

PAUL MOBBS

A practical guide to sustainable IT



A practical guide to sustainable IT

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.

Table of contents

01. Introduction – What is sustainable Information Technology?	9
An introduction to the themes of this series, and an explanation of the reasons why we should view our use of technology as a “system” rather than as separate processes. Too often we focus on the demands or specifications of “the machine” – instead we should look at our information needs.	
02. The ecology of information technology Information Technology	15
Is a tool like all other human tools that we have created through the ages. This unit examines the general issue of how the “human system” uses technology – how digital technologies work for us, how these technologies influence the wider environment we inhabit, and the emerging restrictions on our future use of these tools.	
03. Hardware: Putting the system together	27
An exploration of computer hardware, in its various guises, and how the choices we make about hardware influence the ecological footprint of our use of IT. While many may look primarily at power consumption, the manufacture of digital electronics also has a major impact on the global environment. In this unit we try and measure these impacts, and find some ways to manage our demand for information systems.	
04. Operating systems: Making the system work	51
What makes computer hardware more than an expensive collection of electronic circuits is the software we use to give the machine a purpose. This unit looks at operating systems and computer hardware, the issue of intellectual property rights, and the influence of our choice of software on the ecological performance of the hardware we use.	
05. Desktops: Programs, peripherals and gadgets	69
The computer desktop is the heart of our daily interaction with the power and versatility of information systems. How we configure the desktop, but also how we use the various gadgets and peripheral devices that we connect to the system, has a big effect on the ecological footprint of our use of technology.	
06. Information storage	85
Information is a “resource”, and obeys many of the physical laws that govern our use of matter and energy. How we choose to store and back-up information has a direct bearing on the ecological impacts of our use of IT in general, and also the reliability and resilience of the information that we store.	

07. Local and global networks	101
<p>become so versatile, but that in itself has spawned the exponential growth of both data transactions and the growth of power-hungry online storage systems to support them. Understanding how the network consumes power, and how our use of the network drives consumption, can allow us to keep these demands under control.</p>	
08. Maintenance, security and resilience	116
<p>All tools need looking after if they are to have a long and productive life, and that requires that we understand a little more about how they work and how to care for them. This section considers the basics of system maintenance – both hardware and software – and the basic physical security precautions to take in order to prevent theft and damage</p>	
09. End of life and disposal	133
<p>When does the life of digital equipment end? Is it when it no longer functions (it's broken) or when it can't/isn't capable of performing a useful function any more? This unit examines how much life we can extract from our information systems, and what we should do with them when they have served their useful purpose.</p>	
10. Reuse and reclamation	141
<p>Can we recycle digital equipment? And how can we determine what's best to do with old equipment when it no longer serves its primary purpose? To answer these questions we need to consider the options available for reusing or upgrading old equipment, and identify what the limitations are on using old equipment.</p>	
11. Renewable power	149
<p>Information systems consume electricity, and while in most cases the power might be sourced from the electricity grid, there are alternatives. Renewable electricity can offer an alternative, lower carbon source of electricity. For those who live away from the electricity grid, or where the grid is unreliable, it can also offer a means to use IT without the restrictions imposed by the local grid.</p>	
12. Further information and index	159
<p>Annex A. Suggestions and advice for policy-makers A digest of the ideas and information for managers and policy-makers, and those campaigning on access to information and communications technology.</p> <p>Annex B. Suggestions and advice for internet service providers/server operators A check-list of procedures and options for IT managers, service providers and data centres.</p>	

WHAT IS SUSTAINABLE INFORMATION TECHNOLOGY?

Welcome to *A practical guide to sustainable IT*. This publication will guide you on a journey to examine how the information systems which support our lives work, and how their ecological footprint has grown to become a distinct part of the human impact upon the planet.

In developed nations, information systems are now an indispensable part of most people's lives. In developing nations, information systems are a means of enabling communication and exchange where no such infrastructure existed before, and that in turn creates many new opportunities from education to economic development.

But the development of the global communications network, and the equipment attached to it, has come at a cost. The reality is that these services are not "free"; quite apart from the financial cost, the production and use of these systems is affecting the global ecosystem which supports life on earth. For example, by some estimates the impact of information systems and telecommunications has as great an effect on climate as the air transport system.¹ How great those impacts are, or how we can reduce or manage them, is defined by the choices we make when buying, using or disposing of these machines.

If we're going to have access to these technologies in the future, without harming the planet or depleting the resources from which they are manufactured, we must learn to plan our use of information systems to maximise material efficiency while minimising their environmental impact. In short, we must make them more sustainable.²

1. Computing climate change – How much carbon dioxide do computers emit?, The Economist, 24th August 2009. www.economist.com/node/14297036

2. Wikipedia, "Sustainability". en.wikipedia.org/wiki/Sustainability

1.1. A PROBLEM OF WORDS AND DEFINITIONS

How do we define “sustainable”? That’s a question which has been taxing the brains of policy makers, scientists and environmentalists for the last few decades.

The problem is that over the years different groups have applied slightly different meanings to the term – often based around the objectives of the study or policy they were considering at the time. A literal definition of the term means, “the ability of an activity to endure or function for a certain period of time or perhaps indefinitely”. Unfortunately, for computing, telecommunications and related technologies there is no agreed standard on how we should measure these impacts. Narrow definitions of sustainability may only consider one or two features of an activity without looking at how other important factors interact with the wider world. More general definitions may look at a larger number of related factors which collectively define how “sustainable” the activity is when considered as a whole.

Many definitions take their cue from the United Nations World Commission on Environment and Development report (also known as the Brundtland report), published in 1987, which launched the modern policy debate on sustainability:³

Development that meets the needs of the present without compromising the ability of future generations to meet their own needs.

Going from this all-encompassing definition of sustainability to one which fits the peculiar features of information technology is a very big jump. We could take a very narrow definition, for example, looking purely at the electricity consumed by machine hardware; but if we are to give the consideration of sustainability a broader and more meaningful definition then we have to consider the operation of the whole system.

Assessing the sustainability of our use of information technology isn’t just a matter of measuring environmental pollution. It’s an all-encompassing view of how we design, organise and operate the information systems that allow us to carry out our work and live our lives – and doing so in a way which considers not only the impact on the planet today, but also how we develop, use and preserve information resources for ourselves and others in the future.

Tackling this is obviously a very broad agenda. While our use of computers has a direct effect on the environment, if we’re looking at how “sustainable” our use of those systems is we must include the related factors which influence that process. Yes, we must look at the machine hardware – how those machines are made, but also how they are disposed of. In addition it’s vital that we consider the features of the operating systems and programs that we use, and the processing and storage of information – since it is the value of information which defines the purpose for which we build these machines. If we’re considering how sustainable our information keeping is, then the reliability of our record keeping/data storage systems is important – and that in turn can raise issues related to intellectual property rights and other legal considerations.

This guide is organised as a series of themed sections or what we call “units”, each looking at the sustainability issue from a slightly different angle. The objective within each unit is to get you thinking more systematically about how you can physically organise your data, programs or equipment to achieve the tasks you want to carry out.

3. World Commission on Environment and Development. Paragraph 1, chapter 2, Our Common Future, report to the UN General Assembly, 1987. www.un-documents.net/wced-ocf.htm

1.2. IT'S NOT EASY BEING "GREEN"

To date there has been a lot of work around the issue of *Green IT*⁴ – the design and organisation of computers to achieve certain ecological objectives. Examples of recent green IT studies and useful publications are listed in the bibliography at the end of the guide. Green IT studies and programmes often take energy consumption as the lead indicator of sustainability – for example, the *Energy Star* labelling scheme developed by computer manufacturers and the US government in 1992. Others go further, considering the effects of disposing of those machines on the environment. More recent work seeks to examine the machine as part of a “holistic system”⁵ – considering the whole life-cycle of the machines during production and use.

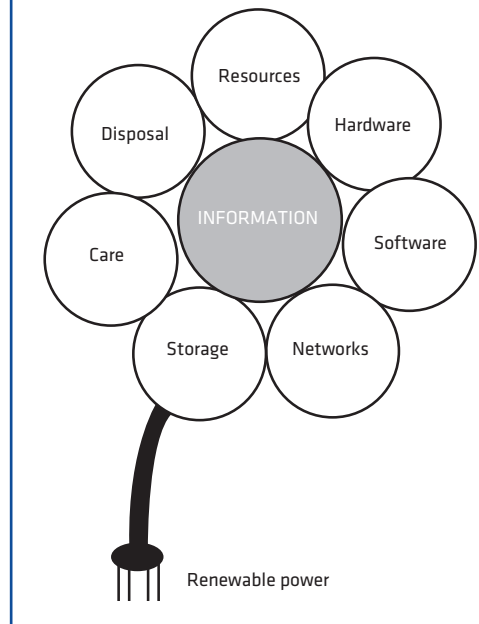
Across the array of green IT publications available there is little connection made to the other equally important issues related to our use of computers – such as systems standards, data formats or information security. That's a pity because there are many common themes between “green” issues and the everyday concerns which arise from our use of computers and information networks. It's by combining the existing “green” ideas of sustainability, with a far broader consideration of the factors which influence our everyday use of IT, that has led us to produce this guide to “sustainable information technology”.

1.2.1. The “information flower”

Rather than base our considerations around a definition which begins with hardware, in this guide we'll look at our use of IT from the point of view of the *information* which the system manipulates. Information is, after all, the vital human resource which we wish to sustain, independently of the hardware with which we choose to process it at any point in time. In order to give this definition of sustainability a practical meaning we've broken down the operations

Figure 1.1.

The sustainable IT “flower”



or processes which define our use of information technology, and from each we have drawn a set of ideas and activities which allow us to review, audit and improve the way we organise and work with information.

To illustrate this, we've organised the guide's themes around the holistic motif of an *information flower* – each petal representing one aspect of the sustainable use of IT, and the stem as the source of electrical power. Each unit in the guide corresponds to an element of the flower:

Resources, or perhaps more appropriately, The ecology of technology – a general outline of how our demand for computers affects the world, but also how those demands have changed us. Our use of energy sources and natural resources isn't random or accidental. We've progressively developed technologies that have required ever-more specialised resources and materials as time has passed. Technological progress has inherently increased the complexity of human

4. Wikipedia, “Green computing”. en.wikipedia.org/wiki/Green_computing

5. Murugesan, San. *Harnessing Green IT: Principles and Practices*, IT Pro (Journal of the IEEE Computer Society), January/February 2008. www.comp.dit.ie/rfitzpatrick/Business_Perspectives_slides/Papers/Harnessing_Green_IT_-_Principles_and_Practice_-_San_Murugesan.pdf

systems, and this has implications for our future as the production of certain essential resources is limited. In this unit we'll look at the resource use inherent in our consumption of information technology, but also at the wider human resource issue, and the ecological limits which are increasingly apparent as the globe's appetite for natural resources grows.

Hardware: Putting the system together – the machine and its impacts on the environment. Today computer hardware comes in many different forms, from the large mainframe and server complexes of large organisations to the mobile phones people carry down the street. In this unit we'll look at the resource implications of computer hardware, and how our choices about design specifications – but most importantly the period of time we plan to use the equipment – affect the ecological footprint of our information processing.

Software – the programs we use on the machine to perform the tasks we require. This is a large subject and so it is divided into two sections:

- **Operating systems: Making the system work** – the programs which make the computer function. Our choice of operating system has an effect on how the hardware functions, and more importantly it can have an impact on the service life of the machine.
- **Desktops: Programs, peripherals and gadgets** – in which we consider how we interact with our computer systems, and the programs we use to perform the information processing tasks that we demand from our computer systems.

Information storage – or more precisely, information storage and backing-up. Information processing is the reason we have created a diversity of machines and networks, and in turn preserving that information is a vital part of making the system secure, resilient and sustainable. In this section we look at how we store information, the options for backing up data, and how our choices create different levels of security and determine our ecological footprints.

Local and global networks – it is by linking computers together that we have created such

a powerful system for human interaction and exchange, but this in turn has amplified the environmental impacts of information technology. As more information and services are digitised, so the energy and resource demands of the system grow. In this section we look at how we network computers together to meet our own information needs, and how all machines have become increasingly dependent upon the global network, both for communication and as a means of storage and distributing/updating software.

Care, or more precisely, Maintenance, security and resilience – looking after our information systems. Like any complex tool or mechanism, computer systems need care and maintenance. While solid-state technology of the silicon chip may have no moving parts, there are many aspects of everyday computer use which require physical care and maintenance. Besides hardware, the computer's software and the file system also need occasional attention to keep the logical elements of the machine working reliably. In this unit we look at how we can care for our information systems, to keep them functioning reliably, but also to get the greatest possible service life from the hardware. We'll also look at how we can plan for system failures, or more serious events, to secure our information even when the hardware fails.

Disposal – in which we examine how we can deal with equipment which has reached the end of its useful life. This is divided into two parts:

- **End of life and disposal** – how to know when equipment has reached the end of its useful life, and how to dispose of it safely. The components inside computer systems can fail, but sometimes they are discarded because of other factors, such as changes in software or operating systems, which necessitate using more powerful hardware. In this unit we'll look at the environmental risks of computer hardware, how to determine if any item of hardware can be used for other tasks, how to dispose of it if necessary, and the problematic issue of erasing data to prevent the disclosure of sensitive information.

- **Reuse and reclamation** – how machines can be upgraded or recycled and given a new lease of life. This is a more ambitious unit, which seeks to take people beyond simply “using” a computer, to thinking more technically about how they can reconfigure their hardware to serve other needs. We’ll look at the options for donating old equipment, but also how components can be upgraded to extend the lifetime of the equipment, or installed with different software to give the machine a new function to support your information infrastructure.

Renewable power – the options for sourcing electricity from lower carbon and renewable resources. Computers require electrical power to

operate, and the generation of electricity is one of the major carbon-intensive industries on the planet, which gives rise to a number of pollutants, besides carbon dioxide. In this unit we’ll look at alternative options to sourcing electricity from the power grid, and we’ll also look at going off grid – both to exploit renewable energy resources and as a means of taking computers outdoors, beyond the reach of the mains grid.

The final unit contains a subject index for all the sections, a glossary of the technical terms used in the guide, as well as a bibliography and directory of online information sources. Green advice for policy makers, and a check-list for internet service providers are then included in the appendices.

SECTION 2

THE ECOLOGY OF INFORMATION TECHNOLOGY

This unit examines the way in which our use of information systems interacts with the human environment. It encourages you to think about the context in which we use computers and the internet, how our choices effect the wider environment, and more importantly how global economic and environmental issues are inextricably tied to our use of information technology. This not only includes the direct impacts of power consumption, or the mining of minerals, but also the changes to the economy that have resulted from the greater technological development and mechanisation of what I call the “human system” – the complex and interlinked network of interactions that we have created for ourselves.

2.1. TOOLS, TECHNOLOGY AND HUMAN DEVELOPMENT

You cannot look at the “ecology” of information technology without reflecting upon the human system which has created these devices. Information and communications technology are part of the human development story. The trends which have created our modern day world are at work in the reasons why we created these technologies, and are behind the explosive growth in their use around the world.

As information and communications technologies have become more widespread, especially consumer-oriented technologies such as mobile phones, they have become “invisible”. They have become such an everyday part of people’s lives that we blindly accept they are there. In fact, we only understand their significance to our lives when the electronic device, or the information network it connects to, breaks down. The implicit association of these technologies with a “modern” lifestyle has in turn become a driver for their adoption in less developed states. The rapid adoption of information technology has not only enabled greater economic development, it has arguably brought with it social and political change as people have used these same media as a conduit for political expression – for example, the events of the *Arab Spring*.¹

Computers and information systems are tools – human-made technologies intended to extend our biological capabilities. We might invest our lives in them, commit our most secret facts and thoughts to them, and rely upon them to organise our work and social lives, but if we are to understand how information systems operate then we have to examine them as we would any other aspect of human society; and that means looking at the role of these tools in the human system, and how that in turn relates to the global environment that humanity is a part of.

Like other tools made by the human species, technology extends the physical and/or mental capabilities of its user, achieving practical ac-

tions which they ordinarily would be unable to accomplish. The difference with information systems compared to, for example, a hammer, is the advanced level of complexity required to create and use these tools. If we’re considering how “sustainable” information technology is then this inherent complexity² has significant implications for the ecological footprint of the technology, and also its future viability.

2.1.1. Convergence

The importance of programmable electronics has been the ability for one electrical device to serve many different purposes. By varying the software program, the electronic hardware can perform a wide variety of complex operations. As a result, not only have digital electronics made the production of technical or consumer goods simpler and cheaper, what we increasingly find is that different kinds of technology are *converging*³ into a single device – removing the need to have separate devices to achieve those same functions. That may seem environmentally beneficial – having one device instead of many; but what has happened is that their success has led to their adoption by an ever greater part of the global population, and so over the last two decades the material and energy footprint of IT has steadily grown – today its carbon footprint is of a similar scale to that of the global air transport industry.

The clearest example of technological convergence is the smart phone. This is a telephone with a video display screen and digital camera built in, capable of playing music radio, television and radio. Most importantly, backing-up these various functions is a powerful computer able to interact with the internet via communications networks – allowing all the information stored or captured on the device to be shared and additional information downloaded. As a result we are no longer restricted to communicating with spoken words; we can converse in text, images, and video – or even in the machine instructions


1. For a detailed exploration of the role of social media in recent protest movements see Mason, Paul (2012). *Why It's Kicking Off Everywhere: The New Global Revolutions*. Verso Books. ISBN 9781-8446-7851-8.

2. Wikipedia, Complexity. en.wikipedia.org/wiki/Complexity

3. Wikipedia, Convergence (telecommunications). [en.wikipedia.org/wiki/Convergence_\(telecommunications\)](http://en.wikipedia.org/wiki/Convergence_(telecommunications))

Figure 2.1.

Global growth in communications technology



(or language) required to manipulate the networks we are connected to. As shown in figure 2.1,⁴ the global growth in the new digital communications services has been exponential for the last three decades.

This same kind of transition took place forty years ago, but then as a result of the electronics revolution created by the technological advance from thermionic valves to the transistor. This not only made the systems cheaper, it allowed the devices to be smaller and more mobile, meaning they could be used in ways they hadn't been used before. For example, the pocket transistor radio or cassette player used pre-existing technologies, but in a way which created a new mass market for electrical goods. Today, it's the switch from analogue transistors to digital microprocessors that is achieving the same revolutionary transformation of consumer behaviour.

2.1.2. Human tools are based on rocks

Within the growing ecological footprint of information technology one of the most important limiting trends is the reliance of human systems on ever-rarer materials. To understand the sig-

nificance of this we have to go back to the very first human technology: rocks.

In our ancient history the first human tools were made of stone, and stone is perhaps one of the most plentiful resources on the planet. Wood, plant matter and animal tissue were equally important resources, but little of this material has survived. For perhaps two to two-and-a-half million years humans relied upon stone tools⁵ to support their daily activities. The reason that the Stone Age ended was not that we ran out of rocks, it was that we found something which was more useful than stone: metals.⁶ Metals are also made from rocks – metal ores⁷ – but these rocks are far less plentiful as a proportion of all the stones available. This meant that in ancient societies metals had very high price, and were only used for very special applications.

What eliminated the rarity of metals was industrialisation, and more importantly the early use of fossil fuels to power the furnaces which made metals. Before the eighteenth century the limitation on metal production, even where the rocks it is produced from were plentiful, was the amount of wood required to fire the furnaces to produce pure metal. In short, the limitation was the quality and scale of the energy available to smelt the ores, not just the minerals resources available to produce metals. This restriction, known as the Law of the minimum,⁸ applies to all other living organisms on the planet too – and in that sense today's technological society is no different from our ancient arboreal past.

In ecological terms, the limiting factor is time – and the shift to coal allowed the human economy to escape the restrictions imposed by nature. Wood represents stored solar energy, and that takes time to grow and mature, and so metal production was constrained by the sustainable limits of local wood production. For that reason coal has been used in metal production since the time of the Ancient Greeks.⁹ Coal represents thousands of years of stored solar energy, meaning it has a higher energy density than wood and so produces

4. Information sourced from the World Bank's global indicators dataset. data.worldbank.org/topic

5. Wikipedia, Stone Age. en.wikipedia.org/wiki/Stone_age

6. Wikipedia, Metallurgy. en.wikipedia.org/wiki/Metallurgy

7. Wikipedia, Ore. en.wikipedia.org/wiki/Ore

8. Wikipedia, Liebig's Law of the Minimum. en.wikipedia.org/wiki/Liebig's_Law_of_the_Minimum

9. Wikipedia, The history of coal mining. en.wikipedia.org/wiki/History_of_coal_mining

more heat for the same weight of fuel. While the coal is easily accessible and reasonably priced, this allows metals to be produced on a much larger scale than when wood was the only fuel source.

