

Problems with the relational model

Problems (1)

To store information about real life entities, we often have to cut them up into separate tables

DreamHome						
Page 1	Customer Rental Details					Date 7-Oct-98
Customer Name John Kay			Customer Number CR76			
Property Number	Property Address	Rent Start	Rent Finish	Rent	Owner Number	Owner Name
PG4	6 Lawrence St, Glasgow	1-Jul-94	31-Aug-96	350	CO40	Tina Murphy
PG16	5 Novar Dr, Glasgow	1-Sep-96	1-Sep-98	450	CO93	Tony Shaw

Problems (1)

To store this information efficiently we'd use 4 tables

Customer	
Customer No	Cname
CR76	John Kay
CR56	Aline Stewart

Rental				
Customer No	Property No	RentStart	RentFinish	
CR76	PG4	1-Jul-94	31-Aug-96	
CR76	PG16	1-Sep-96	1-Sep-08	
CR56	PG4	1-Sep-92	10-Jun-94	
CR56	PG36	10-Oct-94	1-Dec-95	
CR56	PG16	1-Jan-96	10-Aug-96	

Property			
Property No	Address	Rent	Owner No
PG4	6 Lawrence St, Glasgow	350	CO40
PG36	2 Manor Rd, Glasgow	375	CO93
PG16	5 Novar Dr, Glasgow	450	CO93

Owner	
Owner No	Oname
CO40	Tina Murphy
CO93	Tony Shaw

Problems (1)

To answer any questions we have to put them back together again!

select *
from customer, rental, property, owner
where

Customer	
Customer No	Cname
CR76	John Kay
CR56	Aline Stewart

Rental			
Customer No	Property No	RentStart	RentFinish
CR76	PG4	1-Jul-94	31-Aug-96
CR76	PG16	1-Sep-96	1-Sep-08
CR56	PG4	1-Sep-92	10-Jun-94
CR56	PG36	10-Oct-94	1-Dec-95
CR56	PG16	1-Jan-96	10-Aug-96

Property			
Property No	Address	Rent	Owner No
PG4	6 Lawrence St, Glasgow	350	CO40
PG36	2 Manor Rd, Glasgow	375	CO93
PG16	5 Novar Dr, Glasgow	450	CO93

Owner	
Owner No	Oname
CO40	Tina Murphy
CO93	Tony Shaw

Problems (1)

select *
from customer, rental, property, owner
where

Cartesian Product is the most 'expensive' operation

Takes time

Requires lots of memory

Requires lots of processor power

Note

Thankfully, lots of research work in this area to improve the efficiency

Solution - Object Oriented Databases

Store information as a whole item - called an **Object** (in the **Object data model**)

DreamHome						
Page 1	Customer Rental Details				Date 7-Oct-98	
Customer Name John Kay			Customer Number CR76			
Property Number	Property Address	Rent Start	Rent Finish	Rent	Owner Number	Owner Name
PG4	6 Lawrence St, Glasgow	1-Jul-94	31-Aug-96	350	CO40	Tina Murphy
PG16	5 Novar Dr, Glasgow	1-Sep-96	1-Sep-98	450	CO93	Tony Shaw

Problems (2) - Temporal information

The relational model is not very good at storing many 'states' of information (unless explicitly structured)

student

kulD	name	term time address
k0699345	Fred Tester	17 Penrhyn Road
k0545646	Martin Qwerty	20 Richmond Road
k0665665	Mary Smith	34 Kingston Road

Problems (2) - Temporal information

What happens when Martin moves in with Fred?

student

kulD	name	term time address
k0699345	Fred Tester	17 Penrhyn Road
k0545646	Martin Qwerty	20 Richmond Road
k0665665	Mary Smith	34 Kingston Road

Problems (2) - Temporal information

Query: Show all the places Martin has ever lived

Problem: We can't

student

kulD	name	term time address
k0699345	Fred Tester	17 Penrhyn Road
k0545646	Martin Qwertz	17 Penrhyn Road
k0665665	Mary Smith	34 Kingston Road

Problems (2) - Temporal information

Query: When did Martin move house?

