page, so put it in the middle of a ‘Hello and welcome’ message for a personalised touch. There are a few other commands too; include **flastmod** and **fsize** which tell you the date last changed and size of a file, respectively. But one of the most useful is the **exec** command which our example uses to create a simple page that will display different information for Mac users.

Here’s the basic code for the page (missing the HEAD portion):

```
<h1>Welcome to my test page</h1>
```

The “echo” command can help you personalise your site for the Mac (left)...

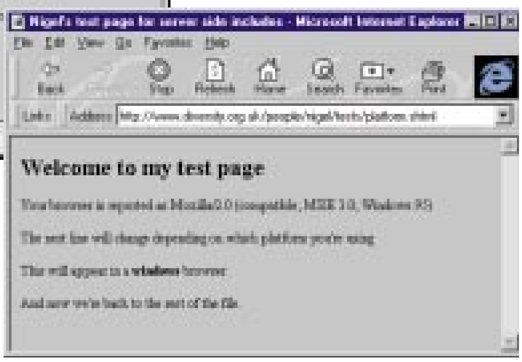

...and for Windows 95 PCs with a minimum of coding (right)

```
1    Your browser is <!—#echo var=”HTTP_USER_AGENT”—><P>
     The next lines will change depending on which platfom you’re
using<P>
2    <!—#exec cmd=”/usr/local/etc/httpd/cgi-bin/platform_ticket”—><P>
```

And now we’re back to the rest of the file.

The first marked line simply displays the information returned by the user’s browser, which is sent to the web server and stored at HTTP_USER_AGENT; the second line calls a program called “platform_ticket” which is a Perl script. Instead of saying “cmd=” you could say “virtual=” and specify the program as a name like “/cgi-bin/platform_ticket”, but on some servers (we used Netscape FastTrack) you might not be able to access the files you need from within the script unless you give the full path to them. Don’t forget you’ll need to include code like this in a file ending in .shtml and enable Server Side Includes on your web server. That option may not be available if you simply have free space with your internet account.

The next thing to do is to create the different versions of the page that you want to display for Windows and Mac users, and save those sections in files called “win.inc” and “mac.inc” (or whatever you’ve referred to in your script). Finally, you’ll need to write the actual script itself, which decides which file to send back to the browser:

```
        #!/usr/local/bin/perl
        # script to output a file based on the browser accessing a site
1       $browser = $ENV{‘HTTP_USER_AGENT’} ;
2       if ( $browser =~ /mac/i ) {
3               open( INCLUDE, “mac.inc” ) ;
4       } else {
           open( INCLUDE, “win.inc” ) ;
        }
5       while( <INCLUDE> ) {
           print $_ ;
        }
6       close INCLUDE ;
```

1. First, we’ve saved the information into a new variable; not strictly necessary, but if you wanted to change the value, you could do so more safely. 2. This line checks to see if $browser contains the string “mac”, and the i on the end ensures that it’s case insensitive. 3. If “mac” occurs in the $browser variable, then the file mac.inc will be opened; INCLUDE is the “handle” given to the file so we can refer to it later. 4. To do more tests, change “else” to “elsif” and add more conditions here. 5. These next three lines read from the file that was opened, one line at a time, printing it out. 6. Finally, we close the file, and the job’s done. So, with just a handful of lines of Perl, and a couple of comments in your web page, you can give people the feeling that they’ve visited a site that’s a little more customised to their needs; and of course you can do more than just display pages based on their browser. For more information about Server Side Includes, visit twister.luton.ac.uk/Manual/ssi.html.

p250 ➢

Questions & Answers

Q We've created an intranet based on a system running Netscape's FastTrack server. So that everyone knows when a page has been updated, we'd like the server to include information automatically. I've heard of Server Side Includes, but can't find much information about them. How do we do this?

A Server Side Includes, also referred to in the Netscape documentation as parsed html, are a way of telling the server to replace certain parts of a page with other information; if you decide to use this, you'll have to tell the server which pages you want it to parse. The usual way is to save pages with the extension **.shtml** and include commands within them, embedded in HTML comments, which allow you to add features such as the date a file was modified or the name of the system requesting the page, and so forth.

The panel on page 248 gives details of how you can use Server Side Includes to create pages that will depend on the type of computer being used to access them. However, if all you want is a simple date and time on the bottom of your pages, Netscape FastTrack server has a "page footer" feature that you can use to achieve the same result, without having to worry about server side includes.

To activate custom page footers, use the server manager, and click on the Content Management button in the top frame, then choose Document footer in the left pane, and you'll see a screen similar to the one displayed here. You can type the footer text that you want in the box, including the tag **:LASTMOD** for the last modification date of the file, and select the date format from the drop down menu.

When you've made your changes, click on OK, then choose the button marked Save and Apply, stop, then restart your web server. All your pages, unless you selected just a portion of the server, will now have a footer added automatically.

Netscape's FastTrack server can automatically add a timestamp to all your pages

Q I am an AOL user and I want to know how I can connect to the IRC network. Do I need to configure my system specially?

A No, you don't need to do anything special if you have the latest version of the AOL software, although older versions won't work with programs that require a 32-bit winsock stack (the part of Windows that links you to the internet).

The best thing to do is to upgrade to the latest version of the AOL software (which you can do at keyword UPGRADE). If you have it already, the system will tell you when you try to download it again. When you've installed AOL, you'll find that in one of the AOL directories there's a file called **winsock.dll**; install your IRC program into the same directory (or alternatively, make sure that it's in the path) and then all you need to do to connect to IRC is to start AOL, then start your IRC program.

One of the best IRC programs to use is called mIRC; you can download it from www.mirc.co.uk or at keyword IRC on AOL, where you'll also find additional information about configuring your system.

Q For the purposes of web-page design, what safe assumptions can I make about resident fonts? Is there a standard set installed with Windows 3.1 and 95 and, if so, where can I learn what it is?

A First, don't forget Mac users! The best thing to do is to assume as little as possible about fonts, as different systems may have widely differing selections. The best bet is to stick to the core TrueType fonts, which includes Arial, Courier New and Times New Roman; Mac users who have installed Microsoft Office applications will almost certainly have these fonts available. Those who don't can download them (and others) from www.microsoft.com/truetype/fontpack/mac.htm.

If you really do want to use lots of fonts to create a specific look for your web site, the best bet may be to provide links to the Microsoft (or other) font pages from the front page of your site, so that those who want to will be able to download them and make sure that they see things as you really intend them.

Q I am using Internet Explorer 4 Final (Build 4.71.1712.6) and MS Dialup Networking 1.2 on Windows 95 4.00.950. When I want to browse the internet it won't dial automatically. I think the reason is because the Save Password checkbox is greyed out. I have tried reinstalling the connection several times following the help exactly, but to no avail. I have also tried using Netscape Communicator but, again, I still cannot check the checkbox.

A Windows doesn't think you're signed on to the system as a valid user. When you install networking, Windows 95 assumes that you need to log on before the system is secure. However, many people simply press the Escape key when they're asked for a password, and since everything else works properly, don't realise the importance of what they're doing.

When you start Windows, if you don't

Setting a blank password will stop Windows prompting for a user name

p252 ➢

have a user created, give yourself a name and then leave the password blank. You'll be asked to confirm it, so leave the box blank again. Now, when you start the system, you won't be asked for a user name.

If you've already created a password and want to avoid having to say who you are each time you start the system, you need to log in properly when you start up, with the correct password (or create a new user, with no password, instead), open the control panels and choose Passwords, then set the password to be blank. You should be able to tell Dialup Networking to save the password, so you'll be able to start your computer and get on the internet, without having to remember anything more than where the power switch is.

Q I'm writing CGI scripts in Perl 5. As Perl and most scripts/libraries are freely available on the net, what's the position regarding copyright?

A Your scripts are your own intellectual property, and unless you decide to give them away, then there's nothing anyone else can do with them. But if you were to distribute, for example, a Perl-based online shopping system which used freely available libraries from the internet, you may find that the conditions of use of those libraries prohibit you from using them as part of your commercial product. After all, why should you sell something someone else wrote?

It's most unlikely that someone would be able to steal your scripts, unless the security on your web server is more or less non-existent. If the scripts are kept in a script directory, it's hard for anyone to see them. All that's passed to the web browser is the results of running the script itself.

You should clearly mark your scripts with your name, and other information, anyway, including an indication of your copyright, and a clear prohibition on use without your permission, if you wish to restrict their use, though of course it can be very hard to prove that someone has used your script, since it can easily be modified to look very different (especially with a language as flexible as Perl).

Q I'd like to set up a dial-in PPP link to the Unix server in my office, so that I can connect to the internet that way, rather than paying to subscribe to an ISP. How do I allocate IP addresses to the link? One of the manuals says something about having to have a separate subnet; what's that, and how do I create one? Is it absolutely necessary?

A A subnet is a section of addresses in a range of TCP/IP network numbers that, when written in binary, has the top section of bits all the same. For instance, if you have a network of 256 addresses, you might use four bits to create subnets, giving you sixteen subnets, each capable of having fourteen machines on it (addresses with the bottom bits all either 1 or 0 aren't used).

By putting your PPP link on a subnet, the other computers on the network will be able to work out easily that traffic for that subnet has to go via a certain route. However, that's not always possible, and it might mean re-configuring lots of your systems just to add a single dial-in link.

An alternative solution is to use proxy-ARP. ARP is the Address Resolution Protocol, and it's the way in which computers running TCP/IP match ethernet addresses to IP numbers. With proxy ARP, you add an entry to the arp table (usually using the /etc/arp command) that effectively means your Unix server is saying to the rest of the network "If you've got anything that's destined for this address, send it to my ethernet card." When the information is received at the Unix server, it'll be passed to the PPP interface.

Since you don't say which version of Unix or PPP you're using, it's hard to be more specific, but it looks like using proxy ARP may be the best solution to your problem. Most modern versions of PPP will be able to make the arp entry automatically when they receive a connection; check your manual.

PCW Contact

Nigel Whitfield is a freelance journalist, maintainer of several internet mailing lists and consultant to a number of non-profit organisations. He can be contacted via internet@pcw.co.uk.

Port of call

Roger Gann takes a look at Intel's Accelerated Graphics Port (AGP) multimedia graphics development. Plus, how overclocking could speed up your CPU — but don't go overboard!

Over the next few months I'll be looking at USB and IEEE 1394, or FireWire, peripheral interfaces. But for now, I'll give you the low-down on Intel's new graphics technology, the Accelerated Graphics Port (AGP). AGP is the third out of four multimedia graphics developments from Intel. It kicked off with MMX, then the Pentium II, and now AGP, and will culminate with Intel's own proprietary graphics accelerator chip, the i740.

Intel's purpose in developing AGP was to improve 3D graphics throughput. Those graphics require far more memory than 2D versions because of the need to store data such as display lists and texture maps, which are particularly memory intensive.

Today's PC applications seldom use a texture map, for the simple reason that's it not practically possible. However, the desire for improved realism, particularly in games, is pushing in this direction. Typical texture maps range from 2Kb to 128Kb. Today's applications average from 12 to 24 different textures of varying sizes in each scene. Thus, each scene generally requires several megabytes of texture-map storage.

To render the many scenes in a 3D application, like a game, at realistic speeds, many texture maps must be readily available in memory. This requirement is incompatible with the current model for 2D graphics on the PC, in which the graphics accelerator stores data in a dedicated frame buffer. Applications may require 5Mb to 10Mb of texture memory, beyond today's typical frame buffer size of 2Mb.

If all the necessary 3D data were kept in a frame buffer based on high-speed RAM, the cost of the additional memory would be prohibitive for mass-market PCs. Dedicating huge amounts of memory solely to the graphics accelerator is impractical. One solution is to store some of the 3D data in system memory instead. This doesn't decrease memory requirements of 3D but allows other applications to use the memory when 3D applications are not running. The challenge is to enable the graphics accelerator to access system memory fast enough to feed the 30fps rendering that users want.

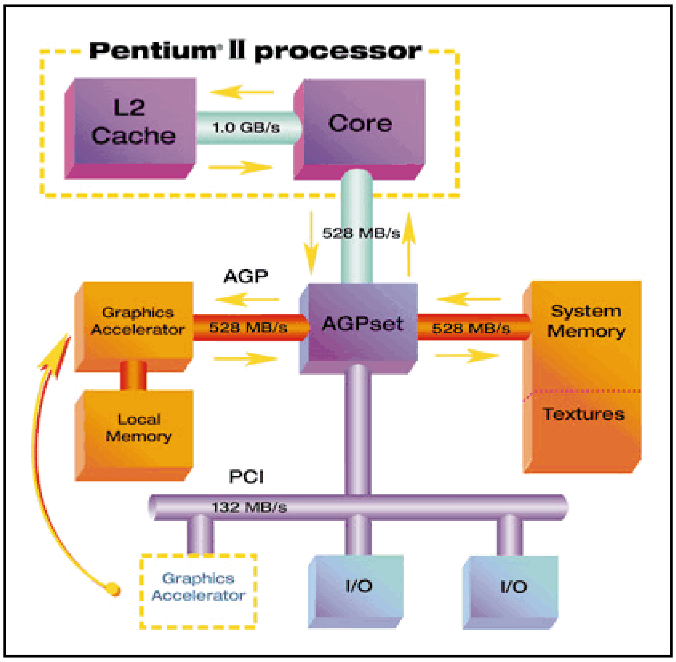

Diagrammatic representation of how AGP works

AGP in detail

The AGP specification is based on that of the 66MHz PCI 2.1 which doesn't have much support at present since all current PCI cards can only run at 33MHz. AGP adds three special extensions via so-called "sideband" signals, provided by some special lines added to the PCI connector:

- pipelined memory read/write operations;
- de-multiplexing of address and data; and
- timing for data transfer rate as if clocked at 133MHz.

All this is designed to improve throughput compared with the PCI bus which, at 33MHz, can deliver 132Mbps at peak rates over its 32-bit data bus.

This bandwidth has to be shared with other peripherals and it relies on a slow software-based method for moving large blocks of data to the frame buffer. Because the typical texture map is stored in small, non-contiguous memory blocks in system memory, many CPU accesses are needed

Intel's 440LX AGPset

- Pentium II systems only.
- Supports AGP.
- Supports ACPI for remote management, quick boot-ups.
- Ultra DMA for better access to IDE hard drives and CD-ROMs.
- SDRAM for faster communication between CPU, graphics accelerator and PCI devices.
- Supports AGP OS and apps software. Drivers and controllers due 1998.

p296 ➢

to retrieve the entire texture map.

The AGP bus runs twice as fast, at 66MHz, and so is capable of a peak rate of 264Mbps. This is in the so-called classic “x1” mode. However, if the “x2” mode is used, which transports data on both the rising and falling edges of the 66MHz clock, the data transfer rate can be redoubled to a theoretical peak rate of 528Mbps. As a result, the AGP peak transfer rate matches that of main memory which, in Pentium or better systems, operates with a 64-bit wide bus at a 66MHz-bus clock.

Thanks to the high data transfer rate between the graphics accelerator and main memory, AGP enables 3D graphics accelerators to use main memory instead of local memory to store texture data. These textures, previously loaded into the local graphics accelerator memory, can now be processed in main memory without a performance impact. Intel calls this technique DIME (Direct Memory Execute).

Most applications could use 2-16Mb for texture storage — by using AGP and DIME, they can get it. The main beneficiary of AGP will be high-end 3D graphics applications and games. Normal business applications will not benefit at all.

AGP hardware

AGP consists of a proprietary expansion slot and the Intel 440LX AGPset, although an “any-processor” Socket 7 solution is imminent from motherboard manufacturer, VIA. The AGP slot is very similar to the existing PCI slot but located further back from the rear of the motherboard to prevent inadvertent insertion of non-AGP cards.