What allowed the Industrial Revolution to take off in the second half of the eighteenth century was the interaction of the technology of coal production and the improved technol-

ogy of metal production. As metal technology improved we see the development of the first steam engines. The first major application of the steam engine was pumping water from mines, which allowed both greater metal and coal production from deeper mines – and it was this self-sustaining process which drove the Industrial Revolution.

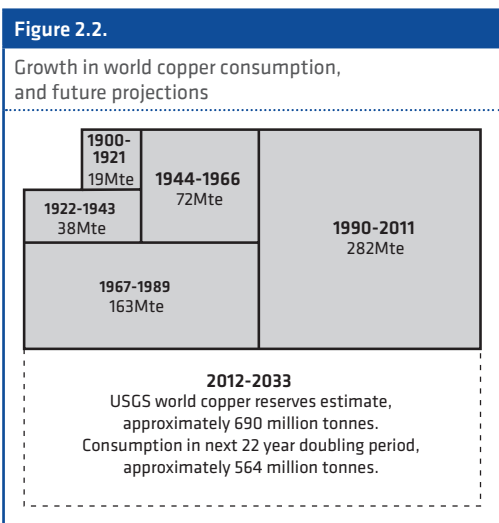
2.2. GROWTH, CONSUMPTION AND RESOURCES

The effect of the information and communications revolution and its interaction with digital electronics has a direct parallel with effects of industrialisation on the use of metals. Independently the elements which make up the technology – computers, communications lines and digital information – existed well before the recent growth of electronic networks. It was only when these technologies were brought together, at prices affordable in a mass market, that the power of the network generated its own self-sustaining growth. It's not just that these trends allowed people to do pre-existing jobs more easily; the process created wholly new applications which caused the level of consumption, and the economy as a whole, to grow. Like the step change in energy and resource use 250 years ago with industrialisation, or the first use of metals over 7,000 years before that, the information and communications revolution is creating a fundamental change in the organisation of human society.

2.2.1. Copper - an indicator of technological development and sustainability

As noted above, advances in human technology consume comparatively rarer metals.¹⁰ Of these, copper provides one of the best case studies of the importance of minerals to the global economy, but also the fragility of that global system due to the natural constraints on human resource consumption. Copper is one of the most important minerals in the technological economy. It's important for micro-electron-

10. Paul Mobbs/Free Range Network (2011). The Limits to Technology. www.fraw.org.uk/fwd/?wslimits



ics, although the bulk of annual consumption is for electric cables, pipes and metal alloys. For example, a quarter of the weight of a mobile phone is made of metal,¹¹ and up to half of that could be copper.

As the world economy grows, so the amount of copper demanded by the world economy grows too. Copper has been in use for at least 7,500 years, but more than 95% of all copper ever mined and smelted into metal has been extracted since 1900.¹² Despite the fact that copper is one of the most recycled metals, and per-

11. U.S. Geological Survey (July 2006). Recycled Cell Phones A Treasure Trove of Valuable Metals, USGS Fact Sheet 2006 3097. pubs.usgs.gov/fs/2006/3097/fs2006-3097.pdf

12. U.S. Geological Survey, Copper Statistics and Information. minerals.usgs.gov/minerals/pubs/commodity/copper/

haps 75% of the copper ever mined by humanity is still in use, in 2011 around 16 million tonnes of new copper were produced from mines around the world. This is because as society becomes more technological, and especially as many developing countries begin to build power and telecommunications infrastructures, the demand for copper continues to grow year on year. As with many natural resources, the amount of copper in the top kilometre of the earth's crust is huge (around 900,000,000,000,000 tonnes, or 5-million years worth of production) but only a tiny fraction of these reserves is economically viable to extract. According to the information from the US Geological Survey (USGS), which studies the global availability of the most important minerals, the amount of copper available in the future is around 690-million tonnes.

In figure 2.1, the growth in mobile phones, like the rest of the economy, is an exponential curve. Any exponential trend has a characteristic called the “doubling time”, a fixed period of time in which the quantity will double. Copper is interesting because its growth over the last century is a well-defined exponential curve with a doubling time of about 22 years. As shown in figure 2.2, we can draw boxes to illustrate how the amounts of copper consumed by the world economy grow with each period of doubling time – from around 19 million tonnes in the first 22 years of the last century, to about 280 million tonnes in the period which is just ending.

The issue for our future use of copper, and the viability of the technologies which depend on the metal, is that the next 22-year doubling period will consume about four-fifths of the remaining copper resource identified by the USGS. However, the production of minerals can't carry on at a constant rate. As the best reserves are used up first, so it becomes progressively harder to find and produce what remains, and more expensive. It is argued that rather than continuing to rise, as the quality of copper ore falls production will reach a peak and then decline.¹³ What this means is that, as they have done for the last decade or so, average copper prices are likely to keep rising as the economy demands more of the metal – and that will impact upon our use of information technology in the years to come.

13. Wikipedia, Peak copper. en.wikipedia.org/wiki/Peak_copper

2.2.2. The limits to growth

Copper is not the only limitation on our sustainable use of information technology, and is not the only useful indicator of the state of human technological sustainability. Forty years ago, a group of scientists produced a study which forecast that there were “limits to growth”¹⁴. This caused great controversy at the time, and since then many politicians and economists have claimed that the projections were erroneous. However, not only does more recent research show that the forecasts are still largely correct when we use the latest data,¹⁵ other work has shown that those who objected to the limits to growth hypothesis were wrong on a number of points.¹⁶ At a recent conference to mark the fortieth anniversary¹⁷ of the publication of *The Limits to Growth*, various scientists presented work showing that the projections of our future sustainability made in the 1970s are still largely correct today. If we are to avoid this outcome then we must seek a radical change in the way that human society produces and uses goods, and manages the unwanted materials from industrial processes.

As noted earlier, energy is an important part of our ability to produce raw materials. What's also important is price, as this affects the price of commodities generally. Over the last half of the twentieth century the price of energy and mineral resources was at an all time low, but in the first decade of the twenty-first century, both energy and mineral prices have been constantly rising. This, it is argued, is a symptom of the ecological limits predicted in *The Limits to Growth*. For example, oil production data shows that, since 2005, global oil production

14. Wikipedia, “The Limits to Growth”. en.wikipedia.org/wiki/The_Limits_to_Growth

15. Graham Turner (June 2008). A Comparison of the Limits to Growth with Thirty Years of Reality, Commonwealth Scientific and Industrial Research Organisation (CSIRO). www.csiro.au/files/files/plje.pdf

16. Ugo Bardi (2011). *The Limits to Growth Revisited*. Springer, ISBN 9781-4419-9415-8.

17. Smithsonian Institute/MIT (March 2012). Perspectives on Limits to Growth: Challenges to Building a Sustainable Planet. si.edu/consortia/limitstogrowth2012 View the conference presentations on the YouTube playlist at www.youtube.com/watch?v=ZiG3l5DaPrE&feature=list_related&playnext=1&list=SP2817969CA87E5B47

has reached a plateau.¹⁸ This is a warning of the imminent arrival of a longer-term decline in production, and the economic difficulties that this will create – as highlighted by the French Prime Minister in April 2011, when told the National Assembly of France that global oil production had reached a peak and that this would have serious implications for the future of the global economy.¹⁹

Information technology will be an important part of how humanity adapts to the restrictions imposed by the limits to growth. That process begins when we accept that we must adapt our personal use of technology to work within these limits.

As a priority we must learn to extend the life of our electronic goods. Making goods last two or three times longer creates a proportionate reduction in the demand for the materials from which they are made, and the energy used to produce the raw materials and assemble the components. To extend the life of existing equipment, or remove the need to buy new more powerful equipment, we have to look at the size and complexity of the software and data that are used on the machine. By making the size of the software programs and the movements of data smaller and less repetitive, it is possible to provide the services that we require using less powerful electronics and it reduces the amount of energy used by equipment.

2.3. THE LIFE-CYCLE IMPACTS OF INFORMATION TECHNOLOGY

The ever-greater use of IT is taking place within a finite environmental system – which means our technology must also have finite limits.²⁰ To understand the sustainability of computers and information appliances we must look at the life-cycle of the devices themselves: from the source of raw materials they are made out of; through the production process; their use by the consumer; and finally their disposal.²¹ This will give us the information we require to redesign the products, and the systems which produce them, to shift from a linear to a cyclical system of resource use as shown in figure 2.3. This represents a challenge to many aspects of the way our industrial systems work today. For example, shifting away from the maximisation of production and short product lives and instead designing goods to have a longer life and be easily repairable.

2.3.1. Measuring what goes into the system

In order to make sense of this complex system, we need more precise information on what each product contains. This is produced by carrying out a life-cycle assessment²² (LCA) of the product. Not only is the composition of the product measured, but also the energy, pollution and waste production that results from raw materials production, manufacturing and use. The availability of this data generally would enable policy makers and companies to tackle the problematic features of industrialisation more easily. Perhaps more importantly for the users of technology, it enables the consumer to preferentially buy goods which meet their demands for higher environmental standards (see box 2.2).

The life-cycle assessment studies carried out over the last decade or so have given us the first snapshot of the energy and resource footprint of IT (see box 2.1). By studying the whole ecological footprint, it has put the impacts of the consumer in context with the impact of the industries creating these goods, and this has focussed the ecological agenda upon the manufacturers. That in turn has allowed academics and campaigners to concentrate on the process-

18. Wikipedia, Peak oil. en.wikipedia.org/wiki/Peak_oil

19. Matthieu Auzanneau, Le Monde (April 2011), *Pétroleum blog*, 8th (in French). Fillon: la production de pétrole "ne peut que décroître" (in French) petrole.blog.lemonde.fr/2011/04/08/fillon-la-production-de-petrole-%C2%AB-ne-peut-que-decroitre-%C2%BB

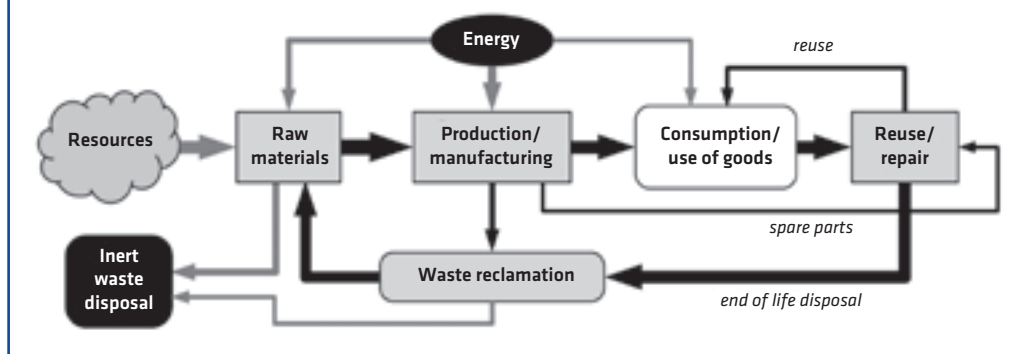
20. Wikipedia, Limits to Growth. en.wikipedia.org/wiki/Limits_to_growth

21. Leonard, Annie (2010). *The Story of Electronics*. www.storyofelectronics.org

22. Wikipedia, Life-cycle assessment. en.wikipedia.org/wiki/Life-cycle_assessment

Figure 2.3.

Ecological design and cyclical resource use



es which produce these goods in an attempt to apply greater pressure for change.²³

With the rising concern about climate change there is an increasing focus on the amount of electricity that IT consumes. The more gadgets we have, especially mobile devices that require charging, the greater the demand for electricity around the globe.²⁴ Though we might focus on the direct use of electricity by IT – because that's the part of the system we can easily measure – in terms of the overall life-cycle of these devices more energy will have been used during their production (for example, see the figures for Apple computers shown in figure 3.1). As the level of direct energy use by electrical goods has reduced, so the energy consumed during production has become more significant.²⁵ For example, the memory chip in a laptop computer can take more energy to produce than the laptop itself will consumer over its three-years-service life.²⁶ An-

other example are video display screens where, although the older glass cathode ray tube (CRT) displays consume more electricity while in use, the newer flat panel displays require as much or more energy to be expended during production.²⁷

2.3.2. The importance of the embodied energy of equipment

The latest digital electronics require some of the rarest metals on earth, and this has increased the energy demand required to produce the devices. That's partly due to the energy invested in producing rarer resources, but also because of the high purity demanded by the semiconductor industry. The laws of physics require that to make something purer through refining you have to use more and more energy to remove progressively more of the unwanted impurities. The metals used to make microchips must be extremely pure as any impurities affect the conductive qualities and speed of the chip.

For those concerned about the ecological impact of the machines they use, this embodied energy²⁸ (the energy used to make the device and all the raw materials it contains) is rarely measured or reported by equipment manufacturers. To put that into perspective with regard to other technologies, weight-for-weight, the amount of fos-

23. Leonard, Annie (2007). The Story of Stuff. www.storyofstuff.org/movies-all/story-of-stuff/

24. International Energy Agency. (2009) Gadgets and Gigawatts, OECD/IEA, Paris. Gadgets and Gigawatts Summary. www.iea.org/Textbase/npsum/Gigawatts2009SUM.pdf

25. Williams, E., Ayres, R., Heller, M. (2002). The 1.7 Kilogram Microchip: Energy and Material Use in the Production of Semiconductor Devices, Environmental Science and Technology, 36(24) p5504-5510. www.it-environment.org/publications/1.7%20kg%20microchip.pdf

26. de Decker, K. The monster footprint of digital technology, Low Tech Magazine, June 16th 2009. www.lowtechmagazine.com/2009/06/embodied-energy-of-digital-technology.html

27. Socolof, M., Overly, J., Geibig, J. (2005) Environmental life-cycle impacts of CRT and LCD desktop computer displays, Journal of Cleaner Production, 13 p.1281-1294.

28. Wikipedia, Embodied energy. en.wikipedia.org/wiki/Embodied_energy

Box 2.1.

The resource footprint of consumer electronics

Computers, mobile phones and other digital electronics are a modern treasure-trove of rare and exotic substances. For example, a quarter of the weight of a mobile phone is made of metals; about half of that is copper, and the rest is mostly aluminium, tin, iron, silver, and gold – as well as minute quantities of platinum, palladium, barium, hafnium, indium, gallium and rare earth metals. You also get metals cropping up in unexpected places, such as the 2 or 3 kilos of lead in the glass of old TVs and computer monitors, or the gold which coats many of the connectors inside our IT equipment.

The diagram on the right shows the relative composition of a computer system (with old-style glass CRT monitor) and a mobile phone. There is little detailed information on the composition of most digital appliances, although that situation is improving with the introduction of life-cycle analysis reporting.

By their nature, devices that rely on extremely pure materials, engineered at microscopic levels of detail, require far more energy to create than “old fashioned” analogue devices. Digital electronics might be more efficient or require less energy during their operational lives, but as they demand more energy during their production they are often no more efficient overall when we look at their life-cycle impacts. For example, a life-cycle study of a 2-gram 32 mega-byte DRAM memory chip estimated that 1,600 grams fossil fuels, 72 grams of chemicals, 32,000 grams water, and 700 grams of gases (mainly nitrogen) were used during its production; and the production of silicon wafers from quartz used 160 times the energy required for ordinary silicon metal. That means the laptop in which this chip would have been installed would use less energy during its working life than was required to manufacture its memory chip.

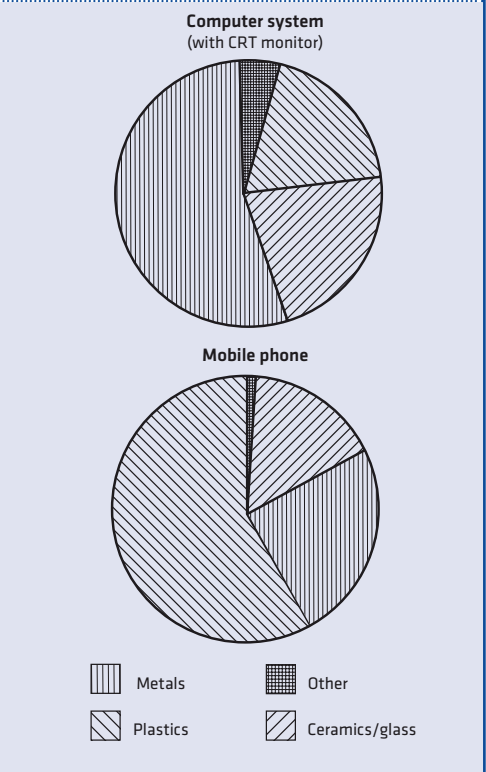
Note also that carbon dioxide from energy use is not the only significant greenhouse gas produced by semiconductor production. For example, nitrogen trifluoride is an etchant used in the production of silicon chips and LCD displays, and although released in very small quantities it is 17,000 times more potent as a contributor to the greenhouse effect than carbon dioxide.

At the heart of most digital equipment is the motherboard. This is a large printed circuit board mounted with chips and other system components and connectors which link it to other parts of the system using wires and ribbon cables. The motherboard is most easily seen inside a desktop computer.

When you look at a computer motherboard, most visible is the aluminium in the cooling fins/heat sinks on the microprocessor and other high-power heat-producing chips. The circuit board itself is clad in a thin layer of copper. The various connectors on the board are made of iron, copper and tin alloys, often with a gold layer of electroplat-

Figure 2.4.

Composition (by weight) of a typical computer system and mobile phone



ing to enhance the conductivity of the mechanical connections. The small round components are capacitors, manufactured using titanium, barium and sometimes other rarer metals. Some of the minute devices on the board are also capacitors, but their small size means they contain much higher quality and rarer materials such as niobium or tantalum minerals extracted from the ore coltan which is associated with the blood metals trade in Africa.

The board itself, and most of the connectors, are made from laminated materials or thermoplastic resins which are made from oil. Most of these components are fixed to the board with solder made from alloys containing tin, copper, silver, bismuth, indium, zinc, antimony and some other metals. Many circuit boards are also given protective lacquer coating, made from oil-based resins, to prevent moisture corroding the circuit board or its components.

sil fuels required to produce a computer chip are 600 times greater than the weight of the chip; by comparison the fossil fuels required to produce a car are only one to two times its weight, and for an aluminium can around four to five times.²⁹

Another aspect of the limited supply of these critical materials is that, as scarcity makes prices rise, the shortage of supply is an incentive to their illicit production. Unlike conventional agricultural resources, where supply can be drawn from a wide area and can shift with the global market, mineral resources can only be produced where they are found. Similar to the issue of oil and the Middle East, as pressure is put on global mineral resources, it is a handful of states who are becoming the focus of the world's trade in high-tech minerals. As a result of the problems with supply, some governments are arguing for strategic policies to protect these "critical raw materials"³⁰ to ensure the world has access to these resources in the future.

For example, an important metal in the production of miniaturised digital electronics is tantalum. Half of the world's tantalum supply is mined in Australia, and it is produced as a by-product of other metal-mining operations in many states, but between 1% and 10% may be mined illegally in central Africa. This trade in turn supplies the finance that perpetuates the armed conflict³¹ in these areas, and the human rights abuses that are the result.³² While it may soon be possible

to buy equipment which is accredited as "blood metal" free, the basic laws of economics mean that everyone is benefiting indirectly as a result of this illicit trade. By increasing supply within a tight market, it lowers the price of these resources for all. The only way to address the issue of conflict resources is to have a globally policed system which is able to accredit the trade in all resources.

The debate over green IT shows the value of life-cycle analysis, and also demonstrates the complexity of how we must manage the transition towards more sustainable industrial processes. Some of the leading hardware manufacturers are now commissioning reports on the impacts of their activities – and life-cycle analysis is a key part of providing this information. But while life-cycle assessment allows us to look at the impacts created by our demand for certain goods, for those whose work or lives have become dependent upon these technologies, it can also throw up some problematic questions on the impacts implicit in their use of technology. Knowing that the activities which you undertake on a daily basis require the expenditure of such resources puts our individual lifestyles under a critical microscope. Although the scope and standards of life-cycle analysis reports vary, if we utilise what information is available (see box 2.2) we can make better purchasing decisions, and so begin to address the impacts created by our demands for these technologies.

