



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.

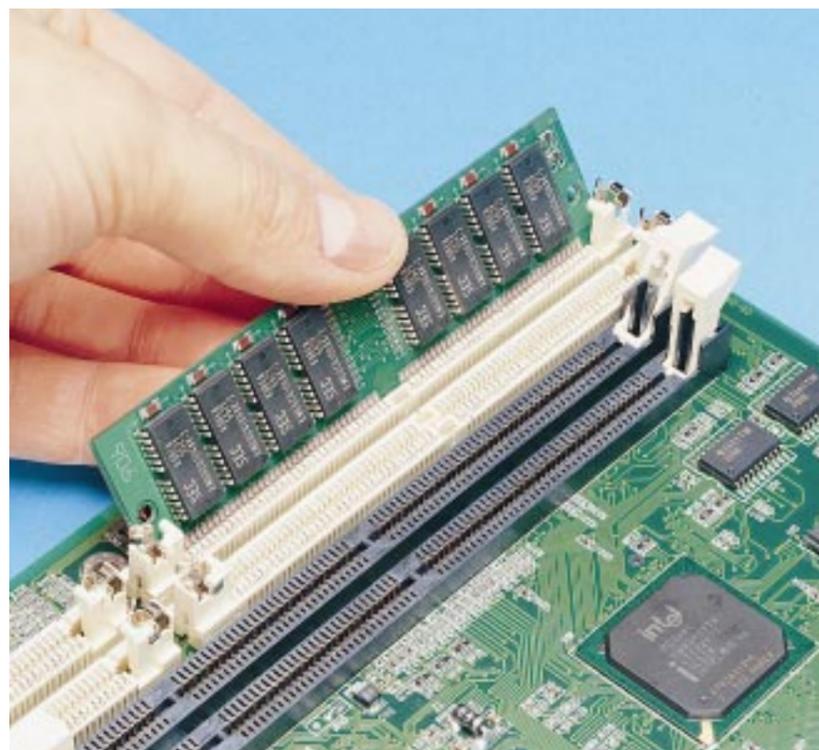
There's an old computer industry adage which goes: "You can never be too rich, too slim or have too much installed memory." Well, actually I just made that up. But the gist of it is undeniably true. Like death and taxes, insufficient computer memory is a perennial problem, and the way things are going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

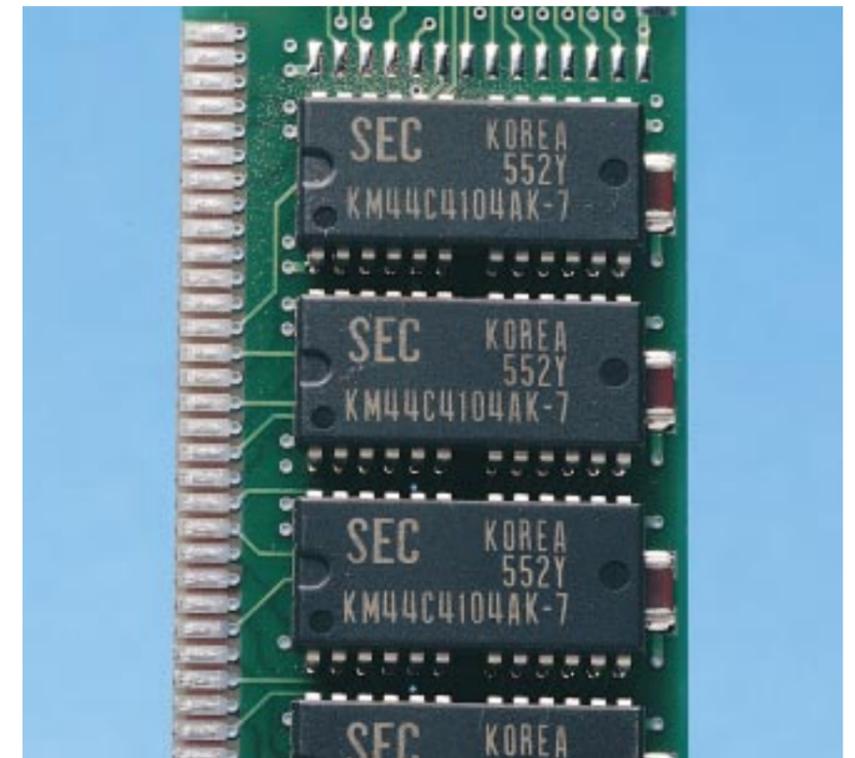
The sweet spot for your machine depends mainly on the operating system you're using. For example, if you're only using DOS, 1Mb is probably enough (or 4Mb if you want to play games). If you want to use Windows 3.1x, 4Mb should be viewed as an entry-level quantity of RAM and 8Mb as the sweet spot.

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,

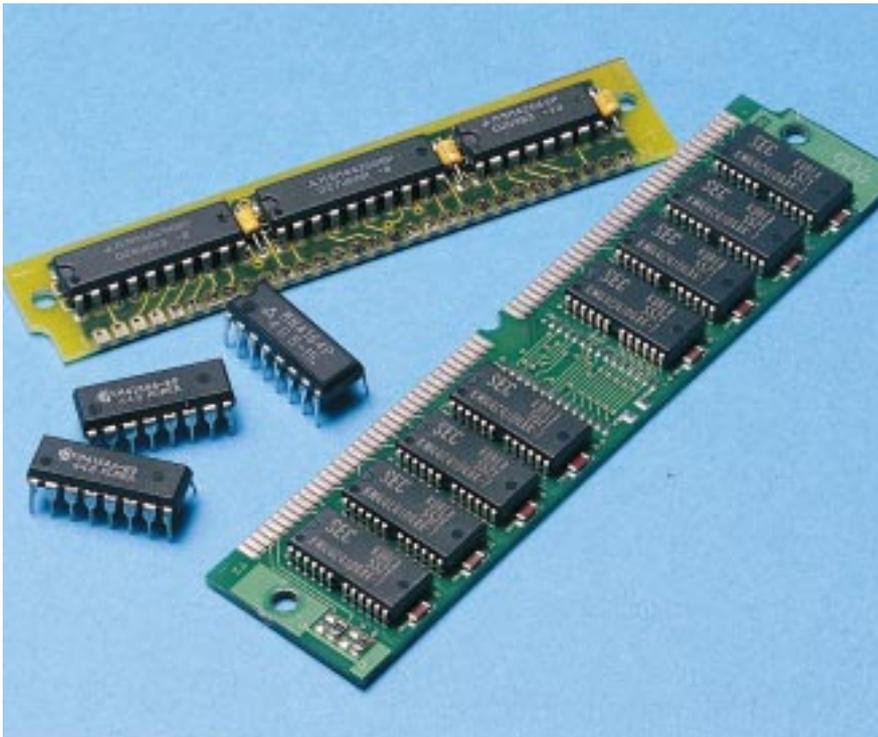
SIMMs are not always easy to identify. Here, the last digit (-7) represents 70ns

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.



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

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

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

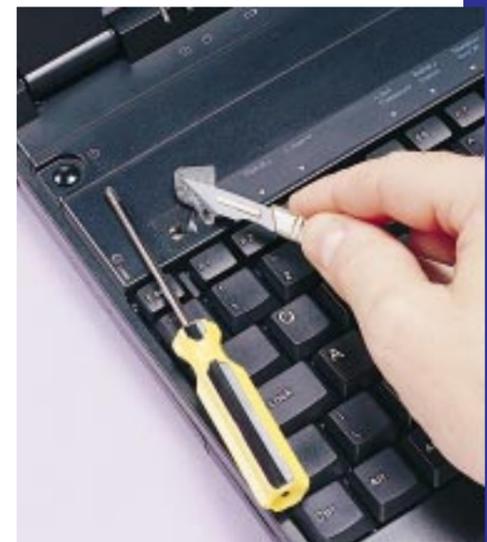
surrounding the drive, *very carefully* disconnect the thin ribbon cable, extract the old drive and replace it with the new one. With the drive installed, you then load CMOS Setup to get the notebook to auto-identify the drive, and away you go.

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



Screws can be hidden under adhesive labels

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.

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.

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



Hewlett-Packard P SureStore CD-Writer 6020

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.

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.

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

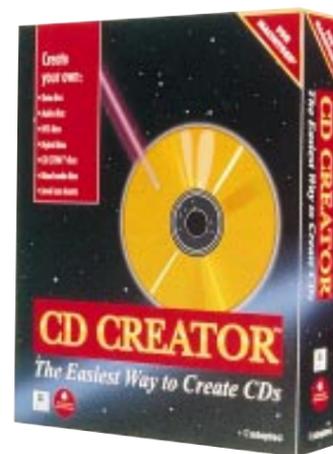


Sony Spresca 920

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.

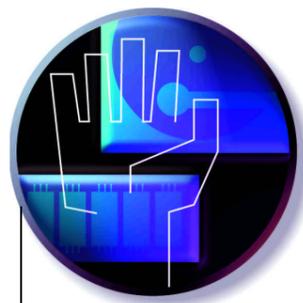


CD Creator for PC and Mac

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
Spresca 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.

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



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

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 gamers. 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.

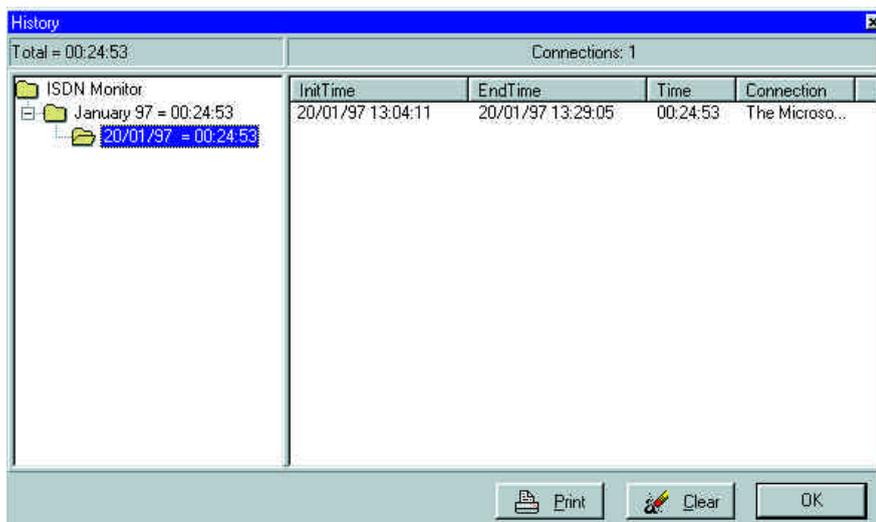
Codenamed Tillamook, this CPU will be among the first to be built on an Intel-developed pop-out module for notebooks, called MMO. It will include a processor, Level 2 cache, a voltage regulator, a clock and a PCI chipset called the 430TX. Instead of plugging a new processor onto the motherboard, users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



Keeping a tight control on ISDN line charges is crucial. This little utility, ISDN Monitor, specifically 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.

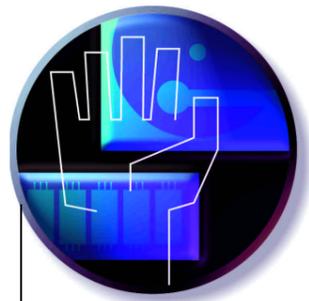
The final step is to roll out TCP/IP to the workstations, a relatively easy task if you're using Windows 95, less so if you're still using Windows 3.1x. Your internet service provider will have allocated you a range of IP addresses, say 192.168.0.1 to 192.168.0.25, and you then manually dole out one 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



Serial solution

Roger Gann shows you how to maximise comms speed by making the most of that unsung PC hero, the serial port. Plus, do *you* know the difference between baud rate and bps?

This month I'm taking a look at that unsung hero of the internet, the humble PC serial port — without it, where would modems be? I'll be showing you various ways to ensure you're squeezing every last drop of performance out of it, so that it doesn't bottleneck your modem and mar your surfing "experience".

The first port of call

The original PC serial interface used the 8250 (later 16450) Universal Asynchronous Receive/Transmit or UART chip. The 8250 UART has only a one-byte receive buffer, which means that after the 8250 receives a character, it has to be retrieved before the next character arrives or it's overwritten, i.e. lost. This leads to CRC errors, with consequent retransmissions and an inevitable impact on throughput.

At 1,200bps, a new character arrives about every 8ms — plenty of time for an interrupt handler to fetch the character from the UART and store it in a receive buffer. But at 19,200bps, a new character arrives about every 0.5ms, and at 115,200 baud (the maximum speed of the standard PC COM port) about every 90 microseconds. Depending on your CPU speed, there may not be enough time to reliably service one interrupt per character. As a result, the effective throughput ceiling of an 8250 UART is 38,400bps.

The superior 16550A or AFN UART overcomes many of the limitations of the 8250 and is reliable at speeds of up to 115,200bps. It has two main improvements over the 8250. Firstly, it has a pair of 16-byte send and receive buffers, thus giving it 16 times the buffering capacity of the 8250. Secondly, it has a First In, First Out (FIFO) buffering scheme, which can dramatically

improve performance on modem transfer speeds of 9,600bps or higher.

To see which UARTs are installed on your system, exit Windows and run the MSD utility included with DOS 5.0 and later by entering MSD at the DOS prompt. Once in the program, press C to see your COM ports. The type of UART installed for each of your ports appears on-screen.

Serial port upgrades

If you have an 8250 UART you have two choices: you can either upgrade to a 16550A or buy a third-party card, such as the Hayes ESP II. Note that if you have an internal card modem, it will most probably have its own 16550 UART aboard, so upgrading your PC's UARTs won't be necessary.

Upgrading 8250 UARTs

If you have a pretty old I/O card in your PC, the chances are it's got a 40-pin socketed 8250 UART. In this case, upgrading is simple. The 16550A is pin-compatible with the 8250, so all you've got to do is obtain a 16550AFN UART from somewhere like Maplins or Radio Spares, unplug the 8250 and replace it with the 16550A. However, if the 8250 is soldered in or you've got a more modern I/O card with a multi-function ASIC chip, then it's much less hassle to buy a £20 replacement I/O card.

Hayes ESP

Hayes Optima V.34 modems use a bigger "data dictionary" than their rivals and this enables V.42bis data compression to achieve 8:1 compression ratios, twice as much as normal. On a 28.8Kbps modem this would give you a maximum potential data transfer rate of 230Kbps, way beyond

the capabilities of a "fast" 16550 serial port. So Hayes offers its own serial-card solution, the Hayes ESP Communications Accelerator. This 16-bit ISA card is outwardly 16550-compatible but features a co-processor plus a pair of 1Kb send and receive buffers designed to eliminate buffer overflow problems. The ESP has had a spot of bother on the driver front over the years and the NT drivers have yet to emerge from interminable beta, but the Windows 95 drivers work just fine. The ESP II single port card goes for around £60 plus VAT. If you're serious about high-speed comms, this is the card to use.

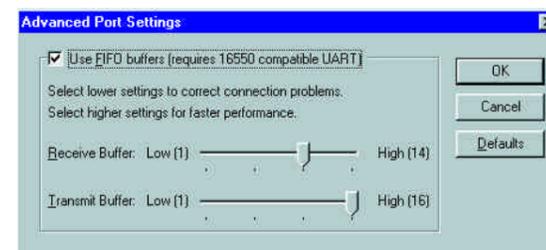
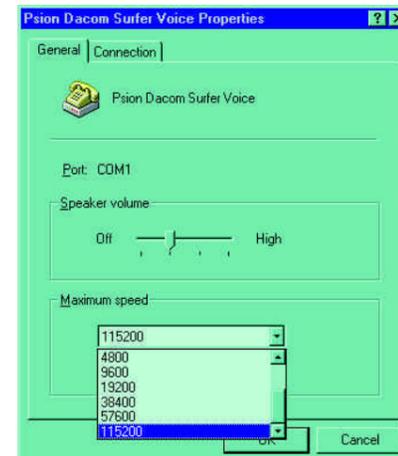
Software tweaks

As well as upgrading your serial port hardware, there are a number of software tweaks you can make to improve data throughput.

Flow control

With today's fast data transfer rates it's essential for the computer to be able to exercise some control over the flow of data to and from the modem, in order to prevent data loss. After all, there's no point in sending more data if the last lot hasn't been dealt with. This process is called flow control and there are several ways of doing it. The slowest way is software flow control, which makes use of two special characters, XON (Transmit on) and XOFF (Transmit off). When software flow control is in effect, sending a modem an XOFF character stops the flow of data. Sending the modem an XON character causes data to flow again. This is an inefficient method, significantly reducing bandwidth by adding two control characters for every eight sent.

The alternative is hardware flow control



Top Windows traditionally errs on the side of caution when it comes to modem line speeds, but it's normally quite safe to throw caution to the wind and crank it up to its maximum speed: 115,200bps in the case of V.34 modems

Above Once again, Windows errs on the side of caution when it comes to the FIFO buffers in the 16550 UART. Experiment by upping the Receive Buffer from 14 to 16

or RTS/CTS, which is more responsive and efficient. This uses a set of pins in the serial port, "Ready to Send" and "Clear to Send", and depending on their "state", data is either sent or paused. A minor caveat: you will need a decent, fully-wired serial cable with the CTS/RTS wires in place for hardware handshaking to work. Note that this type of flow control isn't the exclusive preserve of external modems — internal modems can also take advantage of it.

So how do you choose hardware flow control? In Windows 3.1x open Control Panel, then double-click on the Ports icon. The resulting window offers an icon for each COM port. Double-click on the icon of the port you want to configure, then click on the Settings button. Select Hardware from the Flow Control list box, and then click on OK. In Windows 95 click on the Modem applet in Control Panel, highlight your modem and click on the Properties button. Click on the Connection tab and then on the Advanced button. Check the "Use Flow Control" box

and select Hardware (RTS/CTS). Click OK a few times to back out.

You may also need to tell your comms apps to use hardware flow control. For example, if you use WinCIM 2.01, simply add the line

```
FlowControl =3
```

to the [Connector (CIS Connection)] in the CIS.INI file.

You'll also need to tell your modem which type of flow control to use. Check through its "initialisation string" and see if you can see "&K3", the usual Hayes command for turning on hardware flow control. You'll need to do this for DOS and

Windows 3.1x comms apps; under Windows 95, provided you've installed the modem correctly, turning on the hardware flow control via the Control Panel modems applet is sufficient.

Windows 3.1x

Windows 3.1 was the first version of Windows to support the FIFO feature of the 16550 UART for Windows-based applications. (DOS apps running under Windows 3.1, 3.11, or Windows for Workgroups 3.1 don't support the FIFO feature.) However, Windows for Workgroups 3.11 does give DOS comm apps access to the FIFO feature. An upgrade for the Windows for Workgroups communications driver SERIAL.386 is available which can also use the transmit buffer (TX) of the 16550 UART. Previous versions of the communications driver use

the receive buffer (RX) only. Check the Microsoft KnowledgeBase on the web for WG1001.EXE, which contains the updated driver.

You may be able to improve throughput without the need for new hardware, simply by adjusting a couple of Windows SYSTEM.INI settings. Incoming data is held in an internal buffer by COMM.DRV, the Windows 3.1x serial port driver, until it's retrieved by the comms program you're using. However, it's not always possible for the comms program to retrieve this data in a timely manner. If the buffer isn't cleared, then data can get overwritten. Luckily, it's possible to adjust the size of the buffer. By default, this is set to 128 bytes.

By simply adding a line to the [386enh] section of your SYSTEM.INI file, you can create a buffer as large as 10,000 characters. Actually, since the COMM driver maintains a separate buffer for each of your PC's COM ports, you can add up to four lines, each controlling the size of a particular port's buffer.

