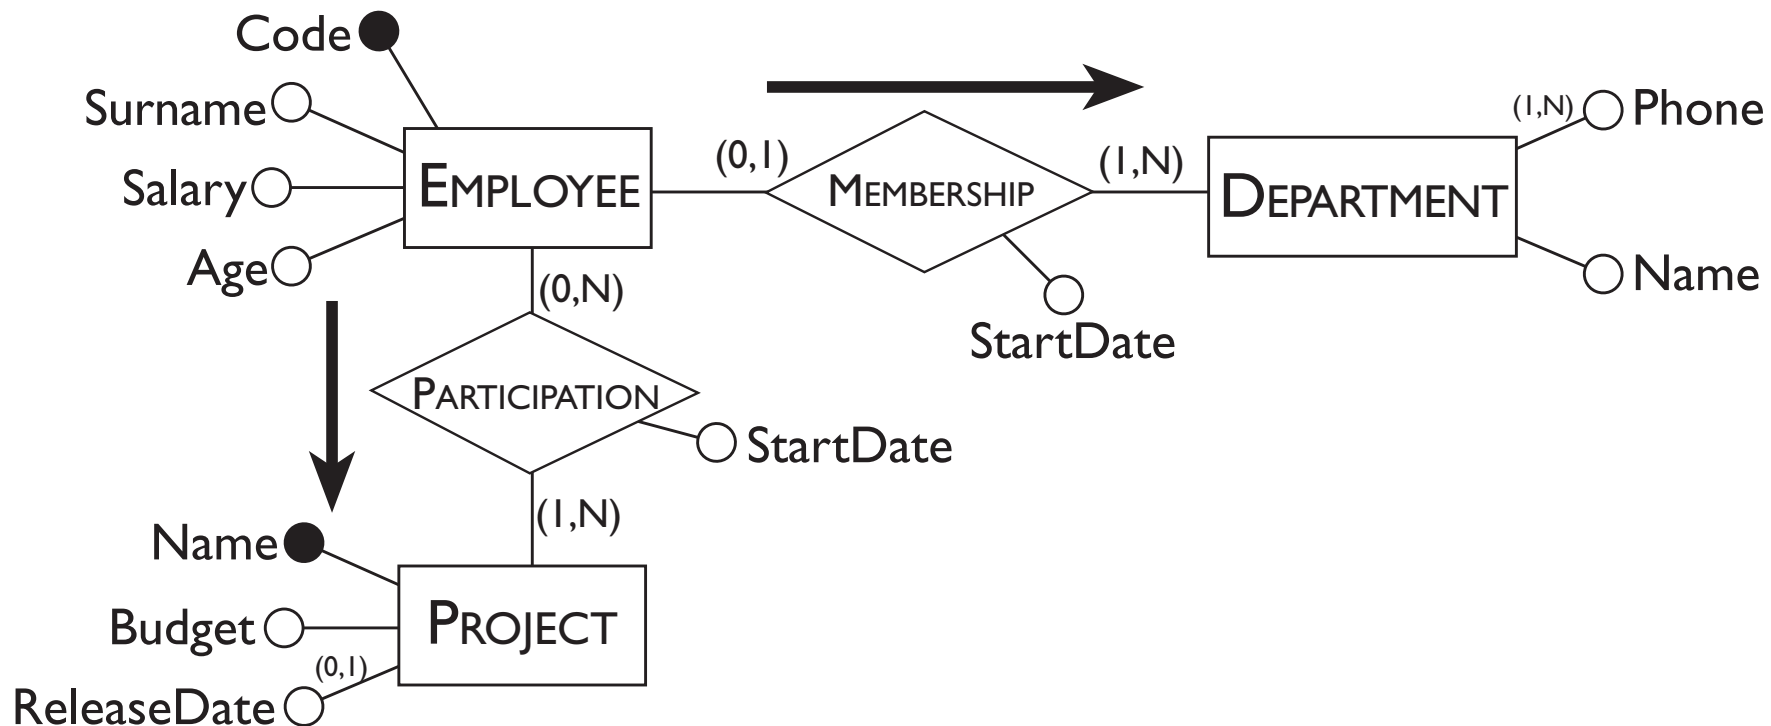



Table of accesses

The cost of an operation evaluated using the table of volumes and the navigation schema can be summed up in the table of accesses.

Concept	Type	Accesses	Type
Employee	Entity	1	R
Membership	Relationship	1	R
Department	Entity	1	R
Participation	Relationship	3	R
Project	Entity	3	R


Restructuring tasks of an E-R schema



Analysis of redundancies

- A redundancy in a conceptual schema corresponds to a piece of information that can be derived (that is, obtained by a series of retrieval operations) from other data.
- An Entity-Relationship schema can contain various forms of redundancy.


Examples of schemas with redundancies



Taking a decision about redundancies

- The presence of a derived piece of information in a database presents
 - an advantage: a reduction in the number of accesses necessary to obtain the derived information;
 - some disadvantages: a larger storage requirement (which is often a negligible cost) and the necessity for carrying out additional operations in order to keep the derived data up to date.
- The decision to maintain or delete a redundancy is made by comparing the cost of operations that involve the redundant information and the storage needed, in the case of presence or absence of redundancy.

An example of analysis of redundancy



In this schema the attribute **NumberOfInhabitants** is redundant.

Load and operations for the example schema

Table of volumes

Concept	Type	Volume
Town	E	200
Person	E	1000000
Residence	R	1000000

Table of operations

Operation	Type	Frequency
Operation 1	I	500 per day
Operation 2	I	2 per day

- **Operation 1:** add a new person with the person's town of residence.
- **Operation 2:** print all the data of a town (including the number of inhabitants).

Table of accesses in presence of redundancy

Operation 1

Concept	Type	Accesses	Type
Person	Entity	1	W
Residence	Relationship	1	W
Town	Entity	1	R
Town	Entity	1	W