29. Arizona State University (undated). Life Cycle Assessment of IT hardware.
www.it-environment.org/about%20project%20-%20LCA%20of%20IT%20hardware.html

30. European Commission (2010) Critical raw materials for the EU, CEC.
ec.europa.eu/enterprise/policies/raw-materials/files/docs/report_en.pdf

31. Global Witness (2009) Faced with a Gun, What Can You Do? www.globalwitness.org/media_library_get.php/980/1277197135/report_en_final.pdf

32. Sourt, C. The Congo's Blood Metals, The Guardian, Friday 26th December 2008 www.guardian.co.uk/commentis-free/2008/dec/25/congo-coltan

2.4. IS INFORMATION TECHNOLOGY SUSTAINABLE?

The answer to that question is both “yes” and “no”. Under its current design, the way we build and utilise information systems is completely unsustainable for many different reasons. In part that's linked to the unsustainability of the human system in general, as described in *The Limits to Growth*. However, that's not to say that we couldn't address many of the present problems with IT to keep these technologies functioning in the future. The difficulty is that doing so will require the structure of the IT industry, and the products it designs and markets, to change to a wholly different model of working – and, without consumer and legislative encouragement, that is likely to result in both a cultural and economic resistance from the industry as it exists today.

The bottom line of sustainability is the ability of a system to keep functioning – and on that simple analysis information technology has a number of serious problems which need to be addressed. Some relate to the supply of minerals, while others, such as energy consumption, are implicit in the nature of the materials involved in the process. There are ways to address many of these issues. However, they're not “business as usual”, which is why it requires a major institutional and ethical change within the information and communications technology (ICT) industry. As consumers of these goods we have a role to play in this process; partly by lobbying for better reporting and environmental standards, but also by changing our own practices to minimise the impacts of the work we carry out using these technologies.

To make the diminishing stock of raw materials last longer we need to extend the life of all electrical goods. At present digital electronics is only achieving a fraction of the lifetime that could be achieved if they were designed for

a longer life. The difficulty for the electronics industry is that longer life will lead to lower turnover, and that in turn means that the nations who have specialised in the mass production of electrical goods will grow more slowly. Another great step forward would be designing devices in ways that maximise recycling and reuse, and to remove as much of the toxic content of electrical goods as possible, so that the end-of-life reclamation of IT equipment does not create intractable toxic waste residues.

While making gadgets last longer has an impact on manufacturers, perhaps the greatest impact will be upon the software community. They too focus on short product lifetimes, planned obsolescence and restricting backwards compatibility to ensure that users must upgrade. However, this “culture of obsolescence” is predominantly the preserve of the proprietary software industry. The most sustainable life-cycle for IT involves the use of open standards and open intellectual property. This enables a progression of development rather than continual reinvention, and is far more likely to lead to extended lifetimes because the pressures to continually upgrade are not so great. For this reason the free and open-source software movement, and fledgling open-source hardware movement, offer a greater potential to develop a more sustainable IT industry in the future – and we need to work to support them, both as users and if possible developers.

In the end, this is a design issue; it is a matter of how we choose to build human systems. If we respect the physical boundaries to the natural world and work within these limits then we can make a truly sustainable culture. The difficulty is that in recognising these limits globally we must begin the process by first applying these limits to ourselves.

Box 2.2.

Sustainability check-list

The initial concerns about the use of IT were related to energy consumption, and were initially addressed through the labelling of more efficient goods – such as the US Environmental Protection Agency's (EPA) Energy Star logo. As the ecological issues related to IT have become more numerous, there are various standards which have been developed to accredit goods which are made to a higher environmental standard:

USEPA/DoE Energy Star Program

The longest-running energy-labelling scheme for both commercial and domestic appliances (including PC hardware), the scheme provides both audits/labels for products as well as accrediting and giving awards for excellence by organisations. It's significant amongst schemes because it puts as much emphasis on the domestic use of hardware as it does on large corporate installations. As well as product labels, their website provides a series of downloadable guides on different aspects of energy efficiency and reducing energy consumption. For more information see the website: www.energystar.gov

Electronic Product Environmental Assessment Tool (EPEAT)

This is an accreditation tool for electrical equipment managed by the Green Electronics Council. It maintains an online database of products that have been accredited using various environmental criteria, awarding each either a gold, silver or bronze classification. While there has been criticism of this scheme because it doesn't evaluate more radical measures – such as the elimination of common PVC (polyvinyl chloride) and flame-retardant plastics – it is currently the main labelling scheme

used by most IT equipment manufacturers. To access the products database and learn more about the scheme see the website: www.epeat.net

Restriction of Hazardous Substances Directive (RoHS)

This is a European law introduced in 2006 which seeks to restrict the use of certain metal and flame-retardant compounds in all consumer goods. Goods produced or marketed in the European Union (EU) have to meet these standards. While there is no specific logo for goods, those which are compliant have “RoHS” printed on the packaging or on the body of the product. The RoHS legislation has recently been merged with the EU-wide controls under the Waste Electrical and Electronic Equipment Directive (WEEE) – this is discussed in detail in unit 9. By being WEEE compliant you not only tackle some of the toxics issues, but you help to simplify the way that the waste industry collects and processes the equipment at the end of its life.

Greenpeace Guide to Greener Electronics

This is the most wide-ranging rating scheme for electronics, based on energy use, carbon emissions, the use of recycled materials and the management of chemical hazards. As well as the general report, Greenpeace produces a “report card” for each major hardware manufacturer detailing their performance generally, and which products meet Greenpeace's reporting criteria. For more information see: www.greenpeace.org/rankingguide

In addition to the above sources, you can often find more in the “environmental reporting” or “corporate social responsibility” of most IT equipment manufacturers' websites.

SECTION 3

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, embedded¹ 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 micro-processors 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 modding⁵ 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 limitations⁶ 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 reports⁹ 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 firmware¹¹ (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 film¹² which highlighted how the batteries in the early models failed after 18 months. As a result of a

Figure 3.1.

Carbon emissions from Apple products



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

3.2.2. What is open hardware?

Today a new open-hardware¹⁶ 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. Dannonitser, 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 Manifesto¹⁸ 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 systems¹⁹ – 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 expected 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 a 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 unit²⁰ (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 es-

sential 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” look 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:

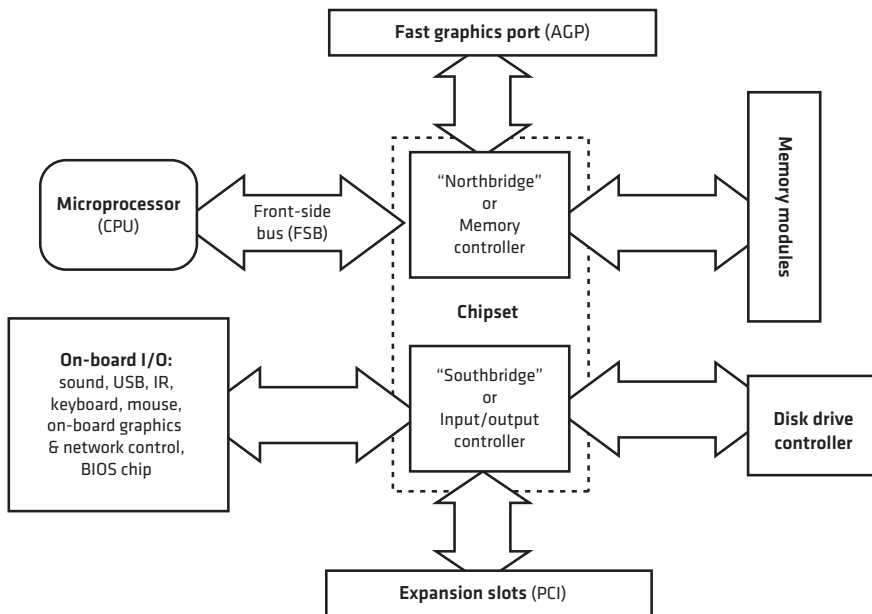
20.Wikipedia, Central processing unit. en.wikipedia.org/wiki/Central_processing_unit

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 microprocessor²³ 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



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 metal 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 storage²⁵ or random-access memory²⁶ (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

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

the different types. Look under the “speeds” subsection for a list of the different classifications of DIMM module.

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 ASCII²⁸ 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 Unicode²⁹, 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 port³⁰ (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 drives³¹ have been the main method of secondary storage³² 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 than 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 Attachment³⁵ (E-IDE or parallel ATA); and Serial Advanced Technology Attachment³⁶ (serial ATA or SATA). Any motherboard more than four or five-years old will use IDE/E-IDE {7a}; a motherboard less than 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 Interface³⁷ (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 flash³⁸ 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



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 attach 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

Figure 3.6.

A PC card



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.

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 (al-

though 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 together, 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 factor⁴¹ (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 it's 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

Figure 3.7.

A power monitor



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.

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

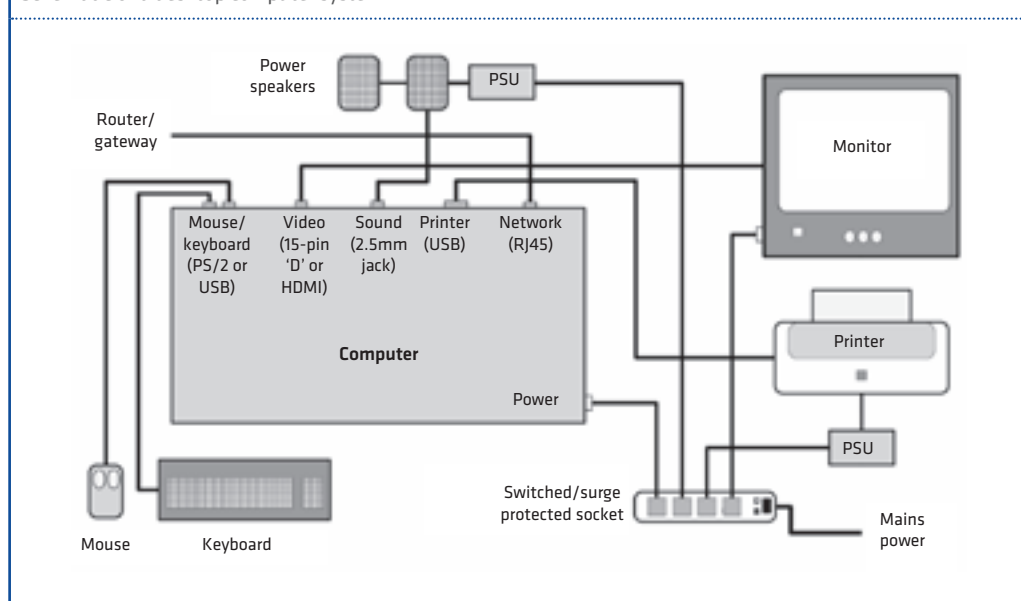
some simple rules which can help protect and preserve the equipment. Figure 3.8 shows a very simple desktop system as an example.

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 lightning 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 if 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 is 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



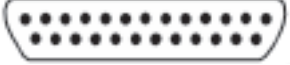



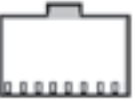







Table 3.1 lists the bare minimum system specification quoted to run various operating systems. Ideally you should double or triple these figures. If you are buying second-hand equipment, and what you're buying is slightly under-spec, you could perform a motherboard/processor upgrade to boost the power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

OPERATING SYSTEMS: MAKING THE SYSTEM WORK

A good computer system isn't simply bought or built – it's designed. Even when you buy a system commercially, the process of defining the correct hardware specification and operating system components still applies. The process begins with the construction of computer hardware; then installing the operating system; and finally you decide what application programs to install, how to protect data on the machine, secure access to the system, and devise back-up systems for data to look after your information.

To a great extent, the specification of the hardware affects which operating system(s) can be installed on the machine. Even so, operating systems¹ – the collection of programs which run the hardware and provide programs and functions for users and networks – are themselves very versatile, and define what the system will provide for the user. A single machine can have many different functions: These can be very general, such as in a work station for office or creative work; or they can be highly specialised, such as a file server which shares information across a local network. What determines the role of a system is the type of operating software used (desktop, server or a thin-client terminal) and how it is installed and configured.

Traditionally “green IT” has concentrated on power consumption and carbon emissions, and has largely been centred on the needs of large data centre operators. In this kind of environment saving a few milliseconds here or there, by changing the way the operating system is configured, can make a big difference to the efficiency overall when multiplied by the millions of transactions a day that flow through a single large computer installation. However, equally significant savings in time and energy can be made for small office and individual computer users – and multiplied by the millions of computer users

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

around the globe that potentially adds up to a very large amount of time and resources too.

If we include the human resource side of the system, not just the hardware, the use of operating systems, and configuring the features of the system to protect data, is a critical “resource” issue – as significant as energy consumption. If data is lost it has to be recreated; if files are corrupted, and there is no effective system of backing-up, the work has to be done again. Secure and resilient system design ensures that data loss is avoided and security is maintained, and in this way we avoid the need to expend more time and energy recreating lost or damaged information.

In this section we're not going to examine the mechanics of how an individual operating system is installed. There are more than enough books and guides already available on this subject, as well as the operating systems' own installation documentation. Instead we're going to look at how operating systems are configured to provide certain functions or services, what useful functions we can make these systems perform, and how we can maximise the benefit of IT equipment – including older equipment – to get the most benefit from it and minimise the ecological impacts. Even if you do not install the system on your machine, by understanding the general principles of how they work you can improve your everyday understanding of how systems operate.

4.1. THE DEBATE OVER FREE VERSUS PROPRIETARY SOFTWARE

When evaluating the potential uses of any computer hardware, we first need to decide which operating system to use. For some this decision has already been made on their behalf as the new systems people buy are usually pre-installed. For PC-type machines bought from most commercial vendors that system will be a Microsoft Windows system. Macintosh computers have their own system that's specific to the hardware, but for PC-type machines there are also alternative free and open source software² (FOSS) operating systems which will efficiently meet the everyday needs of computer users. To understand the difference between these two models of computing – the proprietary software³ world and the world of FOSS – you need to understand more about the objectives of each model.

4.1.1. Intellectual property rights

Intellectual property rights⁴ are the line which demarcates the clash between the desire for freedom of use and access to data, and the need to preserve economic power in the new information age. The options available when installing operating system software are wholly entwined with intellectual property issues – and the choices we make will in turn affect how users are affected by the growing restrictions on intellectual property. Intellectual property rights are the issue which will, more than any other political or technical obstruction, define the economic and legal pressures governing how information systems and the internet will evolve in the near future.

Unlike the technology humans have developed over history, the technology of the information society⁵ is not benign. Devices can be made to be semi-intelligent, able to control the terms of how we use the operating system and/or its stored

information after we have bought them. In this way society is becoming increasingly constrained by the technological standards that manufacturers develop rather than the political decisions of policy-makers. And if, with the development of more intelligent systems, manufacturers choose to lock-up the functions of their products with digital rights management systems,⁶ then society's use of these new tools will be restricted; not to improve our experience of the technology or to protect fundamental civic rights, but primarily to protect the economic interests of their manufacturers.

These new systems, backed up by the protection of intellectual property law, create other restrictions which generate a number of hidden costs to the computer user. Once a user commits to a certain system standard – for example, the choice of using either a Macintosh, Windows or Linux-based PC – the barriers that this creates can make it difficult to change after that decision is made. Primarily there is a cost penalty, either for changing the machine technology, buying new operating software and licensing it, or converting the media that you used with the old technological standard to the new. Technological standards create a means whereby system designers have control over the market for certain goods and services, locking people into using their products in the future. As we move towards a society where information systems are an inescapable part of everyone's lives, we have to pose a very important question about the information systems we create: How do we square the development of a new networked society, where information is intended to flow more freely, with an intellectual property system whose primary purpose is to restrict information flows in order to extract greater value from them?

To examine this idea in more detail let's consider something that we do every day – breakfast!

When we eat a modern, consumer-friendly, pre-prepared breakfast meal we eat brands – *we chew on other people's intellectual property*. At the same time, the acceptance of the brand means that you leave it to others to decide the taste or texture of your breakfast. It's difficult

2. Wikipedia, Free and open source software. en.wikipedia.org/wiki/Free_and_open_source_software

3. Wikipedia, Proprietary software. en.wikipedia.org/wiki/Proprietary_software

4. Wikipedia, Intellectual property. en.wikipedia.org/wiki/Intellectual_property

5. Wikipedia, Information society. en.wikipedia.org/wiki/Information_society

6. Wikipedia, Digital rights management. en.wikipedia.org/wiki/Digital_rights_management

to customise the brand. To make a branded commodity meet the needs of the individual user invalidates the generic standard that the purchase of a national or global brand implies. Instead we are required to adhere to the standards the brand conveys, and conform to the conditions that this entails.

The process of combining one product with another to create a distinctive brand identity is called product bundling.⁷ For example, Microsoft select what they believe you need to have on a computer or home entertainment console, and in the process confer a certain amount of economic power not just on their own products, but on the product or hardware providers that Microsoft retains an economic relationship with. Likewise Sky's satellite TV system is based upon bundling different pay-per-view channels; and increasingly telecommunications companies offer bundles containing land-line, broadband internet and mobile services, as well as the ability to download TV programmes or films.

4.1.2. Free software – the freedom to “unbundle”

The problem with bundling is that it removes choice and the ability for people to experiment and extend their knowledge without having to undertake costly and sometimes difficult changes to their systems. Using the breakfast corn-flake analogy, if there was a separate sachet of chocolate powder it would be possible to choose the strength of the chocolate flavour on the corn flake. I can decide to add a tiny sprinkle, or I could pour on half the sachet over a single bowl. Free software enables this choice to be made. In contrast bundling restricts your ability to mix 'n match a computer system that suits your own needs. Often you get the lowest common denominator of computer system which may be too much for your budget, or too little for your needs.

It gets more interesting when we start to selectively buy the elements that make up our breakfast. If I can unbundle the brand I can choose who sells me my chocolate coating; if I don't like person X's chocolate coating, I can buy person Y's (this was the basis of the anti-trust lawsuits against Microsoft – specifically the way Microsoft incorporates software such as web browsers into

their operating systems).⁸ I could even produce my own coating if no one produced a flavour I liked. Also, if I find that the quality of the corn flake is not consistent enough, I can experiment and switch to another manufacturer's flake without sacrificing the use of the particular chocolate flavour I had grown to love. When we unbundle products we allow a greater number of options as to how the collection of products can be configured by the user. It is this kind of “real choice” that branding or bundling stifles; by managing the menu of consumer selections those involved dictate the market.

The fundamental difference between the proprietary versus the free/open model of computing can also be demonstrated using the corn flake analogy. The FOSS movement would stress the ability to choose the flavour of the flake and the quality of the chocolate coating. For example, the GNU/Linux operating system has many different “flavours”, called distributions⁹ (box 4.1). They are all roughly compatible, and run the same standard set of programs, even though they are made by different organisations. You can even combine elements of the different flavours to make your own customised version of Linux – and a number of Linux developers/distribution producers provide information and extra programs for free to help you do this.

To sum up the intellectual property issue and how it relates to software, what defines the difference between the “proprietary” and “open” models of the Information Society is: The difference between what constitutes “selection” versus “choice”; how this preserves our creativity and freedom of expression and communication; and which systems/software applications allow this without expensive lock-ins to one particular company or service provider. Distinguishing the full range of options you have to solve a particular design, technological or communications problem, and managing the legal or other implications of pursuing those solutions, is the practical method by which you enact this freedom of choice. There is of course no easy way to acquire these skills; it's something you have to work at by learning more about the systems and technologies involved.

7. Wikipedia, Product bundling. en.wikipedia.org/wiki/Product_bundling

8. Wikipedia, European Union Microsoft competition case. en.wikipedia.org/wiki/European_Union_Microsoft_competition_case

9. Wikipedia, Linux distribution. en.wikipedia.org/wiki/Linux_distribution

Box 4.1.

Free software and the GNU/Linux system

The free and open source software (FOSS) movement began in the 1980s as a reaction against the increasingly closed world of computer software. A large part of the computer industry grew out of academia, where the idea that knowledge should be open is generally cherished; being “open” for this community was not a novel concept. Sharing prevents the duplication of effort, and ensures that potential solutions can be peer reviewed and implemented far faster than if small groups are working in isolation. GNU/Linux, the first viable “free” computer operating system to be developed, is a perfect example of the strength of this model of working.

The Linux system kernel was developed by Linus Torvalds during the early 1990s while he was a computer science student in Finland. The internet allowed him to share his ideas with software writers around the world, which led him into organising a collective effort to develop an alternative, Unix-like operating system kernel for the PC. Rather than a wholly proprietary system like Windows, or a partially open system such as the Macintosh operating system, Linux has been developed under a free licence – allowing the code to be shared and modified by users.

