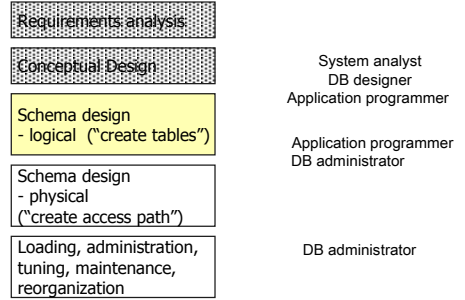


3. Schema Design:

Logical Design using the Relational Data Model

- 3.1 Logical Schema Design
 - 3.1.1 The Relational Data Model – Basics
 - 3.1.2 Keys, candidate keys and more
 - 3.2 From Conceptual to Logical Schema: Mapping ER to RDM
 - 3.2.1 Relationships to relations: a simple step
 - 3.2.2 Mapping weak entities and multivalued attributes
 - 3.2.3 Consolidation
 - 3.2.4 Mapping generalization hierachies and more
- Kemper/ Eickler: chap. 3.1-3.3, Elmasri / Navathe: chap. 9

Context



HS / DBS05-04-RDM1 2

3.1 Logical schema design

- Second phase of DB design
 - Transform the E-R model into the logical schema of a particular data model
 - Easy for Relational Data Model (RDM) can be performed automatically (e.g. Oracle Designer, Visio, and many other tools)
- The next steps:
 - Relational Data Model: Basics
 - General principles for mapping entities and relationships to relations (of the RDM)
 - Define the DB schema by means of a Data Definition Language (DDL) - SQL DDL in Relational DBS
 - Formal analysis of the relational schema

HS / DBS05-04-RDM1 3

3.1.1 The Relational Data Model - Basics

The Relational Data Model

- Simplicity and formal rigor as the guiding principle

KISS - Keep It simple, stupid
Do we need entities and relations?
No, just relations and attributes

- Relations are mathematical objects
- Relations may be implemented by tables
- Introduced by E.F. Codd, 1970, at that time at IBM Research labs, San Jose. Honored by the Turing award for his achievement.

HS / DBS05-04-RDM1 4

3.1.1 Basics of the RDM

Basic ideas

- Database is collection of relations
- Relation R = set of n-tuples
- Relation schema $R(a_1, a_2, \dots, a_n)$
- Attributes a_i have atomic values from domain $D_i = \text{dom}(a_i)$

Important terms

relation (table) →

attribute →

tuple →

Student			
fName	name	email	matrNo
Tina	Müller	mueller@...	13555
Anna	Katz	katz@...	12555
Carla	Maus	piep@...	11222

relation schema:

`Student(fname, name, email, matrNo)` or
`Student(fname:Name, name:Name, email:EmailType, matrNo:number)`

HS / DBS05-04-RDM1 5

RDM

• Notation

Relation $R \subseteq \text{dom}(a_1) \times \text{dom}(a_2) \times \dots \times \text{dom}(a_n)$

Attribute set $A = A(R) = \{a_1, a_2, \dots, a_n\}$

Tuple = $r \in R$

Degree of R: number of attributes

Relation Schema $R(a_1, a_2, \dots, a_n)$

• Terminology

different notations in use

Relation = table (file)

not the same: relations are sets.
tables may have duplicate entries

Tuple = row, record

Attribute = field (component)

table, row, record, field mainly used in the context of an RDBS implementation

Properties of the RDM

Properties

- No duplicate rows: R is a set
- No tuple order
- Database relations are time variant
update, insertion, deletion of tuples
- Integrity constraints must hold for all states over time
- Attributes have a primitive type, no constructed type
- Unique names in the relation and the DB namespace
otherwise dot-notation: R.a, db.S.b
- Attributes single-valued (more or less...)
- Attributes may have no value (NULL value)

HS / DBS05-04-RDM1 7

3.1.2 Keys, candidate keys and more

Key

Important term

- Each tuple identifiable by values of particular attributes ("key attributes")
- Key of $R(a_1, \dots, a_n)$: a minimal sequence of attributes which identify tuples

Example:

(first_Name, last_Name, birthdate, phone)
is a sequence of attributes which identify tuples of relation "Student"
It is not minimal: (first_Name, last_Name, birthdate)
will do

HS / DBS05-04-RDM1 8

RDM: key

Formally: $(k_1, \dots, k_k) \subseteq A = \{a_1, \dots, a_n\}$ is a key iff:

- 1) if v_1, \dots, v_k are the values of k_1, \dots, k_k of a row r_1 and v_1', \dots, v_k' for a row r_2 , and $r_1 \neq r_2$
 \Rightarrow there is at least one i , $1 \leq i \leq k$, $v_i \neq v_i'$
(identifying property)
- 2) there is no subset of $\{k_1, \dots, k_k\}$ with this property
(minimality)

each relation R has a key since R is a set
(in theory ...)