The 440LX features Intel’s Quad Port Acceleration (QPA) which increases overall system concurrency, bandwidth and performance. The QPA includes the direct connect AGP, dynamic distributed arbitration and multi-stream memory access. The AGPset also offers Advanced Configuration Power management Interface (ACPI) for enhanced power management and plug-and-play capability, Ultra DMA for faster storage throughput and SDRAM for increased RAM performance.

At present, AGP support is lacking in Windows 95 and Windows NT 4.0, so all you get with AGP is a faster data transfer rate and no DIME support. This support will arrive via DirectDraw in Windows 98 and Windows NT 5.0 next year. Parts like the direct pipeline from the graphics accelerator to system main memory would be unaffected by the absence of OS support, but a major portion of AGP functionality, like the dynamic memory allocation required to allow video to peacefully co-exist with Windows in the same RAM space, would be missing without this support. So, Intel has designed a DirectX 5.0 software driver that will enable its AGP silicon to work with the current version of Windows 95, rather than waiting for Windows 98. Rumour has it that Intel postponed the launch of the 440LX chipset when it became apparent that the increasingly later release of Windows 98 would largely render the 440LX’s launch meaningless.

Advantages of AGP

- Up to four times the bandwidth of PCI.
- Exclusive graphics bus.
- DIME; Direct Memory Execution of textures.
- CPU accesses to system RAM can proceed concurrently with the graphics chip’s AGP RAM reads.

Overclocking

I’ve had a couple of letters asking about processor overclocking and whether it is a good idea. Overclocking refers to the process of running your CPU at a clock or bus speed for which the CPU has not been specified; typically a higher speed. In most cases this is achieved by changing a few settings on your motherboard.

Sounds great; but is it safe? The surprising answer, in most cases, is yes. Provided you don’t go wild with your overclocking, it will probably run just fine.

The big problem is one of heat: the faster a processor runs, the hotter it runs. Unless the CPU is cooled properly, this can cause the chip to overheat and be damaged by so-called “electro-migration”, which is a kind of slow rotting of the silicon. Most chips can safely run at 80° C, which is too hot to touch, but a cooling fan can drop this to 50° or less, so the degradation needn’t be a major threat.

This is true of Intel CPUs, which seem conservatively rated. CPUs from AMD and Cyrix are less conservatively rated and run much hotter as a result, so these are less suitable candidates for overclocking. There’s far more scope with Intel silicon.

Another problem is that random system hangs can occur after overclocking. Windows 95 is very sensitive for some reason: I installed a faster Pentium in an old PC and Windows 95 wouldn’t load, telling me it couldn’t find KRNL386.EXE. Hence, thorough testing of an overclocked system is an important precaution.

Finally, you should not forget the warranty implications: if you overclock, you will probably void your warranty.

There are two motherboard settings which need to be changed to overclock: the bus speed, and the processor multiplier. The Pentium supports three bus speeds (50, 60 and 66MHz) but some recent motherboards additionally offer 75 and 83MHz. To change the bus speed, look in your motherboard manual for something like “CPU External (BUS) Frequency Selection”: these are the jumpers which need moving. Go easy here and just move the bus speed up a notch. If you do try the 75 or 83MHz settings, remember that this increases the PCI bus, from 30 or 33MHz to 37.5 or even 41.6MHz. This can lead to several problems with PCI devices like SCSI host adapters, some video cards and network cards. Each CPU uses a multiple of the bus speed — the so-called bus multiplier. A P120 uses an X2 multiplier on a 60MHz bus and a P166 uses an X2.5 multiplier on a 66MHz bus. Intel Pentium CPUs support the following multipliers: X1.5, X2, X2.5 and X3. Intel Pentium Pro CPUs support X2.5, X3, X3.5, and X4.

To change this setting, find something like “CPU to BUS Frequency Ratio Selection” in your motherboard manual. For example, a P200MMX in the right motherboard can be made to run at 233 or even 250MHz. On the power supply front, merely switching from STD to VRE is often sufficient, but try increasing the voltage if the original settings do not work.

Overclocking ingredients

- Intel Pentium processor.
- Good-quality motherboard, capable of 75MHz or 83MHz bus speed, supporting a wide range of CPU supply voltages (overclocked CPUs need more juice).
- Good-quality RAM; fast EDO SDRAM.
- Decent CPU fan plus good heatsink.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address (p12) or via email at hardware@pcw.co.uk.

CD drive-in

Roger Gann shows you how to install a CD-ROM drive in your PC. His step-by-step guide tells you how to get what in the right slot and get your discs spinning in record time.

These days it's virtually impossible to buy a PC that does not come with a CD-ROM already fitted. But maybe you already own a PC without a CD-ROM, or perhaps you want to fit a faster drive or a CD-R? No matter what you're installing, whether it's CD-ROM, DVD-ROM, CD-R or CD-RW, the installation procedure is almost identical. As I'll be looking at SCSI in more detail in a month I'm confining myself to the fitting of an ATAPI (or IDE) CD-ROM drive.

Step-by-step

There are two phases to fitting a CD-ROM: installing the hardware, and then getting the operating system to recognise it.

Hardware installation

1. Take the usual safety precautions. Gather together all the instructions and driver disk plus all the tools you need (generally just a Phillips screwdriver). Turn off the PC and remove the casing lid.
2. Configure the CD-ROM drive, if it needs it. You'll probably want to install it on the secondary IDE channel so make sure it is set to be a "Master". (In my experience, though, setting a CD-ROM to be a slave or master seems to make little difference. But do the right thing!)
3. Check out the drive fixings and where it's going to fit. You'll need a free 5.25in drive bay for the new drive. Try not to mount it in the lowest bay, as the disk tray or caddie could foul your keyboard. Mounting it high makes it easier to get at the disk eject button with the tray open. If it's an ATAPI drive, try and keep it close to the IDE connectors on the motherboard (or I/O card) as IDE ribbon cables tend to be short. Make sure you have the right mounting hardware, too, like bolts or rails.
4. Insert the drive in the drive bay and make sure it doesn't foul anything. Tighten (but don't over-tighten) the mounting bolts.
5. With the CD-ROM drive installed in its bay, attach a spare power cable and the 40-way data cable to it. Some connectors are "polarised" with a little notch and so cannot be fitted the wrong way around. If yours isn't, make sure the coloured (red or blue) stripe on the ribbon cable goes to Pin 1 of the connector on the drive, which is normally the side next to the power connector. If you're lucky, the pins will be numbered. Ensure it is correctly orientated at the interface card (i.e. other) end.
6. If you have a sound card, don't forget to connect the audio cable. This is a thin, flexible cable that plugs into the back of the CD-ROM drive and into the sound card. This cable isn't necessary to play sound from a video or game on CD-ROM, but it will be required if you want to play an audio CD through your setup.
7. Power up the PC. You shouldn't have to make any alterations to your CMOS setup but some BIOSes can now treat a CD-ROM as a boot device, so check out Setup's IDE settings just in case. Listen to an audio CD to check that the drive is working.

Software configuration

We now need to get the operating system to recognise the new drive and assign it a drive letter.

- Under DOS/Windows 3.1 this used to involve separate Real mode drivers plus the MSDOS CD-ROM Extension, MSCDEX. But Windows 95 doesn't need external drivers, as it relies on its own 32-bit internal drivers.

p292 ➢

Compaq question: Going into Overdrive

Q I use a Compaq Presario 7150 with P75 processor, 16Mb RAM (upgraded to 4 x 4Mb SIMMs), a 1.2Gb Western Digital HDD, running Win95.

Recently, I installed Microsoft Flight Simulator v6.0 for Win95 and the machine is showing its limitations with regard to speed and graphics quality. Contacting Compaq on four separate occasions regarding further upgrades to the Presario, has been a hopeless task and I have variously been advised: the machine cannot take a larger Intel Pentium (I was thinking of a P133 or P166); I cannot fit an Overdrive upgrade; I *can* fit a P100 but this should be done by a specialist; “unsure of what can be done” and so on (the warranty implications are acknowledged).

So, can I fit a higher-speed Pentium to the Presario? Can I fit an Overdrive (I believe that 75 can go to 120 by this route but if feasible, would this make a difference)? Would extra RAM help, like increasing to 24Mb?

Simon Potter
Hereford

A Checking out the System Upgrade Page at Intel’s web site reveals that a Pentium Overdrive is suitable for a Presario 7150 — the 125MHz version (or PODP3V125, to give it a part number). Yes, it will make a performance difference, but don’t expect your socks to be rocked. Adding more memory is a good idea for Win95 and, given the low price of memory, it might make more sense to opt for this first, but you may have to ditch a pair of 4Mb SIMMs to make room for a pair of 8 or 16Mb SIMMs. Check to see if your CD is cutting the mustard: games run from slow (double-speed) CD-ROMs will suffer.

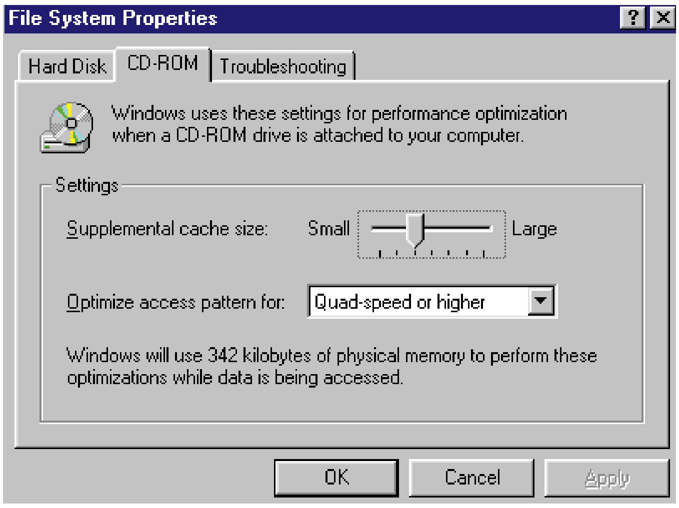

Right Once you have fitted your CD-ROM drive, it is important to adjust the Windows 95 settings to take account of the increased speed

So if you intend to use Win95 exclusively (if you never intend to boot in MSDOS mode to play DOS games) you won’t need to install these drivers.

● If you do install the Real mode drivers, it will add a line to your CONFIG.SYS that will look something like this:

```
DEVICE=C:\CDROM.SYS /D:MSCD0000
```

The **/D:MSCD0000** assigns a system signature to that CD-ROM (you *can* fit more than one CD-ROM drive, after all). You can call it whatever you like, but this signature *has* to be identical to the one specified by MSCDEX: the line in AUTOEXEC .BAT might look like this:

```
MSCDEX /D:MSCD0000
```

Note the matching signature, **MSCD0000**.

● Although Win95 can automatically detect most new hardware, the situation with IDE devices is a little special. Win95 doesn’t actually detect an IDE hard disk, as it takes it for granted that there’ll be a hard disk from which to boot and IDE is the lowest common denominator. And because ATAPI CD-ROMs are IDE devices, they get tarred with the same plug-and-play brush.

Generally, if you try to get Win95 to auto-detect a newly-fitted CD-ROM, it won’t find it. In fact, Win95 can only specifically detect the older types of CD-ROM drives that use proprietary interfaces like Mitsumi, or MKE/Panasonic. The horrible thing is, if Win95 doesn’t “see” your ATAPI CD-ROM drive at boot time, you can’t explicitly add an ATAPI drive manually, using the Add New Hardware applet. It either works first time or doesn’t, simple as that.

Optimising CD-ROM performance

CD-ROM drives are relatively slow devices but cacheing them can enhance their performance. Win95 uses a 32-bit development of SmartDrive, VCache, which ought to be adjusted to suit the type of CD-ROM you’ve got and how much memory you can afford to donate to the cache.

● To do this, open **Control Panel** and click on the **System** icon. Click on the **Performance** tab and the **File System** button. Finally, select the **CD-ROM** tab. This dialog lets you specify the speed of your CD-ROM and the amount of cache memory you want to devote to it. When you’ve made your choices, restart Win95.

● Most high-speed CD-ROM drives, i.e. 16-speed and above, place an extraordinary burden on the CPU during data transfers. Sometimes, in excess of 90 percent of the processor’s work is devoted purely to supervising the transfer of data off the CD-ROM into memory, which leaves only ten percent for running the rest of the computer. However, most modern CD-ROM drives support Direct Memory Access (DMA) transfers, which place little burden on the CPU.

If your PC’s chipset supports bus-mastering DMA, it’s important that you take full advantage. By default, under Win95, support for DMA transfers isn’t turned on; but once it *has* been turned on, there should be a noticeable improvement in PC performance when making intensive use of the CD-ROM.

● As well as requiring hardware support, you also need operating system support, and if you have the OSR2 release of Win95 (version 4.00.950 B) then DMA support is available. If you have the right chipset, additional options appear on the CD-ROM settings in Device Manager. So, click on the Settings tab and check the DMA box. Check the Sync Data Transfer option for good measure.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address (p12) or via email at hardware@pcw.co.uk.

Work and play

Roger Gann looks for the first time at plug-and-play. Marketing gush says you plug it in and you're instantly ready to go. But is it that easy in reality? The hardware hardman checks it

I've been wittering on about hardware and upgrading for the past 18 months or so and managing to avoid answering readers' questions in print. But no more! Starting this month, about a third of the column will be devoted to the more interesting or challenging readers' questions. If you have a toughie and need some help, please drop me a line care of *PCW* or email me at hardware.pcw.co.uk. I can't promise to answer all your problems but I'll do my best.

Okay, back to work. This month I'll be looking at the minor miracle of plug-and-play. For the past year and a half I've been largely ignoring this major improvement to PC upgrading, preferring to show you how to install hardware from first principles, but I feel the time has now come to shed some light on the black magic of plug-and-play.

Bus route

Microsoft, Compaq and Intel originally defined the bus-independent plug-and-play (PnP) standard as long as four years ago. The goal of plug-and-play is to turn a PC into a user-friendly device, as easy to use as a domestic appliance. A user can simply attach a new device ("plug it in") and begin working ("play"). For plug-and-play to work properly, it requires support from three areas: the add-on hardware, the operating system and the supporting software such as drivers and BIOS. The BIOS specification was published on 1st November 1993.

When it works properly, two principal benefits should flow from plug-and-play. First, ease of use: an obvious gain really, but plug-and-play makes it easy to install and configure add-on devices with little or no user intervention. After you've plugged the device in, the BIOS and/or operating system will carry out all the graft.

Second, reduced support costs: this may not seem important to the home user but it is still relevant because hopefully you won't need to hang on the phone quite so often when plug-and-play is in full swing. It can also reduce wasted time and frustration. Multimedia vendors find that most of the sound cards returned to them as defective actually work perfectly. It's just that end-users find installing and configuring them far too complicated.

The plug-and-play standard is not operating system specific and will appear in other operating systems in due course, but for the moment Windows 95 is the only operating system that fully supports it. Both Windows NT 4.0 and OS/2 Warp 4.0 currently lack plug-and-play support, but future versions of both operating systems will support it.

A major benefit of plug-and-play in Windows 95 is "event notification of changes to the system state". This neat trick permits dynamic changes in the installed hardware to prompt an automatic reconfiguring of the operating system and its device drivers. For example, if you were to undock a notebook, not only would the operating system realise it was coming off the network and so should unload all the network drivers, but it would also alert any apps that had files open on the server, that those files ought to be saved.

Going live

Plug-and-play is also a "live" feature and, while you can't hot-swap PCI cards, you can do this with PC Cards. Whenever you insert a new PC Card, Windows 95 detects it automatically, and will try to automatically configure it and install any required device drivers. If it doesn't have the correct drivers, it then prompts the user for them. And all the user has done is slide in a PC Card!

