Networking Basics

Keen to link your computers together in a network? It's easier and cheaper than ever to do this, thanks to recent developments in both hardware and software. Here's a quick update on modern networking basics — and a practical guide to setting up your own network.

WHENEVER YOU have more than one computer being used at the same location, networking them together makes a lot of sense. Not only can you transfer files between them quickly and easily, but they can also share expensive resources like laser printers, hard disc arrays, backup tape drives, CD and DVD burners, scanners, internet connections and so on.

Thanks to recent developments in both hardware and software it's now also easier than ever to hook up your own network, and at surprisingly low cost. In this Tech Update we're going to explain the latest approaches to cost-effective networking, and give you enough understanding to let you tackle setting up your own network with confidence and success.

First, though, let's take a quick look at how computer networking developed.

Networking history

In many ways networking is just another development of data communications, which began at the same time as computers themselves. For example in 1940 when George Stibitz and Samuel Williams, researchers at AT&T's Bell Laboratories, developed their Complex Number Generator — one of the earliest digital computers — they made it possible for users to operate the computer remotely from teletype machines on three different floors in their Manhattan building. They also developed a system to allow users to communicate with the machine even more remotely, via telex machines.

Then in 1964, when computers were still very much in the 'mainframe' era, Drs John Kemeny and Tom Kurtz of Dartmouth College in New Hampshire developed not only the simplified BASIC programming language (Beginners All-purpose Symbolic Instruction Code) but also a system called **time sharing**, which allowed many students to use a single computer at what seemed like the same time, from a large number of teleprinter terminals spread around the campus and elsewhere.

Officially, though, networking is said to have started in 1976, when researchers Dr Robert Metcalfe and Dr David Boggs of the Xerox Palo Alto Research Center (PARC — the same place that developed the GUI or 'graphical user interface', the ancestor of Macintosh and Windows operating systems) presented details of a 'Local Computer Networking System' which allowed data to be transferred between a number of computers and printers, at what then seemed a high speed: 2.94Mb/s (megabits per second). They dubbed the new networking system **Ethernet**, and explained that it involved all data being sent in standardised 'packets'. Each packet of data carried its own electronic 'address', which allowed the packets to be directed around the network by switching circuits.

Xerox Corporation was granted a patent for Ethernet ("Multipoint Data Communications System with Collision Detection", US Patent No.4063220) on December 13, 1977. Then in 1979 Xerox teamed up with Digital Equipment Corporation and Intel to define the first official standard specification for Ethernet. This Ethernet Version 1.0 was capable of transferring data at 10Mb/s over fairly thick (10mm diameter) 50Ω coaxial cable, and became known as 'thick Ethernet' or '10Base5'. The second term was because it allowed a 10Mb/s data rate, it used **base**band transmission rather than a modulated high-frequency carrier, and would give reliable data communication over cable lengths up to 500 metres.

Four years later, the 802.3 working group of the IEEE (Institution of Electrical and Electronics Engineers) released its first *industry* standard for Ethernet, giving it the name *IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications*. This was essentially a polished-up version of the DEC-Intel-Xerox 'thick Ethernet'.

Later developments of Ethernet came fairly quickly, and resulted in Ethernet soon becoming the most popular computer networking system.

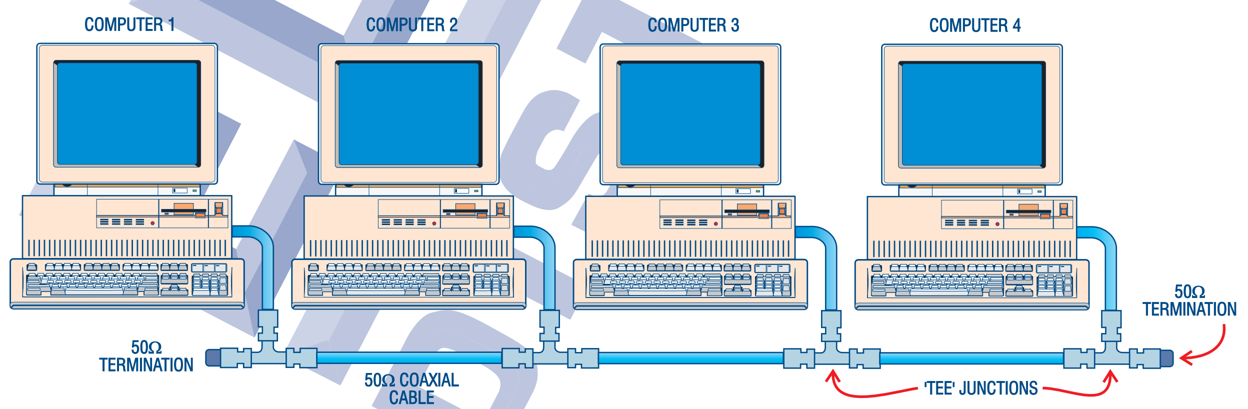

Fig.1: Four computers connected up in a network using the bus or 'daisy chain' topology — in this case using 50-ohm coaxial cable, as used for 10Base5 (thick Ethernet) or 10Base2 (thin Ethernet) networking.