The syntax looks like this:

```
COMxBuffer=num
```

where x is the COM port number, i.e. a number between 1 and 4, and "num" the buffer size. Thus:

```
COM4Buffer=1024
```

sets the size of the buffer for COM 4 to 1,024 bytes. Don't forget that you've got to exit and restart Windows for the change to take effect.

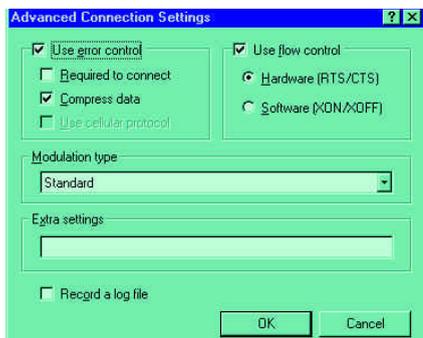
How big should your COMM buffer be? If characters are being dropped from incoming data transmissions, create a new buffer of at least 2,048 bytes. If problems

Universal Serial Bus (USB)

Luckily, in the nick of time, a solution to the inherent inadequacies of the serial port has arrived: the Universal Serial Bus (USB) which, by the end of the century, will probably replace both serial and parallel ports on PCs. With speeds of up to 12Mbits/sec, USB is aimed at simplifying and standardising the interconnection of PCs with peripherals such as modems, printers, mice, keyboards, digital speakers, joysticks, gamepads, telephones, telephone networks, scanners and digital cameras.

USB uses a tiny four-pin connector (two for data, one for power and one for ground). In fact, there are two types of connector: "Series A", a plug for a peripheral and a socket for a PC platform, is for applications permitting the cable to be moulded into its peripheral, while "Series B" is used for applications requiring a removable cable. The two-connector series is keyed differently to avoid mis-matching. PCs will most likely feature a USB "expansion hub" built into the keyboard or monitor, which will allow peripherals to be daisy-chained together — USB can connect up to 127 different devices to a single PC.

Instead of relying on the intelligence embedded in the host PC, the USB detects which devices are added or removed, automatically determines what host resources each peripheral needs, including driver software and bus bandwidth, and makes them available. Future versions of Windows 95 will be equipped with drivers that allow a PC to recognise USB peripherals. The plug-and-play feature enables users to attach or remove peripherals with the system running and without requiring the nuisance of rebooting. According to Dataquest, by the summer, 75 percent of the 22 million PCs manufactured worldwide will be equipped with USB ports.



If you want reliable, high-speed modem connections it's important to make sure you're using hardware rather than software handshaking; use this dialog to choose which

persist, increase the size of the buffer to 4,096 or even 8,192 bytes. Note also that the 16550A by default wakes up in 8250 mode and has to be switched in to 16550 mode via software. So if you've got a 16550 buffered serial port, you have to explicitly tell Windows 3.1x to use it. Add this line to the [386Enh] section of your SYSTEM.INI file:

```
COMnFIFO=0n
```

where n represents the number of the COM port for which you are activating the buffer.

Another tweakable parameter is the COMBoostTime. This SYSTEM.INI entry specifies the time allowed a "virtual machine" to process a COM interrupt: if it's too short, you can lose characters. Its default value is 2 but you could try increasing it to 4. Once again, it's located in the [386enh] section:

```
COMBoostTime=4
```

Rant corner — baud and bps

It's embarrassing, the number of otherwise knowledgeable people who use the terms baud rate and bits per second interchangeably. For example, the brand new CompuServe front-end, v3.0, continues to refer to "baud rate" when it really means bits per second. Baud actually refers to modulation rate: the number of times the line changes state every second. The difference between baud and bps is a result of the different forms of modulation used to encode digital data into analogue waveforms. At 300bps transmission rates, bps and baud happen to be the same; but above this speed, the 1:1 relationship ends. The maximum baud rate possible over normal phone lines is 2,400; the maximum bits per second rate is, at present, 33,600bps. So please stop using baud when you mean bits per second. Purleeze!

Note: If you don't find it here, just add it. You can specify an even higher value but be conservative, since setting it too high can actually slow down communications.

■ Windows 95

Mercifully, Windows 95 makes optimising your modem's data throughput considerably easier than Windows 3.1x. Provided the modem is installed correctly, then there are only a couple of things to check. You should have received an INF setup file with your modem, which correctly describes all the modem's features for Windows 95. But if it didn't come with one and your modem doesn't feature in Windows 95's list of modems, you can always install it as a generic "standard" modem. This is OK-ish but not optimal, and if at all possible I'd recommend paying a visit to your modem maker's web site to see if proper Windows 95 INF files are available for your modem, for download.

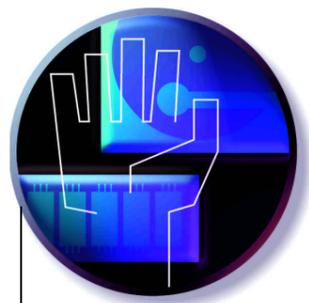
Once installed, open Control Panel and

click on the Modem applet. Highlight the modem and click the Properties button. On the General Tab set the maximum speed to four times that of your modem's speed: if you have a 28,800bps modem, set it to 115,200bps. Windows traditionally errs on the side of caution and always underrates the modem speed, typically setting the maximum speed to 57,600bps.

Next, click on the Connection tab and then on the Port Settings button. Make sure the Use FIFO buffers option is ticked and experiment by sliding the Receive buffer all the way up to 16. Click OK and then click on the Advanced button. Make sure the Use Error Control and Use Flow Control options are checked. Click OK several times to quit. And, er, that's it!

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



Printer potential

Look at your printer: it's working, right, and it does just the job you want? Maybe; but it could do better. Why not upgrade it? Roger Gann explains how.

The humble printer is arguably the most important computer peripheral you can buy, perhaps second only to the PC itself. Printers have traditionally tended to be fixed-function devices, capable of performing the actions they were designed to and nothing else, and largely incapable of being upgraded. If you had a daisywheel printer (remember them?) you could change the daisywheel or add a cut-sheet feeder, but that was about it. Things weren't much better for dot-matrix printers, either.

Expansion potential increased with the advent of laser printers. Many had at least one slot for a font cartridge, a slot for additional memory or perhaps an interface to the printer's "video" engine. It was now possible to upgrade your printer. Today, most laser printers are built with expansion in mind and feature things like SIMM slots, MIO ports for network interface cards, font card slots and comprehensive paper-handling options. Upgrade options for inkjet printers are less extensive, however. Nevertheless, it's surprising just how far you can upgrade a printer. But why would you want to?

There are several reasons. First in line is cost — it can be cheaper than buying a new replacement printer if all you're missing is a feature that an add-on can provide. And add-ons protect your investment in a printer: you only have to pay for the features you wish to add. You might think it doesn't make economic sense to upgrade an old laser printer, but printers like the LaserJet II and III were built like tanks and even today still have plenty of life left in them.

Printer upgrades neatly fall in to two categories, hardware and software.

Hardware Memory

Laser printers, or more accurately, page printers, hold a whole page of data in memory before printing it. The amount of installed memory isn't too important when it comes to printing text but it's crucial if you want to print graphics: a 300dpi printer needs a full megabyte of memory to print an A4-sized image. The move to higher resolutions, to 600dpi, actually quadruples the memory requirement, but the application of memory compression technology has kept the amount of RAM required to just a couple of megabytes. Even so, putting in extra RAM will always help when it comes to printing complex graphics, especially if PostScript is involved.

In the past, extra printer memory took the form of expensive proprietary custom memory cards, but most current printers take standard SIMMs so adding more RAM needn't be too painful, especially as

memory is still relatively cheap. Don't forget that if you have a GDI printer you can add more memory to the host PC to improve printing performance, with the benefit that it goes into the PC's general memory "pool".

Some printers, such as the EcoSys range from Kyocera and the LaserJet 5, can take flash memory as well. This special kind of memory is used to semi-permanently store downloaded print images, e.g. forms or letter headings. Every time you wanted to print a letter you could send a code to the printer to print the scanned-in header first. This way you'd print on plain paper and so wouldn't have to keep stocks of specially-printed headed paper: you'd never run short either. More importantly, it would mean that you could dispense with the need for a second paper tray to hold the headed paper.

Paper handling

It's highly likely that, regardless of what printer you've got, additional paper handling add-ons are available for it, so that you can use different sorts or sizes of paper with

Enhancing your parallel port

This is one of the few printer hardware upgrades that you fit in your PC. I'm not talking about any sort of printer port, I'm talking about the very latest bi-directional (bi-di) parallel port. This is available in several flavours, sometimes called an Enhanced Parallel Port (EPP) or Extended Capabilities Port (ECP). Either way, a "bi-di" parallel port enables a two-way conversation between your PC and the printer, allowing the printer to pass status messages back to the PC, vital things

like "Paper Jam" or "Low toner". These "8-bit" parallel ports are faster than the standard 4-bit parallel port fitted to most PCs, and if you have a GDI printer you'll appreciate a faster parallel port. So, if your printer has a "bi-di" port, it makes sense for the PC to have one as well. The best one I've come across is a VL-Bus multi I/O combo card with an IDE interface plus fast serial ports. Price £29 from Dabs Direct ((0800) 558866) but check the ads in *PCW* for other deals.

Updated printer drivers

The Windows printer drivers that shipped with your printer have most probably been superseded, the ones that came with Windows almost definitely have. Getting the latest driver means you'll be using a less buggy (hopefully!) version that's probably a bit faster and possibly with more controls or features. You can check the version number by opening Printers in the Control Panel, clicking on Setup and then clicking on the "About..." button. The best way to get the

latest drivers is to download them from bulletin boards such as CIX, CompuServe or over the internet, from the manufacturer's web site. For example, the latest PostScript driver, v4.1, is available from the Adobe Systems web site at www.adobe.com and this has numerous extra features compared to the plain vanilla Win95 version written by Microsoft. You can also download the latest PPD and INF files for your PostScript printer to ensure it's correctly installed.

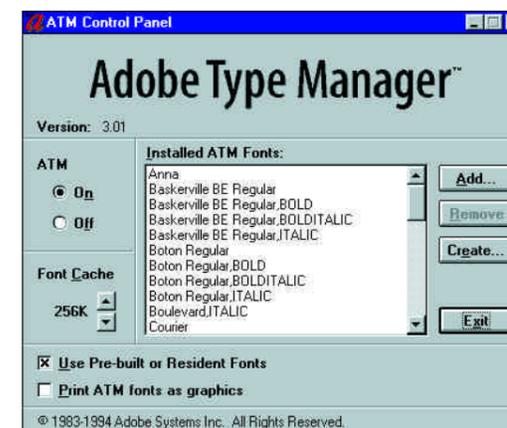
your printer. Take the dot-matrix printer — printing sticky labels can be a risky proposition with a push tractor feed that many dot-matrix printers have as standard. Many allow the fitting of a pull-tractor feed specifically for this task. Or you might want to do a lot of mail merges so you need a cut-sheet feeder instead, something a little more advanced than that rudimentary cut-sheet tray that came as standard with the printer.

Most laser printers are now designed with add-on paper tray "stacking" in mind and so can take a second paper tray that fits to the underside of the printer. Fitting a second paper tray allows you to put your letter-heading in one tray and your continuation sheets in the other. Or, if the printer is a busy one, being shared on a network, the second tray can be used to expand the printer's overall paper capacity: such trays typically have a 250-sheet capacity. Some inkjet printers even have a second paper tray option. Most are a doddle to fit — the printer just sits on the second paper tray and once you've made a few changes to the configuration via the front panel or through the printer driver, you're away.

Colour

You may not be able to upgrade your laser printer to colour (though the NEC SuperScript 610/660 can do spot colour) but both inkjet and dot-matrix printers can commonly be upgraded from mono. By far the easiest to convert are inkjet printers: providing they are "colour ready", all you do is remove the black ink cartridge and replace it with a snap-in tri-colour cartridge. You then specify "colour" in the printer

driver (or install a colour driver) and, hey presto, you've got colour, for about £25 or less. If only all upgrades were this simple.



The ATM Control Panel. The create panel allows you to make customised versions of Adobe Multiple Master fonts

Font cartridges

One of the original printer upgrades was to increase its stock of internal fonts. This took the form of a cartridge or card that you plugged into a special socket on the printer.

Expanding the number of internal fonts might not seem relevant to Windows users; after all, scaleable TrueType fonts are a dime a dozen these days. But there's still a compelling reason not to turn your back on resident fonts — speed. When a font resides in the printer, not only is there no time wasted downloading soft fonts, but the print time is also quicker. If you have to do a lot of printing in the same font, then you should consider making that font a permanent, resident printer font.

Pacific Data Products produces one of the best known font cartridges, the 25 in One. This cartridge contains 172 fonts and is available for the II, III and 4 Series printers, plus IBM and Epson laser printers. It's also available as a SIMM for the LaserJet 5.

PostScript upgrades

The ultimate font upgrade has to be the one that installs PostScript on your laser printer as these typically also install the basic "family" of 35 PostScript fonts. You get the exacting typographical accuracy of PostScript, too.

Installation is simple: you just plug in the PostScript cartridge (or SIMM in the case of the LaserJet 4 and 5) to get instant, high-quality PostScript fonts. This is undeniably simple, but there's a price to be paid as most genuine PostScript printers have relatively powerful processors, something lacking in older LaserJets, and this can make for slow printing. You may need to top up printer memory, too. Expect to pay about £200 to £250 for a PostScript upgrade, from Hewlett-Packard. Pacific Data has the PacificPage Level 2, a PostScript Level 2 language emulation SIMM module for the LaserJet 4 and 4 Plus printers.

Pacific Data Products also sells a range of plug-in accelerator cards, designed to speed up sluggish print performance in a LaserJet fitted with a PostScript cartridge. There are two versions: the PacificPage



IIXL, for the LaserJet II, which consists of the PostScript cartridge plus an accelerator board; and the PacificPage PE/XL, which is for the LaserJet IIP, IIP Plus, III, IIID and IIIP.

Software upgrades

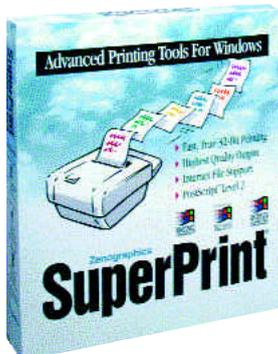
The software driver plays a significant part in the printing process, having a major influence on both speed and print quality. The drivers that come with Windows are pretty good but, as ever, there's always room for improvement, either adding features that don't exist or doing existing tasks better.

Perhaps the simplest printer enhancement you can buy is to add extra TrueType fonts to your Windows installation — these are available from Microsoft, Monotype, Bitstream.

More sophisticated upgrades fall into two categories — accelerators and print enhancements.

Adobe Type Manager

Adobe Type Manager (ATM) predates TrueType and was the original scalable font manager for Windows. For most Windows users, it



SuperPrint's comprehensive colour correction options

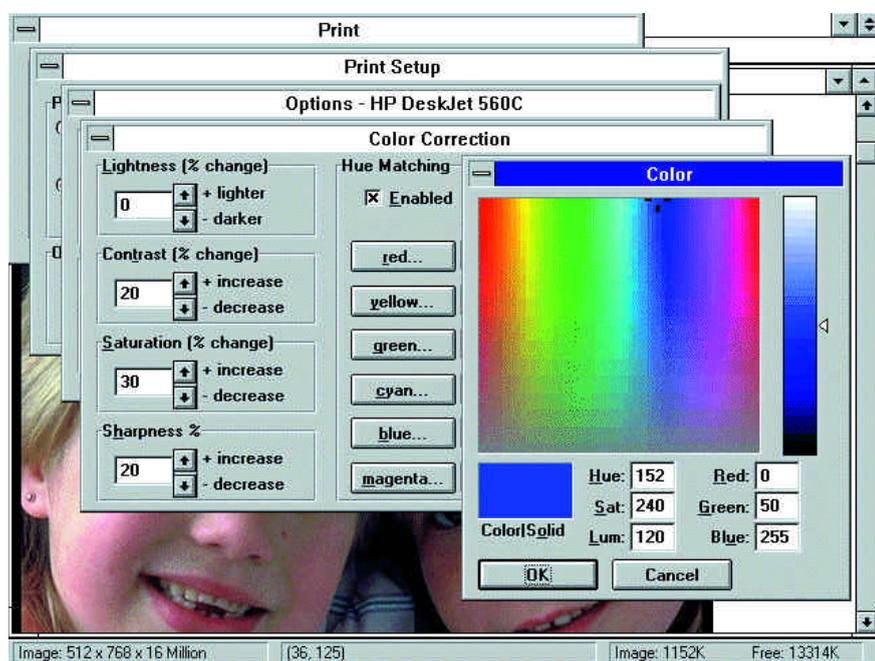
carries out much the same function as TrueType — it allows a printer font to be accurately displayed on-screen at all point sizes, thus removing the need for a raft of separate fixed-point screen and printer fonts. ATM also lets you use superior PostScript soft fonts with a non-PostScript printer, and its Multiple Master fonts feature is useful if you want to fine-tune the appearance of a font. However, much of ATM's thunder was stolen by TrueType

which has the redeeming virtue of costing zilch. As a result, only those that really need to use genuine PostScript fonts will fork out £30 for ATM, i.e. those submitting work to a print bureau or typesetter.

SuperPrint 4.0

SuperPrint's job is simple — it speeds up Windows printing, especially colour printing, which has always been a relatively slow process. It does this by using highly-optimised 32-bit printer drivers and sending the data to the printer in a special compressed "metafile" format. It comes complete with "SuperDrivers" for Windows 3.1x, 95, and NT 3.51, with controls for sharpness, contrast, lightness, saturation, dot gain and hue matching. Like ATM, it allows you to turn your non-PostScript printers into PostScript devices, but this time for graphics as well.

While you can get speed gains when used in conjunction with low-end colour printers, I guess that the true benefits of SuperPrint only become apparent when used in conjunction with high-end output devices where fine control of the output print quality is important. You will need a fairly well specified PC in order to run SuperPrint; basically, the more RAM you have the better, and I'd start at 16Mb. This amount allows its "SuperRIP" to perform full-frame rasterisation of the entire image to be "printed" in RAM; with lesser amounts of RAM it has to break down the image into bands, which is of course slower.





Drive time

Floppy disks have been overtaken and lapped by today's removable storage drives. Roger Gann explains the new technologies and sorts the hatchbacks from the sportsters.

It's all to do with relativity, you know. Ten years ago, when your PC had a 20Mb hard disk, a 1.2Mb floppy was a capacious device capable of backing up the entire drive with a mere 17 disks. Today, the standard hard disk fitted to PCs has become the 1.2Gb drive: a 600-fold increase in capacity. And what of the humble floppy? In the same period, its capacity has increased by less than 20 percent. As a result, it's now at a disadvantage when used in conjunction with any modern large hard disks — for most users, the standard floppy disk just isn't big enough anymore.

In the past, this problem only affected a tiny proportion of users, and for those that did require high-capacity removable disks (typically DTP'ers running Macs), solutions were available: 44Mb 5.25in SyQuest removable hard disks. But these were expensive.

Times have changed, and today, everybody needs high-capacity removable storage. These days, applications don't come on single floppies. They come on CD-ROMs. Thanks to Windows and the impact of multimedia, file sizes have gone through the ceiling. Create a Word document with a few embedded graphics and before you know it, you've got a multi-megabyte data file, quite incapable of being shoehorned into a floppy disk.