How it works

When a new plug-and-play card is installed in a PC with a plug-and-play BIOS and then powered up, the BIOS will detect the new card and set a flag. When Windows 95 boots, it picks up this flag and launches its hardware detection software to identify just what has been installed. Each bus (PCI, ISA, PC Card and so on) is interrogated by its own bus enumerator; they also inspect each card to see whether it's PnP and, if so, what range of hardware resource settings it can accommodate.

Once data has been collected from all the cards, a Windows 95 arbitration algorithm allocates the available IRQs, DMA channels and I/O ports in the most efficient manner. This data is then written to the PnP cards. If a legacy (i.e. non-PnP) card is detected, an attempt will be made to identify the card from the Registry database and its default settings assumed. The PnP cards can then be configured, using the remaining free resources.

If a SCSI bus is detected, a second pass is made once the SCSI chipset has been identified, to enable the SCSI enumerator to interrogate the SCSI to find out what SCSI devices are attached. The same occurs with PCMCIA slots and PC Cards. This process is transparent to the end-user: all they'll see is a Windows 95 dialog saying it's detected new hardware and that it's installing the drivers for it. If the drivers aren't included with Windows 95, the user will be prompted for a floppy disk containing them.

On legacy PCs, the information stored in the centralised Registry is used to notify the

p294 ➢

Question time: Playing with Fireball

Q I recently bought a Quantum Fireball 1280Mb HD and connected it to a PCI motherboard with onboard EIDE and AMI WINBIOS. I used “auto detect C” (I did not alter setup to enable LBA) and then started to load DOS which formatted the disk to 503Mb, so I re-read your Hands On article on hard disks which mentioned the DOS limitation. I approached various sources without success, including two computer service/upgrade/repair houses, one of which had supplied the disk. They apparently don’t have the problem but both offered to reformat the disk for me. While this would have solved the disk problem, I prefer to understand and do things myself rather than remain ignorant for future occasions. So I re-read the article in more depth and noticed the bit about logical block addressing in modern boards — then found this in my BIOS, “Enabled” it, and tried again, but DOS would only confirm the existing format and re-format to the same 503Mb. Is there a way out of this problem? Incidentally, the article mentions Logical Block Addressing and IDE address translation: are they the same thing?

Bill McGregor

A No, they’re not, Bill. Address Translation refers to an IDE drive’s chameleon-like ability to adapt to a drive geometry (its cylinders, heads and sectors) different to its native geometry. So long as its capacity isn’t overstated, an IDE will be happy with just about any geometry you care to throw at it. Let’s say drive X actually has 100 cylinders, two heads and 50 sectors per track (spt) — that’s a total of 10,000 sectors (100 x 2 x 50). Address translation allows you to specify a different geometry: e.g. 50 cylinders, four heads and 50spt. This way of describing the layout of the hard disk, often abbreviated to CHS, is unsuitable for dealing with the large IDE drives we see today and so the LBA method is used instead. If your PC’s BIOS supported large hard disks then FDISK would be able to see more than 503Mb, so the only conclusion I can draw is that your PC, as it stands, is incapable of directly supporting IDE drives larger than 503Mb. Why should this be? It’s down to a clash between the drive geometries supported by the PC’s BIOS and IDE. DOS errs on the side of caution and picks the lowest common denominator from each. Check out the table shown below.

So we know the cause of the problem, but what can be done about it? Most hard disk manufacturers have anticipated this problem and supply their larger drives with special device driver software that permits the PC to recognise the entire capacity of the drive. This is a cheap solution which makes it very easy to completely prepare a drive for use. However, the downside is that this solution can cause problems with Windows 95: unless it comes with a 32-bit Windows driver, Win95 may run its file system in 16-bit MSDOS compatibility mode, which is a bit slower than its normal 32-bit disk sub-system. I’m surprised your Quantum didn’t come with this software so check back with your supplier.

A better solution is to upgrade your PC’s ROM BIOS. For older PCs this is easier said than done. ROMs are large, socketed chips and are relatively easy to swap out. But the cost is prohibitive and getting hold of them is another matter; it can work out cheaper to replace the motherboard with one that has newer ROMs. Your PC may have a BIOS that resides in Flash memory, making it possible to upgrade the BIOS from software for free. I upgraded a Gateway 486 in this manner. Or you can invest in an EIDE interface card: this will come with BIOS on-board, will cost you about £25, and is a direct replacement for the £10 multi I/O cards fitted in most popular PCs. Taking the drive elsewhere and formatting it on another PC may work, but it’s not recommended. It is safer from a date point of view if the drive is configured in the PC in which it is going to be used.

	BIOS	IDE	Lowest value
Max cylinders	1,024	65,536	1,024
Max heads	255	16	16
Max sectors per track	63	255	63
Capacity in bytes	8,422,686,720	136,902,082,560	528,482,304
Capacity in megabytes	8,032.5	130,560	504

user of a potential resource conflict when configuring a peripheral. It is used to perform device detection using the known resource information. Suppose the user wanted to install a new, but non-PnP, device on a legacy PC system such as one that requires IRQ 5. But a legacy network card is already installed on the PC and this already uses IRQ 5. In this situation, the system tells the user that a device is already using IRQ 5 and a different IRQ setting should be chosen.

All the pieces that go to make plug-and-play feasible are now in place. All PCs sold these days feature a plug-and-play BIOS. The Pentium is now the entry-level processor and as a result PCI is commonplace. As all PCI cards adhere to the plug-and-play standard, this means machines less than three years old are the ones that will make a perfect match for the new operating system.

PnP isn’t a panacea for all hardware configuration ills: it won’t resolve duplicated hardware clashes, for example. When I installed an Adaptec plug-and-play SCSI card, PnP didn’t draw my attention to the fact that there were two floppy disk controllers in the system. The automated configuration didn’t disable the controller on the Adaptec; that had to be done with a DIP switch. Never lose sight of the fact that a PC has only a finite number of free IRQs and DMA channels, and when you’ve used them all up, all the plug-and-play in the world won’t help you.

■ *Next month: Troubleshooting PnP.*

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address or via email at hardware@pcw.co.uk.

Sounds easy...

…and it is easier than you think, with the advances that have been made in plug-and-play technology. Roger Gann takes you through fitting a sound card.

It is almost a year and a half since we previously dealt with installing a sound card. A lot has happened since then, so it's high time we returned to the subject.

Back then, I installed the granddaddy of all sound cards, the Mk I SoundBlaster, a mono card with FM synthesis and 11 MIDI voices. This time I've laid my hands on the latest in PC audio hardware, the SoundBlaster AWE64, which represents the state of the sound-card art. It has come a long way from the original card that bore the name. This one offers 64-voice polyphony; 32 controlled in hardware and 32 in software.

Like its predecessor, the AWE32, the AWE64 is backward-compatible with all DOS games and supports the full General MIDI instrument set as well as Creative's SoundFonts. The AWE64 Gold also features gold-plated RCA connectors and an SPDIF digital output, and comes with 4Mb of RAM as standard, although both cards can accommodate up to 28Mb with Creative's SIMM modules.

This latest generation of SoundBlaster also features an all-new synthesis engine using E-mu's multiple-point wavetable interpolation algorithm — the same technology found in its Emulator professional-level product. This technology produces more realistic sounds by interpolating the curve between sample points, rather than connecting them with a straight line (the method used by the AWE32, and the actual standard for most consumer sound cards).

Step by Step — installing a sound card

Installation falls into two distinct phases: fitting the actual card in the PC and then configuring it and its software. Last time, we had the dubious pleasure of configuring the hardware resource settings (IRQs, I/O addresses and DMA channels) by hand. But this time we'll be putting our feet up while plug and play does all the resource allocation donkey work. Most modern sound cards support plug and play and the SoundBlaster AWE64 is no exception. Thanks to this, installing a card should rival falling off logs for simplicity.

Hardware installation

1. Electrical safety is important so take the usual precautions of switching off your PC at the wall socket. Leave it plugged in however as this will provide an earth circuit for the PC. Unplug everything attached to the system unit; typically, the keyboard and screen. All micro-electronics are sensitive to static and the kind of static discharge that gives you a very mild electric shock can be fatal to micro-electronics, so take the precaution of earthing yourself, by, say, touching some metal plumbing, before handling the new sound card. Or better still, invest in an earthing wrist strap.
2. Remove the casing lid by undoing the screws at the back.
3. Take the cover off. It's probably secured by four self-tapping Phillips screws at the rear. Put the screws and cover in a safe place. You can now see the innards of your PC.

At the bottom is the motherboard and this will have as many as eight expansion slots, a mixture of 16-bit ISA and 32-bit PCI slots along its rear edge. Unless you've got a highly integrated motherboard, one or more of these slots will already be filled, by cards such as your graphics accelerator.

4. Choose a free expansion slot. Almost all sound cards are 16-bit ISA devices so your sound card will have to go into one of these slots. If the sound card has an IDE interface for a CD-ROM drive, try to pick a slot close to where the drive is going to go. Sound cards can be susceptible to RF noise generated by other PC hardware and you may have to shuffle your cards around to rid yourself of annoying buzzes or hums. Sometimes it helps to site it as far away as possible from the PC's power supply.
5. Undo the bolt securing the blanking plate at the end of the expansion slot and remove the plate. Hold the card by its top edge and press its connector edge firmly into the expansion slot. This may be a tight fit and you may have to use a modicum of force.
6. Tighten the bolt to stop the card from flapping around. If you have a CD-ROM drive fitted, install the CD audio cable between the drive and the sound card. This will let you play audio CDs through your PC speakers.
7. The AWE64 has a digital output so plug this in to the SPDIF socket on the card and fit the blanking plate that holds the phono output socket in a spare position. Unusually, it comes complete with proper MIDI cables, so hook this to your keyboard.
8. Replace the system unit cover, do up the screws and plug everything back in.
9. Plug in any speakers to the sound card, plus any microphones or line connections. Powered speakers will have a volume control, so make sure this is turned up!
10. Power up the PC and make sure everything is working as it was before (to check that the sound card isn't interfering with anything). If you do have a problem, see the section below on troubleshooting.

Software configuration

If all is well with your PC, the next step is to install the software for your card.

- If you have Win95 your first stop should be the "Add new hardware" wizard in Control Panel. If you have a "full" plug-and-play PC, then Win95 ought to automatically identify the new hardware the next time you boot. In any case, you should let Win95 try to auto-identify your new card.
- When it does, you'll be prompted for the Windows 95 installation disks or the manufacturer's driver disks. You can then check it's correct installation by clicking on the "System" in Control Panel, then on the "Device Manager" tab and selecting the "Sound, video and game controllers" entry on the device tree.
- Finally, click on your sound card and the "Properties" button.
- However, if Windows 95 can't detect it or mis-identifies it, you should then install the Windows 3.1x drivers supplied with the card; Windows 95 comes with a reasonable range of sound card drivers but lacks, for example, drivers for Turtle Beach, Orchid and miro cards. Often, these drivers or more up to date versions can be pulled down from dial-up services such as CIX, CompuServe or AOL, or direct from the manufacturer's web site.
- If you're still running DOS+Windows 3.1x, install the drivers that come with the card.

As well as drivers, the AWE64 comes equipped with a raft of audio utility software supplied on a pair of CDs. I had the 8Mb RAM upgrade and this came with yet another CD containing the SoundFont samples, so you should install this software.

The end is nigh... Or is it?

Recent increases in bus bandwidth and processor speed have made host-based audio processing a possibility. The PCI bus, running at 33MHz and supporting burst-mode 32-bit bus mastering, can easily support many channels of digital audio: a 16-bit stereo 44.1KHz channel will consume just 385,437 PCI cycles/second, or just over one percent of the total bandwidth of a 33MHz PCI bus.

The speed of host processors has increased to an extent that many audio-processing tasks, like music synthesis, can be performed easily, even on low-end PCs. The programmer no longer has to use obscure techniques to force general-purpose CPUs to behave like DSPs: MMX and other similar technologies offer DSP-like instructions and the promise of host-based DSP function libraries. Much of the work handled by dedicated hardware on the sound card can now be offloaded onto the CPU and performed in software.

Host-based is best

Ensoniq has implemented such an architecture in developing its host-based AudioPCI sound card. The main advantage of such an approach is that the hard disk and system memory can be used to store wavetable samples. It is these samples that determine the quality of synthesised music.

Host-based audio processing also allows PC users to select from an array of optional post-processing algorithms, such as chorusing, reverberation and stereo enhancement, without dedicating expensive hardware sub-systems to those tasks.

Basic troubleshooting

■ If you don't hear sounds from Windows, there are at least four things to check:
■ Make sure the sound card works in DOS, using diagnostic software supplied with the card.
■ If it doesn't, the problem is probably physical, so check the card is seated properly, that the volume control (if any) is turned up, and that the connections to external speakers are firm.
■ If you can't play music CDs through your system, check that the CD audio cable is fitted.
■ Check the Control Panel to see whether system sounds are available. If you double-click on the Sounds icon and all your choices are greyed-out, reinstall the Windows sound driver. While you're there, make sure the necessary WAV files are in your Windows\Media folder.
■ Did your sound card come with a mixer? There's probably a simple software control that lets you adjust volume, bass, treble, and other output settings. If the volume is set too low here, that explains the silence.
■ A direct memory access (DMA) conflict can cause digital audio to be distorted or not to play at all. It can also cause your system to lock up, spontaneously reboot, or generate parity errors.
■ To change the DMA setting, check the sound card properties in Device Manager, or if you use Windows 3.1x, check the Setup dialog of the Control Panel's Drivers utility.

Troubleshooting in Windows 95

■ During the installation of an expansion card using the Add New Hardware utility, Windows 95 will scan your computer and try to identify the card or peripheral and its resource requirements, then locate an appropriate driver from the Windows 95 database.
■ If it detects a resource conflict, it will notify you and offer to start the Conflict Troubleshooter. This is an interactive Windows Help file that takes you to Device Manager and walks you through the process of identifying and resolving a particular conflict. Even if installation completes with no sign of trouble, you should still confirm that all is well.
■ Open Device Manager and click the plus sign next to a device class icon on the hardware tree. A list of all the relevant components for that type of device will appear.
■ If there is a problem, Device Manager displays one of two warning flag types:
1. The first is a yellow circle with an exclamation point. Just double-click the offending entry, click the Resources tab, and make a note of the information given in the Conflicts box at the bottom of the Properties dialog box.
2. The second warning flag is a red X, which indicates that the device is not working: either the BIOS reports the device as disabled by hardware (a jumper or switch, say) or a user has manually disabled it in Device Manager; it is disabled if the "Current" box at the bottom of the General tab is unchecked.
■ Things are pretty iffy if Windows 95 launches in Safe mode: this means it has detected a problem that prevents it from running properly. Once in Safe mode it is up to you, the user, to go into Device Manager and correct the problem. You can start up in Safe mode by pressing <F5> when the "Starting Windows 95" message appears during the boot-up sequence.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address or via email at hardware@pcw.co.uk.

Da-do, RAM, RAM

In this sequel to upgrading installed RAM, Roger Gann explains the options and gets down to the real nitty-gritty: a step-by-step guide to the practical process of plugging it in.

Welcome to the second part of my piece on upgrading your RAM. A quick recap: in last month's column I discussed the whys and wherefores of upping your installed RAM. I also looked at the differences between 30-pin and 72-pin SIMMs. This month I will complete my overview of RAM technology and show you how to install extra memory.

Choose the right RAM

It is important to choose the right type of RAM when buying your extra memory: essentially, you should buy more of what is already installed on your PC.