HS / DBS05-04-RDM1 9

Keys and Superkeys

- Frequently: a single artificial attribute, e.g. a number
- Similar to oid in OO systems
- Does not always make sense! (e.g. student administration system!)

Super key

A super set SK of a key of $R(a_1, \dots, a_n)$ with attributes $\{k_1, \dots, k_k\} \subseteq SK \subseteq \{a_1, \dots, a_n\}$ is called a super key of R

- Example `Student(fname, name, email, matrNo, major, birthdate)`
Key: `matrNo`, Super key: `{matrNo, name}`

HS / DBS05-04-RDM1 10

Primary Key, candidate key

Important terms

- A relation must have a key, but may have several.
- The keys of a table are called candidate keys*

Example:

`Student(fname, name, email, matrNo, major)`
Candidate Keys: `(Email)`, `(MatrNo)`

- Primary Key: an arbitrarily chosen candidate key
- Primary key access is "by value", not by location (oid) as in OO-languages

Example:

`Find Student.name where email='katz@inf.fu-berlin.de'`

* Schlüsselkandidaten, nicht Kandidatenschlüssel

HS / DBS05-04-RDM1 11

RDM: Foreign Key

Important term

How can rows of other tables be referenced?

Example:

`R: Student(fname, name, email, matrNo)`
`S: Exam (prof, std, subject, grade, dateTime)`

Exam.std should identify exactly one student.

Foreign key

A set of attributes FK in relation schema S is called a foreign key if

- attributes of FK have the same domain as the attributes of primary key pk of a relation R
- A value of FK in tuple t_s of S either occurs as a value of pk for some t_r in R or is NULL.

HS / DBS05-04-RDM1 12

E-R to RDM mapping: first step


• Case study after simple mapping

```
Tape(id, acqDate)
Movie(id, title, category, price_day,
      director, year, length)
Customer(mem_no, name, f_name, zip,
         city, streetNo)
Actor(stageName, real_name, birth_date)
Format(format, extraCh)
Rec_on (tapeId, movieId) // renaming
Has (tapeId, format) // renaming
Starring (stageName, movieId)
Rental ?? Weak entity ?!, phoneNo: MV attrib.
```

HS / DBS05-04-RDM1 19

Logical Design E-R to RDM mapping

Weak entities



- ▶ For each weak entity WE create relation R
- ▶ Include all attributes of WE
- ▶ Add key of identifying entity to weak entities (relative) key
- ▶ Part of key is a foreign key

Relational Schema:

```
Building(bid)
Rooms(cid, bid, area, ...)
```

Note: weak entities have to be transformed before relationships are transformed

HS / DBS05-04-RDM1 20

Logical Design E-R to RDM mapping

• Example

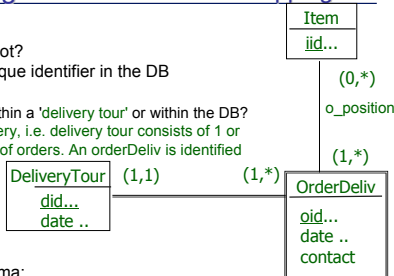
Weak entity or not?

Item: has a unique identifier in the DB

Order: ?

Is oid unique within a 'delivery tour' or within the DB?

Assumed: delivery, i.e. delivery tour consists of 1 or more deliveries of orders. An orderDeliv is identified within a tour.




Relational Schema:

```
DeliveryTour(did, date)
OrderDeliv(oid, did, date, contact)
O_Position(iid, oid, did)
Item(iid, ...)
```

HS / DBS05-04-RDM1 21

Weak entities: case study

Case study



```
Tape ( id, ... )
Customer ( name, mem_no, ... )
Rental ( id, from, until )
c_r ( id, from, mem_no )
```


Why is "Rental" not weak entity dependent on customer?

... application semantics: customer lends more than one tape with the same 'from'-time stamp.

HS / DBS05-04-RDM1 22

E-R to RDM mapping: Multiple values

Multiple value attribute \Rightarrow weak entity with a single attribute



Relations:

```
Customer ( name, mem_no, ... )
Phone ( mem_no, pNo )
```

... or array-type for attribute (PostgreSQL and others).

HS / DBS05-04-RDM1 23

3.2.3 E-R to RDM mapping: Consolidation

Simplification (consolidation) of the relational DB schemas

Rule:

Merge those relation schemas with the *same* key attributes into one relation schema

$$R(k_1, \dots, k_n, a_1, \dots, a_n), S(k_1, \dots, k_n, b_1, \dots, b_m)$$

\Rightarrow

$$RS(k_1, \dots, k_n, a_1, \dots, a_n, b_1, \dots, b_m)$$

Remark: Attribute semantics must match, not the literal name

```
Tape ( id, acqDate )           Tape ( id, ... )
Has ( id, format )           Recorded_on ( tapeId, movieId )
=> Tape ( id, acqDate, format ) => Tape ( id, acqDate, format, movieId )
```

E-R to RDM mapping: Simplification

Transformation

- unambiguous for relations representing 1:N
- 1:1 relationships: merge with one of the "entity-tables"
- M:N relationships: never merge
Representation in RDM always by a **separate table!**

Always merge 1:N relationships?