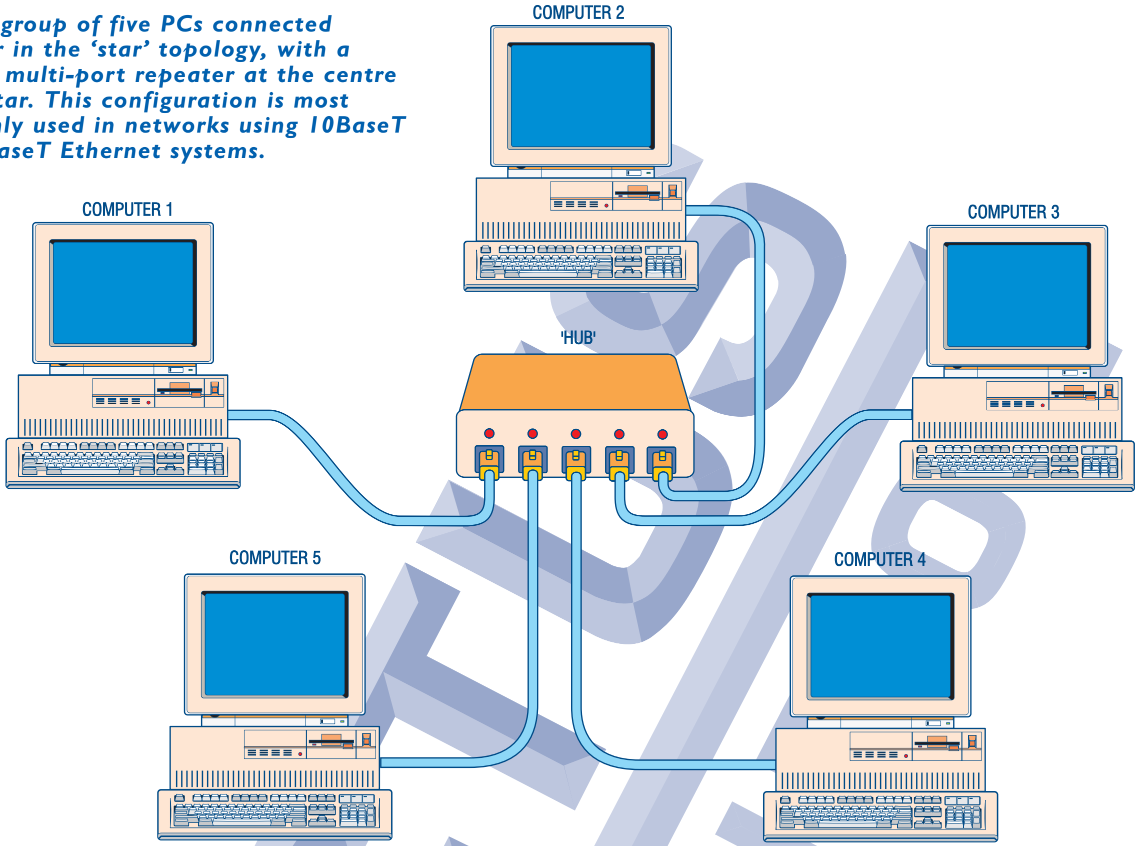

Fig.2: A group of five PCs connected together in the 'star' topology, with a 'hub' or multi-port repeater at the centre of the star. This configuration is most commonly used in networks using 10BaseT or 100BaseT Ethernet systems.

In 1985, the IEEE working group came out with 'thin' Ethernet, also known as 'cheapernet' or '10Base2'. This specified the use of thinner (5mm diameter) 50Ω coaxial cable, which still allowed 10Mb/s baseband transmission as before, but with a maximum cable length of 185 (rounded to '200') metres.

Then in 1990 the IEEE 802.3i or '10BaseT' Ethernet standard was released, which opened the door to much cheaper networking because it allowed 10Mb/s transmission over the low cost 100Ω **unshielded twisted-pair** or 'UTP' Category 3 cabling which had by then become widely used for telephone wiring in buildings. Using this cable also allowed the network to use a 'star' configuration or **topology**, rather than the bus or 'daisy chain' topology needed for thick and thin Ethernet (10Base5 and 10Base2). The two different topologies are illustrated in Figs.1 and 2.

The next big development came in 1995, when the IEEE working group released the 802.3u standard. This became known as '100BaseT' or 'Fast Ethernet', because it allowed 100Mb/s baseband transmission over either two pairs of Category 5 100Ω UTP cabling (100BaseTX), or four pairs of Category 3 cabling (100BaseT4) or two multi-mode fibre-optic cables (100BaseFX). In other words, 10 times the speed of 10BaseT.

Then in 1997 came IEEE 802.3x, which defined **full duplex** or simultaneous two-way data communication over either 10BaseT or 100BaseT. This effectively doubled the data rate again, because before this Ethernet allowed only **half duplex** or 'one way at a time' communication — especially in 10Base5 and 10Base2 coaxial systems.

Then in 1998 and 1999, the IEEE working group released four different implementations of the 802.3z 'Gigabit Ethernet' standard, achieving 1Gb/s transmission or another 10-times increase in data transfer rates. The four versions of this are 1000BaseSX, which uses 850nm lasers and a multi-mode fibre-optic cable; 1000BaseLX, which uses a 1300nm laser and either single or multi-mode fibre-optic cable; 1000BaseCX, which uses 'twinax' shielded twisted-pair (STP) cable; and 1000BaseT, which uses four pairs of Category 5 UTP cabling.

Don't let all of these different 'flavours' of Ethernet worry you, because the two early coaxial-cable based versions (10Base5 and 10Base2) are now regarded as obsolete and are not recommended for new networks. Also Gigabit Ethernet is not only fairly expensive, but in any case its 1Gb/s performance is essentially 'overkill' for most home and small office applications.

That really just leaves 10BaseT and 100BaseT, which both turn out to be very suitable and cost-effective for both home and small office networking. In fact as we'll see soon, the price of everything you need for 100BaseT now makes it virtually the ideal system for new networks, even those with only two or three computers.

Because of this, we're going to ignore the other varieties of Ethernet in the rest of this update, apart from the foregoing mention of their existence.

USB networking

Before we dive further into Ethernet, though, let's look briefly at another approach to networking that can be more attractive if you only have two or three computers to be hooked together and you want do it with the least possible hassle. This is using USB-USB link cables.

As you're no doubt aware, most PCs made since about 1996 are fitted with ports for the Universal Serial Bus or USB — a medium speed serial interface bus designed expressly for connecting computers to peripherals like scanners, printers, keyboards, mice, joysticks, modems, digital cameras, low speed CD-ROM burners and so on. USB is also supported by most of the recent versions of Windows: Windows 98, Windows ME, Windows 2000 and Windows XP (but not Windows NT).

One of the big advantages of USB are that it's 'Plug-n-Play', so cables can be plugged in or unplugged safely even when the power is on. It also offers a fairly high data rate of 12Mb/s, although only over relatively short cables — up to about five metres or so.

USB was designed to connect a single 'host' computer with many (up to 127) slave devices, with a 'tiered star' topology. This means it wasn't really designed for directly linking two or more computers. (For more information on USB, see the Reference Data Sheet USBBUS.PDF, which can be downloaded from our website). However engineers have recently been able to design special adapter cables which can be used to link computers via their USB ports, and software people have written 'driver' software to support networking via these USB-USB link cables, running under any of the operating systems which support USB.

Using one of these cables you can usually achieve data transfer rates of about half the full USB rate, or roughly 5Mb/s.

An example of this type of cable is the XC-4845 USB to Net-Linq cable available from Jaycar stores and Electus dealers. Hooking two PCs together using one of these is very straightforward; you don't need to open up either computer to fit any cards. All that's needed is to install the driver software (supplied on a CD-ROM) on one PC to suit its operating system, and then connect the two computers together via the USB adaptor cable

Here is the Electus XC-4845 'USB to NET-Linq' cable, which can be used to network two or more PCs via their USB ports...