How can you tell? Well, not by looking at it, unless you are a rocket scientist. Unfortunately one SIMM looks much like another; peas in a pod, in other words. You have several options. You can check the documentation that came with the PC when you bought it and this should contain the full specification of your personal computer. Sometimes, the BIOS "signs on" at boot time with a panel of hardware information where the type of memory is listed. Or you can remove an existing SIMM and take it with you to the dealer.

DRAM

Dynamic RAM has to be continually "refreshed" so the charges that hold the bits of information do not fade away. The refresh speed is expressed in nano-seconds (ns) and it is this figure that represents the "speed" of the RAM.

Most Pentium-based PCs use 60 or 70ns RAM. The process of refreshing actually interrupts/slows down the accessing of the data but clever cache design minimises this. However, as processor speeds pass the 200MHz mark, no amount of cacheing can compensate for the inherent slowness of DRAM; other, faster memory technologies have largely superseded it.

Fast-page memory is a type of DRAM which allows for repeated memory accesses with minimum waiting for the next instruction. Ordinary DRAM technology is now obsolete and most modern PCs now ship with faster types of memory.

EDO

Extended Data Out RAM is a fairly recent memory technology and provides a speed gain of between five to ten percent in the memory sub-system. EDO RAM provides a wider effective bandwidth by offloading memory pre-charging to separate circuits. As a result it offers a ten percent speed boost over DRAM, while Burst EDO RAM offers another 10-20 percent.

It is cheap to make/buy and its low power consumption makes it attractive to notebook makers. Fast-page memory, by contrast, has to wait between these charging cycles, thus causing delays. EDO has not been around long but despite its popularity it is already being ousted by SDRAM. In order to use EDO RAM, your PC's BIOS and chipset must explicitly support it.

SDRAM

Compared with EDO DRAM, Synchronous DRAM is capable of transmitting data on every clock cycle, conferring a peak system-performance improvement of about ten percent over EDO.

SDRAM is a solution for fast (i.e. <66MHz) motherboard designs: the same clock drives both CPU and RAM. It offers four times the throughput of conventional DRAM and is easier to run at speeds as fast as 80MHz.

Parity and non-parity

Parity checking is a process that enables the system to detect single-bit errors and halt the system. A ninth memory chip is used to hold checksum data on the contents of the other eight chips in that memory bank. If the predicted value of the checksum matches the actual value, then all is well. If it does not, then the contents of memory is corrupted and unreliable. So the PC is halted.

However, today's memory chips are highly reliable and parity checking is no longer required by most systems. Some motherboards allow you to choose between parity and non-parity RAM. Both EDO and SDRAM use non-parity RAM.

It is easy to distinguish between parity and non-parity RAM on 30-pin SIMMs: if they have two chips or eight chips then it is non-parity, while 3- or 9-chip SIMMs have parity. Seventy-two-pin modules are more difficult to categorise but, in general, if they have four, eight, 16 or 32 chips on them, they have no parity.

Error Checking and Correcting (ECC) memory is more sophisticated and more costly. It automatically corrects single-bit errors, which account for most RAM errors, without halting the system. In addition, the system will halt when 2-, 3- or 4-bit errors are detected.

But ECC memory is relatively slow — it requires more overhead than parity memory for storing data and causes around a three percent performance loss in the memory sub-system. ECC memory is only really needed in critical situations like on servers, for instance.

SIMM Savers

The change-over from 30-pin to 72-pin SIMMs has left many people with perfectly good RAM that just cannot be used in a new machine. Consider those with a 486 motherboard with 16Mb of 30-pin SIMM RAM fitted: they upgrade the motherboard, buying a Pentium design, and then find that because the new motherboard has 72-pin SIMM sockets, all their old, expensive yet otherwise perfectly usable RAM is now rendered unusable. They can always sell it off, I suppose, but that would cost them plenty as the second-hand value of SIMM RAM is laughable, thanks to the rapid decline in the price of RAM.

One answer could be the SIMM Saver, a small circuit board with four or more 30-pin SIMM sockets on it which plugs into a single 72-pin SIMM socket. I was dubious about using one at first but have now got an old 486 PC upgraded with a Pentium OverDrive that has a pair of SIMM Savers holding a total of 32Mb of 30-pin RAM, and the system works perfectly.

A couple of points you should bear in mind, though. Firstly, although they've recently fallen in price, RAM prices have fallen more, and at about £20 a four-way SIMM Saver is still a bit dear. Don't forget that if you have a Pentium motherboard, you will need a pair of these devices. You will also need eight 30-pin SIMMs with which to fill them; because of this, SIMM Savers make more economic sense when used with late 486 motherboards which can take single 72-pin SIMMs. In the end, you might think it better to invest in new memory and flog off your redundant SIMMs.

Secondly, SIMM Savers are relatively tall gizmos and you should check your motherboard to make sure you have plenty of clearance above the SIMM sockets. And, if you are going to buy a pair of SIMM Savers, make sure you buy a "handed" pair — that is, one with front-facing sockets and the other with rear-facing sockets.

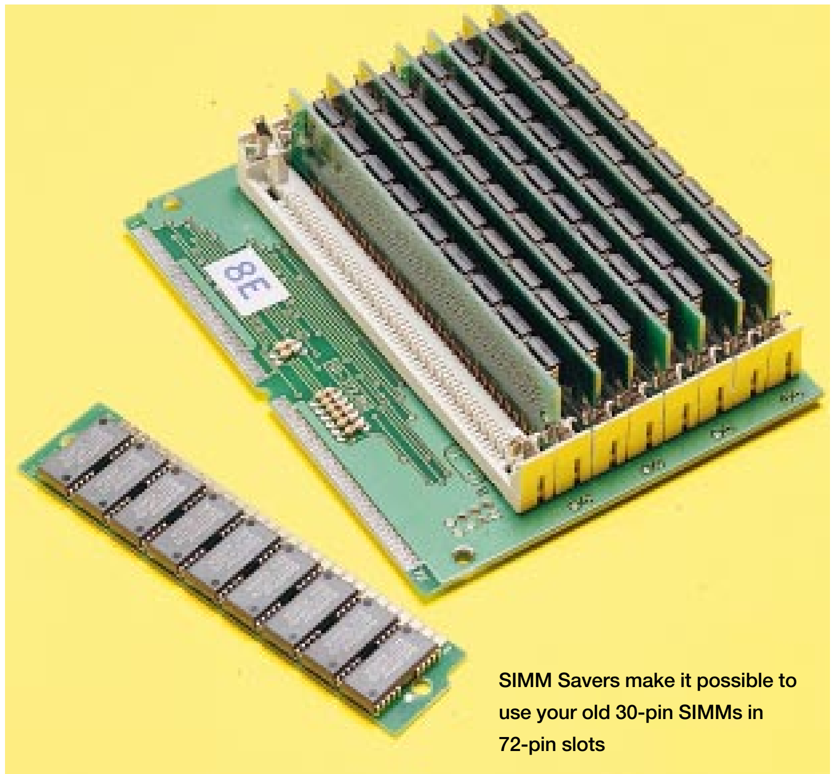

SIMM Savers make it possible to use your old 30-pin SIMMs in 72-pin slots

Cache upgrades

Most modern Pentium motherboards ship with 256Kb of Level 2 "pipeline burst mode" cache on the motherboard. This special type of memory is used to cache memory accesses by the CPU, which would otherwise be delayed by the relatively slow response times of DRAM memory.

Typically this cache is supplied on a COAST (cache on a stick) module, which resembles a SIMM but is a little shorter. This plugs into a COAST socket, which is normally located close to the processor and looks a bit like a PCI expansion slot.

Installing a bigger cache module is as simple as plugging in an expansion card. A bigger cache will improve performance and a 512Kb cache module costs about £60, which makes it a relatively dear form of memory. Note that while the upgrade is not particularly expensive *per se* and is easy to perform, the overall performance gain will be modest; probably less than five percent.

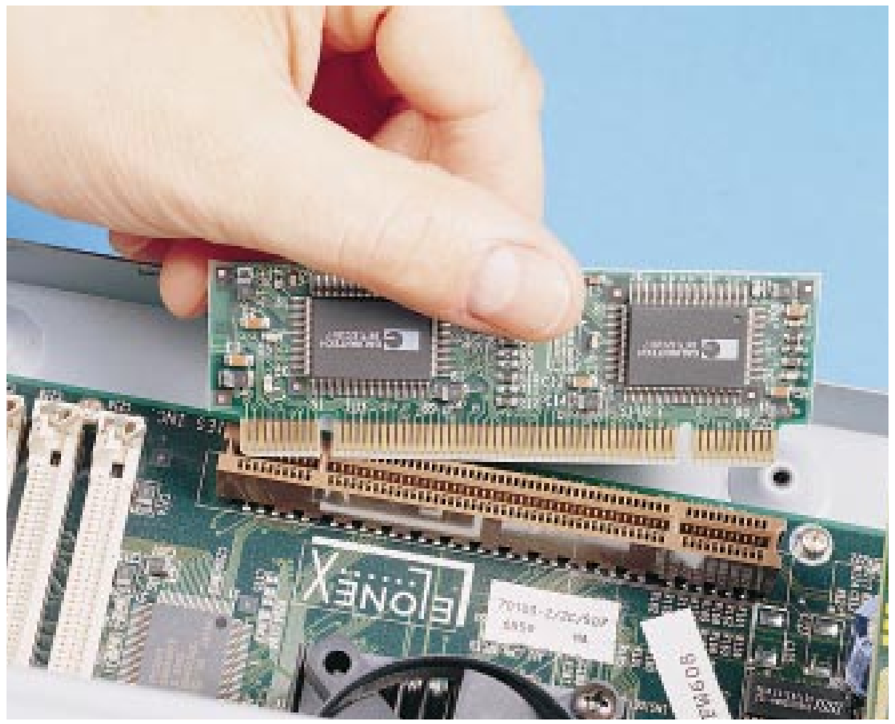

In between RAM and the CPU is the cache. This "cache on a stick" (COAST) module fits into a slot close to the CPU

p252 ➢

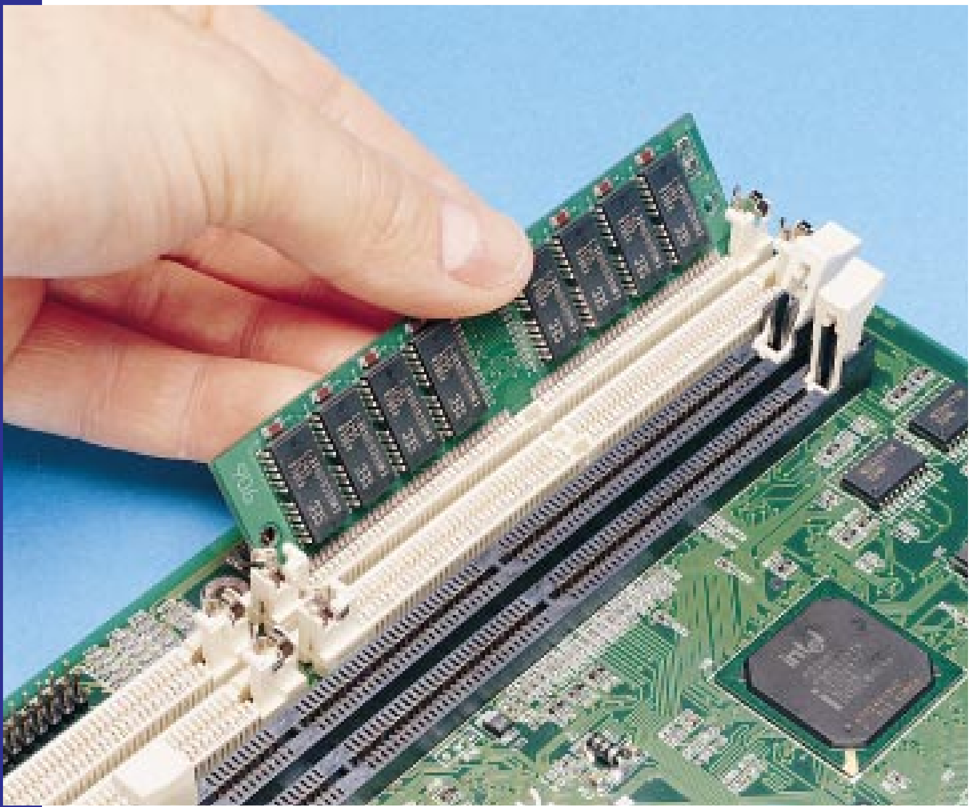

SIMMs will only fit in one way. Make sure the tabs click firmly into place at each end

Step-by-step SIMMs

Before you install your new SIMMs, take a few elementary precautions. Like all other micro-electronics, SIMMs are sensitive to electrostatic charges so before handling them you should ground yourself and discharge any static charges you have built up. You can do this by turning off the power to your PC at the mains socket, leaving the lead plugged in and touching an unpainted metal part on the case of the computer.

If you are feeling ultra-cautious you might even want to purchase an anti-static wristband when you buy your SIMMs.

Be very careful when you handle your SIMMs. Avoid touching the silicon or the contacts along the edge, and do not bend a SIMM or try to force it into a socket.

1. Power down and unplug the PC from the mains and disconnect all other leads. Remove the lid from the PC. Four or five self-tapping screws will be holding it on and you will most likely need a Phillips screwdriver to undo them. Keep them in a safe place.

2. Next, locate the SIMM sockets. If you have any problem identifying them, consult your system manual. Some may be easy to get at but others will be awkwardly located: adjacent to the power supply or under a nest of cables, for example. If necessary, temporarily remove expansion cards and unplug any cables in order to improve access to the SIMM sockets.

3. Now it is time to plug in the first SIMM. You should note that owing to the way you insert SIMMs there is a certain order to the way you put them in. Typically, if you put the rearmost one in first, you will not be able to install the one in front of it, so you must begin with the socket at the front of the bank and work backwards.

Next, work out which way round the SIMM fits. It is "handed" — one edge has a cut-out shoulder which prevents you from fitting the SIMM back to front: the correct orientation will look and feel right; the wrong one will not. If you get confused check out the SIMMs already installed, to use as a reference point.

4. The typical SIMM socket requires you to insert the SIMM at a shallow angle and then rotate it upright to lock it into place.

A small spring-loaded tab will click into the holes at each end of the SIMM. This holds the electrical contacts firmly together, ensuring reliable performance. It is important that these spring-loaded tabs lock into place.

5. Repeat as necessary. On a modern Pentium machine you will probably only have to insert one more SIMM.

6. To insert a DIMM module, flip back the catches at each end of the DIMM slot and slide in the DIMM, vertically. A DIMM module is symmetrical (that is, not "handed") and so it can go in "either way" around. Gradually lift up the catches until they click into place — they will gradually push the DIMM down into the slot as you do this.

7. Reassemble the PC, replace the cover and power up. If the installation was successful, the memory count that is displayed as your PC boots up will show the total amount of memory installed. If not, you may get an error message or a blank screen. If you do, go back and re-seat the SIMMs — that is, remove and reinstall them. Or, try putting them into different slots. If this doesn't work, take them back to your dealer and get them replaced.

Most modern BIOSes will accept the additional memory without further ado but some older versions will beep at this point and automatically enter CMOS Setup. If this happens, there is no need to enter or change any values; simply save the current settings and reboot.

Future RAM technologies

On the horizon are two competing RAM technologies: Rambus DRAM (RDRAM) and Multibank RAM.

RDRAM is a totally new RAM architecture, complete with bus mastering (the Rambus Channel Master) and a new pathway (the Rambus Channel) between memory devices (the Rambus Channel Slaves). A single Rambus Channel has the potential to reach 500Mb/sec in burst mode; a 20-fold increase over DRAM.