The idea of “free” software was developed by Richard Stallman, founder of the GNU Project. The object of GNU was to support the development of free computer software by both professional and amateur programmers. These programs were made freely available for others to use or adapt under the conditions of the GNU General Public License or GPL. The Linux operating system is the product of these two groups; so it's often called GNU/Linux to reflect the contribution of both Linus Torvalds, in the development of the kernel, and the GNU Project, for the development of the licensing system and the other programs with which the Linux kernel was originally developed.

The success of GNU/Linux has been driven by its licence. It enables programs to evolve in a way that serves the needs of computer users rather than the ambitions of the developers – which gives it a significant advantage over “closed” systems such as Windows where adaptation or extension is usually prohibited. The fact that systems/programs can be downloaded for free, or bought at a low cost, has implications for its users too: You can install it on any number of computers without having to pay for extra licences; you can copy the system for friends or colleagues to use; if required, you can take

the system apart and modify it to suit your own needs, and pass on those modifications to others. If these same actions were carried out with the programs developed by “closed” software developers, those involved would be liable for prosecution!

Linux is the kernel of the GNU/Linux system – the core program which runs the computer hardware and makes the memory, disks and other components available for programs to use. Within the whole GNU/Linux “system” there are hundreds of other programs required to provide all its functions, and the user-friendly applications that provide for people's everyday needs. Early on in the development of GNU/Linux, enthusiasts developed FTP sites containing all the files to install a Linux system. Then shortly after that, as the internet was slow when moving large quantities of data, they developed CDs with helpful installation programs which automated the process of installing and configuring the system for less experienced users. Eventually these efforts created commercial or semi-commercial organisations retailing their own version of a Linux system. These different versions of GNU/Linux became known as distributions, or “distros”.

There are a whole range of Linux and Linux-related systems available today. Some are primarily desktop systems, providing an equivalent range of creative and office-based software that is commonly available with proprietary systems; others focus on data servers and more technical computer science related applications; while some have evolved into single-purpose systems such as network routers, firewalls, and even digital video recorders or media studios. Other than the versions meant to be conventionally installed on computers, systems are also configured as “live media” that boot direct from a CD/DVD or memory stick into a working desktop system.

The issue when comparing Linux distributions isn't necessarily “which is the best”. Usually a distribution represents a certain design philosophy that is optimised for certain applications rather than attempting to be “the best distro”. There are many different Linux-based systems to suit different needs, and you need to find the distro that meets your expectations, which is quite simple as you can freely download most Linux distros to try out. For the latest news about Linux distributions available go to the DistroWatch site: distrowatch.com

In the final analysis, in choosing which operating system to use on the equipment you have, the most important indicator must be that it satisfactorily “does what you need to do”. Whether that is achieved using proprietary or free software must be part of the decision process that you undertake when you design the specification of the computer system.

The important issue is that you are aware that you have a “choice” between different operating systems, that you do not have to accept systems as they are installed or provided, and that those choices result in very different outcomes in terms of costs, intellectual property restrictions and compatibility/interoperability.¹⁰

4.2. OPERATING SYSTEMS AND OBJECTIVES

Operating systems are installed to fulfil a purpose; although the system might be the same, it is how the components of the system are selected and configured which give the system the ability to carry out the tasks we require of it. When it comes to defining the roles of different machines, and selecting the software to fulfil those roles, we see another distinction between proprietary and open source systems:

- With proprietary software there is a clear distinction between roles. For example, there is one system for Windows desktop use, and another system for Windows server machines. You purchase software to fulfil the role you require of it.
- With free software, and to some extent the Macintosh OS (because it is based on a Unix-like system), the line is blurred. It is not possible to enable a strict divide between what a desktop is and what a server is. In practice, as free software systems allow the installation of additional software more easily, it is very easy to set-up a desktop machine which can also act as a server on a local network (for example, running a web server to provide an intranet).

4.2.1. Finding the best fit of machine to operating system

Unlike the world of proprietary software, where you have a limited choice of hardware to fit the power demands of the latest Windows or Apple “iOS” operating system, with free software you have the option to design a system to fit the application you need it for. Older hardware, which

would be difficult or impossible to use with recent versions of Windows, can be found new purposes/ have its working life extended using free software. Even fairly new hardware can be configured in different ways to perform different functions, in the process creating different levels of performance or security to match the needs of the user.

For example, if you install a web server on a desktop machine, then it is possible to carry out complex web design, and then test the designs – including scripting and other functions – without a network connection to a dedicated server. Additionally, when using proprietary systems you're penalised for using more than one machine because of the need to license each installation. With free software there's no such financial penalty – *additional installations are free*. In this way we can design machines inexpensively to serve specific purposes, and undertake that task as securely and reliably as possible.

Using more than one machine, and through this consuming more electrical power, might not appear to be very “green”. This ignores the embodied energy of the hardware. If you can make equipment last for longer – as such a large part of the overall impact is tied up in manufacturing the hardware – you reduce the impacts overall. Using an old machine can avoid the need to buy new/more powerful equipment, which saves money as well as energy and pollution. For example, an inexpensive way to back-up a workstation computer, without the time and expense of making lots of data discs or buying a new external hard drive, is to take an old

¹⁰Wikipedia, Interoperability. en.wikipedia.org/wiki/Interoperability

machine and install it as a data server to store copies of your important data. This doesn't use a large amount of energy either, because older equipment often consumes less power, and in any case this type of system need only be turned on for an hour or two while the back-up is taking place.

The first decision you need to make is: What can I do with this equipment? This is a complex decision, and has a direct relationship between what functions you need from the machine and what its internal electronics are capable of supporting (see box 4.2). Breaking down our IT needs into a series of clear roles and specifications requires thought and planning to define what we really want. In the corporate environment the aim is often to provide new infrastructure, usually under contract, because it represents the most economically efficient solution. In contrast, in smaller organisations and the voluntary sector there is the option of adapting existing hardware to work more flexibly, using the skills of those who will use the system, because that represents the least-cost route to finding a solution to their needs.

As outlined in unit two, one of the important objectives of sustainable IT is to maximise the life of the equipment. There are two ways of looking at this:

- Firstly, we can maximise the use of a single machine so that it is using as much of its processing capacity as possible for as much time as possible. This is the objective for dedicated server machines and large data centre rack servers – often because there's a cost penalty to having unused machine capacity sitting idle. Achieving this outcome requires attention to the minute details of system configuration and the organisation of hardware.
- Secondly, we can prolong the life of equipment – usually older equipment – by giving it less intensive processing tasks which match the capabilities of the hardware. For small offices, older machines can take on roles which support or extend the functions carried out by the machines people routinely use – thereby extending the functions available to the user, or freeing up capacity on their main machine to allow them to work more easily. Finding opportunities to do this requires that we

understand more about the different roles machines are able to perform, and then match the available hardware to those roles.

Too often the “creative” IT world focuses on the use of a single, often new or powerful desktop machine when, if we were to think more openly about using a local network, it is possible create some very helpful solutions using older equipment. For example, routinely processing large video files or mastering video DVDs can tie up a system for long periods of time, affecting the completion of other work on that system. But if a second machine were set up with a remote desktop¹¹ accessed over a network you could undertake these same operation without restricting the capacity of the main workstation machine. This also does away with the need for the second machine to have a keyboard or monitor. Likewise, given that laptops use a lot less power than desktop machines, an old laptop is a much cheaper option to run 24-hours a day for file sharing, or to configure as a server running a small local network, rather than leaving a more power-hungry machine running 24/7.

4.2.2. Hardware specifications and practical applications

Let's begin with hardware and the demands of different operating systems. In table 3.1 we looked at the hardware specifications required by different operating systems. In table 4.1 we look at this from a different perspective – what applications can be achieved with a certain level of hardware. Remember that the minimum requirements of different operating systems, given in table 3.1, still apply. For example, to run a lightweight server on old hardware you would use a lightweight Linux distribution or Windows NT, rather than the latest fully functional Linux or Microsoft server systems.

Table 4.1 lists a number of different applications for different classes of hardware. Here's an explanation of what each one encompasses:

Desktops – office-related, graphical and communications tasks

- *High power desktop* – high-end desktop work such as graphic design, video editing and DVD mastering (less powerful machines can

¹¹ Wikipedia, Remote desktop software. en.wikipedia.org/wiki/Remote_desktop_software

Box 4.2.

Hardware specs and potential system applications

Matching the performance of hardware to potential applications is a rather complex subject. The more powerful the hardware, the simpler it is to install an operating system. The skill lies in being able to determine what is the lowest specification you require for a certain task so that you can maximise the lifetime of equipment and keep costs down. Installing systems on older equipment with restricted memory and disk space can be more of a challenge, especially if it requires the use of a text/command line interface (which requires more details knowledge of the operating system) rather than a graphical interface (where the options are usually presented as simple menus).

If you only want a desktop machine, you need a computer with a desktop-level specification. Even then, there's

a considerable difference between simple typing/word processing applications and the type of power required to edit large colour images or process video files. When considering the more creative use of computers, such as backing-up or networking support, you can use a far wider range of equipment to solve a particular problem.

Table 4.1 lists various grades of hardware and the types of application that it can perform. Note that the demarcation between what is "possible" and what is "practical" isn't as clear cut as shown in the table – what matters is that the machine can run the tasks that you design it to perform, at the speed/reliability that you require of the system.

Table 4.1.												
The use and reuse of computer hardware for different applications												
Application	Desktop			Data processing		Support systems			Downcycle		Operating system	
System specification	High power desktop	Medium desktop	Lightweight desktop	Parallel processing	Power server	Download box	Lightweight server	Network support	Educational recycling	Scrap	GNU/Linux	Windows
Recent machine dual/quad core, +2GHz CPU, 2-4GB RAM	✖				✖						H	7/08
Good machine dual core, 1-2GHz CPU, 1-2GB RAM	✖	✖		✖	✖	✖	✖				H/M	V/03
Fairly good single core, ~1GHz CPU, 512MB-1GB RAM		✖		✖	✖	✖	✖	✖			M	V/2K
Old machine 800MHz-1GHz CPU, 256-512MB RAM			✖	✖		✖	✖	✖	✖		M/L	X/2K
Older machine 500-800MHz CPU, 128-256MB RAM			✖					✖	✖		L	X/NT
Very old machine less than 500MHz CPU, 128MB RAM								✖	✖	✖	L	98/NT
Challenge! less than Pentium-3/ 500MHz and 128MB RAM									✖	✖	L	98/95
GNU/Linux system specs:	H, heavy (SuSE, RedHat); M, medium (Debian, Fedora, Mint, Ubuntu); L, lightweight (Damn Small, Puppy, Slackware).											
Windows system specs:	7, Win. 7; V, Vista; X, Win. XP; 2K, Win. 2000; 08, Win. Server 2008; 03, Win. Server 2003; 98, Win. ME/98; NT, Win. NT; 95, Win. 95.											
Note that Macintosh hardware hasn't been considered in the above list. Different versions of MacOS are designed for different generations of Mac hardware, and so the decision about which machine requires which version of MacOS is determined largely by the available hardware.												

do this, but will take much longer to do so). For these installations the latest versions of an operating system are used in order to have access to the latest software and hardware.

- *Medium desktop* – general purpose machine, capable of most office-type tasks, communications and some graphic design/editing. It could run a Linux-based web server for undertaking web design work. The latest operating systems would run slowly on these machines. While using the latest system is not impossible, it may be more productive to use slightly older software if that provided the required applications and hardware drivers.
- *Lightweight desktop* – for internet communications and simple word processing, or it could act as a “thin client” connected to a more powerful machine over a network. A problem when using Windows systems is that hardware drivers may no longer be available, especially for printers. Anti-virus software will also not be available and there may also be problems with network integration. This makes free software preferable as these systems have better hardware support, use open-network standards, and malware isn't such an issue.

Data processing – heavy-duty data processing, storage and network control

- *Parallel processing* – use as a second machine for processing tasks that interfere with your everyday work. For example, video file conversion, DVD mastering, creating DVDs, and any other tasks that runs for many hours or use a lot of power. You could use a screen and keyboard, but it's easier to access the box from your main machine using a remote desktop (for graphical applications) or secure shell (for command line control). As well as single machines, many small machines can be connected together to form a more powerful “cluster”¹² for specialist applications (e.g. video animation, scientific analysis).
- *Power server* – a dedicated server machine for a small office, or to serve content to a busy network, running a range of protocols, databases, or even thin clients¹³ (using older ma-

chines as the clients, and the powerful machine as the terminal server).

Support systems – machines to run networks and store data

- *Download box* – a dedicated machine to run 24/7 internet services such as file sharing or remote data/file access (for non-professional/small-scale applications this is most efficiently done using an old laptop due to its inherently low power consumption). Where peer-to-peer file sharing is used for information distribution the machine could be installed with a standard desktop operating system. For low volume web or FTP access, or to run encrypted tunnels for remote access to the network, a server system is a better option.
- *Lightweight server* – a simple server machine to support services across a network, rather than a single user. Like the “power server” class above, powerful machines could be used as a terminal server service for a thin-client network, but many older machines might struggle to do that for more than one or two client terminals. Instead older equipment could be configured as a file server/encrypted storage server for backing up desktop machines (using FTP or a shared/networked file system), or a simple web server to run an intranet for a small office, or a mail server to forward incoming mail for users on the local network.
- *Network support* – if good quality, it could perform some lightweight server functions, although it may be difficult creating enough hard drive storage to make it worthwhile. If using a system which required minimal hardware it could be used to run the programs required to operate a local network, and perhaps act as a firewall/gateway for a broadband internet connection. Note that as well as installing a full operating system, there are a number of stand-alone installations¹⁴ which will provide dedicated secure network services, Wi-Fi access points and router/firewalls. They boot from CD/DVD, install largely automatically, and can be configured from menus rather than requiring the more expert skills

12. Wikipedia, Cluster (computing). [en.wikipedia.org/wiki/Cluster_\(computing\)](http://en.wikipedia.org/wiki/Cluster_(computing))

13. Wikipedia, Thin client. en.wikipedia.org/wiki/Thin_client

14. For example: Smoothwall www.smoothwall.net; Engard Linux www.engardlinux.org; or Zeroshell www.zeroshell.net

required to configure an operating system to perform these functions.

Downcycling¹⁵ – productive end-of-life uses for old machines

- *Educational reuse* – if the machine can still run a lightweight Windows or Linux system then it can still perform a useful function as a training platform to learn/practice hardware and operating system installation skills. By dismantling, re-assembling and then installing the machine people get the skills/confidence required to do the same types of operation on more expensive equipment, but without making expensive mistakes if they do something wrong. Any quality of machine can perform this purpose, although because of the risk of damage to the machine it should be considered “expendable”.
- *Scrap* – “scrap” doesn't necessarily mean throw the whole thing away. Old boxes can be scavenged for their screws, cables/connectors, fans, heat sinks and other reusable components – to be kept as spares for the repair of other machines. Hard drives can be used with drive caddies to back-up large amounts of data from other machines. Unless they have no further use or have a poor picture quality, one or two old monitors can be kept for use in emergencies if another monitor fails. The only thing that must go for electronic waste disposal and recycling, as they have little value for reuse, are the metal chassis of the desktop box, plastic enclosures and motherboard/electronic circuits.

Looking at the list of potential uses for equipment above, it might seem a rather complex approach to design different machines to perform different jobs. In practice, to perform a task well with a computer, the machine should be designed to serve its purpose as well and securely as possible. For example, a general purpose desktop machine could undertake all the roles listed in table 4.1, but to do so simultaneously would drain some of the CPU's power if it all took place on the same machine (every extra program running on the system takes a little bit of power away from the important job you're doing at the time). It's more efficient to design boxes for specific tasks, optimising the hardware and software to provide both functionality and an efficient level of power consumption for the job concerned (for example, server machines do not need a monitor).

Making a redundant computer into something useful doesn't require an expensive upgrade of its components (as outlined in unit 3) – as is usually the case when installing a fully functional desktop system. Again, it's a matter of design, and matching the minimum specification of the job to the hardware available. Of course there comes a point when a machine is beyond practical use, usually when it can only run a lightweight¹⁶/text-only operating system, or you can't install enough hard disk storage or memory to provide a useful level of service. Even then, old machines can be useful for learning how to maintain or rebuild computers, or installing operating systems. These matters are examined further in unit 10, *Reuse and reclamation*, and unit 9, *End of life and disposal*.

15. Wikipedia, Downcycling. en.wikipedia.org/wiki/Downcycling

16. Wikipedia, Lightweight Linux distribution. en.wikipedia.org/wiki/Lightweight_Linux_distribution

Box 4.3.

System documentation

Even where one user looks after their own system, remembering all the relevant information about the system is very difficult. Where a group of people look after a number of computer systems, it's wholly impractical. It's a good time-saving and security practice to document each system installation. Begin with details such as the model of the computer and any relevant serial numbers (this is the sort of data insurers, the police, or repair agencies may require). Then there's the hardware profile itself (e.g. knowing what type of memory modules used in case you need to replace them). Finally there's the system-specific information – such as the type of system installed, hard drive configuration and critical passwords. If you buy a system pre-installed, then you should be able to pull together much of this information

from the vendor's hardware specifications and the data provided by BIOS or the operating system.

By documenting this information and storing it securely it simplifies the process of system maintenance, especially where there are a number of machines administered by different people. If you store this data as a simple text file (*not on the machine it is related to* – in the event of a fault you will need this data!) it can be used as a working document; adding information on system maintenance, any faults which occur, and any remedial action which was taken. In this way it's possible to keep a log for the machine, which makes system maintenance a far simpler procedure.

The table below outlines the types of data that is relevant to keep. Whether you keep all of this depends

Date installed/technician ID	It's useful to keep a note of the date installed, and the identity of the person installing the system
Operating system/version	Include details of any service pack upgrades/Linux kernel ID
Processor type	Make, model, speed, number of cores, L1/L2 cache
System memory	Note the type, speed & no. of modules, not just the memory size
Hard drives	For each hard drive note the make/model, its capacity, and an outline of the partitions on the drive and file systems used
BIOS serial number	Usually displayed on the BIOS splash screen – useful to keep in case of theft, or if the insurer or repair contractor asks for it
BIOS key	Which key to press to enter the BIOS configuration menu
BIOS passwords	Note the BIOS set-up/administrator and user passwords, if used
Motherboard	Make and model, as well as the serial number if you can find it
Plug-in cards/peripherals	Make and model, and serial number
IP configuration/address	Note whether the interface is configured with a static IP address or DHCP, its gateway/DNS addresses, and its MAC address
Administration/Root password	<u>Critical password</u> – you should log this and keep it safe
User accounts	For each user account, document user ID and any user-specific configurations used with the account (e.g. access permissions)
Boot loader password (Linux)	A boot loader password prevents bypassing of the root password
Hard drive encryption password	<u>Critical password</u> – if you use hard drive encryption, log this
Enabled network services	Note which services are enabled on the machine

upon the technical level of control you wish to have over the system:

Finally, all this information is of course very sensitive. It's not the sort of data that you would print and pin to the wall! For that reason the file should ideally be encrypted. Keeping the information on an encrypted memory stick or hard disk is one option. Another is to

use a symmetric cipher so that it is not necessary to keep a public/private key pair with the file. In the event of someone needing to access that data they could read the file from the network and decrypt it to obtain the information they require. If changes are made it's a simple matter to decrypt, edit and re-encrypt the file, and then replace the copy kept for maintenance with the new file.

4.3. PLANNING AND DOCUMENTING INSTALLATIONS

When installing systems, planning is important – each user needs to access the resources they require to work. That might be on one machine, or it might require configuring machines across the network – for example, allowing access to file servers or shared resources such as printers. It helps to plan these details first so that you can organise how you will install any new machines, or modify the configuration of others. Where a group of people look after a small network, keeping written logs and procedures is essential so that everyone can co-ordinate their efforts and not lose vital information about changes or modifications to machines or system settings. Even when there is only a single machine in use, there is still information that is essential to its operation and maintenance – such as the critical passwords for the machine, and also the passwords to access and configure a broadband router or similar devices.

Each “system” that is installed should have its own log of critical and background data (box 4.3). A “system” includes details of not just the machine, but also any hardware that one or more machines are attached to. While not es-

sential, it's good practice. If a machine should fail, or if a critical password is forgotten by its user, keeping a log securely stored on another machine (or on a back-up memory stick) allows quick and easily access to the data you need to unlock the system – or gives information to technical support staff so they can diagnose problems more easily. Even if a machine is not in use, if you have a log of its vital statistics, it makes it easier to decide if it can be re-installed and put into service (or stripped for the required parts) if another machine fails.

In addition to the hardware, it's also important to keep a log of operational data. For example, the connection details/passwords for broadband access, access/configuration details of email and internet services if they are run by a third party, and even details of insurance or maintenance contracts covering any equipment. By having a secure, central location for this critical information you will save much time and resources at a later date should there be a problem – be that with the hardware, or an important staff member meeting with an accident.

