



# Trigger happy

Mark Whitehorn fires off a few shots of wisdom about the new concepts in client-server systems, and updates you on the Bloor research report on Scalability in Databases.

**T**he wonderful thing about Triggers, is Triggers are wonderful things" (with apologies to AA Milne.)

Client-server systems are becoming easier to install and use. They bring with them a whole raft of new ideas, concepts and jargon that *must* be absorbed by an aspiring database professional, and we will look at some of these newcomers over the next few months.

The information presented here is not meant to be product specific, so it should apply generally to SQL Server, DB2, Informix, Sybase etc. Clearly there will be some implementational differences, but those should be relatively minor.

PC-based RDBMSs like Access combine the front-end and back-end into a more-or-less seamless whole. So when you want to ensure that a business rule is enforced within an Access database, it is common to put the code that enforces that rule into, say, a form. In a client-server system, the front-end application is divorced from the back-end database. This separation may be so complete that the person designing and building the back-end may have no idea what front-ends will be run against their work in the future, so it is necessary to have a mechanism for enforcing business rules within the database itself.

Incidentally, this certainly is not meant to

imply that front-end applications on client-server systems cannot, or should not, enforce business rules. You might well want to apply different business rules for different groups of users, and one way to do this is to supply different front-end applications which apply different business rules. However, the big advantage of rules enforced on the server is that they cannot be subverted by any front-end application and so they are applied absolutely.

There are several ways of implementing business rules within a back-end database; triggers are used by many RDBMSs. Triggers are made up of parts, and the parts may differ slightly between implementations, but typically you may find

that a trigger has: a name, a triggering operation (i.e. the event that sets it off), an activation time, granularity, a body, and conditions. These parts are all entirely logical as soon as you think about what a trigger is supposed to do.

The easiest way of explaining one is to look at an example. Suppose you have a (very sensible) business rule which states that the details of a given customer must not be deleted from the customer table if that customer still owes money to your company.

This trigger has to look for every SQL statement that tries to delete customers from the CUSTOMER table. The trigger must operate before the record is deleted. It must first check that the customer is debt free and only allow the deletion to occur if that is true. Finally, if the delete statement tries to delete multiple records at a time (DELETE \* FROM CUSTOMERS) in a single statement, then the trigger must run multiple times, checking the debt status of each customer that the DELETE statement tries to remove.

The parts outlined above can be tied to a description of this trigger. It might be called Check\_CUST\_Del (name). It looks for any DELETE operation that occurs against the table (the triggering operation). When that occurs, the debt check is carried out (the body) before the DELETE completes (the activation time; the trigger is activated before the delete). If the DELETE runs against multiple rows in the table, the trigger is run once for each row in the table that's affected (the granularity).

If we modified this trigger so that it only applied the check if the customer had placed orders of a value greater than £10,000, then that would be a condition. Defining a condition for a trigger is optional.

## More on Bloor

Those of you who have been worrying about the implications of the Bloor report on Scalability in Databases (see Hands On Databases, August) can stop worrying; the report essentially no longer exists.

So what about the facts described in the report? They do not appear to have changed. Bloor did find that Microsoft's SQL Server did not scale well, did find the other points of concern that were covered in my August column, and the company still stands by those findings.

What has changed is that the report itself is no longer an available part of the sum

## Business rules

Databases have all sorts of generic integrity checks that can be applied: referential integrity, not nullable, and so on. These are "generic" integrity rules in that they have the potential to be applied to all databases. In addition to these generic integrity checks, there are constraints that you may wish to apply to a specific database.

It may be a rule in your company that no-one below the level of a manager can be given a pay rise of more than ten percent. It may be a rule that all orders with a value of more than £10,000 have to be approved by the MD. These are known as "business rules". They can be just as important as the generic varieties but clearly have to be handled by some other mechanism, such as triggers.

total of human knowledge. If you do not have a copy of it, your chances of obtaining all the information therein are slim indeed.

Here are some more facts, so far as I have been able to ascertain them. The report looked at RDBMSs from two companies: IBM and Microsoft. Bloor claims that both companies provided software and support people to set up the RDBMSs for testing; Microsoft disputes this, and claims that it did not supply people to help with the test.

The licence for SQL Server says (in broad legalese) that you must have Microsoft's written permission before publishing the results of testing. Bloor did not obtain this permission, and says that it took Microsoft's assistance in the project as tacit acceptance of the fact that the results would be published.

A couple of months after publication of the report, Bloor received a "cease and desist" letter from Microsoft's lawyers, which states that Bloor is in breach of contract because it failed to obtain the required written permission. Since Bloor had indeed failed to do so, it felt it had no option but to cease distribution of the report. A reasonable person might ask why, if Microsoft considered that the report was inaccurate, it didn't try to prevent the report being published on those grounds rather than on what appears to be a technicality?

The greatest loss from all of this may not be the information contained in the report, but the methods it used for testing databases. Testing for speed is a complex area, too often treated in a trivial way. This Bloor report was (in my opinion) a genuine contribution to our attempts to derive a

## Tips and Tricks: Access shortcuts and ordinal numbers

■ Shane Devenshire is a regular contributor to Tips and Tricks. He lives in Walnut Creek, California. In Access 95 or 97, here are a few shortcuts you might find interesting:

In the File Open, Save or Save As dialog boxes:

1. F4 drops down the Save In or Look In list.
2. F5 refreshes the screen.
3. Del (Delete) deletes a file to the Recycle bin.
4. Shift+Del deletes a file without sending it to the recycle bin.
5. Single-click the file, then single-click it again; this allows you to rename the file.
6. Right-clicking the file will display a shortcut menu with, among other things, Quick View and Send To.
7. With the Details option on, you can click the top of the columns to sort each field.
8. With Details on, you can also best fit the columns by double-clicking on the line between the column titles.

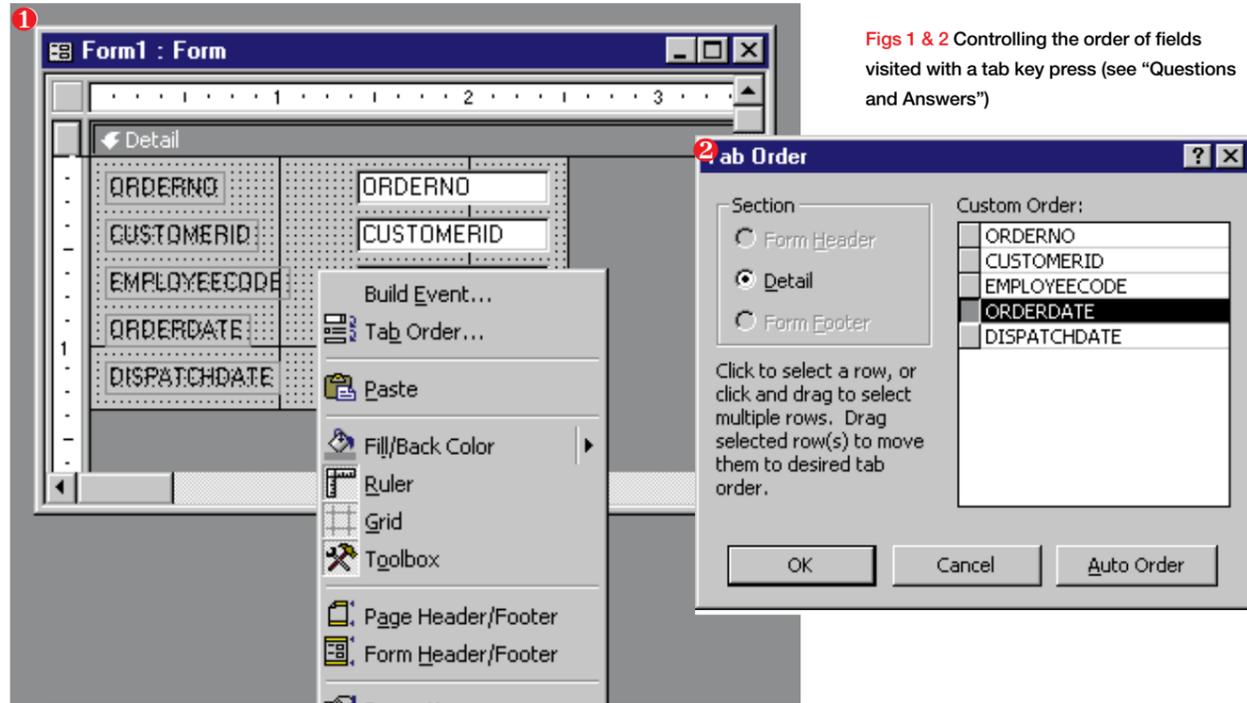
■ James Talbut came up with a problem and a solution, but wonders whether anyone knows of a better version?

"I would like dates to be formatted using ordinal numbers, hence 'Monday 2nd June 1997'. I cannot find any way of using number formats to produce ordinal numbers and have had to resort to the function shown (see *opposite*) which is neither quick nor elegant. Any comments?"

```
Public Function FormatDate(dArg As Date) As String
Dim sRetVal As String
Dim vOrdinals As Variant
Dim iDay As Integer
vOrdinals = Array("th", "st", "nd", "rd")
sRetVal = Format(dArg, "dddd, d")
iDay = Day(dArg)

If iDay < 4 Then
sRetVal = sRetVal + vOrdinals(iDay)
ElseIf iDay < 20 Then
sRetVal = sRetVal + "th"
ElseIf iDay Mod 10 < 4 Then
sRetVal = sRetVal + vOrdinals(iDay Mod 10)
Else
sRetVal = sRetVal + "th"
End If

FormatDate = sRetVal & Format(dArg, " mmmm, yyyy")
End Function
```



**Figs 1 & 2** Controlling the order of fields visited with a tab key press (see "Questions and Answers")

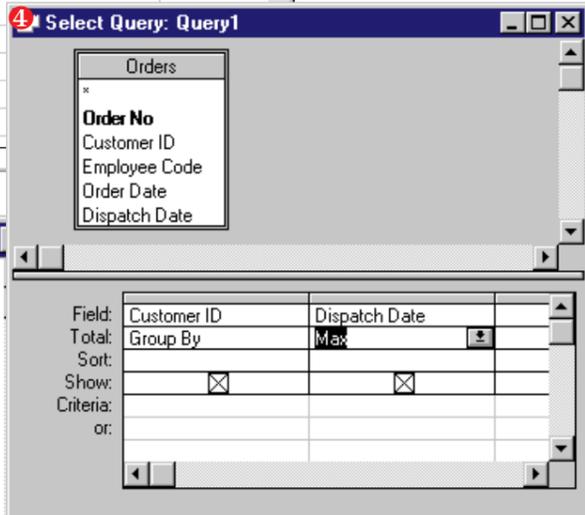
**3** Table: Orders

Order No	Customer ID	Employee Code	Order Date	Dispatch Date
1	274	45	01-Apr-93	10-May-93
2	3884	75	01-Apr-93	16-Apr-93
3	2436	12	01-Apr-93	20-Apr-93
4	4083	21	01-Apr-93	14-Apr-93
5	1460	85	01-Apr-93	09-Apr-93
6	182	10	01-Apr-93	12-Apr-93
7	964	46	01-Apr-93	19-Apr-93
8	7749	147	01-Apr-93	
9	5829	89	01-Apr-93	
10	5975	37	01-Apr-93	
11	2557	159	01-Apr-93	
12	5676	91	01-Apr-93	

**Figs 3 - 6** Keeping track of customer complaints: Finding the last contact, by date (see "Questions & Answers")

**5** Select Query: Query1

Customer ID	MaxOfDispatch Date
1	4/5/96 5:29:15 PM
2	8/29/95 5:34:34 PM
3	9/11/96 5:11:43 AM
4	12/11/95 6:45:47 PM
5	5/2/94 5:29:49 AM
6	1/2/96 3:39:12 PM
7	10/11/95 5:15:17 AM
8	10/23/96 12:08:23 AM
9	11/26/96 2:08:55 PM
10	10/24/94 2:39:19 PM
11	3/12/96 6:51:46 AM
12	9/12/96 9:37:57 PM



**Fig 6** Maximum date

```
SELECT DISTINCTROW Orders.
[Customer ID], Max(Orders.
[Dispatch Date]) AS [MaxOfDispatch
Date]
FROM Orders
GROUP BY Orders.[Customer ID];
```

**Questions & Answers**

**Q** How can I transfer stock between two locations in a firm chosen from a number of available options? All the names of the locations are stored in one table as different fields.

Kieran Gilmurray  
(who also supplied the answer)

**A** This case statement (right) allows the user to do two things:  
1. Determine the value in a field, in this case a quantity of stock residing at a particular location of their choosing.  
2. Determine the actual name of the field itself to transfer the stock to. (A similar piece of code is used to transfer stock from a particular location.)

**Q** When I create a join in Access, that join appears in the query tools. Sometimes, even though I haven't created a join, one still appears in the query tool. Why?

**A** The rule is quite straightforward. A join will form automatically if:  
• two tables appear in a query which have fields of the same name; or  
• the fields are of the same type; or  
• one of the fields is a primary key.  
If you don't want the join, simply click on the line that indicates the join to highlight it and press the delete key.

**Q** What happens if there are no joins between a pair of tables in a query?

**A** The result is a Cartesian product. Suppose that you have six customers in one table and 40 orders in another. Assuming that each order has a reference to a customer, then you can expect a query that shows customer name and order information to have 40 rows. If you don't have a join showing between the two tables, you will get 240 rows (6 \* 40) because each and every customer will be matched to each and every order.  
As a general rule, Cartesian products

A case of transferring stock (see question, left):

```
*** CASE TO BE DETERMINED WHERE THE STOCK IS TO BE ADDED TO *****
TOO = Me.[TOLOC] ' field on form containing location to which the
stock is going.
Select Case (TOO) ' Test input.
Case "BURGESS HILL" ' NB THIS IS THE SAME AS BURGESS HILL
TOOAMT = STOCK_RS![PHY_QUANTITY] ' amount in a field
TOFIELDNAME = "PHY_QUANTITY" ' name of the actual field in a stock
table
Case "PF_COURIERS"
TOOAMT = STOCK_RS![PF_COURIERS]
TOFIELDNAME = "PF_COURIERS"
And as many other "Case" lines as you need...
Case Else ' Must be something else.
MsgBox "You have entered a location that is not present amongst your
available options", kstosign, "Message"
Me.[TOLOC] = ""
Exit Sub
End Select
' ** use the field values and names in a calculation and update the
appropriate fields in your stock table
STOCK_RS.LOCKEDITS = False ' OPTIMISTIC
STOCK_RS.Edit ' Enable editing.
STOCK_RS(TOFIELDNAME) = STOCK_RS(TOFIELDNAME) + Me![quantity]
STOCK_RS(FROMFIELDNAME) = STOCK_RS(FROMFIELDNAME) - Me![quantity]
STOCK_RS.Update ' Save changes.
```

are not what you want, so make the joins explicit by using the Relationship editor.

**Q** How can I control the order in which the fields are visited when the user presses the tab key?

**A** The default "tab order", as it is called, is the order in which the fields were added to the form. You open the form in design mode, right-click the form and select tab order (Fig 1, opposite). This opens up a box where you can visually set the tab order (Fig 2). This same box can be reached in Access 2.0 by opening the form in design mode and selecting Edit, Tab order from the menu system.

**Q** We are designing a simple Access database to keep track of customer complaints. The overall

design was simple and it's about ready to use apart from a single report that I can't figure out. Each customer record has numerous contact records attached. How do I get a report that contains just the last (by date) contact data for each customer?

Ian Batterham

**A** The easiest way is to base the report on a query that finds the last contact data "should this be date?" for each customer. I happen to have an ORDERS table (which lists customer numbers and order dates) which serves to illustrate the method (Fig 3).

What you need is a "group by" query (Fig 4) which groups the customer numbers and finds the "maximum date": in other words, the most recent (Fig 5). The SQL for this is shown in Fig 6.

sensible methodology of testing databases.

I think that if Microsoft had known the report to be inaccurate, it should have attacked on that basis. But to attack the report on a technicality, rather than on technical grounds, seems to me the worst of all worlds, making the company look (in my personal opinion) ridiculous, high-

handed and arrogant.

As I do not want to receive any "cease and desist" letters myself, I should stress that I am not suggesting Microsoft has behaved in an unreasonable way in this matter. I have every confidence that you can continue to think of Microsoft as one of the family: like, say, a brother... a big brother.

**PCW Contact**

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. Write to him at the usual PCW address (on page 12) or email him at [database@pcw.co.uk](mailto:database@pcw.co.uk).



# Return to client-server

Client-server upgrading, currently a hot topic among readers, takes centre stage. There are Access tips and tricks, plus... how can this be? Mark Whitehorn reviewing his own books?!

Way back in May when the subject of client-server was first raised in this column, I said I'd cover the background and then go on to look at how much work was actually involved in installing a server-based RDBMS and connecting to it from a client. Since then, we've looked at the hardware requirements and the pros and cons of different operating systems and RDBMS products. I came to the conclusion that NT was the operating system most people would choose. As for the RDBMS, I am aware that many people will default to SQL Server. Having had a detailed look at IBM's DB2 for NT, I consider it to be a better product overall, particularly in its most recent incarnation (version 5) which should have hit the streets by the time you read this. In addition, IBM tells me that it will be cheaper than SQL Server, which has to make it attractive.

In fact, I liked DB2 so much, I bought the company. (Not strictly true. If I had enough money to even think of buying IBM, I'd simply retire and race cars. I don't even have shares in it). However, I do like it so much that I have devoted two months of my life to writing a book about it (see page 276). The next section is adapted from that book, mainly reduced in wordage since its purpose isn't to tell you all about installing DB2, but rather to show that server-based RDBMSs are almost ludicrously easy to install.

## Installation

You can install DB2 on an NT server, an NT workstation and even on a Windows 95 machine. I reckon the majority of people will be installing on an NT server so that is the install we will describe, although installation on the other systems is very similar.

The installation of DB2 for NT will come

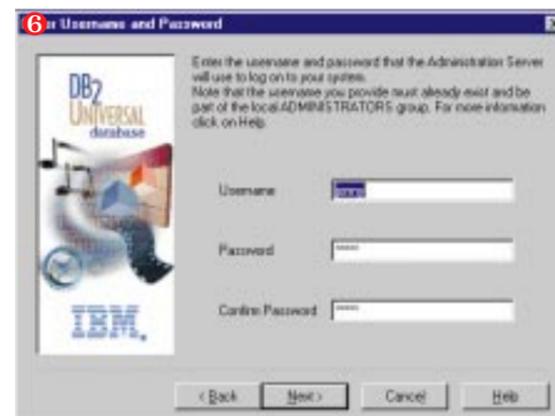
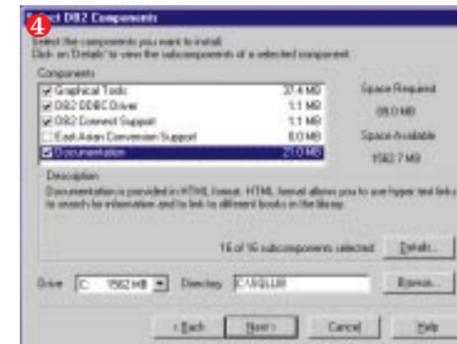
as a revelation to those used to the standard IBM installation routines. It is no longer a long-winded, unwieldy process stuffed full of deeply searching questions and guaranteed to take up days of your precious time. Instead, it is just like that of any other NT program; you are guided by dialog boxes with helpful information and default settings.

So, where do you start? Strangely, you start by ensuring that you have a user account with NT administrator privileges and a user name with eight or fewer characters. Yes, I know this sounds weird, but user names in DB2 have to be eight or fewer characters in order to provide compatibility with DB2 running on other systems.

There are other restrictions besides the one on length: for example, you shouldn't use the characters / or ~, but if you stick to the usual ones (including @ # \$ and \_) you should be

**1 This tells you how wonderful DB2 is and how wonderful you are for using it. Click the Next button. A box says: "Space requirements are being calculated" so select the product you want to install. 2 What you install will depend upon what you've bought from IBM. Choose what sort of installation you want: typical, compact, or custom. 3 With each one comes an idea of the space required and a brief description. Custom is the recommended option as it allows you to see more of what's going on. Select everything that looks appropriate. If in doubt, select it. It can be added later, but you've got to find the CD, shut down DB2, re-run the install etc.**

fine. Since the install process automatically registers the installer as a DB2 administrator, the NT name of that installer has to be restricted in this way. (In case you now begin to wonder, as I did, whether this means that



**4** When asked whether to autostart the Control Center, reply "yes" unless you have reasons to choose otherwise. **5** Choose a folder for DB2; the default suggestion of DB2 for Win NT is a good one. **6** Next, the communications protocols. Accept the defaults at first. Supply a user name and password to enable the Administration Server to log on to the system. This user must have administrator privileges, so enter the details of the user name and password that you're using presently for this installation. Once you're happy, click on the Finish button and the install should proceed smoothly. Finally, reboot the machine. That's the server installation complete.

all NT users who use DB2 have to have short names, the answer is "no"). The easiest thing to do is to create a new NT user with a suitably short name and assign administrator privileges to that user. With that done, you can start the installation.

Login to the NT server as the user you have just created. You are going to have to re-boot the server, so make sure you are the only user on the machine. Close all programs that are running (including those that might have autostarted without you noticing) and put the CD-ROM into the drive. The install program will then start up

and you'll see a Welcome dialog on a big blue background.

## Access Tips & Tricks

### Type conversion

If you convert a number field into a Yes/No type, Access will convert all null values and all non-zero values into No. All other values are converted into Yes.

### Removing primary keys

Assigning a primary key during table design is easy: you just use the key button. Removing a primary key is slightly more devious: you have to call up the Indexes box (using the fifth button from the left on the toolbar), highlight the entry for the primary key and press the delete key on the keyboard.

### Stars in your eyes

Setting the input property for a field to "Password" causes the data to be displayed as a series of stars, rather than as values.

### Text field size

Don't bother setting the maximum number of characters for a text data-type in Access in order to save storage space: it doesn't. You may have been fooled into thinking that it does (I certainly was) because several authoritative books make this claim (including the

excellent *Access 97 Bible* that I reviewed a couple of months ago).

I began to suspect that it wasn't true, so I checked with Bill Marklyn (see book reviews, page 276). He says that Access dynamically allocates storage space to text data-types, so no disk space is saved whether you choose five or 255 characters for the size. All your choice affects is the number of characters that the user is allowed to enter into the field. (However, choosing the correct numeric type (Byte, Integer, Long Integer) does matter in terms of storage space.)

## Become a groupie!

Join the Access User Group. User groups are fun because they are run by enthusiasts rather than marketing droids in suits. The AUG is no exception, and the meeting I attended recently was stuffed full of useful information. Not only did it afford the pleasure of talking to like-minded enthusiasts, it also included talks on, for example, the pleasures (and pains) associated with moving Access applications to SQL Server.

The next major event is a national seminar to be held at Microsoft's UK HQ in Wokingham. Individual membership costs £65 (plus VAT).