Unlike traditional DRAM, Multibank organises its internal 256Kb banks off a narrow central bus which allows access to each bank, individually. This design can complete a burst to, or from, one bank and then begin a burst to, or from, another — all in a single clock cycle. Multibank RAM has already appeared on some graphics accelerators.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address or via email at hardware@pcw.co.uk.

Slam in the **RAM**

Would you like to run bigger programs? Run more of them? Increase your system's power? It's not difficult. Roger Gann explains what you can do with RAM, SIMMs, DIMMs and pins.

There's an old computer industry adage which goes: "You can never be too rich, too slim or have too much installed memory." Well, actually I just made that up. But the gist of it is undeniably true. Like death and taxes, insufficient computer memory is a perennial problem, and the way things are going it looks set to become chronic.

Maybe the forthcoming Network PC and Java revolution might result in lower levels of memory being needed, but for the time being the popularity of vast, monolithic applications and operating systems, typified by Windows 95 and Office 97, make it almost inevitable that sooner or later you will be forced to add more RAM to your system. Users have grown accustomed to the luxury of running multiple apps with in-place OLE editing plus rich multimedia content, and you can't do this on a PC with 8Mb of RAM — not unless you enjoy watching paint dry.

So, this and next month's column will be devoted to memory upgrades. This month's instalment covers the reasons for upgrading RAM and takes a first look at selecting the right sort of memory for your PC. Next month's column will complete this and go through, step by step, the task of actually adding the RAM.

RAM defined

Random access memory, RAM for short, constitutes your computer's work space. When you launch an application or open a file from your hard disk, the files and data are read from the disk and copied to RAM. Because the data is held as a series of small electrical charges, these can be rapidly accessed by the CPU so it can do its job without having to pause while the data it requires is delivered to it. What's more, because the data is stored as an array of electrical charges, it can be accessed randomly (as opposed to serially) and the CPU can locate, change or erase any single byte, in any order, among several million.

I like to use the analogy of a library when describing how a computer handles data. In my computer "library", the shelves are the hard disk storage and the reading desk is the RAM. Now, although the library may have hundreds of shelves and millions of books (the files and programs), if your desk is only large enough to take one book, then that's as much as you can read or work on at any particular time.

But supposing you got a bigger desk

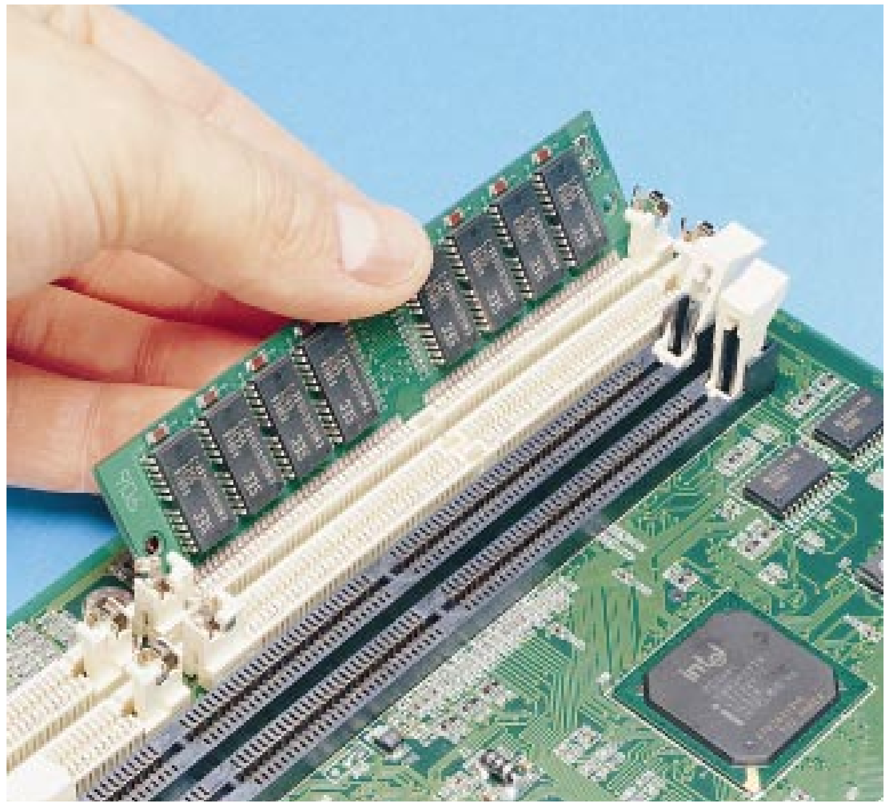

You can get up to 64Mb RAM on a single SIMM: just clip it onto the motherboard to give your system a massive performance boost

(added more memory)? Now you'd be able have several books open at a time. Adding RAM doesn't just give you more workspace, however. It also improves performance. Windows users who try to multitask several applications on 4Mb of RAM will be painfully aware of this.

And that's why memory plays such a key role: if you have too little, your PC can be, literally, unusable. But RAM is just like any other part of a computer, a plug-in component, and so it is relatively easy to increase the amount of memory installed in your PC.

There are two principal reasons you should add memory: to enable you to run bigger or more programs, and to increase system performance. The big question is, "how much"? The answer will vary from user to user but it is possible to lay down some broad guidelines. There exist so-called "sweet spots" in installed memory, above which installing additional RAM doesn't improve performance or functionality.

Life is sweet

The sweet spot for your machine depends mainly on the operating system you're using. For example, if you're only using DOS, 1Mb is probably enough (or 4Mb if you want to play games). If you want to use Windows 3.1x, 4Mb should be viewed as an entry-level quantity of RAM and 8Mb as the sweet spot.

Windows 95's sweet spot is 24Mb, with 16Mb as a realistic entry threshold. There's also a ceiling, above which adding more RAM doesn't make much difference. For example, I have two P166 Windows 95 PCs: one with 32Mb, the other with 64Mb. The latter doesn't run any faster than the former. So, unless you want to handle exceptionally large data files such as 24-bit colour scans in Photoshop, consider 32Mb to be a reasonable maximum value.

One final thing to remember about Windows 95 and memory: under this operating system the concept of "free memory" is largely meaningless. No matter how much RAM is installed, Windows 95 will take most of it for cacheing purposes, leaving you with perhaps just a few megabytes of physical memory free. Finding this out when you've just filled your PC to the gunwales with RAM can be a little disconcerting, to say the least. Luckily, Windows 95's hold on physical memory is not very strong and it relinquishes memory as applications demand it.

How does RAM speed things up?

It's easy to understand how adding memory lets you run more or bigger applications, but how does it speed up the system? Well, extra RAM reduces Windows' tendency to use virtual or disk-based memory. All versions of Windows make use of a swapfile (on disk) in which to dump the contents of real memory, and whenever it runs low on memory it reads this data back into memory when next required, thus simulating more memory than it actually possesses. But reading and writing to a hard disk is far slower than reading and writing to RAM chips. It is this "disk churn" that slows down Windows so much.

Windows 95, which makes particularly aggressive use of virtual memory, exacerbates the problem. It uses dynamic virtual memory to swap programs from RAM to disk and back again as the amount of available memory changes. Load a large application or two (like Word 97 and Outlook) and Windows 95 will obediently make room for it by flushing everything else to disk: in return, you get to watch paint dry, even on a fast, 16Mb system.

The more memory you dedicate to Windows 95, the less swapping you will have to endure and the faster your system will appear. In fact, in certain situations, adding RAM can be a more cost-effective performance boost than fitting an overdrive processor.

The good news is that memory prices are at long last affordable. For more years than I care to remember, RAM seemed to be stuck at about £30 per megabyte, but in the past 18 months or so RAM prices have been in free-fall. In January 1996 the price of a 16Mb EDO SIMM was £480. One year later, you could buy them for one tenth of this (just £48). In the past few months,

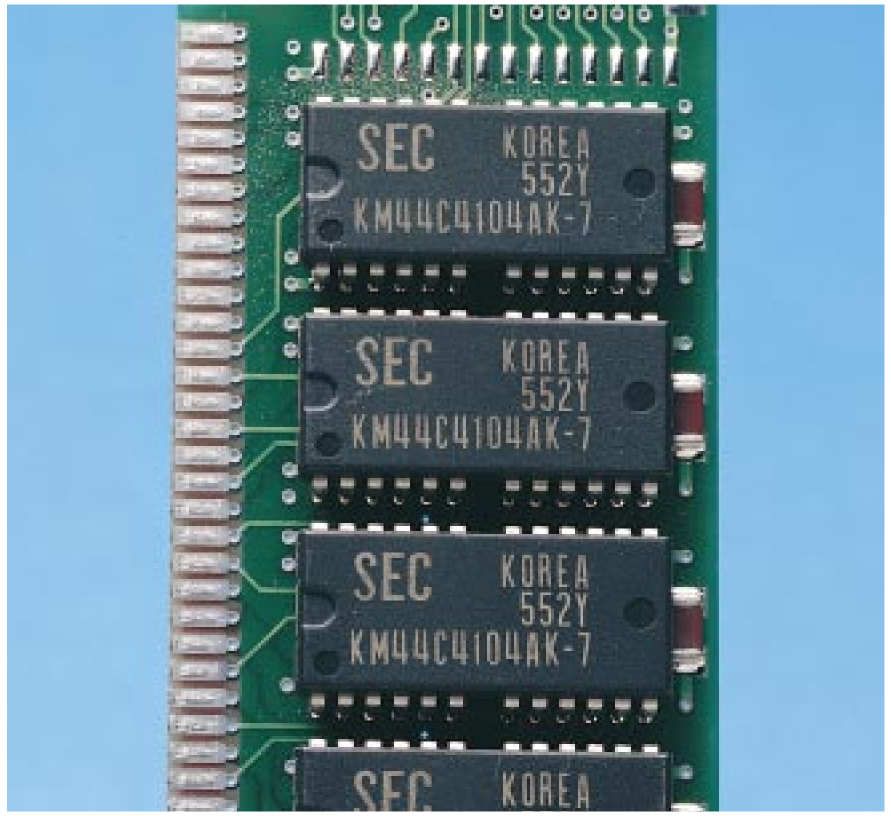

SIMMs are not always easy to identify. Here, the last digit (-7) represents 70ns

though, RAM prices have started to pick up and that same SIMM now goes for about £55. Even so, it's still excellent value for money and if you need extra RAM there's no time like the present to purchase it.

Types of memory

Upgrading memory on older 286 and 386 PCs used to be a nightmare, requiring dozens of tiny DRAM chips, each with 16 spindly legs, to be plugged in to equally tiny DIP sockets on the motherboard. These days, installing RAM rivals falling off logs in simplicity: memory chips are soldered onto Single In-line Memory Modules (SIMMs) which plug into special sockets on the motherboard. You can get up to 64Mb of RAM on a single SIMM, so fitting a large amount of memory requires you to clip just a couple of SIMMs onto the motherboard — easy or what?

SIMMs come in various sizes and speeds: 1Mb, 4Mb, 16Mb, 32Mb and 64Mb, with speeds of 90, 80, 70 and 60 nanoseconds (ns). The lower the number, the faster the SIMM. It's important to get the speed of the SIMM right when upgrading but unfortunately there's no easy way to make sure you're getting the correct SIMMs for your PC.

p270 ➢

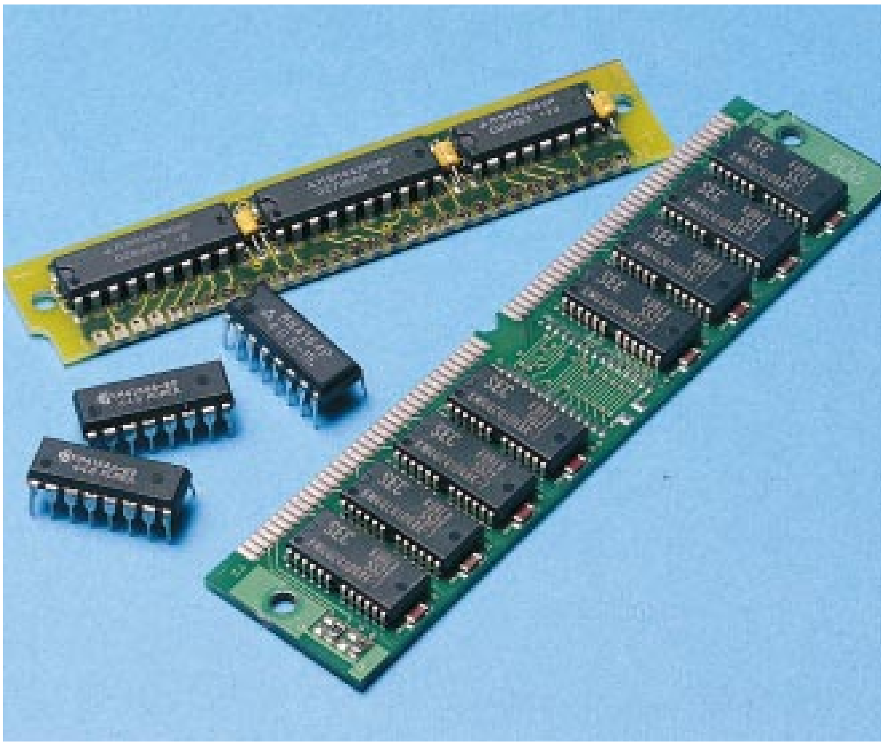

Here are the three major types of RAM: 72-pin (at the front), 30-pin (at the back) and the old-fashioned DRAM chips in the middle

Perhaps the simplest and most foolproof ways to determine the speed of the RAM you require are: read your manual to find the specifications for the system, call your dealer's technical support people, or actually examine the existing RAM in your system. Often, the last digits of the number on the chip or module represent the speed, and usually the chip makers use only two digits to indicate it. For instance, an 80ns chip will have an 80 (or –8), whereas a 150ns chip will have a 15.

Unless you have an eye to the future, don't buy RAM that's faster than you need. It won't hurt, but then neither will it speed up your system. This is because the system defaults to the slowest memory installed. By all means put fast RAM into a slow PC, but never put slow RAM into a fast PC.

Additionally, if you own a PC, avoid mixing chips of different speeds or from different manufacturers in the same memory bank of SIMM sockets: chip companies don't make their chips to exactly the same electrical standards, and the differences could confuse your system. If you must mix chip types, insert the chip with the slowest access time into the first bank. The rest of the chips will then run at the speed of the first.

Making contact

Another thing to watch is the type of contacts: you should choose the SIMM to match the contacts in the SIMM slot. If your SIMM slots have tin connectors, use SIMMs with tinned edge contacts. If they have gold, then use gold. There have been reports of corrosion where tin and gold have been used together.

30-pin vs 72-pin

The most significant difference between SIMMs is in their "pin-out": some have 30 pins, others have 72. The original SIMMs had 30 but for the past couple of years the 72-pin variety has predominated. Of course, the two types of SIMM are incompatible, physically and electronically.

Thirty-pin SIMMs are typically 8- or 9-bit devices with the ninth bit performing an integrity check on the eight data bits. Thirty-pin SIMMs are often advertised as x8 or x9. By contrast, 72-pin SIMMs are 32-bit devices (or 36 bits with parity) and are advertised as x32 or x36.

Many Pentium PCs dispense with parity checking these days (or provide a BIOS option to disable it). Modern DRAMs are far more reliable than their forebears and less liable to corrupt data, so this feature is less than essential. One benefit of disabling parity is that your computer will run faster. If your PC requires parity-checked SIMMs, make sure this is what you get.