Awkward as it is, there's no getting away from the fact that a PC just has to have some sort of removable, writable storage, with a capacity in tune with current storage requirements. We all need removable storage for several reasons: to transport files between PCs, to back up personal data, and to act as an overflow for your hard

disk, to give you (in theory) unlimited storage. It's much easier to swap removable disks than fit another hard disk to obtain extra storage capacity.

In 1981, 5.25in floppies kicked off at 160Kb, quickly went to 180Kb and then to 360Kb with the advent of double-sided drives. In 1984, the 5.25in floppy maxed-out at 1.2Mb. That same year, Apricot and HP launched PCs with the revolutionary Sony 3.5in 720Kb disk drive. Three years down the road, this doubled in size to 1.44Mb and for the past decade or so, that's where it's stayed. Oh sure, there have been attempts to increase the capacity of the humble floppy, but none got very far.

First, there was IBM's bid (in 1991) to foist a 2.88Mb floppy standard on the PC world, using expensive barium-ferrite disks, but it never caught on. And both lomega and 3M had another go in 1993 with the 21Mb "Floptical" disk. But it never took off: it was just too dear and too small.

For a long time, the range of removable storage options wasn't extensive. You could choose between low-capacity floppies, clumsy tape solutions, expensive optical drives, and removable cartridges. But today, the situation is radically different, and finally there is a real choice in removable storage. There is a wide range of capacities on offer, starting around 100Mb and rising to 1Gb. The new removable drives have never been more affordable, and having been designed with the end-user in mind, are easy to install.

Installation options

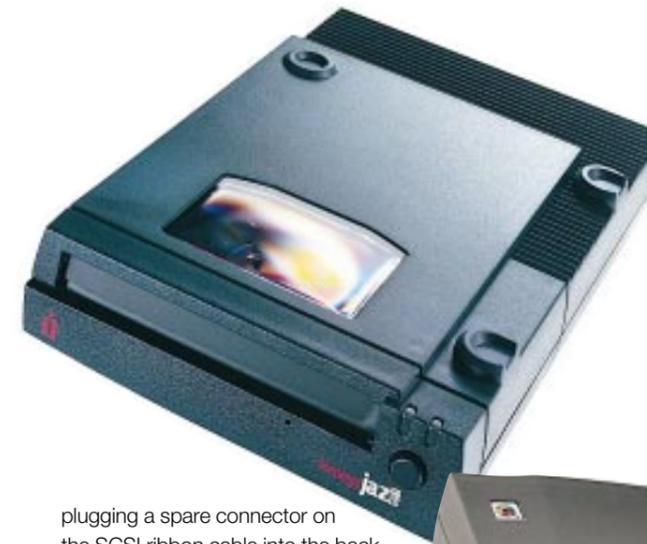
Removable drives are available both in external and internal versions. The former will have either a parallel port or SCSI interface, while the latter will have a SCSI or

IDE interface, although IDE is not all that common at the moment.

Unlike fitting a second hard disk, which can be a daunting task, the great news is that most of the new breed of removable storage devices have been designed with ease-of-installation in mind. The easiest to install are the external drives with a parallel port interface. Not only do you not have to take the lid off your PC, but it also lets you share/move the drive between different PCs. You can use it with a notebook, too. And because it uses your printer/parallel port, installation is no more complicated than plugging in a cable.

The only problem with this otherwise perfect arrangement is that the parallel port, which was never designed to handle these devices, isn't very "fast", and older parallel ports can act as a bottleneck, slowing down the drive's data transfer rate. Luckily, however, most recent PCs are equipped with an enhanced parallel port (EPP or ECP) and these are capable of good data transfer rates. I recently used a parallel port tape streamer hooked up to an EPP port and it was capable of backing up at just under 8Mb/min — a pretty respectable rate. In any event, such bottlenecks would only really impinge on fast devices, such as removable hard disks. Slower devices, such as the Zip drive, don't have a stellar data transfer rate to start with and so are unlikely to suffer at the hands of the parallel port interface.

If the ultimate data transfer rate is what you're after, you ought to consider SCSI versions of these devices. Most internal drives (Jaz, SyQuest, MaxIT and Zip) tend to use a SCSI interface, which is fine if you already have a SCSI card fitted as installation is no more complicated than



Left lomega's Jaz drive can take cartridges as big as 1Gb — enough space for you to back up your entire hard disk

Below SyQuest's EZflyer is cheaper but holds a maximum of 230Mb

plugging a spare connector on the SCSI ribbon cable into the back of the drive. The same applies to external SCSI drives, too. Very often they can be "daisy-chained" to other external SCSI devices. If you haven't got a SCSI card, they can be expensive to buy and fiddly to configure, although lomega does a cut-down SCSI host adaptor called the Jet, which sells for around £80. Given the choice, I'd spend another £20 and buy a better host adaptor, from, say, Adaptec, Bus Logic or AvanSys.

Spoilt for choice

Although all these devices offer the same basic facilities and differ only in terms of capacity and speed, they do employ different technologies. They can be split into categories based on their respective technologies. At the last count there were three removable storage technologies: super-floppy, hard disk, and magneto-



optical. Even though CD-R discs are dropping in price and the Panasonic innovative PD hybrid drive offers some interesting advantages, neither technology currently has much market presence.

Super-floppies

The current crop of super-floppies may resemble conventional floppy disks in physical size and operation, but they don't use the same magnetic recording technology employed on the standard 1.44Mb floppy. Without doubt, the most popular super-floppy is the Zip drive, of which lomega claims to have sold ten million. The secret of the Zip's good performance (apart from its high 3,000 rpm spin rate) is a technology pioneered by lomega (based on the Bernoulli aerodynamic principle) which actually sucks the flexible disk up towards the read/write head rather than vice-versa. The disks are soft and flexible like floppy disks, which makes them cheap to make and less susceptible to shock. They're fast, and have a capacity of 94Mb — big enough for most users. Sadly, Zip drives are not backward compatible with 3.5in floppies and can't be used as boot devices, although there are BIOSs in the pipeline which will permit that. Imagine, you could install Win95 from a single floppy!

The new kid on the super-floppy block is the rival standard being promoted by 3M and Compaq. An LS-120 disk looks very similar to a common-or-garden 1.44Mb 3.5in disk, but uses a refinement of the old 21Mb floptical technology to deliver much greater capacity and speed. In fact, this technology had originally been developed by lomega, but

p304 >

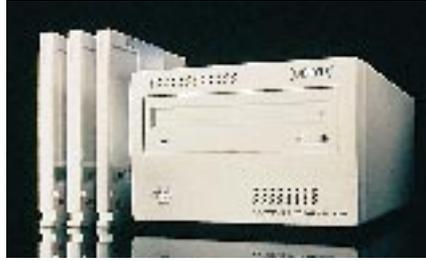
was abandoned and sold on to 3M. Named after the "laser servo" technology it employs, an LS-120 disk has optical reference tracks on its surface that are both written and read by a laser system. These "servo" tracks are much narrower and can be laid closer together on the disk: an LS-120 disk has a track density of 2,490tpi (tracks per inch) compared with 135tpi on a standard 1.44Mb floppy. As a result, the LS-120 can hold 120Mb of data.

As well as being a bit bigger, its other advantage over the Zip drive is that it can also read/write ordinary 3.5in floppies and so can be used as a boot device. While its data transfer rate is faster than that of a standard 3.5in floppy drive, it's not as fast as a Zip drive due to its considerably slower spin speed. And at present, the LS-120 isn't available to buy as an add-on — it's currently only available ready-installed on Compaq's new range of Deskpro PCs.

Hard disks

Above about 100Mb, the most commonly used removable drive technology is derived from that found in conventional hard disks, which not only gives you high capacities but also provides fast performance, pretty close to that of conventional fixed hard disks. These drives behave just like small, fast hard disks.

This approach isn't particularly new: SyQuest cartridges have been using the technology for many years. However, SyQuest drives have tended to be small



The PowerMO 230 is one of the latest generation of magneto-optical drives from Olympus

(44Mb and 88Mb). Recently, SyQuest embraced the 3.5in form factor and launched a 135Mb drive (the EZ-135) but this has been rapidly superseded by the 230Mb EZflyer. It's fast, like a hard disk, and reasonably priced (about £200).

Xyratex, the IBM spin-off, has got in on the act with its 540Mb MaxIT drive. This is a little pricier (about £300) but has twice the capacity and can read older SyQuest cartridges. It's nippy, too.

The drive with the largest capacity is Iomega's Jaz which can take cartridges as large as 1Gb. This, too, is affordably priced (about £375) and offers excellent capacity with great performance. Its huge capacity makes it possible to back up your original hard disk in one go, in just a matter of minutes. It's a good choice for audio-visual work, capable of holding an entire MPEG movie in one swallow. You could even make it your primary drive!

Magneto-optical

As you might expect from the name, these drives use a hybrid of magnetic and optical

technologies. The disks have a special alloy layer that can be modified under the influence of a magnetic field. Changes can then be read by reflecting a laser beam off the alloy layer. The magnetic field in the magneto-optical (MO) disk actually twists the laser's beam of light, and this twist can be detected and used to read the data.

However, this is a "two-pass" process, and as a result, MO operations are relatively slow. This drawback is compounded by the fact that MO heads tend to be heavy, which also makes for slow average access times. Nevertheless, MO disks are cheap at around a tenner and they have top archival properties, often being rated with an average life of 30 years — far longer than any magnetic media.

MO came close to being killed off by the Zip, but a new generation of faster, cheaper drives, spearheaded by Fujitsu, have breathed new life into the format. MO disks are much the same size as 3.5in floppies but about twice as thick. Olympus has been the first to market the latest generation of MO kit with the slick-looking PowerMO 230. And Fujitsu has recently announced a 640Mb MO drive, the £299 DynaMO, which is backwards-compatible with older 128Mb and 230Mb MO disks.

PCW Contacts

Roger Gann can be contacted by post c/o PCW at the usual address, or via email at rgann@mcgilivray.win-uk.net



Two's company

You don't need an entire bank of PCs to create a network. Roger Gann's no-nonsense guide shows that linking two PCs can be not only a simple task but cheap, too.

As we race headlong towards the Millennium (no Matrox puns intended), the PC buzzword is "connectivity". To get the maximum from your investment, your PC just has to be connected to others, either via a local area network or to the internet. The *Hands On Hardware* column is nothing if not topical, so this month I'll be showing you how cheap and easy it is to network a pair of PCs, courtesy of Windows 95.

Why network? The age-old reasons are just as valid for two PCs as they are for 2,000: you can share data and expensive peripherals, you can centralise backup and improve security. Here is an example: maybe you've already got a PC at home in addition to the one you've bought your child, ready for Christmas. There's no sense in having two printers, or two tape streamers or two CD-ROMs in the house. You might as well share them on a simple

network. Or, maybe you just fancy playing Quake in deathmatch mode?

This sounds like using a sledgehammer to crack a nut but it's so cheap to network a pair of PCs that it's almost a no-brainer. With less than £50 you'll be able to fully network both PCs to let them share resources. And yes, that does include VAT!

We're talking about real bargains here. And Windows 95 is the best network client you can have. It comes ready for just about any kind of connectivity you care to throw at it, together with all the networking software you may need. All you have to do is to add the hardware.

Cable capers

Mercifully, your network hardware shopping list won't be very long. All you need is a pair of network interface cards and some cable.

By far the most popular networking cable (or "media" in networking parlance) is thin Ethernet, also known as 10Base-2. This resembles a co-axial cable, the sort used in TC aerials. A five-metre long cable, terminated in BNC connectors will cost you around £10.

You will also need a pair of T-pieces and a pair of terminators, the total cost of which is about £6. Thin Ethernet uses what is termed a daisy-chain or "bus" topology, with all the devices linked by what becomes a single cable running between each PC.

If you were far-sighted you might think that twisted-pair cabling (also known as 10Base-T) was the true path to follow, rather than thin Ethernet and you'd be right.

However, because it uses a star "topology" with all the cables returning to a central hub, the cost of the hub generally

makes it a non-viable proposition when linking only a pair of personal computers. But it is possible to buy what I term a "null-modem" twisted pair cable that you can use to directly connect two PCs, thereby obviating the need for a hub.

Many network interface combo cards will have both 10Base-2 and 10Base-T connectors and they will automatically sense which one you're using. But they tend to be relatively expensive and for our modest project I'm wanting to keep the costs down, so I'll go for the simplest 16-bit Ethernet card with just a 10Base-2 BNC connector, which should set you back less than £20 a kick.

A really good gauge of compatibility is whether or not it is "NE2000" compatible, so look out for this — the NE2000 was an old make of network card, which was originally sold by Novell, the mother of all PC networking.

Connecting the card

Installing a network card isn't a difficult job: no worse than, say, fitting a SCSI host adaptor. These days most network interface cards, even the cheap ones, are software-configurable, which makes life much simpler when first commissioning your network.

All network cards require hardware resources, specifically an IRQ and an I/O port address. Typically, an NE2000-compatible card will, by default, be set to use IRQ 3 and I/O port 300h. There's normally no problem with this I/O address

but the default IRQ is another matter. We all know (don't we?) that IRQ 3 is already spoken for, being the preserve of COM 2. It is now possible to share IRQs on a PC but it's not a particularly good idea, so in the interests of having a quiet life I'd recommend using another, unused, IRQ instead — perhaps IRQ5 which is earmarked for LPT 2. Or maybe IRQ 10. It really does not matter which one you choose just so long as it is free.

The fact of whether you've got an old-fashioned card festooned with jumpers, or a modern jumper-less software-configurable card determines the install sequence. If it's the former, you should use Windows 95's "Add new hardware" wizard to install the card drivers. These may come with the card, or if it really is NE2000-compatible it might rely on the standard NE2000 drivers that come with Windows 95.

At this point, using the card's default settings, the Windows 95 Device Manager will tell you whether there are going to be any hardware resource conflicts. If there are, it's a simple matter to adjust the hardware resource setting that is conflicting, to one that isn't. Note down the new settings and adjust the card's jumpers to make it match. You can then open the PC and install the card in the usual way.

If, on the other hand, you have a modern network card, the correct sequence is to install the card in the PC first and then reboot Windows 95. If your PC is blessed with a plug-and-play BIOS, the new card

should be recognised and you'll be prompted for drivers and then settings. If the BIOS doesn't support plug and play, use the Add new hardware wizard to auto-detect the new card. Install the appropriate drivers when prompted and if there's a hardware resource clash, simply adjusting the values on-screen will reconfigure the card. The final step is to cable up the two PCs. Fit the T-piece onto the BNC socket on the network interface card using the "vertical" side. Plug the Ethernet cable into one side and the terminator plug into the other. Then, just repeat the operation at the other end and that's it.

Putting in the protocol

We've now got to install the network protocol we're going to use, so hunt down the Windows 95 CD-ROM and plunk it in the CD-ROM drive because we'll be needing it. Windows 95 comes with several heavy-duty network protocols but for our little peer network (so-called because the workstations are of equal status), all we really need is NetBEUI. If you use dial-up networking to access your internet access provider, you might already have TCP/IP installed. Don't worry, because you can have more than one protocol "stack" in use at a time.

Fire up Control Panel and click on the Network icon. The Configuration tab should list the network card you've just installed. Click the Add button at the bottom and add a Client (the Client for

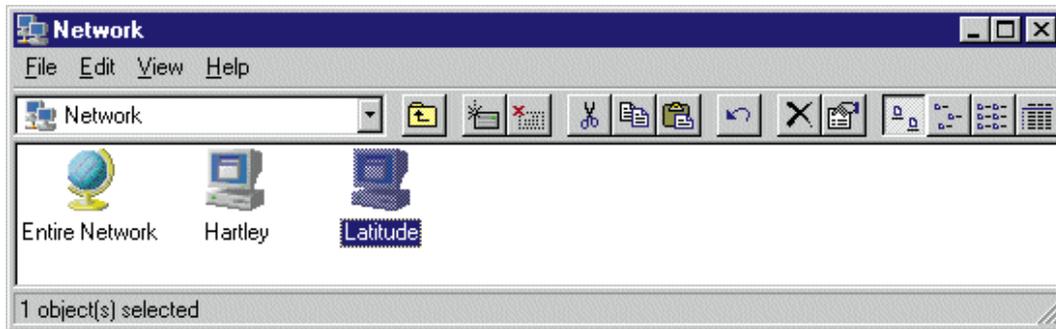
Troubleshooting

What if you can't browse your new network at all and can't see any of the shared resources on the other PC? Often it's much better to try and sort out a problem like this from the comfort of a DOS prompt, unencumbered by Windows 95 baggage. Windows 95 comes with elementary network diagnostics which come in handy at times like these. Open a DOS window and type NET DIAG <CR> on one machine. Select I for IPX (assuming you're using this network protocol). Type N to make this PC act as a diagnostic server, then do the same at the other machine, typing I for IPX and Y. The first machine will be broadcasting test messages which the other PC will be listening out for. If it can't hear anything, there must be a hardware problem of some sort, either with the cable or with the configuration of the card.

```

MS-DOS Prompt
Auto
Microsoft(R) Windows 95
(C)Copyright Microsoft Corp 1981-1995.
C:\WINDOWS>net diag
IPX and NetBIOS have been detected.
Press I to use IPX for diagnostics, N to use NetBIOS, or E to exit this program.
Microsoft Network Diagnostics will use IPX.
Searching for diagnostic server...
The diagnostic server has been located on the network.
Communicating with diagnostic server. This may take several seconds.
Validating reply from diagnostic server.
The diagnostic server's reply is correct.
This indicates that network information is being sent and received properly.
The command was completed successfully.
C:\WINDOWS>_

```



The Network Neighbourhood facility displays all machines connected to your peer-to-peer network: from here, you can configure the shared drive

Microsoft Networks). Click OK. Select Protocol and click Add. Select Microsoft from the list on the left and NetBEUI from the list on the right. Click OK. Finally click Service, select Microsoft on the left and File and Printer Sharing for Microsoft Networks on the right, then click OK.

Back at the main Network dialogue, select the Identification tab, give your computer and your workgroup (i.e. network) a name. Make sure both PCs have different individual names but the same workgroup names. Entering a description is optional.

When you accept all these settings, you'll be asked to restart the PC. When the PC reboots, you'll see a new dialogue box halfway through the Windows 95 boot sequence, prompting you for your name and password. These are "blank" at present so if you enter a name and password it'll ask you to re-enter and confirm it. Remember the password as it could come in handy later! We're now logged on to our little network. But how do we access drives and printers on the other PC?

A caring, sharing kind of network

Before we can see resources like drives and printers on other PCs, they have to be "shared". Once shared, they become available on the network. At its most simple, you right-click the drive, folder or printer icon and, from the pop-up context menu, select Sharing. Click on the Shared As button and give the resource a Share name (e.g. CD-ROM or just a letter if it's a hard disk). Select the type of access you want. For instance: full, read-only or "depends on password".

If you want, you can give a password specific to this resource. In this way you can make entire drives available either for sharing (you share the root) or just individual folders. On a network like this, I'd share entire drives and probably give full-access and read-only for CD-ROM drives.

A printer connected to the host can be shared, too. Once shared, don't forget to go over to the other PC and add the printer in Control Panel, specifying it as a network printer rather than a local printer. When Windows 95 browses the network to find

shared printers, it will discover the shared printer on your other PC. It will then install an appropriate driver, then you'll be able to use that printer as though it were directly connected to your PC.

