PAUL MOBBS

A practical guide to sustainable IT

Unit 3



APC

This unit is one of 12 sections to a "A practical guide to sustainable IT", a hands-on guide to working with everyday technology in an environmentally conscious way. The guide has been written by environmental activist and ICT expert Paul Mobbs, and was commissioned by the Association for Progressive Comunications (APC) with the support of the International Development Research Centre (IDRC). To download the full text of the guide, or any of the other units, please visit: greeningit.apc.org

Author
Paul Mobbs

Copy-editing
Alan Finlay

Layout proofreading
Lori Nordstrom

Publication production
Karen Banks and Flavia Fascendini

Graphic design
Monocromo
info@monocromo.com.uy
Phone: +598 2 400 1685

Commissioned by the Association for Progressive Communications (APC).



Conducted with support from the International Development Research Centre (IDRC).



The views expressed in this publication are those of the individual authors and not necessarily those of APC or IDRC.

A practical guide to sustainable IT

Published by the Association for Progressive Communications (APC) with support from the International Development Research Centre (IDRC).

South Africa
2012

Creative Commons Attribution 3.0 Licence
<creativecommons.org/licenses/by-nc-nd/3.0/>
Some rights reserved.

APC-201206-SU-R-EN-DIGITAL-162
ISBN: 978-92-95096-71-4

HARDWARE: PUTTING THE SYSTEM TOGETHER

Hardware is the physical components which make up a computer system. In this section we'll look at hardware, computers, and how we can design and arrange the components of computer systems to work more efficiently. Today computers come in all shapes and sizes, and the electrical technologies on which they are based can have a serious impact on the environment. How we plan our need for information and communications, and the equipment we use, can in turn have an effect on the level of these impacts. That doesn't just involve what type of machine you use – one of the important factors is to maximise the working life of the device before you dispose of it.

3.1. WHAT IS "COMPUTER HARDWARE"?

Fifty years ago a "computer" was a collection of electrical hardware which filled a small building. Today that term encompasses a very broad variety of equipment, some of which are the size of a grain of sand. The commonest programmable computer devices do not resemble the archetypal vision of a machine with a screen and keyboard. Most computers are small and insignificant chips, embedded[1] in radios, which run the electronics of TVs, cars and washing machines.

If we look at the number of high performance processors which are intended for crunching data, many of those are in computer servers, telecommunications centres and supercomputers – invisibly serving data to other machines rather than "real" people. Even if we think of devices specifically intended to connect people directly to information, there are between five and six-billion mobile phone accounts in use around the world,[2] compared to around two-billion personal computers.[3] And even when we talk of personal computers (PCs), many people accessing the internet and online services are doing so via video game consoles and set-top TV boxes, not just stand-alone PC-style computers.

That's why putting together a guide to sustainable IT is so difficult – the market is so diverse that it's difficult to give a concise guide which will cover all the possible means by which people might use computers and electronic networks.

Even when we talk of PC hardware,[4] there are many different types of machine in use which loosely fulfil the definition. The managers of data centres, made from thousands of microprocessors and consuming perhaps thousands of kilo-watt hours of electricity, have a strong influence over the ecological footprint of IT. The fact that most green IT reports address large computer centres is a measure of their importance. Even so, if we weigh up which group of computer users collectively has the greatest influence on the ecological footprint of IT, then it's the everyday users of personal computers and communications devices. That's partly because their equipment does not function as efficiently as the highly managed computers of data centres; and primarily because it's the mass of computer users at work and at home, connecting to online services, who are driving the demand both for new hardware and the greater network capacity required to support online services.

In this section we'll look first at the issue of open and closed hardware, and using proprietary devices such as mobile phones or games consoles in contrast to more conventional personal computers. Then we'll move on to look specifically at PCs, and the factors which influence the way they work. At each stage what we'll try to do is give an idea of the impact of these devices, and how you can get involved in changing/managing that impact.

1. Wikipedia, Embedded computer. en.wikipedia.org/wiki/Embedded_computer
2. BBC News (July 2010). Over 5 billion mobile phone connections worldwide. www.bbc.co.uk/news/10569081
3. Wikipedia, Personal computer. en.wikipedia.org/wiki/Personal_computer
4. Wikipedia, Personal computer hardware. en.wikipedia.org/wiki/Personal_computer_hardware

3.2. "OPEN" AND "CLOSED" HARDWARE

The simplest way to consider what "open" and "closed" hardware is, is to look at one recent trend which defines the difference between the two – the "jail-breaking" of Apple and Android phones and the modding[5] of certain types of games console. Closed proprietary equipment, such as mobile phones and games consoles, are not intended to be serviced or modified by their owners. In many countries it's an offence under intellectual property law to try to modify the design of these devices in order to adapt or extend the way they operate. This has come to a head recently as the result of computer enthusiasts circumventing or removing the manufacturer's limitations[6] on the functions of machines – removing blocks to certain types of activity, or to allow the running of new or modified software. In response the companies involved have brought court cases against those responsible, which has in turn made modding and jail-breaking campaigning issues amongst many digital rights activists.[7]

For most PCs and similar hardware, the restrictions on modification are not so strict. Much of the PC hardware people use, both laptop and desktop machines, are to some extent designed to be modified. On laptop machines you can in most cases change the memory modules and hard disk drive, and replace the internal battery when it fails. On desktop machines almost all the components can be modified or swapped, allowing much greater freedom over the configuration of the system.

The reason for this relates to the way the first IBM-compatible PC was created. Rather than being designed as a wholly proprietary product, the rush to get a PC into market meant that it was quickly cobbled together from a range of "off the shelf" components available from hardware manufacturers.

While each individual module or circuit board is covered by the same intellectual property restrictions as closed hardware, assembling these different elements together to make a computer system is not. As we'll examine later, this allows greater scope to create a more efficient computer system tailored to the user's needs.

3.2.1. Closed hardware as an obstruction to repair and modification

The major hurdle in managing the ecological impacts of closed hardware are the obstacles to changing the components of the system – or simply being able to repair them affordably.[8] Most mobile phones have software-based power management controls which permit the user to select various modes to reduce the energy consumption of the device, but this represents the limit to what is easily/legally permitted. Other consumer electronic devices, such as set-top boxes and TVs, may also have such options built-in.

The general problem is that as the power consumption of individual devices has dropped, a greater proportion of the ecological impacts occur during manufacturing. If we look at the life-cycle impacts, over a wide range of digital consumer electronics much of the energy and carbon impact, and certainly much of the waste, is created during the production of the device. Therefore the majority of the impact that the device will have on the environment *already exists on the day it is purchased* – our ongoing role as the user of the device, powering it from the electricity grid, is less significant.

For example, let's look at Apple's products as shown in figure 3.1. Usefully, Apple produces life-cycle environmental reports[9] for their products, compiled under international standards for environmental reporting (Apple should be praised for this, and it is something that all manufacturers should be required to compile for their products). Roughly a quarter of the life-cycle carbon emissions of Apple products are the result of their use by the consumer; around

5. Wikipedia, Modding. en.wikipedia.org/wiki/Modding
6. Wikipedia, Hardware restrictions. en.wikipedia.org/wiki/Hardware_restrictions
7. For example, Electronic Frontier Foundation, Jailbreaking Is Not A Crime. www.eff.org/pages/jailbreaking-not-crime-tell-copyright-office-free-your-devices
8. The Story of Electronics, Annie Leonard, Free Range Studios, 2010. www.storyofelectronics.org
   See also The Story of Stuff: Annotated Script, Annie Leonard, Story of Stuff Project, 2010. www.storyofstuff.org/wp-content/uploads/2011/02/SoE_footnoted_script.pdf
9. Apple Computer, Apple and the Environment. www.apple.com/environment/reports/

a half or more are produced during the production of the device (the remainder is expended in transport and during materials reclamation).

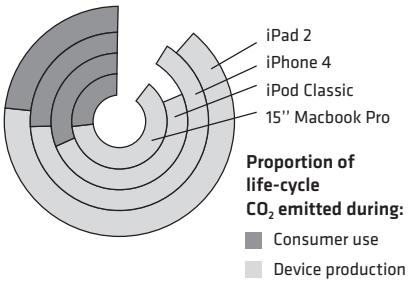
One of the important variables in the calculation of these impacts is the length of product's service life. For the Apple devices shown above they are expected to have a three-year life-span, except for the Macbook which is expected to last four. If you can make these devices work for longer than this then you will decrease the impact of the product overall. Even though using the device for longer will mean that the user's proportion of the total impact will rise, because you use more power, that fact the most of the impacts are created during production mean that you are unlikely to tip the balance the other way. And in any case, by delaying buying a new product you are delaying the creation of new pollution in order to replace it with a new one. For example using an iPod or iPhone for four years instead of three will reduce the overall impact by almost a third.