Most 486 systems support 32-bit memory addressing, which means that with 8-bit (or 8+parity) 30-pin SIMMs you'll need a matching set of four identical SIMMs to create a bank of 32-bit memory. Systems typically have two or four banks of SIMM sockets, each of which must either contain four matching SIMMs or remain empty.

The requirement to completely fill 30-pin SIMM banks often poses problems when upgrading. Say you've filled both of your four SIMM banks with eight 1Mb SIMMs. Unfortunately, you won't be able to upgrade this configuration to 16Mb simply by buying another 8Mb of memory. To expand this configuration, you would have to ditch half your existing memory (four 1Mb SIMMs) and replace them with four 4Mb SIMMs (16Mb), forcing you to jump from 8Mb to 20Mb in the process.

Systems that use 72-pin SIMMs do not suffer from these restrictions. Because each 72-pin SIMM is a 32-bit device, each memory bank consists of a single SIMM socket on a 486. Pentiums, on the other hand, have a 64-bit data path and so 32-bit SIMMs have to be installed in matched pairs to establish a 64-bit data path.

Pentium motherboards with eight 72-pin SIMM slots are fairly rare and most come with just four slots, so the same "memory ceiling upgrade" problems can recur — even with modern motherboards.

DIMM – the latest thing

The latest version of the SIMM is the 64-bit Dual In-line Memory Module, or DIMM. We are just beginning to see these appear in the latest high-end Pentium PCs. DIMMs have 168 pins in two (or dual) rows of contacts; one on each side of the card. With the additional pins a computer can retrieve information from DIMMs, 64 bits at a time instead of the 32- or 16-bit transfers that are usual with SIMMs.

DIMMs are a bit longer and taller than SIMMs, with a couple of clamps at either end, and are fitted vertically into their sockets. PCs typically have just one or two DIMM slots.

Next month I will be examining the various memory technologies available and writing about actually fitting the memory.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address or via email at hardware@pcw.co.uk.

Upgrading by the book

Notebooks are an attractive option, but obsolescence can strike faster than crowsfeet or a receding hairline. Roger Gann takes a look at upgrading your personal power unit.

Notebook PCs are awkward and expensive to upgrade. This is largely due to one thing: all notebooks sold today are essentially one-off proprietary designs. As a result, they're all different and share very few common, standard parts. Sure, they'll have standard Intel CPUs and RAM, but beyond that, they will all use unique casing designs and motherboards.

Size precludes the incorporation of items like standard PCI slots or drive bays, or even relatively small features like SIMM slots. Often, the only way to upgrade a notebook is to make use of its PC Card slots or invest in a docking station that will feature expansion slots and a drive bay or two. You see, from an upgrading perspective, the architecture of a notebook is largely a closed book (*ha-ha!*).

Other factors conspire to make upgrading awkward and uneconomical. The pace of technological change in the notebook industry often makes it less expensive to buy a new machine after a few years than to upgrade an existing one: that new notebook computer you just pulled out of the box could be obsolete within a year. Technology in the notebook computer market is changing so rapidly that product life cycles last six months or less.

You might be hesitant to buy a notebook for fear of obsolescence. But take heart: there's a lot you can do to extend the life of your portable system.

A word of advice: if you want to upgrade, don't delay — do it as soon as possible after buying your notebook. As I've mentioned, notebooks have a depressingly short product life and are essentially proprietary designs. If you do want to buy an add-on for your notebook, a docking station for instance, you may not be able to obtain one 18 months down the line. And because of its unique design, you won't generally be able to use an alternative.

What's possible

So, *nil desperandum*... It is possible, with a little bit of ingenuity, to upgrade a limited number of notebook features. So what exactly is possible? Obviously, all modern notebooks will have one or two PC Card slots and these are the equivalent of the standard PC expansion slot. You can fit a wide variety of cards, from network cards, to SCSI cards to sound cards. And if you're feeling particularly flush, you can even fit tiny PC Card hard disk drives.

So the PC Card option has to be your first port of call when upgrading. It might be a relatively expensive solution, certainly compared to a similar upgrade performed on a normal desktop PC, but it's very quick and easy to install; particularly if you're using Windows 95, which has turned PC Card software support from a nightmare into a dream.

One upgrade which is common among desktop PCs but is almost without exception impossible on notebooks is the CPU upgrade. Unfortunately, most notebooks use CPUs that are hard-soldered to the motherboard and so cannot easily be removed. A few notebook brands do feature socketed CPUs and so, in theory, they are upgradeable, but I would be very careful about dropping in a faster Pentium, say, unless that notebook explicitly supported a faster CPU. Overheating is a very real problem in the close confines of a notebook, and simply fitting a faster Pentium is a recipe for notebook meltdown. Height restrictions (caused by the integral fan) will probably preclude you from fitting a Pentium overdrive, too. So one way or another, CPU upgrading is a no-no, as is upgrading other permanent parts of the system like sound and graphics.

OK, so what's left? The RAM and the hard disk: these are both upgradeable.

Memory upgrades

Perhaps the best way to improve your notebook's performance is to add more memory. Your applications will run faster and smoother, while more RAM will even extend the amount of time you get out of a battery charge.

Nowadays, many notebooks are supplied with 8Mb of RAM which is considered insufficient for running Windows 95: 16Mb is a more appropriate figure. Unfortunately, SIMMs and DIMMs are just too big for the latest slimline notebook designs. So, instead of using what has become a generic commodity, the SIMM, notebook users are forced to buy RAM cards made specifically for their notebook. The recent slump in RAM prices has taken the edge off notebook memory prices: for example, a 16Mb upgrade for a Toshiba Portege or Satellite Pro can be had for about £85; contrast this with what I paid to add a meagre 4Mb of RAM to my Dell, which set me back a cool £140 about two years ago.

Plan your RAM upgrade carefully, because the number of RAM expansion sockets in your notebook will be limited. You may have as few as two spare RAM sockets and if you buy, say, 8Mb and find it to be insufficient for your needs, you may have to junk your existing RAM to install 16Mb. What I' m saying is: if at all possible, buy more RAM than you actually need because you're bound to need it next year.

Installing the memory modules is easy, much easier than the corresponding task on a desktop PC — normally, you remove a clip-on panel, either above or below the keyboard, plug in the modules, replace the cover and power up the notebook. End of story. But one word of caution: these memory modules often use quite flimsy PCBs, so treat them gently.

Hard disk upgrades

Notebooks are no different from desktop PCs when it comes to storage; there's never enough! Luckily, swapping a hard disk is not an insurmountable problem. The vast majority of notebooks use 2.5in IDE hard disks, which are just smaller versions of the familiar 3.5in IDE drives we see everyday. Apart from size, these tiny drives differ from their larger peers in that they have a single connector that carries both power and data. To all intents and purposes they are identical, and installing a larger IDE drive in your notebook is in theory no different to that of installing one in a desktop PC. Sadly, the practice is somewhat different. The problem lies in gaining access to the drive itself.

This is not a problem if the drive itself is removable and slides out of the notebook casing after a catch is released. All you have to do is dismantle the plastic shell surrounding the drive, ***very carefully*** disconnect the thin ribbon cable, extract the old drive and replace it with the new one. With the drive installed, you then load CMOS Setup to get the notebook to auto-identify the drive, and away you go.

Note that some older notebooks have BIOSes that can't handle large drives, i.e. those bigger than 540Mb. And as it's now virtually impossible to get 2.5in drives smaller than 810Mb, it means that older notebooks cannot easily be upgraded. If your notebook has a Flash BIOS, it's worth checking out the manufacturer's web site in case BIOS upgrades are available. For details of commissioning the drive, installing an operating system and so on, refer to my Hands On column in *PCW* August '96.

In many notebooks, the drive isn't removable and is buried inside the notebook. Accessing these types of drive is hard work but by no means impossible. Be careful: if there's any warranty left on your notebook, you'll probably invalidate it by doing this. Notebooks are not easy to open. I can't begin to describe the various dismantling permutations for all the different notebooks, so what follows is only a general guide.

1. The plastic casing will be held together by a number of tiny screws. Some will be visible and others won't because they'll be hidden! For example, the rubber feet on the underside of a notebook often conceal screw holes. Likewise, adhesive labels above the keyboard may disguise the presence of screws. The first job is to locate all these screws and remove them, making a careful note of which screw went where.
2. The casing top should now separate from the bottom. You may have to unplug a couple of thin, delicate cables that link the two halves: examine the plug at one end and slide down the sleeve to release the cable. Watch out, too, for a tiny micro-switch that is depressed when you close the lid; this is vulnerable.
3. The old hard disk will probably be held in a metal clamp, secured by a pair of screws. Release the drive and carefully pull the ribbon cable off the rear. Install the new drive and, as they say in all the best car repair manuals, reassembly is the reverse of disassembly!

A final note: transferring the contents of the old hard drive to the new is a problem in notebooks. VisionTek offers an interesting solution to this problem — Drive Exchange. It consists of a Type II PC Card plus cables and software. You plug the new drive into the PC Card, then insert this into the notebook. The software takes a complete copy of the old drive. When it's done, you simply install the new drive.

Screws can be hidden under adhesive labels

Alternatives

You don't necessarily have to go down the PC Card route to upgrade: there are other ways to connect peripherals to your notebook, especially physically-large devices such as CD-ROM drives or tape streamers.

Parallel port devices

The parallel port is a useful fallback if all your Card slots are spoken for. For example, you'll be able to hook up an external CD-ROM drive — the Micro Solutions' Backpack 4X CD-ROM (around £140 on the street) is a quad-speed, 250ms, 3.5lb, mains-only drive which is solidly constructed and works with both standard and the newer, faster, enhanced parallel ports. Installation with the included software takes less than a minute. A printer cable attaches to the throughput connector on the back of the drive. An 8-speed version is available (about £240) and both models can have sound, too, which adds about £80 to the final bill.

Docking Stations

A desk-based docking station adds desktop PC-like expansion opportunities to notebook computers, as they often feature expansion slots and drive bays to which you'd fit ordinary common-or-garden expansion cards and CD-ROM drives. The downside is that docking stations aren't cheap and can cost several hundred pounds. They are not designed with too much portability in mind, either.

Direct Cable Connection

This is a low-cost way of networking your notebook to a desktop PC to enable you to access its hard disk or CD-ROM drive as though they were directly connected to your notebook. The Windows 95 Direct Cable Connection lets you make what is, in effect, a network connection to another PC via a parallel or serial cable to access shared resources on that machine, such as a CD-ROM drive. Direct Cable Connection is part of Windows 95 and is free — all you need is a cable, which will set you back around £15. DOS and Windows 3.1x users can use the slightly less sexy InterLink facility instead.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* at the usual address, or via email at hardware@pcw.co.uk.

VisionTek 0181 561 5533

Fasten your seatbelts...

...let's go for a drive. Roger Gann gets in gear and sets off to explore CD-R, the storage disc set to leave the CD-ROM drive standing. Choosing, installing and running, it's all here.

With the first DVD-ROM drives now on sale at more reasonable prices, the writing is on the wall for the humble CD-ROM drive. And with so little profit margin left for the CD-ROM drive makers, their attention will switch to producing something more profitable; I think *that* something is the recordable CD-ROM drive (CD-R).

CD-Rs have been around for years, having first appeared in the rarefied world of hi-fi at astronomical prices. They eventually made the crossover to the PC market, but then prices came down a bit. The original HP SureStore 4020I, which then cost £700, today costs under £425. "Bare" CD-R drives can be bought for much less (some, at street prices below £300).

Cooking your own CD-ROM discs is now fairly affordable. But who needs home-brewed CD-ROMs? Anyone with a modern, well specified PC. CD-R is cheap bulk storage: at around £7 each, a 650Mb CD-ROM disc works out at about 1p per megabyte. CD-R is a good choice as a storage or backup device for several reasons. Firstly, the ubiquity of CD-ROM drives; your disc will be readable on many PCs. This also makes CD-R an excellent medium for transferring large files; you can even use it to back up expensive CD-ROM based reference works. Secondly, unlike tape, CD-R is a random-access device, which makes it fast to get at archive material. The discs are more durable, too: they have a longer life span than tape cartridges, they're physically durable, and they can't be wiped by coming into contact with, say, a magnetic field.

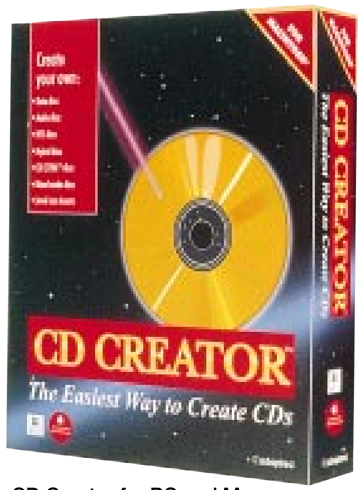

CD Creator for PC and Mac

It's true that CD-ROM drives aren't storage speed demons compared with modern hard disks, but they're fast enough for running infrequently used applications and for accessing historical data. The CD-ROM is thus a perfect medium for today's big multimedia data files. Also, you can store just about any form of data on a CD-ROM: you can mix and match video, Photo-CD images, graphics, sound and conventional data on a single disc.

Choosing a CD-R drive

There's a reasonable choice of CD-R drives available from companies like Hewlett-Packard, Mitsumi, Sony and others. Prices start at under £300 (ex VAT), rising to around £600 (ex VAT). You can choose between external models which sit on your desktop, or internal models that sit in a spare 5.25in drive bay. Although they have the special ability to write to CD-R blank discs, these drives otherwise behave like ordinary CD-ROM drives. Almost all are SCSI devices so you'll need a SCSI host adaptor card, but parallel port drives are also available and work surprisingly well.

Perhaps the most important feature of a CD-R drive is the speed at which it can write, as "burning" a full 650Mb CD-R disc can be a slow process. For example, audio CDs are 1x, 150Kb/sec devices and can hold 74 minutes of music, so this is how long a single-speed burn will take (a 2x will take half this time, and so on); the faster the CD-R, the quicker you can burn your discs.

Hewlett-Packard P SureStore CD-Writer 6020

Although CD-ROM drives are now hitting 16x, with peak data transfer rates in excess of 1,800Kb/sec, you won't find any CD-R drives offering this kind of performance. While most can read at 4x or 6x, the vast majority of CD-R drives are 2x writers. That is, they can write to disc at a speed of 300Kb/sec. Some CD-R drives can manage to write at 4x (600Kb/sec) but these are about 50 percent dearer and you need a well-specified PC to keep a 4x CD-R happy. Finally, check just what mastering software is included in the deal.

Installing a CD-R drive

If you've installed an ordinary SCSI CD-ROM drive, installing a CD-R version won't come as a great shock as it will install in a similar way. If you already have a SCSI host adaptor installed, all you have to do is connect the new CD-R drive to a spare connector on the SCSI cable if it's an internal drive, or plug it into the socket on the SCSI host adaptor card if it's an external), having first given the drive a unique SCSI ID number (say, between two and six) and turning termination on if it's the last device in the SCSI daisy-chain. An external CD-R drive is a good choice if you plan on using it to back up different PCs.

Some CD-R drives will be bundled with their own private SCSI card and there's a compelling reason for installing these even if you have a SCSI host adaptor already installed. The second card ensures the CD-R drive gets the highest possible data throughput. It's important to use a bus-mastering SCSI host adaptor as this will improve throughput, which is the name of the CD-R game. If you have a high-spec SCSI host adaptor to start with, you can get away with just using the one card.

The simplest CD-R drives to install are those which make use of your PC's parallel port. These drives install in minutes: plug in the cables, power-up the drive, install the software drivers and you're away. I keep emphasising the importance of good data transfer rates and that goes double for parallel ports; you must have a PC with an EPP or ECP parallel port. This is important because a standard/4-bit parallel port may not have a fast enough data transfer rate to keep the CD-R satisfied. Parallel port CD-R drives are available from HP and others.

