Relational Databases Designing a Relational Database Summary and Revision There are two main ways of proceeding: - conceptual modelling (e.g. ER) • define a high level model of the data and derive the appropriate set of Relational databases consist of a set of relations (tables) in which tables - columns are named and typed normalisation - rows are identified only by the data they hold • list all the data in one table and then break it down into appropriate no two rows hold exactly the same data separate tables - one or more unique columns are chosen to identify the rows primary key The two techniques give rise to a similar set of tables There are only two ways of relating data - by putting the data in the same row of a table - by relating rows by making them hold a common value foreign keys MSc/Dip IT - ISD L19 Relations Revision (457-480) MSc/Dip IT - ISD L19 Relations Revision (457-480) 457 26/11/2009 458 26/11/2009

ER Modelling I

Identify the principal kinds of object you wish to hold data about - entity

types

- these become tables

Identify the attributes of the entity types

- each attribute holds one or more 'simple' values usually a number or short string
- one or more attributes becomes the primary key of the entity and thus the table
 - most entity types are self-contained (strong) and have such a key
 - although sometimes you have to invent an artificial new attribute to be this key
 - but some (weak) entity types have no primary key, but are dependent for uniqueness on another entity type and so must import part of their key from an owning entity type
- these become columns of the entity types
- multi-valued attributes become tables in their own right

459

MSc/Dip IT - ISD L19 Relations Revision (457-480)

ER Modelling II

Identify the relationships between the entity types

- each of these represents a set of connected entities
- most are between two entity types and so consist of pairs of entities, one from each type
- we need to state for each entity type participating in a relationship
 - if every entity must participate or not (total/partial participation)
 - if any entity can participate more than once
- the latter gives rise to three situations
 - each entity can only participate once on each side (1-1)
 - these become a foreign key column in one side, preferably a totally participating side
 - on one side an entity can participate only once, but on the other an entity can participate many times (1-N)
 - these become a foreign key column in the N side
 - entities on both sides can participate more than once (M-N)

460

- these become a separate table

Normalisation

Start by listing all the data as if it were in one big table and identify a primary key

Identify functional dependencies

- Split the tables up in order to reduce the duplication of data without losing the functional dependencies
 - if the primary key has more than one column in it and if any other column is dependent on **only part of the key** then a new table should be formed



- (<u>A</u>, <u>B</u>, C) and A -> C means every repetition of an A redundantly repeats a C
- make it $(\underline{A}, \underline{B})$ and (\underline{A}, C) note never throw away the key
- if there are columns which are indirectly dependent on the key then again there is redundant repetition of these columns
 (A, B, C) and A-> B and B -> C means every repetition of B

461



redundantly repeats Cmake it (A, B) and (B, C)

MSc/Dip IT - ISD L19 Relations Revision (457-480)

26/11/2009

Implementing a Relational Database

Use a GUI to create the tables, their columns and their constraints

Use SQL Create Tables

Manipulating a Relational Database

Data is added by:

- importing data
- using SQL insert
- using an application with a GUI and a form interface
 - the application will have some way of doing inserts for you

Data is updated

- using SQL update and delete
- using an application with a GUI and a form interface

463

Adding Constraints

The table structure constrains the data structure in some ways

It can further be constrained by further restrictions on the data:

- stating some columns are **unique** (i.e. never repeated in the table)
- stating some columns are not null (i.e. have a value for all rows)
- stating that one or more columns make up the primary key
 which amounts to being both unique and non-null
- stating that the range of values in a column are restricted (check)

462

- stating that one or more columns is a foreign key
 - i.e. takes the values from an equivalent number of columns in a table usually these are a primary key
 - e.g. if we have T1(<u>A, B</u>, C, ...) and we want to put a foreign key to this table in table T2, we have to add two columns, D and E, i.e.
 - T2(...., D, E, ...) and then

MSc/Dip IT - ISD L19 Relations Revision (457-480)

- (D, E) refers to T1(A, B)

26/11/2009



Relational Algebra

- The DBMS maps SQL select queries into relational algebra which is a series of steps
 - each producing a new (virtual) table or view
 - either reducing one table or view into some of the rows (select σ) or some of the columns (project Π)
 - or connecting two tables
 - joining ()) two tables puts together each pair of records from the tables which share a common value to form a single records containing all the data from each
 - this is the reverse of normalisation
 - outer joins create rows out of unpaired data filling them out with nulls
 - a **union** of two tables with the same number and same types of columns creates a table with all of the records which are in each
 - **intersection** and **difference** take two similar tables and produce the table with records which are in both or in only one