The problem with trying to use many consumer electronics goods for longer than their manufacturer's anticipated life-span is that you run-up against the issue of *planned obsolescence*.[10] Analogue/mechanical devices will work right up until individual components fail or wear out – although it might be that the device has had its components engineered to last a certain period of time and then fail. In contrast, digital devices can have their service life programmed into them. They are designed to work for a certain period of time and then, irrespective of whether they are serviceable or not, they will cease to function. The only way to circumvent problems such as planned obsolescence, in order to extend the service life of a product, is to find ways of repairing the hardware, or reprogramming the firmware[11] (the counters or other limiting controls set in the on-board programming of the device) which is preventing the device operating.

For example, returning to the Apple iPod again, in 2003 it was the subject of a short film[12] which highlighted how the batteries in the early models failed after 18 months. As a result of a subsequent court cases brought against Apple, which showed that the design of the battery limited the life of the device, Apple offered replacement batteries for the early iPods.[13] More recently claims have emerged about the latest design of Apple products, in particular the use of tamper-proof screws to hold the case together.[14] It is argued that this makes it harder for people to carry out their own repairs/servicing in order to extend the life of the device. The iPod case, and other examples where equipment has been engineered to have a short life and/or be difficult to repair, were highlighted recently in the documentary, The Light Bulb Conspiracy.[15]

Figure 3.1.

Carbon emissions from Apple products



3.2.2. What is open hardware?

Today a new open-hardware[16] movement is beginning to take off,[17] pushing the boundaries of how people can share skills, information and designs in order to adapt or modify different types of equipment. Open hardware, at its

10. Wikipedia, Planned obsolescence. en.wikipedia.org/wiki/Planned_obsolescence
11. Wikipedia, Firmware. en.wikipedia.org/wiki/Firmware
12. The Neistat Brothers (2003). iPod's Dirty Secret. www.ipodsdirtysecret.com
Film available via YouTube, www.youtube.com/watch?v=F7ZsGIndF7E
13. BBC News (3rd June 2005). Apple deal for iPod battery cases. news.bbc.co.uk/1/hi/technology/4606481.stm
14. Kahney, Leander (January 2011). Is Apple Guilty of Planned Obsolescence?
www.cultofmac.com/77814/is-apple-guilty-of-planned-obsolescence/
15. Dannoritzer, Cosima (2010). The Light Bulb Conspiracy.
www.facebook.com/TheLightBulbConspiracy
View online at documentaryheaven.com/the-lightbulb-conspiracy
16. Wikipedia, Open hardware. en.wikipedia.org/wiki/Open_hardware
17. For example see www.openhardware.org or opensourceecology.org Watch a short video from Open Source Ecology on the principles of open hardware at vimeo.com/30171620

most open, allows the copying and complete modification of the hardware design. While there has been little movement on this in the world of computing, in the fields of intermediate technology and engineering there are a variety of projects now running around the globe which allow the owner or builder of a device or machine complete freedom to vary or improve the design – usually on the condition that all such changes are fed back into the community sharing the design.

Other groups such as iFixit, through their Self-Repair Manifesto[18] and online repair guides, are promoting the rights of owners to maintain and repair their hardware. Reconfiguring computer hardware can be fairly simple once you've got the hang on it – as we'll see later.

In general, the aims of the open-hardware movement are to extend the scope of people's involvement with their equipment. Its key objective is to develop the designs and skills that people can use to motivate their involvement, and to extend these "open", principles to a far greater part of the engineered systems which support our lives.

3.3. UNDERSTANDING COMPUTER HARDWARE

For those who find it difficult to understand how to use computer programs, the idea of taking the case off their machine and fiddling with the workings inside can be very daunting. There's a great mystique surrounding the workings of computer hardware – although this isn't deserved because compared to most other consumer electronics, such as radios or washing machines, the way computer hardware works is fairly simple. Even if you don't bother to modify your system, knowing how the insides of the machine function can be very helpful to finding faults with your system, and it can be immensely useful to understand all the acronyms and specifications you are bombarded with when you buy equipment.

3.3.1. Types of PC hardware

When we talk about computers we're conventionally shown a large box with a screen and keyboard (or a laptop which combines these elements). In fact computers come in all sorts of sizes, and many don't have a screen or keyboard.

- At the smallest end are plug systems[19] – very small/low power consumption computers (such as the GuruPlug or SheevaPlug) which plug into a mains plug socket and run your local network. You access the machine over a network connection, and the machine itself usually runs a Linux operating system.
- Next are mobile phones and palmtop machines, and tablet/pad computers – and we should add games consoles and similar devices into this category too. As outlined earlier, these are made from proprietary hardware components/modules and are very difficult to modify. Even the open mobile handsets which use the Android operating system are largely intended to allow users to change the software on the machine, not the hardware itself.
- Next there are laptop and notepad computers. Unlike wholly proprietary systems, which are not intended to be modified by their users, these allow limited scope for maintenance and upgrading. Usually you have the option to change the memory, hard drive and internal battery, although finding the right memory module or drive can be difficult.
- Next there are desktop machines – the large metal boxes that we traditionally associate with the term "computer". With the exception of Apple machines and some other proprietary designs, almost all of the components in a "PC-compatible" desktop machine can be modified or upgraded. This not only allows you to increase the power of the machine easily, but the hardware itself is far easier to reuse/recycle at the end of its life.
- Finally, there are server machines. These are like the plug-in machines noted above, but they're many times more powerful, and for that

18. Ifixit, Self-Repair Manifesto. www.ifixit.com/Manifesto

19. Wikipedia, Plug Computer. en.wikipedia.org/wiki/Plug_computer

Box 3.1.

Hardware buying check-list

Cheapest is not necessarily best

Often you get what you pay for. Devices which use non-reusable or disposable goods, such as the ink or toner in printers, are often sold at cheap prices because the proprietary components they consumed are sold at a relatively higher price. What you need to consider is the "total cost of ownership" – how much, over the likely life of the equipment, you are going to spend on purchasing, servicing, consumables and power consumption. Often paying more initially can reduce the costs later.

The warranty period is the best guide to the ex pected service life

The longer the warranty period, the greater the likelihood that the machine will have a long service life. In a choice between hardware with a similar specification, a longer warranty can represent a more reliably engineered system which is less likely to go wrong.

Does the manufacturer provide environmental reporting information?

The issue with environmental reporting – such as measurements of power consumption, embodied energy, waste production and carbon emissions – isn't simply that it gives you extra information to decide between different brands or devices. Companies which issue environmental reports are more likely to police the activities of their suppliers to ensure that they conform with the product standard – and that ensures that tight environmental standards are applied in everyday industry practice.

Check for non-branded equivalents for the equipment

Branded goods have the greatest exposure in the media, but often more generic/non-branded equivalents are cheaper. The spares for non-branded goods may also be cheaper and more easily available as they are more likely to be used in a range of consumer goods. Warranty periods/terms are also important in making a choice between branded/non-branded goods.

Is the user able to service and repair the device?

Enquire if the device is a sealed unit, or whether the user can easily replace components which might fail over its operating life.

Check if the device uses proprietary or standard/ generic cables/connectors

If a device uses a generic power supply, and the connectors (e.g. sound/audio or ethernet) use standard fittings, these will be easier and cheaper to replace. For power supplies especially, being able to use a generic power supply means that you can reuse the unit with other equipment when the machine reaches the end of its working life.

Check that repair is an option, rather than whole sale replacement

Check that spares for the most likely parts of the equipment to fail – such as rechargeable batteries, hard drives, cables and connectors – are available at a reasonable cost, even if it requires a service technician to undertake repairs.

Does the equipment contain hazardous substances?

In compliance with recent legislation in Europe and elsewhere, goods should indicate whether they contain hazardous compounds such as mercury, cadmium or lead. Some manufacturers are also producing goods which are free of PVC and flame retardants.

Is there are take-back/trade-in option?

If you rent/lease goods you should expect them to be taken back at the end of the lease. For goods which are bought, check if the manufacturer runs a take-back scheme for the machine, or the machine's consumables such as ink/toner cartridges.

Does the warranty offer on-site service?

If the device fails you may have to send it back to the manufacturer's service department for repair which is likely to cost time and money. Check if the manufacturer can offer a contract with a service company to carry on-site repairs as this may take less time to organise.

reason power-hungry. They are intended to run computer networks, or large internet data centres, and so are not the sort of machine you'd use for general purpose computing. Some servers are stand-alone boxes, rather like a desktop machine. For the last decade or so the trend is to shift away from single machines to rack or blade hardware, where multiple server machines operate in a single large case.

In the rest of this unit (and the unit on Reuse and reclamation later) we'll primarily deal with desktop computers. The principles involved in desktop machines also apply to most laptop machines, although the hardware is more difficult for the inexperienced user to modify. We'll investigate the installation software on the machine in the next section, *Operating systems*. What we'll look at below is how a standard desktop machine works, how to identify the different parts of the machine, and what all that jargon about hardware specifications actually means.