Once you've shared your resources, you should be able to see them from the other PC. Go over to it and click on Network Neighbourhood. You should see the other PC listed next to the Entire Network icon, plus an icon for this PC. Double-click it and the shared resources will be listed (typically, a drive letter like C:). Right-click the shared object and select Map Network Drive from the pop-up context menu. Give the new drive a letter like D and check the Reconnect at logon box. When you start Windows 95, you'll be able to access the other PC's hard disk as if it were Drive D:.

PCW Contacts

Roger Gann can be contacted by post c/o PCW at the usual address, or via email at rgann@mcgilivray.win-uk.net



Out with the **old**

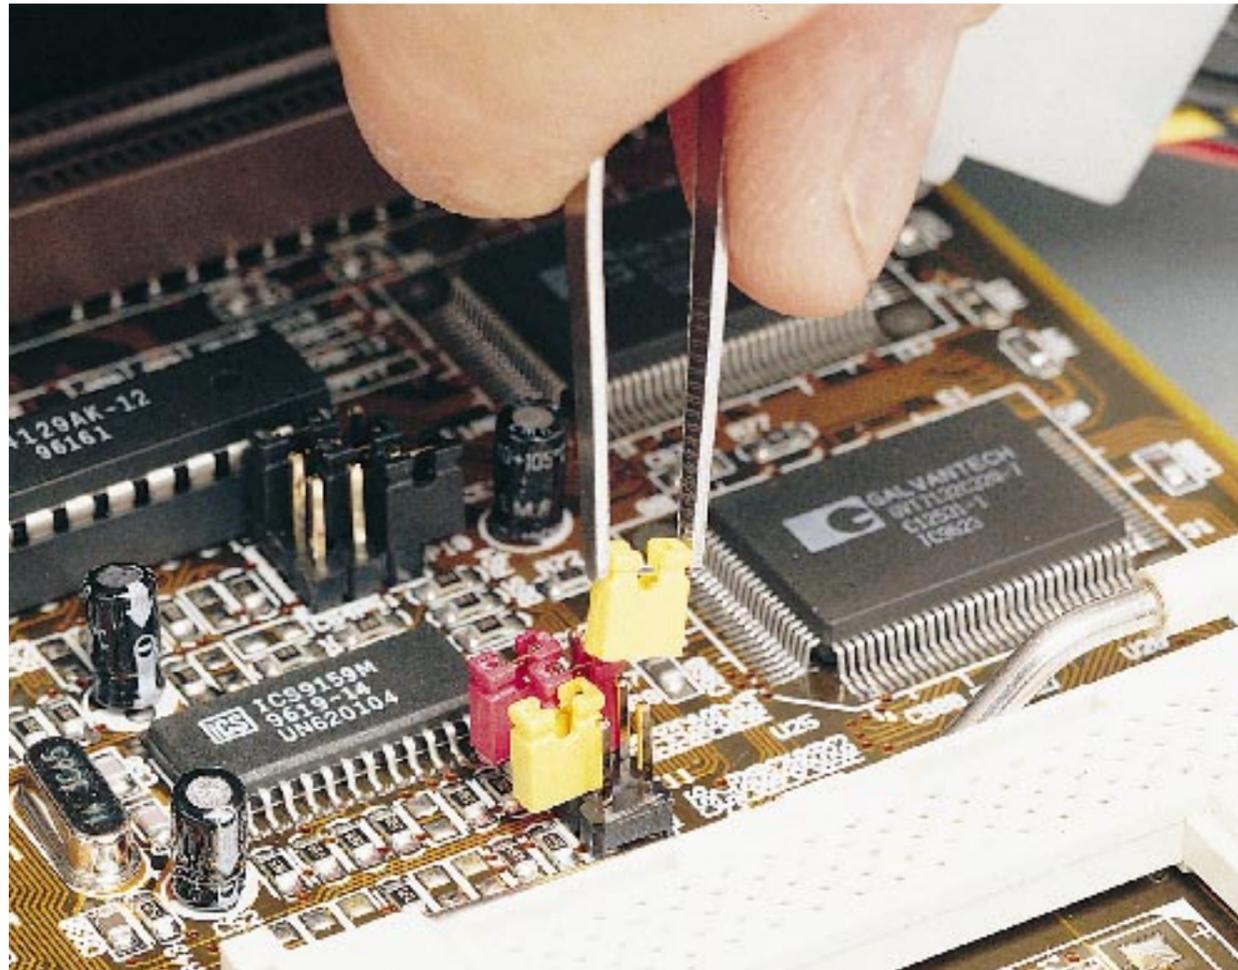
Fed up with watching paint dry while your PC chugs away? Why not upgrade the motherboard — Roger Gann shows you how.

Over the past few months I've looked at upgrading your PC in various ways: adding more memory, fitting an overdrive processor, adding an EIDE hard disk, and so on. While

all these upgrades will improve your PC's performance, the overall gains garnered from these upgrades aren't particularly breathtaking. This is because all you're doing is moving a performance bottleneck

elsewhere in the PC to some other, relatively old, component.

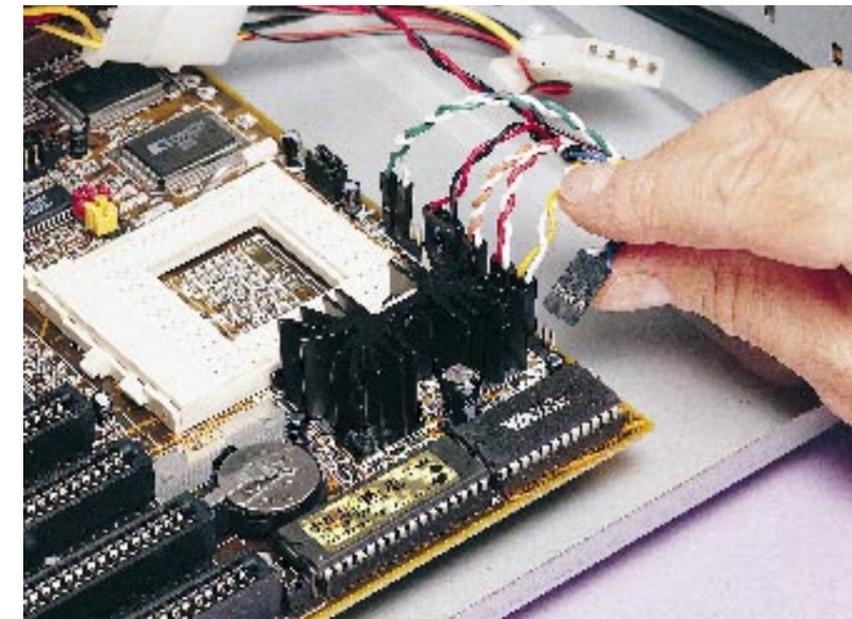
A more fundamental upgrade is required if you are looking for tangible performance gains: I'm talking about replacing the



Coloured jumpers: we're in luck as this motherboard uses coloured jumper sleeves to differentiate the various types of jumper, e.g. clock speed and processor type. Make sure the new motherboard is correctly jumpered for the type and speed of processor you're using

motherboard in your PC with a more powerful one. Not only would such a move allow you to use a more powerful processor, but you'd benefit from other advances, too, such as pipeline burst mode cache or the latest RAM technologies, such as EDO or Synchronous DRAM. You would also benefit from the latest PCI bus technology and lifesavers such as plug and play, especially if you're a Windows 95 user.

Sometimes you have no choice but to swap motherboards. The fastest Pentium overdrive for the 486 architecture stops at 83MHz, yet Pentiums are now running at 200MHz, and if you want more performance you'll have to upgrade your motherboard. But this is not as expensive as you might think. A DX4/100 overdrive will cost £120, while a 100MHz Pentium motherboard will cost only £70 more. Don't lose sight of the potential extra costs imposed by the new motherboard,



Multicoloured leads: pay close attention to the multicoloured leads and plugs before you unplug them from the old motherboard — it'll help you to identify each cable. If you're lucky, the plugs will be labelled

such as new memory if you've got old 30-pin SIMMs, or a new PCI graphics card if all you've got is VL-Bus.

A motherboard upgrade sounds complex, but don't worry. It is relatively hassle-free and your old PC gets a power boost without the need to duplicate or replace files, connections or peripherals. Complex reconfiguring is not necessary and, outwardly, your PC remains as its old familiar self.

The downside

Although it is physically possible to upgrade just about any PC's motherboard, you should pause before taking the plunge. Swapping motherboards is rather more involved than fitting an expansion card and not something for the technically timid. Choosing a motherboard can be a big problem because, apart from Intel with its Endeavour and Atlantis main boards, most are brand-less designs from companies of

which you may never have heard.

Another problem is that, thanks to the pace of change, a particular model of motherboard might only be in production for a few weeks before being replaced. Remember, only four months separated the launch of the 166MHz and 200MHz Pentium processors, so the odds are that when you see a motherboard listed in an advertisement, its make and model won't be listed. Motherboards are effectively

generic products, rather like potatoes. As for technical support, well — you are almost always on your own. Then there is the problem of overdoing it. There is no point in performing such an upgrade on a 386 because the rest of its hardware, like its graphics card and its hard disk, will be so slow as to be inadequate. The “weakest link in the chain” analogy applies here.

Another significant problem is the form factor of the motherboard: that is, its shape and size. While most clone manufacturers use the so-called “Baby AT”-sized motherboards which are readily interchangeable, most big name manufacturers tend to use their own proprietary motherboard designs, which aren't. Not only are they an odd, non-standard shape, but they may also be highly-integrated designs incorporating all I/O functions directly on the motherboard and not on expansion cards. You expect this from the likes of IBM and Compaq but you'll find that Amstrad, Viglen and Elonex also fall into this category.

The only way to upgrade the motherboard in these makes of machine is via the manufacturer, which makes them difficult to get hold of and pricey. So before jumping in at the deep end, make sure your new board will fit okay.

Step by step

Step 1 — Preliminaries

■ Before doing anything, it's a good idea to print out your existing CMOS settings. You can do this by taking a screendump of each

screenful of settings. Ensure your printer is on and connected to the PC, and press the PrtScr key.

The reason for doing this is to prevent potential problems later on when you install your new motherboard and you can't remember the previous settings. Don't forget, we're only changing the motherboard — all the other components are staying put.

Probably the most important settings are those that set the hard disk's geometry, because if you get these wrong, the PC won't reboot at the completion of the upgrade.

You could argue that this is unnecessary as most modern BIOSes are capable of auto-configuring the hard disk's geometry, so what's the point? Problems can arise if you previously used drive translation, say to reduce the number of cylinders to less than 1,024. When auto-configuration kicks in, it will only use the drive's default geometry, not the translated one. The drive will not then boot or be recognised. But if you have the original values, it's easy to key them in.

■ Take a moment to examine your new motherboard and read through its undoubtedly sparse manual. Check if there's anything important of which you should be aware. Most motherboard manuals are invariably techie, but you should try to identify the positions of important jumpers.

■ If you're fitting a new processor as well, install it and move all the jumpers now,

while you've got perfect access to every corner of the board. The same applies if you're fitting new RAM.

■ Finally, assemble your toolkit. You'll need a crosshead screwdriver, an electrician's screwdriver and maybe a pair of fine needle-nose pliers.

Step 2 — Dismantle the case

■ Switch off the PC and unplug everything attached to the system unit: typically the keyboard and screen. All micro-electronics are very 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 touching some metal plumbing before handling the new motherboard. Better still, invest in an earthing wrist-strap.

■ Remove the casing lid by undoing the screws at the back.

Step 3 — Remove the expansion cards

■ Our next task is to remove all the expansion cards, so undo the screw holding each card's blanking plate and carefully extract each card, putting it somewhere safe.

I guess that most people will need to remove just three cards: a graphics accelerator, a sound card and a multi I/O card which will have the serial ports, parallel ports and the hard disk/floppy disk controller. This will have a couple of grey ribbon cables plugged into it, so unplug

p291 >

them, taking careful note of which side of the ribbon has the coloured edge.

■ Tuck the now loose cables away behind the disk drives and don't forget to unplug the hard disk's activity LED cable. More modern motherboards have done away with separate I/O cards and instead integrate the I/O ports on the motherboard, so in this case unplug all the cables.

If you have a sound card, unplug the thin CD audio cable before removing the card. With all the cards out, there's no need to remove anything else. The hard and floppy disk drives can stay put. Some compact casing designs use drive cages that hang out over the motherboard and make installing/removing a motherboard awkward, so remove these as well.

Step 4 — Remove the cables

■ The next step is to remove the rest of the cables from the motherboard. Located somewhere near the keyboard socket at the rear you should see a pair of large, white or cream, power leads perhaps labelled P8 and P9. These can be quite reluctant to unplug, so you may have to use measured force to get them to shift.

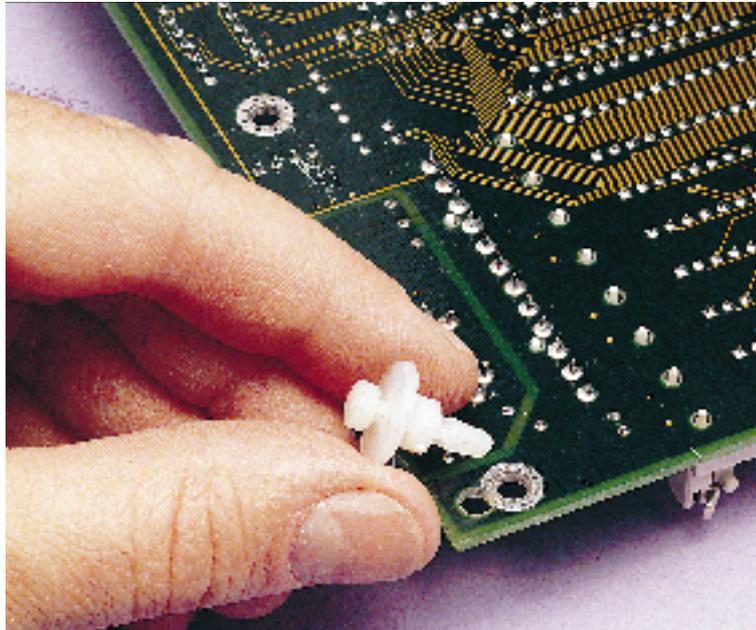
■ There will also be a group of thin multi-coloured cables plugged in at the front of the motherboard. These are for the power and turbo LEDs, the keylock and reset button. Occasionally, the little black plugs will be labelled. If they're not, it makes sense to identify them now while they're still connected. There's often a name printed on the motherboard to identify that connector.

Step 5 — Remove the old motherboard

■ The motherboard will probably be held by a couple of bolts. Locate them, undo them and put them in a safe place.

■ The motherboard will now be resting on several white plastic stand-offs, or spacers. These fit into tapered slots at the base of the chassis. You should now be able to slide the old motherboard out to the left but you may have to wiggle it about a bit first.

■ Once the old motherboard is out, use the pliers to extract the stand-offs and re-fit them in the corresponding holes in the new



Plastic standoffs: be sure to make a careful note of which holes the standoffs occupied in the old motherboard before transferring them to the new one. You'll need a pair of pliers to remove them

motherboard. There are normally several spare holes in the motherboard so don't fit spacers to all, as there won't necessarily be corresponding slots in the system case.

■ If you're able to reuse any SIMM modules, now is the time to remove them and transfer them to the new board.

Step 6 — Install the new motherboard

■ Give the new motherboard a once-over to make sure you've correctly set its jumpers.

■ Slide the new motherboard into position at the base of the chassis, so that the stand-offs engage in the right slots in the casing. The position of the keyboard socket and the hole in the casing for it will help you to locate the motherboard correctly.

■ Refit the bolts to secure the motherboard in place.

Step 7 — Reconnect the cables

■ Refit the P8 and P9 power cables to the motherboard: the norm is that the black cables on each plug should go together. The most fiddly bit is refitting the little multicoloured cables. Typically, these will be positioned along the front edge of the motherboard and I can guarantee that they won't be in precisely the same position or order as they were on the old motherboard.

■ There will probably be connectors for these cables: keylock, reset, power, turbo LED, switch and speaker. Sometimes they can be grouped together to form two or

three plugs. If this is a problem because the motherboard connectors are not arranged like this, it's okay to split the plastic plug into two separate plugs using a sharp knife.

Step 8 — Refit the expansion cards

■ Replace the previously-removed expansion cards, the graphics card, the I/O and hard disk controller.

■ Reattach the hard disk data ribbon and LED cables.

■ Replace the cover, plug in the screen and keyboard, and power up your PC. If

something is not right, you'll know it because the BIOS will signal an error.

As the BIOS can't depend on the video working, it will simply emit a series of coded beeps through the speaker. If you haven't fitted the memory properly, or if the video card isn't working, you'll hear a pattern of repeated beeps. Check the motherboard handbook to decode the pattern. In my experience, the most common problem has involved a badly-fitted or defective video card. A badly-fitted SIMM which is not quite seated is another candidate.

Step 9 — Edit the CMOS settings

■ The new motherboard's CMOS settings will effectively be blank, so the first time you boot up you'll typically be invited to enter CMOS setup to key things like the date and time, hard disk details and so on.

■ Try out the hard disk auto-configuration option and, using the notes you made earlier, compare the old settings with the suggested new ones: if they're the same, accept them; if not, edit the values manually. Don't forget to specify your floppy disk settings. Delay tweaking the advanced CMOS settings for another day, save the settings, and reboot. All things being equal, your PC should now boot up as normal.

PCW Contacts

Roger Gann can be contacted by post c/o PCW at the usual address, or via email at rgann@mcgilivray.win-uk.net



The perfect host

If you want to add peripherals to your PC, SCSI is the ideal choice of interface. Roger Gann shows you how to install a SCSI host adaptor.

Although much of SCSI's thunder has, in recent years, been stolen by Enhanced IDE which offers great performance, ease of installation and value for money that's hard to beat, the basic strengths of the Small Computer System Interface (SCSI — pronounced *scuzzy*) remain undiminished.

SCSI remains the interface of choice for advanced PC peripherals such as very large, hard AV drives, magneto-optical drives, writable CD-ROMs, DAT tape drives and scanners. It offers the ultimate in data throughput, good flexibility, and lately a degree of compatibility that had eluded it from the outset, making it relatively

painless to add a multiplicity of diverse peripherals to your PC.

For serious multimedia work, SCSI is essential, guaranteeing the very high data throughputs required when handling video and audio streams. While it's true that Mode 4 PIO EIDE hard disks can deliver 16.6Mb/sec and that the latest eight-speed IDE CD-ROMs can deliver 1.2Mb/sec, this method of transferring data requires such a high degree of CPU intervention that it affects the CPU's ability to further process the data it's transferring, such as video or MPEG data streams. This problem doesn't

occur with SCSI because data transfers are controlled by the host adaptor and hence steal few, if any, CPU cycles.

As its name implies, SCSI is a system interface rather than a device interface, which is designed to operate with specific devices such as disk or tape drives. System-level interfaces on the other hand are more general purpose, designed like expansion buses to match virtually any kind of device to a computer system. What makes SCSI so special is that it's "smart", and has some of the advanced arbitration features of the EISA and Micro Channel buses used in PCs.

The original SCSI specification featured an 8-bit data path at a maximum speed of 5MHz, providing a maximum possible data transfer rate of 5Mb/sec. "Wide SCSI" doubles the width of the data path to 16 bits (and allows for further widening to 32 bits) with transfer rates from 10Mb/sec to 20Mb/sec. "Fast SCSI", meanwhile, doubles the data rate across the SCSI connection by doubling the speed to

10MHz. A system that takes advantage of all SCSI-2 possibilities (i.e. Fast and Wide) can achieve a data transfer rate as high as 40Mb/sec, though these are fairly rare. A final benefit of SCSI is its ability to extend the "reach" of a single expansion slot: each SCSI host adaptor can support up to seven devices, and you can have more than one SCSI host adaptor.