465

MSc/Dip IT - ISD L19 Relations Revision (457-480)

26/11/2009

Case Study: Dream Home Property Management

(from "Database Systems" by Connolly and Begg)

The database must support the following requirements:

1. Data Requirements:

- The company has a number of branches and staff working at those branches.
- The company maintains a number of properties (each managed by a specific member of staff), their owners and a list of clients who rent these properties.
- A client will view properties and may eventually commit to renting one, in which case a member of staff organises a lease.
- 2. Functional (Transaction) Requirements (extract):
 - List the addresses of all of the properties and those managed by a particular member of staff.
 - List all the clients viewing properties owned by a particular person.

166

- Identify clients who have inspected all properties with 3 rooms.
- Add the details of a new lease.

MSc/Dip IT - ISD L19 Relations Revision (457-480)

26/11/2009



- Owner id, name, address, tel
- Lease number, date_start, date_end, rent, payment_method, deposit

467

MSc/Dip IT - ISD L19 Relations Revision (457-480)

Creation of Tables I

- 1. For each entity type, create a table, viz: Staff, Client, Branch, Property, Owner, Lease
- 2. Add columns for each attribute, e.g.: Staff(staffNumber: number, name: text, tel: text, sex: text, title: text)
- 3. Identify keys:

Staff – staffNumber	Client – id	Branch – number
Property - number	Owner - id	Lease - number

But note, alternative candidate keys: Branch - tel or fax

Property – address

4. Identify other constraints on columns, e.g.: uniqueness – none obvious here

non-null - staff name, property address, etc.

checks – staffNumber is between 1000 and 1999 sex is 'F' or 'M'

468

Creation of Tables II

 5. For each 1-N relationship put a foreign key column in the N-side entity table, each of which refers to the primary key column: Manages - a column <i>managedBy</i> in <i>Property</i> Holds - a column <i>heldBy</i> in <i>Lease</i> Rents - a column <i>onProperty</i> in <i>Lease</i> Owns - a column <i>owner</i> in <i>Property</i> WorksAt - a column <i>worksAt</i> in <i>Staff</i> With - a column <i>branchOf</i> in <i>Property</i> 6. For each M-N relationship, create a new table: Views - with foreign key columns <i>client</i> and <i>property</i> and primary key also both columns, <i>client</i> and <i>property</i>. 7. For each relationship attribute, put a column in the table with the foreign key: In the <i>Views</i> table, add a column, <i>date</i>. 	<pre>create table Staff(staffNumber: Number constraint pk_staff primary key, name: Text not null, tel: Text, sex: Text constraint ck-sex check sex in ('M', 'F'), title: Text, worksAt: Number constraint fk_worksat references Branch(number)) create table Branch(number: Number constraint pk_branch primary key, address: Text not null, tel: Text unique not null, fax: Text unique not null)</pre>
MSc/Dip IT – ISD L19 Relations Revision (457-480) 469 26/11/2009	MSc/Dip IT – ISD L19 Relations Revision (457-480) 470 26/11/2009

The Resulting Tables II

create table Client(
 id: Number constraint pk_client primary key,
 name: Text not null,
 address: Text not null,
 tel: Text not null)

471

The Resulting Tables III

The Resulting Tables I

<pre>create table Lease(number: Number constraint pk_lease primary key,</pre>
date_start: Text not null,
date_end: Text,
rent: Number not null ,
payment_method: Text not null ,
onProperty: Number constraint fk_on references Property(number),
deposit: Number,
heldBy: Number constraint fk_heldBy references Client(id))
create table Owner(id: Number constraint pk_owner primary key, name: Text not null, address: Text, tel: Text not null)
create table Views(
client: Number constraint fk_vclient references Client(id),
property: Number constraint fk_vprop references Property(number),
date: Text,
constraint pk_views primary key (client, property))

472

MSc/Dip IT - ISD L19 Relations Revision (457-480)

The Role of Normalisation

The Universal Relation would be:

U(leaseNumber, staffNumber, staffName, staffTel, staffSex, staffTitle, branchNumber, branchAddress, branchTel, branchFax, date_start, date_end, leaseRent, payment_method, deposit, clientID, clientName, clientAddress, clientTel, propertyNumber, propertyAddress, type, rooms, propertyRent, ownerID, ownerName, ownerAddress, ownerTel, viewingDate)

The primary key is hard to find – the minimal set is: clientID, propertyNumber since everything is determined by these.