Hardware requirements

Unlike CD-ROMs, you can't fit a CD-R drive into any old PC and expect it to work consistently. Successful CD-R burning needs a PC capable of providing a fast, consistent, data flow during the burning process. The biggest problem you'll come across in CD-authoring is running out of data during a burn. A CD-R disc has to be written to in a contiguous, track-by-track fashion, writing the data in a smooth, uninterrupted stream. The write session has to be closed properly. If, for any reason, there's a glitch in the data flow during the burn and the unit is momentarily left with nothing to write, the data buffer empties and a "buffer under-run" occurs. That £7 CD you're trying to write to then becomes useless and you'll have to start again.

So what's the solution to this problem? A good start is a fast, well-specified PC. Some CD-R makers recommend a 486 with 8Mb of RAM or better but in my experience it's the "better" that you need. I'd say nothing less than a P90 with 16Mb RAM, especially if you want to do track-at-a-time or multi-session recording. Make sure you buy a CD-R drive with a decent-sized on-board buffer of about 1Mb. You'll need a big, fast, hard disk, too. Although you can successfully burn CD-Rs using Enhanced IDE/PIO Mode 4 hard drives, most makers recommend a SCSI hard drive with at least 1.3Gb free; this is the minimum space needed to create a full 650Mb CD, especially if your mastering software needs a disk image of the CD-R to be created first.

Sony Spressa 920

For serious CD-R authoring you're looking at fitting a second SCSI hard disk, dedicated solely to the CD-R mastering task. Avoid drives which feature thermal recalibration on-the-fly as this interrupts data flow. It's best to use so-called AV (audio-visual) drives, which are designed to deliver fast, sustained transfers. Another tip to ensure a smooth data flow is to defrag your source drive first. During the burn session it's probably better *not* to put Windows multitasking to the test, so don't run anything else at the same time. And run ScanDisk too, just to make sure all the files can be read. If, after this, you still get continuous buffer under-runs, you'll probably have no alternative but to drop the write speed, from quad- to double-speed, or from double- down to single-speed.

Buffer under-runs aren't the only reason discs get ruined. The final, lengthy stage in the recording process closes out the disc with a table of contents which consumes about 13Mb of space per session. If the CD-R drive and software aren't set up properly, or if the proper SCSI termination isn't in place, it's possible to get through most of a recording session but have the closing process fail and the disc ruined at the last moment. Beware if you're intent on making good use of a CD-R's multi-session capabilities: the 13Mb overhead soon mounts up if you burn many sessions onto one disc. It's for this reason that multi-session CD-R is unsuitable for incremental daily backups; you may back up just a few megabytes of data but then sacrifice another 13Mb to close the session.

Mastering software

When buying a CD-R pay close attention to the mastering software that comes with it. This is the utility that lets you arrange and format the data to be burned onto a CD and is as important as the drive itself. Just as fax modems are bundled with "lite" versions of full-blown fax programs, so the odds are that you'll come across a lite version of one of the mastering packages sold by Corel, Incat Systems and Elektroson in the box. These will provide enough functionality to let you perform basic mastering chores but will lack the bells and whistles found in full versions. Mastering involves marshalling many files, so these packages will include File Manager-like interfaces for routine tasks like backup and archiving. The full versions will offer multi-session and full audio CD capabilities. They will all support the ISO 9660 standard, a CD-ROM data format readable by PCs running DOS or Windows, by Macs running System 6.x or 7.x, and by almost every flavour of UNIX. This is a limited standard in some ways, but Microsoft's Joliet CD-ROM formatting standard expands on 9660 by allowing semi-intelligent truncation of long filenames, and it is beginning to be supported by the new generation of CD-R software.

Some packages will first create a physical image file of the entire CD-ROM on the hard disc; a time-consuming process. This is a complete, bit-for-bit mirror image duplicate of all the files to be burned to a CD-R disk in a recording session. If you have a fast PC, you might be able to dispense with the time and bother of a complete physical image by using a virtual image file method instead. This is smaller, and carries a set of pointers to the files on your hard disk to be sent to the CD-R drive.

Finally, if you're bothered about buffer under-runs, most packages allow you to run a dress rehearsal of a burn session but without actually writing to the disc. This is a slow process, but could save you time and money in the long run.

PCW Contacts

Roger Gann can be contacted by post c/o *PCW* or via email at hardware@pcw.vnu.co.uk.

CD Creator; Adaptec 01276 854500
SureStore CD-Writer; HP 0990 474747
Spressa 920 CD Writer; Sony 01932 816000

Still waters…

… run deep: Intel has been quietly working away on processor developments, some with curiously watery codenames. Roger Gann anticipates a deluge of new products to come.

A river runs through it — Intel's future processor strategy, that is. For some time now, Intel has codenamed its future processor development projects after local rivers. The watery names have replaced the traditional Pnn scheme that Intel used for a decade.

Last year wasn't particularly busy for Intel. Even though it sold Pentiums by the shipload and hoovered up the cash, its new product release schedule was quite relaxed. After the introduction of the 150 and 166MHz Pentiums in January '96, the only new products were the 200MHz Pentium (July) and the 150MHz Mobile Pentium.

This year promises to be hectic. The launch of the Pentium with MMX Technology will be swiftly followed by OverDrive versions and a special version for notebooks, and we'll see the launch of an MMX-enabled version of the Pentium Pro. So this month I'll take a break from the toil of upgrading and instead gaze into my crystal ball at Intel's road map of future processor products.

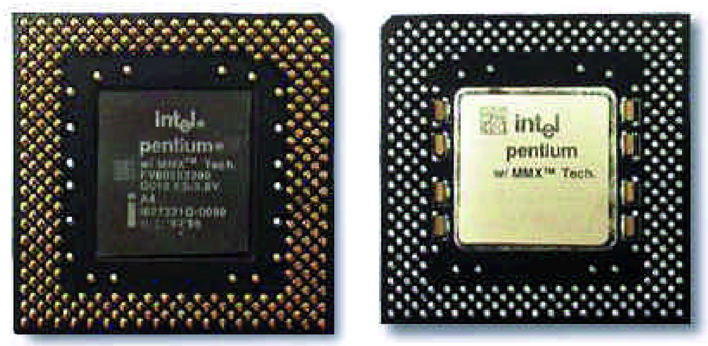

The new MMX Pentiums are pin-compatible with their predecessors, but to gain maximum benefit you will need MMX-enabled software

Klamath

The Pentium now has the benefit of MMX technology, but what about the Pentium Pro? You won't have long to wait: Klamath is a cost-reduced version of the Pentium Pro intended to bring Pro performance to mainstream PCs this year. Although still based on the P6 core used in the Pentium Pro, Klamath will add the MMX multimedia extensions already seen in the new Pentiums. Klamath will initially run at 233MHz and later at 266MHz. These new, faster clock speeds are made possible by Intel's new 0.28-micron CMOS process. This will be slightly faster than the 0.35-micron BiCMOS version of the Pentium Pro, allowing Klamath to outperform it but at a lower cost. The new processor will come with a large, 512Kb, L2 cache.

Klamath's most visible difference will be its packaging: it won't resemble existing Pentium and Pentium Pro chips. The new processor will be housed in a new package, the Single Edge Connector (SEC) cartridge, a plug-in design that will require a complete motherboard redesign and which will function in much the same way as a CPU daughtercard. The SEC cartridge will probably ship in a metal case that provides for thermal transfer as well as electro-magnetic shielding.

Although the Pentium Pro was the most powerful processor Intel makes, it added little to its profits: Intel shipped fewer than three million Pentium Pros during 1996, compared with more than 60 million Pentiums. For Klamath to sell in the vast quantities that Intel would like, it must be cheaper to fabricate than the existing Pentium Pro.

The Pro broke new ground by coupling the microprocessor chip with a secondary or Level 2-cache chip in a single package. This solution, although providing good performance, was expensive and consumed too much of Intel's manufacturing capacity. The SEC design allows Intel to employ cheap, off-the-shelf, conventional SRAM chips (normally used for external caches) for its "internal" L2 cache; a move that offers considerable cost savings. The downside of the SEC is that its new packaging means there'll be no Pentium Pro MMX OverDrive upgrade. If you want Klamath, you'll need a new motherboard, too.

Initially, the SEC daughtercard will be designed to direct I/O traffic. Later on, Intel is expected to incorporate microcontrollers, DSP technologies and even some analogue devices into the card, making the prospect of full-screen, full-motion MPEG-2 or DVD-playback a distinct possibility.

It's now looking like Klamath will ship towards the end of Q2 this year, later than originally expected. So, it's more than likely that AMD will beat Intel to the draw by launching its rival K6 processor early in April. A minor fly in the ointment in the impending launch of Klamath is that its new motherboards will use the old 440FX PCI chipset which lacks support for advanced features such as SDRAM, Intel's Advanced Graphics Port (AGP) and 33Mb/sec EIDE transfer rates. An enhanced chipset, likely to be called the 440LX, won't be available until the second half of this year.

Two types of user will derive the most benefit from Klamath. First in line are the users of 32-bit operating systems such as Windows NT 4.0. These business users will see a big performance boost from the chip. Secondly, home users will just love Klamath as it will be great for 3D gaming. The Pentium Pro's floating-point unit is much faster than the Pentium's and thus most 3D programs will be dramatically faster.

Klamath will be a desktop-only processor: there will be no notebook version of the MMX-enhanced P6 processor, simply because it runs too hot. It's likely that Intel will skip a beat and jump to a notebook version of its next-generation processor, code-named Deschutes. So, although Pentium MMX notebook processors will be introduced and enhanced this year, there will be no Intel Pentium Pro-class notebook processor until 1998.

Deschutes

The Klamath architecture will support up to two processors and will be aimed at the workstation and entry-level server markets. The existing Pentium Pro supports Symmetric Multi-Processing (SMP) to the tune of four processors, maximum. This limitation will be overcome by Deschutes, a version of Klamath specifically designed for network servers. Intel is rumoured to be developing an eight-processor motherboard using the Deschutes architecture. And, some reports claim that the new architecture will feature a new bus interface that will allow manufacturers to build systems with up to 32 processors.

The processor itself will be a refinement on Klamath and will probably be a shrink of the Klamath design using the forthcoming 0.25-micron process. The die size of Deschutes will shrink markedly as a result: from 690 mils (0.69ins) for Klamath, which will use 0.35-micron technology, to 400 mils (0.4ins) for Deschutes; enough to allow high-volume, low-cost production.

Similar to the Klamath design, Deschutes will offer a CPU based on the Pentium Pro core but with separate components, such as a larger Level 2 cache and a modular CPU-to-cache bus on an SEC daughterboard. CPU speeds will range from 266MHz, 300MHz and 333MHz for Deschutes.

Some sources have predicted that Intel may also be moving to develop a 75MHz bus to run the Deschutes. Although that would require a new motherboard and related components, they noted that the faster bus might also be necessary. In common with Klamath, Deschutes will be configurable, capable of handling multiple cache memory types such as static RAM or multibank DRAM, the Advanced Graphics Port controller or a communications director for managing high-speed communications. Deschutes might appear towards the end of the year but 1998 is a far more likely timeframe.

In order to realise the full impact of Deschutes, you'll need a network operating system that can handle lots of processors. The current release of Windows NT Server will scale to eight processors but will need to support even more than this in future to take full advantage of Deschutes.

Katmai

Okay, so Katmai isn't a river, it's a volcano. But allow me a little artistic licence here!

Katmai will most likely surface late in 1998 or early '99, about six months after Deschutes. It will essentially be a Klamath MkII, with revamped and enhanced MMX technology, dubbed MMX 2 in the US. This version of MMX technology will significantly improve 3D graphics performance, making it an absolute must-have for serious gamesters. It concentrates on boosting 3D graphics performance — something not specifically addressed by the first version of MMX.

To recap, MMX cleverly re-uses the CPU's floating-point registers for its own highly optimised multimedia tasks. Obviously, programs can't use both FP and MMX instructions within the same routines as both share the same register set. This is rarely a problem, since most programs don't use FP at all. Those that do, typically use these calculations to generate data, while MMX is typically used in separate routines that display data.

For 3D graphics, Intel recommends that geometry calculations remain in floating point while MMX is used to accelerate 3D rendering routines. Each time the processor swaps the FP registers between FP and MMX, a time-consuming "context-switch" is required. Hence Intel's admission that 3D rendering doesn't benefit dramatically from MMX, although applications that take advantage of Microsoft's Direct3D programming interface will see some improvements. MMX2 will have additional instructions and a larger cache compared with the first release of MMX.

Merced

Also referred to as P7, Merced will be the first CPU to be released which supports IA-64, the new 64-bit architecture jointly developed by Intel and Hewlett-Packard under the codename Tahoe. Merced, which appears to have slipped into 1999, is likely to serve primarily the workstation and server markets, while Willamette, a cut-down version of Merced, will drive high-end personal computers.

It is early days, however, and very few hard facts are known about this processor. Intel has been tight-lipped about this particular project. It seems that the processor will comprehensively address many of the limitations of its 32-bit forebears. The chip is a joint Intel/HP venture and will not only support x86 binaries but also HP's PA-RISC binaries for its UNIX workstations.

Presumably, x86 support will be provided by way of emulation, in much the same way that PowerPC Macs can run 68K code. The key in both cases is sheer speed; the PowerPC was so fast anyway that the speed penalty imposed by emulation was unnoticeable. The same will probably be true of Merced.

Tillamook

Conspicuous by its absence from the MMX launch was the 200MHz version of the Pentium MMX for notebooks. Don't worry, though: Intel is planning something special for later in the year.

Codenamed Tillamook, this CPU will be among the first to be built on an Intel-developed pop-out module for notebooks, called MMO. It will include a processor, Level 2 cache, a voltage regulator, a clock and a PCI chipset called the 430TX. Instead of plugging a new processor onto the motherboard, users will plug the MMO processor module into the MMO socket.

So, for the first time, notebook users will have a CPU upgrade path. Fabricated using the new 0.25 micron process, the new CPU will not only be smaller than the Pentium MMX on which it is based but it will draw significantly less juice; down from 2.45v to 1.8v. A 233MHz version will follow.

And just for the record, Tillamook is a town, not a river. So there.

PCW Contact

Roger Gann can be contacted by post c/o *PCW* at the usual address, or via email at hardware@pcw.vnu.co.uk

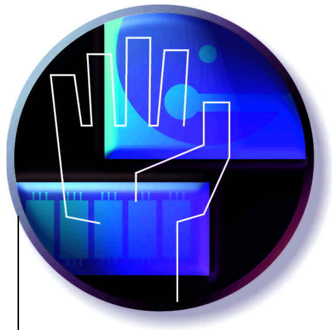

Fraught in the net

Getting connected to the internet can be a risky and potentially expensive business. Roger Gann takes the worry out of it with a back-to-basics brief on what's what and where to get it.

I guess a reasonably high proportion of *Personal Computer World* readers regularly access the internet from their PCs at home, using a V.34 modem and a normal phone line. Such an act would have been rocket science a few short years ago, but today it's a routine event. For single users, this method, a modem plus a conventional phone line, is the most cost-effective internet access solution. But what if you wanted to give internet access to a group of users, say those on a small network? Well, there's absolutely nothing to stop you from scaling up the single-user solution and applying it to everyone on the network, buying them all a modem, giving them all a phone line and their own ISP account. It's feasible, but it's a less than desirable solution. Not only would it be a nightmare to configure and administer, but it would also be expensive to implement.