Installing the card

As is often the case, this is a two-part job: installing hardware first, followed by the software. If you're not running Windows 95 then you'll need to know just what hardware resources (i.e. IRQs, DMA channels, I/O ports and ROM addresses) are already spoken for in your PC. This is so that you can identify any potential hardware clashes in advance and adjust the settings on the SCSI card before you fit it.

I thought I'd mastered the arcane mysteries of SCSI a long time ago, but only last week, when I was fitting an external Zip drive, my system was brought to its knees for an afternoon by a common or garden SCSI card — and this was a PC running Windows 95! Time spent gathering hardware resource information is time well spent and can prevent much frustration.

The most common brand of SCSI host adaptor is Adaptec and that's what I'll be fitting: an Adaptec AHA-154CP SCSI-2 card. The default resource settings of this card are IRQ 11 and I/O port 330h and usually these settings are, generally speaking, "free". So all things being equal, they shouldn't clash with anything. However, if problems do arise, this card has a set of easy-to-get-at DIP switches which configure the I/O ports, address space of the SCSI ROM BIOS, and other factors. Other SCSI minutiae are configured by hitting CTRL+A during the boot sequence, which pops up the SCSISelect menu. By contrast, the PCI version, the AHA-2940 PCI Fast SCSI-2 card, is a DIP switch and jumper-free zone and is completely software configurable — perfect for Windows 95.

Hardware installation

Step 1

Take the usual safety precautions: power down and unplug the PC from the mains and disconnect all other leads, such as printers and mice. If you're unsure which bit goes where, mark where each lead goes when you unplug it.

Remove any static electricity you might be carrying by earthing yourself before handling any electrical components: touch a metal pipe, for instance; or if you've got one, wear an earthing wrist-strap.

Step 2

Take the lid off of the PC. It'll be held on by four or five self-tapping screws and you'll most likely need a Phillips screwdriver to undo them — don't lose the screws!

OK, you can now see the inside of your PC. If you're going to do detective work to discover the hardware resources used by the existing cards, do it now. Choose a free 16-bit ISA expansion slot. Undo the bolt securing the blanking plate at the end of the slot and remove the blanking plate.

Step 3

Hold the card firmly by its top edge and

press its gold edge connector firmly into the expansion slot. Watch the clips on the external SCSI connector: they often get in the way, or get trapped.

Do up the bolt to secure the card. If you already have a floppy interface, flick DIP switch 5 to disable the floppy interface of the Adaptec card.

Step 4

If you're installing any internal SCSI devices, set their ID and termination before fitting them in their drive bays.

Plug in the power leads and make sure your internal SCSI ribbon cable has enough connectors for all your devices.

Hook up the ribbon cables to the internal SCSI devices and, as always, note the coloured edge of the ribbon cable — this has to be aligned with the same side as Pin 1 of the sockets at each end.



Every SCSI device must be assigned a unique SCSI ID, from 0 to 7

SCSI Troubleshooting

It's often at this point in the procedure that you'll come across problems. If it's an IRQ clash, this is easily changed by hitting CTRL+A and using SCSISelect. If it's an I/O or ROM BIOS address overlap you'll have to flip some DIPs, but as the switch block is on the top edge of the card there's no need to yank it out to get at them. If the problem isn't a blatant hardware resource clash, look for the obvious errors: two SCSI devices with the same ID number, a loose connector or termination plug, or a SCSI device that's not turned on. Once you've excluded the obvious, you can look for the subtle.

Double-check your work

Verify that your SCSI devices are compatible with your host adaptor. With ASPI systems, make sure that you've installed both host and peripheral drivers. Watch for driver error messages by using the F8 interactive boot option.

Juggle with the cabling

Sometimes the cable lengths between external peripherals are critical so try swapping a short cable for a longer one (or vice versa), or moving to a different connector on the ribbon cable. Make sure you snap in place the retaining clips or

wires on each SCSI connector. This is particularly important with SCSI connections because they work like old-fashioned Christmas lights: if one goes out, they all go out.

Juggle with the SCSI IDs and device order

Although it shouldn't make any difference which ID is used or where on the cabling "chain" a device is plugged in, changing these is worth a try.

Check termination

Try moving the terminator from the end device, or even adding one in the middle. For example, instead of terminating the last device, try terminating the second-last. Work your way back. Or leave it off entirely.

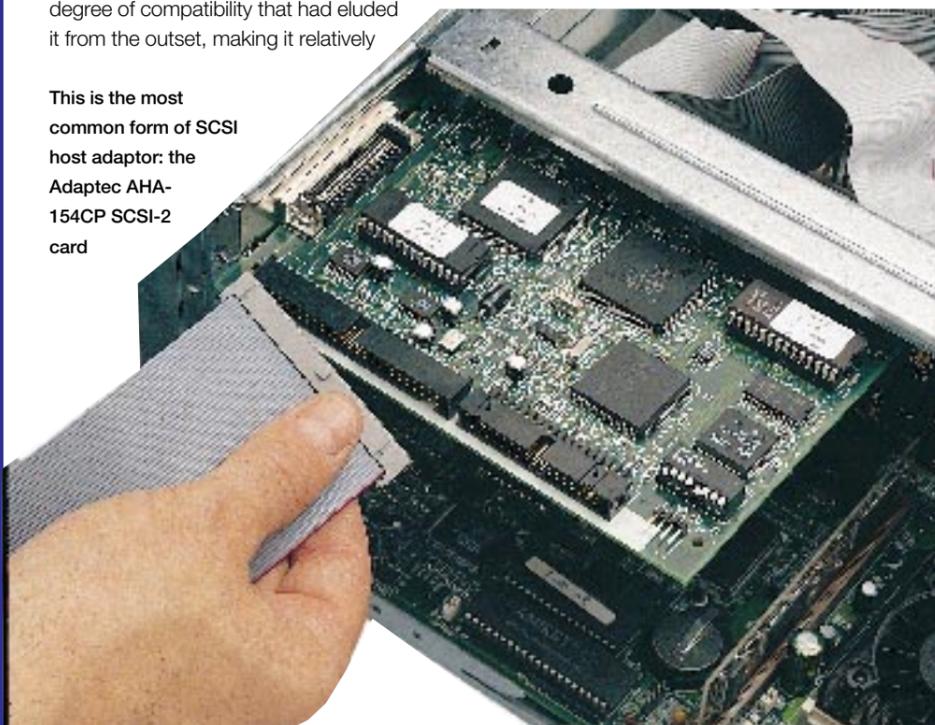
Divide and conquer

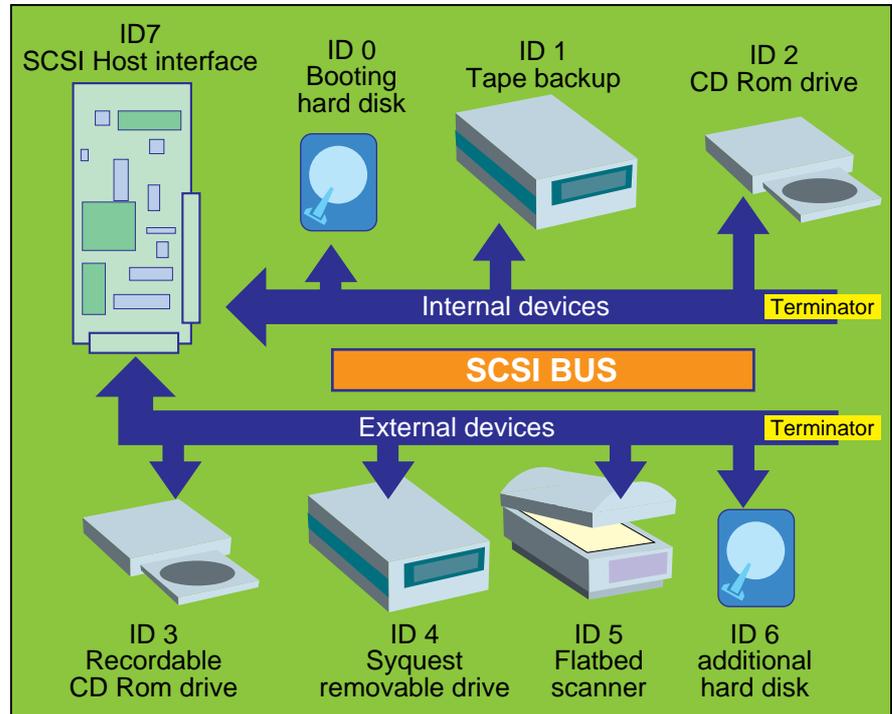
If all else fails, break the SCSI chain in two and try each half separately.

Use the resource kit

Finally, check out the Windows 95 Resource Kit Help file on the Windows 95 CD-ROM. It's a gold-mine of useful troubleshooting advice and information.

This is the most common form of SCSI host adaptor: the Adaptec AHA-154CP SCSI-2 card





■ If you are installing a bootable SCSI hard disk as well, don't forget to alter your CMOS setup to "No hard disk installed" as the BIOS on the SCSI card takes care of the hard disk from now on. If you're not, disable the SCSI BIOS, otherwise it'll hunt for a drive to boot from, won't find it, and the "timeout" will delay booting up for a minute or so.

Step 5

■ Put the lid back on, do up the screws and plug everything back in, including any external SCSI devices such as scanners or Jaz drives.

Power up the PC and make sure everything is working — i.e. that it will boot up correctly. One new thing you'll notice at boot time is the SCSI BIOS "signing on" after the video BIOS, listing all the SCSI devices it's found.

Step 6

■ If you've installed a new SCSI hard disk as well, now's the time to partition it using AFDISK (Adaptec's version of FDISK) and high-level format it with FORMAT. If you need to low-level format the drive first, this can be done via SCSISelect.

Installing the software

■ Apart from bootable hard disks, all other SCSI devices require some sort of software support so we now need to install some software. This used to be a fraught affair but Adaptec has gone some way, with its EX-

SCSI package, to making the software installation and configuration virtually automatic.

■ If you're running DOS+Windows 3.11, depending on your devices and hardware it'll add these device drivers:

```
DEVICE=C:\SCSI\ASPI4DOS.SYS /D
DEVICE=C:\SCSI\ASPIDISK.SYS /D
DEVICE=C:\SCSI\ASPICD.SYS /D:
ASPICD0
```

The ASPI4DOS line loads the ASPI driver for the card while ASPIDISK adds support for removable disks such as SyQuest, Zip or Jaz. The third line installs support for a SCSI CD-ROM drive.

■ Windows 95 doesn't need these Real Mode drivers and instead has its own

Hardware terms

AV	Audio Visual
ROM	Read Only memory
DAT	Digital Audio Tape
EISA	Extended Industry Standard Architecture
IDE	Integrated Drive Electronics
EIDE	Enhanced Integrated Drive Electronics
SCSI	Small Computer Systems Interface
PIO	Programmed Input/Output
IRQ	Interrupt Request
DMA	Direct Memory Access
CMOS	Complementary Metal-Oxide Semiconductor
BIOS	Basic Input/Output System

SCSI Basics

SCSI ID

Just as multiple IDE drive installations require a master and a slave, so every SCSI device (and that includes the host adaptor) must be assigned a unique SCSI ID number.

There are eight SCSI IDs (0 to 7) and typically the host adaptor will take ID 7. Generally speaking, so long as the IDs are unique you can use any ID, the only difference between them being that higher numbers get higher priority during arbitration, so ID 6 will be serviced before ID 5. However, bootable devices normally have to be either 0 or 1. Note that if you already have an IDE drive in your PC then you must boot from this drive — you have no choice.

Most external SCSI devices use a pushbutton or rotary switch on the rear panel to set the ID, while internal SCSI devices will probably use jumpers or DIP switches. The very latest Plug and Play SCSI devices feature SCAM (Self-configuring Auto Magically) and these automatically allocate SCSI IDs without user intervention.

Termination

Just as thin Ethernet network requires termination at each end of the network chain, so does SCSI. A terminator is a bank of resistors that absorbs the excess signals on the SCSI line and prevents them from reflecting back across the cable and causing problems. They can take the form of a plug or a resistor pack that plugs in to the device's circuit board, or DIP switches.

So, if you have no external SCSI devices: the host adaptor will be at one end of the SCSI "chain", with an internal SCSI device (a hard disk, say) at the other. Both devices will need to be terminated and if you fit a CD-ROM in the middle of the chain, between the hard disk and the SCSI card, then it won't need to be terminated because it's not at the end. However, if you fit it after the hard disk, then it will. Similarly, if you plug in an external device this will have to be terminated and the host adaptor un-terminated. Sorting out termination is important as it can prevent your SCSI devices from working properly. Luckily, modern SCSI devices often feature "auto-

termination" which can automatically work out which devices require termination.

Cabling

You have to pay special attention to SCSI cables. For a start, external cables must be a minimum of 0.3m long, while the total SCSI chain must be no longer than about 6m. Internal SCSI devices (i.e. those inside the system unit) all attach to one 50-pin ribbon cable that typically has connectors for two devices and the SCSI card. If you plan on adding a third internal device you'll need a ribbon cable with more connectors. If you're re-attaching a third internal drive (a CD-ROM, say), you'll need a new ribbon cable with connectors for as many drives as you plan to install — an 8-way internal cable should cost no more than £25.

Sadly, external connectors are less standardised: there are no less than three types of connector; the 25-pin plug, as featured on Macs and Zip drives, the 50-pin Centronics-style connector, and the newer SCSI-2 50-pin high-density connector. External cables tend to be dear, so expect to pay up to £30.

internal 32-bit Protected Mode SCSI drivers. However, the drivers supplied with the card are likely to be newer than the ones on the Windows 95 CD-ROM, so you should use these in preference. EZ-SCSI

4.0 includes a useful Windows 95 SCSI utility, SCSI Explorer, which lets you change some esoteric SCSI settings, things like automatic drive power down and write cacheing.

• PCW Contacts

Roger Gann can be contacted by post c/o PCW at the usual address, or via email at rgann@mcgillivray.win-uk.net



E-IDE add-ee-oh

Tread lightly through the seven steps to installing an Enhanced IDE hard disk: your PC's performance depends on it. Roger Gann is your guide.

In last month's column I showed you how to upgrade your PC to Enhanced IDE, which not only allows it to see the full capacity of today's enormous hard disks, but gives you other benefits, too, such as lightning-fast data transfer rates and a second IDE channel for other devices such as CD-ROM drives and tape streamers. This month I'll be completing the job by taking you step-by-step through the process of installing an Enhanced IDE hard disk and preparing it for use.

Fitting a bigger hard disk is perhaps the most popular upgrade that users undertake and it's not too difficult to see why. The arrival of Windows, and Windows 95 in particular, has led to an explosion in the amount of disk space occupied by a PC's operating system, its applications and data. Given plenty of RAM, many older PCs are more than capable of running the new disk-hungry 32-bit operating system, but their small, 170Mb hard disks just aren't big enough to hold the new operating system as well as all your apps.

The hard disk I used in the upgrade was the latest Quantum drive, the 2.5Gb Bigfoot. This is different to any other Enhanced IDE drive you can buy simply because of its form factor. Unlike its rivals, which fit a 3.5in drive bay, the Bigfoot harks back to a bygone age and fills a 5.25in bay.

It's one of the cheapest EIDE drives you can buy, but it's not the fastest. This is partly due to its slower, 3,600rpm spin speed and the diameter of its platters, which means its heads have more ground to cover, resulting in a so-so average access time of 15ms. However, the larger platter size does have some compensation: its tracks are

correspondingly longer and the heads thus don't have to move about quite so much. As a result, on small record transfers, its data transfer rate is up there with the market leaders.

Step by Step

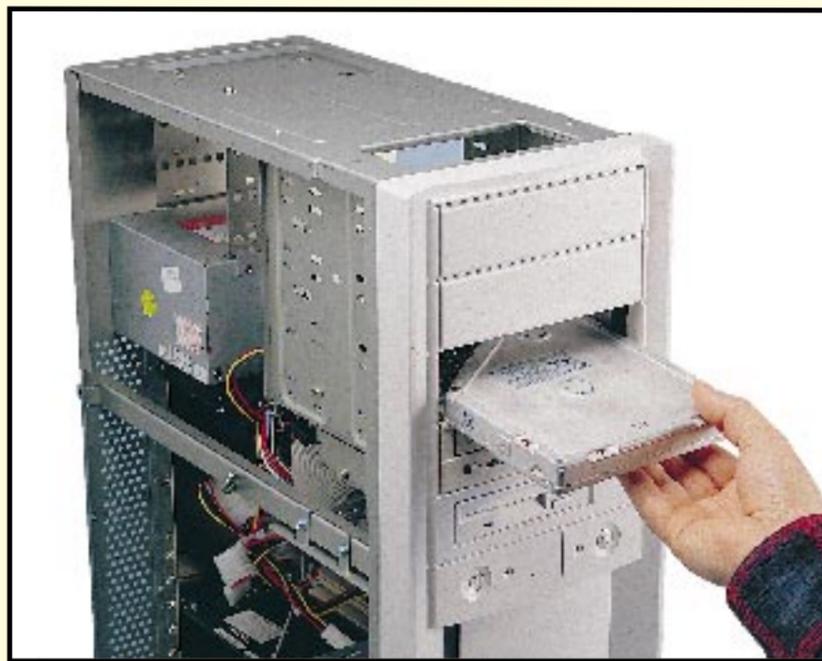
Preliminaries

Take the time to prepare a bootable floppy and make sure it actually boots beforehand. Copy these DOS utilities on to it: FDISK, FORMAT, SYS and SCANDISK. Don't forget, we won't be

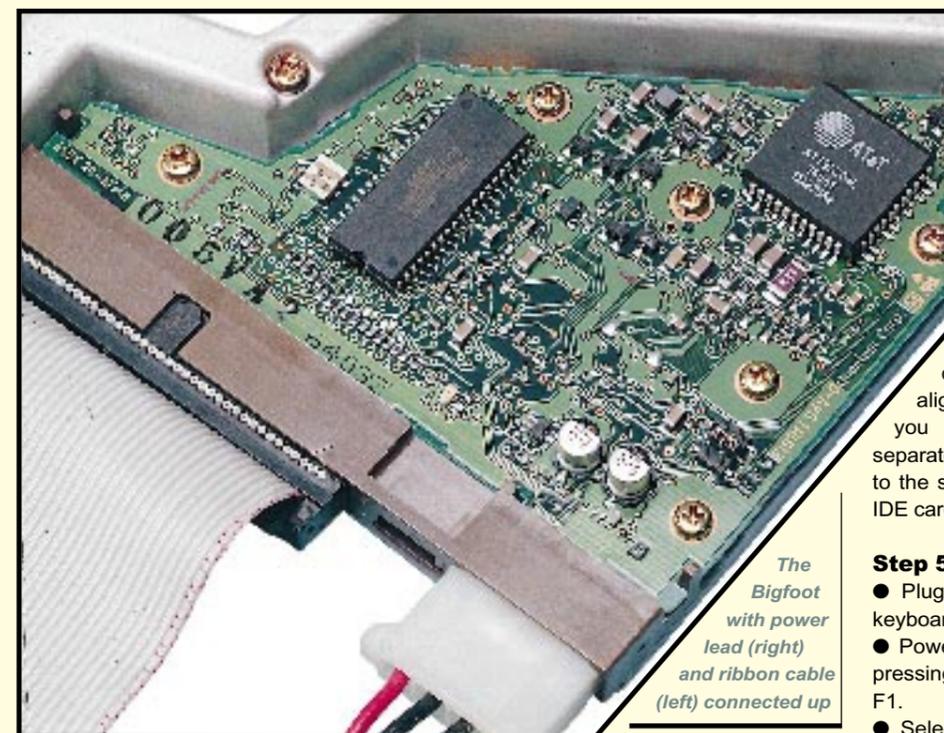
using the special driver software that often comes with large drives, as we'll be relying on our new EIDE interface card's on-board BIOS to support our new, large, hard disk directly.