4.4. POWER MANAGEMENT SETTINGS

Following installation, the operating system is configured to optimise its functions to meet the intended purpose(s) of the machine. When configuring a new/newly installed desktop system, after configuring user accounts, the next step is to configure power management (box 4.4). Power management is important, not just because of the power consumption issue. When you configure power management you're also setting up features such as screen savers. These are basic security measures to control access to the system. Setting a short lock-down period on a screen saver or the power management menu doesn't just save energy – it makes the system more secure by preventing casual access.

4.4.1. How power management functions
Power management offers various options. The primary one is giving the machine time limits to define what happens after it has been idle for a certain period. There are five levels of power management which determine how much power the system consumes:

- **1. Fully operational** – the machine is working normally. Note that the level of “normal” power consumption can be also modified in various ways. The hard drives can be powered down after a certain period of inactivity, and the priority given to different parts of the operating system. The user applications can also be changed to manage the load on the operating system. This is normally only implement-

ed on laptop systems, but there's no reason you cannot implement the same policies on a desktop machine. The trade-off for optimising power management is to lose some power for user applications, which often means that users playing high power-hungry video games will deliberately disable power management.

- **2. Stand-by, sleep or suspend** – this is a low-power standby mode in which the machine's operation is halted, hard drives and the video display are powered down, but power is maintained to the memory. It does not consume as much compared to running at full power, but it's significantly more than if the machine is put into hibernation or shut down. When configuring screen savers, setting the suspend period to a few minutes after the screen saver activates will save a lot of power. The problem is that those who use their computer as a high-powered media player will object because their music will stop playing. Consequently it's often less intrusive to leave the suspend option for ten minutes to half an hour after the screen saver activates. A problem with suspend is when power is interrupted, or the laptop battery runs down. On a laptop the normal procedure is to transfer to hibernation mode or to execute a soft shutdown. On a desktop, which has no battery back-up, a power cut or inadvertently turning off a machine in sleep mode can cause data loss from any open applications. For this reason, of the two options, hibernation is more secure.
- **3. Hibernate** – this halts the machine, backs-up an image of the working machine on the hard drive, and then powers down to an off state. The difference between hibernation and sleep modes is the amount of time it takes to restart. A machine in sleep mode will wake up after five or ten seconds of the wake button being pressed. Restoring full operation from hibernation can take thirty seconds or more, depending on the speed of your BIOS self-check (it is quicker than having to power-down and reboot, but some users find the wait problematic). Hibernation can sometimes cause problems if an error occurs when the hardware is powered down, or the machine is switched-off at the mains. When the machine restarts, instead of booting the operating system, it copies the memory im-

age back from the hard drive and resumes operation from the point it was told to hibernate. Errors can cause the reboot to crash, and the machine has to be reset and rebooted from scratch.

- **4. Soft shutdown** – telling the machine to “shutdown” powers down the system. This means that the system is “off” but the power is still connected at the power supply (hence “soft” shutdown as only the software has been stopped). With soft shutdown you restart the machine by pressing the power button, and then the system reboots from scratch.
- **5. Hard shutdown** – the human operator physically switches off the machine at the mains. This disconnects all power so that the machine consumes nothing.

On laptop systems, not only are the major components integrated into a single unit, they are optimised to use as little power as possible – but to do this you trade off some of the processing power available with desktop machines. For demanding jobs the power saving limitations on laptops mean that jobs can take longer, unless you deliberately turn off the power saving features on the laptop (at least while its running on AC power).

However, for most everyday office-related tasks a lot of processing power isn't required, and so enabling these features will have little effect on the speed of operation.

You can also enable similar power saving features¹⁷ to those used on laptops on most desktop systems in order to reduce their everyday power consumption:

- For *Windows XP* to *Windows 7*, the standard power management system offers three modes: *High Performance*, *Balanced*, and *Power Saver*. If you're not editing large colour images or manipulating large databases and spreadsheets, Power Saver mode will make little difference to the way the system appears to work – and you can always switch back temporarily if you need to undertake more powerful data manipulation. There are

17. Wikipedia, PC power management. en.wikipedia.org/wiki/PC_power_management

Box 4.4.

Power management and the machine hardware

Power management is an issue that developed alongside the growth in mobile equipment during the early 1990s – especially laptops, where prolonging battery life was an important consideration to the usability of the device. As the technology has evolved, it has become an important factor in managing the power consumed by desktop and server machines too.

New microprocessors, particularly the multi-core type, can be configured to run at different speeds. The amount of power the chip consumes is proportional to its operating speed and the supply voltage of the processor. Varying the speed of the CPU from the operating system performs two roles: It prevents excessive heat generation when the power of the chip is not required; and, more importantly, it reduces power consumption so that the amount of energy the machine uses can be scaled to the demands on the system. On laptops it's often used to deliberately extend the battery life of the machine at the expense of operating speed.

Depending on the age of your hardware, there are two power management systems in use:

- The *Advanced Power Management (APM)* system was developed in 1992. Support for APM began with Windows 95, but was not included with Vista (the Linux 2.6 kernel still supports APM in most distributions). Any hardware which only supports APM will not be able to function with Vista or Windows 7 unless you have a specific driver for those systems.
- The *Advanced Configuration and Power Interface (ACPI)* was developed in 1996, and has superseded APM in most PC hardware. ACPI is more versatile, and offers more options to configure power consumption by the processor, display, hard drives and peripheral devices. ACPI has been supported since Windows 98, and is supported in most Linux distributions.

With Windows systems the power management software for specific types of hardware must be loaded along with the other driver software during system installation. Power management is configured from the *Control Panel*: In *Windows XP*, you'll find the *Power Options* icon on the *Performance and Maintenance* menu; in *Vista* the icon is on the *System and Maintenance* menu; in *Windows 7* it's on the *Hardware and Sound* menu. The power options dialogue allows you to configure the power management features available with the hardware you are using.

On the Macintosh (OSX), you'll find the power management icon in the hardware section of the *System Preferences* menu.

With Linux-based systems the APM/ACPI service is usually enabled automatically – but to function you must be using hardware which conforms to the APM/ACPI standard. This is a particular problem with certain brands of laptop and desktop hardware developed for use primarily with Windows. As is usually the case when installing Linux, you need to check the compatibility of the hardware before you start. As GNU/Linux has become a more popular operating system, power management is no longer the obstacle that it was five or more years ago. To be certain that all hardware functions are available, check the information and reviews of motherboard, peripherals and laptop systems at linux-hardware.net or www.linux-drivers.org

The power management controls of the operating system interface with the APM/ACPI firmware in the machine's motherboard. The motherboard will then control both the devices connected to the motherboard, as well as sending signals to the monitor and any compatible USB devices. Depending upon the motherboard's BIOS settings, it's possible that power management has been disabled or limited. Although power management is often enabled by default, you should check the configuration of power management in BIOS, and change as required (if necessary identify the make/model of the motherboard, then search for and download the board's user manual to check power management options).

also a number of third-party¹⁸ power applications which allow you to refine these settings, and which manage the programs that run on the system to reduce the demands on the processor, memory and hard drive.

- Most Linux distributions offer similar power management features to Windows. As with the FOSS world in general, there are also specialised projects¹⁹ which develop power management systems for use across the Linux community. Many of the popular Linux distributions are now producing their own detailed system documentation on implementing power saving features/policies in server and desktop machines – and these guides are available free on the distribution website. For the more technically minded, there are also free programs which work across most distributions to audit and improve power performance. The leading application is *powertop*,²⁰ which not only monitors how the system/running programs consume energy, it also suggests ways of improving system configuration to reduce power consumption.
- MacOS also has its own built-in power management modes and there are third-party applications which can monitor and optimise the operation of the system.

4.4.2. Advanced power management options

For servers, where you don't always have a person on-hand to press the “power” button, starting the machine can be a problem – especially if there are hundreds of them. For machines being used as servers there's an option called “wake on LAN” (WOL). If a data packet is sent to a machine in suspension or hibernation containing the card's unique address number (some also require a security password), the machine will restart, provided you have a WOL-compatible network card. Later the machine can be remote-

ly put into suspension/hibernation, or it can suspend operations according to its own power management configuration, until it is needed again. Using WOL allows a single operator to have control over hundreds of server machines – often using load monitoring programs to shut them down or wake them up in order to match demand while minimising power consumption. For small system users, it can also be configured to wake-up file servers or data processing machines so that they don't use as much power when not required. The standard power management menu allows you to configure time limits to put the machine into either sleep, hibernation or soft shutdown depending upon how long it has been idle for; or, for laptops, when the battery is running down. You will also have the option to spin down the hard drives and shut down the display a certain time before entering suspend/hibernation. This reduces power consumption without the inconvenience of locking out the user while the system re-initialises. As noted above, of the two options, provided the hardware is fully ACPI compatible, hibernate is the more secure state.

In most operating systems strict criteria for power management are not configured as a default they have to be set by the user or the system installer. That is why configuring power management should always be one of the final steps in system installation/configuration. In addition, it's also possible to get power optimisation programs which monitor system usage and suggest ways of changing power settings and system timing signals to reduce power consumption – although these often need more advanced knowledge of the operating system in order to enable them. There are various proprietary programs available for Windows, a number of which allow remote system access over a network, so that power setting can be configured to optimise power consumption (for example, Data Synergy's PowerMAN or 1E's NightWatchman, and Verdiem produce a free application called Edison). There are also proprietary applications for MacOS. For Linux there are applications developed for large servers, and, as suggested, there are also free programs such as *powertop*.

18. Bemowski, Karen. Windows IT Pro (September 2010). Buyer's Guide: Power Management Software for Windows Workstations. www.windowstpro.com/article/buyers-guide/Power-Management-Software-for-Windows-Workstations-Buyers-Guide Download the product table from www.windowstpro.com/content/content/125800/101910-updated-125800-table.xls

19. For example, the Less Watts projects www.lesswatts.org

20. *powertop* www.lesswatts.org/projects/powertop/

Box 4.5.

Operating systems check – list

Always match operating systems to hardware specifications

A machine which doesn't have enough system resources to run its installed software will struggle, and waste the user's time; a system that is too powerful for the software installed upon it represents a waste of resources – *in effect the user is wasting the system's time*. What we should find, by planning our needs and matching them to a machine specification, is a balance between these two extremes.

Define a role for a machine and install a system that meets that role

Operating systems work well when they perform a defined purpose. That could be a general purpose like a desktop machine, or a specific purpose like a data server. Sometimes you can combine both – such as a desktop with a web server installed to allow web development. In practice it's better to manage the load on machines in a way that optimises their role rather than obstructing it. Find a role for the hardware you have, and only buy new hardware when what you have can no longer perform the tasks you require reliably.

Old machines can be useful if you assign them a suitable role in the system

Some roles require little processing power, such as a printer server, file server, or running network services such as DHCP and email. A new server will do all this and

more – but if you have a suitable old machine the ecological option is to use that equipment until it can no longer reliably fulfil the role.

Always document the machine hardware and the system installation

Knowing your hardware is essential to maintaining it, and especially for dealing with faults and bugs when they crop up. More importantly, data secured behind password access controls, and especially encrypted data, can be lost if the password is forgotten. Securely stored system documentation, listing system and hardware configurations, can avoid this.

Configure power management options to reduce consumption, but avoid settings that are so obstructive that the user disables them

Power management can make savings in energy use, provided that it does not become so obstructive that the user disables it. After a period of idleness, switching to hibernation is more secure than sleep mode as it is less likely to create data loss in the event of a power failure/ the machine being switched off.

Give feedback on power management

The best way for a user to understand the load of their system is to get direct feedback on the desktop – allowing them to change their behaviour and not over-tax their system.

4.4.3. Providing feedback on system load

Finally, for the user at their computer desktop, there is very little indication as to the level of “stress” that their machine is experiencing at any point in time. The machine can appear to slow for no apparent reason, or the mouse can stop moving, leading them to click buttons and press keys, which will often exacerbate the problem. The solution is to install a small power monitor application on the task bar of the desktop to give feedback. This will show when the system load is reaching a problematic level.

More technical applications will also break down the display to show the overall system load, processor(s) speed/load, network activity and hard drive activity.

By giving feedback you allow the user to moderate their demands on the system, making it use less power, and avoiding overloading the system to the point where applications crash or the system hangs. There are free load monitors available for Windows systems, and Linux desktops usually come pre-configured with applications such as *system-monitor*.

DESKTOPS: PROGRAMS, PERIPHERALS AND GADGETS

When we talk of “computers”, most people will think of a desktop computer system, with a mouse, keyboard and video display screen. In the previous units we've looked at machine hardware and software. In this unit we'll look specifically at the influence of desktop computers on the sustainability of information technology. Surprisingly what we find is that the variable which has the greatest impact is not the machine – *it's the machine's user*.

The way we use desktop computer systems, and how we use and generate data, has a great effect on the ecological footprint of the system. For example, the programs we use, and how well we understand those programs and their functions, have an impact on the environment. Certain peripherals, such as printers, also have an impact as a result of the way we use them, and if we look at recent trends, it's the effects of information capture gadgets – such as digital cameras and sound recording – that are driving the storage capacity demand of IT. When we consider the wider global network of computers, the increasingly “virtual” nature of the relationships people maintain is driving the demands on network connections and data storage.

Over the next four units we'll examine the ecological effects of desktop computers and the ways we can tackle them. In this unit we'll look specifically at the issue of desktop computers; in the next we'll look at information storage and the measures we can use to keep, back-up and secure information; then we look at computer networks and how our use of data influences the efficiency of data movement and resource consumption; and finally we'll look at physical security and maintenance.

5.1. THE ROLE OF THE “USER” IN THE ECOLOGY OF INFORMATION TECHNOLOGY

In unit 2 we looked at the “ecology” of information technology; the effect that our demand for information systems has on the environment generally. There was a very important part missing from that analysis – *the user of the equipment*. When we think of desktop computers, an integral part of the system is the human user who sits in front of the computer. The type of job that they are doing, and how efficiently they work, affects how intensely the computer is utilised and how great the impact of that activity is.

There are many facts and figures about the impacts of computing in terms of emissions, but because they often consider different system boundaries, or because they are based on certain industries or activities, the results given vary widely. Often the units these impacts are measured in differ too, making direct comparisons difficult. For example, some give the emissions as an equivalent of elemental carbon (C), some the emissions of carbon dioxide (CO₂), and others give measurements adjusted to reflect the value of different greenhouse gases as an “equivalent” of carbon dioxide (CO₂eq).

To make comparing the impacts of using desktop systems less complex, we'll take rough figures from the range of information available to look at the scale of the impacts from different parts of a desktop computer system:

- Firstly, the *internet* - the energy used by the global digital networks that the computer is connected to. This is a difficult figure to find because it assumes certain things about the user of the system. To give a general figure, let's take the annual estimate for the internet's carbon emissions, 300 million tonnes of CO₂, and calculate that as an average figure for two billion internet users connecting for eight hours a day. **That gives a figure of around 0.05 kilos of carbon emitted per hour of use.**
- Secondly, *power consumption*. How much energy or carbon is produced for each unit of

electricity consumed² depends upon how efficient the power generation system is and the types of fuel it uses. Coal power plants generate around one kilo of carbon dioxide for each kilowatt-hour of power consumed, while natural gas produces around 0.44 kilos, and wind turbines about 0.01 kilos. Therefore, assuming a mix of fuels which give around 0.6 kilos per kilowatt-hour, and assuming the desktop system consumes an average 250 watts of power, **the PC emits around 0.15 kilos of carbon per hour of use.**

- Thirdly, *embodied energy*. How much energy it takes to manufacture a PC³ and how much greenhouse gases are released as a result vary widely. The results usually depend on what the researchers carrying out the work measure, and where they sourced their data. Let's apply a very broad rule, and assume that the total impact of the hardware is around 1½ times the impact of the power consumed, in which case **the embodied energy of the PC hardware is 0.23 kilos of carbon per hour of use.**
- Finally, *the computer user*. The computer user runs on food rather than electricity. In the most developed nations, where the average diet consists of a lot of processed and out-of-season fruit and vegetables, the amount of energy and carbon is much greater than in less developed states, where more local and basic foodstuffs are consumed. Research in the United States (US) puts the annual carbon emissions from the average diet at around 5,500 kilos of carbon per year.⁴ Other

1. Guardian Online (August 2010). What's the carbon footprint of... the internet? www.guardian.co.uk/environment/2010/aug/12/carbon-footprint-internet

2. Wikipedia, Comparison of life-cycle greenhouse-gas emissions.

en.wikipedia.org/wiki/Comparisons_of_life-cycle_greenhouse-gas_emissions

3. Kris De Decker (June 2009). The monster footprint of digital technology. www.lowtechmagazine.com/2009/06/embodied-energy-of-digital-technology.html

4. Brighter Planet (2010). The American Carbon Footprint: Understanding and reducing your food's impact on climate change. attachments.brighterplanet.com/press_items/local_copies/52/original/carbon_foodprint_wp.pdf

estimates put European dietary emissions at about half that figure, and in less developed states perhaps a quarter of the value in the US. Given that desktop computers dominate the business world in the US and Europe, let's take a figure for food emissions as somewhere between Europe and the US – 4,000 kilos of CO₂ per year. Therefore, averaged over a year, **the computer user emits 0.46 kilos of carbon per hour of use.**

While green IT guides often look at the carbon emitted by the system hardware, it's the system's "wetware"⁵ (the human component) that's also a significant source of carbon emissions. In the example above, the desktop computer system emitted half the carbon emissions of the food used to "fuel" the system user; if we had used the example of a laptop computer, the user might have emitted twenty times more than the computer system they were using. Yes, the efficiency of power generation and how we make our IT equipment is very important, but if we look at the factors which influence the use of IT then the role of the computer user is very important too.

As the computer user is a significant part of the process, it's important that our efforts to reduce energy and resource use include the needs of the computer operator. For example:

- *Training* – learning how to use all the features of the system allows us to more efficiently use the software on the machine, and helps us to work more fluidly;
- *Software and hardware* – if the computer can't easily run the software programs, or the operating system regularly crashes and requires rebooting, then that affects the efficiency of the computer user;

- *The use of networks* – if the user is overloaded by emails⁶ not directly related to their work, or are swamped with spam email, that affects the rate at which they can use the system; and
- *The way we create and manipulate files* – with a little training or practice we could minimise the resources our work required, increasing the efficiency with which we use the system.

One of the most important factors is data loss. Information is a resource, just like metals or other commodities. It costs time and energy to create the information we store on our computers, and in that process the emissions created by the computer user are as significant as the machine itself. The loss of data – especially if such losses are avoidable – represents a waste of energy and an excess of carbon emissions. Surveys on the causes of data loss⁷ show that hardware or system malfunctions account for 44% of data losses, human error 32%, software errors 14%, computer viruses 7%, and natural disasters⁸ – such as fire or flooding – only 3%. Another problem, when machines become clogged with files, is that we can lose track of data, and spend time on the computer blindly searching for something stored on the hard drive.

By avoiding data loss, or losing track of information and spending a lot of time searching, we can reduce the repetitious and unnecessary time spent searching for or recreating lost data. How we configure the operating system and programs we use, and how we run systems for backing up and data storage, can have an impact on the extent or likelihood of data loss.

5. Wikipedia, Wetware (brain). en.wikipedia.org/wiki/Wetware_(brain)

6. Guardian Online (October 2010). What's the carbon footprint of... email?
www.guardian.co.uk/environment/green-living-blog/2010/oct/21/carbon-footprint-email

7. For example, QASys (2009). Statistics About Leading Causes Of Data Loss.
www.aztecscsystems.com/downloads/OASys_Protect_FAQ.pdf

5.2. MANAGING THE ENERGY USE OF DESKTOP SYSTEMS

Let's consider a more traditional approach to green IT – the energy and resources consumed by the desktop system. In section 4.4 we examined the power saving features of operating systems. Changing the way programs are executed or the system is configured can easily save 5% to 10% of the power consumed by a desktop machine – perhaps more if the system is switched into the most efficient power saving mode. Unlike laptops, desktop systems often have no power-saving features enabled by default – it requires the user to manually activate these features.