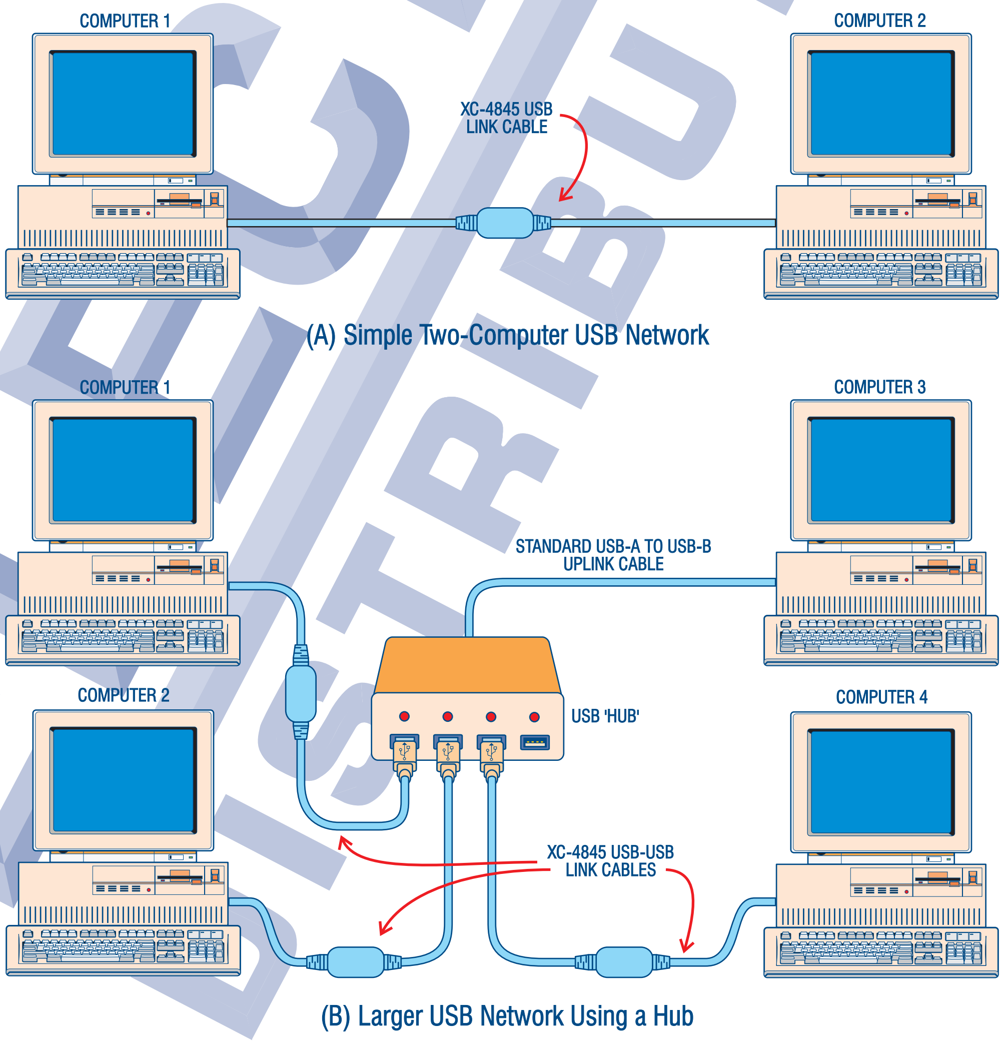

Fig.3: How the XC-4845 USB to NET-Linq cable shown at above right can be used for networking two or more PCs using their USB ports.

— which plugs into a spare USB port on each machine, or on a 'downstream' USB port on a hub. The basic idea is shown in Fig.3(a). The two machines will then form a simple network, with all the usual facilities for file transfer, printer sharing and so on.

It's also possible to connect additional PCs (up to 17 in all) to this simple USB network, simply by using additional XC-4845 cables and a USB hub. Fig.3(b) shows how this is done. Note that the PC connected to the hub's upstream port (Computer 3) doesn't need an XC-4845 USB link cable — just a normal USB cable. But each of the other PCs does need its own XC-4845 cable.

USB networking is really only suitable if the computers are all in the same room, fairly close to each other and running USB-compatible operating systems. It also becomes a bit expensive if you want to connect up more than about three PCs. In fact for networking three or more computers, it's really easier and cheaper to use 10BaseT or 100BaseT Ethernet cables and equipment — especially if the PCs are in different rooms. So let's now turn our attention to this approach to networking.

Fig.4: The basic format used in Ethernet data packets or 'frames'. Note the destination and source address fields, used in directing each packet around a network.

Ethernet basics

As we've seen, the most popular system used for local area networks or 'LANs' is Ethernet, developed from the system developed at PARC in 1977 by Metcalfe and Boggs. And although there are many different implementations of Ethernet, they're all based on the same basic logical bus topology, the same system of sending data in relatively small 'packets' or **frames** which carry the address of both the sending and receiving device, and the same method of allowing a lot of data transceivers to share that common bus. This is the method explained in the original Xerox patent, called **carrier sense multiple access with collision detection** (CSMA/CD).

Without going too deeply into the technicalities, CSMA/CD essentially works by having each device 'node' on the network listen for bus activity (i.e., **carrier sensing**) before it tries to transmit a packet of data. This is like a subscriber on one of the old telephone 'party lines' picking up the receiver to listen if someone is already using the line, before they try to make a call themselves.

If a device doesn't detect any bus activity, it begins to transmit the packet of data. But of course it's possible for another device to begin sending its data at much the same time, in which case there'll be a 'collision' — two packets of data are present on the bus at the same time, and the data gets 'mixed up'.

While the devices are transmitting their data, they're also sensing the bus so they can monitor for any collisions (this is the **collision detection** function). If they sense the extra bit transitions produced by a collision, both devices stop transmitting their data packets and the first device sends out a special 'jam sequence' code — basically a short collision alarm message. Then both of the devices that were trying to transmit wait short but random periods of time before listening for bus activity and trying to transmit again. This is known as **backing off**, and the random back-off delays are to try and ensure that they shouldn't have a collision with their next attempts.

What does an Ethernet data packet or frame actually look like? Fig.4 gives you an idea. The main things to note are the destination and source address segments or fields, which can each be between 2 and 6 bytes long (16 - 48 bits), and the data segment itself which can be between 46 bytes and 1500 bytes long. We'll see how those address fields are used in directing the packets around the network shortly.

By the way, if you were wondering why the Ethernet standards have a specification for the maximum cable length between any two device 'nodes', it's because of the way the CSMA/CD collision detection scheme works. It always takes a certain time for signals to 'propagate' along a cable, and this delay time is directly proportional to the cable length. But for the collision detection scheme to work, each device must be able to receive the collision warning signal from any other device before it's too late to respond. So the maximum total 'round-trip' propagation time must be less than the Ethernet's CSMA/CD **slot time**.

Since cable delay time (in each direction) is the main contributor to total propagation time, this sets a limit to the maximum cable length between any two PC nodes on the network.