Operation 2

Concept	Type	Accesses	Type
Town	Entity	1	R

Table of accesses in absence of redundancy

Operation 1

Concept	Type	Accesses	Type
Person	Entity	1	W
Residence	Relationship	1	W

Operation 2

Concept	Type	Accesses	Type
Town	Entity	1	R
Residence	Relationship	5000	R

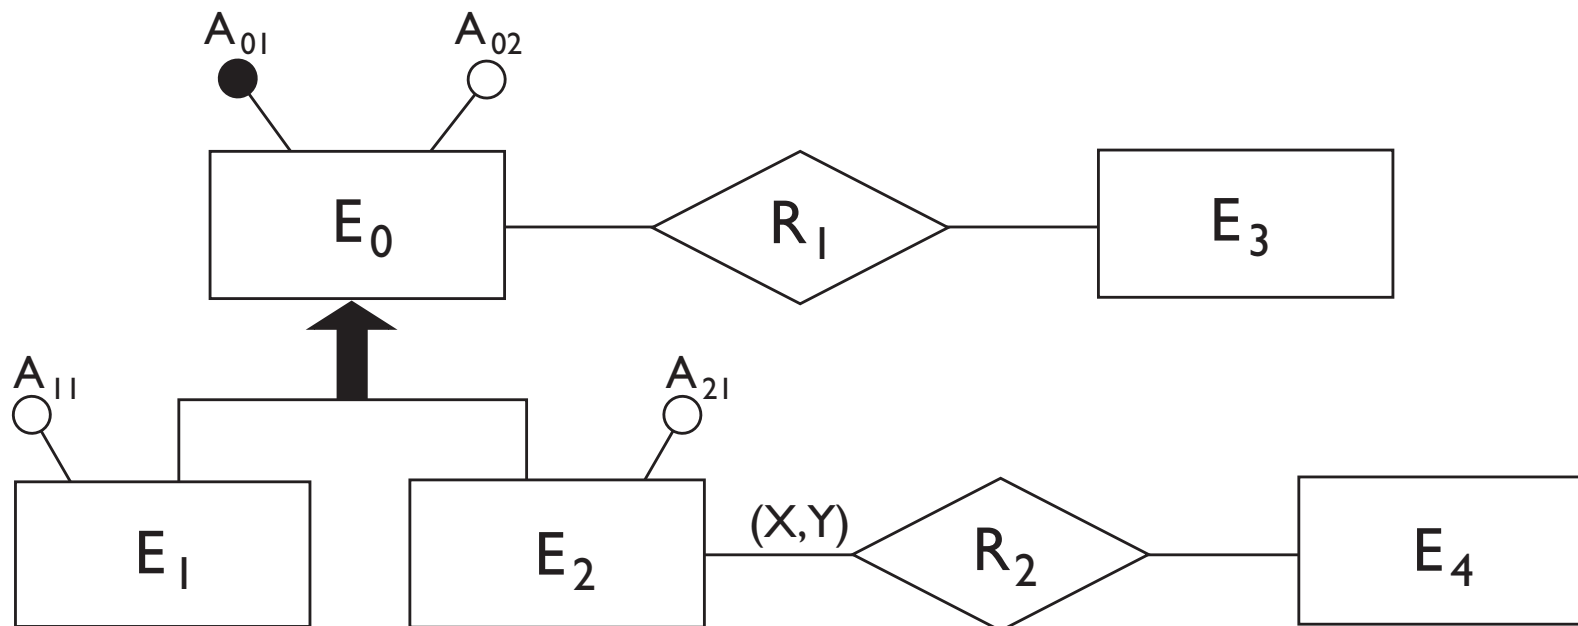
Comparing the cost of operations

- **Presence of redundancy.**
 - Operation 1 requires a total of 1500 write accesses and 500 read accesses per day.
 - The cost of operation 2 is almost negligible.
 - Counting twice the write accesses, we have a total of 3500 accesses a day.
- **Absence of redundancy.**
 - Operation 1 requires a total of 1000 write accesses per day.
 - Operation 2 however requires a total of 10000 read accesses per day.
 - Counting twice the write accesses, we have a total of 12000 accesses per day.
- It worth maintaining the redundant data




Removing generalizations

- The relational model does not allow the direct representation of generalizations of the E-R model.
- We need, therefore, to transform these constructs into other constructs that are easier to translate: entities and relationships.

Example of a schema with generalization




Possible restructurings of the previous schema



General rules about generalization removal

- Option 1 is useful when the operations involve the occurrences and the attributes of E_0 , E_1 and E_2 more or less in the same way.
- Option 2 is possible only if the generalization is total and is useful when there are operations that refer only to occurrences of E_1 or of E_2 , and so they make distinctions between these entities.
- Option 3 is useful when the generalization is not total and the operations refer to either occurrences and attributes of E_1 (E_2) or of E_0 , and therefore make distinctions between child and parent entities.
- The various options can be combined.

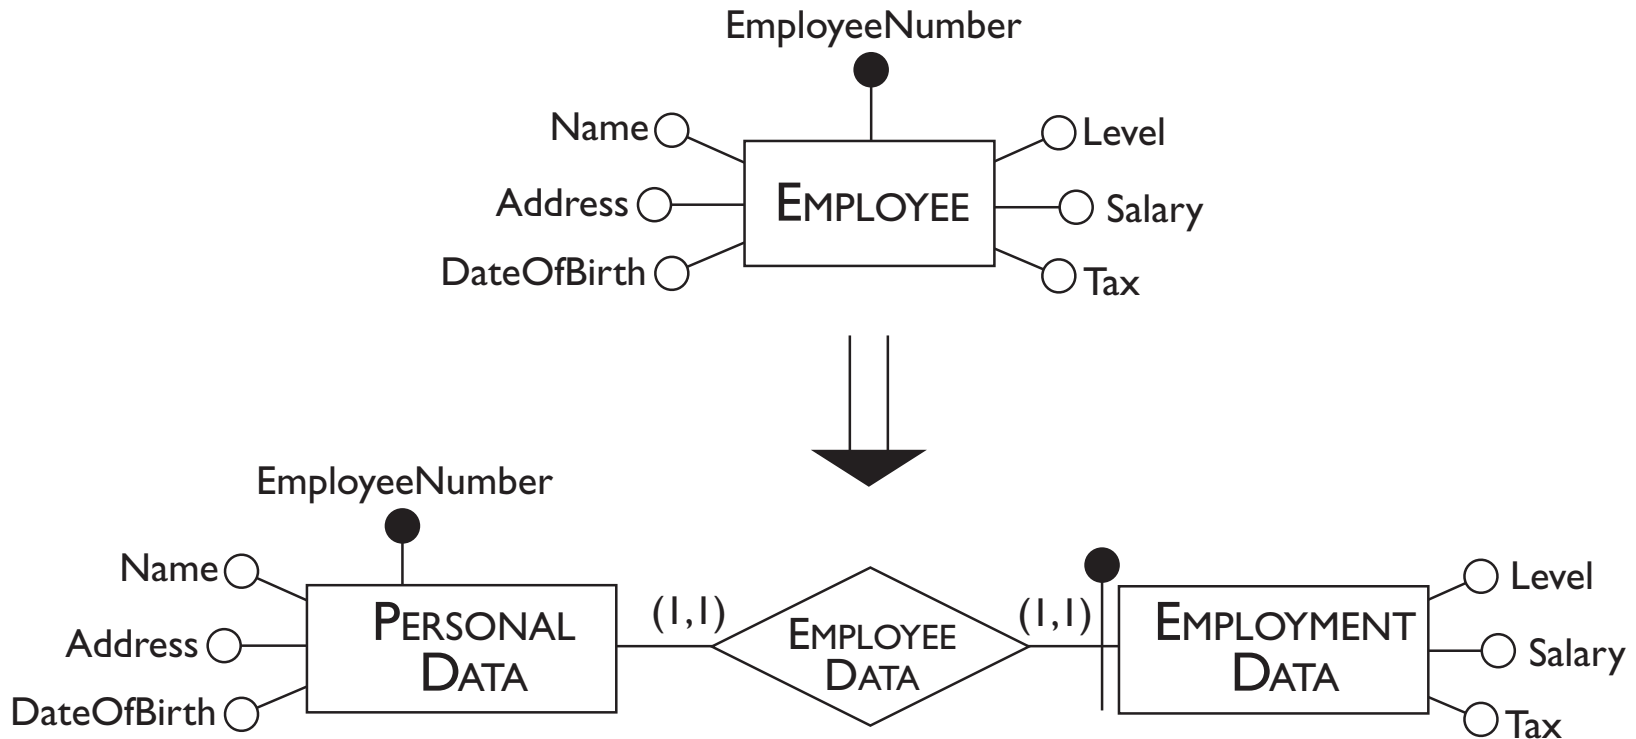
Possible restructuring of the previous schema