For more details contact UK Access User Group, Stokesley House, 53 Prestbury Road, Cheltenham GL2 2BY. Tel 01242 256549. Fax 01242 226021.

## What do you think?

As always, feedback from readers is very welcome, particularly with regard to the client-server products you are using or thinking about using. If the majority of you want to talk about Microsoft's SQL Server, that's fine — but I need to hear your thoughts and needs. You never know, I might write a book about SQL Server, too...

### ■ More on Access currency formats

Two issues ago I said: "If you choose currency as data-type for a field, Access will use whatever Windows has been told is the default currency (via the control centre). However, once a data-type has been assigned to be, say, dollars, Access *won't* change this to pounds if the data file is moved to a Windows machine where the default currency is pounds." I've been playing with this over the last month and have discovered that you can control currency formats more precisely, if you wish.

When you choose a currency format during table design, it can either be declared generically as "Currency", or explicitly as, say,

```
#,##0.00" mk"; -#,##0.00" mk"
```

(a format I get if I tell my NT machine that I



## And on the bookshelf this month...

OK, so it *is* really weird for an author to review his own books, but, having just had two published, one of which is being promoted by *Personal Computer World*, it seems to me to be overly self-effacing to ignore them completely.

### Inside the Relational Model – with examples in Access

**Authors** Mark Whitehorn & Bill Marklyn

**Publisher** Springer-Verlag

**ISBN** 3-540-76092-X

**Price** £19.50; soft cover, accompanied by a CD-ROM of examples and samples

As the title suggests, this one is about the relational model. I believe that, in order to build good relational databases, you need to understand the relational database model. Many of the textbooks that I have on the topic are authoritative but, sadly, lacking in realistic examples and (dare I say it) a sense of humour.

The subjects covered in the book have been carried in the *Hands On Databases* column over the past four years and this book draws that information together in an integrated way. There is also much new information, as well as a CD-ROM with all the sample files.

The audience at which it is aimed includes those who are building a database and who want to know more about the process in order to reduce the potential for encountering pitfalls later. The section on SQL is also a useful primer for anyone new to the language.

### DB2 for Windows NT – Fast

**Authors** Mark Whitehorn & Mary Whitehorn

**Publisher** Springer-Verlag

**ISBN** 3-540-76200-0

**Price** £19.50; soft cover, accompanied by a CD-ROM

The focus of the column has, by popular demand, been turning towards the topic of client-server databases, and the second book is a "get-you-up-and-running-quickly" guide to IBM's latest client-server product, DB2 for Windows NT. I have always been impressed by IBM's database engine (a robust, elegant and speedy implementation, just like a database engine should be) though dismayed at the tenacity with which the company clings to its command-line interface. This latest version of DB2 has a GUI; yes, a proper, user-friendly GUI with all mod cons.

am Finland). The former format (if you see what I mean) will convert to the local

currency as the table is moved between machines, and the latter won't. During table design, you can set the format manually to either of these as you see fit.



## PCW Contact

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. Write to him at the usual PCW address or email [database@pcw.co.uk](mailto:database@pcw.co.uk).



# Primary purpose

Mark Whitehorn attends to the use of primary keys in data tables. You can also read about graded data and how to get into the currency control business with Access.

Last month I wrote: "If you choose currency as data-type for a field, Access will use whatever Windows has been told is the default currency (via the control centre). However, once a data-type has been assigned to be, say, dollars, Access won't change this to pounds if the data file is moved to a Windows machine where the default currency is pounds." I've been playing with this during the past month and have discovered that you can control currency formats more precisely if you wish.

When you choose a currency format during table design, it can either be declared generically as "Currency" (Fig 1) or explicitly as, say, "#,##0.00" mk";-#,##0.00" mk" (Fig 2) — a format I get if I tell my NT machine that I am Finland. The former format (if you see what I mean) will convert to the local currency as the table is moved between machines, and the latter won't. During table design, you can set the format manually to either of these as you see fit.

### Grades

Some data inherently cries out to be graded. OK, I'm thinking about student marks here: 70 or above = A, 65 - 69 = B and so on. But you might be a travel agent thinking: Cost £500 or above = Luxury, Cost £400 - £499 = Business Class... down to, Cost less than £50 = Economy.

Whatever you need to grade, always store the value to be graded (mark, cost and so on), never the grade. You can use a query that looks at the original data and calculates the grade you need. I have included an MDB file called UOD on our cover-mounted CD-ROM this month. This stores students' examination marks in a

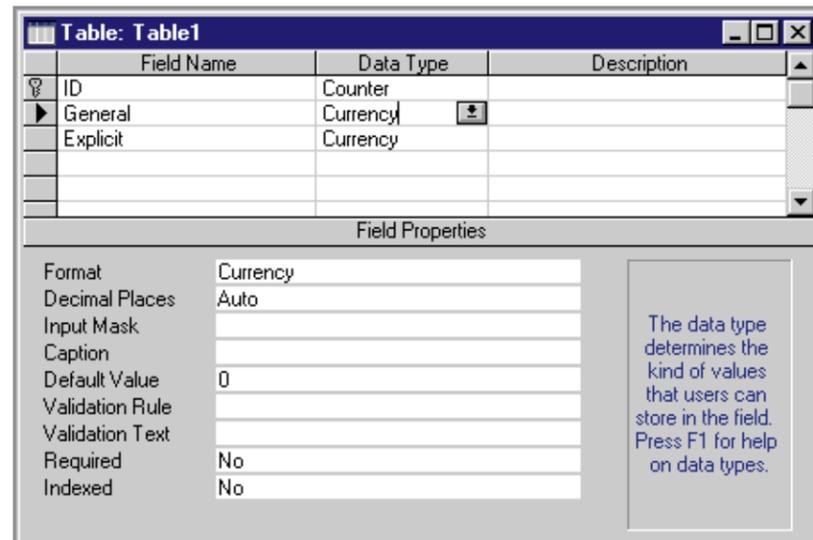


Fig 1 (above) Declaring a currency format generically during table design... and declaring it explicitly (Fig 2, left)

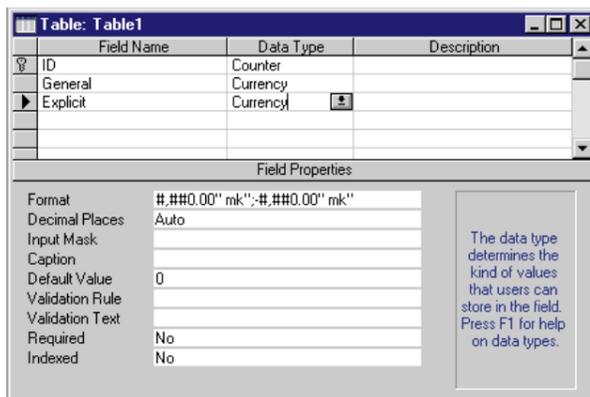


table called ATTEND (that also stores which student attends which course). The query called Grades uses a formula:

```
Grade: IIf ([Exam]>=70, "A", IIf ([Exam]>=65, "B", IIf ([Exam]>=60, "C", IIf ([Exam]>=50, "D", IIf ([Exam]>=40, "E", "F")))))
```

to assign a grade (Fig 3). The formula may

```
Grade: IIf ([Exam]>=70, "A", "B")
```

which says that if the value in Exam is greater than or equal to 70, assign an "A", otherwise assign a "B".

### Database cure

"I am writing an occupational health database. Each nominal" [I wonder what a

FIRSTNAME	LASTNAME	COURSENAME	Exam	Grade
MIKE	WELLINGTON	Cytogenetics	76	A
MIKE	WELLINGTON	Intro. to Polymorphism	67	B
MIKE	WELLINGTON	Automotive Design	72	A
SALLY	JONES	Simple Ergonomics	56	D
MIKE	WELLINGTON	Ecology	23	F
MIKE	WELLINGTON	Population Genetics	45	E
*				

Fig 3 (above) Assigning a grade

Fig 4: tblMinPrevOcc

FldRefNo	PrevOccCode	LenPrevOcc
10003	1234	4
10230	4567	2
10003	7896	5
10458	1256	2
....	....	..
10230	5693	1

nominal is? — MWJ "may have more than one previous occupation, stored in the table 'tblMinPrevOcc' (Fig 4).

"I want to run a query which will tell me how many nominals have had a previous occupation: i.e. how many unique values are there in the fldRefNo field? I thought I could do this with the following statement: SELECT COUNT(DISTINCT fldRefNo) AS fldPrevOccCnt FROM tblMinPrevOcc; but I get a syntax error.

"I feel that the design of the database may be at fault but am not sure what is the best solution. I am using Access 2.0. Any suggestions?"

Mark Capaldi

The design is certainly worth examining, since there is no obvious primary key for the table tblMinPrevOcc. You could use all three fields, but that is getting slightly messy and you would then exclude the possibility that the same person had, on more than one occasion, done the same job for the same length of time, which is unlikely, I know, but still possible. A better alternative might be to add a Counter field (but see "Candidate keys", later).

However, I digress. Here is a solution to

your problem. Create a query called (for the sake of argument) "Unique" which is as follows:

```
SELECT DISTINCT tblMinPrevOcc.FldRefNo AS fldPrevOccCnt FROM tblMinPrevOcc;
```

which is very similar to your original. This shows (but does not count) the unique entries in the table:

FldPrevOccCnt
10003
10230
10458

You can then base a query such as SELECT DISTINCTROW Count(Unique.fldPrevOccCnt) AS [Total No Of Nominals] FROM Unique; on this query called UNIQUE. This second query will return the number you want:

Total No Of Nominals
3

A sample is on our cover-mounted disc this month as NOMINALS.MDB.

### Candidate keys

Mark Capaldi's question (see "Database cure") obliquely raises another: given that all tables should have a primary key, how do you choose the field, or combination of fields, to use as the primary key?

The answer, delightfully, is often not as simple as it seems. "What do you mean, it isn't delightful?" I hear you cry. If all the decisions we had to make during database design were easy, where would be the fun? And, more to the point, how could we justify our huge salaries?

Fig 5: Customers; but which candidate key?

CustomerName	Address	ManagingDirector	TelephoneNo
Fred Bloggs Ltd.	23 Park Rd, London	Brian Bloggs	098 765 432101234
Helen Shipping	46 Troy Lane, Paris	Helen Ofroy	01234 567 8901233

Most people (he wrote hopefully) know that a primary key is a field or combination of fields, the values which uniquely identify the records in the table.

Let's assume that you work for a company called PenguinShipping. This company manufactures items which it then sells to customers. As the database designer, you decide that one table for CUSTOMERS and another for ORDERS is appropriate.

ORDERS		
OrderNo.	Value	ShippingDate
1	£50	1/9/97
2	£75	2/11/97

Assuming that each entry in OrderNo is different, that field is an obvious candidate to be the primary key because the value therein uniquely identifies the record to which it belongs. I use the word "candidate" advisedly in this case since this is actually the technical term used to describe any field, or combination of fields, which could be a primary key.

Clearly there are other candidate keys in this table. We know that a primary key can be made up from multiple fields, so given that OrderNo already uniquely identifies the records, OrderNo + Customer must also be a candidate key. By the same token, OrderNo + ShippingDate is also a candidate key, but Customer + ShippingDate is not a candidate (unless we happen to know that two or more orders are never shipped to the same customer on the same day).

However, it is often impossible to identify candidate keys simply by studying the data. Additionally, you often need to have an understanding of what the data means, how it is used, and what are the company rules for that data.

For example, suppose I tell you that PenguinShipping actually makes ships and that the Value field is in £m. ShippingDate is actually the date on which any given ship is launched. If I also tell you that the port authorities forbid more than one launch per month, you can suddenly identify ShippingDate itself as a candidate key, since there can never be more than one record with the same ShippingDate.

Clearly, once you have identified the candidate keys in a table you have to pick one of them to be the primary key. As it stands, OrderNo is almost certainly the most reasonable candidate in the table above, but consider the one shown in Fig 5.

Suppose that all customer names are guaranteed to be unique, as are all

Fig 6: Orders

OrderNo.	CustomerName	Value	ShippingDate
1	Fred Bloggs Ltd.	£50	1/9/97
2	Helen Shipping	£75	2/11/97

Fig 7: Adding a numeric field

CUSTOMERS			
CustomerName	Address	ManagingDirector	TelephoneNo
Sam's Sweets	16 Green La., London	Sam Sugdon	0123432101233454
Bon Bons	12 Apple Lane, Cheapstow	Billy Brown	0143210987654342

ORDERS			
OrderNo.	CustomerName	Value	ShippingDate
1	Sam's Sweets	£24	1/9/97
2	Bon Bons	£45	1/9/97

addresses and telephone numbers. Here we have a plethora of candidate keys. Not only can we use any one of those three, we can use any combination of any keys except ManagingDirector on its own. Assuming that common sense prevails, the most obvious choice looks like CustomerName.

This means that the ORDERS table can now use a foreign key to reference the CustomerName field as in Fig 6, and all is well. But, suppose you now build exactly the same type of database, only this time you work for PenguinConfections. This company doesn't ship ships, it sells sweets. In this case, electing to use the CustomerName field is less likely to be satisfactory. Which raises the question, why should the product make any difference to the database design?

The answer lies in the number of records in the ORDERS table. PenguinShips will be lucky to complete ten ships a year. This means that the ORDERS table will remain tiny. PenguinConfections, we hope, will ship many more orders than this. Assuming that it ships, say, 10,000 orders per year, then the ORDERS table will hold 10,000 records. Imagine all of those customers' names written dozens of times each — wasting space, cluttering up the disk and slowing down the database. The obvious answer is to add a numeric field to CUSTOMERS and use that as the primary key (see Fig 7).

This is a great solution (and one that many people would employ almost instinctively) even though it renders the ORDERS table less readable, since we now have a number rather than a name. However, it is worth remembering that this

solution brings its own overheads because you have added a field to CUSTOMERS.

Under certain circumstances this might be a mistake. Suppose that your CUSTOMER table is mega-huge (because it stores all your potential customers) and even though your ORDERS table has 10,000 entries it only actually references a small proportion of the records in the CUSTOMER table. Adding a CustomerNo field to CUSTOMERS will reduce the size of the ORDERS table but the concomitant increase in the size of CUSTOMER may negate the gain.

So it is yet another case of horses for courses. In practice, I tend to do the following: if a table has an obvious, relatively short, single field candidate key then I use that as the primary key. This is not uncommon (Part Number, Student matriculation number, National Insurance number and so on). If such a field is not obvious, by default I tend to add an artificial numbering field such as Counter/Autnumber in Access.

#### Client-Server

I haven't forgotten that this is a continuing topic, but there is not enough space this month. I will return to it next time, bringing laughter and song.

#### PCW Contacts

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. Write to him at the usual PCW address or email [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk).

UK Access User Group, Stokesley House, 53 Prestbury Road, Cheltenham GL2 2BY. Phone 01242 256549; fax 01242 226021



# And that's magic...

Mark Whitehorn conjures up tricks to use in Access. He finds it is a product that knows where its towel is, without any flannel, and has meta-data which needs to be stored.

**Y**ou can use Shift F2 virtually anywhere in Access. It opens a "Zoom box" where you can edit the text upon which your edit cursor was sitting. Use it whenever the text you are trying to read sits in a small box. It is mega-useful when entering a complex condition in a query, say. It works for users of your application who can be encouraged to use Shift F2 when entering data into a form.

You can also use Shift F2 while programming. If you place the edit cursor on a user-defined function, pressing the magic key combination will jump you to the definition of that function. Doing the same on an Access function like CurrentDB (i.e. not user-defined) will open up the object browser at the appropriate place.

## Committing records

Another useful tip to feed to your users is that Shift Enter will commit an edited record to disk without the need to move onto the next record. (Or click on the edit record symbol on the left of the form.)

## Access currency formats

If you choose currency as data-type for a field, Access will use whatever Windows has been told is the default currency (via the control centre). However, once a data-type has been assigned to be dollars, say, Access won't change this to pounds if the data file is moved to a Windows machine where the default currency is pounds. This makes perfect sense (since the absolute value of items would alter as data files were moved from country to country) but needs to be remembered if you are distributing an application into other countries.

## A free toy for data mining in the data dictionary

Access is a product that really knows where its towel is. Well, in the literal sense it has no idea of the physical location of its textile dehumidifying implement; in fact, it probably hasn't got one. What it does have is the database equivalent, which is called meta-data, and it certainly knows where that is — in the data dictionary. Your problem is trying to read that data and I have a solution.

A database consists of the user's data stored in tables. A DBMS also has to store information which "describes" the database itself (the names of the tables, the names of the queries, the explicit joins between the tables, the data integrity constraints etc). Such information is known as meta-data. One of Codd's rules that define a Relational DBMS

says that the meta-data can't be stored any old way; it must be stored as data in tables, just like the user's data. Access obeys this rule and stores the meta-data in a series of hidden tables, the names of which all begin with MSys (presumably standing for Microsoft System). The data in the tables is the meta-data, the tables themselves can be referred to as the data dictionary.

Codd says that storing the meta-data in the same way as the user data ensures that the users of the database can access the meta-data in the same way in which they access their normal data. So any user who has learnt to query their data to find out, say, how many customers live in Hereford, can use the same techniques to find out how many forms there are in the database.

All of this is fine and dandy in theory but in practice it rarely works, for a very simple reason: the database designers are free to use whatever data format they like within the tables to represent the structure of the database, and they seem to choose very abstract formats. In Access, for example, data about the queries is stored in a table called MSysQueries (Fig 1).

The first 15 rows in the table shown, define one query in the database. In order to find out what sort of query it is, you have to find the value 1 in the Attribute field within the first 15 rows and then look up the value in the field called Flag.

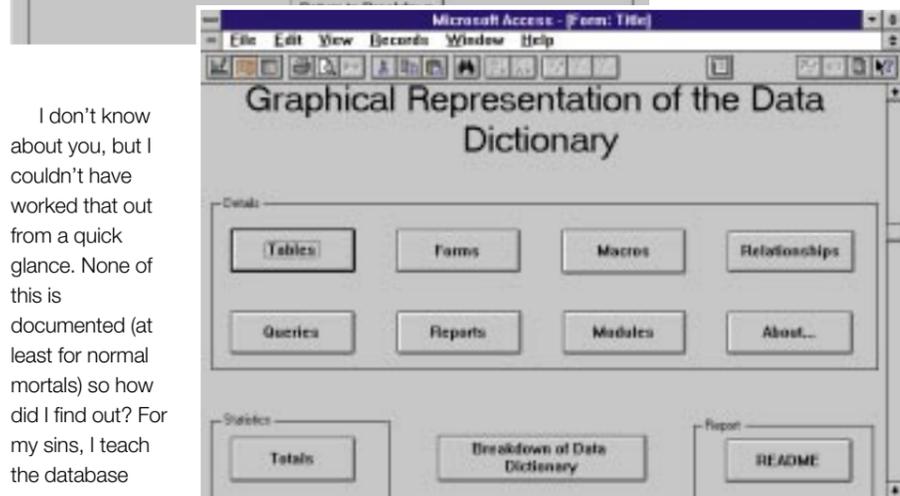
Flag	value	Meaning
1		Select
2		Make
3		Append
4		Update
5		Delete
6		Crosstab
9		Union



Fig 1 Inside the data dictionary. Intuitive, isn't it?



**Fig 2 (left)** Some of the useful information inside Black.mdb



**Fig 3 (below)** The GUI front-end to Black.mdb

I don't know about you, but I couldn't have worked that out from a quick glance. None of this is documented (at least for normal mortals) so how did I find out? For my sins, I teach the database course at Dundee

University. One of the projects I set the students this year was to build an Access application that displayed the meta-data in a user-friendly fashion. In order to do this, the students had to "reverse engineer" the data they observed in the data dictionary, working out what all the bizarre formats meant before generating queries that made sense of them.

**Student demonstration**

It seemed a shame for this work to be wasted, so I am including one of the better ones on our cover-mounted CD-ROM (in Black.mdb). This was written by one Johnny Black and includes some notes he made about the formats used within the data dictionary.

Three riders: first, remember this was a student project, so neither Johnny nor I can guarantee that all his conclusions are correct. Second, the structure of the data dictionary varies between different versions of Access. Black.mdb was written specifically for version 2.0 and may well not provide accurate information if imported into later versions. Third, this work was completed to a tight time schedule and had I given the students more time could

obviously have been improved. With those riders, Black.mdb should still serve as an excellent starting point for anyone who wants to examine the data dictionary.

**Operating theatre**

Having covered the hardware aspect of client-server computing, we can move on to a more contentious area: the operating system and the back-end DBMS.

Assuming for a moment that you do not want to stray into the realms of mainframes just yet, the obvious operating systems onto which a PC database application could be up-sized are UNIX, Novell NetWare and Windows NT.

UNIX is possible, but my experience is that most PC users tend to shy away from it, which is a shame because it's a very stable platform for RDBMSs.

Next comes NetWare, which some people may consider, especially if they are using it as their NOS for file and print. The problem is that NetWare was optimised for precisely those functions, not as an application server. So although products like Oracle 7 are available for NetWare, in practice they sell like cold cakes.

I love NetWare dearly as a NOS but

cannot recommend it as an operating system for RDBMSs.

Finally there is Windows NT. This is not as stable as UNIX, and we could debate for some time whether it is really a stable enough platform for a mission-critical RDBMS. It is apparent however that many people are migrating to it and that the major RDBMS firms are providing products for it. If you choose to use Windows NT, at least you will have lots of company, so that's my general recommendation unless you have good reason to choose otherwise.

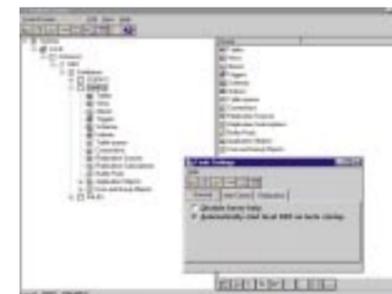
Each of the big three has an offering for NT: IBM has DB2 for NT, Oracle has Oracle for NT, and Microsoft has SQL Server. Two of these products have a history of running on other platforms, and all three have perceived strengths and weaknesses.

DB2 is famous for running on mainframes, being very big, powerful and secure, and being a pig to drive. Oracle is famous for running on UNIX, being very big, powerful and secure, and being a pig to drive. SQL Server is famous for running on NT and having a great user interface.

If you are currently an Access user, the obvious choice is surely SQL Server from Microsoft. After all, you are already happy with Access: SQL Server comes from the same company so you can expect the same quality of software. You can also expect a high level of compatibility between the front-end and the back, so is there even a question here?

The answer is "Yes", there is a serious question, but a relatively complex one with several facets. Access and SQL Server were both developed within Microsoft but by different development teams. My experience is that the teams can produce differing standards of products, so the fact that both of these come from Microsoft is largely irrelevant. Indeed, SQL Server was originally a Sybase product and still has large chunks of Sybase code buried within it which has no connection with Microsoft at all. This code includes the query-optimiser which is "less-than-optimal".