Example: `Person(id, name, ...)`
`Room(rNo, size, ...)`
`Sits_in(id, rNo, since, netSocket#, ...)`

Merge would result in a relation with many NULL values

`Person(id, name, rNo, since, netSocket#...)`

HS / DBS05-04-RDM1 29

E-R to RDM mapping: Simplification

Merging 1:N relationships

- makes sense in most cases
- if relationship has many attributes do not merge if many NULL values expected
- if attributes of relationship are used infrequently by applications, do not merge
This is an efficiency argument: avoid unnecessary data transfers


Never merge M:N relationships

~~`Person(id, name, ..., hobby, casualty...)`~~
~~`Hobby(hobby, kind, class_of_risk)`~~
~~`Has_H(id, hobby, casualty)`~~

Key changed, redundancy


HS / DBS05-04-RDM1 26

3.3.3 E-R to RDM mapping: Generalization



First alternative: One "big" A-table with attributes from all specializations

`A(x, y, s, t)`




HS / DBS05-04-RDM1 27

E-R to RDM mapping: Generalization

Second Alternative:

- different relations for A's, B's and C's
- make a one-to-one correspondence between every tuple from B and the appropriate A's
- ..and the same for the C's




B and C are separate relations
 In this example:
 - disjoint specializations
 - not exhaustive

`A(x, y)` Key of A as
`B(x, s)` foreign key in
`C(x, t)` B and C

HS / DBS05-04-RDM1 28

E-R to RDM mapping: Generalization

Third alternative



`AB(x, y, s)`
`AC(x, y, t)`


If subtypes are exhaustive, separate tables which include A's attributes, are a reasonable choice.

- First solution: may be many Null values
- Second solution: joining data from different tables (A and B or C) to get all about C or B may be time consuming
- Third solution: only valid for complete specializations

HS / DBS05-04-RDM1 29

E-R to RDM mapping

Recursive relationships



`Employee(eid, ..., managed_by_eid, ...)`

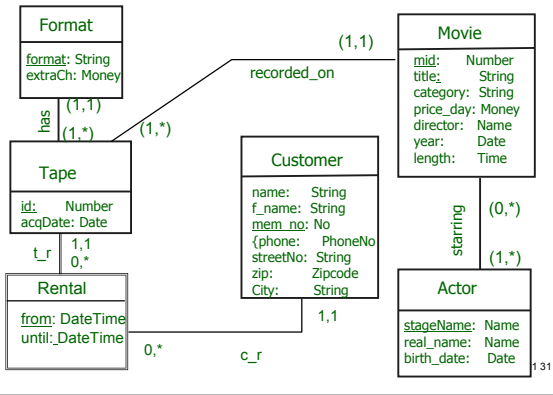
- Transformation step depending on cardinalities just like non-recursive relationships

N-ary relationships

- Like binary N:M relationships: Transform into one table with a composite key from the entities involved and attributes from the relationship, if any

HS / DBS05-04-RDM1 30

Case study revised



Logical Design E-R to RDM mapping

```

Tape(id, m_id, format, acqDate)
Movie(id, title, category, price_day,
      director, year, length)
Customer(mem_no, name, f_name, zip,
         city, streetNo)
Actor(stage_name, real_name, birth_date)
Format(format, extraCh)
Rental(t_id, from, until)
Phones(phoneNo, mem_no)
Starring(stageName, movieId)
    
```

HS / DBS05-04-RDM1 32

Summary ER to RDM (system independent)

- Represent **entities** by relations (tables)
- Treat **weak** entities and generalization in a special way
- Represent **relationships** by relations, keys depend on cardinality of relationship
- Simplify the abstract schema by folding relations having the same key.

Note:

- Lost nearly all constraints in the abstract relational schema
 - Most cardinalities are lost
 - No existential dependencies (as for weak entities)
 - No value restrictions for attributes (not part of E-R)
 - Key constraints survived
- Concrete SQL / DDL allows to specify most of them

Summary

- Relational data model
 - Representation of data as relations (tables)
 - Very simple structure has pros and cons (which?)
 - the most important data model today
- Important terms & concepts
 - Relation: set of n-tuples
 - Relation schema defines structure
 - Attribute: property of relation, atomic values
 - Superkey, candidate key, primary key identify tuple
 - Transformation rules for entities, relationships
 - Simplification

HS / DBS05-04-RDM1 34