This is why the original 10Base5 Ethernet had a maximum node-to-node cable distance of 500 metres, because of its 10mm diameter 'thick' coaxial cable. It's also why 10Base2 Ethernet has a maximum node-to-node distance of 185 metres, because of its thinner 5mm diameter coax (RG-58). And of course why 10BaseT and 100BaseT have a maximum node-to-node distance of 100 metres, because of their use of UTP cable.

10BaseT basics

10BaseT Ethernet uses low cost 'Category 3' (or better) four-pair UTP cabling of nominal 100Ω impedance to connect up the computers and other equipment as a network. All connections are normally made via 8-pin RJ45 *modular* plugs and sockets, which are very much like the 4-pin RJ11 and 6-pin RJ12 modular connectors now used for telephones, only slightly larger.

The basic connections used for 10BaseT UTP cabling

10BaseT and 100BaseT Ethernet both use UTP (unshielded twisted pair) cables, fitted with RJ45 modular connectors. Here are some samples of the 'patch' cables sold by Electus...

and the RJ45 connectors are shown in Fig.5. As you can see only two pairs of wires are actually used; the other two pairs are connected, but not used.

Each computer connects to the network via a **network interface adaptor card** or 'NIC', which is basically a plug-in card containing a controller to manage the data interface between that computer and the network, and a data transmitter/receiver combination or 'transceiver'. In most modern NICs the controller and transceiver are combined in a single LSI chip, which makes the card very low in cost.

(By the way some of the latest model PCs have a network interface built directly into the motherboard, so they don't need the addition of a separate NIC. They have all the hardware necessary for networking already present.)

An important point to remember is that regardless of whether the PC's network interface is in the form of a NIC or built into the motherboard, it has a fixed and unique network or 'MAC' (media access control) address which is hardwired into it during manufacture. This address is a six-byte (i.e., 48-bit) binary number, which is used by the network to identify the node at which that PC is connected. The interface or NIC will only 'accept' data packets which carry this address in their distination address field (Fig.4), and will always include the same address code in the source address field of any data packets it transmits out to the network.

Often diagnostic software will represent the NIC's MAC address as a 12-digit hexadecimal number, such as 054F17B3A8. Not that *you* normally need to worry about your PC's NIC address, of course — the networking software looks after all that.

To allow the NIC's controller to manage the exchange of data between the computer's processor and all other computers on the network, you have to install networking 'driver' software in the computer. NICs usually come with a floppy disc or CD-ROM which provides a range of software drivers, to suit different PC operating systems and/or network operating systems.

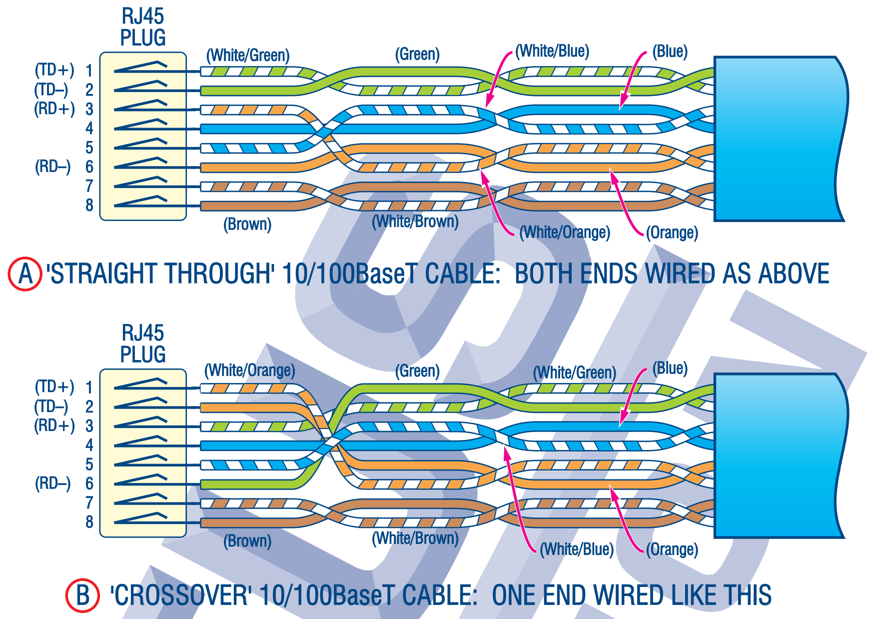

Fig.5: The connections for RJ45 plugs used on Category 3/5 UTP cables as used for 10/100BaseT networking. Normal 'straight through' cables have both ends wired as in A; 'crossover' cables have one end (only) wired as in B, with the transmit and receive pairs swapped.

The simplest possible 10BaseT network is shown in Fig.6. As you can see it's simply a pair of PCs, each fitted with a NIC and with the two NIC ports connected together with a Category 3 or higher rated UTP 'patch lead'. If the two PCs are running virtually any modern operating system (i.e., Windows 3.11 or later) it's simply a matter of installing the appropriate software drivers in each computer and they'll be operating as a fully functional 'peer to peer' network. Data can be transferred between the two in either direction at 10Mb/s, simultaneously if necessary.

More speed: 100BaseT

10 megabits per second is quite fast, of course. Fast enough for most home and small office networks, in fact. But if you *do* need even faster networking, this can be achieved quite easily by using 100BaseT (strictly 100BaseTX) instead. This is very similar to 10BaseT, and still uses four-pair UTP cable; the only difference is that the connecting cable(s) must now be Category 5 rated — so that they can cope with 100BaseT data rates of up to 100Mb/s. The actual cable connections are still exactly the same.

For 100BaseT you also need to use NICs which are rated for 100Mb/s operation, as well. Luckily most currently available low cost NICs are capable of operating at either the 10BaseT or 100BaseT rates, so this isn't likely to be a problem. Most of these so-called '10/100BaseT' cards are designed to plug into a PCI bus card slot in the computer, to take advantage of the higher speed.

These are the only two real differences between 10BaseT and 100BaseTX — the latter needs Category 5 cabling and NICs capable of the higher speed. So the simple two-PC network shown in Fig.6 could operate at either 10Mb/s or 100Mb/s in either direction, depending on the rating of the cables and NICs. The maximum length of cable

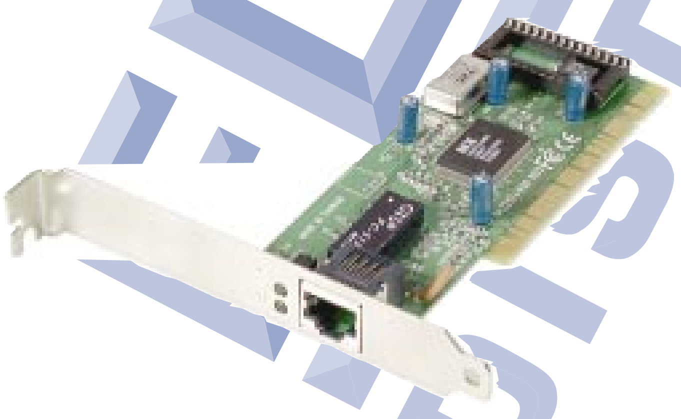

The YN-8062 network interface card or 'NIC' sold by Electus. Low in cost, it offers 10BaseT or 100BaseT operation and plugs into a standard PCI card slot.

between the two PCs would be 100 metres.