There is one golden rule to remember when using computers: *Computers are stupid!* At the simplest level, computer "hardware" is a collection of electrical circuits which need to be told what to do by a computer program. The central processing unit[20] (CPU) of the machine is designed to carry out a set of instructions contained in a computer program. Each instruction is initiated by giving the CPU one or more numbers which represent the instruction (called the opcode) and any data it needs to carry out that instruction (call the operand). The CPU then executes the instruction and stores the result (if any) inside its own internal memory (called registers) or inside the main memory system. Without the instructions supplied by the "software"program the hardware is just an expensive collection of plastic, metal and chemically altered silicon connected to a power supply.

What makes the computer a useful tool is the *operating system*[21] that controls the hardware, and that provides support for the application programs which the user runs. There are various forms of operating system, and some are more complex than others, but understanding the role and functions of the hardware is an essential part of understanding the role and functions of the operating system. By understanding the principles of how the functions of the computer are organised you will be able to work more easily and efficiently, and quickly identify faults and problems when things start to go wrong.

3.3.2. The motherboard

Irrespective of what type of computer system you use, all its components are arranged on a circuit board – usually a single large board called the *motherboard*.[22] On a desktop computer you can see the motherboard quite easily if you remove the left-side cover (for tower cases) or the top cover (for desktop cases). On laptops the motherboard is buried deep inside and you can't get at it without dismantling the whole machine (so best not to try!). Even mobile phones and MP3 players are usually designed with a single large board which contains all the main components of the system. That's because it's the simplest way to manufacture and assemble the device.

The easiest way to learn how to manipulate the hardware of a computer system is to get an old desktop computer, take it to pieces and then put it back together again. When assembling a new computer, or recycling an old one, you don't work inside the case of the computer – it's too cramped a space to work without risking damage to the components. Instead you assemble the separate parts of the computer system on a bench/desktop and then, when you know everything works as it should, you dismantle the system and then install each part into the case (we'll return to the process for doing this in the unit on *Reuse and reclamation*).

Figure 3.2 shows a "logical" layout of a computer motherboard. It doesn't really look like this, and the precise design of motherboards vary, but if you look at most motherboards (there are various illustrated guides available on the 'net) you'll probably be able to pick out the components listed in this diagram. Figure 3.3 shows what a motherboard might "physically" looks like, and the important parts are labelled to help you (these are referenced in the text below in {curly brackets}). Let's go through the diagram and explore what each piece does within the system:

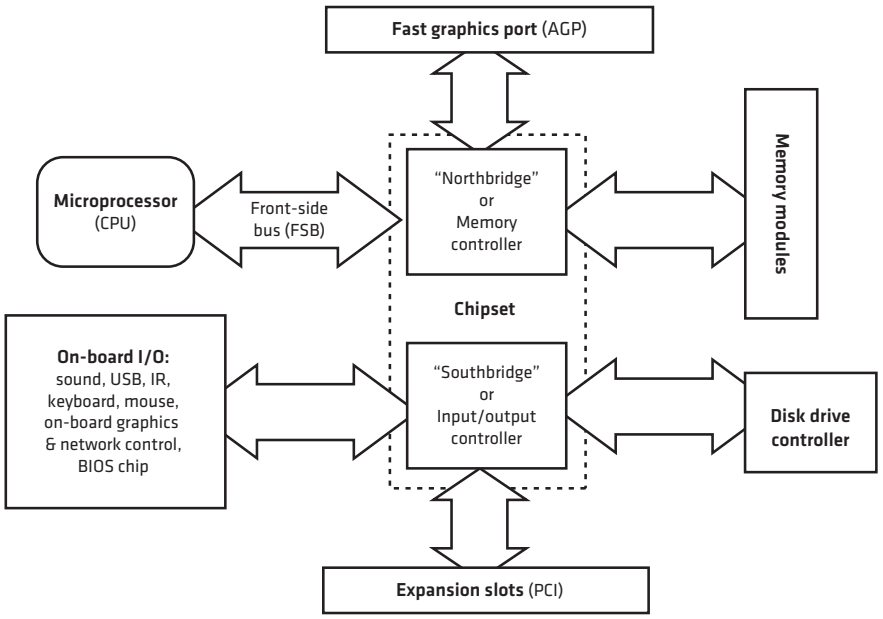
20. Wikipedia, Central processing unit. en.wikipedia.org/wiki/Central_processing_unit

21. Wikipedia, Operating system. en.wikipedia.org/wiki/Operating_system

22. Wikipedia, Motherboard. en.wikipedia.org/wiki/Motherboard

Figure 3.2.

Logical organisation of a motherboard



3.3.3. The processor and chipset

Firstly let's look at the microprocessor[23] or CPU, the chipset and the front-side bus. The CPU is the heart of the machine. In the 1950s the central processing "unit" would have been a cabinet the size of a delivery van, containing the racks of circuits which made up the CPU. Today it's a single chip, about the size of a postage stamp – and it has thousands of times more processing power than the van-sized CPUs of the 1950s.

There are many types of CPU available, and often what you pay for isn't simply a higher speed, it's the amount of extra memory or parallel processing systems that the chip contains. For much of their early lives Macintosh computers used a different type of processor to IBM-PC style machines. That changed a few years ago when Apple adopted the same processor family as the PC, the x86.[24] The first generation of x86 processors was produced in 1978, and formed the heart of the first IBM-compatible PC, which all later designs emulated. We're now up to somewhere around the ninth generation of x86 processor, and the range of processors has broadened to include designs for desktops, low-power versions for laptops, basic chips for cheap mass-market machines or more advanced versions for those who need lots of computing power. More specialised computers, such as the server machines that run the internet, often use more specialised chips that are designed to perform a narrow set of functions at very high speed; unlike the standard x86-family of PC chips which are made to perform a wider variety of functions, from processing graphics data or doing complex decimal arithmetic.

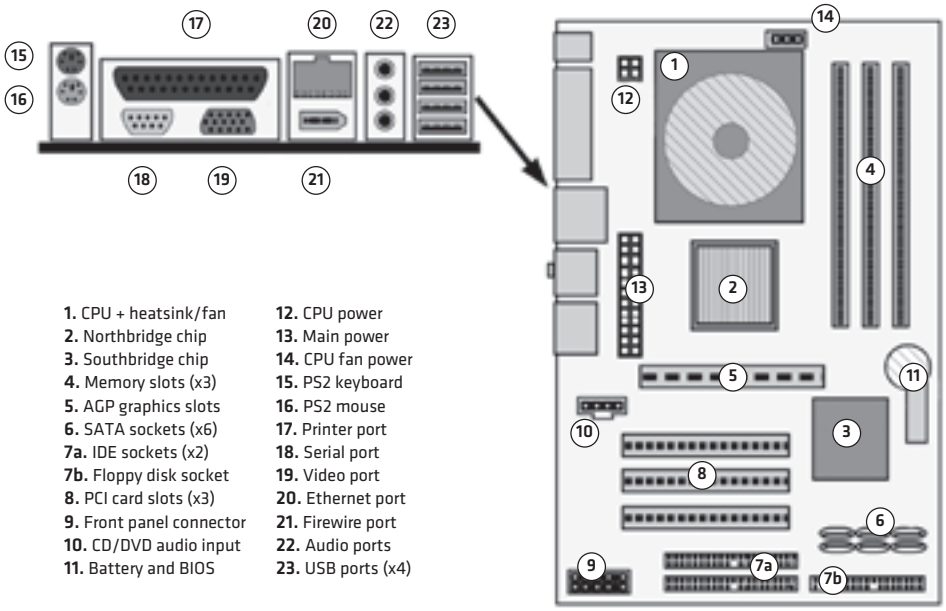
If you look around the processor connector on the motherboard {1} you'll see a network of small wires on the board leading away towards other chips on the circuit board. That's the front-side bus (FSB). Near to the processor there will be another large chip on the board (usually about an inch/2cm to 3cm square) – often it has its own small metal heatsink on top, and sometimes a fan, to keep it cool. That's the northbridge {2}, a chip that controls communication between

23.Wikipedia, Microprocessor. en.wikipedia.org/wiki/Microprocessor

24.Wikipedia, x86. en.wikipedia.org/wiki/X86

Figure 3.3.

Physical organisation of a motherboard



1. CPU + heatsink/fan
2. Northbridge chip
3. Southbridge chip
4. Memory slots (x3)
5. AGP graphics slots
6. SATA sockets (x6)
7a. IDE sockets (x2)
7b. Floppy disk socket
8. PCI card slots (x3)
9. Front panel connector
10. CD/DVD audio input
11. Battery and BIOS
12. CPU power
13. Main power
14. CPU fan power
15. PS2 keyboard
16. PS2 mouse
17. Printer port
18. Serial port
19. Video port
20. Ethernet port
21. Firewire port
22. Audio ports
23. USB ports (x4)

the memory, the CPU and the rest of the machine. Some of the latest machines don't have a northbridge because, due the problems of trying to increase the speed of the system, many of its functions are now carried out by the CPU itself. A little further away from the northbridge you should see another fairly large chip – that should be the southbridge {3} (sometimes, especially on fairly new motherboards, this has a heatsink on top as well). This controls the disk drives and the other peripheral parts of the system.