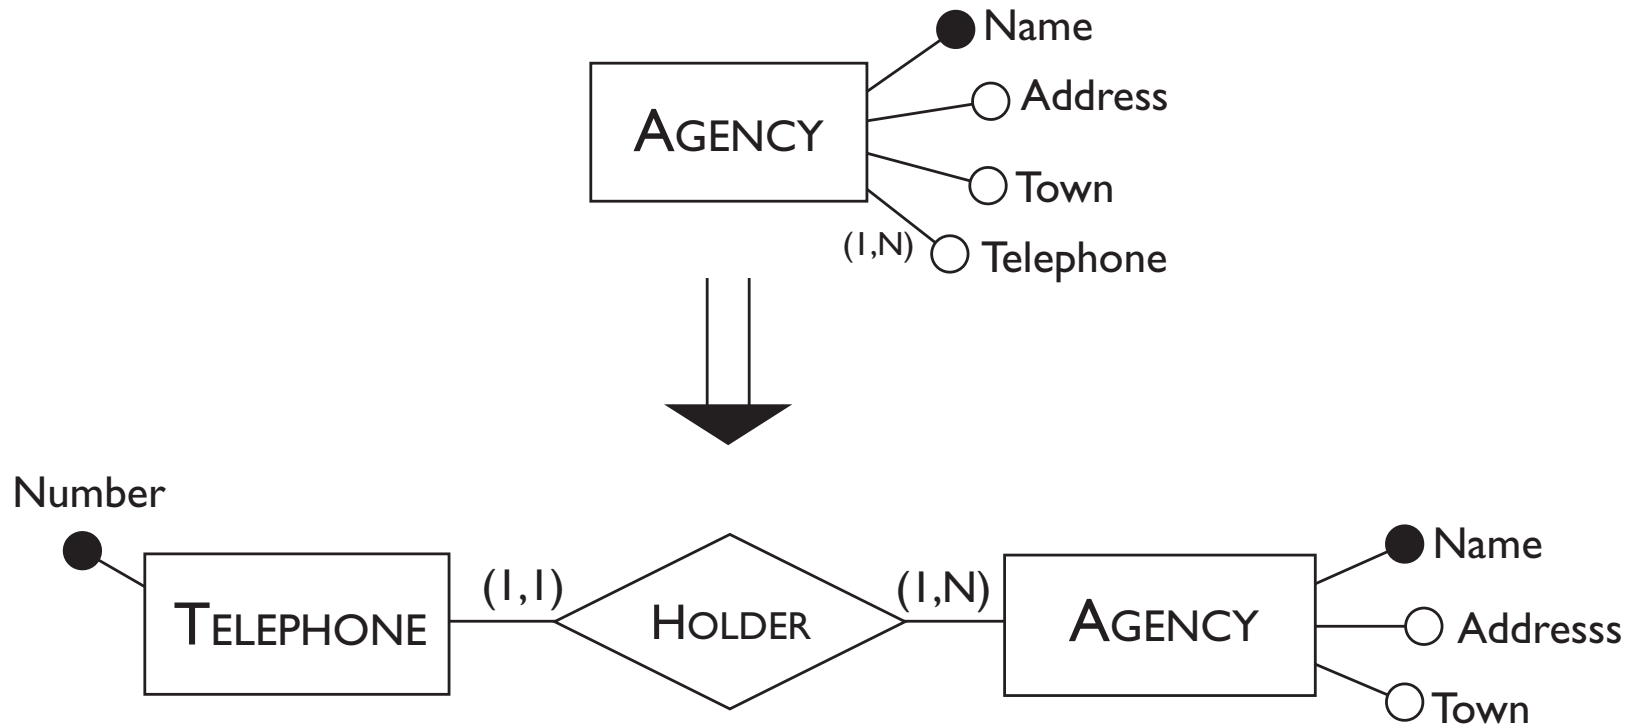
Partitioning and merging of entities and relationships

- Entities and relationships of an E-R schema can be partitioned or merged to improve the efficiency of operations, using the following principle.
 - Accesses are reduced by separating attributes of the same concept that are accessed by different operations and by merging attributes of different concepts that are accessed by the same operations.
- The same criteria as those discussed for redundancies are valid in making a decision about this type of restructuring.


Example of partitioning of entities




Example of deletion of multi-value attributes



Example of merging of entities



Example of partitioning of a relationship




Selection of primary identifiers

- The criteria for this decision are as follows.
 - Attributes with null values cannot form primary identifiers.
 - One or few attributes are preferable to many attributes.
 - An internal identifier with few attributes is preferable to an external one, possibly involving many entities.
 - An identifier that is used by many operations to access the occurrences of an entity is preferable to others.
- At this stage, if none of the candidate identifiers satisfies the above requirements, it is possible to introduce a further attribute to the entity. This attribute will hold special values (often called *codes*) generated solely for the purpose of identifying occurrences of the entity.

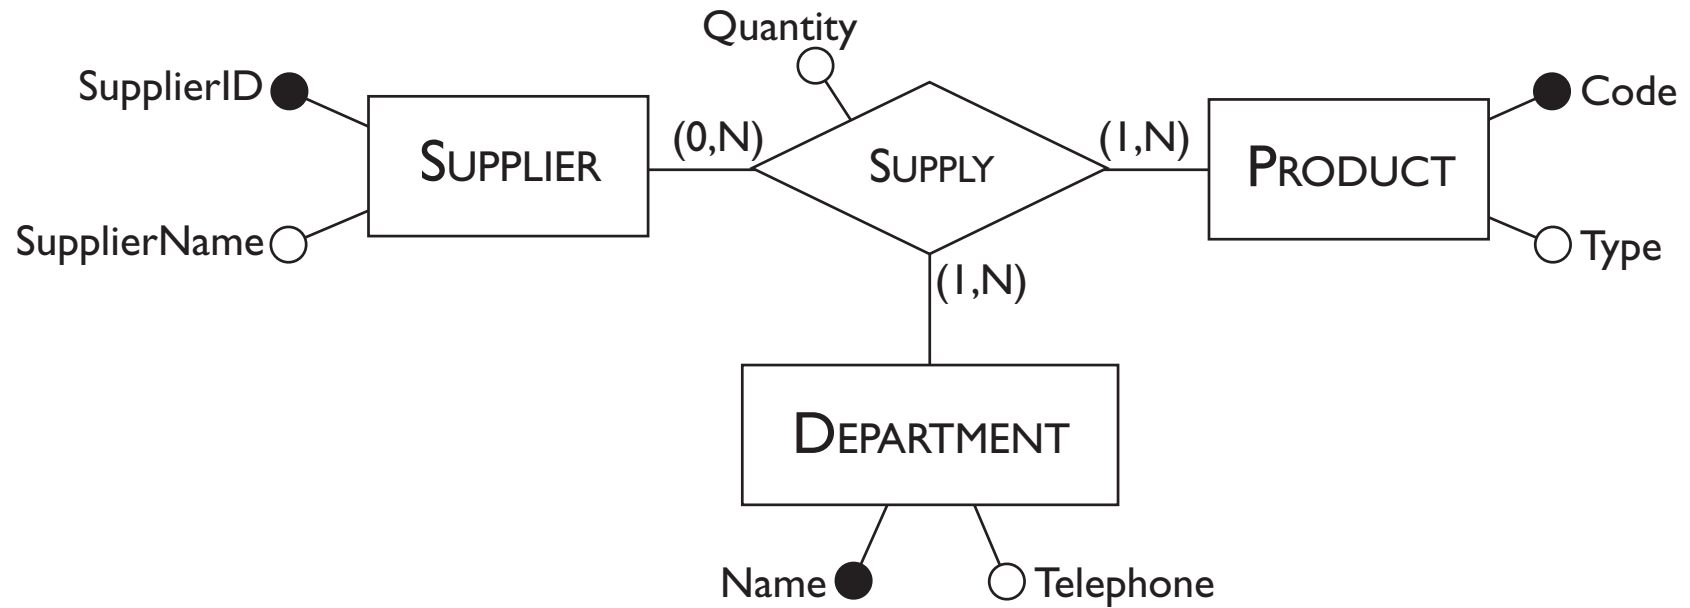
Translation into the relational model

- The second step of logical design corresponds to a translation between different data models.
- Starting from an E-R schema, an equivalent relational schema is constructed. By equivalent, we mean a schema capable of representing the same information.
- We will deal with the translation problem systematically, beginning with the fundamental case, that of entities linked by many-to-many relationships.