Adding a hub

What if you need to expand the simple 10/100BaseT Ethernet network of Fig.6, by adding more computers? That's easy. All you need to do is install a NIC and the appropriate driver software into each additional PC, and connect them up via a multi-port 10BaseT or 100BaseT **hub**. We'll look closer at these shortly, but for the moment just think of a hub as a kind of data distribution box which sends beefed-up copies of any data packets arriving at any one of its ports to all of the other ports.

The general arrangement of this kind of hub-based network is shown in Fig.7, and as you can see the network now has a physical 'star' configuration or topology. Each of the PCs is connected to a port on the hub, so they are all 'peers' in the network.

The length of each cable should be kept less than 100 metres, so the maximum total cable length between any two PCs is less than 200 metres. This applies regardless of the proportion of each cable which is open 'patchlead' and the proportion that is wired permanently within walls or suspended ceilings. The other main requirement is that all cabling should be rated for Category 5, especially for 100BaseT operation.

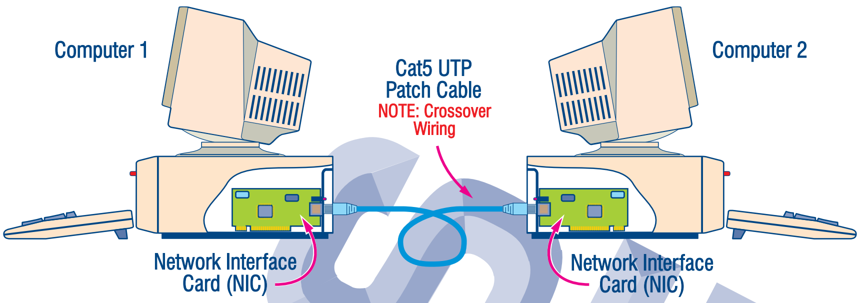

Fig.6: The simplest possible 10/100BaseT Ethernet network, made by fitting a NIC to each of two PCs and linking them with a UTP patch cable (Category 5 rating for 100BaseT). The cable needs to have 'crossover' connections, though — see text. Once the necessary driver software in installed in each PC, they're a network.

The network interface card

As we've already noted, all except the very latest PCs need the addition of a network interface card or NIC to adapt them for Ethernet networking. As you can see from the photo on page 5, a typical 10/100BaseT NIC is a very small and low cost card which plugs into a PCI slot, with just about all of the circuitry built into a couple of ICs. One of the ICs includes the NIC's hard-wired MAC address.

As you'd expect there's a single RJ45 socket on the card's mounting bracket, to accept the network cable. The only other items of interest are a couple of LEDs just near the RJ45 socket, which are used to indicate activity at either the 10BaseT or 100BaseT rates.

The NIC shown is available from Electus as the YN-8062, and has a retail price of around $30. It comes complete with a floppy disk containing networking driver software to suit a wide range of PC operating systems (Windows for Workgroups 3.10 and 3.11, Win95, Win98-2000, WinME, Win NT4 and Macintosh) and networking systems (Lantastic V6.0, NDIS 2.0, Netware 3.1/4.1/5.0, SCO UNIX 5.0 and a TCP/IP packet driver), and also a setup/configuration program.

Incidentally Electus also stocks a 'combo' NIC which features both RJ45 and coaxial cable port connectors, making it compatible with both 10BaseT and the older 10Base2 'thin Ethernet' networks. This card is worth remembering if you need to replace a NIC in an existing 10Base2 network. It has the Cat. No. YN-8060, and is only a couple of dollars more than the YN-8062. But note that it's only capable of 10Mb/s operation, so you can't use it in 100BaseT networks. It too comes with a floppy disk containing all necessary driver and configuration software.

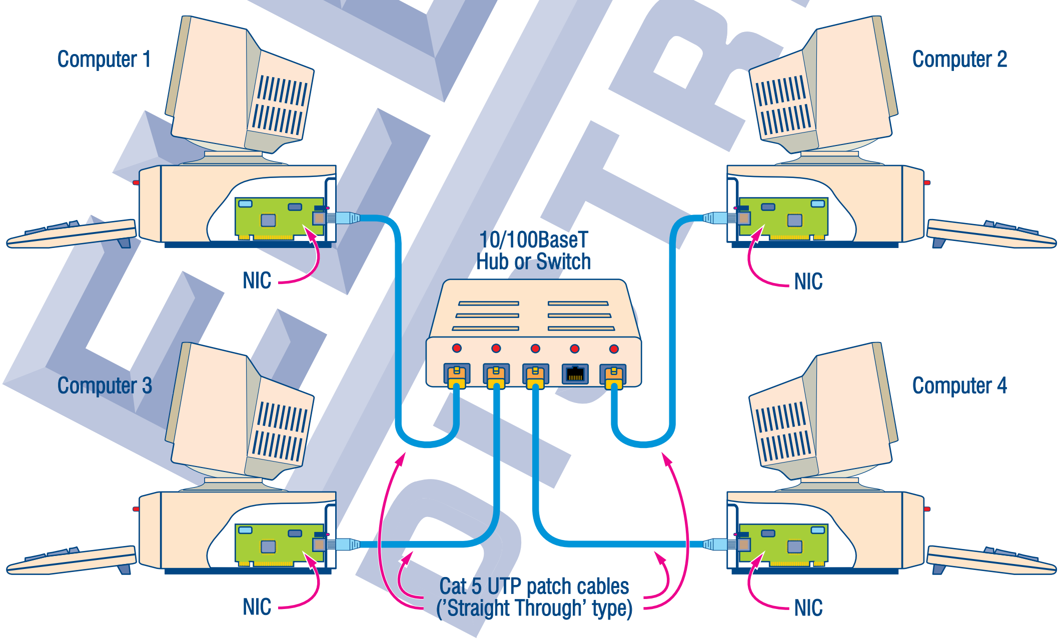

Fig.7: To add more PCs to the network, the simplest approach is to use a hub as shown here. However for faster operation you can use a switch instead, as explained in the text.