3.3.4. Assessing the speed of the system

Whenever you see adverts for computers they always stress the speed of the processor as the most important selling point. That's not wrong, but it's not necessarily true (see box 3.2). While the speed of the chip is important, the design of the other components in the system also has a big contribution to the overall speed. The main difference when you buy a very expensive computer isn't just the speed of the CPU or the size of the memory – it's the quality of the other chips on the motherboard. This increases the data transfer rate between the CPU, the memory, the disk drives and the accelerated graphics display. The more expensive the motherboard, the faster the chipset of the motherboard will function, and so you'll squeeze more data through the system every second. Even so, and irrespective the system specification of your computer, you need to match what your hardware can provide to the demands of the operating system (see box 3.3).

When you look at a motherboard the CPU appears to be the largest component on the board. That's because it has to have a large fan fixed to the top of the chip to remove the heat it generates. In older designs of CPU, if the fan failed the chip would burn itself out. Modern designs of chip are temperature controlled, and will slow down as they heat up – so if your computer appears to be working very slowly it's possible that the CPU fan has broken, or it needs the dust cleared from the cooling fins to improve the rate of heat loss. As the amount of heat that the CPU and motherboard generate increases with the CPU's power, some manufacturers are even developing miniature refrigeration systems to fit inside the computer's case. For the large server farms and data centres which run the internet,

Box 3.2.

Processors and motherboards

A PC motherboard is designed to take a specific type or range of x86 CPU chips. There are two main manufacturers of x86 CPUs: Intel and AMD. Each produces a range of chips with different specifications but there are two general differences between all CPU chips:

- **32-bit versus 64-bit** The bits refer to the width of the data bus. What this means is that a 64-bit chip will move twice as much data during each operation as a 32-bit chip. For 64-bit chips you have to use a 64-bit operating system; a 64-bit processor might work with a 32-bit operating system but you'll be wasting its processing power as none of its 64-bit features will be available. Even if you use a 64-bit operating system, some programs might not be able to use the extra power available from the chip, and so for some tasks you're not going to be processing data much faster than on a 32-bit CPU.
- **Number of cores** A recent innovation. As the complexity of CPUs has grown we're reaching the physical limits of how big they can be. The solution has been to include more than one CPU within the same chip, allowing programs to be run in parallel. The greater the number of "cores" in the CPU, the more processing power it has. Multicore CPUs are available in 32-bit and 64-bit versions.

What makes the greatest difference to the speed of a system, apart from the CPU, is the specification of the motherboard's chipset – in particular the speed of the front-side bus (FSB). The higher the FSB speed, the faster the data can move between the CPU, the memory and peripherals. Another feature to look out for is the amount of cache memory that the CPU has, usually called the Level 1 (L1) or Level 2 (L2) cache. This speeds up the reading of data from the memory, and the bigger the cache memory the faster the CPU can operate.

Different generations of chip use different sockets or slots to connect the chip to the motherboard. In practice, choosing the correct CPU chip to fit to a motherboard isn't an issue you're likely to deal with. Usually your computer will arrive ready built; or, if you upgrade a computer, you can buy a motherboard with the CPU and memory already fitted more cheaply than buying each separately. The only time you need to worry about finding the right processor for a motherboard is when the old one has failed – in which case you must get the manual for the motherboard to find which CPUs it is able to accommodate.

Every motherboard has a technical manual written for it; if you can identify the design of the motherboard (the maker's name is usually fairly prominent, but finding the code which identifies it can be more difficult) then you can search online and find a copy of the manual. The manual will tell you which CPUs/memory modules the board can work with.

getting rid of the heat generated by the motherboard is a major problem, and they can use as much electricity pumping air to cool the machines as the motherboards consume directly. For hot climates, maintaining the CPU fan, or perhaps buying a larger, more efficient fan to cool the CPU, is essential to keep the machine functioning normally. You must also organise a good circulation of cool and relatively dry air (humid air can corrode the workings) to prevent overheating. Overheating will, if it takes place for long periods, shorten the life of the components on the motherboard.

What few people seem to realise is that all that extra heat, and powering the fans or other technology to remove the heat from inside the case, represents wasted energy. Many desktop computers now burn hundreds of watts when operating at full power (laptops are naturally less power hungry because of the need to maximise the battery life). For those living off the power grid, using small power generating systems like photovoltaic panels or wind to charge batteries, the power consumption of IT equipment is becoming a pressing issue. If you want to conserve power then you need to use a laptop and other battery-operated/mobile digital equipment as these are optimised to use a little power as possible.

3.3.5. System memory

Next, let's look at the memory {4} – also called the primary storage[25] or random-access memory[26] (RAM). Over the last decade or so memory chips have become extremely complex devices. Depending upon the speed of the motherboard, the memory modules must match the transfer speed of the front-side bus or you'll get errors (or more likely it'll refuse to work and the motherboard will signal a hardware error). Confusingly, there are different classes of memory (DRAM, SDRAM, etc.) which operate at different data rates (DDR, DDR1, DDR2 etc.) and which come in different sized packages to fit different sized slots on the motherboard. If you look at the Wikipedia page for *DIMM*[27] (dual in-line memory module) you can get a list of all the different types. Look under the "speeds" subsection for a list of the different classifications of DIMM module.

Table 3.2.

Bytes and magnitude

Magnitude	Symbol	Size, bytes
1 byte	B	1
1 kilobyte	kB	1,024
1 megabyte	MB	1,048,576
1 gigabyte	GB	1,073,741,824
1 terabyte	TB	1,099,511,627,776
1 petabyte	PB	1,125,899,906,842,624

Usually a motherboard has at least two memory slots. Sometimes the board requires that they are used in a certain order, and so problems arise when you use more than one memory module. If the memory modules do not have exactly the same speed/data rate specification the board may have problems reading data because it can't synchronise the reading and writing cycles. It's all very complicated, and on older hardware, where you might have problems finding large memory modules, it can be quite annoying. The general rule is to either use one large memory module, big enough to provide the memory space you need; or to find matched pairs of modules and always keep them together.

For all operating systems, the amount of memory that the system has is one of the most critical limitations. The table in box 3.3 lists a number of different operating systems and the absolute minimum system requirements to run them. Although the system designers say that the system "will run" with these specs, that doesn't mean that it will do so comfortably. These are the bare minimum, and so in practice it's best to double or triple these figures (e.g. if it states a Pentium 2 processor, use a Pentium 3 or 4; if it says 256MB of memory use at least 512MB or more). In practice you're stuck with the processor that you have – for the cost of buying a new processor on its own you can often buy a whole new motherboard with a processor already attached. However, no matter what speed your processor runs at, adding more memory to the system will usually help

25.Wikipedia, Computer data storage. en.wikipedia.org/wiki/Primary_storage

26.Wikipedia, Random-access memory. en.wikipedia.org/wiki/Random-access_memory

27.Wikipedia, DIMM. en.wikipedia.org/wiki/DIMM

Box 3.3.

System requirements

Each operating system is usually accompanied by a wad of documentation called release notes; a list of the features of the new system, installation instructions, and any significant changes from the previous version. It should list the system requirements – a specification of the type of processor(s), the minimum speed of the processor, and the amount of memory and hard disk space required to run the operating system.

When installing a computer, or recycling old parts to make a new one, you need to have an idea of what the system is capable of doing. To a great extent that is determined by the memory, but also by the processor.

Table 3.1 lists the system requirements for various operating systems. **Consider all these figures the minimum for installation. Wherever possible use the fastest processor, the largest amount of memory and the largest hard disk drive to make installation and use as simple as possible.**

Note also, earlier versions of the Linux distributions listed in the table will require less RAM and HDD space to function.

Table 3.1.

Systems and system requirements

Operating system	CPU Type, MHz	RAM MByte	Hdd MByte
Microsoft Windows			
Windows ME	P2+, >233	32	320
Windows XP	P3+, >600	64	1536
Windows Vista	P4+, >1000	512	10240
Windows 7, 32-bit	P6+, >1000	1024	16384
Windows 7, 64-bit		2048	20480
Gnu/Linux: Graphical installation (e.g. desktop)			
Fedora 15	P4+, 1000	768	6144
Ubuntu 11.04	P4+, >1000	384	5120
OpenSuSE 11	P3+, >600	256	3072
Debian 6	P4+, >1000	512	5120
Slackware 13.37	P1+, >166	≥128	5120
Gnu/Linux: Lightweight and live distributions			
Knoppix 6.4.4	486+, >90	≥64	(live)
Puppy Linux 5.2	P1+, ≥166	128	(live)
Slackware 13.37	486+, >90	≥64	≥512
Damn Small Linux	486+, ≥90	≥32	≥50
Gnu/Linux: Text-only install (e.g., servers)			
Fedora 15	P3+, >600	>256	≥1024
Ubuntu 11.04	P3+, >600	>128	≥1024
Debian 6	P1+, >100	>256	≥1024

386/486+ - 386/486 processor, or better
P1-P6+ - Pentium-I through to Pentium-VI processor, or better, including 32-bit or 64-bit multi-core CPUs
> - 'greater than' – the minimum for an installation
≥ - 'greater than/equal to' – the bare minimum required for operation (for ease of installation double/triple this figure)
(live) - Live distribution – runs from CD/DVD so hard disk not required

to speed-up the operation of the system. That's because the more data the processor can hold in the memory, the less it needs to access data from the hard disk at a far slower rate.