An E-R schema with a many-to-many relationship




E-R schema with ternary relationship



E-R schema with one-to-many relationships



E-R schema with external identifier



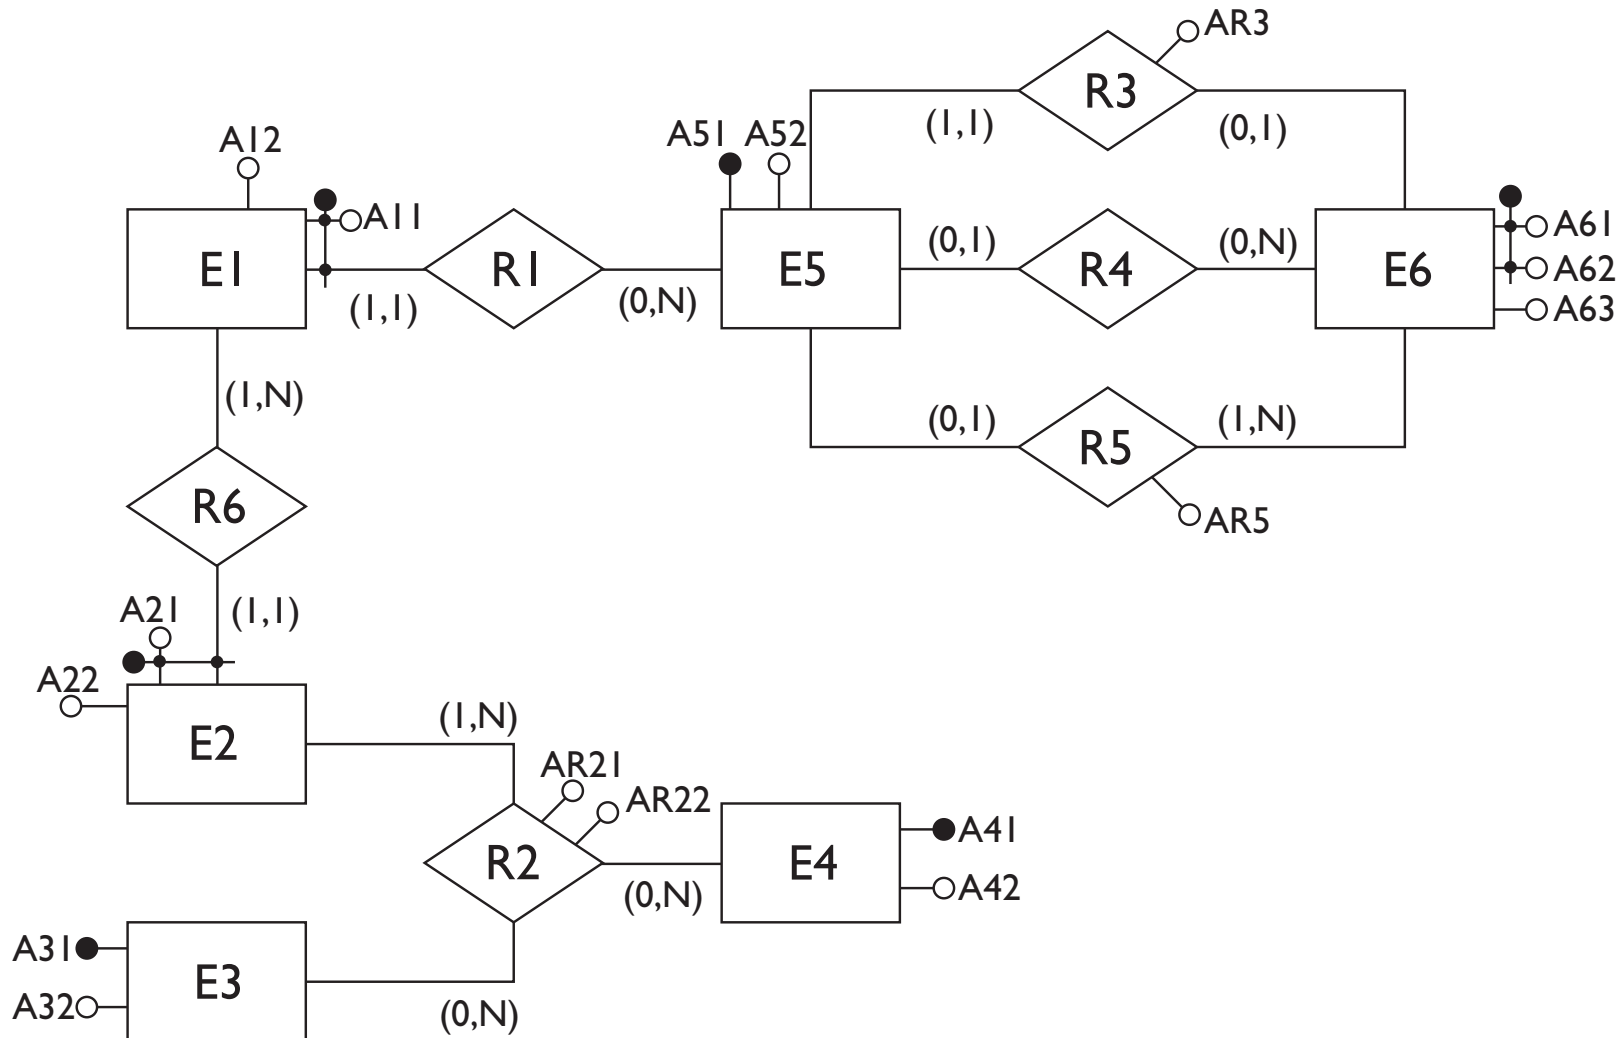
E-R schema with one-to-one relationship



E-R schema with one-to-one relationship



An E-R schema for translation



Result of the translation in the relational model

E1(A11, A51, A12)
E2(A21, A11, A51, A22)
E3(A31, A32)
E4(A41, A42)
E5(A51, A52, A61R3, A62R3, AR3, A61R4, A62R4, A61R5, A62R5, AR5)
E6(A61, A62, A63)
R2(A21, A11, A51, A31, A41, AR21, AR22)

Translations from the E-R model to the relational (1)

Type	Initial schema	Possible translation
Binary many-to-many relationship		$E_1(\underline{A_{E11}}, A_{E12})$ $E_2(\underline{A_{E21}}, A_{E22})$ $R(\underline{A_{E11}}, \underline{A_{E21}}, A_R)$
Ternary many-to-many relationship		$E_1(\underline{A_{E11}}, A_{E12})$ $E_2(\underline{A_{E21}}, A_{E22})$ $E_3(\underline{A_{E31}}, A_{E32})$ $R(\underline{A_{E11}}, \underline{A_{E21}}, \underline{A_{E31}}, A_R)$
One-to-many relationship with mandatory participation		$E_1(\underline{A_{E11}}, A_{E12}, A_{E21}, A_R)$ $E_2(\underline{A_{E21}}, A_{E22})$