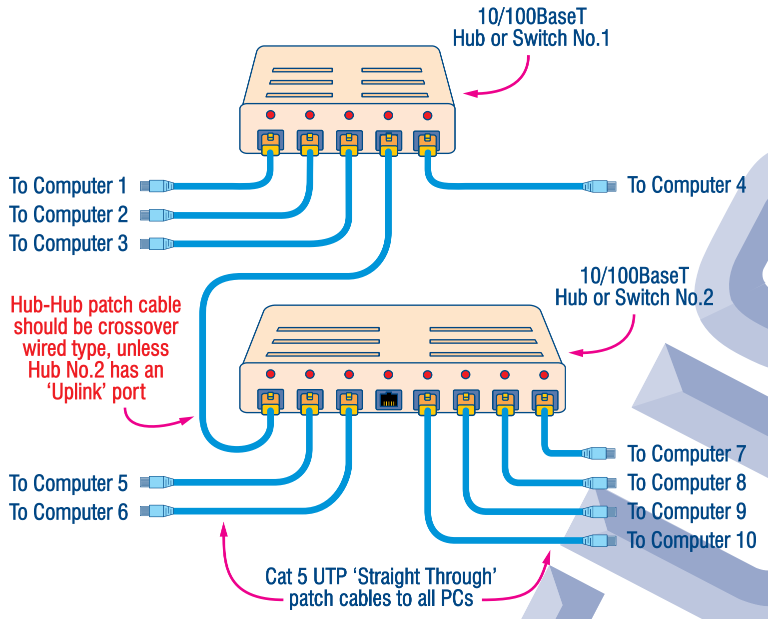

Fig.8: When more PCs need to be added to a 10/100BaseT network based on a single hub (or switch), the easiest approach is to add another hub or switch and 'cascade' them as shown here.

Cascading hubs

As we've seen the simplest way to hook up more than two PCs in a network is using a hub, as shown in Fig.7. Hubs are relatively simple and low in cost, and they're essentially multi-port 'repeaters' or signal booster amplifiers. They have no built-in 'intelligence'; any packet of network data which arrives via the *input* connections of any of their ports is simply beefed up and sent out again via the *output* connections of all ports (including the one they came from).

As low cost hubs are available with up to eight ports, this means that the simple hub-based network scheme shown in Fig.7 can be used to link up as many as eight computers. Electus dealers sell two suitable Nexus 10BaseT hubs for this, the YN-8070 Five-port model and the YN-8072 Eight-port model. Both are very compact and run from a small plug pack.

What about networks with more than eight PC's? When you want to expand past eight, the simplest solution is to add a second hub connected in 'cascade' with the first. This is shown in Fig.8. The second hub connects to one port of the first, in place of one of the original PCs.

There's an important point to note about this cascading of hubs, though. For the cascading to work properly in terms of network operation, there must be a 'crossover' connection between the two hubs. In other words, the input connections at one end of the hub-to-hub cable must connect to the output connections at the other, and vice-versa.

Many hubs are designed so that one of their ports can have these connections swapped over inside with a switch, so that by using the switch you can configure that port as a 'Uplink' port for cascading. With the switch in the 'Normal' position, the port is configured just like all the others, to connect to a NIC in a PC. Both of the hubs sold by Electus have one port which can be configured in this way, for example.

If you *do* have a hub with a port which can be configured for uplinking, you simply use this hub as the cascaded or 'expansion' hub (Hub No.2 in Fig.8), and link the two hubs together using a normal 'straight through' Cat5 patch cable.

But what if neither hub can be configured to have an uplink port? There's still no real problem — all you need to do in this case is link the two together using any of their normal ports, but using a Cat5 patch cable with 'crossover' wiring. In this case the necessary reversing of the transmit and receive pairs is provided by the cable, rather than the wiring inside a hub. Fig.5(b) shows the connections for one end (only) of a crossover cable. The other end is wired as shown in Fig.5(a).

So just remember: when you're cascading hubs to expand your network, either use an uplink port on the lower-level hub and link them with a normal 'straight through' cable, or use normal ports on each and link them with a 'crossover' cable. By the way this same rule also applies with switches, which we're going to discuss in the next section.

One final point about cascading hubs and switches: the patch cables used to interconnect the hubs or switches should be no longer than five metres long. This is again to keep network propagation delays within the 10BaseT or 100BaseT specifications, for reliable operation.

Switches vs hubs

While you can keep on expanding a network by using cascaded hubs, as shown in Fig.8, what you find if you do this is that before long, the network seems to become noticeably slower in operation. This slowdown is due to network 'congestion', which arises because of the way hubs are 'dumb' and simply boost any data signals which arrive at any of their port inputs, and broadcast them out again to *all* ports.

So in a network hooked up via hubs, every PC's NIC receives *every* data packet sent over the network, and it's up to the NIC and its driver software to inspect every packet and 'ignore' all of those which are intended for another PC. It's like a telephone system hooked up as an enormous party line, with no switchboard. And because of the huge amount of data sent to each NIC all the time, the network slows down.

To solve this problem, or at least minimise its impact,

This very compact 10BaseT five-port hub is very suitable for small networks. One of its five ports can be configured as an 'uplink' port, with the connections internally crossed over. Power comes from a small 'plug pack'. The Nexus unit shown is sold by Electus dealers as YN-8070.

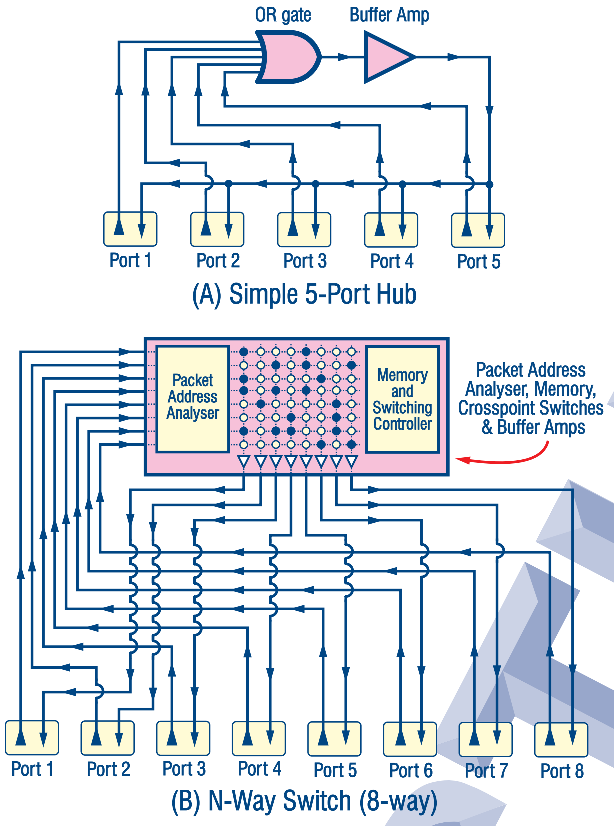

Fig.9: Diagrams showing the difference between a hub and a switching hub or 'N-way switch'. The latter is much smarter, directing the data packets.

the simplest approach is to use switching hubs or **switches** instead of normal hubs. Switches may look very similar to hubs, but they're the next step up in terms of 'intelligence'. Instead of simply boosting every incoming data packet or frame and broadcasting it to all ports, a switch checks the source and destination fields in each frame (see Fig.4), to work out where that frame has come from and where it should be sent. It also has an internal memory where it stores the hardwired network address of each NIC that is currently connected to its ports (or the addresses of the NICs that are connected elsewhere in the network, via its uplink port). And finally, it has a 'crosspoint' switch array, with a set of electronically controlled switches so that it can direct any incoming data frame to the output pair of just one port — the port which connects to the NIC corresponding to its destination address.