In my June issue column I suggested that many client-server databases are also mission-critical to a greater or lesser extent. There is a huge gulf between a word processor and a database which stores information crucial to the survival of a



**Fig 4** Using DB2 for NT is actually a pleasant experience, at least in the new version. Clockwise (from top left): Choosing fields for a table; setting the primary key; setting data types; the main control area for DB2

company. It generally doesn't matter much if the former falls over occasionally: the latter must be much more stable and, if it fails, must fail gracefully so it can continue to work when restarted.

Companies like Oracle and IBM have always worked in mission-critical environments and understand that, in this environment, the stability and data-integrity of the product is of paramount importance. They understand that software mustn't be shipped until it is really rock-solid.

Now ask yourself what you think about Microsoft's track record in producing software that never falls over? Ask yourself how many times you have heard about bugs turning up in new releases of Microsoft products which will "be fixed in the next version". I'm not suggesting that Microsoft can't learn about mission-critical software. What I am suggesting is that I want to see a good track record before I commit my data to a given product. I believe IBM and Oracle have already demonstrated such a track record.

"Ah, but what about drive-ability?" you ask. "You're really suggesting that we use a product with a Neolithic user-interface." (Those of you who haven't experienced the joy of playing with the current version of DB2 for NT may be surprised to discover that it still has a command-line interface. No, I'm not joking, or exaggerating. Yes, I do mean you don't get cute little buttons, icons and slider

bars. You have to type a string of characters into a textbox in order to drive it.)

But no, I wouldn't inflict a command-line on anyone. The good news is that I have been playing with the new version of DB2 Universal Database for NT (due for release in October) which has a really intuitive, easy-to-use interface. In fact, it closely resembles Access (Fig 4).

Of course, if we are allowed to talk about next versions, Microsoft is likely to say the next version of SQL Server will fix any current problems that exist within the product. Apart from considerations about track records in this area, a good question to ask here is, "Which company has the easiest job?" IBM, which has to bolt on a respectable interface to an already mission-critical piece of software, or Microsoft, which has to convince us that SQL Server is of mission-critical quality.

**Scaling up**

Scalability is loosely defined as: "How much the performance of a given product alters as conditions are changed."

A good example is the number of users attached to an RDBMS. Suppose you have an RDBMS running with 20 concurrent users. If it were perfectly "scalable", doubling the number of users would no more than double the response time for each given user. But life isn't usually like that, and doubling the number of users often makes the response time more than

double for each user (because the RDBMS spends proportionately more time swapping between the users). So no-one expects perfect scalability but it is possible to compare products and see which scale most efficiently. Bloor Research has recently published a report called *The Realities of Scalability* which examines the scalability of DB2 for NT and SQL Server (it also looks at the scalability of DB2 for AIX).

The scalability of SQL Server comes in for criticism: "Put simply, Microsoft SQL Server for Windows NT at high numbers of users performs dramatically worse than either of the other two databases." As for the stability: "This database had a number of failure states that could repeatedly [sic] be generated. With large numbers of users, it was found that the database would grind to a halt."

DB2 for NT, as you might expect given its history, does not come in for such criticism: "Despite the occasional unexplained server or database crash, it proved impossible to consistently generate a fatal error in DB2 for Windows NT. It seemed to be capable of taking most things thrown at it."

**Perspective**

All of this has to be seen in perspective. Suppose you work for a small to medium-sized enterprise. You have developed an Access database on a standalone PC and it has been riotously successful. Now your boss wants you to upgrade it so that ten people can use it. There is no suggestion, either from the Bloor report, or from me, that SQL Server would be inappropriate. However, I can see no reason why the new DB2 wouldn't be better and it is actually cheaper, so it is well worth considering.

Now suppose you are developing an application that may become mission-critical in the future, or which needs to support many users (say >1,000) and/or lots of data (say >100Gb). Under these circumstances, for the reasons given above, I would strongly recommend that you look carefully at the alternatives to SQL Server rather than view it as the default option.

**PCW Contacts**  
 Mark Whitehorn welcomes readers' correspondence and ideas for the Database column. Write to him at the usual PCW address or email [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk).  
 Bloor Research 01908 373311



# Classy chassis

Mark Whitehorn turns his attention to “professional” quality hardware for client server databases and what you can expect for your money. Plus, tips and tricks for Access.

**L**ast month I was trying to “sell” you a very expensive piece of kit. I hoped to convince you that standard PC chassis were inappropriate as servers for mission-critical databases and that you ought to buy something expensive from a reputable manufacturer.

A fair question here is, “What do I get for my money?” In other words, a “professional” server is likely to cost twice as much as a PC of the same specification, so where does the money go? The following list is not exhaustive but it gives some idea of the extra hardware features you can expect from a “professional” server:

- RAID disk arrays.
- Fault monitoring and prediction on hard drives, processors and memory (anyone who remembers HAL at work in 2001 — A Space Odyssey will appreciate this).
- Pre-failure warranty, which allows the components identified as liable to failure, to be replaced before they do so.
- Dual peer PCI buses for high I/O bandwidth.
- Array controllers which allow RAID storage to be added while the system is in operation.
- Tape backup systems which provide rapid backup (40Gb per hour).
- Redundant power supply and UPS options.
- Redundant controllers on standby in case of failure.

Most of these are self explanatory. However, a brief word about RAID (Redundant Arrays of Inexpensive Disks) may be helpful. “Brief” is the operative word, as there are several flavours and the area is relatively complex. A RAID allows you to store your data on several disks, which holds a number of advantages as far

as database servers are concerned. For a start, a RAID can be set up in such a way that if even one of the disks fails, your data remains safe because each piece is stored redundantly on more than one disk.

Clearly, this has the disadvantage that although you might have five 2Gb disks and therefore 10Gb of disk space, you won’t be able to store 10Gb of data. The data is more secure, though. Additionally, there is a rather elegant spin-off from the redundant nature of RAID which is that the data can often be retrieved more rapidly.

Furthermore, RAID can be hot-swappable, so if a disk fails you can extract the duff one, slot in a good one and the data which was on the failed disk will be reconstructed from the others. The database can continue to perform, albeit more slowly, until the new disk is fully online.

It doesn’t take a genius to work out that the money is going into two main areas: speed, and keeping the server up and running. You may never need the RAID array; that redundant controller may spend all its time idling. Nevertheless, it would make your day having arrived one morning to discover that even though one of the disks has failed, your system is still running.

## Remember it well

A database server needs RAM, and then it needs more RAM. A database server will use RAM to hold the RDBMS engine and to cache data, transactions and indices. The more RAM you give it (within reason) the faster it will work. Think in terms of 64Mb as a minimum and 100Mb as a more reasonable starting point. Make sure the server you buy has free memory slots so the

## Time, gentlemen, please...

Following reader Gareth Wade’s query about adding time lengths to give a running total, (*Quickies*, PCW April) I have since received a couple of solutions from others.

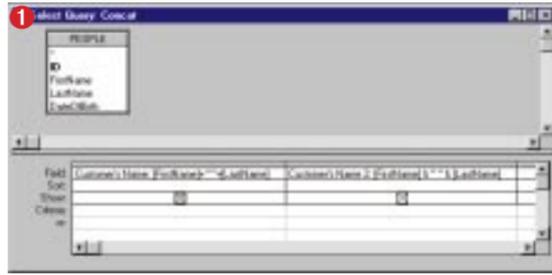
■ MA Roberts writes: “Regarding the ‘Database Quickie’ from Gareth Wade (PCW April) I have had a similar problem to Gareth’s when adding together time values. The short time format field in Access does not work satisfactorily for adding values that will exceed 23:59. One workaround I am using is to convert the time field into a number type. When adding time records together in a number-type field, Access converts the time to a fraction of the 24-hour day. It is therefore necessary to multiply this field by 24 to get a value for hours and decimal minutes. This field may then be used to obtain a value for time (in hours and decimal minutes) multiplied by an hourly rate such as might be used in a time-sheet style application.”

■ And Vidar Eggen writes from Oslo: “I had the same problem when wanting to report on the added flying time of pilots as well as aircraft at my flying club. What we really need in Access is the [tt:mm] format of Excel. Until then, I use this basic formula:

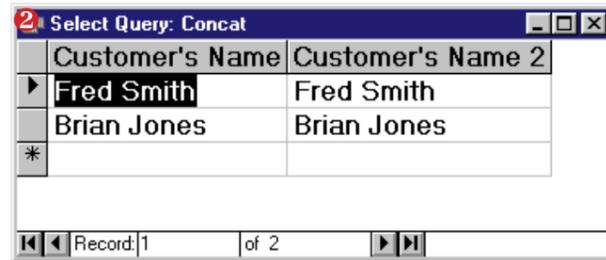
```
=Str$(Sum(Hour([Time]))+Sum(Minute([Time]))\60) & ":" &
Format$(Str$(Sum(Minute([Time])) Mod 60);"00")
```

“The output of this formula will be ‘123:45’ in text format. The reason for converting the sums to a string is to make sure that five minutes is printed as ‘05’, not ‘5’. If quoting from this, please rephrase my language into readable English!”

(Vidar’s English seems fine to me and is considerably better than my Norwegian! — MW)

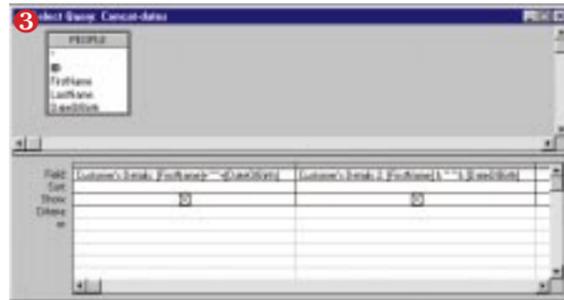


machine and turned it loose on a large database. Rather to my surprise Access used all the four processors on the machine and was able to query 1,000,000 indexed records in a couple of seconds. Interestingly,



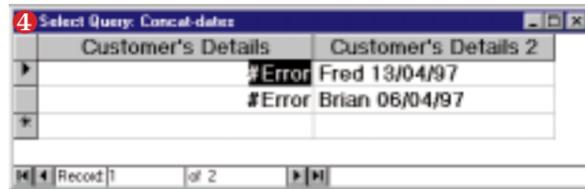
Excel used only one of the available processors and its performance was startlingly unimpressive. So the bottom line is, what do I use as my test database server? I am

memory can be expanded later without replacing all the RAM it currently holds.



**Purchasing processors**

In general, buy several processors. Simple operating systems can only use a single processor, but complex ones like UNIX, VMS and NT4 Server can use multiple processors. Mind you, just because an operating system can make use of several CPUs doesn't mean that the programs which run on top of it can too. But all the respectable RDBMSs can and do, and this isn't restricted to client-server RDBMSs.



currently running a Compaq ProLiant 5000 with four 166MHz Pentiums, 384Mb of RAM and 9Gb of disk space. I look upon it not so much as a server, more as a non-intrusive sleeping aid. Okay, it is more expensive than the tablets, but it has fewer side effects.

**Listing 1**

```
"Report generated " & Now() & " for " & [Customer]
```

**Listing 2**

```
Between Date()-7  
And Date()
```

**Listing 3**

```
Between #1/1/1981# And Date()
```

**Tips and tricks for Access**

By popular demand, I am introducing tips and tricks this month. It is aimed at Access users and will include information which ranges from simple to complex and should provide something for everyone.

**1. Queries**, by default, use the field name as the column name in the answer table. If you ever want to alter this, insert the text you would prefer to use, followed by a colon, before the field name. Thus:

```
Customer's Name: [FirstName]
```

will replace "FirstName" with "Customer's Name" in the answer table. Which leads us neatly to the second tip.

**2. The "+" operator** can be used to concatenate text strings, as many people know. Thus:

```
[FirstName]+[LastName]
```

in a query will join the two strings together. But this unhelpfully yields names like "BrianJones". So, simply add in a literal space like this:

```
[FirstName] + " " + [LastName]
```

which then gives: "Brian Jones".

However, you can also consider using the lesser known "&" operator. Amazingly,

```
[FirstName] & " " & [LastName]
```

gives exactly the same result: "Brian Jones" (Figs 1 & 2). You're thinking, "Hello, he's flipped. If it gives the same result, why bother to use it?" Well, the advantage of getting into the habit of using "&" is that it automatically converts all data types into text. This has no effect if the two fields are already text (as in the case above) but it does wonders with expressions like

```
[FirstName] & " " & [DateOfBirth]
```

which will produce "Brian 06/04/1967" rather than the "#Error" produced by the + operator when it tries to concatenate data from two disparate data types (Figs 3 & 4). In other words, if you habitually use "&" instead of "+" you won't have to use type conversion functions. You could even use it in a report like Listing 1.

**3. When adding fields to a query grid**, you can drag and drop them as usual. However, you can also select multiple contiguous fields using shift-click on the first and shift-click on the last, and then dragging and dropping as usual. You can also add non-contiguous fields by CTRL-clicking on them.

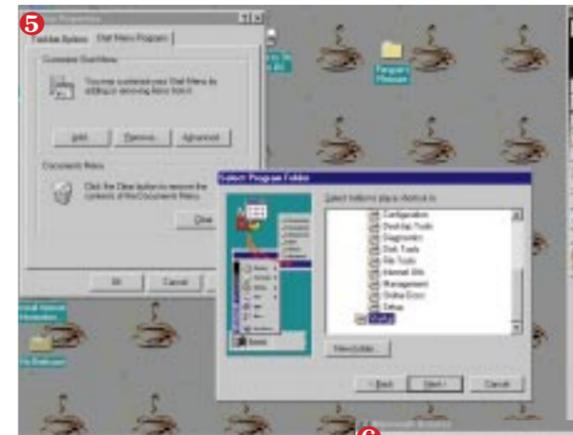
**4. When looking at a table of data**, you can resize any field to fit the widest data in the field by double-clicking on the right-hand column border (this has to be done at the top of the table). You can hide the column by dragging the right-hand column border across to the left-hand border. You can open this up again with the mouse but it takes practice, as you often simply widen the adjacent column instead. It is often easier to use the menu system Format

**Listing 4**

```
Set button OnClick event to "=Clicked([Screen].[ActiveControl].[Caption])"
```

**Listing 5**

```
Function Clicked(capStr as string)  
    Debug.Print capStr  
'Or assign it to the parameter for the query  
End Function
```



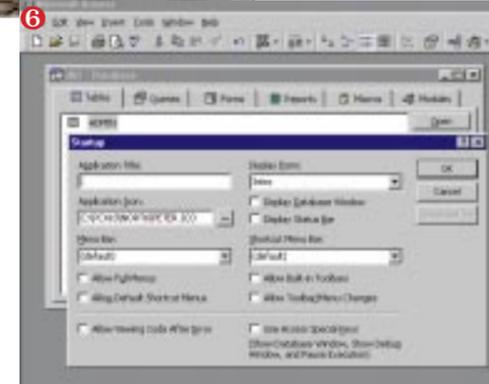
delay would be seven days, so the default could remain as Date() and the validation rule could be set as Listing 2. The "Between... And" operator is a useful one to get to know. You can use it with fixed dates like Listing 3 or for data types other than dates: Between 1 And 45

Show Columns, or Format Unhide Columns, in Access 97.

**5. Record selector.** Most of us have discovered the delights of the "Find Specified Text" button (the binoculars) which will search a table for information. By default, this will search the current field and the dialog box allows you to specify the entire record. However, it is often easier to click the record selector (at the left of the record) before starting the search and this automatically sets the search to run against all fields. This works in form views as well as in datasheet views.

**6. "Between... And".** As regular readers will know, I am obsessed with the idea that the data in our databases should be as "clean" as possible. Take something as simple as a DateOfBirth field in a maternity database. The default (for the child, not the mother) might be the function Date().

But perhaps not every child's record might be entered on the day (staffing levels being what they are) so the simple answer is to allow any entry. But that opens the door to really whacky entries. Suppose we know that all entries, no matter how delayed, are always brought up to date on a Sunday, say. This means that the longest



**Gaining automatic Access**

A reader has written in for help: "I'm struggling with Access 97 for NT, trying to build a database for other people to use. I have seen other databases where a form opens up automatically, which is what I want but how do I do it?"

"Additionally, since many of my users are not used to databases, is there any way of limiting the options which are available to them when Access is running?"

"Finally, is there a way of getting Access to start as soon as NT fires up? Can I set up the machine so that whenever the user switches on the PC, it goes into NT, lets them log in, and then goes into Access?"

The simple answer is yes. Let's start from the outside and work inwards. To fire

up a copy of Access, you can add it to NT's StartUp menu. Click on Start and choose Settings, TaskBar and the Start Menu Programs tab. Click Add and then browse in the normal way to find Access. Then click on Next, select the StartUp folder and finish. This seems more complex than it used to be in Windows 3.1 — weren't operating systems supposed to be getting easier to use? (Fig 5)

Access 97 has a new startup form which allows you to control many options when it kicks into life. You can get to this by right-clicking the border of the Database window and choosing StartUp, or go via the menu with Tools, StartUp. Fig 6 shows the sort of options that are available.

The subject of passing captions from buttons has exercised the brains of several readers. For instance: "I was intrigued by the buttons and captions capture question. Your method in the February issue appears to be the way to go and can be speeded up using [Listing 4] (using Access 7) which calls [Listing 5]." Kelvin

Malcolm Fraser, too, sent in a solution and it's on our cover-mounted CD, this month, in a file called CAPTIONS.MDB. Grateful thanks are also due to Steve Foster, Richard Todd, Malcolm Bacchus, Wilf Davies, Alan Berry and others who provided elegant solutions.

**Beyond the pale™**

I discovered the following rumour on the web site of Bloor Research, at [www.bloor.co.uk](http://www.bloor.co.uk). It appears that the word Metadata™ has been trademarked™ to a pharmaceutical™ company in the US. At first sight this sounds (if you'll forgive the mixed metaphor) unbelievable, but I suppose those people who issue trademarks in the US may not be as familiar with planet database as those of us who live there. To make matters worse, the company that was awarded the trademark™ has started sending "Cease and Desist" letters to anybody it finds using, er, well, that word™. I'm currently trying to trademark the word "pharmaceutical™" in the hopes of wreaking some small revenge™.

**PCW Contact**

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column, at [database@pcw.co.uk](mailto:database@pcw.co.uk).



# Serves you right

Well, he promised, and here it is. As a new, regular section of his column, Mark Whitehorn introduces the subject of client-server computing explaining how, why and wherefore.

**A**s promised last month, a section of this column will now be devoted to the topic of client-server computing.

There are several definitions used in client-server databases. For our purposes I will use the definitions shown below unless stated otherwise. What we can discuss over the next few months is:

1. **Server** Suitable hardware, OS and RDBMS.
2. **Client** Suitable hardware and software.
3. **Component positioning** Where you should place the business rules and the data (yes, I know I said it sits on the server, but there are some exceptions!).

We can also look at how an existing single-user system can be upgraded to client-server: that little lot should keep us busy for a few months.

## Background

Most companies upsize to a client server because they need to change a single-user database into a multi-user. Although this can be done by moving the data onto a file server (see my previous columns), such

solutions are usually limited in both power and the number of users. For many companies, the only really effective way to provide multi-user access to a database is by moving to client-server.

In many cases, the change to client server implies another subtle change: the database changes from a useful but non-essential part of the business to a mission-critical system. I am not suggesting that all multi-user databases are mission critical, simply that in my experience many become so. The following is a useful conversation to have with your boss before you start.

*"Suppose that we go ahead and build this client-server database, and suppose that it is as successful as we all hope. Now imagine: once it has been in place for six months or so, it begins to fail sporadically. How dangerous will that be to our business?"*

If the answer is: *"It will be annoying, but we can live with it because (insert appropriate answer here)..."*, you don't need to read the rest of this section. If the answer ranges from: *"Well, that would be difficult to quantify..."* to *"Such failures, if prolonged, would seriously damage the*

*company"*, you need to think very carefully about hardware for the server and be prepared to spend some serious money.

Incidentally, you need to keep a lookout for coded, political answers like: *"We are sure that any system you build will be reliable."* This appears not to answer the question (hence the political reference) but in fact it does. It decodes as: *"We can't afford an unreliable system, and you will be fired if we get one. But, of course, we don't want to spend any extra money."* If you receive such an answer, as far as you and your career are concerned, this is now a fully mission-critical system.

If you think your client-server database is or will become mission critical, bear the following in mind. Most people are used to the fact that a typical PC costs around £1,000. What you have to do is make them realise that while that's fine for a word-processing machine, it's totally inappropriate for a mission-critical database server. For a start, database servers, by their very nature, work harder and need to be of a higher specification than a normal PC. For another thing, if a word-processing PC fails, those important letters can be typed up on another PC. If your database server fails, what are you going to run the company accounts on?

## The nitty gritty

So, as you will have guessed, we have reached the part of the column where I try to get you to spend a small fortune on your server hardware. Please understand that as I try to part you from your hard-earned cash, I don't have shares in any of the hardware vendors who may be mentioned here. Rather, I just want to make sure you keep your job.

p270 >

## Client server definitions

■ **Standalone database** — Runs on a single machine. That machine is typically a PC running Windows and an RDBMS (such as Microsoft's Access). The data and the data-processing engine all reside on this one machine, so multi-user access to the data is not possible.

■ **Client server** — The data and the data-processing engine are moved across the network and run on a dedicated machine called a database server. Since multiple-client PCs can access this server, the system becomes multi-user. The clients will still

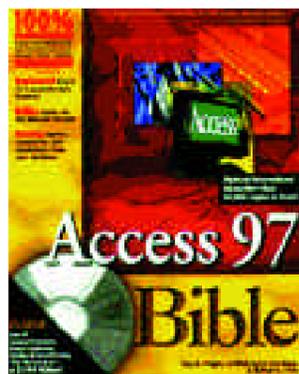
typically run Windows and some sort of interface to the database.

■ **Client** — A PC running Windows (of whatever flavour).

■ **Server** — A computer (not necessarily Intel-based, but probably so) which runs a dedicated RDBMS back-end. This will not be Microsoft Access, since it cannot run as a database server, but instead will be something like Microsoft's SQL Server, Oracle's Oracle, or IBM's DB2. The server will also run on a server operating system such as Windows NT, OS/2 or UNIX.

**Book review: Access 97 Bible**

Cary Prague has established himself as one of the more authoritative Americans to write about Access. His earlier books have been excellent and this one (over a thousand pages long), written in conjunction with Michael Irwin, is no exception.



It is only when the authors stray into areas which are more to do with the relational model than Access that the content becomes a bit flaky. As an example, they attempt to distinguish between a one-to-many join and a many-to-one join. Since these are, as the authors acknowledge, essentially the same animal viewed from a different direction, it

relationship. Surely a misprint?

However, this minor carping on my part should not, under any circumstances, prevent you from buying this book for the treasure-chest of Access gems with which it is stuffed. For novice Access users, it will be an invaluable road-map. For experienced users, it is a wonderful source of tricks and tips. As is so often the case with books of this type, one single example or tip can save you, say, a couple of hours' work. This is a "must-have".