Step 1

● Yes, I know it's boring and repetitious, but when you're dealing with electrical appliances you have to take safety precautions: so power down and unplug the PC from the mains, and disconnect all other leads, such as printers and mice. Before handling any electrical components, discharge any static electricity you might be carrying by



The Quantum Bigfoot is different to other EIDE drives as it will fill a 5.25in drive bay



The Bigfoot with power lead (right) and ribbon cable (left) connected up

the drive. Most modern ribbon cables are polarised to prevent you from inserting them the wrong way round, but others aren't, so look for a red or coloured stripe on the ribbon cable. This indicates Pin 1. Look at the socket on the drive (and on the interface card) for the Pin 1 label to correctly orientate the plug.

● Plug the other end of the ribbon into the primary E-IDE channel interface, making sure to align the coloured edge with Pin 1. If you have an IDE CD-ROM, use a separate 40-way ribbon cable to connect it to the secondary IDE channel on the E-IDE card.

Step 5

● Plug in the mains leads, video and keyboard cables, and power up the PC.
● Power up and enter CMOS setup by pressing the appropriate key, e.g. DEL or F1.

● Select the hard-disk option from the CMOS setup menu. Many modern BIOSes now offer auto-detection; they'll interrogate the drive's firmware to find out its values. Choose this option, if available, otherwise select Drive Type 47 or User defined. This allows you to plug in the drive geometry values manually; these will be detailed in the documentation and, most probably, on the drive itself. You'll be asked for the number of cylinders, heads and sectors per track, plus exotica like Write Pre-compensation and Landing Zone.

● Save the changes and quit Setup. Reassemble the PC, put the lid back on, do up the screws, plug all the cables back in and power up the PC.

Step 6

● We now have to partition the hard disk. This is done using FDISK, so boot from your previously prepared bootable floppy and load FDISK.

● Select 1 (Create DOS Partition) and 1 again (Create Primary DOS partition). Choose to make the entire drive one partition if that is what you want (but see overleaf).

● Quit FDISK and

earthing yourself: touch a metal pipe, for instance, or the PC's chassis; if you have an anti-static wrist strap, put it on.

● Undo the four or five screws and take the lid off of the PC.

Step 2

● Check out the drive fixings and where it's going to fit in the PC. You'll need a free drive bay for the new drive. The Bigfoot doesn't need any special mounting hardware but its 5.25in size means you'll have to put it in an externally accessible 5.25in drive bay, rather than a 3.5in internal bay, which is a minor niggle.

● Try and place it close to the EIDE interface as IDE ribbon cables tend to be short. Note that due to the very high data transfer rates made possible by E-IDE, you have to be careful about the ribbon cable used; for example, you mustn't use one longer than 18ins. Make sure you have the right mounting hardware, too, things like bolts or rails.

● Offer the drive up to the drive bay and make sure it doesn't foul anything else. Many modern drives list the geometry details on the label, so make a note of this before installation. If this drive replaces an older drive, remove it.

Step 3

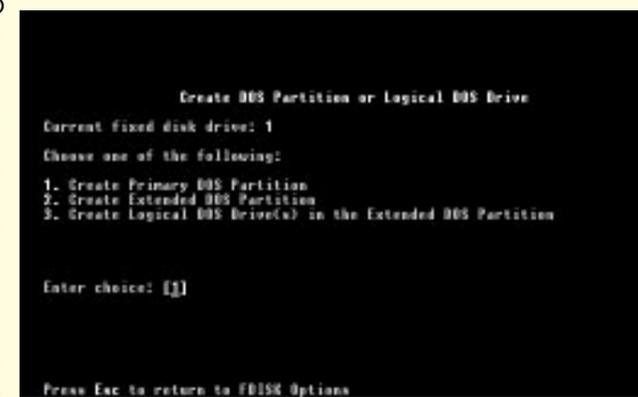
● As we're fitting a single drive, there's no need to move any jumpers on the drive and it should work in its

default configuration. However, if you're adding it as a second drive, one drive will have to be nominated as a master and the other as a slave. This will entail some jumper-shuffling on both drives, so have their respective user guides handy. Luckily, most modern drives now feature explanatory labels which describe the jumper settings and drive geometry, and this is true of the Bigfoot. At a pinch, if you can spare the second IDE channel, you can always leave the original drive as a master and simply plug it in to the secondary IDE channel.

Step 4

● Slide the drive into a vacant 5.25in bay and tighten up the mounting bolts. Insert a spare power lead (which can only be fitted one way).

● Now fit the 40-way data ribbon cable to



FDISK presents you with a list of simple menu options

Hard disk health and efficiency

If you read the section, *The importance of partitioning* (opposite), it's easy to see how storage efficiency, the ratio of wasted space (or overhang) to usable disk space, drops dramatically as cluster size increases.

Luckily, there is a very simple solution to this problem, and that is to partition the drive into smaller volumes, thus reducing the cluster size and raising storage efficiency in the process. This is easier said than done, however: repartitioning can normally be done on the fly and the process will zap all the data on your hard disk, so you'll have to back it up first and restore it later. This is doubly irritating for purchasers of new PCs (which are typically supplied unpartitioned) who might want to repartition right away: these PCs often come with lots of pre-installed software as well, such as Office Suite bundled software, which is sometimes supplied without master floppies. Another twist of the knife.

Partitioning is carried out using the old DOS utility, FDISK. Simply delete the original single partition and create a Primary DOS Partition just under one of the size thresholds listed here (i.e. 255Mb or 511Mb, rather than 256Mb or 512Mb). Next, create an Extended Partition occupying the rest of the drive, and from this carve out logical drives of 255Mb or 511Mb size. You then have to reboot and format each partition to make it usable.

So far, I've only come across one PC manufacturer, Gateway, that partitions new PCs hard disks. The P5-200 comes with a 2.5Gb hard disk and the one I looked at was

split into a 500Mb primary partition, containing all the program files, and a 2Gb extended partition. This was completely empty and so would be very easy to carve up with FDISK.

There are a number of alternatives to this painful route. Curiously, disk compression software such as DriveSpace and Stacker alleviates the overhang problem because it allocates file space on a per-sector basis. Stacker 4.0 is particularly efficient in this regard, as it allocates file space on a per-byte basis and so overhang is eliminated. It might seem daft to want to compress a huge disk merely to get over a DOS limitation, but there is a way around this. If you've got the Windows 95 Plus! Pack, it's possible to run DriveSpace 3 on a no compression setting on a drive partition up to 2Gb in size. This will give you high storage efficiency and minimal speed losses from running a disk compression utility.

Or, you can use PartitionMagic 2.0, which does allow you to adjust partition sizes on the fly. The latest version covers DOS and Windows 95 as well as OS/2. Written by US software developer, PowerQuest, it's available here in the UK from POW! Distribution (see PCW Contacts). Not only does PartitionMagic allow you to grow and shrink partitions at will, simply by grabbing an on-screen slider, it also overcomes the problem of file overhang and will dynamically resize disk clusters to smaller, more efficient sizes. It also increases the number of root directory entries to 1,024, which is useful under Windows 95 where long file names can

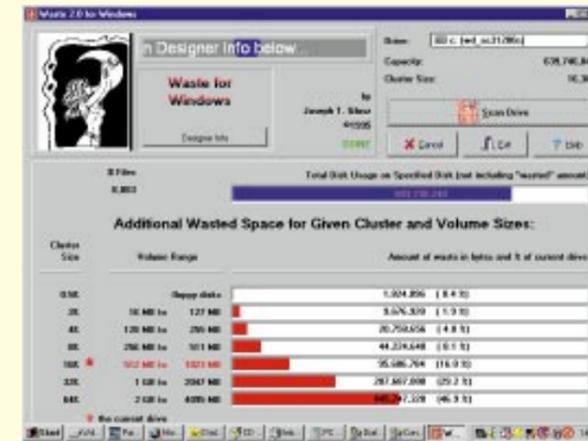
The importance of partitioning

But, that's not quite what you see, we've partitioned our new hard disk in a particularly daft way; as a single volume. While this undoubtedly makes for a simple life (after all, you only have to deal with just one drive letter), because the disk is so large, it also makes for very inefficient use of its space.

I illustrate this with an old chestnut I always trot out: a friend ran out of space on a 500Mb drive so he bought a new 1.1Gb drive, partitioned it as a single volume, and copied the contents of the 500Mb drive over to the new one. On this new drive, the data took up not 500Mb, as you might expect, but closer to 700Mb. *Shurely shome mishtake?*

The cause of this chronic loss of space can be traced back to the origins of MSDOS and Microsoft's decision to use a File Allocation Table or FAT file system. Originally, each entry in the FAT referred to a single 512-byte sector on the hard disk. Unfortunately, the FAT isn't open-ended and has a fixed maximum of 65,536 entries. This one-to-one FAT entry to disk sector mapping was okay for disks up to 65,536 x 512 bytes or 32Mb in size. In fact, up until MSDOS 4.0, 32Mb was the maximum size of hard disk supported by DOS. But how does DOS cope with disks bigger than 32Mb?

Simple: it breaks the one-to-one FAT entry to disk sector relationship by making each FAT entry represent *more than one* disk sector. So, for drives up to 128Mb in



size, each FAT entry represents a 2Kb cluster of four 512-byte sectors. And from there on up, as disk size doubles, so does cluster size: up to 256Mb it's 4Kb/eight sectors; up to 512Mb it's 8Kb/16 sectors, and so on.

Drive Capacity	No. of sectors	Cluster size
<128Mb	4	2Kb
128Mb to 255Mb	8	4Kb
256Mb to 511Mb	16	8Kb
512Mb to 1Gb	32	16Kb
>1Gb	64	32Kb

Note that both Windows NT 3.51 and OS/2 Warp use superior file systems (NTFS and HPFS) and so aren't afflicted with this problem.

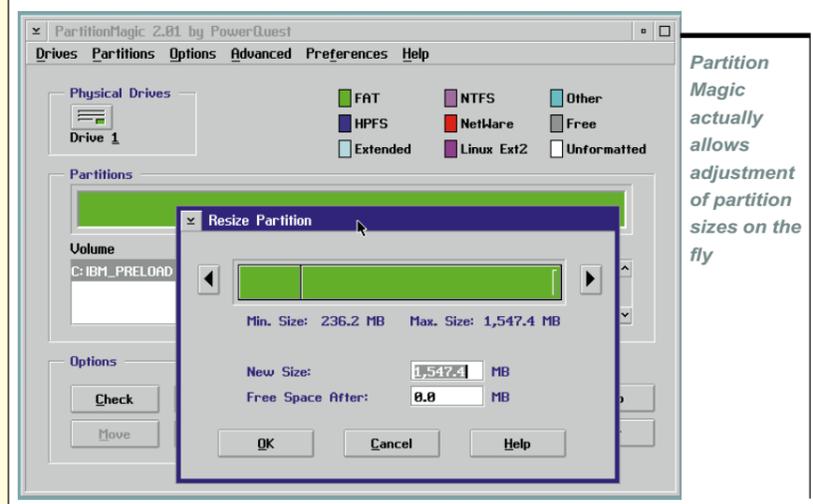
Here's how this arrangement becomes inefficient: under the FAT file system, disk space is allocated to files in

This month's cover CD has a selection of hard-disk utilities including *Waste for Windows*, which analyses your hard disk for wasted space

whole clusters, and a file can occupy less than a cluster. So, if you have a 4Kb cluster, a 1Kb file will consume one cluster and the remaining 3Kb of disk space in that cluster is unusable, which is tragic but acceptable.

But if you have a large hard disk, with, say, 32Kb clusters, this overhang of wasted space rockets to a massive 31Kb, which is very bad news indeed.

A 1Kb file is an extreme example, but on average, every file will waste half a cluster. So if you've got a 1.2Gb hard disk, this means that every file will be wasting 16Kb of disk space. When my aforementioned friend moved his data from a 500Mb drive to a 1.1Gb drive, he crossed two jumps in cluster size, from 8Kb to 32Kb, and in so doing lost 200Mb of disk space; about 20 percent of



Partition Magic actually allows adjustment of partition sizes on the fly

the PC will reboot.

Step 7

● We now have to high-level format it, using FORMAT: boot from your system floppy and use FORMAT C: /S to format

the drive and transfer the system files to it.

Once this has been done, remove the floppy and reboot the PC to ensure it actually boots from the hard disk. And that's it. Your new hard disk is ready and you can start installing software on it.



Magic card trick

A bigger hard disk is increasingly affordable, but it could cause compatibility problems. Roger Gann shows how you can use EIDE to overcome them.

Given the bloated nature of modern 32-bit operating systems and applications, these days it isn't so much a question of "if" but of "when shall I fit a bigger hard disk?" Luckily, big new hard disks have never been more affordable. Almost uniquely among computer components, they have simultaneously grown in size and speed while their prices have plummeted. Today, you can pick up a good-quality 1Gb drive for £160 or less – just two years ago I paid £700 (sob!).

The drop in price is great, but the overall increase in size can be a problem, at

least from the compatibility point of view. It's now becoming difficult to buy disks smaller than 540Mb and older machines can't directly support these larger disks without resorting to software trickery. This is no big deal for DOS+Windows 3.1x, but for Windows 95, which uses 32-bit protected-mode drivers, the presence of real-mode drivers can result in reduced performance.

A much more elegant solution is a low-cost hardware upgrade to give support for larger drives and let Windows 95 handle the drives directly. In this month's column

I'll take you through the process of upgrading to Enhanced IDE (EIDE) and I'll complete the piece next month by showing you how to install a new large hard disk.

Size solutions

There are several solutions to the IDE size problem (see page 295). Most large EIDE hard disks come complete with special device driver software that permits the PC to recognise the entire capacity of the drive. They work by placing a BIOS overlay in the boot sector, which is the same for all IDE drives, regardless of size. The overlay is loaded as soon as the drive boots, and it provides the support for large partitions that is missing from an older BIOS.

The down side of this solution is that it won't give you a second IDE channel, which is useful for connecting up slow IDE devices such as CD-ROMs and tape streamers. Also, it's likely that Windows 95 will run its file system in the slower MSDOS compatibility mode. Even under Windows 3.1x, a special driver is still required, to restore both 32-bit Disk Access and 32-bit File Access. Another more worrying problem is that if for any reason you lose the special driver, DOS won't be able to "see" the big hard disk.

A better option is to buy an inexpensive EIDE interface card. These cost about £35 and are a direct replacement for the £10 multi I/O cards fitted in most popular PCs. Examples include the VL-Bus Promise EIDE2300-Plus and the ISA Bus Future Domain IDE-16042. Both include two IDE channels and a floppy-drive interface. The Promise also features enhanced serial and parallel ports, which is a useful upgrade if you have a fast modem or a printer that needs a

"bi-di" printer port. And, of course, they both include an on-board BIOS that supports enhanced IDE: this supplements your existing BIOS.

The only drawbacks I can think of is that you'll possibly lose an expansion slot, and the additional BIOS steals a little (16Kb) Upper Memory.

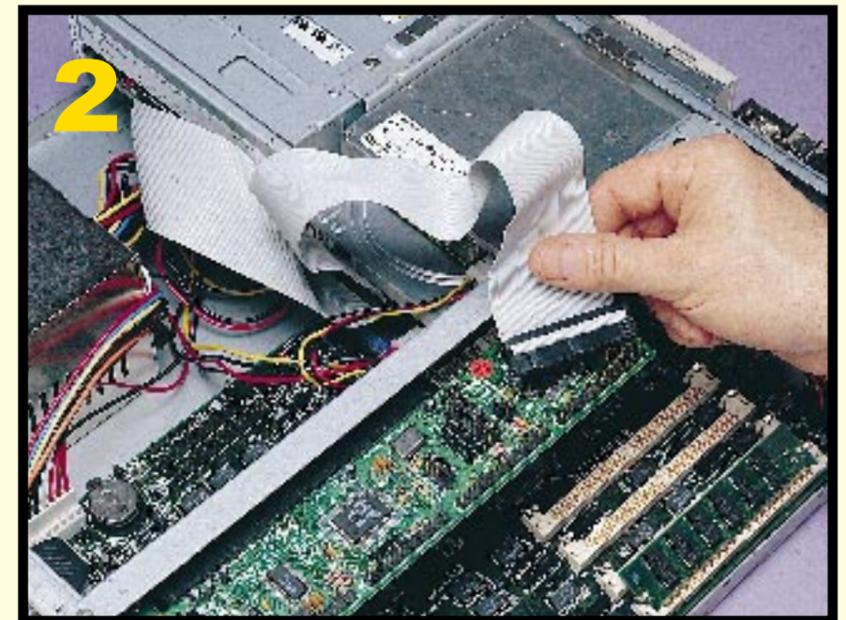
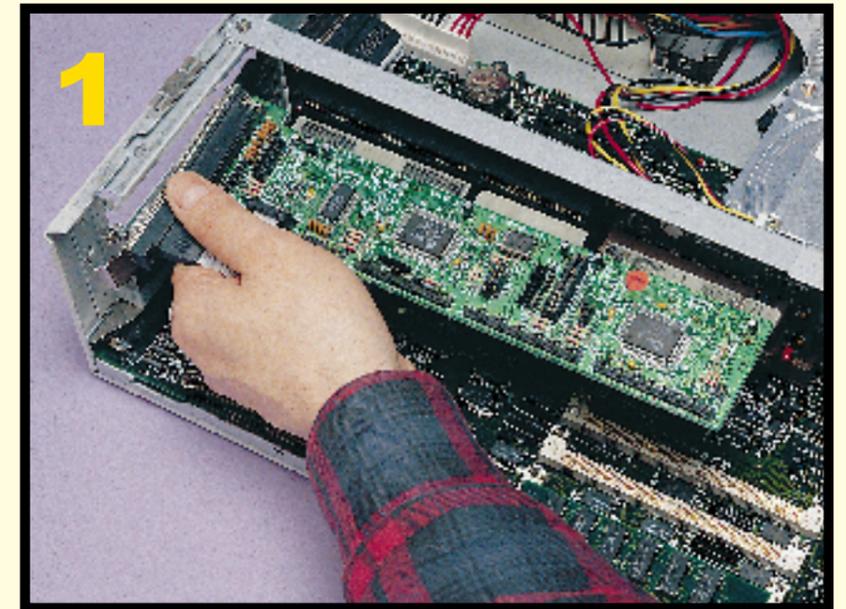
Another good reason for going down the EIDE card route is that you gain a second IDE channel. This doubles the number of IDE devices supported to four: you could, in theory, hang four IDE hard disks off such a card, with a master and slave drive attached to each channel. However, a more likely scenario would be a pair of Enhanced IDE hard disks on the primary or EIDE channel and slower IDE devices, designed for the ATA Packet Interface (ATAPI) standard such as CD-ROM drives or tape streamers attached to the secondary or IDE channel.