So if you like, a switch is much 'smarter' than a hub, and much more like a telephone switchboard. By directing each packet of data only to the NIC it's intended for (or at least the hub or switch to which that NIC *is* connected), a switch cuts down the network congestion significantly. Because each NIC now only receives data that is specifically addressed to it, there's now much less risk of data collisions too. As a result the network can operate at much closer to full speed.

The basic differences between a hub and a switch are shown in Fig.9. As you can see there's a lot more in a switch, to give it the intelligence to remember the NIC addresses, work out the destination of each data packet and send it only to the port leading that NIC's address.

If you want to take advantage of the significant increases in network efficiency and speed offered by switches, Electus offers two Nexus models — both rated for 10/100BaseT operation. The YN-8082 is a five-port switch, while the YN-8084 provides eight ports. Both have one port which can be used either as a normal port to connect to a NIC, or as an uplink port to connect to another switch or hub for cascading with a straight-through cable.

Note that both of these switches are capable of 10/100BaseT operation (i.e., up to 100Mb/s), whereas you'll remember that the two low-cost hubs currently stocked by Jaycar are only capable of 10BaseT (10Mb/s) operation.

Do you have to use either hubs OR switches, or can you mix the two? You can always use a combination of the two if you want. Replacing your top-level hub with a switch will generally give a big improvement, even if you still use hubs at the lower levels. As switches are usually twice the price of a hub with the same number of ports, this kind of mix-and-match combination can be very cost effective.

For example a combination of an eight-way switch at the top level with a couple of five-way hubs at the next level down, as shown in Fig.10, will usually give quite speedy operation for up to 14 PCs. Replacing the two lower-level hubs with switches will certainly give an additional speedup, but often not enough to justify the extra cost.

In short the decision about using switches instead of hubs is really up to you, because it depends largely on the amount of money you have available to invest in networking and the degree of network congestion slowdown you're prepared to tolerate.

Note though, that if *are* using a switch at the top level and your network includes a PC shared as a network resource by all the other PCs — such as a print server, a server for a CD-ROM drive tower or hard disk RAID (redundant array of inexpensive drives) array, or a high speed internet connection — this machine or machines should be connected directly to a port on the top level switch. Don't connect servers to the lower level hubs or switches, or most of the potential efficiency of the switch won't be realised.

Wireless networking

You've probably heard about wireless networking — networking computers using radio waves instead of cables. So just before we finish up we'd better give you a basic understanding of what this is all about.

Practical wireless networking dates from 1997, when the IEEE published its 802.11 standard for wireless

As you can see, a five-port switch looks very similar to a five-port hub. However it has the capability to give much speedier network operation. This Nexus 10/100BaseT unit is sold by Electus dealers as the YN-8082.

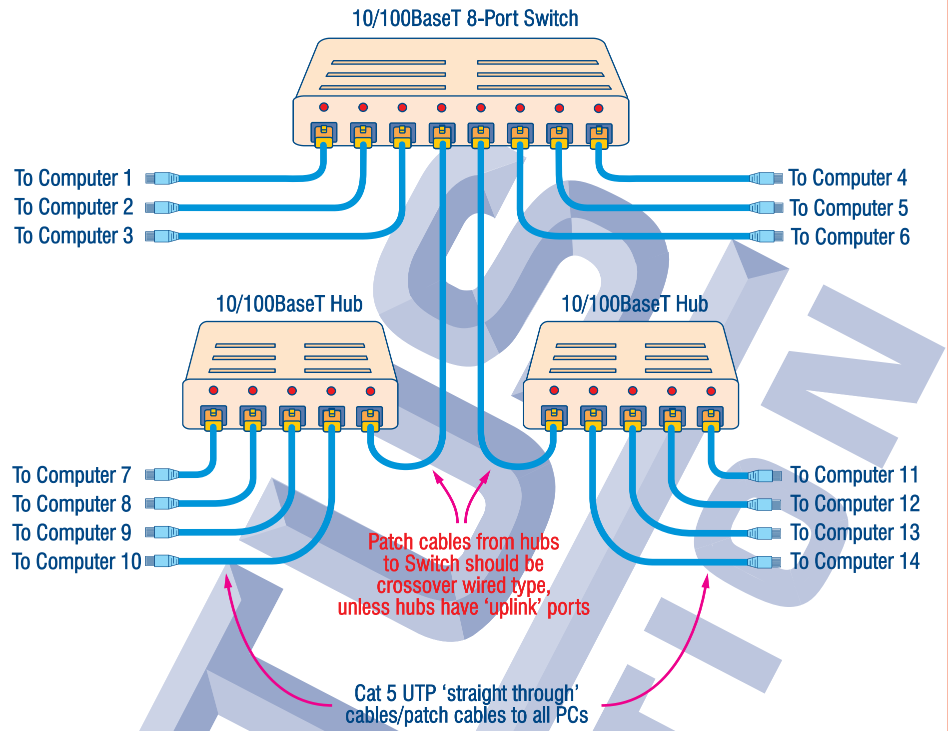

Fig.10: Can you mix switches and hubs? You certainly can. Here's a network of 14 computers, with an eight-port switch at the top level and two five-port hubs connected to it in cascade. Note that servers and other network resources should be connected to the top-level switch, for the most efficient operation.

Ethernet LANs. This originally used data transceivers sharing frequencies in the 902 - 908MHz UHF band, but the equipment was quite expensive and subject to interference. More recently '802.11b' transceivers using frequencies in the 2.40 - 2.43GHz microwave band have appeared, offering better performance at a lower price. Even more recently some of the '802.11a' transceivers use even higher microwave frequencies at around 5GHz, offering higher speed but at a higher price.

Apart from using radio waves instead of copper-based cables, 802.11 wireless networking works in very much the same way as normal wired Ethernet like 10BaseT. The main differences are in operating speeds and price.