Before we move on, system memory requirements are often quoted in mega- or giga-bytes, and some new hard disks have a capacity in excess of a terabyte. What do these numbers mean?

Computers use binary numbers, or number base two. Humans use decimal, or number base ten. Counting in magnitudes of 1,000, the human-readable sequence would be *1, 1,000, 1,000,000*, etc. The nearest whole value to 1,000 in binary is two to the power ten (2),10 or 1,024. Table 3.1 shows the magnitude of the different prefixes applied to the size of data storage. As the difference between each magnitude is a multiple of 1,024 this means a megabyte contains 1,024 kilobytes, a gigabyte contains 1,024 megabytes or 1,048,576 kilobytes, etc.

Also, what is a byte? Today PCs use 32 or 64 "bit" processors. A bit is one binary digit – a single binary 0 or 1. When a CPU reads data from the memory it reads 32 or 64 bits at a time. The greater the width of the data bus, the more data the computer can read in one action, and the faster the CPU processes data. The first PCs had an eight-bit data bus, and eight bits were known as a byte. As one byte/eight-bits is what you need to represent a single character of information using the ASCII[28] encoding system, we came to associate one byte with one character of text. However, today's computers need to represent far more characters, including those from languages other than English, than are available using ASCII. As a result ASCII has been superseded by Unicode[29], which requires at least two bytes (16-bits) or more to represent a single character.

3.3.6. The video display

Next, let's look at the accelerated graphics port[30] (AGP) {5}. A modern video display is made up of points of data called *pixels* – the more pixels, and the more colours they use, the finer the image your display can reproduce. The first computer terminals didn't do pixels at all – they just displayed characters. In terms of the amount of data involved, 24 lines of 80 characters, at one byte per character, is 1,920 bytes. Now think of a modern wide-screen video display. It usually has 1,366 pixels across the screen and 768 from top to bottom. Each pixel is capable of representing in excess of a million colours, which means each pixel takes 24-bits, or three bytes of data to represent its value. 1,366 pixels x 768 pixels x 24-bits (three bytes) is 3,147,264 bytes; and given that the screen refreshes its image at least 50 times a second that represents a data flow in excess of 150 megabytes per second. In physical terms, that's the content of a data CD disc every 4¾ seconds! In fact, the AGP slot is designed to handle over 500 megabytes per second, equivalent to a blank DVD disc every 8¾ seconds – that's far more data than the motherboard's PCI slots can handle (which usually have a capacity of 128 megabytes per second).

The CPU would find it difficult to shift 150 to 500 megabytes a second to the video display and do all the other tasks it needs to perform. If you want to play videos or an animated video game, without the display jumping and hiccuping all the time, then the video display system has to be handled by purpose-designed graphics card working separately from the CPU. By making the AGP port an expansion slot you can fit the graphics card of your choice into the motherboard. Then, instead of the CPU transferring data from the memory, the video display is managed by the northbridge of the chipset without the CPU being involved. All the CPU needs to do is update the display image held in the main memory. On older machines, without AGP, the size and quality of the display is limited by the speed of the system. With AGP, and an accelerated graphics card, the system can send far more data to the display and produce a much higher quality picture.

3.3.7. The disk drive controller

Next let's look at the southbridge and disk drive controller {3}. The front-side bus works at very high speeds to shift data from the memory to the CPU and AGP port. That would be a waste of time when communicating with the hard disk, and especially the keyboard, because the data can't be supplied by the hardware at anything near that rate. If you can get hold of an old motherboard, and look at the little copper wires on its surface, you'll see that there's one set of wires linking the CPU, the northbridge and the memory; a second set links the northbridge and

28.Wikipedia, ASCII. en.wikipedia.org/wiki/ASCII

29.Wikipedia, Unicode. en.wikipedia.org/wiki/Unicode

30.Wikipedia, Accelerated graphics port. en.wikipedia.org/wiki/Accelerated_graphics_port

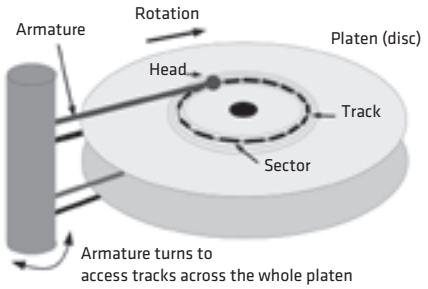
the southbridge; and a third links the southbridge to the disk drive connectors and PCI card slots. The southbridge runs its own "slow" data bus which controls the input/output devices on the motherboard. It collects together a large quantity of data in its memory buffer and then, when there's a gap available on the front-side bus, it moves the data at high speed into the main memory. This allows the CPU and memory to continue operating at high speed without continually being interrupted by data travelling to and from the peripheral devices.

Hard disk drives[31] have been the main method of secondary storage[32] on personal computers for over 20 years. Before we had hard drives, PCs relied entirely on cassette tapes, or on floppy disks.[33] First, the truly "floppy" 5¼-inch disks made of a plastic film sandwiched in a card sleeve, then the more sturdy 3½-inch disks held inside a rigid plastic shell. The very first PCs used a floppy disk to "boot up" the computer with an operating system. In fact most operating systems allow you to create a removable "boot disk [34] to restart the system in the event of a hard drive error, although today you might use a CD or a USB stick since floppy drives are now obsolete. When complete, the system disk could be removed and other disks, used to store data, were inserted as required. Hard disks have the advantage that they are inside the machine, and so are not prone to loss or damaged like floppy disks; they also operate at a far higher speed and have a much larger capacity than floppy disks. Although you might still find floppy disk drives on old computers, using the drive is difficult as most operating systems no longer include the programs to make the hardware work. Even so, many motherboards still include a floppy drive connector {7b} (it looks like a shorter version of the hard drive IDE connector {7a}).

A hard drive is a spinning aluminium disc or "platen" coated in a magnetic material. The heads travel across the surface of the disk to write data by magnetising spots on the surface; or read the stored data by sensing the magnetised areas passing beneath the head. The disk is enclosed in a hermetically sealed metal housing to prevent dust

Figure 3.4.

Hard disk drive schematic



particles getting inside. The heads hover above the spinning disk at less then the width of a human hair – in terms of the scale and speed involved, it's been likened to a 747 airliner travelling at full speed about six feet off the ground. Particles of dust (or, worse still, a human hair!) would cause damage to the disk and the heads, which is why it's sealed away. The metal housing also protects the sensitive magnetic materials inside from stray magnetic fields that might affect its operation.

There are two standards now in common use for PC drives: Integrated Drive Electronics (IDE) or Parallel Advanced Technology Attachment[35] (E-IDE or parallel ATA); and Serial Advanced Technology Attachment[36] (serial ATA or SATA). Any motherboard more than four or five-years old will use IDE/E-IDE {7a}; a motherboard less then two-years old will use SATA {6}, although some retain IDE sockets for compatibility; those motherboards manufactured between these times will use either, and a large number use both. There is a third standard drive interface – Small Computer System Interface[37] (SCSI) – but that's not commonly found on PC machines. It was mostly used on server machines because it allowed a large number of hard drives to be connected to the system. Early Apple Macintosh machines also used it.

Compact flash[38] memory cards are used instead of electro-mechanical hard drives on notebook and tablet machines, and the latest disk-

31. Wikipedia, Hard disk drive. en.wikipedia.org/wiki/Hard_disk_drive
32. Wikipedia, Computer data storage. en.wikipedia.org/wiki/Computer_data_storage
33. Wikipedia, Floppy disk. en.wikipedia.org/wiki/Floppy_disk
34. Wikipedia, Boot disk. en.wikipedia.org/wiki/Boot_disk
35. Wikipedia, Parallel ATA. en.wikipedia.org/wiki/Parallel_ATA
36. Wikipedia, Serial ATA. en.wikipedia.org/wiki/Serial_ATA
37. Wikipedia, SCSI. en.wikipedia.org/wiki/SCSI
38. Wikipedia, CompactFlash. en.wikipedia.org/wiki/CompactFlash

Box 3.4.

Hard disk drives

Most PC hard drives, until recently, used the IDE standard. The motherboard has two sockets, called IDE0 and IDE1. A 40-wire ribbon cable plugs into a socket and each cable has either one or two connectors – so a maximum of four IDE drives can be attached to the motherboard.

Although all IDE drives work the same way when connected to the ribbon cable, for two drives to share the same cable one has to be classed as the "master" drive and the other as the "slave". This is done by changing the position of a small jumper connector on the back of the drive. If the drives are not correctly selected, then there will be a clash when the motherboard tries to talk to them and they won't work. Not all drives have the same jumper order to select the drive, but the correct order is usually displayed on the top of the drive. There is another option, called "cable select", but in many cases that doesn't work.