The received wisdom is that it's a good idea to give IDE devices with wildly differing data transfer rates separate IDE channels. At its peak, an EIDE drive using Mode 4 can transfer data at rates as fast as 16Mb/sec; a typical quad-speed CD-ROM drive knocks it out at 600Kb/sec. The lowest common denominator rule applies again here and there will be occasions when the data transfer rate of the hard disk will be pulled down to that of the CD-ROM, especially where both are in use simultaneously, for example when you're installing software from CD-ROM. With 8X CD-ROM data transfer rates now approaching 1.2Mb/sec and beyond, this speed differential is set to become less of a problem, but it's still something most of us have to watch out for.

Yet another benefit of EIDE is speed. A new EIDE interface card will support the lightning fast data throughput of the new drives. There are four types, or modes, of Polled I/O or PIO transfers and the fastest drives now support Mode 4 which, as we've seen, can offer data transfer rates as high as 16.6Mb/sec. The most your existing Mode 0 IDE can manage is 3.3Mb/sec.

A word of caution here: these are maximum burst rate figures, and the actual data throughput for Mode 3 and Mode 4 drives will be significantly lower than those numbers indicate. Note too that to upgrade to EIDE you'll need a motherboard with a local bus. These speeds demand a 32-bit VL-Bus or PCI local bus; the 16-bit 8MHz ISA bus can't support these levels of throughput.

So, as a prelude to upgrading your hard disk, we'll replace your existing IDE interface card with a new EIDE interface



(1) Here, a new EIDE disk controller is fitted into place. This one is designed for a VL-bus

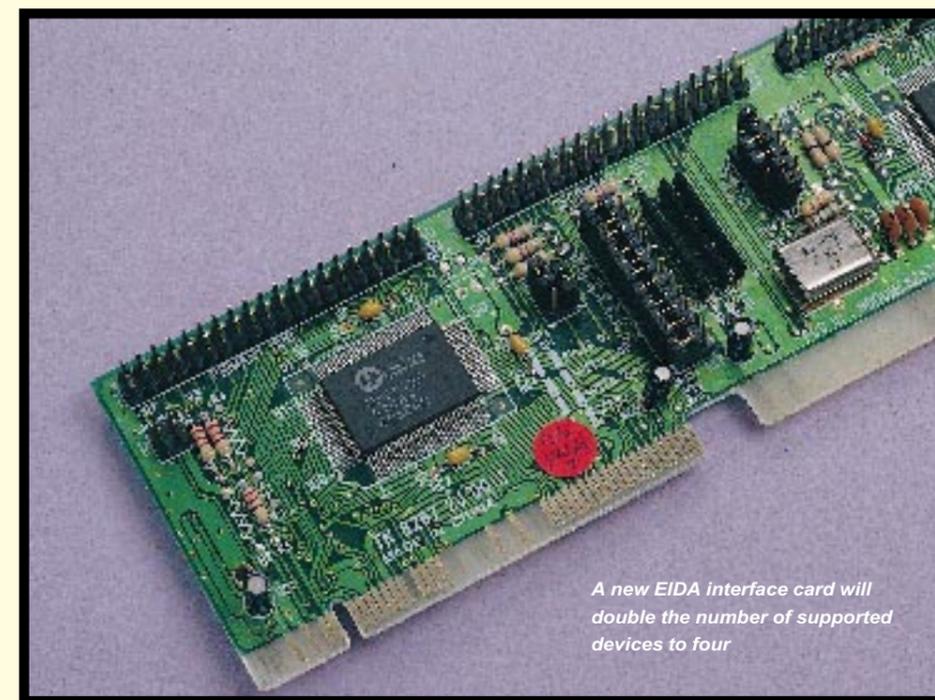
(2) It is possible to daisy-chain the hard disk and CD-ROM drive together on the same channel, but this is not recommended due to their different data transfer rates

Step-by-Step

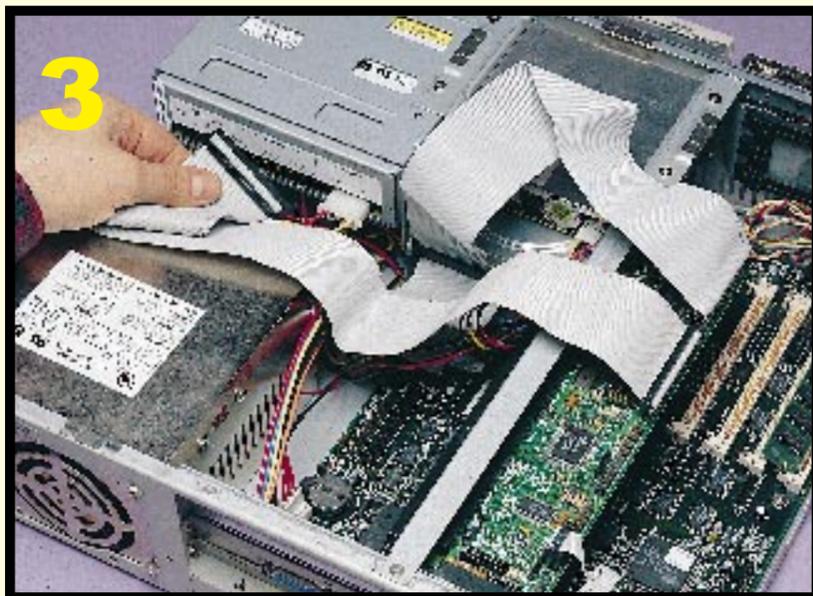
INSTALLING ENHANCED IDE

Step 1

● Make the usual preparations: clear a work space and get out your Phillips screwdriver. Power down and unplug the PC from the mains and disconnect all other leads, such as printers and mice. Before handling any electrical components, discharge any static electricity you might be carrying by earthing yourself: touch a metal pipe, for instance, or the



A new EIDE interface card will double the number of supported devices to four



PC's chassis. Undo the four or five screws and take the lid off the PC.

Step 2

● Locate the existing IDE interface card and disconnect the ribbon cables from it. You may also have to unplug any external serial ports. Undo the screw holding the card in and extract the card from its slot. Put it somewhere safe. If your PC features I/O integrated in the motherboard, there's nothing to remove, but you'll have to disable at least the on-board IDE interface, and more if the new card has a floppy interface and serial and parallel ports.

Step 3

● Configure the new EIDE card before you fit it. These early cards are definitely not plug-and-play; they are often festooned with jumpers and may require adjustment to avoid clashes. For example, the EIDE2300Plus defaults to using DMA Channel 1, which is the same as the default setting of a SoundBlaster card. If you have one of these you should use the alternative, DMA 3, and move the appropriate jumpers. Also, don't forget that the second IDE channel steals another IRQ, typically 15, so make sure this IRQ is free. When you're done, find a spare slot for the new EIDE interface card and insert it. Tighten up the fixing screw.

If you're adding a PCI EIDE controller, note that most PCI systems can't map IRQs 14 and 15 to a PCI slot. Since those interrupts are typically needed by a disk controller, you should have a workaround, which comes in the form of a ribbon cable and paddle card. Luckily,

(3) *The CD-ROM drive is generally fitted into the secondary IDE channel on the EIDE card*

this can plug in to the otherwise unusable ISA "shared" slot, so you don't lose a slot. This card allows the controller to tap interrupts 14 and 15 through the ISA connection. Some newer PCI motherboards offer a separate connector, called a legacy connector, near one of the PCI slots to provide the needed IRQ. For these systems, you simply connect the ribbon cable from the EIDE card to the legacy connector rather than to a second ISA card.

Step 4

● Now refit the hard disk drive's 40-way ribbon cable. Modern ribbon cables are

Explanation of acronyms and terms used

● Hard drive and peripheral interface standards

BIOS	Basic input/output system.
EIDE	Enhanced integrated drive electronics.
IDE	Integrated drive electronics.

● Other terms

ATAPI	AT Attachment Packet Interface.
CMOS	Complementary metal-oxide semiconductor.
DMA	Direct Memory Access.
IRQ	Interrupt Request Line
PIO	Polled (or programmed) input/output.

The IDE ceiling

While a PC's standard BIOS is theoretically capable of supporting hard disks as large as 8Gb and IDE drives can reach 130Gb in size, the real size limitation is just 504Mb. This low figure is the result of a clash between the drive geometries (the number of cylinders, heads and sectors per track) supported by each hardware standard. Unfortunately, these differ substantially and DOS defaults to the lowest common denominator in each case, hence the 504Mb limit. The table below highlights the problem:

	BIOS	IDE	Lowest value
Max. sectors per track	63	255	63
Max. heads	255	16	16
Max. cylinders	1,024	65,536	1,024

The maximum size of a hard disk drive is derived from multiplying these values, (sectors x heads x cylinders) by 512 (the number of bytes in a sector), so you'll end up with the following capacities, in bytes:

8,422,686,720
136,902,082,560
528,482,304

Or, to express these values in megabytes, we divide them by 10,242 or 1,048,576:

8,032.5Mb
130,560Mb
504Mb

And that's how we wind up with the half gigabyte IDE limitation. Two years ago it wasn't a problem as IDE drives were no bigger than 420Mb and only SCSI drives were larger. As they used their own BIOS, the limit didn't apply to them. However, with today's large IDE drives, the 504Mb limit is a very real problem for older PCs.

polarised with a notch to prevent you inserting them the wrong way round, but some aren't, so look for a red or coloured stripe on the ribbon cable. This indicates Pin 1. Look at the socket on the new interface card for the Pin 1 label to correctly orient the plug. Be sure to plug it in to the E-IDE/primary (and not the IDE) port on the EIDE interface card, making sure to align the coloured edge with Pin 1.

Because of the very high data transfer

rates possible thanks to EIDE, you have to be careful with the ribbon cable and you mustn't use one longer than 18 inches. Don't forget to plug in the 34-way floppy ribbon cable on to the floppy interface on the EIDE card, making sure to get it the right way round. Reconnect any serial or game port cables, too.

Step 5

● If you have an IDE CD-ROM, now's the

time to give it its very own IDE channel. Use a separate 40-way ribbon cable to connect it to the secondary IDE channel on the EIDE card. Don't forget to set it to be a "master" IDE device if it has been set to a "slave" previously (though in my experience, this doesn't seem to matter too much). Do this by moving a jumper at the back of the drive.

Step 6

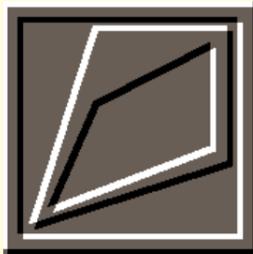
● Reassemble the PC, put the lid back on, do up the screws, plug all the cables back in and power up the PC. Check the CMOS setup and get the BIOS to auto-identify the hard disk just to make sure that all the cabling has been installed correctly. Some EIDE cards may let their own BIOS handle the hard disk totally and insist therefore that you set the Drive Type on the PC's BIOS to "Not Installed". Your PC should reboot at this point. If it doesn't, re-check the cabling, especially the Pin 1 process.

Step 7

● Now's the time to install the card's Windows 3.1x and Windows 95 drivers. For Windows 95, fire up Control Panel, click on the "Add New Hardware" icon and opt for the manual install. Select "Hard Disk Controllers" and click on the "Have Disk" button. Place the driver disk in the floppy and click on "OK". When that's done, restart Windows 95. ■

PCW Contacts

Roger Gann can be contacted either by post c/o PCW or via email at rgann@mcgilivray.win-uk.net



Processor push-ups

Give that lazy old PC a kick up the socket with a processor upgrade. Roger Gann tells you how to fortify the 486s.

PCW Chip Photography by Graham Pearson

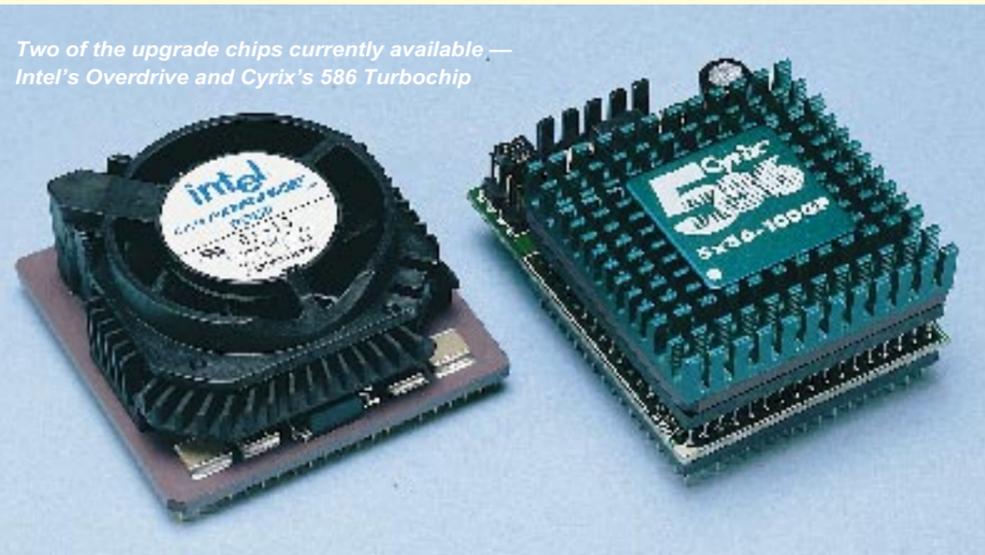
My steely gaze alights this month on processor upgrades. As sure as night follows day, the PC you buy today will inevitably appear to run slowly tomorrow, weighed down by the ever-increasing burden of running the latest 32-bit software and operating systems. Luckily, for the vast majority of PC owners it's possible to revitalise their sluggish PCs by upgrading their existing processor and replacing it with a more powerful one. Luckier still, this is a relatively simple task.

This month I'll be looking at upgrading your CPU and divulging the odd tip or two. I'll concentrate on upgrading 486s as not only are they the most plentiful but they also have the most pressing need of upgrading. Although the 486 is now technically obsolete, the world is awash with them: according to Intel, there are already 50 million upgradable 486s out there. This is no bad thing from an upgrading point of view, because the 486 family was designed from the outset to be upgradable.

Sockets

The key to processor upgrading is that most, if not all, 486 processors are socketed. This means that it's a simple task to extract the old processor from its socket and plug in a replacement. Even those

Two of the upgrade chips currently available — Intel's Overdrive and Cyrix's 586 Turbochip



Overdrive vs "loose"

You don't have to look too far to see that plenty of dealers these days are selling both ordinary and Overdrive processors. Internally, these chips are almost identical but they carry quite different price tags: an Intel 486DX2/66 Overdrive might go for £85 while just £22 will buy you a ST486DX2-66-GS CPU made by the French chip giant, SGS.

The two chips are, to all intents and purposes, identical but one is almost four times the price of the other. Surely only a fool would buy the dearer of the two? Well, when confronted with bargains like this, you have to bear in mind the reason for the price difference: you get more if you plump for the Intel.

The cheaper CPUs are invariably sourced from the OEM (original equipment manufacturer) market and so aren't meant to be sold in the retail market — they are intended to be installed by PC manufacturers and not end-users. They are supplied in bulk, with little or no packaging, no instructions and no support: when you buy one of these, you're on your own.

With the Overdrive you get full instructions, tech support, software, a three-year warranty and a money-back guarantee. You also get special versions of the chip that will allow a 3.3v Overdrive CPU to run in a 5v socket, for example.

If you know what you're doing, CPUs that are sold "loose" are a bargain. But if you're new to the processor upgrade game it's best to play safe and go down the Overdrive route.

Supercharging an Overdrive

Intel has gone to considerable lengths to make the potentially fraught matter of processor upgrading as simple as possible. It *must* be as uncomplicated as possible if it's to be a retail product, likely to be installed by a non-technical end-user.

At its simplest, all you have to do is take out the old CPU and bung in the new one. At worst, you might have to move some motherboard jumpers. The down side of this approach is that it limits your choice somewhat: if you had an old 25MHz 486SX and wanted to fit a Pentium Overdrive, you could only buy the 63MHz version because the faster, 83MHz version is meant for 33MHz motherboards. Or, if you were looking at a DX4 you'd have to fit the 75MHz version rather than the 100MHz one.

Nevertheless, most motherboards (even quite old ones) were designed to take a variety of processors and so were able to run at a variety of speeds, typically 25MHz or 33MHz. These so-called "clock speeds" are normally determined by a set of jumpers on the motherboard, so if you could adjust the clock speed from 25MHz to 33MHz you'd be able to fit the 83MHz Pentium Overdrive rather than the slower 63MHz version. Sure, it would cost you more, but you'd be getting a faster PC at the end of the day and thus it's definitely worth paying the extra £30 or so. However, setting the motherboard speed jumpers can be fiddly and you'll need the assistance of your motherboard handbook to tell you which ones to move, but the hassle is worth it in the long run.

What performance gains can you expect?

Simply fitting a more powerful processor doesn't commensurately increase the overall power of your PC — it'll have a greater impact on some tasks than others. Yes, you do get gains from upgrading your processor but nowhere near as great as you might have hoped for.

Synthetic benchmarking may reveal integer and floating-point performance improvements of as much as 100 percent as a result of such an upgrade, but the real-world improvements you can expect will be much lower: in the 20 to 30 percent range.

Other factors such as installed memory, hard disk and graphics have just as great an influence on overall performance and will dilute the apparent gains to be had from installing a processor upgrade. So if you do a lot of processor-intensive tasks, such as spreadsheeting or multimedia, then provided the price is right, a processor upgrade is worth considering as a PC's mid-life booster.

handful of 486SX processors that were soldered down were, typically, additionally provided with an empty upgrade socket.

The original Pin Grid Array (PGA) socket held the CPU in by friction, which meant that the chip didn't just lift out: the unit and its 168 pins had to be levered out. Normally, the pin holes in a traditional processor socket tapered slightly to ensure a good electrical contact; the chips were effectively tightly held in by friction, hence their reluctance to be removed. Luckily, every Overdrive kit comes with a high-tech crowbar (chip-puller) to prise old 486s out of old PGA sockets.

To make CPU upgrading a lot easier, even though it's a task you'll probably only ever do once or twice, Intel designed the Zero Insertion Force (ZIF) CPU socket. A ZIF socket is a little larger than a normal PGA socket and uses a lever, or handle, to clamp the chip pins tightly. To remove the chip, you unclip the lever and lift it. This unclamps the CPU which then lifts out very easily — with zero force, in fact. Most 486s and all DX4s and Pentiums feature ZIF sockets.

There are several sorts of ZIF socket. The type of socket will determine what processor upgrade options are available, so it pays to take a peek under the bonnet to see just what sort of a socket you've got.

There are four types of ZIF socket installed in 486-based PCs:

Socket 1 A 169-pin, blue, ZIF socket; this can only accept 486 Overdrive processors.
Socket 2 A 238-pin, blue, ZIF socket; this can accept both 486 and Pentium Overdrive processors.

Socket 3 A white, 237-pin, ZIF socket; this is essentially the same as Socket 2 but can take the low voltage 3.3v DX4 and Pentium Overdrive. The "missing" pin is used to correctly orientate the square chip.
Socket 6 A 235-pin socket for DX4 processors.

Pentium-based machines have even larger ZIF sockets:

Socket 4 A white, 273-pin, ZIF socket; this is used for Pentium 60 and Pentium 66 processors.

Socket 5 A white, 320-pin, ZIF socket; this is used for Pentium 75 processors and above.

Socket 7 A white 321-pin ZIF socket; the current standard socket for the entire Pentium range. It can take future Pentium upgrade processors.

Socket 8 The largest ZIF socket of all, for the Pentium Pro.