Wireless networking equipment designed to the 802.11b specification (also known as 'Wi-Fi') can theoretically give data rates up to 11Mb/s — a whisker faster than 10BaseT. However in practice typical 802.11b network links rarely achieve rates higher than about 5Mb/s. In other words, 802.11b typically delivers around half the speed of 10BaseT.

Electus also sells this 'combo' NIC (YN-8060), which is compatible with both 10BaseT and 10Base2 networks.

The current price of an 802.11b 'access point' (the actual wireless transceiver, combined with a NIC or a USB controller to hook up to a PC) is typically around $700, while the PC-card NIC/transceiver cards needed for the wireless-linked PCs are typically around $200 each. So even to hook up a pair of PCs using 802.11b wireless networking will currently cost around $900 — more than 10 times the cost of a pair of 10/100BaseT NICs and a crossover cable.

When you bear in mind that 802.11b only delivers about 5Mb/s, this cost may be hard to justify. However if you want networking without the hassle of cables, it may still be worth considering.

This year equipment designed to the newer 802.11a specification (also called Wi-Fi5) has also become available. This does offer higher speed than 802.11b, typically achieving 22Mb/s over short distances (about 5m) and 13Mb/s over medium distances (about 10-12m). However the cost is even higher than 802.11b equipment, with an access point typically costing $1200 and PC Card NIC/transceiver adaptors between $400 - $500 each. So at present you'll pay quite a premium for that increase in speed.

No doubt the cost of 802.11b wireless networking equipment will come down in time, making it more attractive. However until then, most people are likely to find 10/100BaseT wired Ethernet the most appealing proposition from a cost point of view.

Incidentally both 802.11b and 802.11a offer lower data security than wired Ethernet, because of the wireless links involved. Both varieties of 802.11 employ data encryption to increase security, called Wired Equivalent Privacy (WEP). However this does appear to be vulnerable to hacking. Of course this isn't likely to be much of a problem for home networks, or even a network for a small business.

Networking Jargon Explained

10Base2: 10Mb/s Ethernet networking using a 5mm diameter 50-ohm coaxial cable as the transmission medium. Also known as 'cheaperNet' or 'thin Ethernet' (IEEE spec 802.3a-1985).

10Base5: The original 10Mb/s Ethernet, using a 10mm diameter 50-ohm coaxial cable as the transmission medium. Also known as 'thick Ethernet' (IEEE spec 802.3-1983).

10BaseT: 10Mb/s Ethernet using two pairs of 100-ohm UTP cable rated to Category 3 or better.

100BaseT: 100Mb/s Ethernet, usually using two pairs of 100-ohm UTP cable rated to Category 5. Also called 'Fast Ethernet'.

1000BaseT: 1Gb/s or 'Gigabit' Ethernet, using four pairs of Category 5 UTP cable, or twinax, or lasers and optical fibres.

Bridge: A piece of Ethernet networking equipment which can be used to link two separate 'domains'. Essentially a switching hub or switch, but with only two ports.

Category 3 cabling: Unshielded twisted-pair cabling which is manufactured to specifications such that it will support 10Mb/s data transmission.

Category 5 cabling: Unshielded twisted-pair cabling which is manufactured to specifications such that it will support 100Mb/s data transmission.

CSMA/CD: Carrier Sense Multiple Access with Collision Detection — the basic principle used in Ethernet networks, to allow many computers to share the data transmission medium.

Ethernet: The data networking system developed at Xerox in 1977 by Metcalfe and Boggs. Now the most popular system for LANs.

Fast Ethernet: Another name for 100BaseT Ethernet, running at a data rate of 100Mb/s (ten times the rate of 10BaseT).

File server: A computer connected to a LAN which provides file storage resources to the other 'client' machines. It may have a RAID array of hard disks, a rack or 'tower' of CD-ROM drives, or both.

Gigabit Ethernet: Ethernet networking technology which offers a data transmission rate of 1Gb/s (1000Mb/s).

Hub: A 'dumb' signal booster and interfacing device for Ethernet networking, which simply takes any incoming data arriving at any of its ports, and rebroadcasts it to all of its port outputs.

IPX/SPX: Logical protocols used in Netware.

LAN: A *Local Area Network*, linking a group of computers together in roughly the same physical location.

MAC: Media Access Control, the method used to regulate the way each network device can access the transmission medium. Uses the unique hardware address given to each NIC when it's manufactured.

NIC: The *network interface adaptor card*, which is used to provide each computer with its network data interface and transceiver (and also its unique hardware address on the network).

Print server: A computer, connected to a network, which is dedicated to managing a printer resource.

RAID array: A *Redundant Array of Inexpensive Drives*, with multiple hard disk drives connected in such a way that they offer a large amount of highly reliable data storage. Often managed by a 'file server' computer, to provide bulk data storage for a network.

RJ45 connectors: Compact modular connectors with 8 contacts, developed from the smaller RJ11 and RJ12 telephone connectors. (The 'RJ' simply stands for *registered jack*.)

Repeater: A 'dumb' signal booster and interfacing device for Ethernet networking. Very much like a hub, but with only two ports.

Router: An even more 'intelligent' version of a switching hub or 'switch', used to interconnect multiple networks — or a network and the internet.

STP: Shielded twisted-pair cable, as opposed to unshielded twisted pair cable.

Switch: An intelligent *switching hub* which analyses the source and destination address codes in the header of each data packet or frame arriving at one of its ports, and directs the packet only to the output port which connects to the NIC with that destination address. This reduces network congestion and hence increases speed.

TCP/IP: *Transmission Control Protocol/Internet Protocol*, the main logical data exchange protocol used on the internet and for LAN-internet communication.

Thick Ethernet: The original 10Base5 Ethernet of 1983, using 10mm diameter coaxial cable for 10Mb/s networking.

Thin Ethernet: 10Base2 Ethernet, using 5mm diameter coaxial cable for 10Mb/s networking.

USB network bridge cable: An adaptor cable which allows networking between two PCs, or between multiple PCs, using their USB ports.

UTP: *Unshielded twisted pair* networking cable, usually with four pairs of insulated wires and with each pair of wires twisted together to minimise interference. Available with a Category 3 rating to support 10Mb/s transmission, or a Category 5 rating to support 100Mb/s transmission.

WEP: Wireless Equivalent Privacy, the data encryption system used in 802.11a/b wireless networking. It is intended to offer a level of security equivalent to that of wired Ethernet.

Wi-Fi: Another name for 802.11b wireless networking technology, using frequencies in the 2.4GHz UHF band. Offers data rates up to about 5Mb/s.

Wi-Fi5: Another name for the newer 802.11a wireless networking technology, which uses frequencies at around 5GHz and offers data rates up to about 22Mb/s.

Wireless networking: Linking computers via radio waves rather than copper wires or optical fibres. (IEEE specs 802.11a/b)