The Functional Dependencies

Functional dependencies include:

- staffNumber -> staffName, staffTel, staffSex, staffTitle, branchNumber, branchAddress, branchTel, branchFax
- branchNumber -> branchAddress, branchTel, branchFax
- clientID -> clientName, clientAddress, clientTel
- leaseNumber -> everything except viewing date
- propertyNumber -> propertyAddress, type, rooms, propertyRent, ownerID, ownerName, ownerAddress, ownerTel, managingStaffNumber, managingStaffName, etc.
- ownerID -> ownerName, ownerAddress, ownerTel
- clientID, propertyNumber -> viewingDate, leaseNumber

MSc/Dip IT - ISD L19 Relations Revision	n (457-480)
---	-------------

26/11/2009

MSc/Dip IT - ISD L19 Relations Revision (457-480)

26/11/2009

Second Normal Form

473

Second Normal Form splits off those relations dependent upon part of any candidate key:

Client(clientID, clientName, clientAddress, clientTel)

- Property(<u>propertyNumber</u>, propertyAddress, type, rooms, propertyRent, ownerID, ownerName, ownerAddress, ownerTel)
- ClientProperty(leaseNumber, staffNumber, staffName, staffTel, staffSex, staffTitle, branchNumber, branchAddress, branchTel, branchFax, date_start, date_end, leaseRent, payment_method, deposit, <u>clientID</u>, <u>propertyNumber</u>, viewingDate)

475

Third Normal Form

171

Third normal form splits off transitive dependencies, split *Property* into: Property(<u>propertyNumber</u>, propertyAddress, type, rooms, propertyRent) Owner(<u>ownerID</u>, ownerName, ownerAddress, ownerTel)

From *ClientProperty* split off :

Staff(staffNumber, staffName, staffTel, staffSex, staffTitle,

branchNumber) // aka worksAt

Branch(<u>branchNumber</u>, branchAddress, branchTel, branchFax) Lease(<u>leaseNumber</u>, staffNumber, date start, date end, leaseRent,

payment_method, deposit)

leaving:

ClientProperty(clientID, propertyNumber, viewingDate, leaseNumber)

476

A Difference!



Smith.	
select Client.name	$JS \leftarrow \boldsymbol{\sigma}_{name = "Jane Smith"} (Owner)$
from Client, Views, Property, Owner	$JSprops \leftarrow JS \longrightarrow Property on \\ owner = id$
where Owner.name = "Jane Smith"	JSviews ← JSprops ▷ Views on

479

2 List all the clients viewing properties owned by a particular person e.g. Jane

and Property.owner = owner.id and Views.property = JSvClie

Property.number

and Views.client - client.id

JSviews ← JSprops ▷ Views on property = number JSvClients ← JSviews ▷ Client on client = id

 $\pi_{\text{clientname}}$ (JSviewingClients)

3. Identify clients who have inspected all properties with three rooms.

select C.name	3 rooms $\leftarrow \sigma_{rooms=3}$ (Property)
from Client C	3 roomsNos $\leftarrow \pi_{\text{number}}$ (3 rooms)
where not exists	reduceView $\leftarrow \pi_{\text{client, property}}$ (Views)
((select P.number from Property P	3roomClientNos ← reduceView ÷ 3roomsNos
where P.rooms = 3)	3 roomClients \leftarrow Client \triangleright 3 roomClientNos
minus	on $id = client$
(select V.property from Views V	result $\leftarrow \pi_{\text{name}}$ (3roomClients)
where C.id = V.client))	

Implementing the Transactions

4. Add the details of a new lease.

The Relational Algebra doesn't deal with updates.

insert into Lease **values**(1234, "1/1/05", "31/12/05", 240, "monthly", 432, 500, 16543, 132)

which creates a new Lease numbered 1234 made by staff member 16543 for client 132 on property 432.

480

MSc/Dip IT - ISD L19 Relations Revision (457-480)