Translations from the E-R model to the relational (2)

Type	Initial schema	Possible translation
One-to-many relationship with optional participation		$E_1(\underline{A_{E11}}, A_{E12})$ $E_2(\underline{A_{E21}}, A_{E22})$ $R(\underline{A_{E11}}, \underline{A_{E21}}, A_R)$ <p>Alternatively:</p> $E_1(\underline{A_{E11}}, A_{E21}, A_{E21}^*, A_R^*)$ $E_2(\underline{A_{E21}}, A_{E22})$
Relationship with external identifiers		$E_1(\underline{A_{E12}}, \underline{A_{E21}}, A_{E11}, A_R)$ $E_2(\underline{A_{E21}}, A_{E22})$


Translations from the E-R model to the relational (3)

Type	Initial schema	Possible translation
One-to-one relationship with mandatory participation for both entities		$E_1(\underline{A_{E11}}, A_{E12}, \underline{A_{E21}}, A_R)$ $E_2(\underline{A_{E21}}, A_{E22})$ <p>Alternatively:</p> $E_2(\underline{A_{E21}}, A_{E22}, \underline{A_{E11}}, A_R)$ $E_1(\underline{A_{E11}}, A_{E12})$
One-to-one relationship with optional participation for one entity		$E_1(\underline{A_{E11}}, A_{E12}, \underline{A_{E21}}, A_R)$ $E_2(\underline{A_{E21}}, A_{E22})$

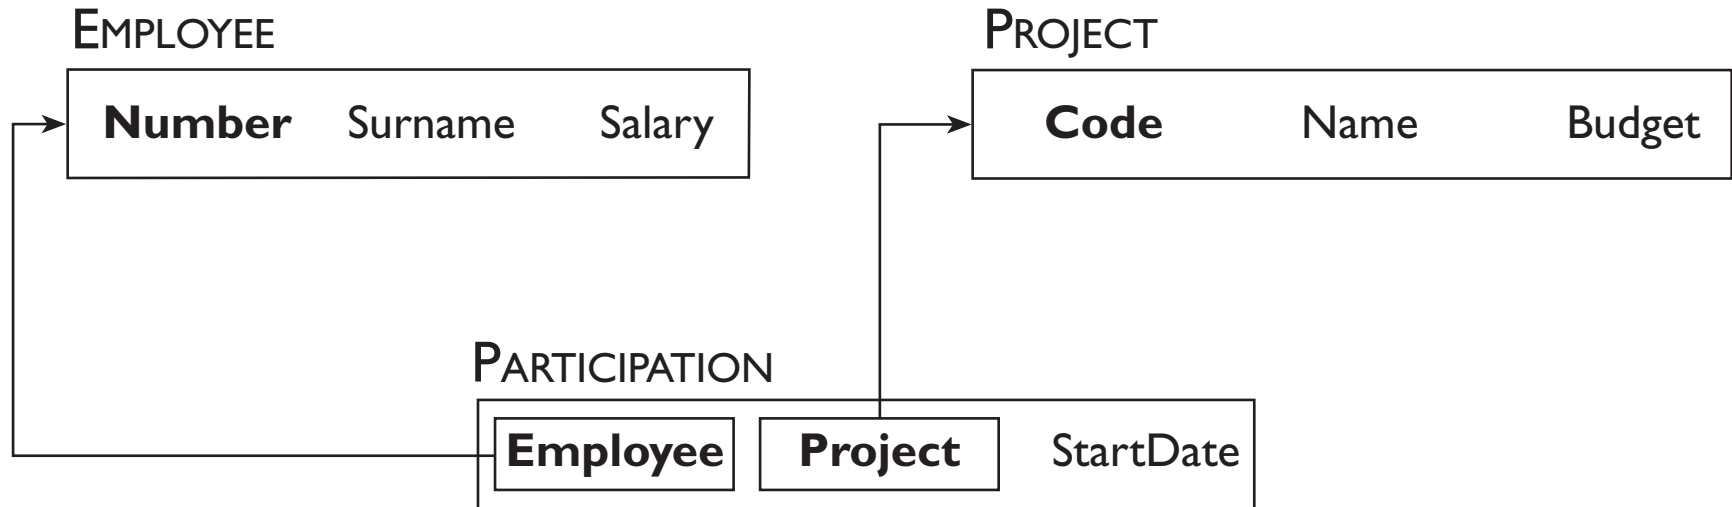
Translations from the E-R model to the relational (4)

Type	Initial schema	Possible translation
One-to-one relationship with optional participation for both entities		$E_1(\underline{A_{E11}}, A_{E21})$ $E_2(\underline{A_{E21}}, A_{E22}, A_{E11}^*, A_R^*)$ <p>Alternatively:</p> $E_1(\underline{A_{E11}}, A_{E12}, A_{E21}^*, A_R^*)$ $E_2(\underline{A_{E21}}, A_{E22})$ <p>Alternatively:</p> $E_1(\underline{A_{E11}}, A_{E12})$ $E_2(\underline{A_{E21}}, A_{E22})$ $R(\underline{A_{E11}}, \underline{A_{E21}}, A_R)$