It tells you how to use Access' GUI and macro language but stops short of programming in VBA (Visual Basic for Applications). Subjects are covered in detail and with accuracy.

seems like needless obfuscation. It isn't helped by their statement during this rather bizarre interlude, that a many-to-one relationship is (in theory) a one-to-one

■ **Access 97 Bible** by Cary Prague & Michael Irwin. £42.99 (IDG Books, ISBN 0-7645-3035-6) from Computer Manuals 0121 706 6000

Also, bear in mind that, as before in this column, I will name names and quote figures, but you must only regard them as approximates. Do not base your entire business strategy on the figures I quote, because I don't know what your business requirements are. It often takes a couple of days' consultancy work to provide accurate figures for a given company. But as I hate reading evasive articles, I'll give ballpark figures which I think are reasonable for the average small-to-medium-sized enterprise (SME).

**Buying a server**

Buy it from a reputable company. I know they charge more but you usually get what you pay for, not only in terms of reliability but also in terms of compatibility (important with the kind of server OS you will be running) and support.

Good names here are Compaq, IBM, HP, Olivetti, Apricot, NetFrame etc. Note the "etc". Just because I haven't named a supplier, doesn't mean it produces poor products. On the other hand, don't assume just because a supplier can make a reasonable PC, it can also make a good server. Buy from a manufacturer with a good track record in making servers.

**Questions and answers**

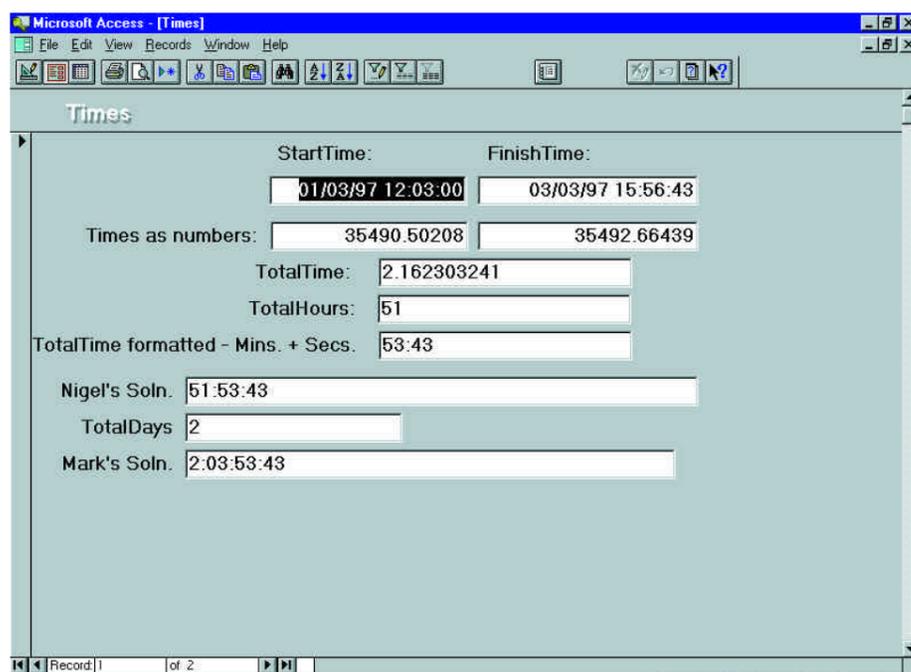
I have received a few readers' letters during the past couple of months, so let's deal with these.

**Q.** "Regarding the problem posed by Gareth Wade in your April column, which was about the problem of displaying times greater than 24 hours in Access. It does not have the [h]:nn:ss format that Excel uses for

displaying hours greater than 24. As far as I am aware there is no format, as such, that will solve this problem. Access has other uses for square brackets.

"I experienced a similar problem when totalling the times in a relay race that took place over two days. I realised the total time for a team could be greater than 24 hours and would then revert back to 00:00. I fixed it in a hurry by displaying just the minutes

**Fig 1** There are several ways in which times can be manipulated



**Fig 2** A design view of the same form

[see also Fig 1]

and seconds with nn:ss format and dealing with the hours separately. As I am sure you are aware, dates and times are stored as double precision numbers that represent the number of days after 30th Dec, 1899, so `Int([TotalTime]*24)` displays just the number of hours. The two parts can easily be combined into a single text box by making its control source

```
=Int([TotalTime]*24) & ":" & Format([TotalTime], "nn:ss")
```

"While not actually being a formatting solution, this nevertheless seems to address the problem."

**Nigel Collins**

**A.** Nigel's solution is very neat. I have included a form which shows his solution, as well as some of the intermediate steps (Figs 1 & 2). The first two text boxes just show data from the two fields in a record. These are the start date/time and the finish date/time of some mythical event. These text boxes are formatted as General Dates. The next two show the date/times formatted as numbers with fixed numbers of decimal places. As Nigel points out, the dates/times are really double-precision numbers that represent the number of days after 30th December 1899.

In Fig 2, you see that the syntax of the control source for these text boxes is of the type

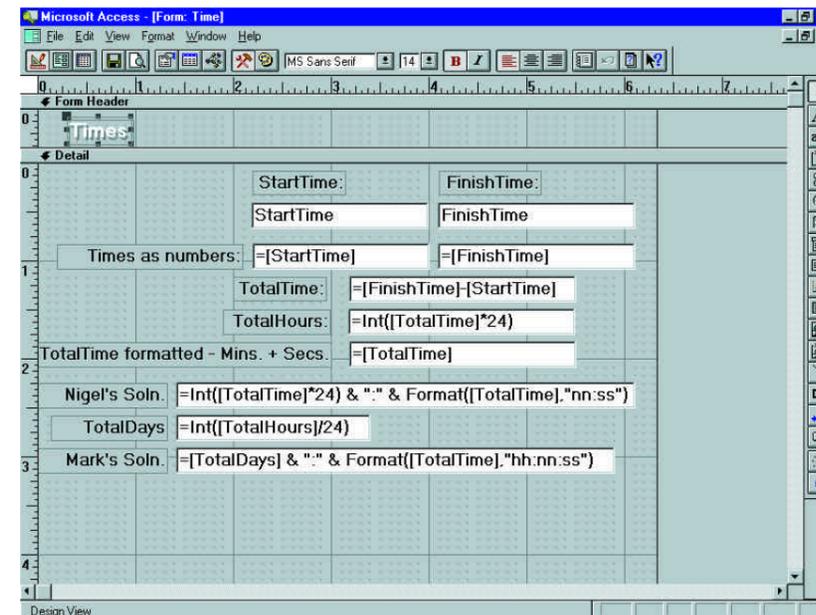
```
"=[StartTime]" - not "StartTime"
```

Unless you use this syntax, Access will not allow you to choose a format like Fixed for the text box (although you can still type that format in by hand).

The third row on the form shows one text box which is calculating the difference between these two numerical values, while the fourth is calculating the total number of whole hours between the two dates.

The fifth shows the difference formatted as minutes and seconds. Formatting in this way effectively tells Access to ignore the hours and show only the minutes and seconds. (Note that this uses "nn:ss" and not "mm:ss" as you might expect.)

Nigel's solution is shown next and elegantly combines these two methods of showing the time difference. This led me to realise that we can also calculate the total number of days relatively simply and hence display the information in days, hours, minutes and seconds if so desired, as



shown in the final two text boxes in Fig 2. This is clearly not better than Nigel's solution, it's just different.

**Q.** "In your March Databases column, you mention the subject of storing hierarchical data in an SQL database.

"The scheme you outline, which essentially uses a pointer to the next higher item in the hierarchy, is conceptually simple and very easy to understand. Sadly, as you noted, it is a pain to program with. You may be interested to know of a couple of articles written by Joe Celko. They were published in the weekly magazine Computing (19th January 1995 and 26th January 1995 editions).

"The concept he describes uses nested sets to represent hierarchical levels. Celko gives an example, using a simple organisation chart to demonstrate the tree structure. He also outlines a set of queries which allows you to do important things like find all the leaf nodes (those which have no further dependents) and identify the hierarchical chain of command from any individual, from the lowest to the top level.

"Please do follow this up and publish something in your column. Although relational databases are very good and very useful, there is nevertheless still a lot of data in the world which is hierarchical."

**Richard Howells**

**A.** Joe Celko is an American writer whose work on SQL I know and regard highly. Those interested in more efficient ways of

storing hierarchical information would benefit from obtaining and reading these articles. Or send me some email, and if enough people are interested in the subject, I'll follow it up for you.

**Q.** "I run a photographic model agency and we keep details of our models' characteristics, such as hair and eye colour, on our Access database. Sometimes we receive requests for "all girls with blonde hair". This is easy, as I can prepare a query for "blonde". However, I cannot seem to be able to request "all" hair colours. There appears to be no way of presenting a wild-card search as one of the input fields on the query through "[Input hair colour:]".

"Can you think of any way of handling this criterion without resorting to non-Access code such as Basic?"

**Mike Illes**

**A.** I presume you're using parameter queries, in which case this syntax may be useful:

```
Like [Input hair colour:] & ""
```

If you enter Blonde, it will find all Blondes. If you enter nothing, it will find all records. It also does a "fuzzy" search, so that you can mis-spell Blonde and still find the appropriate records.

**PCW Contact**

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column, at [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk).



# In the round

The subject of rounding in Access gets a discreet revival, and your contributions are sought for the most unusual RDBMS application. Mark Whitehorn briefs you on what's required.

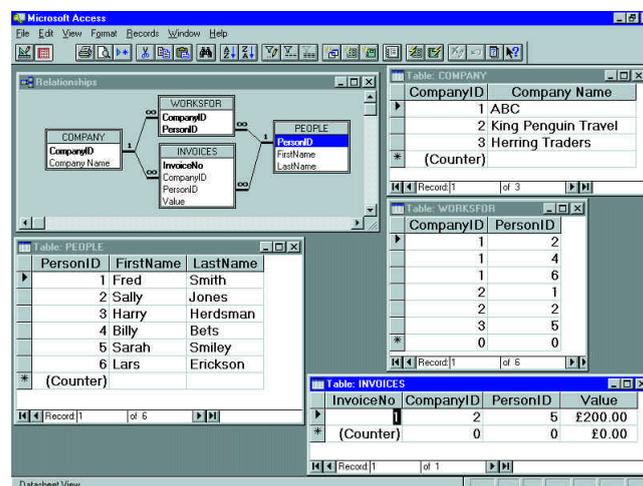
**T**his column has harboured several discussions about rounding functions in Access. I thought all had gone quiet until two intriguing emails arrived, one from Roger Moran and the other from Ray Hall. While some of us (myself included) find this topic fascinating, I am loath to devote much more space to it since it may be of limited interest to some people. So, I have included their emails in full as memo fields in the DBCMAY97.MDB file on the CD-ROM. This ensures that the information is available to those people who wish to look at it, but doesn't soak up bandwidth for those who aren't. See the form called "Rounding" if you are interested, and thanks to Roger and Ray for their contributions.

## Competition time!

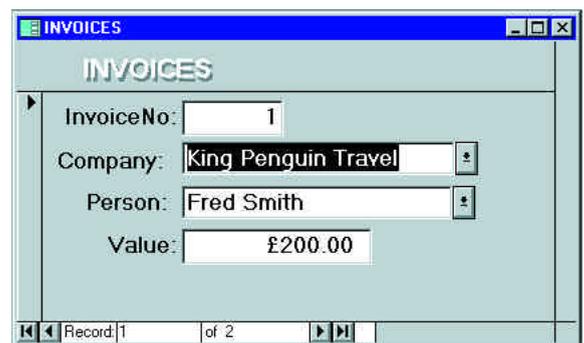
From Andrew Leaman: *"I am interested in databases but find it difficult to think of useful applications. Could you possibly supply a list of typical applications, starting at the simplest and working up to the more complex?"*

Starting with the most basic database is easy: an address list, maybe a list for sending Christmas cards. As to the more complex uses, these are almost without number and range. Banking systems, air traffic control systems, process control systems in factories — all have at their heart some form of database. In fact, it is possible to argue that almost all computer applications are essentially databases; some of them just have rather odd front-ends.

Take a word processor, for example. It stores and manipulates data. You can query the database (with the search facility),



**Fig 1** Taken from DBCMAY97.MDB and showing the tables from Greg Barstow's question



**Fig 2** The two combo boxes on the form from DBCMAY97.MDB

generate different forms (normal view, outline view, etc.) and generate reports (print-out). You see? A document is really a database and a word processor is really a Data Base Management System.

I wouldn't want to take this argument too far, but what about a competition? We'll offer one of our much sought after book/record tokens for the most unusual example of an application developed with a recognised RDBMS. It doesn't have to be developed by you or your company, but if it happens to be so, that's fine. Just to start the ball rolling, I've heard of a really odd

application, developed in Australia, which made the international news recently... But I'll leave it open for a reader to suggest that one since I'm sadly excluded from the list of potential prize-winners.

## Sets and subsets

Greg Barstow writes: *"I do freelance for a range of companies. Each company has a number of people who can commission work from me for the company concerned. I need to generate invoices for each piece of work, and I want a form in Access which lets me pick the company from a combo*

box (which I can do easily). Then I want the next combo box (which allows me to pick the person to whom the invoice should be sent) to show me only the people who work in that company."

This is a good generic question. Essentially it asks: "Given a long list of options which can be unequivocally sub-setted by a choice in another list, how do I show this elegantly on a form?" There are many applications. If you have different sales people who work on different product lines, or different aircraft which are serviced by different engineers, this is an area which may be of interest to you.

For one possible solution, see DBCMAY97.MDB. Fig 1 (p285) shows the tables involved and a small quantity of sample data. It also shows a tempting but incorrect set of relationships between the tables. Those who enjoy conundrums can work out why this particular set of relationships works but is non-optimal. The answer is on page 288 (the relationships in the MDB file on the CD-ROM are correct).

The form (Fig 2) has two combo boxes. The upper one is straightforward. It looks up values in the table COMPANY:

```
Select [CompanyID],[Company Name]
From [COMPANY];
and writes the value from
COMPANY.CompanyID into
INVOICES.CompanyID.
```

The lower combo box is more devious. It pulls data from a query called TheRightPeople rather than directly from PEOPLE. The purpose of this query is to find the people who work for the company which has been selected in the upper combo box.

The SQL for this query (Listing 1) is, like a great deal of SQL, impenetrable at first glance. Translated into English (more or less) it says:

- Find the value which is in the combo box above.
- Use that value to find the correct record in the COMPANY table.

**Listing 1 SQL for TheRightPeople query**

```
SELECT DISTINCTROW COMPANY.CompanyID,
[FirstName]+ " "+[Lastname] AS Name, PEOPLE.PersonID
FROM PEOPLE
INNER JOIN (COMPANY INNER JOIN WORKSFOR
ON COMPANY.CompanyID = WORKSFOR.CompanyID)
ON PEOPLE.PersonID = WORKSFOR.PersonID
WHERE ((COMPANY.CompanyID=[forms]![invoices]![companyID]));
```

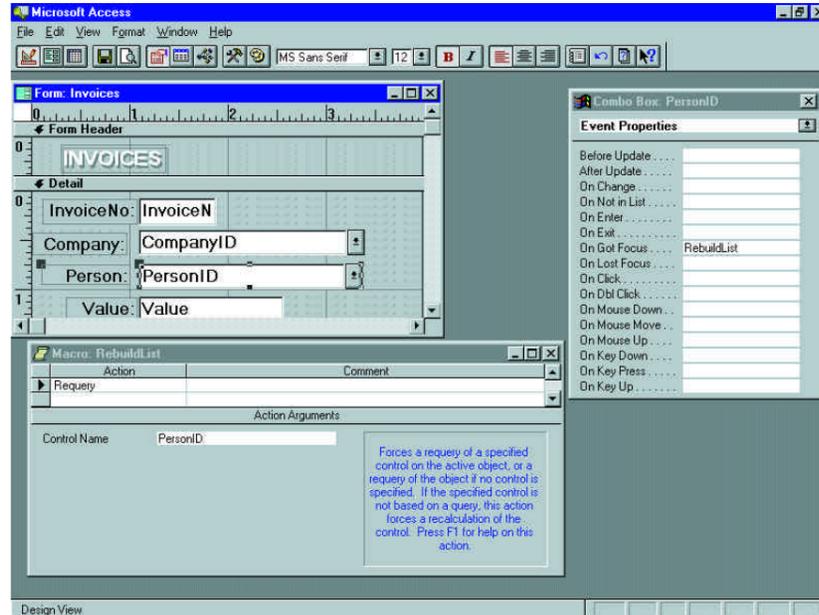


Fig 3 The extensive, and highly complex, macro used to force a re-query of the lower combo box on the form

in the COMPANY table.

- Then use that to find the people who work for the company. This has to be done via the table called WORKSFOR, since that is the table which stores the information about who works for which company.
- Finally, collect the relevant person's ID number and assemble their name neatly, attaching the first and last name together so that it looks tidy in the combo box.

If we run through that again, we can add in the relevant bits of the SQL statement:

```
Find the value which is in the combo box above.
[forms]![invoices]![companyID]
Use that value to find the correct record in the COMPANY table.
WHERE ((COMPANY.CompanyID=[forms]![invoices]![companyID]))
```

Then use that to find the people who work for the company; this has to be done via the table called WORKSFOR, since that is the table which stores the information about who works for which company.

```
FROM PEOPLE
INNER JOIN (COMPANY INNER JOIN
WORKSFOR
ON COMPANY.CompanyID =
WORKSFOR.CompanyID)
ON PEOPLE.PersonID =
WORKSFOR.PersonID
```

Finally, collect the relevant person's ID number and assemble their name neatly, attaching the first and last name together so that it looks tidy in the combo box.

```
SELECT DISTINCTROW
COMPANY.CompanyID,
[FirstName]+ " "+[Lastname] AS Name,
PEOPLE.PersonID
FROM PEOPLE
```

The nett result is that, when the second combo box is opened, the only names that appear belong to people who work for the company you have just chosen in the upper combo box. At least it would, except that it doesn't work automatically all the time because Access often buffers information so it doesn't automatically re-query the source for a control. The lower combo box has the query called TheRightPeople as its source. If this query has already returned an answer table, then simply selecting a different company in the upper combo box doesn't cause TheRightPeople to be re-run; hence the list of people may not be up to date when it appears in the lower combo box.

The answer is to force a re-query every time you use the lower combo box. This can be done in code or with a macro, and

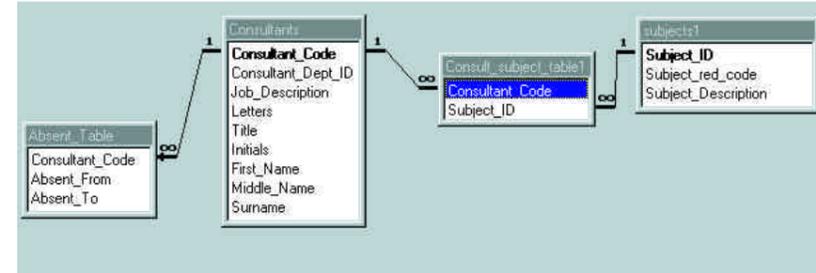


Fig 4 The tables used in David Ruffel's database

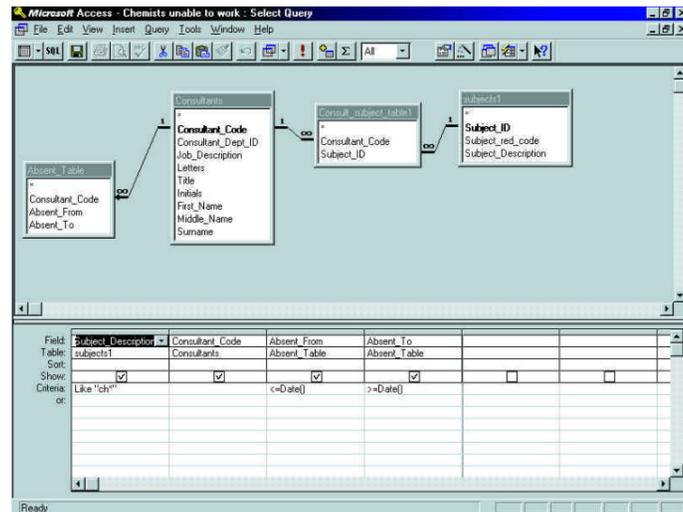


Fig 4a The GUI version of a rather impenetrable SQL statement

since I usually demonstrate everything in code, I thought for the sake of variety I'd use a macro (Fig 3).

David Ruffel emailed in a question which seems to have general application. He maintains a list of consultants who are experts in various areas — mathematics, computing, etc. There is a many to many relationship between the consultants and their subject areas, hence three tables are needed to model this relationship, while a fourth table contains information about times when the consultants are absent (Fig 4). Finding those consultants who can provide information about a given subject, say, Chemistry, presents no problem

(Listing 2, p288).

However, David also maintains a table of the dates during which particular consultants are unavailable. What he needed was a query which found not only the consultants who were experts in a particular area, but also those who were available on a particular date.

To my twisted mind, the easiest way to solve this one is to use one query to find all of the Chemists who are unable to work. This can be accomplished with Listing 3 which looks pretty horrible if you aren't used to SQL, but is much more understandable in Access' GUI (Fig 4a). This produces an answer table which looks like Fig 5.

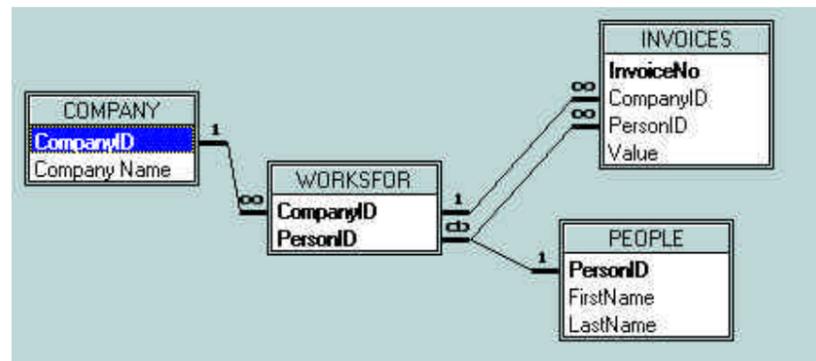


Fig 7 Showing a better set of relationships to use in DBCMAY97.MDB

The ConsultantCode identifies the Chemist who can't work on the given date (in this case I am using the Date() function to return today's date). Then, a simpler query can use the information in this answer table to identify all of those who can work (Listing 4).

This solution may not be optimal and I haven't tested it extensively, but you might want to use something akin to it if you have a similar problem.

The sample file is on the CD-ROM as an Access 7 file called QP4.MDB.

The format of this file (Access 7) brings me to another email, from Dave Milor.

**Versioning**

"Is there any specific reason why you write your articles with reference to Access version 2 but are using the Windows 95 interface? I have heard that there are problems with version 7 regarding speed and possible bugs."