Over the last few years a new standard, serial ATA or SATA, has been introduced. This is an improvement over the previous IDE system, not only because it gets rid of some of the barriers to the size of drive, but it also gets rid of the ribbon cable. Having two connectors on a ribbon cable was always difficult to arrange when putting a system together – and it also interfered with the air flow needed to cool the drives. SATA uses thin cables which are far easier to connect-up than ribbon cables, and also does away with the need to select the drive using a jumper. Instead each drive plugs into a socket on the motherboard, so the maximum number of drives is limited by the number of SATA sockets available on the board.

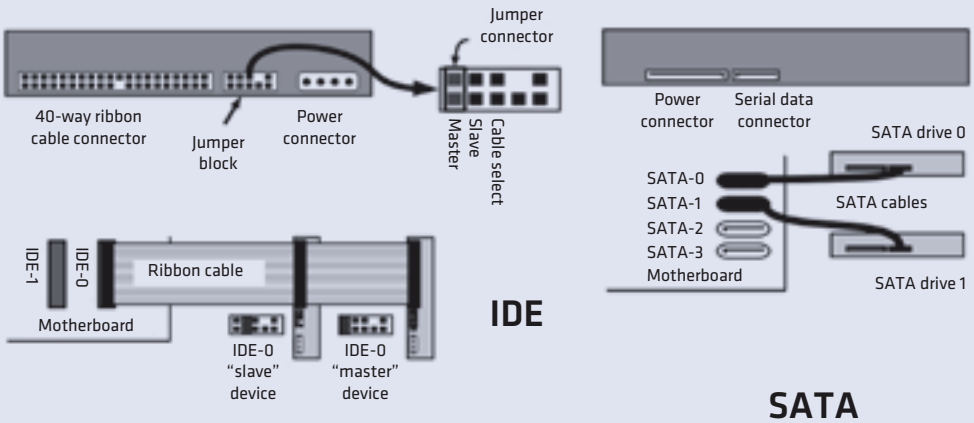
Note that the miniature 2½-inch hard drives used on laptops operate in the same way as the standard 3½-inch hard drives used in desktop PCs, and come in both IDE and SATA versions. If you want to upgrade a laptop hard drive, the manual for the machine should tell you what type of drive is used, and what capacity of drive the machine is capable of accommodating.

One very important point – adding extra hard disks can add a large power drain to the computer system. For that reason you need to have the minimum practical number of hard drives to create the amount of data storage required. For example, a single one terabyte (TB) hard drive will use less power than two 512 gigabyte (GB) hard drives. If reducing power consumption is important, then use just one large hard drive rather than two or three smaller drives. The exception would be where you are using another hard drive as a back-up for the primary hard drive – but as we'll examine in the section on *Information storage*, you can use a removable drive caddy to avoid the need to have that drive connected to the system at all times.

Finally, hard drives also produce a lot of heat, and in hot climates they can become very hot – which can significantly shorten the life of the electronic components. If a drive is running too hot to touch when in use you should consider adding an extra fan to the system to draw air across the hard drives. That might use more electricity, but that's not much when compared to the embodied energy of the hard drive if the excess heat causes it to fail.

Figure 3.5.

IDE and SATA drive configuration



Note: Not all drives use this order!!

less low-power notebook computers. These use IDE/parallel ATA electronics. However, there is an issue as to how long-lived this technology is, especially if the card is used heavily for very long periods. Flash drives, rather like the USB sticks commonly used with computers, work very fast and don't use much power.

The largest amount of drive space you can organise on the machine is limited by the number of hard disks you can attached to the system, and how big those drives can be (see box 3.4). A problem with IDE was that the motherboard could only handle a certain size of hard drive – beyond a certain limit the motherboard couldn't address the disk. That can be a problem when using some older motherboards because only small hard drives are accepted by the motherboard. This limits the maximum storage available, and is a problem because small hard drives (a few gigabytes or less) are hard to find these days. SATA doesn't have these same restrictions as it is designed to handle very large storage capacities, and as a recent standard there's no problem getting hold of large SATA compatible hard drives.

3.3.8. PC cards, front panel connectors and BIOS

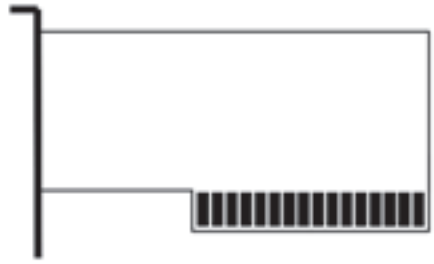
The AGP {5} and PCI {8} "slots" are edge connectors – they're designed to have a circuit board slotted into them to extend the functions of the motherboard. For most everyday functions you won't need to bother about the slots. It's only when you need something out of the ordinary, such as an internal modem, or something non-standard like a data-logging or specialist interface card, that you'll need to use these slots.

An important consideration about PC cards is that they consume power. Some, such as the high-performance AGP card used by the latest video gaming machines, use a very large amount of power (a good indication that a card consumes a lot of power is when it has its own cooling fan). If you can get by without installing a card in the machine, this will minimise power use. If you must install a card, such as a parallel interface card for an old-style printer cable, then check with the manufacturer's specifications to find a the card which has a low power consumption.

Next, let's look at the front panel connectors {9}. A PC has a power and reset switch, a hard disc drive (HDD) and power light, and, recently, a sleep light. These all connect at a small group of terminal posts, usually in one corner of the motherboard. There's also an internal speaker that makes a beeping noise to warn of errors – although some motherboards have the sounder fixed on the board and so don't need an external speaker.

Figure 3.6.

A PC card



The small button cell battery on the motherboard {11} is there to keep power to the machines *Basic Input/Output System*[39] (BIOS) chip when the machine is switched off. This stores all the configuration settings for the motherboard. If removed, after a few hours you'll lose these settings – which is a good way of clearing the access passwords on the motherboard. Also, if this battery runs down, the BIOS settings might be lost or corrupted, in which case you'll need to replace the button cell. Most motherboards, especially laptops where the battery is buried deep inside the machine, have a pair of terminals which can instantly reset the contents of the BIOS memory if connected together. To find them (very useful to know for your laptop if you forget the user/set-up password) you'll need to download the manual for the motherboard/laptop and look up the BIOS settings information.

3.3.9. Power supplies

Next let's look at power supplies. These days PCs only have one large power connector {13}. In ancient times, before the Pentium-II and the introduction of the ATX motherboard,[40] the AT power supply used two connectors side-by-side

39. Wikipedia, BIOS. en.wikipedia.org/wiki/BIOS

40. Wikipedia, ATX. en.wikipedia.org/wiki/Atx

– but you're unlikely to run into one of those these days. The advantage of the ATX supply is that it gives the computer the ability to turn itself off when it shuts down. However, the ecological cost is that when the ATX supply shuts down the system it's still consuming a small amount of power to keep the electronics in standby mode, ready to start the machine when the user presses the on button. Therefore, when you've shut a system down, it should also be isolated from the mains power supply. Not just because of the power consumption (on older desktop machines usually the power drain is about five to 15 watts – about the equivalent of a small low energy light bulb), but also because power surges might damage the electronics. Note, if you use a modem connected to a phone line, disconnect that too, or get a surge protected phone adapter, to prevent damaging power surges coming down the phone line. The router unit on a broadband phone connection should, unless you get a direct lightening strike, prevent power surges reaching the computer.

The ATX power socket on the motherboard has a lug on one side, and the plug on the end of the cable has a small hook that engages with the lug to stop the connector working loose. Since the development of the Pentium-IV, the CPU now needs its own power supply, and this is provided by the four-pin connector {12} on more recent ATX power supplies (it also has the lug/hook arrangement to stop it falling out). Also, before you turn the power on, the large fan on top of the CPU needs a power supply. Depending on the type of fan, you'll find either a two- or three-pin connector {14} somewhere near the CPU where the fan can plug into the motherboard. Sometimes the northbridge chip also has a fan, in which case that should plug-in somewhere nearby too (although some are directly wired into the motherboard – meaning they're a pain to remove if they fail and have to be replaced).

As computers have become more powerful, so their power consumption has risen. Early PC power supplies were rated at less than 200 watts. Most PCs today have power supplies rated at 600 to 800 watts, and some of the latest high performance PC (used for video games and number crunching) can have power supplies of over 1,000 watts – with water cooling of the chips to increase the rate of heat removal. On top of that, the monitor might consume anywhere from 50 watts (LCD) to 250 watts (plasma), and printers, speakers and other peripherals can consume another 50 to 200 watts. Put to gether, a high-powered PC can consume about the same as a small electric heater when running at full power, which is why they produce so much heat.