An E-R schema with a many-to-many relationship



Graphical representation of a translation of the previous schema



E-R schema with one-to-many relationships



Graphical representation of a translation of the previous schema

PLAYER




TEAM

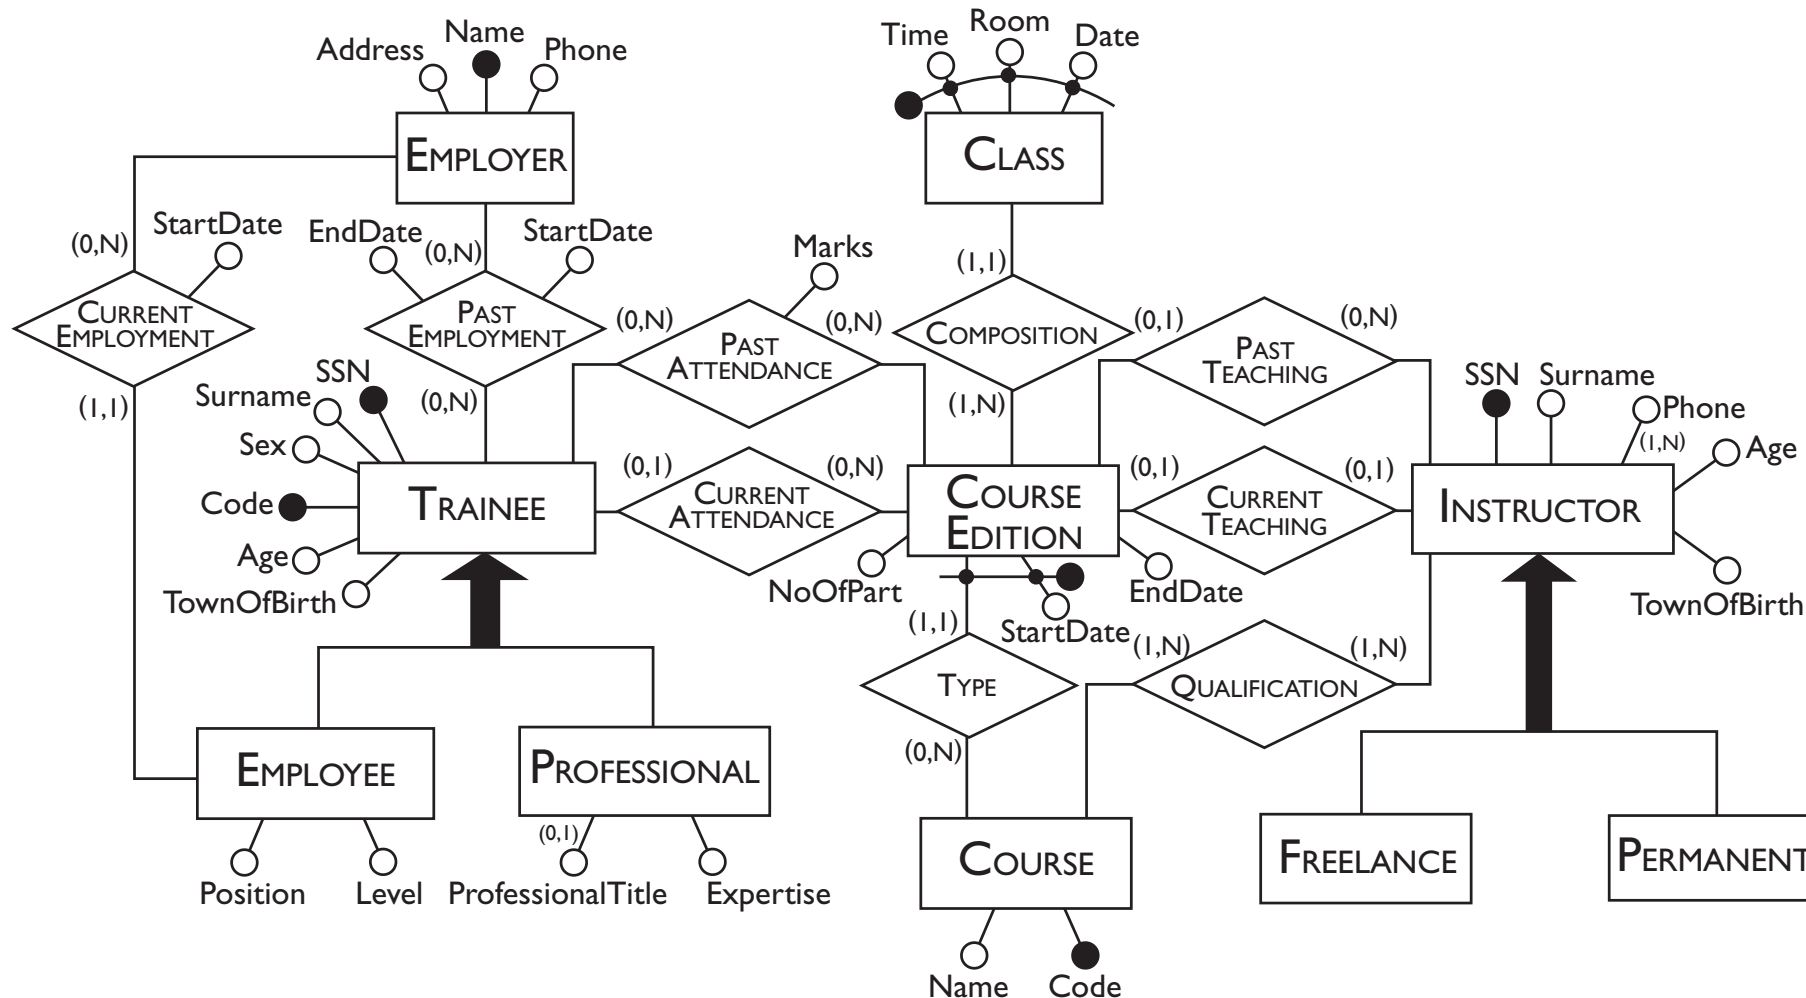


CONTRACT

Graphical representation of a relational schema



The E-R schema of a training company



Operational requirements

- **operation 1:** insert a new trainee including all his or her data (to be carried out approximately 40 times a day);
- **operation 2:** assign a trainee to an edition of a course (50 times a day);
- **operation 3:** insert a new instructor, including all his or her data and the courses he or she is qualified to teach (twice a day);
- **operation 4:** assign a qualified instructor to an edition of a course (15 times a day);
- **operation 5:** display all the information on the past editions of a course with title, class timetables and number of trainees (10 times a day);
- **operation 6:** display all the courses offered, with information on the instructors who are qualified to teach them (20 times a day);
- **operation 7:** for each instructor, find the trainees all the courses he or she is teaching or has taught (5 times a week);
- **operation 8:** carry out a statistical analysis of all the trainees with all the information about them, about the editions of courses they have attended and the marks obtained (10 times a month).

Database load

Table of volumes

Concept	Type	Volume
Class	E	8000
CourseEdition	E	1000
Course	E	200
Instructor	E	300
Freelance	E	250
Permanent	E	50
Trainee	E	5000
Employee	E	4000
Professional	E	1000
Employer	E	8000
PastAttendance	R	10000
CurrentAttendance	R	500
Composition	R	8000
Type	R	1000
PastTeaching	R	900
CurrentTeaching	R	100
Qualification	R	500
CurrentEmployment	R	4000
PastEmployment	R	10000

Table of operations

Operation	Type	Frequency
Operation 1	I	40 per day
Operation 2	I	50 per day
Operation 3	I	2 per day
Operation 4	I	15 per day
Operation 5	I	10 per day
Operation 6	I	20 per day
Operation 7	I	5 per day
Operation 8	B	10 per month

Access tables for the analysis of the redundancy

The attribute NumberOfParticipants in COURSEEDITION can be derived from the relationships CURRENTATTENDANCE and PASTATTENDANCE.

Operation 2 with redundancy

Concept	Type	Acc	Type
Trainee	E	1	R
CurrentAtt'nce	R	1	W
CourseEdition	E	1	R
CourseEdition	E	1	W

Operation 5 with redundancy

Concept	Type	Acc	Type
CourseEdition	E	1	R
Type	R	1	R
Course	E	1	R
Composition	R	8	R
Class	E	8	R

Operation 2 without redundancy

Concept	Type	Acc	Type
Trainee	E	1	R
CurrentAtt'nce	R	1	W

Operation 5 without redundancy

Concept	Type	Acc	Type
CourseEdition	E	1	R
Type	R	1	R
Course	E	1	R
Composition	R	8	R
Class	E	8	R
PastAtt'nce	E	10	R

Analysis of the redundancy

- From the access tables we obtain (giving double weight to the write accesses):
 - presence of redundancy: for operation 2 we have 100 read accesses and 100 write accesses per day; for operation 5 we have 190 read accesses per day, for a total of 490 accesses per day;
 - without redundancy: for operation 2 we have 50 read accesses per day and 100 write accesses per day; for operation 5, we have 290 read accesses per day, for a total of 440 accesses per day.
- Thus, when the redundancy is present, we have disadvantages both in terms of storage and access time.
- We can therefore delete the attribute NumberOfParticipants from the entity COURSEEDITION .

Removing generalizations

- For the generalization on instructors:
 - the relevant operations make no difference between the child entities and these entities have no specific attributes;
 - we can therefore delete the child entities and add an attribute Type to the parent entity.
- For the generalization on trainees:
 - the relevant operations make no difference between the child entities but these entities have specific attributes;
 - we can therefore leave all the entities and add two relationships to link each child with the parent entity: in this way, we will have no attributes with possible null values on the parent entity and the dimension of the relations will be reduced.