I use Access 2.0 whenever possible simply because Access maintains compatibility in only one direction. If I provide a solution in Access 2.0, anyone using that version or later can read it. However, if I supply an MDB file in the most recent version, Access 97, only those people with that version can use it. All of the machines that I now use are running either Windows 95 or NT 4, which is why the screenshots of Access 2 appear as they do. When run under 95 or NT, Access 2 "acquires" the look and feel of these particular systems.

With regard to bugs, Access 7 certainly has them; but then, so does every bit of software I have ever seen (including my own). I haven't come across any which would make the product unusable.

The speed issue is more complex. Given any RDBMS, speed considerations can be split into two areas. First, there is how fast the interface runs. This is non-trivial, since a slow interface makes development work painful and will upset end-users. Second, there is data-processing speed: essentially the speed with which queries run. This is very different, and also clearly non-trivial. This second measure of speed is the one which people like myself love to benchmark, but we shouldn't ignore the interface speed, even if it is more difficult to quantify.

So, what about speed in Access 7? The interface speed is worse than Access 2.0, but the data processing speed is a bit better with certain queries. Access 97 (the

### Listing 2 Finding Chemistry consultants

```
SELECT DISTINCTROW subjects1.Subject_Description, Consultants.Title, Consultants.Surname
FROM subjects1
INNER JOIN (Consultants INNER JOIN Consult_subject_table1
ON Consultants.Consultant_Code = Consult_subject_table1.Consultant_Code)
ON subjects1.Subject_ID = Consult_subject_table1.Subject_ID
WHERE (((subjects1.Subject_Description)="Chemistry"));
```

### Listing 3 Finding the Chemists who are unable to work

```
SELECT subjects1.Subject_Description, Consultants.Consultant_Code, Absent_Table.Absent_From,
Absent_Table.Absent_To
FROM subjects1
INNER JOIN ((Consultants LEFT JOIN Absent_Table
ON Consultants.Consultant_Code = Absent_Table.Consultant_Code)
INNER JOIN Consult_subject_table1 ON Consultants.Consultant_Code = Consult_subject_table1.Consultant_Code)
ON subjects1.Subject_ID = Consult_subject_table1.Subject_ID
WHERE (((subjects1.Subject_Description) Like "ch*")
AND ((Absent_Table.Absent_From)<=Date())
AND ((Absent_Table.Absent_To)>=Date()));
```

### Listing 4 Identifying the Chemists who can work

```
SELECT DISTINCTROW [Able Chemists].Consultant_Code, Consultants.First_Name, Consultants.Surname
FROM Consultants
INNER JOIN [Able Chemists]
ON Consultants.Consultant_Code = [Able Chemists].Consultant_Code
WHERE ((([Able Chemists].Consultant_Code) Not In (Select Consultant_Code from [Chemists unable to work])));
```

Fig 5 Answer table from Listing 3

Subject Area Description	ConsultantCode	Away From	Until
chemistry	5	Thursday, January 02, 1997	Thursday, June 19, 1997

next version on) definitely requires a better machine to run the interface (in other words, it is even slower). However, the data-processing is markedly faster, even given machines of the same spec.

#### A client-server future

In the February issue I asked if the column was too biased towards Access and, more generally, what did people want me to cover in future issues. The area which was overwhelmingly popular as a new topic was practical information regarding client-server databases.

This email from Tom Cliff was typical:

*"I've been building databases for a couple of years and all of them have been on standalone PCs. Some of these have grown and now need to be migrated to a client-server system, because the volume of data has increased or they need to become multi-user. What I need is a good overview/*

Fig 6 Entry into the INVOICES table

InvoiceNo	CompanyID	PersonID	Value
1	1	3	£200.00
2	1	4	£0.00

*description of what is involved. How much work is it? How do I actually do it?"*

So, we'll start next month having a look at the hardware you need, then move on to actually installing a back-end RDBMS, setting up a database on it, connecting to that database from clients, and so on.

#### Answer to conundrum

The relationships shown in Fig 1 ensure that the values inserted in INVOICES.PersonID are drawn from the available list of people (which is kept in table PEOPLE). They also ensure that the values inserted into INVOICES.CompanyID are drawn from the list of available companies. What these

relationships *don't* forbid is an entry into the INVOICES table like Fig 6. This is bad, because person 3 doesn't work for company 1.

A much better set of relationships to use are those shown in Fig 7 (p287). These ensure that the INVOICE table can only ever contain "meaningful" combinations of CompanyID and PersonID.

#### PCW Contact

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column, at [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk)



# All buttoned up

Mark Whitehorn presents another slant on last month's buttons and captions query, which is far from sewn up. Plus, the complexities of data entry in Access, and CUSTOMER care.

In the February issue I published a letter from Glen Rowe. He wanted multiple buttons on a form, all with different captions. Whenever a button was pressed, a query would run which returned values based upon the caption on the button. For example, if you press the button labelled "Penguins" you see the records which relate to Penguins: pressing the "Fish" button yields information about Fish. I produced a solution by building a query which snatched the caption from the button which had just been pressed.

I also wrote the following: "The obvious solution at first is to try to pass the button's caption to the query as a parameter. As far as I know (and I stand to be corrected) this can't be done."

James Talbut replied: "Passing a parameter to a parameter query in Access can be done, but it's not very pleasant and means the you need to use DAO (Data Access Objects), which can be good or bad depending on the circumstances."

The solution that James then provides does, as he suggests, pass a parameter to a parameter query. What it doesn't do is pass the *caption* of the button as a parameter (which is still, as far as I know, impossible). In turn, this means his solution is somewhat more awkward than the one shown in the February issue because every button you add to the form must have its caption *and* its OnClick property altered. With the February solution, all you had to do was clone the button and alter the caption. However, James's solution to the initial problem is still well worth studying since it illustrates a very different way of solving the conundrum.

He continues: "Before I include all the bumf to show you how it's done, I'll just

**Fig 1** Query, "Person":

```
SELECT DISTINCTROW People.ID, People.Name, People.Value
FROM People
WHERE (((People.Name)=[Person]));
```

**Fig 2** Form Code, behind "People"

```
Private Function Button_Click(sName As String)
Dim qdfPeople As QueryDef
Dim rsPerson As Recordset
Dim sRowSource As String

Set qdfPeople = CurrentDb.QueryDefs("Person")

qdfPeople.PARAMETERS("Person") = sName
Set rsPerson = qdfPeople.
OpenRecordset

While Not rsPerson.EOF
If Len(sRowSource) > 0 Then sRowSource = sRowSource & ";"
sRowSource = sRowSource & rsPerson.Fields("Name") & " - " &
rsPerson.Fields("Value")
rsPerson.MoveNext
Wend
List10.RowSourceType = "Value List"
List10.RowSource = sRowSource
End Function
```

mention the implications.

"To actually access the resultant data you need to use DAO to step through a recordset. If you have a large amount of data this can be inefficient (OK, it's always inefficient, but it's more noticeable with a large dataset), particularly if you are intending to perform some secondary operation/selection on the data. However, the query itself is still run by the JET engine so the main data manipulation routines are still run as quickly as Access can do them.

"To control the parameters you need to use the Parameters collection of the QueryDef object for the query you are

interested in. The most simple example I could come up with looks like this:

**Table: "People"**

ID	Name	Value
1	Fred	1
2	Jack	4
3	Fred	9
4	Harry	16
5	Anastasia	25
6	Bob	36
7	Martin	49
8	Stephen	64
9	Harry	81

(See Fig 1.)

Form, "People":  
Six command buttons, each with their OnClick property set to:  
=Button\_Click("Fred")  
with Fred replaced by the caption on that particular button. The buttons are labelled Fred, Bob, Harry, Martin, Anastasia and Stephen. At the bottom is an empty list box called List10.  
(See Fig 2, page 285.)

"And that's it. Hope it's useful."  
Very interesting, James. I have constructed an example file (in Access 7.0) based on this example called TALBUT.MDB which is on the CD.

Next an email from Norway: "I have a problem that I've worked on for several months now, without finding any easy solution," writes Tore Saetre, of Bergen. "I've developed and am maintaining a Microsoft Access database that keeps track of customers. The database contains two main tables: Customer and Sales. The problem is that I need to have a field in the Customers-table that shows how many orders the customer has in the Sales-table.

"My first idea was to make a query that counts orders in Sales for each customer. The query lists each customer and how many orders the customers has. I tried to move that value to the Customer-table with an update query that updates a NoOrders-field in Customer through a link to the first query. The problem is that the first query can't be updated and the second query won't update my Customers-table."

I receive many questions like this one, questions which hinge upon the desire to store redundant data in tables. My initial reply to Tore was that, in general, storing derivable data is a bad idea because if the data changes (as you make more sales), the data in the CUSTOMER table goes out of date. It is normally preferable to use a query to calculate the information you need: whenever you want to see this data, you can run the query which shows you the customer details and the number of sales.

Tore was only partially convinced and replied: "But data in the sales table is only changed and added to once a month. The rest of the month the user browses through CUSTOMER's 4,000 records and needs to see as much data on each customer as possible. Running a query for each customer is too time-consuming. I see why it can't be done in a ordinary customer/sales database, but I still hope to find a solution to this problem.

**Fig 3 (right)** Small samples from both the CUSTOMER and ORDERS tables in the database "Tore"

CUSTOMER No	TITLE	FIRST NAME	LAST
1	Mr	John	Knight
2	Mr	Alfred	Clark
3	Ms	Margaret	Wintert
4	Mr	Duncan	Walker
5	Mr	Joseph	Whyte
6	Mrs	Norah	Cooper
7	Mrs	Helen	Lynch
8	Ms	Grace	Falcon
9	Mrs	Mary	Robb
10	Mr	George	Peders
11	Ms	Elizabeth	Chalme
12	Mr	John	Walker
13	Mrs	Ann	Dick
14	Mrs	Helen	Taylor
15	Mr	Thomas	Falcon
16	Mrs	Lilias	Murray
17	Ms	Ethel	Holmes
18	Mr	Joseph	Ferrie
19	Mrs	Agnes	Angus
20	Mr	David	Whyte

OrderNo	CustomerNo	Order Details
1	2	Bar Bar
2	5	Foo
3	6	Baa Foo Baa Foo
4	4	Baa Baa Foo
5	6	Baa Foo Foo
6	6	Baa Foo Baa Foo
7	5	Foo
8	6	Baa Foo Baa
9	6	Baa Foo Foo
10	55	Baa Foo Baa Foo
11	535	Baa Baa Foo
12	54	Baa Foo Baa Foo
13	35	Baa Foo Baa
14	6	Baa Foo Foo
15	36	Baa Baa Foo
17	5	Foo
18	526	Baa Baa Foo
19	26	Baa Foo Baa Foo
20	5	Baa
21	6	Foo Baa Foo
22	6	Foo Baa Foo

**Fig 4 (left)** The form called "Customers & Orders" shows customer details and order details

"For an experiment, do you have an example of how I can run a query to see customer details based on the customer currently displayed in a form? Do I need to run a code, macro or function and how do I get the query to filter everything but the record on-screen? Thanks for your help."

Okay, we'll solve this one both ways. There is a sample database on the cover CD (in Access 2.0) called TORE.MDB. This has two tables ,CUSTOMER and ORDERS. CUSTOMER contains simple details of 4,000 mythical people, ORDERS contains about 27,500 orders. Each customer has at least three orders to their name (or, in this case, to their CustomerNo) and some have considerably more. Fig 3 shows small samples from both tables.

The form called "Customers & Orders" (Fig 4) shows customer details and order details. The main form is based on CUSTOMER, the sub-form (which is called Sub1) is based on a query which simply

performs a join on the two base tables. This form displays one record for each of the 4,000 customers.

The form called "Customers & Orders Count" (Fig 5) again shows one record for each customer. However, instead of listing the details about each order, it simply shows how many orders each customer has placed. This form is based upon the query called "Customer & Order Count".

These two forms are variants which answer Tore's request for an "example of how I can run a query to see customer details based on the customer currently displayed in a form".

I ran this database on a 486/100 with

**Fig 5 (below)** The form called "Customers & Orders Count" shows one record for each customer

16Mb of RAM using Access 2.0. "Customers & Orders" opens almost instantaneously, and you can scroll through the records at the rate of about 400 per minute. "Customers & Orders Count" takes about 25 seconds to open, but once open you can scroll through the records at the rate of about 1,600 per minute.

To put this into simplistic terms, the first one does the calculations for each record as you ask to see it (which is adding about 0.1 seconds to the scroll time between each pair of records). The second one does all of the calculations for all of the records before showing any of them to you.

These forms are simply based on queries: no code or macros are required. As you can see, using the first form, it is possible to see customer and order details with essentially no speed hit at all for these numbers of records.

The query called "Customers & Orders Count" (the one which takes 25 seconds to

run) can also be used, if required, to create the base table that Tore requested. I have included a Make-Table version of the same query in TORE.MDB. If you run this, it will create a base table called CUST which is just like CUSTOMERS except that it includes a field which shows how many orders each customer has placed.

In Tore's case this could be run once a month, and then CUSTOMERS replaced by CUST; in that case, the 25-second speed hit disappears. The downside is that we are now reliant upon redundant data in the database. If anyone forgets to renew the tables at the end of the month, then all of the users of the database start to work with inaccurate data. You pays your money, you takes your choice — speed or security.

**Quickies**

Here's a quick one from Gareth Wade: "When running a simple Report, I need to add all the time lengths and give a running

**Fig 6**  
Private Sub LastName\_AfterUpdate()

```
End Sub
Simply add a line so that it now reads;
Private Sub LastName_AfterUpdate()

    Me!LastName = StrConv(Me!
LastName, 3)

End Sub
```

**Fig 7**  
Me!LastName = StrConv(Me!LastName, 3)

**Fig 8**  
Private Sub LastName\_AfterUpdate()  
Dim Length As Integer  
Me!LastName = StrConv(Me!LastName, 3)  
If Left(Me!LastName, 3) = "Mac" Then  
 Length = Len(Me!LastName)Length - 3  
 Me!LastName = StrConv(Me!LastName, 3)  
 Me!LastName = "Mac" Me!LastName  
End If  
If Left(Me!LastName, 2) = "Mc" Then  
 Length = Len(Me!LastName)  
 Me!LastName = Right(Me!LastName, Length - 2)  
 Me!LastName = StrConv(Me!LastName, 3)  
 Me!LastName = "Mc" + Me!LastNameEnd If  
End Sub

total. Easy enough; but when Access gets to 23:59, it reverts back to 00:00. What format do I use to carry on past that magic 24-hour mark? I have tried all the manuals and help files but can't find the answer."

I don't know the answer to this one.

Anyone else care to solve it?

■ And this one from Malcolm Rowley: "I find I am put off by the complexity required in Access to solve simple problems, e.g. data entry character formatting.

In Alpha I simply have to specify the word/character format I require, e.g. Upper, Lower, Word (first character upper case, the rest of the data left as is) in the field rules, and this is applied to all data entry that takes place and can be re-applied to all existing data. This is ideal for data entry of names and addresses; the only difficulty being the McCartneys (etc.) of this world. To have any simple form of formatting that stores the data in an access table in a particular format, do I really have to resort to code? Is there a simple way of accepting data entry in Access and storing it as first character upper case, the rest left as is?

For example: macdonalds becomes Macdonalds / macIntyre becomes MacIntyre / 27 park avenue becomes 27 Park Avenue. A simple solution would be appreciated."

The answer is that you have to use code, but only one line so it may just meet your requirement for simplicity! Suppose you have a form called Capital which displays a field called LastName in a text box called LastName. Switch to design, double-click on the textbox in question, select the Event properties, click on the one called "After Update", click on the ellipsis button which appears (three dots) and choose Code Builder. Between the two lines which appear (Fig 6) "Me" means the current form, "LastName" is the name of the textbox, and "StrConv" is a function (only available in Access 7.0 and above) which converts strings. The "three" tells the string conversion function to do the type of conversion that you asked for (capitalising the first letter of each word).

So, Fig 7 means "make the contents of the textbox called LastName equal to the same as it is now, but with all of the first letters of the words capitalised". This will convert: penguin penguinsson to Penguin Penguinsson / 23 the larches to 23 The Larches / mcdonald to Mcdonald and so on.

■ On a slightly different tack, I haven't used Alpha for well over a year, but I agree that

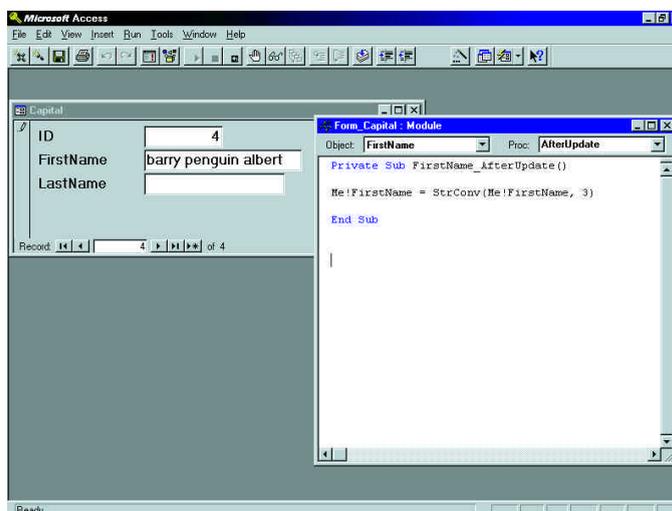


Fig 9 (left) This single line of code will capitalise all distinct words placed in the FirstName field (see Fig 10)

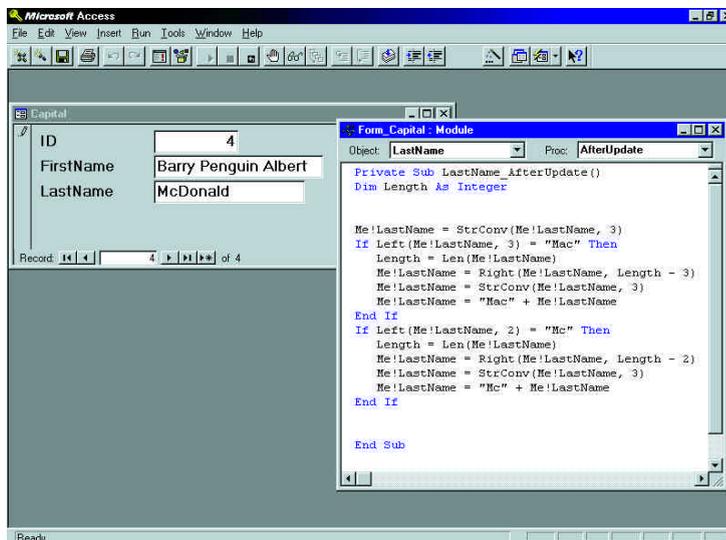
Fig 10 (below) This more complex code capitalises two of the more complex surname types in the LastName field

processes like this one are easier (and certainly more intuitive) in Alpha than in Access. In fact, to quote from a review I wrote at the time of its release: "...of these, the Field Rules are the most impressive, if only because Alpha copes with them better than any other RDBMS I have seen."

There is clearly a trade-off in design terms, which these two products exemplify. In both products, the designers identified a host of commonly needed functions and made them easily available from the interface. This approach has pros and cons. The good news is that commonly performed processes (like designing tables and building queries) are easy: the bad news is that if you happen to want a function that the designers didn't provide, you suddenly have to go to a more complex area, such as coding.

One difference between the two products is where the line was drawn. In this case, the Alpha designers decided to include capitalisation in the "easy" set and the Access designers didn't.

I have always found that, whatever product I use, I eventually reach a stage where I have to use code simply because it is impossible for the designers to put



everything into the "easy" set. On a happier note, once you do start to use code, your horizons expand considerably. You can, for example, begin to deal with unusually capitalised names like those you mention. If you expand the code to that shown in Figs 9 and 10, this will change macdonalds to MacDonalds, and mctavish to McTavish. (Also see the Access 7.0 sample database called ROWLEY.MDB.)

I am not suggesting that this is perfect code. As usual in the sample code I give, there is no error trapping. In addition, there are some people who prefer the letter after Mac or Mc to be left in lower case. This code is simply provided as an example of what can be done.

## PCW Contact

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column at [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk)



# Animal magic

Mark Whitehorn casts a beady eye over hierarchy in the zoo and shows a way of handling categorisation data. In addition, he passes on some improved solutions to past problems.

**A** reader called Andrew writes: "First of all, may I say how much I have enjoyed your column in *PCW*" (*ah, I love flattery — MW*). "But unfortunately it has got me thinking about a problem which I have never satisfactorily resolved for myself." (*Sounds like work! — MW*). "It involves a recursive link field in a table:

Field	Name	Type	Description
1	Code	Inc	Primary Key
2	Parent	LongInt	Parents Key
3	Category	String	Category Title

"The idea is to provide a hierarchy by using the table to emulate a tree structure which could be used for categorising animals, or departments in an organisation.

"I have been searching for an efficient way of finding all the categories that are below the current item. My best effort has been to run a query that finds all records whose parent is current code, add these to a result table and then query the answer to find all the records where there is a match on any of the codes. I keep this going until there are no records returned (when I've reached the bottom of the tree). Although this works, it does not seem very efficient. Do you have any ideas?"

What I have done is to look at ways of handling data from a table like the one Andrew describes.

**Fig 1** shows such a table. It describes a hierarchy which is used to name the animals in a zoo.

Each record simply represents a single link in the structure, so the data is stored in a reasonably economical manner. However, you

need to look at several records to trace back up the structure.

So, for example, record 12 tells us that Harry's parent record is number three, which in turn tells us that Harry is a King. This in turn leads us to record two, and hence to record one, by which time we know that Harry is a King Penguin, which is a species of bird. The storage is economical, but a little unwieldy.

**Fig 2** illustrates the information in a more legible state, and **Fig 3** shows the query. Note that the answer table is not showing one record for each record in the base table. This is because it takes several records to define one "branch" of the tree. The SQL is a little tortuous, but essentially it

is merely an expanded version of this:

```
SELECT DISTINCTROW Animals.Category
AS First, Animals_1.Category AS
Second
FROM Animals AS Animals_1
RIGHT JOIN Animals ON
Animals_1.Parent = Animals.Code
WHERE ((Animals.Parent=0))
ORDER BY Animals.Category,
Animals_1.Category;
```

which works for two levels of hierarchy.

This naming hierarchy is a little anarchistic in that names can appear at several levels. Suppose we have a more defined structure which must have, say, four levels. This will use the same base-table structure, but it provides a little more

**Fig 1** Table of animals' names

Code	Parent	Category
1	0	Bird
2	1	Penguin
3	2	King
4	2	Baby Blue
5	1	Sparrow
6	1	Crow
7	3	Fred
8	4	Jim
9	6	Sally
10	4	Jenny
11	6	Fred
12	3	Harry
13	3	Brian
14	5	Spadger
15	1	Starling
16	15	Simon
17	15	Silvia
18	2	Humbolt
19	2	Rock Hopper
20	19	Rocky
21	0	Dog
22	21	Terrier
23	22	Border Terrier
24	23	Brown
25	24	Shaun
26	18	Barry

scope when it comes to defining forms. Two possibilities, one of which makes use of combo boxes, are shown in Fig 4, but neither is entirely satisfactory.