However, how much power the computer appears to use might not actually be consumed by the computer. That's because the quality of power supplies varies, and lower quality power supplies can have a very low power factor[41] (a value which expresses the amount of power drawn versus that actually used by the equipment). Measuring power consumption, and especially measuring the power factor, can give you an idea of how much power is consumed by all or part of the computer system. Whether you choose to upgrade or replace parts of the system then depends on the likely cost savings and the cost of upgrades (see box 3.5). Recently a new 80 Plus standard was introduced for power supplies, guaranteeing a very high power factor and thus much lower power consumption. When replacing power supplies, look for an 80 Plus certified unit, or browse the web for accredited suppliers.[42]

41. Wikipedia, Power factor. en.wikipedia.org/wiki/Power_factor

42. Plug Load Solutions, 80 Plus Certified Power Supplies and Manufacturers. www.plugloadsolutions.com/80PlusPowerSupplies.aspx

Box 3.5.

Power consumption and power factor measurement

For mains powered systems, the simplest way to measure power consumption is to use a plug-in power meter. These are available for both 230V and 110V supplies, and can give a variety of data about the power consumption of the computer system. If you plug the monitor into the socket which powers the system you can measure the consumption of all the parts, or you can isolate just one component (the monitor, printer or computer box) to see how much power it consumes.

What is important to understand is that the power a PC system consumes varies all the time. When idling it may consume only 15% of the maximum load; if you're encoding video or some other power-hungry exercise, and using the speakers at full volume, it will use much more. Most power monitors will measure the supply voltage and the instantaneous power consumption. What's more useful is to use the power meter function, which measures the supplied power in kilo-Watt-hours (kW-h) from when the unit is turned on. This allows you to plug in the meter and measure consumption over an average day, or week, in order to get a more realistic value for the energy routinely consumed by the system.

Another important test these meters can perform is to measure the power factor. The mains electricity supply is an oscillating wave. When that waveform interacts with the electronics of the power supply most of it is used, but not all of it. In poorly made or damaged power supplies a large proportion of that power might be dumped back into the power grid unused; you pay for the power, but it's not used by the machine. The difference between the amount of power supplied and the amount used by the equipment is called the power factor. It's expressed as a number between 0 and 1: A high value, 0.7 or better, is very good; a figure between 0.5 and 0.7, is adequate; a figure below 0.5 is poor and you should consider buying a better power supply.

When measuring the power factor you shouldn't turn the machine on, take the measurement, and then switch off. As the components warm up the power consumption can wander slightly, and so its advisable to leave the meter in place for an hour or so, measure the initial power factor, measure the total power consumed after at least an hour or two, and then measure the power factor again before switching off. This provides a more reliable figure for consumption.

Whether or not it's worth replacing a power supply is a matter of need or cost. For equipment that's powered from an off-grid supply, minimising consumption is very important irrespective of the cost, you may choose to replace the power supply with a higher quality alternative (off-grid power is inevitably more expensive than grid power). For a mains-power system it's a matter of cost. Let's say the system uses 0.25kW-h, of which half is the computer box, and runs for eight hours a day, five days per week, 46 weeks per year. That's a total consumption of 460kW-h/year. Taking a typical UK electricity tariff as an example, at £0.14/kW-h that's £64/year. If you could improve the power factor from 0.5 to 0.7, that saves around 20%, or £12/year but only half of that figure, £6/year, is due to the computer box and its internal power supply. Whether or not it s worth replacing the power supply depends upon the costs of replacement. A high efficiency power supply might cost £50. Assuming a usable life of four years, it will never repay the cost of the new supply. Let's say that the power factor was 0.4; increasing to 0.7 would save almost £10/year, which makes it almost worthwhile. That's why the "adequate" power factor is 0.5; a result any higher is unlikely to merit changing the power supply, unless you were able to do so at a negligible cost.

There are of course other ways to save energy and costs. A laptop computer will generally use a tenth of the power of a desktop, its monitor and speakers. An LCD monitor will generally use around half to two-thirds of the power of a plasma display. Again, these are the kind of decisions you need to take before you buy or obtain equipment; changing components after the purchase is unlikely to pay for itself in the amount of power saved, even if you re-sell the unwanted parts.

Figure 3.7.

A power monitor



3.4. PUTTING IT ALL TOGETHER

By walking through the components which make up a desktop computer system you should hopefully have a better idea of what's in the box, and what each of those parts does. If you wish to take these ideas further, we'll cover maintaining and assembling computers later. Hopefully, what you have gained from studying the innards of a computer system is an understanding of the significance of system specifications, and what those specifications relate to. This should enable you to purchase, or plan the upgrade of systems with a little more attention to the details which influence the power of the system, and the energy consumed by the machine. Box 3.6 contains a check-list of points to consider when purchasing a system, or deciding how to upgrade it.

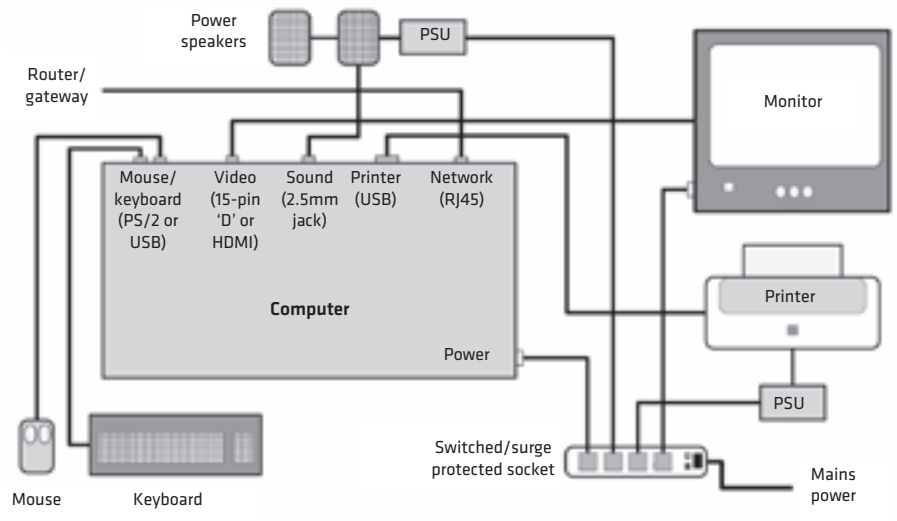
Once you have a system you have to connect all the parts together. With a laptop that's done already – all you do is plug in the power and network connectors, and perhaps an external speaker. With desktop systems it's a little more difficult. You have to connect many different cables to the main computer unit, and there are some simple rules which can help protect and preserve the equipment. Figure 3.8 shows a very simple desktop system as an example.

As a general rule you connect the thickest and heaviest cables first because, if accidentally tugged, they might snag and damage the other more delicate lightweight cables. It's not damage to the actual wire that's the problem. If a lightweight connector is tugged hard at right angles to its connector it can be wrenched out, damaging the plug, the socket, or both. Damage to any of the essential sockets on the motherboard might necessitate replacing the whole board.

All computer equipment needs good ventilation. In a small or confined space the machine will very quickly raise the ambient temperature well above the human comfort zone; and in a very small space it might be unhealthy for the machine too. If you put any equipment inside a cupboard ensure that there are large ventilation holes top and bottom, either in the back or side of the cupboard, to allow air to circulate. If you install the computer unit, the monitor, or a

Figure 3.8.

Schematic of a desktop computer system



laser printer in a small recess or under a shelf, again, make sure there is 15-20cm of free air space above and around the machine to ensure the free circulation of air. And no matter how tempting, it's never a good idea to leave papers, books and files on top of computer units or monitors, or in any position that obstructs the ventilation ducts of the equipment. Not only does insufficient ventilation make the machine run hotter, shortening the life of its components, the thermal controls of the CPU and chipset will slow down the computer to reduce the heat load generated.

Next, electrical power. The electrical mains brings with it certain problems. Electrical noise spikes, from heavy industrial machines or loose electrical connections, can accumulate damage to the components of the power supply and shorten its life. In the worst case a prolonged power surge can disable the power supply, and in the very worst case (such as a lightening strike) damage the motherboard as well. The simplest solution to this is to plug all the mains plugs into a single switched and surge-protected socket. Surge protection minimises electrical noise and hopefully blocks the power surges. More importantly, having a single point to switch off the whole system means that with one switch it's possible to disconnect all the "standby" power drains of the computer unit, printer and other power supplies. Arguably, assuming you always remember to switch off when the machine is not in use, the electricity saved over the lifetime of the system will easily pay for the cost of your surge protected socket.

When laying out power cables always leave a little slack. Ideally you place the switched socket near or next to the computer – partly so that it's always convenient to turn off, but mainly so that you can plug in all the mains cables without leaving loops of cable which might get caught and pulled. If you have a lot of cable left over you can coil the cable loosely, and tie with a cable tie to keep it together. If the power point for the machine is not nearby, tie or fix the cable from the switched socket to the leg of the desk/table. In the event of the cable being pulled hard it prevents the socket and its attached cables being dragged away.