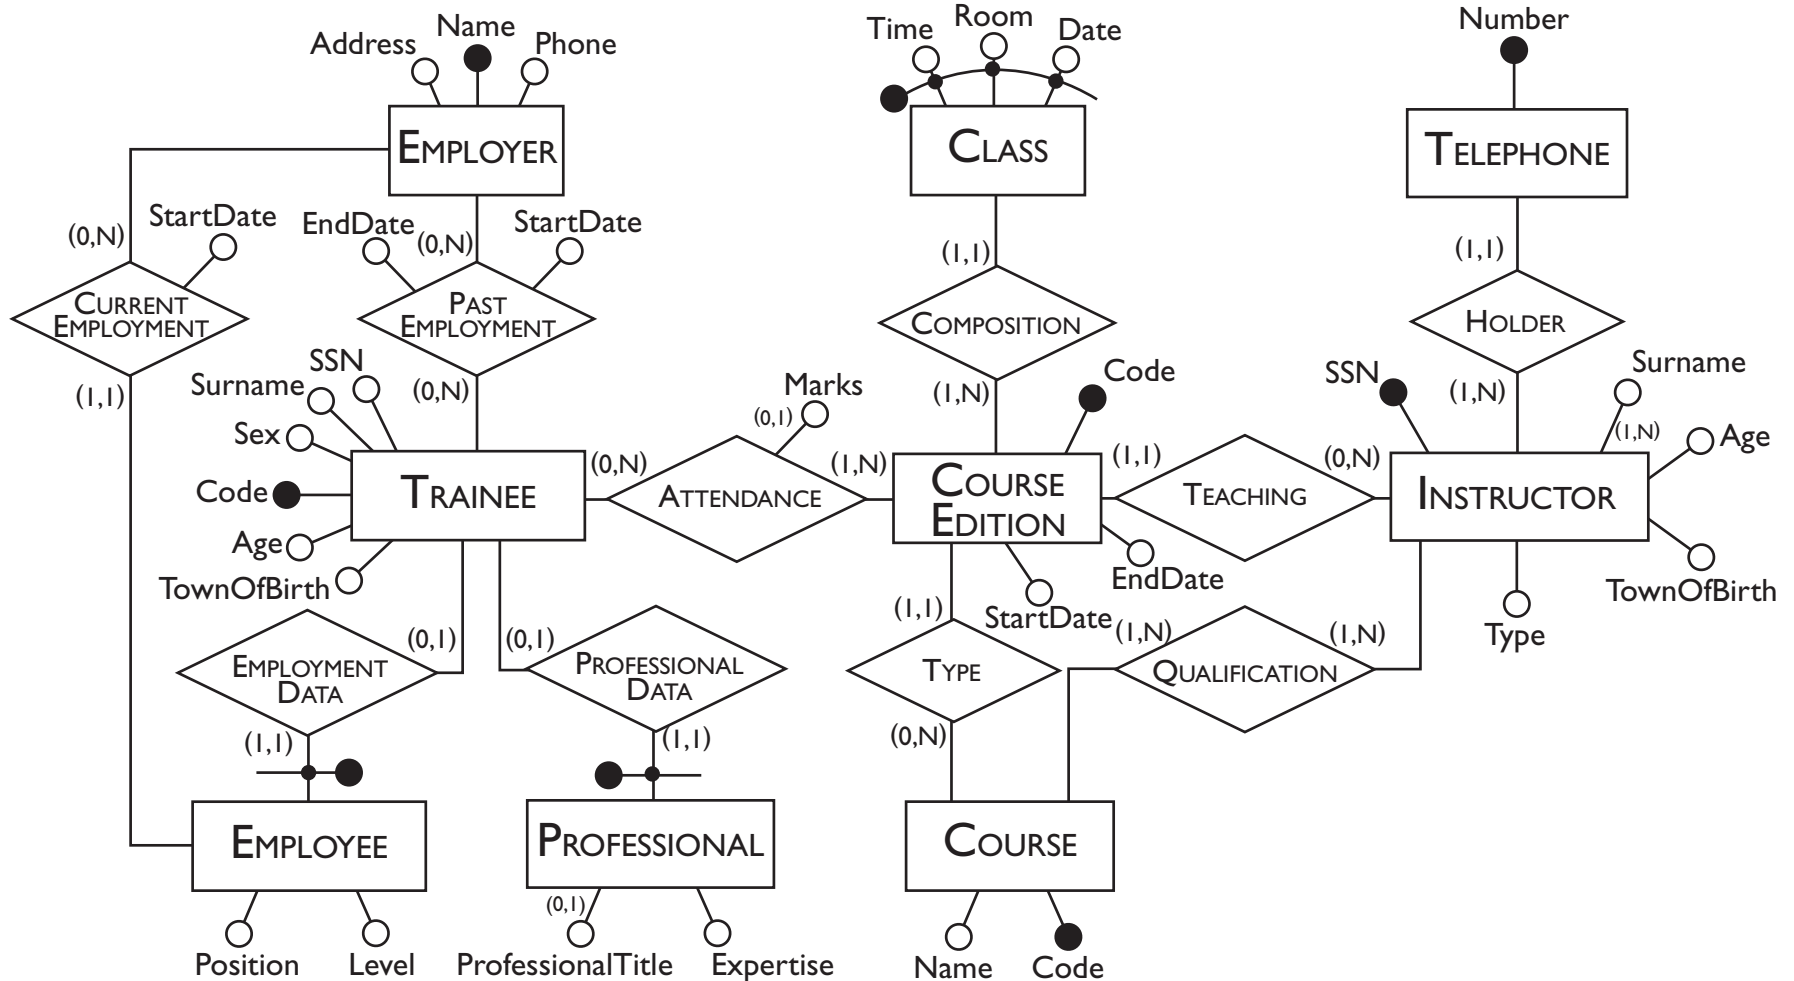
Partitioning and merging of concepts

- The relationships PASTTEACHING and PRESENTTEACHING can be merged since they describe similar concepts between which the operations make no difference. A similar consideration applies to the relationships PASTATTENDANCE and PRESENTATTENDANCE.
- The multi-valued attribute Telephone can be removed from the INSTRUCTOR entity by introducing a new entity TELEPHONE linked by a one-to-many relationship to the INSTRUCTOR entity.

Choice of main identifiers

- TRAINEE entity:
 - there are two identifiers: the social security number and the internal code;
 - it is far preferable to choose the latter: a social security number can require several bytes whereas an internal code, which serves to distinguish between 5000 occurrences, requires a few bytes.
- COURSEEDITION entity:
 - it is identified externally by the StartDate attribute and by the COURSE entity;
 - we can see however that we can easily generate for each edition a code from the course code: this code is simpler and can replace the external identifier.

The previous E-R schema after the restructuring phase



Translation into the relational model

COURSEEDITION(Code, StartDate, EndDate, Course, Instructor)

CLASS(Time, Room, Date, Edition)

INSTRUCTOR(SSN, Surname, Age, TownOfBirth, Type)

TELEPHONE(Number, Instructor)

COURSE(Code, Name)

QUALIFICATION(Course, Instructor)

TRAINEE(Code, SSN, Surname, Age, TownOfBirth, Sex)

ATTENDANCE(Trainee, Edition, Marks*)

EMPLOYER(Name, Address, Telephone)

PASTEMPLOYMENT(Trainee, Employer, StartDate, EndDate)

PROFESSIONAL(Trainee, Expertise, ProfessionalTitle*)

EMPLOYEE(Trainee, Level, Position, Employer, StartDate)

Logical design using CASE tools

- The logical design phase is partially supported by all database design tools:
 - the translation to the relational model it is carried out by these systems almost automatically;
 - the restructuring step is difficult to automate and the various products provide little or no support for it.
- All the systems are able to generate automatically the SQL code for the creation of the database.
- Some systems allow direct connection with a DBMS and can construct the corresponding database automatically.