Problem: Can't answer

student

kulD	name	term time address
k0699345	Fred Tester	17 Penrhyn Road
k0545646	Martin Qwertz	17 Penrhyn Road
k0665665	Mary Smith	34 Kingston Road

Problems (2) - Temporal information

Query: When did Martin tell us that he'd move house?

Problem: Can't answer

student

kulD	name	term time address
k0699345	Fred Tester	17 Penrhyn Road
k0545646	Martin Qwertz	17 Penrhyn Road
k0665665	Mary Smith	34 Kingston Road

Solution - Temporal databases / temporal SQL

Databases explicitly structured to store old information

Databases explicitly structured to note when information changes

Problems (3) - Truth and probability

The relational model only stores information that is **100% true**

student

kulID	name
k0699345	Fred Tester
k0545646	Martin Qwerty
k0665665	Mary Smith

By having these rows in the table...

...we are asserting that this information is **true**

Problems (3) - Truth and probability

student

kulID	name
k0699345	Fred Tester
k0545646	Martin Qwerty
k0665665	Mary Smith

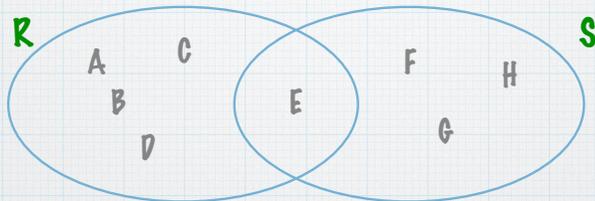
Fred is a student
Martin is a student
Mary is a student

All these statements are **true** - all other information is false

The Closed World Assumption (CWA)

The Closed World Assumption (CWA)

Unless explicitly included, everything else is **false**



Is P in R? No.

Problems (3) - Truth and probability

All other information is **false**

Query: Is Robyn a student?

student

kulID	name
k0699345	Fred Tester
k0545646	Martin Qwerty
k0665665	Mary Smith

Robyn isn't here

...so Robyn isn't a student

Problems (3) - Truth and probability

But what if we want to store **negative** or **false** information

Robyn is to **NEVER** be a student

student

kulD	name
k0699345	Fred Tester
k0545646	Martin Qwerty
k0665665	Mary Smith

Can't do it

Problems (3) - Truth and probability

But what if we want to store information that is **probably true**

Its **90%** probable that Robyn is a student

student

kulD	name
k0699345	Fred Tester
k0545646	Martin Qwerty
k0665665	Mary Smith

Can't do it

Solution - Deductive databases

Databases explicitly structured to store rules and facts

These can store complex statements of truth and false

Problems (4) - Some queries can't be expressed in SQL

Find all the line managers for every employee

EMPNO	ENAME	JOB	MGR	HIREDATE	SAL	COMM	DEPTNO
405	MARCH	ADMIN	938	13/06/1997	18000		2
535	BYRNE	SALES	734	15/08/1997	26000	300	3
557	BELL	SALES	734	26/03/2000	22500	500	3
602	BIRD	MANAGER	875	31/10/1997	39750		2
690	AHMAD	SALES	734	05/12/1997	22500	1400	3
734	COX	MANAGER	875	11/06/2002	38500		3
818	POLLARD	MANAGER	875	14/05/2000	34500		1
824	REES	ANALYST	602	05/03/2000	40000		2
875	PARKER	PRESIDENT		09/07/2002	60000		1
880	TURNER	SALES	734	04/06/2001	25000	0	3
912	HAYES	ADMIN	824	04/06/2001	21000		2
936	CASSY	ADMIN	734	23/07/2002	19500		3
938	GIBSON	ANALYST	602	05/12/1997	40000		2
970	BLACK	ADMIN	818	21/11/1997	23000		1

Black - Managers: Pollard, Parker

March - Managers: Gibson, Bird, Parker

Solution - Extensions to SQL

Extensions to the language to allow this kind of loop processing

Problem: Clutters the language