If anyone wants to try improving them, be my guest. The work so far is in the MDB file called DBCMAR97 on our CD-ROM.

Finally, it is possible to count those categories which exist under any given one, using a GROUP BY query based on the query (in this case called Project) which displays the data in the manner seen in the top right of Fig 4. So

```
SELECT DISTINCTROW Project.Group,
Count(Project.Priority) AS
CountOfPriority
FROM Project
GROUP BY Project.Group
ORDER BY Count(Project.Priority)
DESC;
```

produces the answer table for the data shown:

Group	CountOfPriority
Prototyping	4
Design	4
Repair	2
Year 2000	1
Modification	1
Consultancy	1

#### Rounding out

On to a little housekeeping. Over the past few months I've published several problems with solutions and asked for comments and/or improvements. This month, we'll tidy these up. Although most of the problems originate in Access, the solutions are usually applicable in most RDBMSs.

Last November I published an algorithm for rounding in Access. In last month's column I used several answers (of differing efficiency!) from readers. James Talbut came back with a modified version of his: "I'm probably too late, but anyway, I think I have fixed that rounding code:

```
Function Round(dNumber As Double,
iNumDigits As Integer) As Double
Dim dFactor As Double
Dim dTemp As Double
dFactor = 10 ^ iNumDigits
dTemp = dNumber * dFactor
If 2 * dTemp = Int(2 * dTemp) And
Int(dTemp) <> dTemp Then dTemp =
dTemp + dTemp / 2 /
Abs(dTemp) Round = (dTemp \ 1) /
dFactor
End Function
```

"The solution was simple. I just had to stare at it until my eyes ached. My fudge

factor was either 1 or -1, depending on the sign of dTemp. Unfortunately, this was compounding the way in which Access alternates between rounding up or down. The solution was simply to fudge by 0.5 or -0.5. It's still hideous, for something as simple as a rounding function. But as a technical exercise I think it's kind of neat. And it's probably still quicker than converting to and from a string."

I agree. It's ridiculous that we have to write something like this just to round a number. Why doesn't Access simply include a rounding function? Microsoft, please note for future versions.

	1	2	3	4	5
Bird	Crow	Fred			
Bird	Crow	Sally			
Bird	Penguin	Baby Blue	Jenny		
Bird	Penguin	Baby Blue	Jim		
Bird	Penguin	Humbolt	Barry		
Bird	Penguin	King	Brian		
Bird	Penguin	King	Fred		
Bird	Penguin	King	Harry		
Bird	Penguin	Rock Hopper	Rocky		
Bird	Sparrow	Spadger			
Bird	Starling	Silvia			
Bird	Starling	Simon			
Dog	Terrier	Border Terrier	Brown	Shaun	

Fig 2 The data is stored in a reasonably economical manner

#### Case-sensitive joins

In last December's edition of my column, I published an email from Andrzej Glowinski who bemoaned the lack of case-sensitive joins in Access. I published a solution but he found it too slow, so I asked for other ideas as well as solutions in other RDBMSs. I was impressed with the intriguing answers, reproduced here.

On the subject of case sensitivity and the lost ninth data type, Stephen Parry writes: "There is another solution to Andrzej's problem with case sensitivity in joins, indices and sorting. It does, however, rely on a partially documented feature of access, data type 9 (aka BINARY). The BINARY field data type has been in Access

since its earliest beginnings but has been omitted from key places in the package and its documentation throughout. Various system table fields use it and it appears as the data type for certain field types, in tables attached from other database products.

"It behaves much like TEXT in all respects except any comparison operations (joins, indexing, sorting) operate on ASCII code value, i.e. case sensitively. Very useful; but how do you create one? 'Not easily' is the answer to that. BINARY does not appear as an entry in the field type list in table design view. If you try to use DB\_BINARY with the CreateField method in

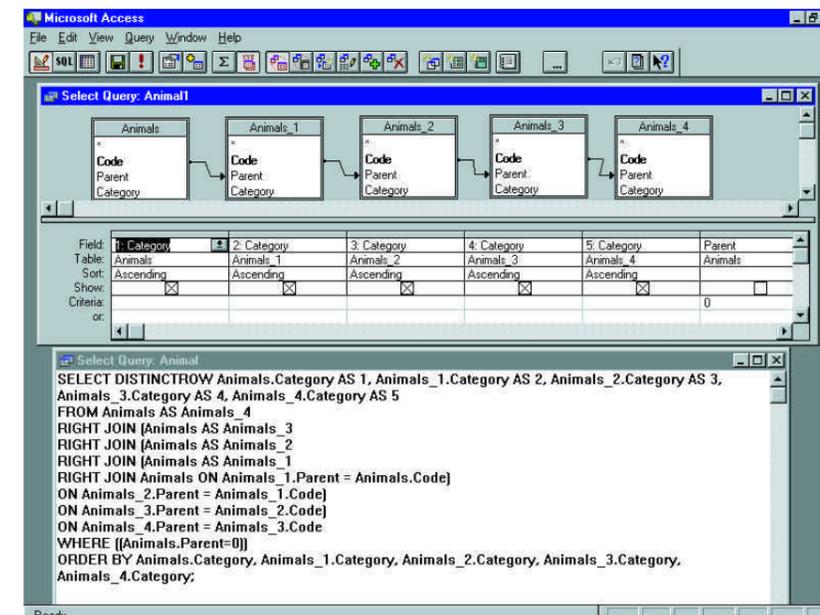


Fig 3 This SQL looks a bit complex, but it is just an expanded version of the code listing on p295

"To add a binary column to an existing table, I suspect an ALTER TABLE statement should work, although I have not tried it. This would allow you to add a new binary column to an existing table. You could then use code or query to copy the data from an existing column and then remove the original column. Thus, you have a long-winded means of changing a column's data type to binary.

"Of course, this all begs a fundamental question: why the (beep!) is something as useful as this so well hidden? I suspect that with sustained use, various 'holes' might be found in the functionality of this data type.

"Microsoft obviously intended a more complete but as yet unrealised solution to problems like this, but needed a quick solution to the Paradox and dBase connectivity problems, as well as a quick fix for some internals. For now, however, I will certainly continue to use it.

"By the way, I have noticed that in Access 7, nasty ol' Microsoft has encrypted its Wizards and other .MDA add-in files in this release. I have derived many useful programs from hacked versions of the Access 2.0 wizards, e.g. a database object directory comparison tool and a scripting tool based on the documentor, which just outputs a mammoth text file without generating a report or a temporary table. Hence, I was miffed to find that A7 has all its wizards compiled/encrypted in some way. Have you seen any way around this? I vaguely recall seeing an MSDN article on

securing add-in code and possibly one giving the unsecured contents of certain wizards, but I have not been able to find either again."

I knew of the binary data type but hadn't thought of it as a solution to this problem. As to decrypting the Access 7.0 wizards, has anyone else found a way?

On the subject of case sensitivity, the following came from Neil Howie: "A bell rang at the back of my mind, so I set about creating a similar query in Paradox 7 but taking advantage of the fact that if you have maintained secondary indexes, you can set a case-sensitive option.

"Adding an auto primary index to each table and setting up secondaries on the Names fields let me set up the join in the query window and display the required result with the greatest of ease.

"For further investigation I took a long doc, converted it to text and wrote a little routine to extract the words to create two files, 12,000 records long. I modified the second to capitalise the last letter of nine words out of ten at random, then pushed these into Paradox. Because of repetitions, it produced 15,000 matches in 18 seconds, so I guess on your enquirer's 200,000++ records it would still not be a practical proposition if he had to do it too often.

"However, what is now worrying me is that executing the same query in Delphi (using Paradox's SQL) takes almost 50 seconds. What am I doing wrong?"

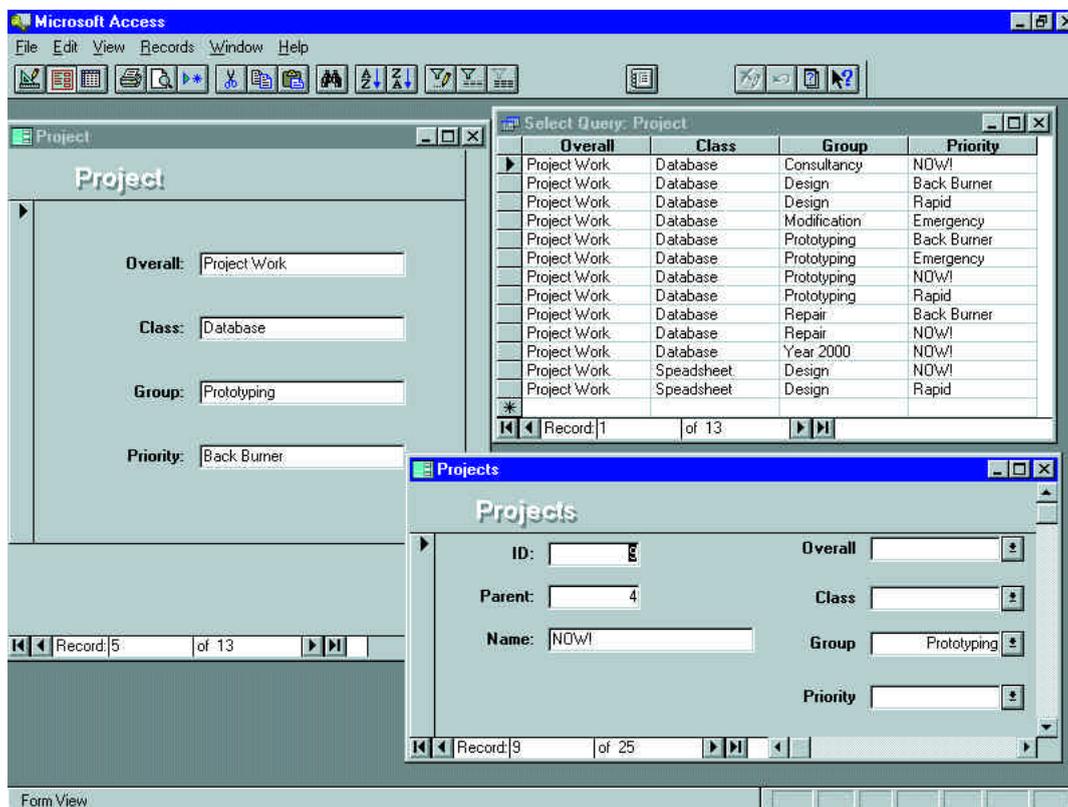


Fig 4 A defined structure with four levels provides more scope when defining forms

### Making your mark

Last month I published a problem from Andy, a teacher who was using Access 2.0 to store his pupils' marks. He wanted not only to store the marks that his students achieved in their tests, but also their positions in the various class groups.

I suggested there was a conflict here between relational theory and expediency, and asked for suggestions and restraint (in the hope of avoiding a holy war). I am delighted to report that (nearly) everyone kept their heads and the majority of answers were helpful rather than religious. Paul Mapstone's answer (below) was the most complete. It is also applicable to any situation in which a rank order is required.

"My advice on this is not to store derived data (i.e. the Position column) unless you have to, for performance reasons. This is a good example. It is fairly straightforward to calculate the Position column in standard SQL using a correlated subquery as follows:

```
SELECT A.[Pupil ID], A.[Test ID],
A.Score, A.Position,
(select Count(*) + 1
from [Test Scores] AS B
where B.[Test ID] = A.[Test ID]
and B.Score > A.Score ) AS Rank
FROM [Test Scores] AS A
WHERE A.Score is not null
ORDER BY Score DESC
```

"Column Rank in the above query should

correctly return the required Position. This works because the Position is equal to the number of people who have a better position + 1. Alternatively, you can use the SQL'92 outer join syntax (which Access seems to partially support) as follows:

```
SELECT A.[Pupil ID], A.[Test ID],
A.Score, A.Position,
Count(B.Score) + 1 AS
Rank
FROM [Test Scores] AS A LEFT JOIN
[Test Scores] AS B
ON A.[Test ID] = B.[Test
ID] and A.Score < B.Score
WHERE A.Score is not null
GROUP BY A.[Pupil ID], A.[Test
ID],
A.Score, A.Position
ORDER BY A.Score DESC
```

"We need the outer join, as the inner join will eliminate the top result. Either of the above queries could be used as the basis of any required reports etc, but if your teacher really wants to store the rank in the Position column (and suffer potential update anomalies), simply save one of the above queries with the name 'Rank query' and use it in the following UPDATE query:"

```
UPDATE [Test Scores]
SET Position = dlookup("Rank",
"Rank query",
"[Pupil ID]=" & [Pupil ID] &
" and [Test ID]="
```

```
& [Test ID])
WHERE Score is not null;
```

An excellent answer. Paul's first paragraph touches on the heart of the conflict. Storing derivable data usually has no benefits and several major disadvantages (such as causing potential update anomalies). However, storing such data can occasionally yield a major performance benefit. Of course, a real purist would never consider mere performance as a justification for breaking one of the central tenets of the relational model. Non-purists, on the other hand, wouldn't even hesitate. I sit uncomfortably on the fence, sticking to the purist line whenever possible and worrying every time expediency forces me to break what I know to be a sensible rule.

Paul's solutions are on this month's cover-mounted CD-ROM as PAUL.MDB and PAUL95.MDB. The first is an Access 2.0 version which crashes Access every time I try to run the update query. There cannot be anything fundamentally wrong with the solution because the problem does not occur in Access 95. So perhaps it is my machine?

**PCW** Contact

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column at [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk)



# Making the distinction

Mark Whitehorn on the peculiarities of DISTINCTROW in Access. Plus, if you're developing a large application, see how code and components can be versatile and non-specific.

I've written this before, but it bears repeating. This column has evolved to a state where it is almost exclusively about Access. When I write about general issues (normalisation, SQL, and so on) I try to keep it as generic as possible but the fact remains that, essentially, all the questions I receive (far more than appear on the page) refer explicitly to Access.

There are probably two reasons. Either everyone who reads this column is using Access, or people who are using another RDBMS read the column and think "Huh. That's all about Access. No point in writing in about anything else."

I'm happy to accept any suggestions for the future direction of topics covered and platforms used. So if you want changes please let me know and minority interests can be represented.

## Distinctly exact (or exactly DISTINCT)

In my series about SQL (*Hands On Workshop, PCW Oct'96-Jan'97*) I said that "the generic DISTINCT becomes DISTINCTROW in Microsoft's Access". By that I meant that if you build a query in Access using the GUI, then by default it will generate an SQL statement which starts:

```
SELECT DISTINCTROW...
```

In many cases this will return the same record set as an SQL statement which starts:

```
SELECT...
```

or even

```
SELECT DISTINCT...
```

The differences between SELECT and SELECT DISTINCT are explained in my workshop series on SQL, and since that was to do with generic SQL it didn't seem

to be the time or place to go into the peculiarities of DISTINCTROW. But this seems like an excellent place, so we will do so! In order to demonstrate this we need a couple of tables:

CLIENTS		
ClientID	Name	Location
1	Fred	Wellington
2	Sally	Wimpy
3	Jean	Halifax
4	Fred	Lancaster

ORDERS	
OrderNo	ClientID
1	1
2	4
3	4
4	4
5	3
6	3
7	3
8	1
9	1

Note that we have two clients called Fred and that Sally has yet to place an order with our company.

Both

```
SELECT Name
FROM CLIENTS;
```

and

```
SELECT DISTINCTROW Name
FROM CLIENTS;
```

return four records:

Name
Fred
Sally
Jean
Fred

whereas

```
SELECT DISTINCT Name
```

```
FROM CLIENTS;
```

returns

Name
Fred
Jean
Sally

DISTINCT (as explained in my SQL Workshop) forces SQL to remove the duplicates in the answer table.

In this case (and in all cases where a single table is queried) there is no difference between SELECT and SELECT DISTINCTROW. However, if we bring two tables into the query, the two forms can be distinguished.

```
SELECT Name
FROM CLIENTS INNER JOIN ORDERS
ON CLIENTS.ClientID =
ORDERS.ClientID;
```

returns:

Name
Fred
Fred
Fred
Fred
Jean
Jean
Jean
Fred
Fred

which is one record for each order in the ORDERS table because the join between the tables has essentially generated nine records, of which we have asked to see only the name of the customer. By contrast,

```
SELECT DISTINCT Name
FROM CLIENTS INNER JOIN ORDERS
ON CLIENTS.ClientID =
ORDERS.ClientID;
```

returns only two records:

Name  
Fred  
Jean

The SELECT statement on its own returns nine records, then the DISTINCT part removes the duplicates.

The "problem" with this answer table is that it implies that there are only two people placing orders, whereas we know it is three because there are two people called Fred. If we expand the SQL statement to include the location field,

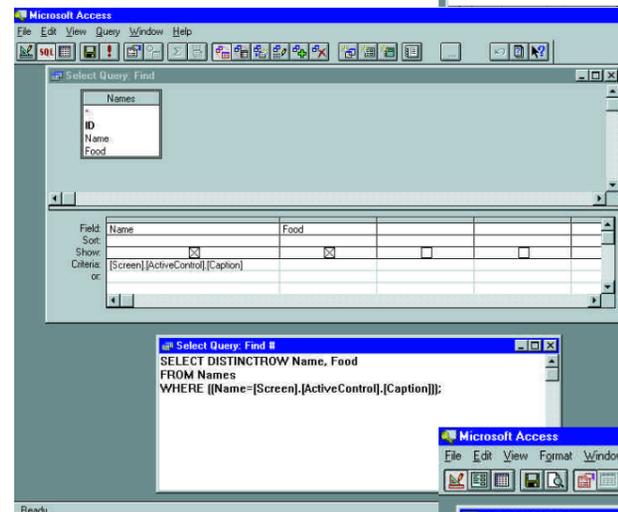


Fig 1 (above) The names used for the button-caption example

Fig 2 (left) The query called Find, as both GUI and SQL

all on our cover-mounted disc as DISTINCT.MDB.

I'll leave you with this brainteaser: If you include the primary key value from CUSTOMERS in all of these queries, does the difference between SELECT DISTINCT and SELECT DISTINCTROW disappear? As they say in all the best text books, explain your answer!

Hundreds and thousands Reader, Glenn Rowe, recently contacted me with an interesting database problem: "I'm developing an application which stores data about many countries. The data is stored in a table, one field of which stores the country code (GB for Great Britain... and so on). The interface will contain many buttons; one for each country. Each button will have the country code as its caption and, when pressed, will return the records

SELECT DISTINCT Name, Location  
FROM CLIENTS INNER JOIN ORDERS ON  
CLIENTS.ClientID = ORDERS.ClientID;  
then we get:

Name Location  
Fred Lancaster  
Fred Wellington  
Jean Halifax

DISTINCT is very literal; it returns unique records in the answer table, whether or not they come from unique records in the original table.

I used the term "problem" above, but of course this is only a problem if you don't know what DISTINCT is supposed to do. In fact, DISTINCT is doing exactly what it was designed to do.

My guess is that Microsoft felt that naïve users of a database might not appreciate this level of subtlety. If they saw a single name in an answer table, they would expect that it represented a single person. So the default in Access was set to DISTINCTROW which in this case, as you will by now have guessed, produces a separate record in the answer table for each unique customer who has placed an order.

Thus:

SELECT DISTINCTROW Name  
FROM CLIENTS INNER JOIN ORDERS  
ON CLIENTS.ClientID =  
ORDERS.ClientID;  
returns

Name  
Fred  
Jean  
Fred

By the way, these tables and queries are

for that country. (These are, of course, arranged hierarchically on many forms, not all on one!)

The problem, as far as I can see, is that I'm going to need hundreds/thousands of queries and as many snippets of code. Building it will be a nightmare, as will debugging and maintenance. Surely this process can be centralised in some way?

The question is a good general one

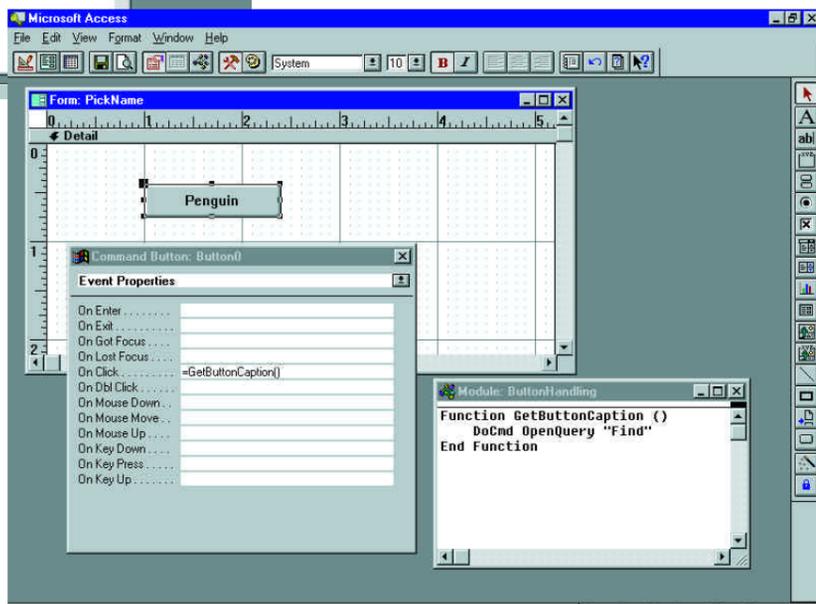


Fig 3 The form called PickName and the associated function which calls the query called Find

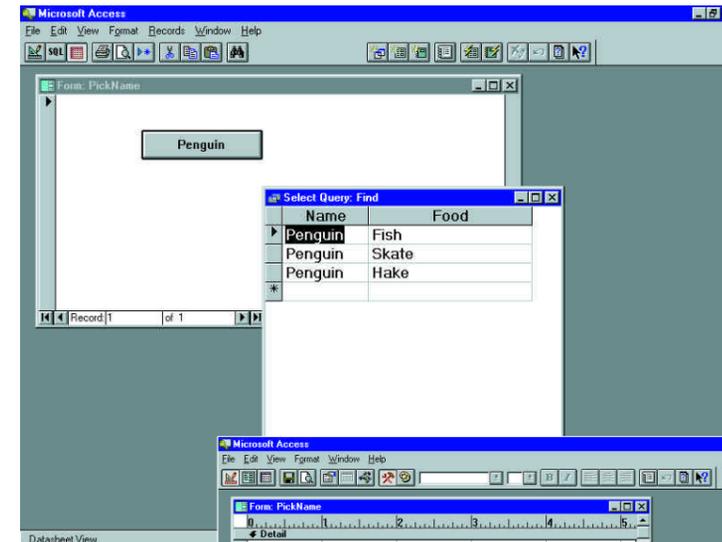


Fig 4 The answer table which appears when the button labelled Penguin is pressed

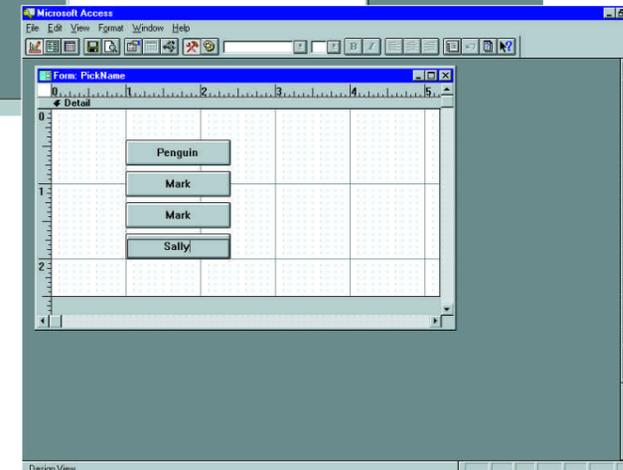


Fig 5 Cloning the button

about the way in which code and components can be versatile and non-specific. The obvious solution, at first, is to try to pass the button's caption to the query as a parameter.