The upgrade choice

So given the right socket, the range of upgrade processors you can fit is quite wide. If you've got a 486SX/25 with a Socket 3 you'll be able to choose between four Overdrives: the 486SX2, which is a speed-doubled version of the SX; the 486DX2, which is a speed-doubled CPU with maths co-processor; the DX4, which is a speed-trebled processor with maths co-processor; or, a 63MHz Pentium Overdrive.

Note that Intel is in the process of phasing out the slower Overdrives and you'll have to move fast to snap up any of the 486SX2 or DX2 models before the summer. That leaves the 75MHz and 100MHz DX4 and the 63MHz and 84MHz Pentium Overdrives. Their prices start at £112 for the entry-level DX4. The 100MHz DX4 and the 63MHz Pentium Overdrive are £145 each and the top-end Pentium Overdrive is £209, but street prices are much lower.

And what about Pentium owners? Although it's always been possible to upgrade Pentium 75s and above, simply by replacing them with faster, "loose", OEM (original equipment manufacturer) versions sold in the retail market, this luxury was denied to owners of older 60MHz and 66MHz Pentiums. This has now been rectified with the recent release of the latest clutch of Pentium Overdrives, which cater for 60, 66, 75 and 90MHz Pentiums. It costs £246 to buy a 120/133MHz Pentium Overdrive — 125MHz in the case of 75MHz upgrades (reviewed in the April issue of PCW).

Alternatives to Intel

But don't think it's Hobson's choice when it comes to CPU upgrades — there are alternatives to Intel. Cyrix has left the upgrade market *per se* but its processors turn up in third-party upgrade products — the new 5x86, for instance, turns up in the PowerLeap/586 (contact Future Upgrades; see PCW Contacts panel). This is as powerful as the Pentium Overdrive 83 but at £95 is less than half the price of its Intel rival. The catch is that the 5x86 is

How to upgrade your processor — see over page



so modern, it's only suited to the more recent motherboards. Kingston Technology does a range of processor upgrades as well, the "Turbochip" range (contact Datrontech; see *PCW Contacts* panel).

Step-by-Step

UPGRADING YOUR PROCESSOR

When handling processors it's particularly important to remove any static electricity you might be carrying by earthing yourself — don't forget, you'll be handling something expensive that's easily zapped by static!

Step 1

- Power down and unplug the PC from the mains and disconnect all other leads.

Step 2

- Take the lid off the PC. It'll be held on by four or five self-tapping screws and you'll most likely need a Phillips screwdriver to undo them. Keep them in a safe place.
- Locate and identify the 486 processor. It might be a good idea to remove some or all of the expansion cards to give yourself more space.

Step 3

- Note the orientation of the printing on

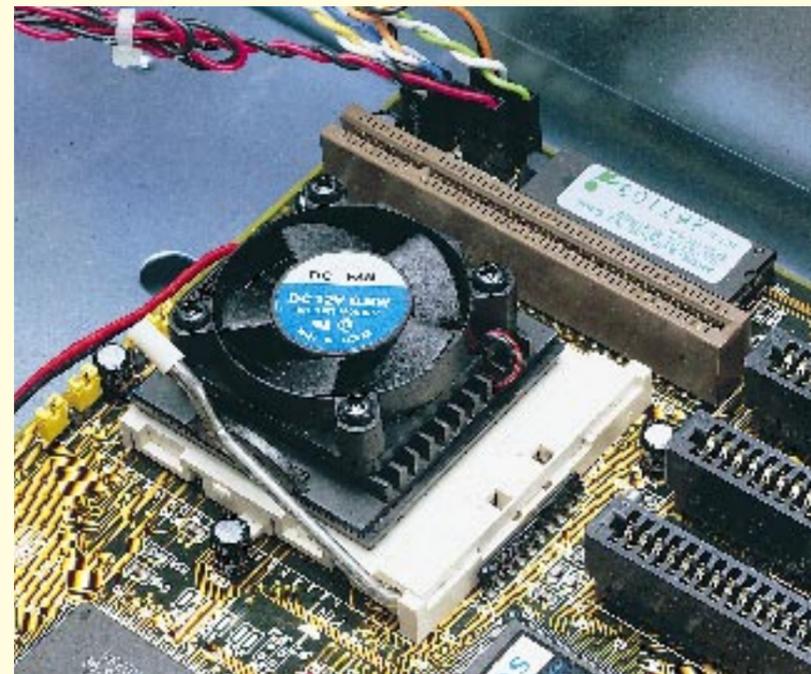
The bevelled corner of the new chip must be lined up correctly. Here, Pin 1 is marked on the motherboard

the old processor: it may help you orientate the new chip.

- Remove the processor. If it's in a ZIF socket, simply unclip the release lever and lift it up — you can now lift out the processor. If it's in a PGA socket, use the special lever supplied in the upgrade kit to prise it out — *be careful to insert it between the socket and the chip and not to put it under the socket*. Use gentle pressure and lift all four sides of the CPU evenly. Put the old CPU somewhere safe.

Step 4

- Because all current processor upgrades are square (i.e. symmetrical) it's quite possible to attempt to fit it into the socket in any of four ways, so take care to correctly orientate the new processor. This is easy: the socket will have a bevelled corner, or mark indicating the position of the corner pin; and the CPU will have a corresponding dot, or bevelled corner.
- Marry the two marks up and insert the chip, quickly checking beforehand that none of its pins are bent. Note that some Overdrive chips use an additional pin to orientate the chip in the socket, which makes it impossible to fit it the wrong way. If the ZIF socket is of the large Pentium Overdrive sort, insert your 486 Overdrive in the centre of the socket so that a line of



pin holes remains visible on all four sides of the ZIF socket.

- Simply drop in the CPU, lower the lever and clip it into the locked position. If it's a PGA socket then you'll have to use some force to insert the chip — make sure the chip is level and goes in "straight" and evenly. Be very careful, too, not to overflex the motherboard. You may need to shift a few motherboard jumpers at this point to identify the new CPU, so have that motherboard manual handy!

The lever on the ZIF socket allows you to lock the new chip into place

Step 5

- Reassemble the PC, replace the lid, do up the screws and plug everything back in.
- Plug in all the cables and power up the PC. You'll soon know if the new processor has gone in okay — if it hasn't, the PC won't boot.

Explanation of acronyms and terms used

Low level

CPU Central processing unit.

Hard drive and peripheral interface standards

BIOS Basic input/output system.

Other terms

Clock speed The speed, in MHz, at which a microprocessor runs.

- If you have a "verbose" BIOS, which signs on with lots of information, it may tell you what sort of processor it recognises. To be extra sure, install the supplied diagnostic software that came with the upgrade just to make sure everything is A-OK.

Further reading

See *First Impressions* for a review of the MAke-it 586, from Improve Technologies, vs Kingston Technology's TurboChip 133.

PCW Contacts

Roger Gann can be contacted either by post c/o PCW or via email at rgann@mcgilivray.win-uk.net

Datrontech 01252 303333
Future Upgrades 01732 465566



Stepping up to CD

Roger Gann takes you through the basics of installing a CD-ROM drive.

PCW Step by Step Photography by David Whyte

A CD-ROM is no longer a luxury, only needed to play exotic and resource-hungry games. With today's large applications and even larger operating systems, it's easy to argue that it's now a necessity: who wants to install something that comes on 30 floppy disks?

So this month I'll show you how to fit a complementary peripheral: a CD-ROM drive. More and more software is being supplied, as standard, on CD-ROM. Corel-Draw 6.0, for example, comes on four CD-ROM discs.

It's quite possible to justify investing in an inexpensive CD-ROM drive for the sole purpose of installing software. As little as £50 will buy you a quad-speed IDE drive, so they've never been more affordable. Besides, you need something on which to play our brilliant cover disks!

Step-by-Step

HOW TO INSTALL A CD-ROM DRIVE

As with fitting a sound card, installing a CD-ROM entails fitting the drive then installing the software drivers.

Hardware installation

Step 1

- As ever, take the usual safety precautions before opening up your PC — gather together all the relevant instructions and installation disks, together with all the tools you need.
- Power down the PC, remove the casing lid, and enter within.

Step 2

- Now's the time to configure the CD-ROM drive, if it needs it.
- If it's being connected to a (primary) IDE interface then it most probably doesn't require configuring at all: most will work just fine with their default settings, which is usually "slave".
- However, if you're installing it on the secondary IDE channel, then set it to "Master".
- As far as a SCSI drive is concerned, check to make sure that its SCSI ID is unique and doesn't clash with any other SCSI devices installed.
- If it's the "last" device in the chain you may have to "terminate" the connection (and "un-terminate" the previous device).

Step 3

- Check out the drive fixings and the

place where it's going to fit — you'll need a free 5.25in drive bay for the new drive.

- If possible, try not to mount it in the lowest drive bay, as the disk tray or caddie could foul your keyboard. Mounting it high also makes it easier to get at the disk eject button with the tray open.
- If it's an IDE/ATAPI drive try to keep it close to the IDE hard disk, as IDE ribbon cables tend to be short.
- Make sure you have the right mounting hardware, too: things like bolts or rails.
- Insert the drive into the drive bay and make sure it doesn't foul anything else.
- Tighten the mounting bolts but don't over-tighten them.

Step 4

- With the driver installed in its bay, attach a spare power cable and the 40-way ribbon cable to it. Some connectors



The front panel of a CD-ROM

- 1 Loading bay/access door
- 2 Headphone jack
- 3 Volume control
- 4 Status indicator
- 5 Play Next Track button
- 6 Stop/Open/Close button

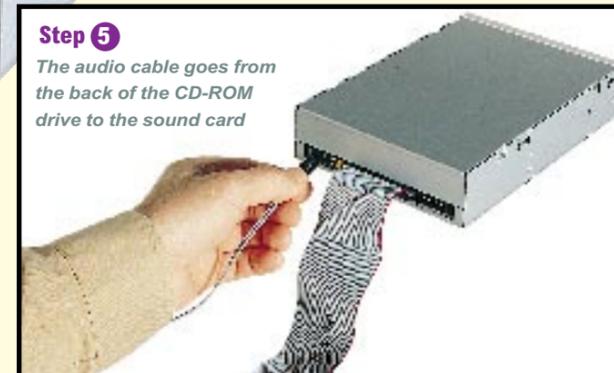
The rear panel of a CD-ROM drive

- 1 Digital output connector
- 2 CD/Audio connector
- 3 Master/slave configuration jumper
- 4 Interface connector
- 5 Power connector



Step 4

Here, the cable's red stripe indicates alignment with pin 1



Step 5

The audio cable goes from the back of the CD-ROM drive to the sound card

are "polarised" with a little notch and can't therefore be fitted the wrong way round, but 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 (normally the side next to the power connector).

- If you're lucky, the pins will be numbered. Likewise, make sure it is correctly orientated at the interface-card end. If you're connecting to a SCSI card, you'll have a similar but wider (50-way) ribbon cable.

Step 5

- 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. While it's not necessary to have this cable to be able to play sound from a video or game, it will be necessary if you want to play an audio CD through your setup.

Step 6

- Power up the PC. The POST diagnostics won't report anything new, but if you've fitted a SCSI CD-ROM, the SCSI BIOS will recognise the new addition and list it, along with all the other SCSI devices. Either way, you don't have to make any alterations to your CMOS setup.
- The drive's activity LED should flash, signifying the drive is receiving power.

Software installation

With the hardware installation complete, we now have to get the operating system to recognise the new drive and to assign it a drive letter.

- To do this, we need to install driver software, though if you use Windows 95 and don't need to boot in DOS mode, you can rely on Windows 95's own 32-bit internal drivers.

Your CD-ROM drive will come with a driver disk. This will normally install two files: a driver, specific to that drive, which is loaded in CONFIG.SYS; and a generic DOS CD-ROM Extension, loaded in AUTOEXEC.BAT.

- Because the actual name of the driver will vary from drive to drive, I'll have to generalise.

The line in CONFIG.SYS 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 (see below).

- If you're running a SCSI system, you'll use a generic CD device driver. If you were using an Adaptec SCSI host adaptor, you'd add a line like this instead:

```
DEVICE=C:\ASPICD.SYS /D:MSCD0000
The line in AUTOEXEC.BAT might look like this:
```

```
MSCDEX /D:MSCD0000 /M:10
```

Note the matching "signature", MSCD0000. The "/M:10" switch specifies the number of sector buffers to allocate to the CD-ROM — a crude kind of cacheing.

- There are other switches you might want to use:

/S: Enables sharing of CD-ROM drives on Windows for Workgroups servers.

/L: Specifies the drive letter to assign to the first CD-ROM drive, so using /L:G would assign the drive letter G to the CD-ROM drive.

- Although Windows 95 has its own 32-

Choosing a CD-ROM drive

Apart from your budget, there are three factors you should consider before buying a drive.

1. The first is choice of interface. There are now just two to choose from: either IDE (also known as ATAPI), or SCSI.

Most PCs will already have an IDE interface card, so fitting an IDE will probably be the cheapest and simplest solution. However, if your PC lacks an Enhanced IDE interface, with dual IDE channels, then you really should invest in one — they cost between £20 and £40.

While it's perfectly possible to slave an IDE CD-ROM to an IDE hard disk, it's not really recommended as it can slow down your hard disk. The "lowest common denominator" rule means that with a CD-ROM drive slaved to a hard disk on the same IDE channel, you'll reduce the hard disk's data transfer rate to that of the CD-ROM drive: that is, 600Kb/sec if you've got a quad-speed drive — an IDE hard disk's data transfer rate will normally be at least four or five times this. With a dual channel EIDE interface card you can connect the hard disk to the faster EIDE, or primary channel, and the CD-ROM to the slower IDE, or secondary channel.

For more demanding use, I'd recommend opting for the dearer SCSI alternative. This is because IDE typically uses Programmed I/O to manage the data transfer, which requires a high degree of CPU attention — many so-called "six-speed" drives can only hit their specified

data transfer rate by taking 100 percent of the processor's time, which leaves zero for anything else.

SCSI, on the other hand, offloads the task of supervising the transfer of data from the CPU to the host adaptor and so is very suitable for computer-intensive tasks such as software MPEG decoding. So, if you're going to use your CD-ROM for serious video playback, I'd look at SCSI. Most current CD-R writers are SCSI, which is another reason for picking this interface.

2. The second factor is data transfer rate. Most drives are now either quad-speed (600Kb/sec) or six-speed (900Kb/sec). Older (double-speed) drives are no longer made).

Most CD-ROM multimedia discs are still cut to run on humble double-speed drives, but an increasing number are now being mastered to run on quad-speed drives. The discs therefore won't run well on older drives. So from a multimedia point of view, a quad-speed drive will be adequate for the foreseeable future.

Consider buying a six-speed drive *only* if you need its superior average access speed and data transfer rate for database searching chores, where the extra turn of speed will be noticeable, especially if the drive is shared on a network.

3. The final consideration is a practical one to do with loading — do you want a tray-loader drive or one that requires a disc caddy? It's horses for courses, but I prefer the former simply because it's quicker.

bit protected mode CDFS CD-ROM file system and doesn't need real-mode DOS drivers, it helps to install the older drivers first because the "Add New Hardware" auto-detection wizard will first look at the contents of the startup files before it tries to sniff out the devices.

If it finds a CD-ROM device driver 

Explanation of acronyms and terms used

● Low level

CPU Central processing unit.

● Hard drive and peripheral interface standards

BIOS Basic input/output system.

EIDE Enhanced integrated drive electronics.

IDE Integrated drive electronics.

SCSI Small computer system interface.

● Other terms

ATAPI AT Attachment Packet Interface.

CDFS Compact Disk File System.

CMOS Complementary metal-oxide semiconductor.

MPEG Moving Picture Expert Group.

MSCDEX Microsoft Compact Disk Extension. (A TSR program giving DOS access to CD-ROM drives.)

VCache Virtual cache.

Wizard A Microsoft term describing online help which guides you through tasks on the fly.

there, it knows one should be physically present. In actual fact, Windows 95 can only specifically detect the older types of CD-ROM drives that used proprietary interfaces, such as Mitsumi or Panasonic.

● When it comes to IDE and SCSI CD-ROM drives, they will automatically be installed in the same manner as would an IDE or SCSI hard disk. This will occur as long as the interface card is properly installed.

Cacheing CD-ROM drives

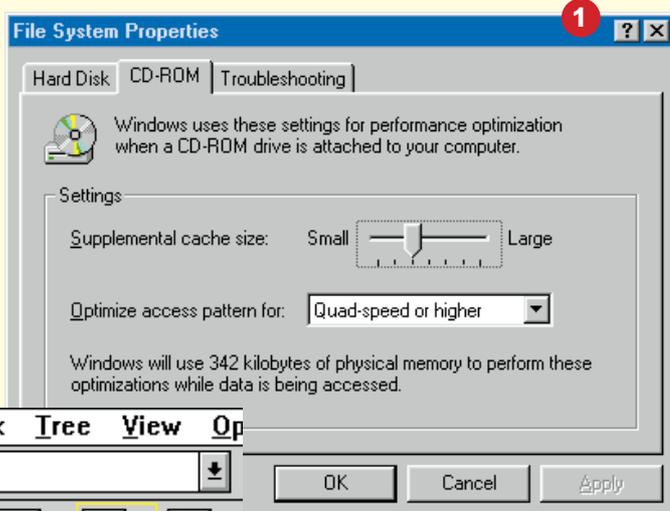
CD-ROM drives are relatively slow devices, but their performance can be enhanced by cacheing.

Ever since MSDOS 6.2, SmartDrive has been able to cache CD-ROM drives. To do this, just make sure that MSCDEX is loaded prior to SmartDrive. Windows 95 uses a 32-bit development of SmartDrive, VCache, which ought to be adjusted to suit

PCW Contacts

Roger Gann can be contacted either by post c/o PCW or via email at rgann@mcgilivray.win-uk.net

Fig 1 Once you have installed your CD-ROM in Windows 95, check to see that its settings are correct



2

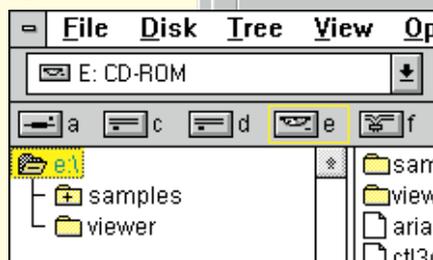


Fig 2 If all has gone well, then the Windows 3.1 file manager should have an extra drive icon

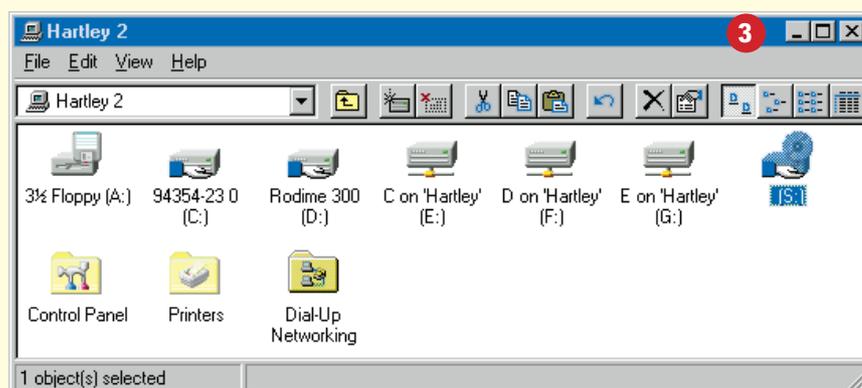


Fig 3 If all has gone okay under Windows 95, the presence of your new CD-ROM drive will be denoted by a new icon in the "My Computer" group