Logical design with a CASE tool

The screenshot shows the Logic Works ERwin/ERX interface. The main window displays an Entity-Relationship (ER) diagram with the following entities and attributes:

- Employee Project**: Emp_Id: NUMBER, Name: VARCHAR2(20)
- Employee**: Emp_Id: NUMBER, Dept_Id: NUMBER, Name: VARCHAR2(20), Salary: NUMBER, Age: NUMBER
- Department**: Dept_Id: NUMBER, Name: VARCHAR2(20), Telephone: NUMBER
- Project**: Name: VARCHAR2(20), Budget: NUMBER, Deadline: DATE
- Manager**: Emp_Id: NUMBER, Dept_Id: NUMBER, Room: VARCHAR2(20), Name: VARCHAR2(20)
- Building**: Name: VARCHAR2(20), City: VARCHAR2(20), Address: VARCHAR2(20)

Relationships are indicated by lines connecting the entities. The 'Employee Project' entity is connected to the 'Employee' entity. The 'Employee' entity is connected to the 'Department' entity. The 'Project' entity is connected to the 'Employee Project' entity. The 'Manager' entity is connected to the 'Employee' entity. The 'Building' entity is connected to the 'Department' entity.

On the right side, the 'ORACLE Schema Generation Report' is displayed. The report title is 'Schema Generation Report'. The 'ORACLE 7.0 Schema Generation' section includes the following options:

- Referential Integrity: Primary Key (PK), Foreign Key (FK), ON DELETE, UNIQUE (AK)
- Schema Option: Pre-Script, Create Proced, TABLESPACE, DATABASE


The 'ORACLE Schema Generation Report Preview' window shows the following SQL code:

```
CREATE TABLE Employee (
  Emp_Id          NUMBER NOT NULL,
  Dept_Id         NUMBER NOT NULL,
  Name            VARCHAR2(20) NULL,
  Salary          NUMBER NULL,
  Age             NUMBER NULL,
  PRIMARY KEY (Emp_Id) );

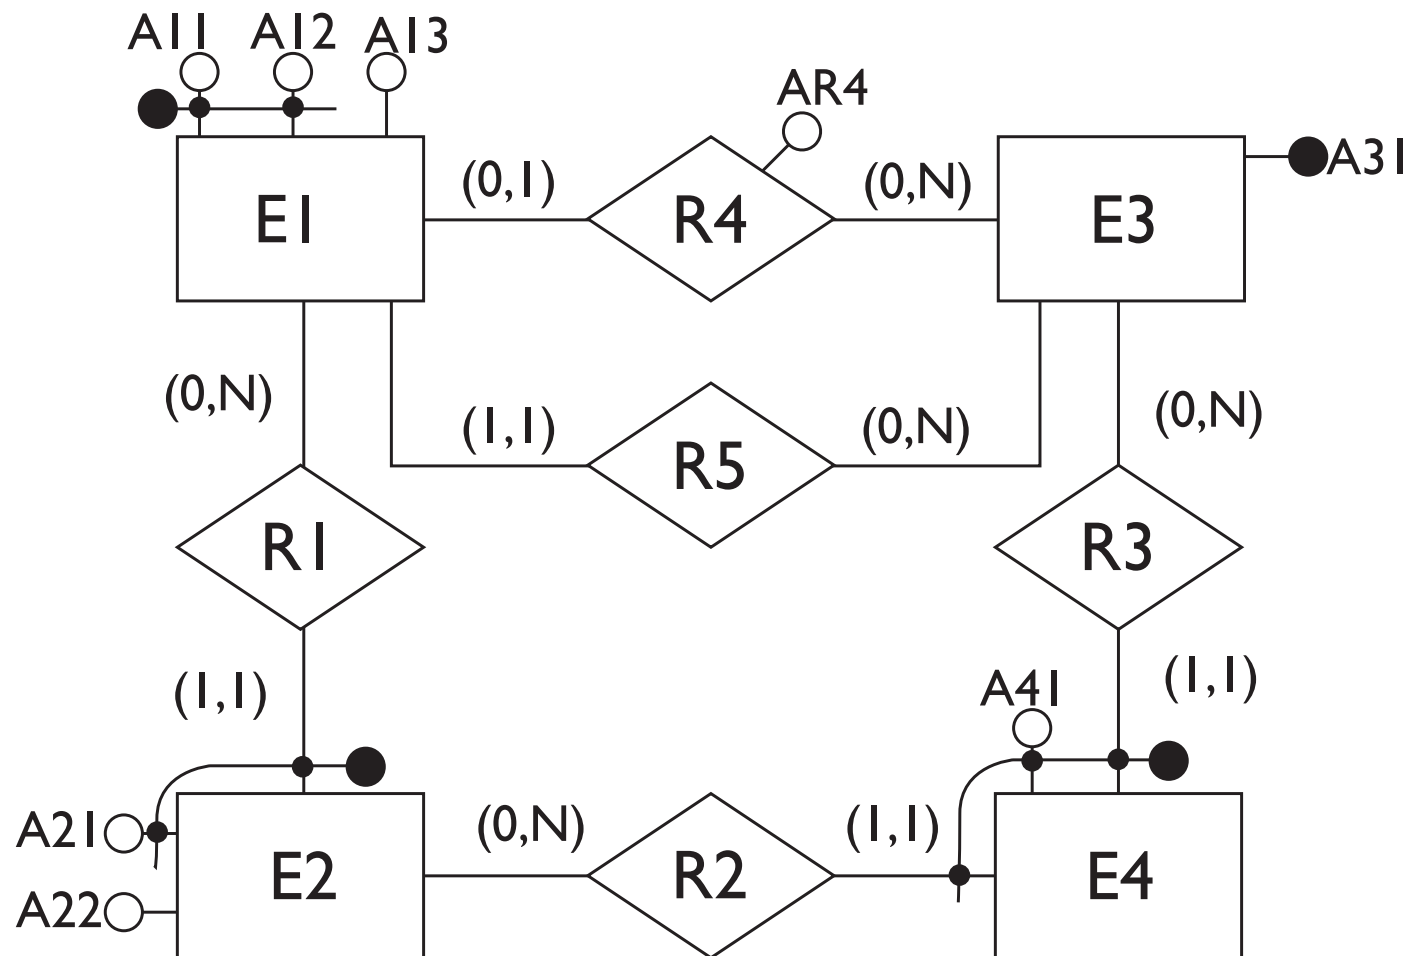
CREATE TABLE Project (
  Name            VARCHAR2(20) NOT NULL,
  Budget          NUMBER NULL,
  Deadline        DATE NULL,
  PRIMARY KEY (Name) );

CREATE TABLE Employee_Project (
  Emp_Id          NUMBER NOT NULL,
  Name            VARCHAR2(20) NOT NULL,
  PRIMARY KEY (Emp Id, Name) );
```


An E-R schema on the personnel of a company



An E-R schema with external identifiers



An E-R schema with generalizations



An E-R schema to translate

