

# Databases - 5

## Problems with the relational model Functions and sub-queries

### 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-98	
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-98
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	Paddress	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)

Page 1		DreamHome			Date 7-Oct-98	
Customer Rental Details						
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

kuID	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

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

## Problems (2) - Temporal information

Query: Show all the places Martin has ever lived

Problem: We can't

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

Problem: Can't answer

## 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

kulD	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

kulD	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

kulD	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

kuID	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

## Functions

Transform a value or set of values using some rule

Built into the SQL standard

**Problem:** Microsoft uses many of its own function names to maintain compatibility with Excel and Word functions

## Functions Categories

String	concatenation, length, substring
Arithmetic	max, min, power, round, trunc
Date	add, subtract dates
Aggregate or group	average, sum, count

## String Function - String Concatenation & Combines fields with additional text if required

select "Employee "&ename&"  
 earns £"&sal  
 from emp

Query1

Expr1000
Employee MARCH earns £18000
Employee BYRNE earns £26000
Employee BELL earns £22500
Employee BIRD earns £39750
Employee AHMAD earns £22500
Employee COX earns £38500
Employee POLLARD earns £34500
Employee REES earns £40000
Employee PARKER earns £60000
Employee TURNER earns £25000
Employee HAYES earns £21000
Employee CASSY earns £19500
Employee GIBSON earns £40000
Employee BLACK earns £23000

## String Function - Substring

mid (string, starting point, no of chars) - returns part of a string

select ename, mid (ename, 2, 4)  
 from emp

02 Sub string

ename	Expr1001
MARCH	ARCH
BYRNE	YRNE
BELL	ELL
BIRD	IRD
AHMAD	HMAD
COX	OX
POLLARD	OLLA
REES	EES
PARKER	ARKE
TURNER	URNE
HAYES	AYES
CASSY	ASSY
GIBSON	IBSO
BLACK	LACK

Record: 15 of 15

## Arithmetic Function - min

min () - returns smallest value in a column

SAL
18000
26000
22500
39750
22500
38500
34500
40000
60000
25000
21000
19500
40000
23000

select min (sal)  
 from emp

03 ...

Expr1000
18000


Record: 1 of 1



## Arithmetic Function - max

`max ()` - returns largest value in a column

```
select max (sal)
from emp
```




SAL
18000
26000
22500
39750
22500
38500
34500
40000
60000
25000
21000
19500
40000
23000

## Aggregate Function - avg

`avg ()` - returns mean value in a column

```
select avg (sal)
from emp
```




SAL
18000
26000
22500
39750
22500
38500
34500
40000
60000
25000
21000
19500
40000
23000

## Aggregate Function - sum

`sum ()` - returns total of all values in a column

```
select sum (sal)
from emp
```




SAL
18000
26000
22500
39750
22500
38500
34500
40000
60000
25000
21000
19500
40000
23000

## Aggregate Function - count

**count ()** - returns total number of values in a column

**select count (sal)**  
**from emp**



SAL
18000
26000
22500
39750
22500
38500
34500
40000
60000
25000
21000
19500
40000
23000

## Group by

Aggregate functions can be applied to subsets of the table by using the **group by** syntax

**select \* or expression**  
**from relations**  
**[where expression]**  
**[group by expression]**

## Group by

Aggregate functions can be applied to subsets of the table by using the **group by** syntax

e.g.

**Calculate avg()**

**Calculate avg()**

**Calculate avg()**

ENAME	JOB	MGR	HIREDATE	SAL
BLACK	ADMIN	818	21/11/1997	23000
CASSY	ADMIN	734	23/07/2002	19500
HAYES	ADMIN	824	04/06/2001	21000
MARCH	ADMIN	938	13/06/1997	18000
GIBSON	ANALYST	602	05/12/1997	40000
REES	ANALYST	602	05/03/2000	40000
POLLARD	MANAGER	875	14/05/2000	34500
COX	MANAGER	875	11/06/2002	38500
BIRD	MANAGER	875	31/10/1997	39750
PARKER	PRESIDENT		09/07/2002	60000
TURNER	SALES	734	04/06/2001	25000
AHMAD	SALES	734	05/12/1997	22500
BELL	SALES	734	26/03/2000	22500
BYRNE	SALES	734	15/08/1997	26000

## Group by

Aggregate functions can be applied to subsets of the table by using the **group by** syntax

```
select job, avg(sal)
from emp
group by job
```

job	Expr1001
ADMIN	20375
ANALYST	40000
MANAGER	37583.333333
PRESIDENT	60000
SALES	24000

## Group by

Find the highest salary for each job category

```
select job, max(sal)
from emp
group by job
```

job	Expr1001
ADMIN	23000
ANALYST	40000
MANAGER	39750
PRESIDENT	60000
SALES	26000

## Nested sub-queries

When one of the conditions of a **WHERE** clause is a query itself, this is called a nested sub-query, i.e.,

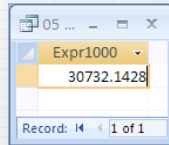
```
SELECT select-list
FROM table(s)
WHERE object operator (SELECT select-list
FROM table(s)
[WHERE condition]);
```

## Nested sub-queries example

Find all the employees who earn more than the average salary

Start by finding the average salary

```
select avg(sal)
from emp
```



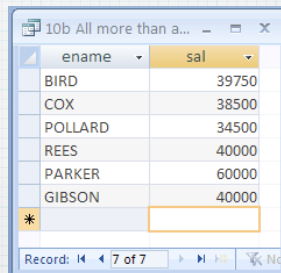
A screenshot of a database query window titled '05 ...'. It shows a single row with the value '30732.1428' under the column 'Expr1000'. The status bar at the bottom indicates 'Record: 1 of 1'.

## Nested sub-queries example

Find all the employees who earn more than the average salary

Now find all the employees that earn more than this:

```
select ename, sal
from emp
where sal >= ( select avg(sal)
               from emp )
```



A screenshot of a database query window titled '10b All more than a...'. It shows a table with two columns: 'ename' and 'sal'. The rows are: BIRD (39750), COX (38500), POLLARD (34500), REES (40000), PARKER (60000), and GIBSON (40000). There is an asterisk (\*) in the first column of the last row. The status bar at the bottom indicates 'Record: 1 of 7'.

## Nested sub-queries example (2)

Typically, you need to use a nested sub-query where you need to use an aggregate function and an attribute at the same time

Find the employee who earns the least money

Attribute required

Aggregate function required

## Nested sub-queries example (2)

Find the employee who earns the least money

First attempt may try something like this:

```
select ename, min(sal)
from ...
where ...
```

many values

single value

This can't work

## Nested sub-queries example (2)

Find the employee who earns the least money

Better attempt:

```
select min(sal)
from emp
```

Find the minimum salary

## Nested sub-queries example (2)

Find the employee who earns the least money

Better attempt:

```
select ename, sal
from emp
where sal = (select min(sal)
            from emp)
```

Find the person who has the minimum salary

ename	sal
MARCH	18000
*	

Record: 1 of 2 of 2