Of all the power-saving features, perhaps the most valuable is the *hibernate* or *sleep mode*. This is often used with laptops – for example when you close the lid the system is suspended into memory. For desktops it can save a large amount of power by configuring the system to switch to a power-saving state when you are away from the computer. There are two steps you should undertake to implement this:

- Configure the power management system to suspend (also called sleep) or hibernate the computer when you are away from it for a certain length of time. Of the two, hibernate is the most secure because it backs up any open programs to the hard disks rather than to the memory – unless the programs used have the option to “save a back-up copy” at regular intervals to prevent data loss. The downside is that restoring from hibernation takes longer than from suspend, and for that reason people avoid using it. How long you set before the machine automatically enters the suspend/hibernate mode depends upon your work patterns. Ideally set this to 20 to 30 minutes or less.
- A better option is to place an icon on the task bar or desktop that initiates suspend or hibernation when you click it. This is more predictable because you decide when the mode is activated, and if used routinely it will save more because the computer need not be inactive for a period before power saving is enabled. The default time limit in the power-saving configuration should still be set as a back-up, in case the user forgets to click the icon; but

putting an icon on the desktop allows more timely control over the power consumption.

Many people configure screen savers on their desktops. These were developed for use with old-style cathode ray tube (CRT) monitors in order to stop the screen's phosphor coating degrading through over use – and were useful because they prolonged the life of the CRT monitor. This feature is now redundant with the development of flat-screen monitors. Moreover, screen savers don't “save” energy. In fact, when they create complex animated or fast-moving images (for example, music visualisation software), they can use far more energy than a blank screen.

Even so, screen savers are important. Setting your screen saver to start after four or five minutes without any activity, and requiring a password to unlock it, is an effective security measure against casual intrusion on a computer system. As in the case of power saving, putting an icon on the desktop/task bar to lock the desktop enables the user to activate the screen saver when you leave the computer.

With newer flat-panel monitors, it's important to configure the system so it turns the display off when not in use. This is usually part of the power management configuration (recently manufactured CRT monitors have this feature too). After a certain period of idleness the display will be switched into standby mode, saving a large proportion of its power consumption. How long before the screen goes blank is set by a slider or list box in the power management menu. On most Linux systems, and some proprietary Windows applications, another important saving can be made by dimming the display brightness after a certain period of idleness – usually around ten or twenty seconds without movement of the mouse or pressing a key. This reduces the level of power consumed by the monitor's back light. This feature is also enabled from the power management menu (usually a check box to “dim display brightness”).

For both CRT and flat monitors the brightness of the monitor is also an issue. Brightness isn't usually a feature solely related to the workings of the monitor. It is influenced by the

relative contrast of the background we view the display against. For example, if there is a bright sunny window behind your monitor, or the sun shines across the monitor, it appears that the display is dimmer than it actually is. In these situations users are often tempted to turn up the brightness and that in turn will increase power consumption. A brighter monitor can also cause greater stress on the eyes of the user. Locating the monitor in a position where it is viewed against a darker background, and where bright lights or sun light do not shine directly across the surface of the screen, achieve the opposite: the monitor will appear to look brighter. This allows – the monitor brightness to be reduced, and that in turn can reduce the stress on the user's eyes.

Sometimes, particularly where the monitor is small, users can have problems reading the screen. In these cases they might turn up the contrast to increase the definition of the text on the screen, but it doesn't create a great improvement in readability. Where users have difficulty viewing what's on the screen, or with particular applications, the best solution is to experiment with the colour scheme. Particularly where the user's eye sight is poor, or they have a small colour blindness or sensitivity, changing the default colour scheme, and changing the font size of characters, can make a big difference to the readability of the display. As well as changing the general desktop theme or colour scheme, many applications allow you to vary their default colours to produce a more readable interface.

5.2.1. Peripherals, gadgets and power consumption

Assuming that power management is enabled, the machine suspends or hibernates when not in use, and video display power management has been optimised, the most significant remaining power usage is going to be the peripherals attached to the computer. Computer networks, especially with the development of broadband, are a problem because they're on all the time – we'll examine this in unit 7. We'll also look at printers in section 5.4 below. What remains are all those other devices which we plug into the machine – speakers, USB hard drives, cameras, scanners, web cams etc.

USB peripherals usually draw their power from the computer, and when the computer

shuts down/hibernates the power to the USB ports is cut. As a result they won't draw power when the machine is not in use. The exceptions are the yellow coloured “sleep and charge” USB ports included on more recent motherboards. These are intended to power/charge mobile phones and other devices, and keep the power turned on even when the machine has been put in a “soft” shut down (not turned off at the mains). When using USB hard drives and other similar devices, don't plug them in the yellow port or they will keep running even when the machine has powered down.

When the machine shuts down or goes into a suspend/hibernation mode, the power to devices which have their own power supply will not be cut. Ideally these devices should have their own internal electronics to monitor the activity of the machine and go into a standby mode when they receive no signal. Most new video displays and PC speakers have an “auto power on/off” feature and will turn on and off depending upon whether they receive a video/audio signal.

Of course all mains-powered devices, even if they power down with the system, are still drawing power in their stand-by mode. All recently manufactured devices should, when in stand-by mode, meet the requirement of the International Energy Agency's One-Watt Initiative.⁸ This requires that devices marketed after 2010 should have a maximum standby power demand of one watt, and after 2013 that figure should be halved. While it's not worth deliberately replacing peripherals for newer ones which include more efficient stand-by power supplies – *the power saved is negligible compared to the waste of embodied energy caused by discarding a functional device* – when you replace equipment ensure that it meets these new standards.

While we can minimise power use from peripherals, the most effective way of eliminating stand-by power consumption is to have the ability to switch off the entire installation by having easy access to the mains switch. This allows the user not only to shut down the computer, but also to isolate the power supplied to all its associated peripherals. It is also safer as it guards against the likelihood of power surges damaging

8. Wikipedia, One Watt Initiative. en.wikipedia.org/wiki/One_Watt_Initiative

Box 5.1.

Changing the desktop environment/configuration

On any desktop system, the size of the video display has a significant impact upon the power consumed by the system. As we saw in unit 3 when discussing video hardware and the AGP slot, the higher the resolution and colour depth of the display, the more data it requires to generate the display. In turn, the programs which create the graphical desktop and organise display data for other programs add to the demand on the system, and that increases with the scale of the display.

With Windows and MacOS you are locked into a single display system; while you can vary the theme of the desktop (e.g. in Windows 7 you can configure a display theme which looks like Windows XP) you can't change the nature of programs generating the display. With Gnu/Linux you are free to change the programs which generate user interface in order to change the amount of processing power required to generate the display and that can reduce power consumption.

Lightweight Linux distributions come with a lower power desktop environment already configured. General purpose distributions are usually configured with one of the more powerful desktop environments (usually Gnome or KDE) but can change the display environment if required. Often you can select which environment to use when you log into the system. The main effect of changing the user interface is to reduce the amount of processing power required to generate the display, reducing the amount of memory required for the display (which allows other programs to run faster). On older hardware this can significantly increase the performance of the system.

Usually the selection of which desktop environments are available is made when the system is installed. Depending upon which Linux distribution is used, other desktop environments can be loaded very quickly from online software repositories. The "heaviest" display manager is KDE, followed by Gnome. These are the standard desktops with most distributions and are optimised to provide a feature-rich desktop. XFCE is less resource hungry than Gnome, although it provides many of the same features. Finally LXDE is the lightest, and on older hardware it definitely improves the usability of the desktop. However, while XFCE/LXDE will work well with most applications, fast video programs or the animated content of web pages may not work well or at all.

Finally, with both Windows and Linux another way to squeeze some extra speed out of the system, or slightly reduce power consumption further, is to run the display hardware at a lower resolution. For example, most new high-definition (HD) displays are 1920 pixels wide and 1080 pixels high. If you're primarily carrying out office work or internet browsing on the machine, rather than playing video games or watching DVDs, you might be able to run the display at a lower setting to save processing power (e.g., 1280x1024, 1152x864 or 1024x768). You may also be able to reduce the colour depth of the display. Provided the video hardware can support the change without causing problems, the effect of reducing display resolution and colour depth is to free up memory and reduce the amount of power used by the video display electronics.

ing the equipment, or electrical faults causing fires while the machine is unattended. As outlined in section 3.4 and figure 3.8, the easiest way to isolate the system from the mains is to use a switched and surge-protected plug strip.

The machine and all its peripherals are plugged into the strip, and at the end of the day the user can shut down the system and then perform a “hard” shut down by disconnecting the power using a single switch.

5.3. PRINTING

There is a modern marketing myth, developed in the 1970s⁹ when the first word processors were available, that one day paper would be eliminated from offices by the development of better computer technology. In 1980, before the IT revolution swept across the office environment, world office paper consumption was 70-million tonnes a year; by 1997 it was almost 150-million tonnes. In Britain 120 billion pieces of paper are printed each year. The average British office worker prints 22 pages every working day and research suggests that 44% of this is easily avoidable and over 21% of this printed matter is disposed of before the end of the day.¹⁰

The paperless office hasn't arrived for many reasons. These are partly *technological* – not everyone uses computers well or feels comfortable reading from screens. They are also partly *psychological* – research shows that the physicality of paper is an important part of how we comprehend its content and mentally map that information. Perhaps most importantly, the change in the scale of office administration as part of the modern service-led economy has made the market for paper grow compared to the 1970s. It's this last point that is the least explored. Yes, the amount of office paper generated per person may have fallen with the advent of new technology; but as the global economy has grown, so the numbers involved in office and administration have grown too. The overall effect has been to eliminate the savings from new technology as more people use these technologies and generate a greater demand for paper.

From an ecological point of view, the debate over paper and printing also misses an important aspect of the technology of the printed word. *Not all printing is bad!* (see box 5.2). It takes a lot of energy to make paper, and a comparatively small amount to print information upon it. However, once printed, providing it is appropriately stored and cared for, that information can be available to anyone without investing further energy for years or decades. Digital information requires a computer system, powered up with electricity, in order to access the data, and more equipment to store that data securely. Each time it is accessed it requires energy, the longer it takes a person to read through the more energy is consumed, and if it is kept live online it requires a powerful server machine and network to hold it ready for access. Depending on the type of publication, there will inevitably come a point when it is more efficient to print information which will be required, in an unchanging form, for the long term – for example, libraries of books, technical information and important reports. Where information is short lived, or changes rapidly, or is only valid for a short period, then it is more suited to storage and viewing on digital media.

The two printer technologies in widespread use are laser and ink-jet printers. How you interpret the ecological impact of these two technologies is complex, and depends upon how much printing you do. There are some life-cycle studies of different printing technologies, but they do not produce a clear result

9. Business Week (June 1975). The Office of the Future. www.businessweek.com/technology/content/may2008/tc20080526_547942.htm

10. Global Action Plan (2007). An Inefficient Truth. www.it-energy.co.uk/pdf/GAP_An_Inefficient_Truth_Dec_2007.pdf

in terms of which is the better option.¹¹ While laser printers are cheaper to run for large volumes, and so have a lower energy expenditure per sheet printed, they can produce a larger volume of solid waste if they require replacement of the toner cartridge (refilling with toner powder is usually only an option on larger/more expensive laser printers). In contrast, while ink-jet printers produce less solid waste, they use more energy per sheet and can have a higher embodied energy because of the smaller volume of printing during their operating life.

If we look for some general rules and advice about printing, there are a number of ways we can maximise the efficiency of using printing or copying devices:

- Purely on the basis on the embodied energy of manufacturing, having a large number of users share a printer is more efficient than having a printer for individuals/small groups. The capital savings also means that buying a more expensive, more efficient printer, is cost effective for larger groups.
- Small laser printers which will duplex are now available. This means they print both sides of the sheet automatically. Even though they cost more, routinely double-sided print jobs will cut paper use and, due to the high price of paper, pay for itself fairly quickly. While you can print both sides by printing odd and even pages separately and reloading the paper, if there is a fault during the printing of the second side this can often waste paper – so a duplexing printer is the best option. Duplexing can be enabled from the printer configuration interface in order to reduce the amount of paper used in routine print operations; users can still print single-sided, but they must deliberately select this option before printing.
- Most word processors and other applications create documents from a standard template, and most allow the configuration of this template file to be edited. Editing the template can enable a specific set of printer settings in order to ensure the most efficient

use of paper and ink/toner. In addition, by minimising the size of the margins, and not using an excessively large font size, it is possible to increase the amount of information on the page.

- The use of ink or toner use can be minimised if the “draft” mode is used for printing – not only with draft documents, but for any publication that is not intended for formal use (memos, circulars, etc.). Again, this is something that must be set in the printer configuration menu on the operating system or software application.
- Most printer control software allows control over who has access to printers. Filtering which users have access can limit the use of more expensive/high quality machines, and allow general printing to be carried out from a cheaper-to-run general purpose printer which uses lower quality paper. This will save money and environmental impacts (for example, limiting access to printing on colour printers and allow all access for printing on the cheapest printer).
- On a plain printed page only 5% of the area is covered with ink/toner. If printing graphs or slides using solid blocks of colour, the amount of the paper covered is far higher. Using a hatched infill on graphs/images rather than solid colour especially when using ink-jet printers – reduces the amount of ink or toner required for printing documents significantly.
- Recycled paper requires half the energy, and produces less pollution, than using non-recycled paper. If possible, rather than bright-white paper which has to be chemically bleached to get it clean, using lower grade off-white paper will save more. While there are various various “sustainable forestry” certification schemes for paper, these products still use new rather than recycled pulp, which uses more energy and emits more pollution. Using 100% recycled paper, or paper with a majority of recycled pulp, is always the better option.
- Waste white office paper is worth more as a product for recycling, and is more efficiently recycled if it is segregated from other lower quality/coloured papers. If it is possible to organise a separate paper collection for higher

11. Fraunhofer IZM/PE Europe (2007). EuP Preparatory Studies “Imaging Equipment” (Lot 4). www.eup-network.de/fileadmin/user_upload/Produktgruppen/Lots/Final_Documents/Lot4_T5_Final_Report.pdf

Box 5.2.

Printing information is not necessarily a bad thing

Paper is a very energy-intensive product to manufacture. There are various estimates for how much energy it takes to produce, and the impact depends upon many variables – from whether the paper plant is run by coal or hydro-power, to how far the paper pulp is transported. For the sake of this comparison let's take a general figure of 10 kilo-Watt-hours per kilo of paper. The weight of a single A4 sheet of 80 gram per square metre paper is around five grams. The energy required to produce one A4 sheet is therefore 50 watt-hours.

For the sake of simplicity we'll ignore the embodied energy of the computer, and to balance that we'll ignore the energy used in printing. We'll also ignore the energy required to store and transmit the digital information. Let's assume it takes three minutes to read a single side of A4 paper. Let's also assume that the desktop computer consumes a quarter of a kilo-Watt, or 250 watt-hours of electricity. Therefore in a fifth of an hour, or 12 minutes, the computer will consume the same amount of energy as making a sheet of A4 paper. If reading a page every three minutes, the user is consuming a quarter of the energy reading from the screen than if they were reading from the paper – which is obviously better for the environment.

However, let's optimise this system. Let's assume the paper is double sided. That uses extra energy for printing, but not significantly more because much of the energy is contained in the paper itself. Reading from the screen will now take almost half the equivalent energy as

reading both sides of the paper. Now let's assume that three people read the information, sharing the same piece of paper and reading it one after another. They are now using the less energy reading the paper copy than it takes to read the information from the screen individually. And, of course, if ten people read it, then the paper version is using far less energy.

In reality the figures are not so simple. For example, if we were reading the information on a laptop computer, that uses a tenth of the power of a desktop system; you'd have to share the piece of double-sided paper between twenty-five to thirty people before the paper version was better than reading it on a laptop. Quite simply, we must “think before we print”. What we can say is that where often-used information is “static”, meaning it is infrequently revised or changed, it is probably better that this information is kept as a printed volume which people can pull from a shelf and read at any time. For example, reference materials and technical manuals. Where information is in “flux”, and in continually being updated, or it is only required for a single use by a few individuals, then online distribution is probably the best option. For example, news bulletins and memos/emails.

What this analysis shows is that there is no one simple solution when it comes to the ecological footprint of our decisions on the use of technology. It is the peculiar characteristics of the factors in each decision which determine what is the best option overall.

quality paper, segregate the collection of white waste printer/photocopier paper.

- Most digital photocopiers are little different in function to a laser printer, although they are usually more energy efficient than a laser printer when producing a large number of copies. For making a large number of multiple copies a photocopier is likely to be

more efficient than a laser printer. However, if you are routinely producing large print runs on a photocopier, getting the job produced by a conventional commercial printer using wet inks and a conventional printing press will be more efficient. Using conventional printing also allows a range of lower-quality/recycled papers to be used.

5.4. E-READERS

Print is no longer a purely paper-based medium. As well as traditional computer screens there are now a number of devices which are specifically designed to replace printed books. Unfortunately the ecological impacts of reading information on the latest tablet devices is a complex issue to interpret.¹² For example, while designed for reading print-based information, they are also capable of undertaking other tasks as well. It also requires information on the carbon footprint of the e-reader, and that is not always provided by the manufacturer. There are some studies, but they are often partisan. For example, one positive study¹³ of e-readers didn't consider the lending or resale of printed books to other readers when it calculated the impacts of e-books versus conventional books; another report¹⁴ more favourable to the print industry notes the uncertainty in the figures which underlie comparisons between print and e-books. Of course, with any electronic technology, the golden rule is always to maximise its use and keep it functioning for the longest time possible before replacing it – which appears to be the bottom line of almost all studies on the comparison between conventional and electronic books.

Another aspect of e-reader technology not often discussed is intellectual property and in particular the use of digital rights management to control the use of the devices and the files they contain. For example, Amazon's Kindle service allows books to be shared by up to six devices registered to the same Amazon account, but not *between* Amazon accounts. There is a facility to loan an e-book you have bought to another Amazon user, but only once, and only if the publisher permits it. One of the great features of the traditional printed book is that it can be shared freely by many people, and can be given/sold back to shops who will re-sell the book. The resale/recycling of second-hand books means that the impact per person of printing the book becomes progressively less as it is passed around. With e-readers the sharing of the data is usually blocked; on desktop systems equipped with e-reader software there is greater scope to share files, but only while the operating system permits that to take place. Consequently it's hard to minimise the resource impact of e-books by sharing the files between many readers.

Additionally, all the transactions of book download and loans are logged, and that potentially has an effect upon our privacy and freedom of thought.¹⁵ When you download files to the e-book reader from a service, the details of the book and the account are logged, and data is available to the company and its corporate partners to mine as a source for marketing or customer profiling; or by security services profiling members of the public from their online browsing habits. While intended for use by com-

12. Palmer, Brian (August 2010). Should You Ditch Your Books for an E-Reader? The Slate. www.slate.com/articles/health_and_science/the_green_lantern/2010/08/should_you_ditch_your_books_for_an_ereader.html

13. Cleantech Group (2009). The environmental impact of Amazon's Kindle. www.tkearth.com/downloads/thoughts_ereaders.pdf

14. Green Press Initiative (2011). Environmental Impacts of E-books. www.greenpressinitiative.org/documents/ebooks.pdf

15. EFF (2010). Digital Books and Your Rights: A Checklist for Readers. www.eff.org/wp/digital-books-and-your-rights

mercial interests, this same technology could enable control over public debate traditionally associated with more repressive states.

On certain services it is also possible for the service provider to read and manage the content of the e-reader device. For example, it is possible for the online services to delete books on

a device if there are rights issues about selling the book, meaning that a book could be digitally “banned” from e-readers by the service provider. Deletion of books on customers' e-readers without owner permission has happened on the Kindle service – perhaps ironically one of the books deleted was George Orwell's *1984*.¹⁶

5.5. FILES, FOLDERS AND ORGANISING DATA

As the speed of data networks has increased, and with the advent of downloadable TV and radio programmes, the amounts of data streaming into computers has grown exponentially. At the same time we have more digital capture devices – from mobile phones, to digital audio recorders, video/stills cameras, and image scanners for digitising images and printed materials – which routinely generate giga-bytes of data (see box 5.3). All of this data fills our own computers, and we're file sharing and emailing large quantities of it to friends and contacts via the internet – replicating multiple copies across computer networks and creating a greater demand for storage.

Once data has been downloaded to a computer, the next issue is how those files are organised. Again, the growing size of files has been driving the demand for larger hard drives on desktop computers. That's partly due to the downloading and editing of video, although the larger capacity of digital cameras and portable music players means that the data imported to the system are becoming ever larger as consumer gadgets store larger quantities of files. As the sources of data have grown, so the total number of files on a hard drive has increased substantially too. Perhaps a routine example of the growth in files would be saving a web page with all its associated formatting and image files which for websites with large amounts of advertising can produce 40 to 100 files along with the main HTML page file.

The traditional method of organising data is by using directories or folders, describing the types of file contained in the name of the folder – as shown in figure 5.1. As we work and create files it is important to store these in subdirectories/folders to hold them separately from

the mass of data held on the system. Ideally the folder name should identify the contents, but as time passes and as old files get shuffled around the system, keeping track of all these files, folders, names, and what they relate to becomes progressively harder.