As far as I know (and I stand to be corrected) this cannot be done. However, we can achieve exactly the same result by building a query which snatches the caption from the button which has just been pressed. In other words, you can't attach a piece of code to the OnClick event property of a button which says "Run a query and pass to it the caption of this button". Instead, you get the button to run a query and get the query to locate the caption of the button which has just been pressed.

In order to do this, we can make use of an object called "Screen" and one of its properties, "ActiveControl". These are defined in the manual thus: the Screen object refers to a particular Microsoft Access form, report, or control; the ActiveControl Property is used to refer to the control that has the focus. Or, to put that another way, Screen.ActiveControl always points to the control which has the focus.

Now consider a button on a form; if you press it, that button, by definition, has the focus. If the button runs a query which has Screen.ActiveControl.Caption as its criterion, then Property will return the caption of the button.

The beauty of this scheme is that when you design the query, you don't have to know which button is going to be pressed. You don't even need to know which form the button will be on. As long as the button calls up the query when it is pressed, the query will seek out the button, read its caption, and use that caption as a criterion.

To demonstrate this, I have used people rather than countries. This choice simply reflects the fact that I don't know enough country codes to fill even a sample table.

The table shown in the screenshot Fig 1 contains names of some individuals and the food they like to eat. The query called Find (Fig 2) has, in the criteria line:

[Screen].[ActiveControl].[Caption]

The form called PickName (Fig 3) has a single button with the caption "Penguin". The OnClick property of this button is set to =GetButtonCaption () which is the name of a function. The

## Round and round we go

In the November issue, I published the following correspondence from Paul le Glassick: "...Incidentally, to get around the lack of an =Round equivalent, I use the Format\$ function. This converts numbers to text, but rounds properly as we know it. With a representative sample of nearly 3,000 records, the following nesting of functions gave correct rounding when calculating VAT:

```
RoundNo = Val(Format$(CCur(Number), "0.00"))
```

where RoundNo is the result and Number is the number or calculation to be rounded."

Paul wanted to know if there was a better way. Simon Hawkins suggested:

```
RoundNo = int (( Number * 100) + 1) / 100
```

"Whether this is a faster method than using the Format function, described in the article, I am not sure. Also, the above may need altering to deal with negative amounts. Hope this is of some use."

This has the great advantage of elegance. However, when I tested it in the form

```
Function Simon (Number) As Integer
```

```
Simon = Int((Number * 100) + 1) / 100
```

```
End Function
```

it returned 1 from 1.49 (which seems right), but 2 from 1.4999 (which seems wrong), and then 2 from 2.49999 (which seems right), and even 2 from 2.499999999 (which is still right, but conflicts oddly with the result from 1.4999!). In other words, it is slightly inconsistent. Or maybe it's my 486 processor. In the form:

```
Function Simon2 (Number) As Double
```

```
Simon2 = Int((Number * 100) + 1) / 100
```

```
End Function
```

it returns 1.51 when given 1.5 (which seems wrong).

This reply came from James Talbut: "I don't like converting things to and from strings. There is an operator in VB that appears to round correctly, and that is the '\ operator (integer division). Making use of this in a function is simple:

```
Function Round(dNumber As Double, iNumDigits As Integer) As Double
```

```
Dim dFactor As Double
```

```
dFactor = 10 ^ iNumDigits
```

```
Round = ((dNumber * dFactor) \ 1) / dFactor
```

```
End Function
```

"This function is simple, quick, and produces the same results as the version using Format\$ that you published. Interestingly it gives a different result to that of the ROUND function in Excel. For some reason Round(2.15, 1) does not give 2.2 (as it would in Excel), it gives 2.1 as does the formula you published."

All of these suggestions work, up to a point, but none are perfect (see the form called "Rounding" in Fig 6). So the plot thickens. Anyone got any further thoughts?

By the way, just as we were going to press, Simon came back with: "OK. Classic case of using my memory instead of looking the code up (and testing it!). The function should read

```
TxtOutput =  
Int ((TxtInput *  
100) + .5) / 100
```

"This will round to two decimal places. Sorry about the earlier confusion."

Number for rounding	Paul	Simon (Int)	Simon (Double)	James
1.1	1.1	1	1.11	1.1
1.2	1.2	1	1.21	1.2
1.3	1.3	1	1.31	1.3
1.4	1.4	1	1.41	1.4
1.49	1.49	2	1.5	1.49
1.499	1.5	2	1.5	1.5
1.5	1.5	2	1.51	1.5
1.50001	1.5	2	1.51	1.5
2.49999	2.5	2	2.5	2.5
2.499999999	2.5	2	2.5	2.5
2.5	2.5	3	2.51	2.5
2.50000000001	2.5	3	2.51	2.5
2.6	2.6	3	2.61	2.6
2.7	2.7	3	2.71	2.7
3	3	3	3.01	3
3.49999	3.5	4	3.5	3.5
3.5	3.5	4	3.51	3.5
3.50001	3.5	4	3.51	3.5
4.5	4.5	5	4.51	4.5
4.5000001	4.5	5	4.51	4.5

**Fig 6** The rounding functions in operation. Note the purely fortuitous grouping of biblical names!

function, also shown in Fig 3, is composed of a single line which simply calls the query.

So, the steps are simple: when you press the button, the function is called; the function runs the query; the query looks for the active control (which is the button), captures its caption, and uses that caption as a criterion. The result is the answer table shown in Fig 4.

To keep it short and sweet, there is no error trapping in the demonstration code. It is also rather crude. For example, unless you close the query, pressing a different button won't re-query the table. However, all of this can be cured by expanding the code in the function from its current, rather minimalistic, one line.

Having answered that question, I realised that, essentially, the solution relied on the fact that the function could be called as soon as the button was pressed. This meant that the solution couldn't be applied to a situation where several selections needed to be made. I know there are manifold ways of handling multiple selections and that combo boxes are often the best solution, but I was intrigued to see if an economical solution (in terms of code and number of queries) could be found while retaining the button-caption-snatching idea.

The form MultipleSelections shows this in operation. A little of the elegant simplicity of the earlier solution is lost but much can be retained. Each button runs a function which places the button's caption in the text box above it. This text box is explicitly named in the function, so each button can be cloned vertically but each separate column of buttons is calling a separate function. Thus, if we made ten selections, we would need ten functions. However, by using Screen.ActiveForm (a close relative of Screen.ActiveControl) in the functions, a degree of flexibility has been retained so that forms which are cloned from this one should also work, as long as the same names are used for the text boxes.

The "Go" button fires a query called (rather imaginatively, I feel) "Find2". The query looks at the text boxes in the form, takes the values from there and runs the query. By using Screen.ActiveForm this query should be versatile enough to work with different forms.

### PCW Contacts

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column at [database@pcw.vnu.co.uk](mailto:database@pcw.vnu.co.uk)



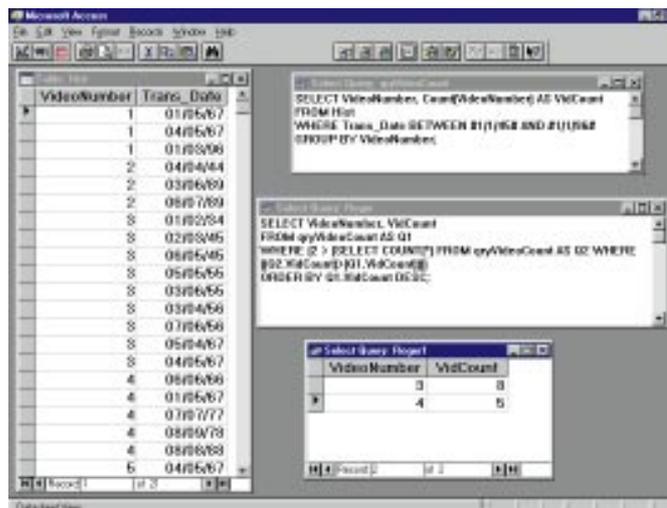
# The video **top ten**, again...

A recent query regarding video rentals has prompted some keen response from readers. Mark Whitehorn presents some of your solutions, plus some general tips and teasers.

In the October issue I published a couple of questions from readers. The first concerned a video rental company which wanted to find out the most popular videos rented out over a specific period, essentially the "Top Ten" videos for that period.

Several people sent in further suggestions, including the following from Roger Moran: "By 'standard' SQL, I assume you are referring to the SQL/92 standard, and explicitly excluding all the various vendor-specific extensions such as the 'TOP' command in Microsoft Access SQL."

No, but the misunderstanding is my fault for not being accurate in my choice of



Here is Roger's solution. The only change I have made is to include the dates into the first SQL statement so that the query runs without putting up dialogue boxes

## Dear Santa...

Last year, I asked you to make our current crop of RDBMSs a better match to the relational model, and also for more and better database design tools. I don't know about the rest of the kids, but frankly, I was just a little disappointed.

None of the RDBMSs has shown a marked improvement, and the tools situation is still dire.

Sigh. Perhaps I'm asking too much. After all, the vendors don't seem to be interested in what I asked for either, so you must face an uphill struggle in trying to persuade them. This year, I'll try for really achievable requests, but first, a bit of background. Many database professionals are having their serious work disrupted by trivial requests from the management about putting data onto this "web" thing. Now, we all know that the web is just a passing fad. Just like CB radio, once the fuss has died down, I'm sure most people will go back to the telephone. But while it is fashionable, we have to look keen. So, what we need are sensible tools for handling data from databases on the web.

Borland has already made a good start with IntraBuilder, and Microsoft is following with the improved web-publishing abilities of the Office products, but there is still a huge way to go. Neither of the products addresses the problems of signalling between the browser and the database to allow proper transaction control. This is not to say that transaction control is impossible, just that it is currently all left to the programmer. There are other problems as well. With present technology, if the browser sends a query which returns, say, 140 records, those are usually buffered at the web server and sent to the browser in chunks of, say, ten. This is crazy! Every time the user wants to see another ten, there is a pause as the server is contacted and the next ten are sent. This problem isn't going to be solved by a single product; it needs to go right to the standards that these technologies use. If you could wangle your way onto the appropriate standards committee to address this one, we would all be eternally grateful.

By the way, I'd still really like a mouse mat with a penguin on it. I mentioned this last year, but it must have got lost in the sleigh. Perhaps you could check under the seats?



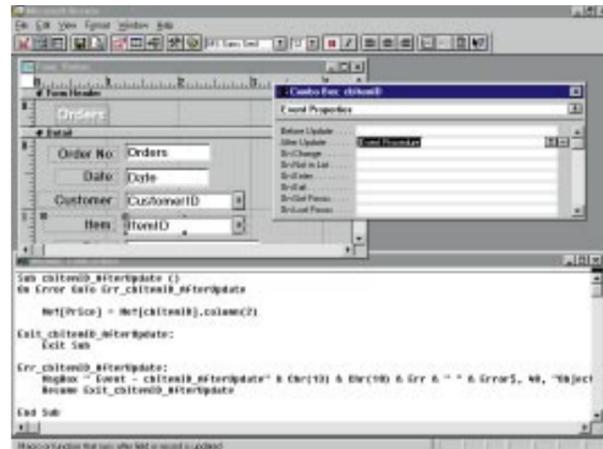
words. Eamonn Mulvihill's original request was for a solution which worked with Borland's Local-SQL which he said didn't support nested SELECT statements. I should have said something like "lowest-common denominator SQL" rather than "standard". Roger Moran continues: "Since MS Access seems to be your first choice when describing problems in your column, I will use MS Access SQL syntax in the following, although the SQL is easily ported to any ANSI-compliant SQL/92 dialect.

"We start with an aggregate query to provide us with a count of transactions per video within a given date range. This is trivial:

```
SELECT VideoNumber,
       Count(VideoNumber) AS VidCount
FROM Hist
GROUPBY VideoNumber
WHERE Trans_Date BETWEEN
[StartDate] AND [EndDate];"
```

I know it will sound picky, but although the sentiment behind this is quite clear, I think the syntax for Access actually needs

(right) Tony Wall's solution, showing the form and the code behind the combo boxes



(below) Geoff's form, which also totals the orders



number of records in qryVideoCount (Q2) whose VidCount is greater than that of the 'current' record in Q1. If that number is less than ten, we know the 'current' record should be included in our 'Top Ten' list. We apply the test to each record in Q1 and the result will be the 10 desired records.

table has so little data in it.

Alistair Logie sent in this solution for FoxPro users faced with a similar problem: "I read with interest your PCW Databases column about SQL selections of the top ten rented videos. One solution in FoxPro would be to use two SQL statements and a Cursor:

```
SELECT VideoNumber,
COUNT(VideoNumber) AS VidCount ;
FROM hist ;
INTO CURSOR mytable ;
GROUP BY VideoNumber ;
ORDER BY VidCount DESC

SELECT VideoNumber, VidCount ;
FROM mytable ;
WHERE RECNO() < 11
```

"I don't know if this is usable on other databases. It does rely on being able to use the absolute record number, RECNO(), as a selection field. The attraction of using a Cursor as a temporary file is that FoxPro creates and erases it itself."

The second question in the October issue covered a more general problem. Suppose you have a set of Customers and Items for sale. Given that each item has a specific, unique price, it is clear that the price fits into the ITEMS table. If each Customer has a unique price for each item, it is equally clear that the price needs to go into a separate table which lists each customer, item and the unique price for that key value.

What if the situation is somewhere in between, where most prices are standard but there are some exceptions? I produced a solution, but asked for suggestions, improvements or comments. The first comes from Tony Wall, and his database is on the CD as TONYWALL.MDB.

"I have enclosed an alternative solution to the Customer & Prices problem from the October issue. Whether it will work efficiently with the 400 customers and 300+ products is another matter.

"In the first instance I took your Default List and Exception List and used these to create a union query as follows:

```
SELECT *
FROM [PriceList] as PL
WHERE NOT EXISTS
(SELECT *
FROM [ExceptionList] as EL
WHERE PL.[CustomerID] =
EL.[CustomerID] and PL.[ItemID]
=EL.[ItemID]);
UNION SELECT *
FROM [ExceptionList];
```

"The aforementioned query works, and is written in standard SQL, but it cannot be considered very efficient, since for each record in qryVideoCount (between 1 and 12,000 in Eamonn Mulvihill's situation), we must execute the sub-query which determines whether the record is to be included in the 'Top Ten' list. In a real world situation, it would probably be quicker to write some procedural code to pull out the ten records from qryVideoCount.

"Hope this information is useful." Yes, it certainly is. It won't do as a solution for Eamonn's problem because of the nested SELECT, and I agree that it is likely to have speed problems on a large data set. However, Roger is quite right that his is a solution comfortably within SQL-92 and I have great admiration for its elegance!

You can clearly tweak the second SQL statement very easily to alter the number of records that it returns, simply by changing the numeric value in the line:

```
WHERE (10 > (SELECT COUNT(*) FROM
qryVideoCount AS Q2 WHERE
```

I have changed it to two in the sample database (ROGER.MDB), because the HIST

```
to read as:
SELECT VideoNumber,
Count(VideoNumber) AS VidCount
FROM Hist
WHERE Trans_Date BETWEEN StartDate
AND EndDate
GROUP BY VideoNumber;
```

"In MS Access, we can use this query directly as the datasource for the second stage. In other products which do not support daisy-chaining queries in this fashion, it is a simple matter to convert the above query into an append query which populates a temporary working table, and then use that as the datasource instead.

```
"Let's assume we have saved the above query with the name [qryVideoCount]. The 'Top Ten' query will look as follows:
SELECT VideoNumber, VidCount
FROM qryVideoCount AS Q1
WHERE (10 > (SELECT COUNT(*) FROM
qryVideoCount AS Q2 WHERE
((Q2.VidCount)>(Q1.VidCount))))
ORDER BY Q1.VidCount DESC;
```

"Here's a basic description of what's going on in the above query. For each record in qryVideoCount (Q1), we find the

"This is used as the basis for the combo box data source on the orders form for retrieving the correct price."

Tony then sent another mail message, saying that the performance was terrible with a larger data set. His suggested improvement is well worth looking at, and is in TONY2.MDB. Geoff Wyss sent in another; see GEOFWYSS.MDB.

"In the October 96 issue of PCW, you posed a problem about merging prices from two sources to make up an Order from a default price list and an exceptions list. The attached database file shows two possible solutions which might be of interest, one using a Union query, the other using an IIF statement. Click on form Orders2. There is a ReadMe button on the form which gives an explanation."

For this month's problem, I have held off producing a solution myself. This is mainly due to cowardice, since this problem brings us once again into the holy wars of "strict relational" vs. expediency. I tend towards a purist view, but in this case I can see reasons for a more heretical solution. Before committing myself to be burned at the stake, what do you think? Let's try not to get too wound up about this one. It is Christmas, after all.

"I am a teacher, and recently used Microsoft Access 2.0 to produce a database for recording and reporting on pupils' performance in tests. To keep the database as flexible as possible in use, I store test results in a single table with four fields: each record in the Scores table includes an identifier for the pupil, an identifier for the test, the percentage score achieved by the pupil in the test, and the position of the pupil out of his/her teaching group. Records are added to the table as data are entered by the class teacher. (The pupil and test identifiers together make up the primary key, as each pupil appears in the table many times, as does each test, but each pupil can only take a given test once.)

"It is useful to have the computer calculate the position of the pupil automatically, once the data for a full teaching group has been entered, but I have not found a function in Access to do this easily. The pupils' teaching groups are recorded in a separate table. It is straightforward to produce a query listing for a particular test, the pupils in one teaching group and their percentage scores, and this query can be sorted appropriately. But then to work out the

### Teasers

A couple of teasers from Jeff Jenkins:

1. Strange database design  
A database I was working on stored amounts as positive numbers and had a separate field to indicate the sign of the number. This field had values of one or two to indicate positive or negative (don't ask why, I didn't design the database). The problem was how to extract the amounts as signed numbers. The solution I came up with was a mathematical one, selecting "amount\*(3-(sign\*2))". This works as "(3-(sign\*2))" evaluates to 1 or -1. A bit of a rigmarole just to get a signed number out.

2. Summing only the positives  
I needed to provide a total of all the positive numbers in a particular field in the database (e.g. all the accounts with a credit balance). I couldn't find an sql command to do this, so I resorted to another mathematical solution: "SUM( ( amount+ABS( amount ) ) / 2 )" Adding the absolute (unsigned) value of a number to itself comes out to zero for negative numbers, and dividing by two negates the effect of adding a number to itself, so the sum works. You can also sum up all the negative numbers with:

```
"SUM( amount - ( ABS( amount ) ) / 2 )"
Incidentally, the first one also works as
"SUM( ( ABS( amount ) +amount ) / 2 )"
but the negative version
"SUM( ( ABS( amount ) - amount ) / 2 )"
gives a positive total of all the negative numbers.
```

position of the pupils within the group and store this in the appropriate field, I have resorted to a complex sequence of operations involving creating and deleting temporary tables.

"Using this method to calculate pupils' positions slows down the program significantly and, because of the use of temporary tables, has led to problems when two teachers using the network attempt to enter test scores at the same time. Is there a direct way to fill in the 'position' field in the Scores table? I have tried different techniques, but have not found a wholly satisfactory solution." Andy

Andy's test database (not containing real data about real children) is on the CD as ANDY.MDB.

### \*PCW Contacts

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. He's on [m.whitehorn@dundee.ac.uk](mailto:m.whitehorn@dundee.ac.uk)



# Modelling job

Where did you put that data? Where did you put that file? Mark Whitehorn looks at the pros and cons of database modelling.

**L**ast month, we started to look at three "database models" you can employ:

1. Everything on a stand-alone PC;
2. PC front end with data on file server;
3. Client-server using a database server as the back end.

In this context it is useful to think of a database as having four different parts:

1. User interface section;
2. Data processing engine;
3. Conflict resolution section (to deal with conflicts introduced by multiple users accessing the data at the same time);
4. The data itself.

We looked, last month, at the pros and cons of the first model. As I wrote at the time, I will give size estimates where appropriate but please don't take them as gospel. (See last month's *Hands On Databases* for other qualifications).

## PC front end with data on file server

If you have a need for multiple users to access the same data, then it is not beyond the realms of possibility that you already have a network. Given a network, you have the option of moving to our second database model.

In this model, much is left the same as before. You would still run Access, Paradox, or the RDBMS of your choice on your PC so the data processing engine bit of the database stays there.

Only two things change. One is that the data files are moved to a file server and the second is that the individual RDBMSs running on the individual PCs need to communicate with each other. They need to do this in order to resolve the multitude of

potential conflicts which suddenly arise when more than one person accesses the same data simultaneously.

Now is not the time to go into all such conflicts but consider a simple example. You and I both work for the same company and we are trying to update the company's customer records. I open up the record for A Smith to increase his credit rating from £2,000 to £3,000. While I am doing so, you delete his record. What happens to his record when I finish editing it and send it back to the file server?

The answer to this question depends on the RDBMS you are using. Access, for instance, maintains a lock file in the same

directory as the data file and this is used to store information about who is doing what at a particular time. Thus, if I had opened the record to update it before you tried to delete it you would receive a message saying that the record was in use by "Mark" and that you wouldn't be able to update the record until I had finished with it. Other RDBMSs use other mechanisms, some less efficient than others, for dealing with these potential conflicts.

The important point is that with this database model the control of the user interface and the data processing engine remain on the PC while the data and conflict resolution are moved to the file

## Delphi Programming for Dummies

by Neil J. Rubenking

Delphi is a multi-purpose programming tool and is commonly used for the generation of user interfaces to databases. This book is one of a series and if you've seen any of the "for Dummies" books you'll already have a good idea of what this one is like.

It has lots of diagrams, some very silly cartoons and lots of icons in the margins to identify Technical Stuff, Warnings, Tips and so on. It is also written in a style which doesn't display the correct reverence for what is a serious programming tool — which is just fine by me. It's not that I don't respect Delphi but books without a sense of humour can be really tedious. This one is fun.

The database section is pretty skimpy, comprising a mere 22 of the 376 pages. It would be impossible to recommend this book to a database-naïve user who wanted to learn how to create databases in Delphi but that description doesn't fit most of the readers of this column. If you are happy that you understand the database end of things, this is an excellent introduction to Delphi.