Once all the mains cables are connected up and tidied away, you can move on to the lighter-weight cables. Again, work from thickest to thinnest. The next is probably the video cable, then USB and network cables, and finally the audio cables. Again, leave loose coils in the cables which might get tugged – especially the mouse, keyboard and network cable (or the phone line is you are using a modem) – to prevent accidental damage. Web cams are usually fastened to the top of the monitor; if you don't have a compatible bracket you can get Velcro pads, sticking one to the top of the monitor and the other to the bottom of the web cam, to stick the two together. Headsets and headphones also need to be carefully located to prevent the cable from being tugged. You can buy small metal cable clamps with sticky foam backs – stick them to the side of the case and then clamp the cable in the jaws to reduce the likelihood that it will be wrenched out. Newer systems tend to have USB sockets on the front; older ones do not. If you regularly plug in MP3 players, cameras or other devices, get a short USB extension cable, plug it in around the back and then trail it around the front of the computer. This prevents continual rummaging behind the computer which will, eventually, cause cables to work loose or break.

Finally, let's look at the dazzling array of port connectors which emerge from the motherboard through the back of the computer case. It's very difficult to plug a connector into the wrong socket as they are all constructed very differently. More often the question arises as to whether the motherboard has the connector you require.

There is no standard set of ports on motherboards, although with more expensive boards you are often presented with a wider range of sockets. Depending upon the design of the motherboard you'll have a variety of standard ports (video, USB and Ethernet), but new motherboards might not be compatible with older peripherals. If you're in the situation where you need a port but there's none provided, or if there are not enough provided, then you'll need to install a PC card in one of the slots to provide the extra interface(s) required. For very old ports, such as the 9-pin serial connectors used by various old types of technical equipment, you can buy short adapter cables which will connect the serial cable to a USB socket (although not all software may be able to work with this system).

Table 3.3 lists various ports and shows the design of their associated connectors (the index numbers from figure 3.3 are also listed in curly brackets). Over the last few years we've seen some ports – like the parallel printer port, and

more recently the PS/2 keyboard and mouse ports – fade away as the world has switched over to the ubiquitous USB. Some, like the DVI video port, were often included but seldom used. Irrespective of their popularity, table 3.3 lists commonly occurring motherboard connectors so that you know what's what.

To return to a point made earlier – even if talking about it feels challenging, *hardware cannot perform any useful function unless told to do so by a computer program!* Now to make the hardware into something useful we have to load an operating system onto the machine and that is an equally complex subject.

Box 3.6.

Purchasing and upgrade check-list

Upgrading is always better in the short-term, buy new for the long-term

Providing you can buy affordable components, upgrading a system is often cheaper and less ecologically damaging than buying a new system. Most systems, if cared for, will last three to five years. With some upgrading you could get that up to between five and seven years of "useful" life. Unless you have to buy more powerful machine because you are required to use a certain operating system, there are a number of ways to increase the power of an existing system and prolong the life of its components:

- Add more memory
Adding memory allows more programs to be run in RAM rather than having data continually move to and from the hard disk. Doubling the capacity of the RAM, or better, should make an appreciable difference to a slow computer.

- Add a new hard disk
For an existing system you can extend the storage space by adding another hard disk. That's the easiest option, but it will use more power than installing a new, larger hard disk (but always check the largest size of drive the motherboard can address). The drawback of installing a single larger disk is that you have to back-up all the data on the system, re-install the operating software, and then put all the old data back again.

- Upgrade the motherboard
Many parts vendors sell packs with a compatible motherboard, processor and memory which is the cheapest way to upgrade the whole system. You should be able to reuse the case, drives and other parts of the system with the new motherboard. If you must install a new power-hungry operating system, this is often the cheapest, most ecological option.

- Reuse the monitor
The life of the monitor is likely, if it's cared for, to be much longer than the computer unit. If you buy a new computer, you can save money by reusing your existing monitor.

- Upgrading laptops
Laptops are a more technical problem. Memory modules are not always easily available. Most 2½ hard drives will work with a range of laptops, but always check with the user manual first. The most problematic part of laptops is the battery. If regularly discharged, it can lose a significant amount of storage capacity in 12-18 months. If this happens, there are suppliers who specialise in battery replacements. Whether buying laptop memory modules, hard drives or batteries, it's worth shopping around as prices vary considerably.

Don't purchase on the processor speed alone (both new systems and upgrades)

As noted earlier, while the speed of the process is important, having a high-quality chipset, with a larger amount of cache memory and a fast front-side bus will provide far more power overall especially when combined with a large memory capacity.

Buy the machine you need, not what you are being sold

Machines sold for the home market are increasingly being optimised for video gaming they have power-hungry video cards, and a large processor and memory capacity. If you want a machine for browsing the internet and basic office tasks, buying that extra power in not necessary. Simple usage requires only a low-end specification; video gaming, large quantities of video editing and graphic design, require high spec machines.

Match the power of the machine to the operating system you wish to run

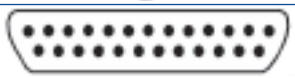
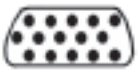
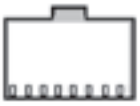
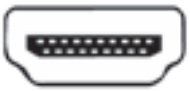
Table 3.1 lists the bare minimum system specification quoted to run various operating systems. Ideally you should double or triple these figures. If you are buying second-hand equipment, and what you're buying is slightly under-spec, you could perform a motherboard/processor upgrade to boost the power at a lower price than a new machine.

Beware when buying hardware for Linux, always check for compatibility!

Always check compatibility before buying computers or peripherals for Linux. Laptops from certain manufacturers are a problem, as are some brands of scanner, printer and imaging devices. This is much less of a problem than it used to be as more manufacturers now sell Linux-compatible hardware, or provide driver software for these devices on their websites.

Table 3.3.

PC ports and connectors

Port	Description	Connector
PS/2 Keyboard {15}	Round six-pin PS/2 socket coloured purple on later boards. It was the keyboard connector until recently when it was superseded by USB.	
PS/2 Mouse {16}	Round six-pin PS/2 socket coloured lime green. It was the standard mouse/pointing device connector until superseded by USB.	
Parallel printer port {17}	25-pin D-type socket coloured magenta on later boards. Superseded by USB.	
RS232 serial port {18}	9-pin D-plug coloured dark green on some boards. Superseded by USB.	
VGA video port {19}	VGA graphics output 15-pin D-socket coloured dark blue on some boards.	
Audio lines {22}	Between two and six, 2.5mm stereo audio jack sockets (microphone which is usually mono). Audio input/output from the sound card. There are various colour codings in common use: Lime green line out/front speaker out* Light blue line input Pink microphone input Orange centre/subwoofer* Black rear speaker out* Grey side speaker out* (* used with surround-sound audio)	
Ethernet port {20}	RJ45 Ethernet connector. Usually with indicator lights to show when connected/in use.	
Firewire port {21}	IEEE 1394 "firewire" socket. Alternative format to USB, used on some video cameras/digital AV equipment.	
S/PDIF coaxial output port	RCA (or phono) coaxial connector used to connect digital sound output to home theatre system. Usually yellow coded.	
USB port {23}	Universal serial bus (USB) standard type-A socket.	
Digital video port	29-pin Digital Visual Interface (DVI) socket used to connect digital video output to home theatre system/video capture devices.	
High definition multimedia interface (HDMI)	HDMI-type connector used to connect high-definition (HD) video monitors/video games displays.	
Wide RS232 serial port	25-pin D-plug. Was once used for connecting modems, but superseded by Ethernet/USB devices a decade or so ago.	
S-video port	Four-pin DIN socket provides video output to a compatible TV/ video display.	

A practical guide to sustainable IT

This practical guide to sustainable IT offers a detailed, hands-on introduction to thinking about sustainable computing holistically; starting with the choices you make when buying technology, the software and peripherals you use, through to how you store and work with information, manage your security, save power, and maintain and dispose of your old hardware. Suggestions and advice for policy makers are also included, along with some practical tips for internet service providers.

Written by IT expert and environmentalist Paul Mobbs, the purpose of the guide is to encourage ICT-for-development (ICTD) practitioners to begin using technology in an environmentally sound way. But its usefulness extends beyond this to everyday consumers of technology, whether in the home or office environment. We can all play our part, and the practice of sustainable computing will go a long way in helping to tackle the environmental crisis facing our planet.

This is also more than just a "how to" guide. Mobbs brings his specific perspective to the topic of sustainable IT, and the practical lessons learned here suggest a bigger picture of how we, as humans, need to live and interact in order to secure our future.

The guide is divided into 12 sections (or "units"), with each unit building thematically on the ones that have come before. They can be read consecutively, or separately. The "unit" approach allows the sections to be updated over time, extracted for use as resource guides in workshops, or shared easily with colleagues and friends.

The guide has been developed on behalf of the Association for Progressive Communications (APC), with funding support from the International Development Research Centre (www.idrc.ca). It is part of a APC's GreeningIT initiative, which looks to promote an environmental consciousness amongst civil society groups using ICTs, and amongst the public generally. Other publications and research reports completed as part of the GreeningIT initiative can be downloaded at: greeningit.apc.org

APC

IDRC | CRDI