The difficulty with this approach is that the directory names only have meaning to the human operator. They have no special meaning to the operating system, and may even be vague or confusing for other users too. Organising files into a directory tree also requires care and consistency. Without attention to the details of which files go where, important data can be lost in a sea of files and subdirectories very quickly. The point is the user can soon spend progressively longer periods of time searching for information – and that's because the design of the storage system does not match the complexity of the collection of information being held.

What has developed over the last decade are systems which enable us to catalogue the files on our computers with far more relevant data than a simple file or folder name. Just like web search engines, desktop search or indexing systems¹⁷ read the content of the users files and create an index of the information held. Of course the problem are those files which do not contain textual information that is intelligible by the machine – such as videos, sound files and

16. Johnson, Bobbie (July 2009). Why did Big Brother remove paid-for content from Amazon's Kindles?, Guardian Online. www.guardian.co.uk/technology/2009/jul/22/kindle-amazon-digital-rights

17. Wikipedia, Desktop search. en.wikipedia.org/wiki/Desktop_search

Box 5.3.

Digital cameras, mega-pixels and data overload

One example of the way the increasing data capacity of equipment is driving data storage is digital cameras. The obsessive selling point for all new digital cameras is the number of mega-pixels inside the image sensor: The higher the mega-pixels figure the better the images are assumed to be. This is not always the case since the true image quality is dependent upon the quality of the optics and the use of filtering and correction software to render a true image under different lighting conditions.

What has happened as the camera mega-pixel count has risen is that the file size of photographs has risen, but that doesn't necessarily translate into "better" photographs. If we look at the mega-pixel rating and the size of print that can be produced, most cameras are taking impractically large images. They are much higher resolution than commercial digital print machines (which usually operate at around 300 pixels per inch resolution); are far higher than can be displayed on the computer monitor (which operate at 75 to 150 pixels per inch); and are excessively higher than is required for producing everyday documents or web-based content (where images of a few hundred pixels square are often all that is required).

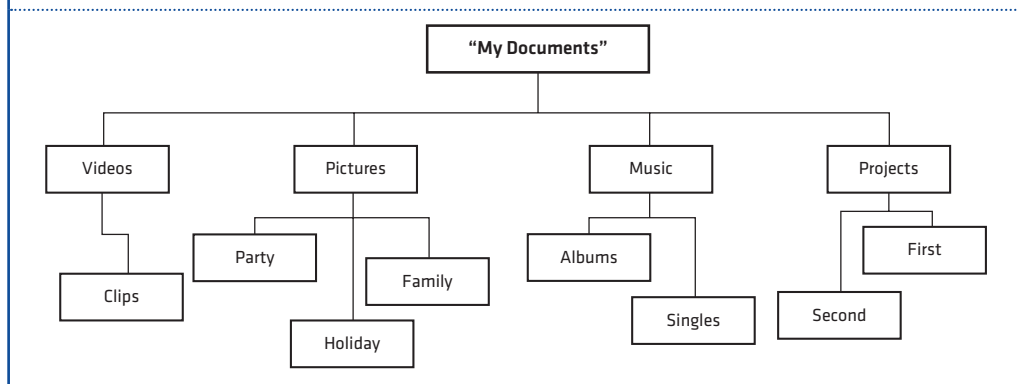
Buying a camera purely because it has more mega-pixels is once – it exceeds six to eight mega-pixels an excessive use of resources for most everyday purposes. Instead, by matching the resolution of the image to the application you

Table 5.1.	
Camera mega-pixel (MP) rating and maximum photo print size	
MP rating	Maximum print size
3	6.8x5.1 inches, 17x13cm
6	10x6.7 inches, 25.4x17cm
10	12.9x8.6 inches, 32.7x21.8cm
12	14.3x9.3 inches, 36.3x23.6cm
16	16.4x10.8 inches, 41.7x27.4cm
To calculate the relative size of the image take the pixels width and height and divide by the image resolution. E.g. a 3MP image with 2048x1536 pixels at 300 pixels per inch gives an image of 6.8x5.1 inches.	

require the image for, you can produce smaller image files and as a result fit more images on the camera or the storage device where you keep the image files. If using a camera that has a high mega-pixel count, turning down the image resolution will reduce the size of the files produced without significantly affecting the image quality.

Figure 5.1.

A traditional directory tree



images. To standardise the organising of files on a computer, a collection of embedded labels called metadata¹⁸ (box 5.4) has been developed to describe the contents of each file. Rather like the track title and artist information used with MP3 music files, by labelling files across the system with information on their content it makes the job of locating and using the files/data stored on the system much easier. In practice describing the content using metadata produces more accurate results.

Most desktop operating systems have a desktop search/indexing system:¹⁹ *Windows Vista* and *Windows 7* come with *Windows Search*; Macintosh systems use an application called *Spotlight*; and Linux desktop environments usually have a search facility built-in, although many use dedicated search applications such as *Beagle*. In addition there are a variety of proprietary and free desktop search applications which can be installed. The benefit of a more sophisticated search system is not only its ability to index data. Most desktop search applications allow you greater control over the content and formatting of file metadata, making the management of metadata across the system far easier. Note that programs to import sound, image or video files often have metadata editing options too, and you should learn to use these when creating files in order to make indexing and finding files easier.

Finally, one of the problems about the large hard disks installed on many new desktop machines is the physical volume of data. The size of new hard drives, increasingly rated in terabytes (a terabyte is a million megabytes), allows users to fill up the storage system with files which are poorly organised, and, worst of all, which are not backed-up. The physical scale of the data stored can also make backing-up an even bigger headache. When considering the efficiency of desktop systems, having large amounts of files stored on the hard drive can complicate how we manage that data. For example the more files, the more effort the desktop indexing system has to expend in keeping track of those files, and that increases the processor power and drive space required to map the system. More problematically, in the event of a system failure, a drive full of superfluous data can hamper the recovery of more important files.

With larger hard drives the temptation is to allow data to stack-up on the system, rather than moving the important files to a back-up storage system and deleting the outdated/unwanted files. One of the most important skills to learn in organising the files on a computer system is to know when to delete them, or to move them onto other forms of data storage that are less demanding of power and resources. We'll examine this subject in detail in the next unit.

18.Wikipedia, Metadata. en.wikipedia.org/wiki/Metadata

19.Wikipedia: List of search engines. en.wikipedia.org/wiki/List_of_search_engines#Desktop_search_engines

Box 5.4.

Using metadata to label files

Metadata is often described as “data about data”. In addition to the image or textual information in a file, additional strings of information are used to identify the content of the file, when it was created, and copyright or other information on ownership. Almost all commonly used file formats support metadata in order to describe the content of the file and other statistics relevant to it. When creating images or capturing other data, many of the latest gadgets will automatically add metadata tags to describe the conditions/setting under which the data was processed and stored. For example, the type of metadata that people commonly see is the track and artists information displayed when playing MP3 music files; by labelling the sound file with text which identifies the music, the computer is able to order the files more logically when they are queued for playing.

The value of metadata is that by allowing the operating system to “know” what is inside the file it allows the search for files to be automated. For example, instead of searching the content of a file for keywords, which often will result in many false matches for common words, if the metadata contains a string of keywords which spe-

cifically describe the file, the system can present you with far more relevant results. Taking a minute or two to add metadata to a file can, over the period that the file is kept, save far more time when you are trying to manage hundreds or thousands of data files.

Most applications have a “properties” option, usually in the file menu, where you can add metadata labels. Figure 5.2 shows the properties menu from *OpenOffice* other applications are very similar. Usually there will be a default list of properties allowing a title, keywords and other comments to be defined (top box). It is also possible to define additional labels that provide a more narrow description of the content (bottom menu box) and which provide descriptions that are relevant to your own work.

In addition to providing information for desktop searching and indexing, many applications and online services now use metadata to configure imported/uploaded files – automatically inserting into the relevant fields the labels which describe the data. This makes uploading easier and also ensures that the labels used properly describe the content.

Figure 5.2.

Document properties



Box 5.5.

Check-list for desktop systems

- While the impacts of the machine are important, the time the user spends sitting in front of the machine is often more significant in terms of the ecological impacts – therefore improving our personal skills and methods of using the system is an ecological issue.
- Always configure the operating system's power management settings to minimise energy use, both by the machine and by the monitor attached to it.
- Besides configuring the shutdown of the monitor after a period of inactivity, if the function is available, configure the system to dim the brightness of the monitor when it is not actively in use.
- Operating systems that allow the configuration of a lightweight desktop environment, or the lowering of the display resolution, can be used for older hardware to increase the performance of the system and/or reduce energy consumption.
- Besides configuring suspend/sleep or hibernation when the machine has been left idle, include icons on the desktop to activate the suspend/hibernation function easily, or to lock the desktop when not in use.
- Try to use a monitor, speakers and USB devices which are able to power-down automatically when the machine hibernates or is shut down.
- On newer machines, try not to inadvertently plug USB peripherals into the yellow “sleep and charge” ports unless they require power when the machine is shut down.
- Use a printer that meets the demands you require. Ink-jet printers are better for low volume, but for higher volumes laser printers are better suited to the task and in most cases it is better to share a printer between multiple users rather than each user having their own printer.
- For information which is not time-sensitive or does not date quickly, and is to be widely shared, printing a hard copy on recycled paper for many people to read often requires less resources than each individual reading that information from a computer system.
- Where possible use 100% recycled paper, and ensure that white office paper is collected separately where a segregated collection can be arranged.
- For larger print runs a photocopier is usually more efficient than a laser printer; and for regular large print runs using the services of a commercial printer is more efficient than using a photocopier.
- Take care with data capture devices – such as digital cameras, scanners and sound recording applications – to store or encode the data at a quality or resolution that's compatible with the types of application – you require the data for. Storing excessively high-definition data requires far more hard drive space.
- To simplify finding or locating information on the computer, use an organised system of directories/subdirectories to hold the files, and also configure a desktop search/indexing system to make finding individual files easier and quicker.
- When creating documents or files, always try and complete the metadata properties provided with that file format to allow easier and more accurate indexing by desktop search systems.

INFORMATION STORAGE AND SECURITY

Managing information storage securely and reliably is one of the ways that we can reduce energy use in IT. Apart from the impacts of creating and storing information, a significant factor in the energy and carbon footprint is the human user; and if the user has to spend time searching for information, or recreating lost information, that represents a waste of time and resources. Learning to manage our use of storage devices is where we begin to tackle this problem. Then we can move on to looking at solutions that protect sensitive data from disclosure, important data from being lost on the hard drive, and all data being accidentally lost, wiped or corrupted.

This unit examines various methods to improve the security and reliability of computer systems and the information that they contain by considering the work of backing up.

The simplest rule to secure information on a computer system is to back up.¹ Backing up works on the principle of *redundancy*; if one copy is lost or damaged you can revert to the back-up copy. There are many different “pathways” for backing up. Which is the best suited/most convenient pathway to use will depend upon how your data is organised, what the objective of the process is, and whether you are trying to secure just a few files or an entire computer hard drive.

1, Wikipedia, Backup. en.wikipedia.org/wiki/Backup

6.1. DATA SOURCES AND REDUNDANCY

The purpose of backing up is to preserve important data; if we have to spend money and time backing up superfluous files then that represents an unnecessary waste of resources. Therefore the first step in developing a system to back-up data is to identify what *needs* to be backed up, and avoid including data that does not need backing up.

6.1.1. Identifying what to back up

If we characterise the types of files manipulated or stored by computers, and the need to back up that data, we can broadly identify six categories:

- **Freely available downloaded data** – for example web pages, commonly available files and data. While this data may be important, it is not irreplaceable. Unless it forms an important part of the information we keep, backing up this kind of information is not a priority because we can easily and cheaply – using less time and resources than keeping our own back-up copy – download another version if we need it.
- **Paid-for downloaded data** – for example downloaded music, subscription publications and paid-for data/files. As this is not irreplaceable data, the question is whether backing up would cost more than the value of the data you may potentially lose. Another difficulty is that music or other downloaded files may only work on the “authorised” computer they were downloaded onto; backing up digitally-locked files such as this only has value if you can transfer them to another machine later (e.g. the iTunes service allows you to authorise up to five machines to play your downloaded music).
- **Stored data (offline)** – for example software installation discs, bought CDs/DVDs or your own data back-ups. As with downloadable data, if you already have an offline copy there is no need to back up this data again unless there is a reason to do so (for example, if you have edited/changed the contents of a stored file since it was read from the storage medium).
- **Stored data (online)** – for example websites, files kept in online systems (e.g. social media)

or information stored in online services (e.g. web-based email and file storage). One of the problems of backing up is the scale of the data that is now routinely stored on many different computer systems. With the growth of online services, this data is increasingly not under the control of the person who created the file or data. While backing up to an online system is a way to avoid backing up using physical media (DVD, USB sticks, etc.), the question is what would happen if the online system was unavailable. For example, if you run a website and the server is hacked, do you have an up-to-date back-up copy of the site to quickly restore the online service? Or if you routinely store data online and lose your internet service, could you access that data if it is urgently required? All data stored in online systems should, if it has value, be backed up locally too.

- **Replaceable personal data** – for example emails, circulated reports/files and information shared between groups of users. This type of data isn't irreplaceable because copies will be held on other computer systems. The issue is how much effort it would take to reconstitute this information if you were to lose it. In most cases keeping back-ups of this information is a simpler option than trying to recover it from many other sources/locations later.
- **Irreplaceable personal data** – for example draft work, personal images or recordings, and files which are not held on any other storage medium or computer system. This is the most important data to back up because there are no redundant back-up copies to replace this data in the event that the computer system is lost, damaged or fails.

What this list illustrates is the importance of valuing the data we propose to back up. If we're trying to minimise the ecological impacts of IT, being able to plan how we back up and minimise the use of resources as part of that process is an essential part of a green IT strategy. There is one key rule to take note of in this process: *The more common and easily accessible a resource the less we need to back it up; the more unique or expensive a resource the greater the need to create a back-up copy.*

Box 6.1.

Backing up mobile devices

Today there are a variety of mobile computing devices in use, from mobile phones to personal digital assistants and fully functional laptop PCs. Mobile devices need special attention when it comes to backing up their contents because they are more likely to be lost, stolen, or damaged.

Many smart phones have built-in back-up software to copy the files they contain to a PC. Connect the phone to a PC and the contents of the phone can be copied either to a single large file (full back-up) or a directory containing the files on the phone (selective back-up). It's important to use the internal back-up software of the device in order to capture the operational and configuration files the device requires, but which are not normally available for the user to manipulate. Android phones are technically already backed up to "the cloud" – the data storage system operated by Google. It's not technically possible to make a local back-up, although there are third-party applications which allow an Android phone to back up to other online services. There are also an increasing number of commercial services, usually run from pre-paid mobile applications, which will automatically back up the data on your phone. In the event of loss or damage to the phone, the data can be recovered from the service and downloaded to a new phone.

For ordinary mobile phones there are methods to back up the contents of the SIM card, but these tend to be expensive as they often require a physical device to read the card. Some mobile operators give the option of backing up the address book and other features of the phone to the operator's system – and these can be downloaded to a new phone in the event of loss or damage.

Backing up iPods and similar music-playing devices is more complex, in part because it runs into problems of digital rights/copyright. In most cases you are loading data from another machine so keeping a back-up copy is unnecessary – unless files are only stored on the device and nowhere else. Most music download services will register a number of playing devices, and so again in the event of loss or damage the old device can be unregistered and the new device registered in order to transfer files onto it.

Finally, laptop and notebook PCs can be updated just like any other desktop system. The machine can perform a back-up over a network, either to another PC or to a network server, or connected to an external storage device (such as a hard drive, USB stick, external hard drive or memory card) to copy the files that require backing up. iPads can back up to another computer via a cable or Wi-Fi connection, or you have the option to store data in Apple's cloud storage system.

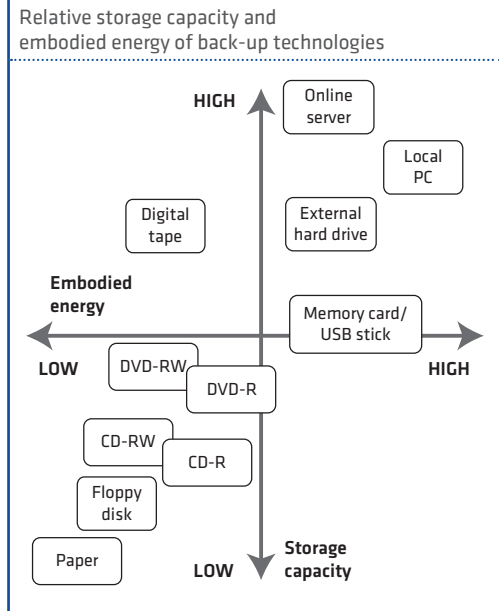
What's important is that we organise files in a way that assists the process of making back-ups. This means organising information in a way that separates the important data – which requires regular backing up – from other less important data. Separating important data from the less used and superfluous files on the computer system is also good practice because it allows low value files to be regularly deleted from the system to free up space on the hard drive. Organising data by its importance is easily done using a clear structure of directories/folders on desktop (or server machine) to isolate the most important or regularly edited files from other data which has a low value or is infrequently changed. For example:

- Users should store their current work/important files inside a single identifiable “work” directory. This simplifies the process of deciding which files to back up – backing up all current files can be easily carried out by copying the “work” directory to another machine on a network or to a removable storage media;
- Other user files that are used infrequently, or which have already been archived, can be segregated into an “archive” directory;
- Low value or superfluous data, such as web pages or downloaded files, should be kept separately from the user's recent and archived data – and these folders can be regularly sorted and their contents deleted to free up hard drive space;
- Where digital rights/copyright is an issue, data which presents a legal problem if copied or backed up (for example, commercial music or video files) can also be segregated from the bulk of user data, and excluded from the back-up process to avoid any legal difficulties which result from making copies of that data.

6.1.2. Ecological impacts of backing up

How we decide to back up will have an impact on the ecological footprint of our IT needs. For example, if we use reusable back-up media, such as digital tape or external hard drives, that can over time have a lower impact than options where we use the storage media once. The difficulty is that there is little detailed information on the impacts of different back-up technologies, and there has been no comprehensive life-cycle analysis of backing-up options to compare one option directly to another. While there are some studies

Figure 6.1.



which show that downloading music is less ecologically damaging than buying music on a CD,² or that buying software online is better than getting it on DVD,³ once you back up those files onto a CD/DVD or other offline storage media to keep them secure, most of these benefits are lost.

Much of the recent research on the impacts of computing to date has focussed on the use of servers and cloud computing rather than the impacts of everyday data storage technologies such as CDs, external hard drives or tape storage. Just as many green IT studies do not consider the embodied energy⁴ of the equipment involved, studies on the electronic distribution of data do not consider the human resources involved in creating or purchasing data – and the relative value of backing up that data via different means to prevent its loss or corruption.

2. Koomey et. al. (August 2009). The energy and climate change impacts of different music delivery methods. download.intel.com/pressroom/pdf/CDsvsdownloadsrelease.pdf

3. Accenture/WSP (October 2009). Demonstrating the Benefits of Electronic Software Distribution: A study of greenhouse gas emissions reduction. www.digitalbychoice.com/en-gb/483648_CarbonFootprint.PDF

4. Wikipedia, Embodied energy. en.wikipedia.org/wiki/Embodied_energy

While we can't make decisions with certainty, what we can do is form general rules on the impacts of different storage options based upon the characteristics of the technologies involved:

- For long-term storage (years rather than months) passively held data has a lower impact than actively maintained data (e.g. storing data for long periods on optical discs like CDs or DVDs requires less energy to maintain than storing data live using online services).
- Where data is not updated on a regular basis, the embodied energy of complex or semiconductor-based technologies is higher than other storage options (e.g. keeping archived data on magnetic tapes or discs, or on optical CDs/DVDs, has a lower impact than using external hard drives, memory card/USB sticks or online storage).

- Where data is regularly updated, meaning that any static back-up would quickly become out of date, then online storage, external hard drives and memory card/USB storage are a better option.

Figure 6.1 illustrates the relative embodied energy of different storage technologies and their relative storage capacity. Certain technologies offer a higher storage capacity, but often this is associated with a higher ecological impact. Even similar technologies can differ – for example a local PC or server machine has a higher impact than an online service because in most cases online services will be optimised to operate more efficiently than a comparable small server or PC. There are also differences between compatible technologies – for instance, because a re-writeable CD-RW or DVD-RW can be reused many times, it has a comparatively lower impact than a single use CD-R/DVD-R.