No, a much more elegant solution is to integrate internet access into your network. The good news is that, depending on your precise needs, you may not have to fork out for expensive new kit in order to do this. In fact, for most companies the expense will be relatively trivial. There are basically two choices available: email-only or full-blown internet connectivity.

Router alternative

Installing TCP/IP on a lot of workstations is enough to make a network administrator cry. Luckily, Bay Networks has an alternative internet access solution for those wedded to the NetWare standard. Instant Internet 3.1 is an IP/IPX gateway server package, a hardware/software combination that allows NetWare users to access the internet without having to worry about configuring TCP/IP on top of an IPX stack. It effectively permits workstations to access the internet using just the NetWare IPX stack — only the Instant Internet server uses TCP/IP. The downside is that it's not particularly cheap — a 50 concurrent user ISDN version is priced at £4,329 and it can be awkward to configure. www.baynetworks.com/

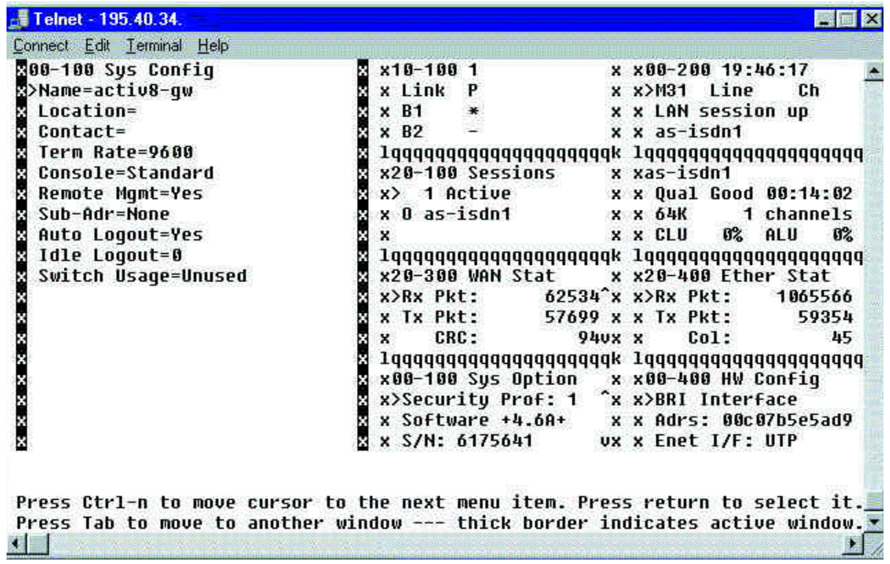

Complicated or what? You can telnet in to your ISDN router once it's been configured, but as you can see, configuration isn't straightforward

Email only

I suspect most companies, even small ones, running even modest networks, rapidly appreciate the benefits of email, often provided as a standard feature in their NOS (e.g. Windows 95 or Windows for Workgroups 3.11) and soon find themselves making heavy use of internal email. So it makes a lot of sense to consider extending the existing email connectivity to embrace the internet — sending an email is probably the cheapest method of global communication available.

For email purposes you don't need instant or continuous access to the internet, nor do you need a particularly fast connection: emails tend to be reasonably compact and the throughput offered by a 33.6Kbps modem will be perfectly adequate for most email traffic. However, if you get much email with large attachments, a faster connection is maybe worth considering. So, your hardware costs will be negligible. You'll also need a mail system that supports internet mail. A good, ubiquitous example of such a system is good old MS Mail, which offers basic email services and is supplied free with both Windows 95 and Windows 3.1x. It's possible to configure MS Mail to initiate a connection to your internet service provider via the modem at regular intervals, say every hour or so, to deliver new outgoing emails and to receive new incoming messages. To do this you simply add the optional "Internet Mail" service to the Exchange Inbox client in Windows 95.

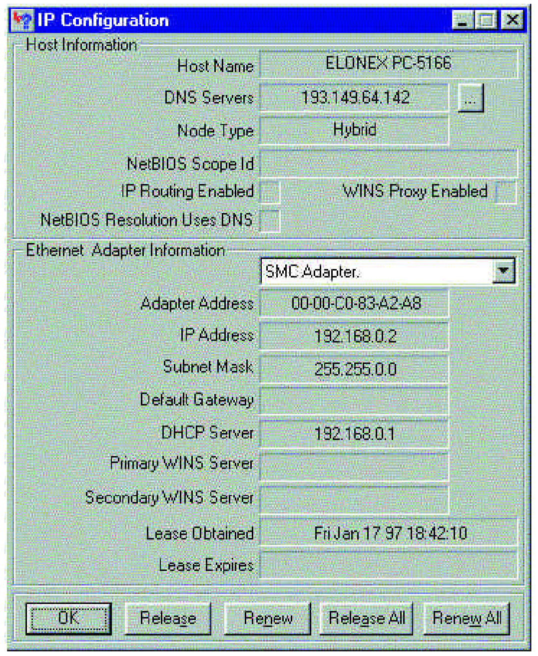

This undocumented goodie is supplied with Win95 and displays your PC's IP configuration, useful if you use DHCP on the server to allocate IP addresses

Such a setup is just dandy for standalone machines collecting their own personal email but it can't handle multiple internet email addresses on the network. What MS Mail needs is a bolt-on external mail package, one that can cope with POP3 (Post Office Protocol 3) and SMTP (Simple Mail Transport Protocol) mail systems offered by most ISPs. With such an add-on, every time MS Mail connects with the ISP's mail server, it downloads everybody's mail in one lump and then sorts it, placing the appropriate messages in the appropriate mailboxes. The result is transparency between internal and external email.

There are a number of these add-ons on the market. The European Microsoft Windows NT Academic Centre (EMWAC) based at Edinburgh University has developed EMWAC Internet Mail Services for Windows NT (along with a whole slew of other useful internet goodies). IMS is a suite of server programs which lets you use Windows NT as a mail server for internet mail. IMS isn't perfect — it requires Windows NT, of course, and can be a bit daunting to set up, but it does have the redeeming feature of being freeware. All the EMWAC goodies can be found at emwac.ed.ac.uk/.

Another alternative worth investigating is SLMail, developed on Bill's doorstep by Seattle Labs. This is available for all the Windows platforms, and a Windows 95 version (limited to six accounts) is available free from www.seattlelab.com/ (and is also on this month's cover CD). Like IMS, SLMail95 offers a complete implementation of the Internet SMTP and POP3. With SLMail installed, your PC can function as a post office for popular internet mail clients such as FreeAgent, Pegasus, Eudora, Navigator and MS Mail/Exchange (with the Internet Mail add-on). The full-blown Windows 95 version costs $189: the Windows NT version costs $325. So, the software cost needn't be great either.

So we've now got our internet mail software in place but there's still one variable to determine — does your existing email account with your internet service provider permit multiple email addresses? The type of email account varies from ISP to ISP. Some, for example, allow up to 99 mailboxes from the one common-or-garden "tenner a month" account. This is certainly a cheap and cheerful solution but the principal drawback is cosmetic, the lack of a personalised email address — it would be, say, rgann@company.demon.co.uk rather than the more impressive rgann@company.co.uk. Private domains cost extra of course, and you'll have to talk to your service provider about what types of account it has on offer and proceed from there.

In terms of hardware configuration, a network email system will be similar to a normal standalone internet connection: you'll just need a phone line and a fast modem. Commissioning the connection would be much the same too. Note that as only the server accesses the internet, it will be the only machine on the network that needs to run TCP/IP. The workstations would continue to collect their email in the usual way, via NetBEUI or IPX. This makes life much easier for the network administrator and also prevents the workstations from accessing the internet during connections.

Full internet connectivity

In some ways granting complete internet access to the network isn't much more complicated than setting up global email. Just two extra things are required — an IP address for each workstation (which will almost certainly require a "corporate"-style access account with your ISP) and rapid access to the internet. A permanent leased-line connection offers instant connection while a V.34 modem takes about 20 seconds to connect. The former is expensive and fast; the latter is cheap and slow. A popular compromise solution is offered by ISDN which offers the best of both worlds — a fast 64Kbps service (with the possibility of 128Kbps) with the economy of demand dialling, plus fast connections, typically just a few seconds, certainly a lot faster than the time taken by your browser to load! For occasional internet traffic, ISDN is very cost effective, offering the speed of conventional leased lines at a fraction of the cost. However, once you spend more than about four hours per day online, a leased line becomes viable.

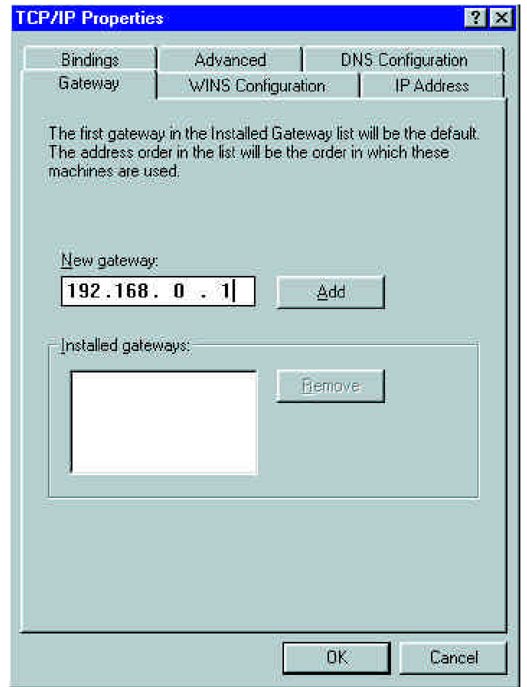

The ISDN router now becomes your gateway to the internet — so don't forget to add the IP address of the router to the Gateway tab of TCP/IP properties

You would use an ISDN terminal adaptor or TA to connect a single PC to an ISDN but for networks you'd fit a single ISDN router instead. This device sits anywhere on the network, watching the TCP/IP traffic that

p296 ➢

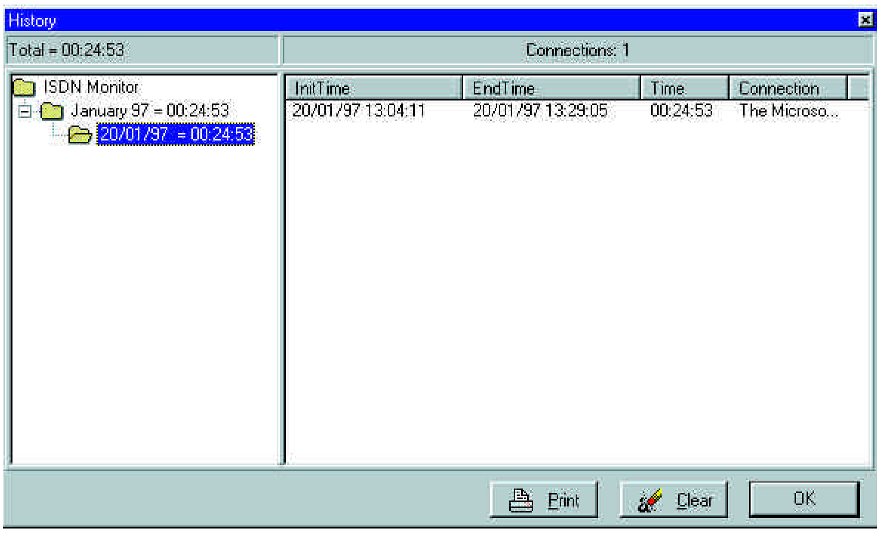

Keeping a tight control on ISDN line charges is crucial. This little utility, ISDN Monitor, specifcally monitors ISDN usage. It's on the cover disk!

passes by. It ignores traffic destined for IP addresses on the internal network but as soon as it detects a network packet aimed at an external internet address, it initiates a call to the ISP and sets up an internet connection. The line is dropped when the connection is no longer required or times out after a specified period of inactivity.

Users are spoiled for choice when it comes to ISDN routers, although you shouldn't necessarily buy the cheapest. Prices vary from about £750 to £1,500, which may sound dear but don't forget to divide the price by the number of users to get the cost into perspective. How do you go about choosing these exotic bits of kit?

ISP at your service

A good place to start is your local friendly internet service provider — what routers do they recommend or support? For example, many ISPs use Ascend routers at their end and so to simplify things they naturally recommend that you use something like an Ascend Pipeline 25 or 50 at yours. However, there are other equally good if not better alternatives, for example the £1,415 Gandalf XpressConnect LANLine 5250I or the excellent £750 Shiva AccessPort. If at all possible choose one that has an analogue phone socket (or two) — this lets you use the ISDN line as a normal phone line for fax machines or ordinary modems, for example.

Connecting an ISDN router to the network is the easy bit: you just plug the ISDN lead into the ISDN phone socket and plug in the network cable. Beware — this is usually a 10Base-T RJ45 which is the same as the ISDN (or "WAN" socket), so don't get these two mixed up. Routers tend to be disarmingly small boxes, no larger than a modem, and so can often be tucked away in a comms cabinet or next to the server.

Configuring an ISDN router is a different matter and can be a particularly daunting task. They come with a range of tweakable settings wide enough to make any network techie jump for joy. However, if you've cut your internet teeth by manually setting up a connection using, say, Windows 95 Dial-Up Networking, you'll already be halfway up the learning curve. Even communicating with the router can be painful. Once configured, most routers let you telnet in to them for management and configuration tasks, but not usually the first time around.

Most routers therefore typically feature a serial port which lets you hook up a dumb terminal or a PC running a terminal program which lets you access the router's text-based configuration menu. One router I installed wouldn't even talk to good old Windows Terminal: I had to download ye olde Procomm Plus for DOS, circa 1987AD, before it could be configured. Luckily, some routers are a bit easier to configure — the Digi Retoura ST has an LCD control panel and keypad for direct configuration, while the latest Ascend Pipelines have Java applet-based configuration firmware, making them configurable from a browser. Most user-friendly of all is probably the Windows-based Shiva Monitor software supplied with the AccessPort.

Once connected to your router you'll need to enter such crucial details as the ISDN number it has to dial, the IP address of the router or gateway at the other end, the IP address of this router, the DNS IP address, logins, passwords and security levels. This is just a basic list of data you can enter; more sophisticated users will want to configure things like the PAP and CHAP security protocols, data compression and bandwidth on demand.

Checking that the installation works is fairly straightforward and most include comprehensive diagnostics in the firmware as standard. Once the installation has passed this test you can use conventional internet tools, such as ping or even a web browser, to make sure that calls are being initiated and terminated as required.

The final step is to roll out TCP/IP to the workstations, a relatively easy task if you're using Windows 95, less so if you're still using Windows 3.1x. Your internet service provider will have allocated you a range of IP addresses, say 192.168.0.1 to 192.168.0.25, and you then manually dole out one of these IP addresses to each workstation. This tedious chore can be greatly simplified if you're running Windows NT on the server, as you can install the DHCP (Dynamic Host Configuration Protocol) service. This lets you define a pool of IP addresses that are then dynamically allocated to each workstation as they log on to the network. So instead of specifying a "static" IP address for each workstation, you'd obtain the IP address automatically from the server.

Silence is expensive

A final word of warning: ISDN is completely different to the usual PSTN phone system we hook our modem up to, and it's easy to rack up excessive ISDN phone bills unless you keep a close eye on it. The main problem is ISDN's inscrutable "silence": the absence of the usual audible clues such as dial tones and negotiation whistles and noise means you're entirely dependent on software to tell you if the line is up and data is flowing correctly, which is a scary prospect. If you don't want phone bills that resemble telephone numbers (ho-ho, weak joke) it's essential to get BT to supply you with completely itemised bills for your ISDN line so you can see exactly what your router gets up to when you're not looking.

PCW Contact

Roger Gann can be contacted by post c/o *PCW* at the usual address, or via email at hardware@pcw.vnu.co.uk