One of the examples used in the book is a function called Hailstone. Given a seed number, this function generates a series of numbers which bounce up and down "like hailstones in storm winds" before converging on the number 1. Apparently there is no way of either predicting how long a given seed number will take to reach 1, or proving that all numbers will eventually reach 1. I was so intrigued that I modified it to run in Access and it is included in the main part of my column for your amusement (see *Hailstones*).

By the way, I saw one of this series in a bookshop in France: "*Windows pour les Nuls*"; a gift of a translation for any of the series connected with databases!

■ *Delphi Programming for Dummies*. IDG Books £18.99 (ISBN 1-56884-200-7).

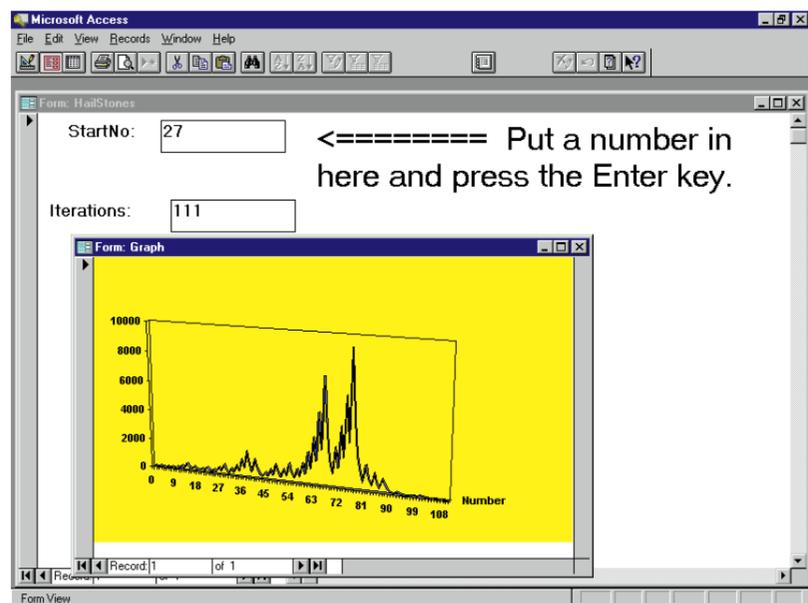


Fig 1 The "Hailstones" algorithm in action

server.

The big advantage of this model is that it provides multi-user access to the same data at relatively low cost. The big disadvantage is that the model is inefficient in two main ways. Firstly, it tends to load the network, soaking up bandwidth like a sponge. Remember that the data is at one end of the wire and the processing is at the other. Every time you query the data, it has to be moved to the client since that is where it is crunched. In a badly designed system this can mean that every query against a 100Mb table requires the entire table to be shipped to the client. Intelligent indexing can reduce this considerably (since the indexes can be shipped for searching and only the relevant records sent out to the client) but in practice, the effectiveness of this depends on the particular RDBMS.

Secondly, the processing is at the client end, so each client needs sufficient resources to cope with the data. If you decide that an increase in the database size warrants an increase in memory of 16Mb, you will need to add that to all the clients: given ten clients, that's 160Mb.

These restrictions mean that the number of simultaneous clients and the size of the data are relatively restricted. Think in terms of ten clients and 1Gb.

**Client server model**

This model is simply a modification of the previous one. The user interface stays on the PC: the data, the processing and the

conflict resolution moves to a server. This is typically not a file server but a server dedicated to running applications like RDBMSs, hence they are generically known as application servers. Machines which are dedicated to running RDBMSs are often called database servers or SQL servers.

But wait. Why go to the expense of a dedicated database server when you have that nice NetWare 3.x file server already in place? Can't you run a database engine on that as well? The answer is that you can, but you probably don't want to. The reason lies in the NOS (network operating system).

NOSs like NetWare 3.x are optimised for File+Print. Application servers run a NOS which is optimised for running multi-user applications. This doesn't mean that you can't do it, just that performance will be compromised, so a dedicated database server is generally considered to be better. If you really want to run one server for both application and File+Print, then NT is probably better than NetWare, all things considered.

The Client-Server model is typically *not* limited by bandwidth. Since the processing and data are now snuggled together in one place, queries no longer mean that masses of data have to move across the network. Instead when the GUI, running on the client, is used to construct a question only an SQL description of that query is shipped across the network to the server.

This SQL will typically be a very short ASCII string. The database engine on the

server processes the query and simply sends the answer (rather than the entire table) to the client. Conflict resolution is also handled centrally, with associated benefits in terms of speed and sophistication. Centralising the processing means that the whole system is easier and usually cheaper to update. If the database slows down you can throw hardware (memory and processors) at the server. You don't have to add it to the clients because they are simply handling the user interface.

On the ticklish subject of size, this model will typically support as many clients and as much data as you can afford. To put that another way, consider this model seriously if you think you will need more than ten clients or >1Gb of data in the foreseeable future.

**Hailstones**

The hailstones algorithm is very simple and most easily expressed in pseudo code.

```
Start with a seed number
Repeat
    If the number is even,
    divide by two
    Else multiply by 3 and add 1
Until number = 1
```

I can offer no justification for implementing this in Access except that it is fun and that the behaviour is really wacky. Give it a number like 26 and it takes a mere ten iterations to get down to one. But 27 takes a monstrous 111 iterations, and 28 a more reasonable 18. A copy of this is in the MDB file on our cover-mounted disk this month. (See also, the *Delphi Programming for Dummies* panel on p285).

**More on meters**

The meter problem (which was discussed earlier this year) resulted in my publishing, on our PCW cover disk, a database with several of the solutions and a kit which allowed you to test any other solutions against the published ones.

I asked only for solutions which were genuinely faster. None have arrived but Paul Bloomfield, a supplier of one of the original solutions, contacted me to say that he had got very different answers from those published. It turned out that he hadn't got the PCW cover disk, so had built his own database for test purposes. I sent him the one I had used and this is his reply:

"Our speed differences are down to one factor: the data. My test database had 2,000 records from 200 meters (i.e. average ten readings per meter), whereas I see yours had

only three meters and so 666 per meter, so the query is doing 66 times the work!

Perhaps there are lessons to be learnt here in problems of scaling. There is no such thing as good SQL per se, the query must be written to fit the data, or the data structured to fit the query!"

I agree. It is un-nerving to realise how many factors we need to consider when working with databases. Still, it keeps us all in employment!

**Case-sensitive joins**

Andrzej Glowinski writes: "I have recently started using MS ACCESS (Win 3.1) at work, developing applications in the areas of large medical (clinical) information resources. We also use an Oracle server and my systems work fine using this but as

soon as I try to build stand-alone versions all hell breaks loose. This is because some ACCESS designer/implementer, in their wisdom, has forced all internal joins to be case insensitive — and there is NO MECHANISM for altering this behaviour!"

I replied to this as follows: "I hate to be contentious, but Access does provide a mechanism which allows joins to be case sensitive.

Assume that we have two tables: NAMES and ORDERS. Each table has a field called NAME which contains case sensitive data. We thus have two fields:

```
NAMES.Name
ORDERS.Name
```

which need to be joined.

Create a query containing both tables, join them within the query (not in the

relationships editor) and run the query. Precisely as you suggest, the data in the tables is joined in a case insensitive manner.

Now add a field to the query:

```
CaseCompare: StrComp([NAMES].[Name], [ORDERS].[Name], 0)
```

This uses the function StrComp which can be rendered case sensitive by setting the third argument to be 0.

The function returns 0 if the strings match in case as well as letter order, so if you set the Criteria for that field to be 0 and re-run the query, the join is now case sensitive. The SQL version reads as:

```
SELECT DISTINCTROW NAMES.Name,
ORDERS.Name, NAMES.Foo, ORDERS.
Baa, StrComp([NAMES].[Name],
[ORDERS].[Name],0) AS CaseCompare
FROM NAMES INNER JOIN ORDERS ON
NAMES.Name = ORDERS.Name
WHERE ((StrComp([NAMES].[Name],
[ORDERS].[Name],0)=0));
```

I am aware that this doesn't allow you to sort case-sensitive material, although it would almost certainly be possible to write a function to do this. Rather more interestingly, it doesn't allow you to apply referential integrity or set a primary key on material which is case-sensitive (since Access treats 'PENGUIN' and 'Penguin' as equivalent in a primary key). However, it does at least allow you to perform the join." (A copy of this is available in the MDB file on the PCW cover-mounted disk).

To which I received Andrew's reply: "Thanks for your message about ACCESS. We have explored the mechanism you propose quite extensively, both on local (native) ACCESS tables and on attached ones, primarily ORACLE. The performance hit is just unacceptable — I'm talking of joins across tables with 200,000++ items in them, so the total loss of indexing (which, effectively, is what you get) is very significant. You can mess around with the order of execution of the query but it is just too unpredictable to be of much use."

I can't argue with any of this. Indeed, depending on how large the records are, this database may be getting to the size where a move to another model is inevitable. One question is, do any other PC-based RDBMSs have a better way of supporting

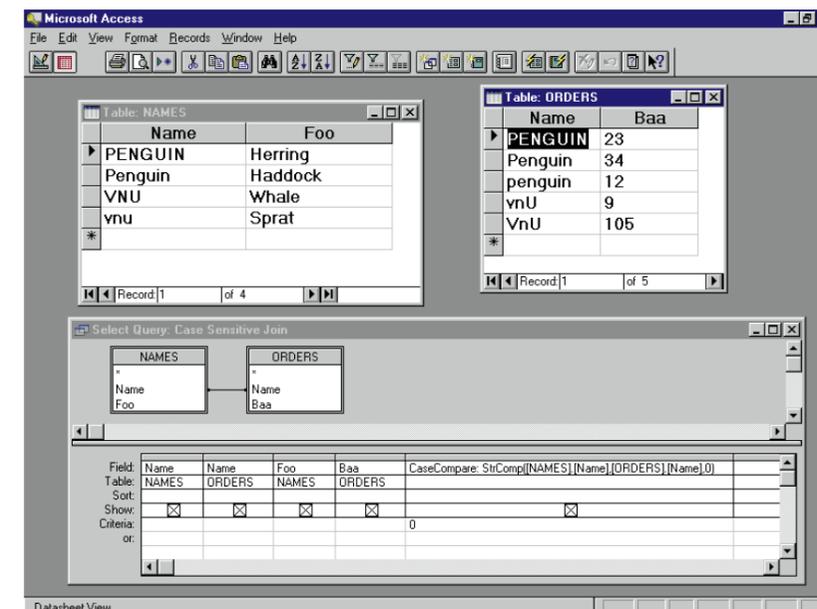


Fig 2 (above) The tables used to demonstrate the case sensitive join, and the query itself

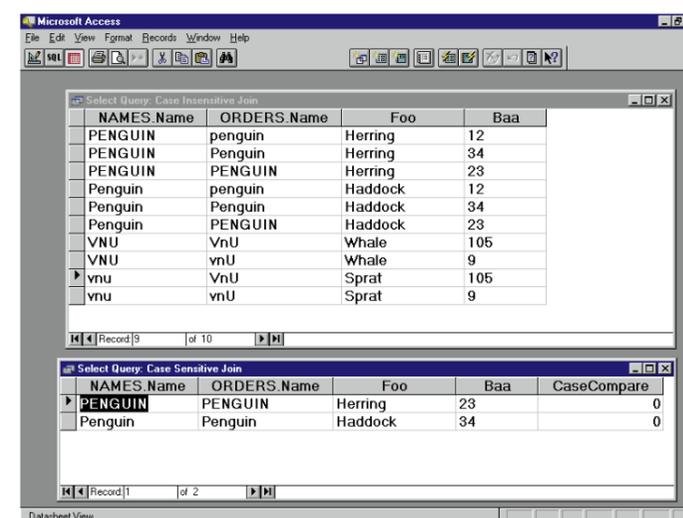


Fig 3 (left) The results of a case insensitive (top) and case sensitive join (bottom)

**PCW Contacts**  
 Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. He's on [m.whitehorn@dundee.ac.uk](mailto:m.whitehorn@dundee.ac.uk)



# One is not enough

A reader worries that as the work gets too big for his company's current database system to handle, which way now? Mark Whitehorn is on hand to dispense this, and other, advice.

**I**would like a discussion of the relative merits of a client/server database and a networked database application in which there is a server. When working on database applications, I am always considering small standalone systems for one PC. Most of our work so far has been in-house and thus manageable. As time goes by and we are involved in increasingly big projects, I am beginning to worry that the one-PC approach will fall on its face.

I believe I have two options:  
 1). Put my data on a big server running something like NT and run multiple copies of my Paradox application all pointing at the tables on the server. Then let Paradox handle the problems. I realise I'd have to be careful about record locking and the like.  
 2). Up-size my data to a client/server application using something like Interbase running on a server (UNIX?) and do it all through SQL, although my knowledge here is very hazy."

Alasdair Macdonald

I received this email and it seems a broad enough question, to warrant some expansion. After all, it is one of the biggest decisions that you are likely to make, and is an area where mistakes are both common and expensive. There are essentially four database models you can employ:  
 1). Everything on a standalone PC.  
 2). PC front-end — data on file server.  
 3). Client server using a database server as the back end.  
 4). Mainframe

The fourth seems inappropriate for discussion in a PC magazine, so we'll ignore it. What I can do here is to outline the

strengths and weaknesses of each of the other three, together with approximate performance and size estimates.

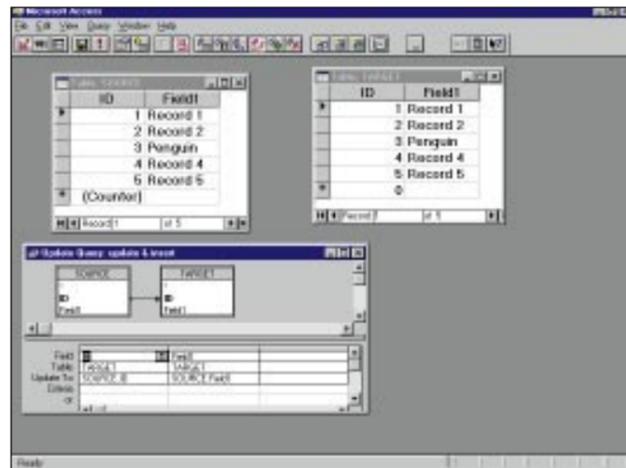
This is a complex area. The decision to choose one of these alternatives will be based upon the interaction of many factors, including response time required, number of users, file size, data size, available resources (including hardware, software and money) and type of data access

(Fig 1, top right) The update query described by Paul [see page 276]. The two tables are the same

(Fig 2, right) Here I have made changes to both of the tables, but have yet to run the query. (I have also altered the view of the query to show it in SQL)

required (read only, read write).

These factors interact in complex ways. Suppose that your system definition makes multi-user access to the data essential: you can instantly rule out a standalone PC. If the number of users is guaranteed to be small (say, three), then on a given hardware



(Fig 3) The state of the tables after the query has been run

reasonably rapidly. If the data is constantly updated, the indices will slow down the updating, and yet removing them will slow down the querying! In a

platform you could allow those users to access a greater volume of data than if there were 50 of them. If the number of users did then increase, the system might still work, but the response time would drop.

To make matters worse, the interactions between these factors are often non-linear: for example, doubling the number of users on a given system might have little impact on response time. Doubling it again might bring the same system to its knees.

It's easier not to give any actual figures, but this is likely to leave you gnashing your teeth and wondering "What exactly counts as 'lots'? Three? Twenty five? Five hundred?" On the other hand, if I do quote hard figures, like saying that you shouldn't consider using a standalone PC for more than 1Gb of data, there'll be someone out there happily using a 200MHz P6 with 1024Mb RAM to access 1.5 Gb.

I will quote figures because it seems far more useful to do so, but please just regard them as general figures from which to start discussions. Please don't take them as gospel, and please don't build your entire database strategy around them alone.

### Work alone on a standalone

The simplest database model is to install everything on a standalone PC. You use an RDBMS like Access, Paradox, dBase and FoxPro to manipulate the data.

Only one person can use it at a time, and I wouldn't use this sort of system for more than about 1Gb. Factors which affect this figure are the hardware (more memory equals larger data files) and the manner in which the data is used. For example, if it's rarely updated, then it can be heavily indexed and queries should run against it

nutshell, if the data is rarely updated, heavily indexed, and you have very impressive hardware, you can go above this limit. With a 286 with 640Kb, don't even think about it.

The major advantage of this sort of system is that it is cheap, and easy to manage. A database can be thought of as four different parts:

- 1). User interface section.
- 2). Data processing engine.
- 3). Conflict resolution section (to deal with conflicts introduced by multiple users accessing the data at the same time).
- 4). The data itself.

In a standalone PC-based database there is only one user, so the conflict resolution section isn't required and the other three appear as a single, seamless entity to the user. Simple. In fact, it is so simple, why would you ever want to go to anything more complex?

One of the major reasons for moving to a more complex database model is that this one cannot handle multiple users. For one thing, there is only one keyboard, so we can expect fist-fights if we try for multi-user. Also, this model doesn't allow for conflicts between the requirements of different users to be resolved. If you need more than one user to access the data at the same time, it is time to split up the components described above and partition them between different machines. This leads us to the second database model, which we'll cover in the next issue.

### Target, aim, fire!

"Suppose we have one table, called Source, containing new records, and another table called Target. We wish to insert into Target the records in Source that are not already in Target, and update those records which are already in both tables to

## Database Systems

by Paul Beynon-Davies

This book looked promising. Many of the subjects covered are of interest, and the style, while a little formal for my tastes, is perfectly respectable. However, reading it in more depth reveals a series of unnerving flaws. For a start, the book is heavily cross-referenced. I like cross-references, but I do like the references to point to the correct place. Far too many here do not. Exactly half of the cross-references in chapter six, for example, are incorrect. With a failure rate like that, they are too frustrating to use. And it is not only the cross-references that are flawed.

The same chapter covers SQL and the author demonstrates retrieval, ordering and grouping for which he uses a base table, eight SQL samples and eight answer tables. Amazingly, three out of eight answer tables are incorrect, a state of affairs likely to induce severe confusion in a novice reader.

This book is aimed at students, but cannot be recommended to them or anyone else, which is sad because in many ways it's a fine book. It is simply crying out for a careful revision.

Hopefully the next edition will be improved, but steer clear of this one.

MacMillan Press ISBN 0-333-63667-8. £19.99

the values of the records as in Source.

Instinctively, programmers will achieve this by two queries:

- 1). Update
- 2). Insert

In Access it can be achieved with a single Update query with a LEFT JOIN.

```
UPDATE SOURCE
LEFT JOIN TARGET ON
SOURCE.ID = TARGET.ID
SET TARGET.ID = SOURCE.ID,
TARGET.Field1 = SOURCE.Field1;
```

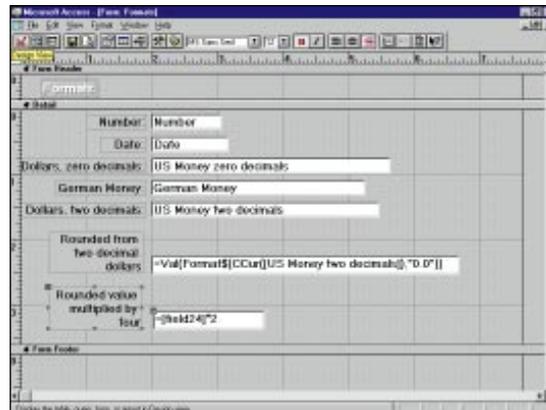
Note that both tables have fields called ID and Field1, and both are of the same data dimensions in each table. Remember that if ID is a counter field in Source, it must be a long integer in Target. This works, since Access matches those records in Source which are not in Target with a Null Record in Target which can then be updated. From an SQL point of view, this technique may not

work in other DBMS but is jolly useful in Access 2.0 and Access 7.0.

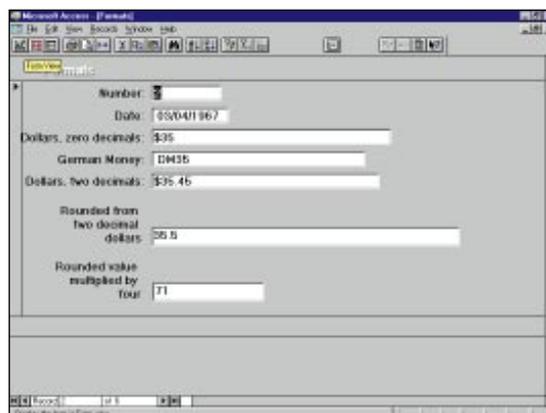
I enclose a sample Access 2.0 database [on the cover disk as pbdemo.mdb] which demonstrates this. To see it in action, play around with appending/changing records in Source and then run the query to see the effects in Target. The uses of this technique are numerous, and variations on the query to supply selection criteria make it powerful."

Peter Blackburn

It is worth noting that this query will not delete records from Target which have been



(Fig 4, right top) Here I have used some of the formats suggested by Paul, and his rounding mechanism



(Fig 5, right bottom) And here is how the data appears. Notice that the value shown by the rounding mechanism is still a number; it can be mathematically manipulated as shown in the lowest "text" box.

This form is also in the sample database on the cover disk

deleted from Source. This is not a criticism. If it did delete those records, it would effectively be replacing Target with a copy of Source. It is simply a characteristic of this type of query. I can't help feeling that this might help in the solution of last month's problem concerning Mark Squire's problem about Customers and Items.

### Currency codes: help wanted

"Your July column covered the formatting options for dates in Access.

It is often overlooked that the Format property in Table, Query, Form and Report design is not restricted to just those formats on the list. The variety of codes available is the same as those used in Excel's Format, Cells, Number dialogue box.

Thus, a code of dd/mm/yyyy will show dates with century, \$#,##0 will show amounts with dollar signs, "DM"#,##0 will show amounts in deutschmarks. This latter is especially useful since Access picks up the default currency format from Control Panel. It can then be overridden to show different currencies on a single Form, etc.

The ability to override the currency format has an additional benefit. Since currency fields are held to a fixed four decimal places internally, they are likely to be more precise than single or double numbers which can give puzzling rounding errors. As you know, Access has no exact equivalent to the =Round function in Excel. Currency fields used for other than currency can be formatted as #,##0, or #,##0.00 if two decimal places are required.

Incidentally, to get round the lack of an =Round equivalent, I use the Format\$ function. This converts numbers to text, but rounds properly as we know it. With a representative sample of nearly 3,000 records, the following nesting of functions gave correct rounding when calculating VAT:

```
RoundNo = Val(Format$(CCur(
(Number), "0.00"))
```

where RoundNo is the result and Number is the number or calculation to be rounded.

Do you know of a better way of doing this?"

Paul le Gassick

The answer is that I don't. Anyone else? (See Figs 1-5.)

### PCW Contacts

Mark Whitehorn welcomes readers' correspondence and ideas for the Databases column. He's on [m.whitehorn@dundee.ac.uk](mailto:m.whitehorn@dundee.ac.uk)