6.2. DEVELOPING BACK-UP SYSTEMS AND PROCEDURES

Deciding how often to back up is a balance between practicality and the cost of data loss. Where back ups can be automated – for example copying a whole hard drive to a digital tape – the process is less demanding of time and so can be carried out more frequently. Where individual users back up their files to removable storage, such as a USB hard drive, it requires more time and so would be carried out less frequently. What's important is that some form of back up is put in place, and then procedures are agreed to ensure that these systems are used on a regular basis.

6.2.1. Backing-up pathways

Figure 6.2 illustrates various methods of backing up. There are three different roles in this process which are defined by who has control over the system hardware:

- *System administrators* are responsible for back-ups which require special actions or security privileges – for example, backing up a hard drive to digital tape. Where a single user looks after their own system they would carry out this role, but for larger installations where

there are defined roles, these actions are the responsibility of the individual who has responsibility for the IT infrastructure.

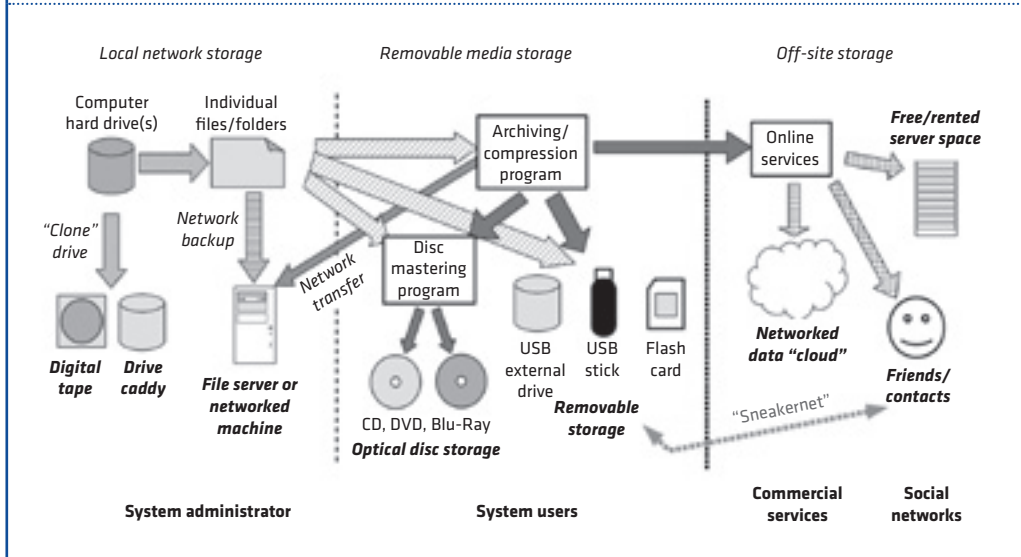
- *System users* are responsible for the files which exist within their own storage space, whether that be on their own machine or on a networked server. While their machine can be backed up centrally, encouraging users to undertake their own back-ups of essential files adds an extra level of security to the process – and helps discriminate between essential data files and other superfluous data.
- For off-site storage the system user/administrator is involved, but it also relies upon other agencies in order to provide these services. These might be free services, contracted services, or informally organised storage offered by virtual communities or social networks. The main concern with all off-site storage is how secure the data is, and whether it can be accessed when required.

6.2.2. System-level back-ups

We'll work through figure 6.2 from left to right. The process begins with the hard drive inside

Figure 6.2.

Back-up pathways



the computer. Information on a computer hard drive is vulnerable to hardware failure, corruption of the operating system (for example by malware)⁵ or user error – and of course theft of the machine or a disaster which befalls the building the machine is located in. There are various ways to protect the information stored on a machine, and each option gives a differing level of protection.

Firstly, it is a simple procedure to duplicate the contents of the hard drive, most straightforwardly by operating a second hard drive within the same machine. There are systems which can do this automatically, such as RAID⁶ hardware which automatically duplicates data on two hard drives. This is effective at preventing loss due to hardware faults but will not prevent losses from malware, user error or physical damage to the hardware. Additionally, running hard drives in parallel increases energy consumption in the machine. Therefore, for most small servers and desktop computers, the simplest option is to copy an *image*⁷ of the hard drive to a back-up storage system:

- Using a drive caddy,⁸ a second hard drive can be inserted into a machine – but as it is removable it doesn't require a power supply at all times, only when it is being used to record a back-up. This is often a cheaper option than using USB hard drives, especially if you are reusing old hard drives from other machines. When the machine is booted up the operating system can read and write from the hard drive, allowing individual files or the entire hard drive to be copied across. When the operation is complete the machine can be shut down and the second hard drive removed and securely stored in another location. In the event of a loss of data the removable drive can be re-inserted and the required data restored.
- Using a digital tape drive,⁹ a copy of the hard drive can be written to tape and then removed and securely stored. As with the removable hard drive, data can be read back from the tape, although the process for doing this can take longer.

Removable media represents an extra level of data security because the data they contain

5. Wikipedia, Malware. en.wikipedia.org/wiki/Malware

6. Wikipedia, RAID. en.wikipedia.org/wiki/RAID

7. Wikipedia, Disk image. en.wikipedia.org/wiki/Disk_image

8. Wikipedia, Drive enclosure. en.wikipedia.org/wiki/Disk_enclosure

9. Wikipedia, Tape drive. en.wikipedia.org/wiki/Tape_drive

Box 6.2.

Removable storage media for backing up

For offline storage and backing up of data there are a number of different technologies available. Which is the best-suited depends primarily upon cost, the storage capacity, the viable lifetime of the storage media, and whether or not the media can be reused. Selecting which technology to use should be considered part of developing backing-up procedures.

For day-to-day backing up of small volumes of data – such as the daily/weekly backing up of current work – cheap reusable USB sticks or flash memory cards are the simplest option. For the amount of data stored they are faster to use than DVDs, and although more expensive than DVDs when you factor in the number of times they can be reused, it is a fraction of the cost of a DVD.

For longer-term archiving of data, CDs, DVDs and Blu-ray discs are the cheapest and easiest option. For infrequently used data, or for regular back ups of user files, they are a simple way of preserving files. Optical discs can be a reliable means of storing data for a de-

cade or more if the contents of the discs are indexed to indicate their contents clearly, and are stored in optimal conditions. While you can reuse CD-RWs and DVD-RWs, they can take more time to wipe and re-record than using other reusable technologies such as USB sticks.

For larger quantities of data, external hard drives, digital data storage (DDS) tapes and hard drives mounted in a drive caddy are the best option. Digital tapes are good for cheaply backing up a single large volume of data, but take longer to record and to read compared to hard drives. External hard drives are a simple option to mirror a user's files on a large drive, and are very simple to use because they are organised in the same way as the computer's internal hard drive. Internal drives mounted in a drive caddy are cheaper to use than consumer-oriented external hard drives, but require a little more knowledge to use as it may involve re-configuring the hardware or operating system permissions in order to access the drive.

Table 6.1.					
Profile of removable storage media					
Media	Capacity, gigabytes	Cost*, £UK	Cost per gigabyte	Reusable**	Life, years***
CD-R disc	0.7	0.2	0.3	no	30-50
CD-RW disc	0.7	0.6	0.9	x1,000	15-30
DVD-R disc	4.4	0.3	0.07	no	30-50
DVD-RW disc	4.4	0.5	0.1	x1,000	15-30
DVD-R dual layer (DL) disc	8.0	1.5	0.2	no	30-50
USB memory stick	0.1-256	1-600	2-4	x100,000	10-15
Compact flash/Smart media	0.25-128	6-750	2-3	x100,000	10-15
Micro-drive	0.1-8	20-60	4-7.5	yes	5
Blu-ray disc	25	2	0.1	no	30-50
Blu-ray dual layer disc	50	7	0.2	no	30-50
Digital data storage (DDS) tape	2-160	2-25	0.2-1	yes	10-12
USB external hard drive	320-2,000	60-300	0.12-0.18	yes	8-10
Hard drive in drive caddy	500-3,000	60-200	0.06-0.12	yes	5-12
* Cost per unit – for CD/DVD/Blu-ray discs this is based on the price for a single disc when purchased in packs of 10 to 25. Price is the recommended retail price, averaged across a number of brands, sourced from amazon.com					
** An x figure represents the number of times the media can be reused before errors are likely to make it unusable. If "yes" then the media can be reused over the expected working life of the technology. For USB sticks/memory cards, this refers to the number of write operations which, because of the way the storage is configured, gives perhaps a fifth of this figure as complete write, delete and re-write operations due to the way the internal electronics of flash media function.					
*** Figure represents the optimal storage lifetime of the technology before the media becomes unreadable – can be much less if not optimally stored and cared for, especially optical discs.					

can be stored securely – for example, inside a fire-proof safe. This protects against theft and disasters. The difficulty is that restoring just a few files from removable bulk storage devices can take time. For this reason it is often more convenient to use a file server machine connected to the local network and allow users to back up files from across the network. The data held on this machine can also be backed up to a bulk storage device, such as a tape or removable drive. If restoring files on other computers on the system, it will be relatively faster to read back files from the central server than having to locate and install removable media. However, if there is a disaster, then the back-up of the central file server can be used to restore all the important files held on the network.

Where there are a number of machine sharing a network, using a file server is likely to be more efficient than individually backing up each system on the network – both in terms of the administrator's time and the hardware required. *Windows 7* has a built-in software application for network-based back-ups; for *Windows 7* and earlier *Windows* systems there are also various third-party applications that run across a network to automate backing up to a central server. For *Linux* systems there are a number of free network back-up and archiving applications, such as *Amanda* or *Bacula*. For *Mac OS* there is also a built-in application, *Timemachine*, which will back up across a network or to an external hard drive, and third-party applications are also available. Another advantage of network back-up software is that it can incrementally back up the contents of a desktop PC while it is in use, meaning that PC need not be left on to perform back-ups when the operator is not using the machine.

The concern with any backed-up data, particularly removable media, is the security of the information stored. With file servers it is possible to configure additional security measures, such as hard drive encryption,¹⁰ to preserve the security of the data stored on the machine.

Similar encryption options are possible with removable media. The difficulty is that encryption uses more processing power, and that in turn increases power demand overall if routinely used for servers or desktop PCs. Also,

if the encrypted media degrades or is corrupted, compressed or encrypted data is also more likely to suffer a catastrophic loss of the whole block of data rather than the corruption of one or two files within the back-up.

6.2.3. User-based back-ups

Now let's look at the computer user. The benefit of the user backing up is that, because they know which files are the most important, it is possible to target which sections of the hard drive are backed up. This can generate a much lower demand for data storage, making it a popular option. Because of this a wider range of back-up technologies are available for use. How, and how often the user backs up should integrate with the types of activity carried out on the system, the sensitivity of the data, and the frequency with which that data is changed. What's important is that those carrying out the process know how to create reliable back-ups, how to configure the programs used to create the right kind of data format, and ensure that the back-up media is tested afterwards to be certain it is usable.

Box 6.2/table 6.1 outline various storage media and their characteristics. Which is the most appropriate depends on the skills of the user, on the costs of the media, and most importantly on the scale of data to be stored. Some media can be reused to reduce costs and ecological impacts. Due to their high environmental impact, USB sticks and memory cards should be regularly reused many times for as long as possible. Although single-use optical discs cannot be reused, their benefit is that they have a long storage life. For these reasons USB sticks and memory cards are better for routine daily/weekly backing up, while optical discs are preferred for the long-term archiving of data and the offline storage of infrequently used data. While there is little hard data available, CDs, DVD and Blu-ray discs have a similar environmental footprint even though they have very different storage capacities. For the greatest efficiency, use the largest capacity disc that's suitable for the scale of back-up operation required.

Creating CDs, DVDs and Blu-ray discs requires the use of a disc mastering application.¹¹ These come as standard on all current operat-

10. Wikipedia, Disk encryption. en.wikipedia.org/wiki/Disk_encryption

11. Wikipedia, Optical disk authoring. en.wikipedia.org/wiki/Optical_disc_authoring

ing systems, although third-party applications are often used as they give more functionality. The mastering program bundles up the data files into an image of the disc to be created, then writes that image to the blank disc in a single operation. Third-party applications usually provide extra, often proprietary options to configure the format of the disc, and to enable the creation of other disc formats such as audio CDs and video DVDs.

CDs/DVDs created on Windows machines have traditionally been a problem due to the use of Microsoft's proprietary format for creating data discs - which restrict their compatibility with non-Windows systems. Since Windows Vista, users have the option of creating a *live file system* format, which is only compatible with Windows Vista/Windows 7; or a *mastered* format, which is more widely compatible with Mac and Linux machines. MacOS and Linux machines usually create discs using the ISO9660 international standard, and so are more broadly compatible across different machines and operating systems.

A problem with keeping an archive of many CDs or DVDs is that it's not possibly remember what is on every disc - and it's very difficult to handwrite the contents on the disc itself. The solution is to capture a directory tree listing of the directory and file names on the disc and store it as a text file. Then, instead of searching the actual discs one at a time looking for a file, search the text files containing the directory trees using the "find" tool of the word processor/text editor. Microsoft keeps a guide to capturing directory trees on its website.¹² For Linux and MacOS, the Unix command *tree -if 'path_to_directory' > 'file_name.txt'* will create a directory tree which can be captured as a text file.¹³ While it takes a few minutes to make a directory tree and store it, keeping a tree of each disc in the archive can cut the amount of time it takes to find a file. That's because the search process can be automated by word searching the text files to find which disc contains the data, rather than manually searching each disc to find the data.

12. Microsoft (April 2011). How to add the Print Directory feature for folders in Windows XP, in Windows Vista, in Windows 7. support.microsoft.com/kb/321379

13. See the manual page for tree command at linux.die.net/man/1/tree

6.2.4. Off-site and online storage

The risk with holding data in a single location such as a home or office is that disasters can happen. Buildings can catch fire, flood, or the equipment can be stolen. For this reason keeping data in another location is advisable, especially irreplaceable data. The simplest method is to create two back-up copies and store one of those in a different location. The difficulty is organising how the data will be transported to the other location, whether it is secure enough to hold the data, and accessible if it becomes necessary to retrieve the data.

Another option for desktop users is to manually back up their machine to a file server in a more secure location, either within the same building via the local network, or in a different building using an internet connection. This can be done in many different ways:

- A formal network archiving program, which will bundle up data and move it to a server;
- A network service, such as file transfer protocol (FTP), to allow more secure access to storage space on a local or remote server machine - for sensitive data this transfer can be made using an encrypted connection;
- A networked file system, such as shared folders on a Windows network or a networked file system on a Mac/Linux network, to share files directly between computers; or
- A more secure virtual private network (VPN),¹⁴ which allows files to be shared to other linked computers across the internet (both Windows 7, MacOS and Linux systems are supplied with the software required to configure VPN connections, allowing computer systems to securely back up to a remote server).

The problem with network back-ups is that, even for a handful of machines, this represents gigabytes of data being routinely transferred across the network - and that requires a lot of energy (we'll investigate this in unit 7). The shift to faster networks, which use comparatively more power compared to the older/slower standard, is in part driven by running more operations across local networks - such as data back-ups. While the centralised/automated backing up of PCs from a server is very simple to organ-

14. Wikipedia, Virtual private network. en.wikipedia.org/wiki/Virtual_private_network

ise, the use of the network for backing up can create a heavy drain on the network's capacity.

Seeking the least ecologically damaging route to backing up requires us to value to the data we are copying against the impact this action creates. In many cases backing up across a small wired network, using 100 megabit rather than gigabit speeds, will use less energy than a wireless connection. Backing up more than two or three gigabytes of data to an online service that keeps files live 24/7 is likely to consume more energy overall than reusable storage media.

In devising a policy for backing up it is necessary to weight these different factors to produce the least ecologically damaging option.

6.2.5. "The cloud"

The greatest movement in computing at present is the storage of data in "the cloud". Cloud computing¹⁵ has evolved with the ever greater use of mobile handsets and computers. Ten years ago this wasn't viable, but with the greater availability of broadband services, and the development of cheap high-capacity servers, storing large quantities of data online is now a viable option.

Apple's mobile devices, Google's Android operating system, and Microsoft's new Windows 8 system are tied to the use of cloud storage. Online services, such as banking or social media, also use data stored on many machines which are part of the cloud. Unlike an identifiable server, where you "know" where your data is stored, how the cloud handles and stores data isn't managed by the user - it's automatically determined by the rules that govern the cloud system. Cloud storage operators often run multiple data centres, and files can be spread across one or many of those locations depending upon which represents the most efficient way to move and hold the information concerned.

However, the important issue about "the cloud" is that it's not just a data storage system; it represents a whole business model for online commerce. If we look at the services using cloud storage, enforcing intellectual property rights in the digital domain is often

an important element of the way they operate. This is best explained by looking at Apple iTunes, or Amazon's Kindle e-book system. In order to enforce intellectual property rights in the digital world it is necessary to track the use of data. The most secure way to do that is to link the storage of data on an individual's computer to the data stored on the organisation's cloud servers. This is enabled by having to register the device(s) which use these files with the company's information systems - for example Apple's music files or Amazon's e-book. In this way services can be provided, the movement of data tracked, and the use of intellectual property policed.

While there has been great interest in the efficiency of cloud storage, there has been little debate about the enforcement of tighter intellectual property rights over digital data and the effect this has on society. Human culture has traditionally been shared, and that's been the key to the development of knowledge and learning. In contrast, lodging information in "the cloud" enforces strict boundaries on people's ability to share and re-work/re-imagine the elements which make up our culture.

One of the most debated points about the cloud is the way in which these systems log large quantities of data about individual's patterns of data use and communication. This enables all sorts of profiling activities in order to identify people's attitudes and interests. Primarily this data has commercial value to marketing and public relations agencies - and that has a whole number of negative ecological impacts given that such a large part of the internet's capacity, and increasingly direct mail, is tied up conveying sales information.¹⁶ More controversially, this same data has increasing relevance to the state security role of police and security services and by extending the powers of the state over people's data these systems could potentially be used to police the freedom of thought, expression and communication enabled by digital communications.

Finally, through various online services people are lodging more of their data online. As a result the demand for processing power and data storage is increasingly being driven by the

15. Wikipedia, Cloud computing. en.wikipedia.org/wiki/Cloud_computing

16. Taylor and Koomey (2008). Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising. imc2.com/Documents/CarbonEmissions.pdf

Box 6.3.

Basic rules for making back-ups on removable media

However you choose to back up files, there are certain basic rules to ensure that the process works well, and that the data stored is readable in the event of the loss of corruption of files:

When regularly backing up current work using reusable media, always store the back-up in a secure location – if storing sensitive information, the back-up should be stored with a level of physical security that reflects the importance of the data it contains.

If data is to be archived, and is of great value, two copies of the archive media should be created, and one of those should be stored in a different location. This ensures that the data it contains can be recovered even if the copy kept locally is lost, damaged or destroyed.

Always check the readability of back-ups after creating them

- When making CDs, DVDs and Blu-ray disks, always ensure that the disc mastering program performs a verification read after writing the disk to check its content.
- Back-up applications should automatically check that the files moved are readable, but if manually copying files to a USB stick or external drive, copy them back into an empty directory to ensure that the back-up is readable.

Mechanical external/internal hard drives and micro-drives can have their lives shortened if shaken or dropped, and must be stored in a location that is well away from strong magnetic field (electric motors, video displays, etc.). USB sticks and memory cards should also be carefully stored in secure locations as mechanical stress (e.g. being dropped or subjected to heat and cold) can break the internal electrical contacts and render them unusable.

Magnetic media, such as digital tapes, must be kept at an even temperature, protected from extremes of heat

and cold, and shielded from magnetic fields (ideally, if stored for a long period, they should be kept inside a metal container to exclude all magnetic fields).

When archiving data onto CDs/DVDs, always label the disc with an identifiable index – for example, the date the back-up was made and a few words to describe the content.

When making a large number of archive CDs/DVDs, create a listing of the directory tree on the disc and store it as a text file – these files can be searched when you want to find a stored file, rather than manually searching the discs.

Indexing the CDs/DVDs according to the date they were made, using the format year-month-day-disc number (e.g. 2012-03-12-01). This creates a unique index key to identify the disc, and keeping the discs in date order allows them to be accessed quickly when required.

Optical discs require careful storage to maximise their working life

- Store the discs in a CD/DVD folder to protect them from damage, or if making a large number of discs you can store them on the spindle containers that the blank discs are supplied in.
- Never expose the surface of the disc to strong light for long periods of time, particularly sunlight or near to bright fluorescent lights (ultra-violet light damages the polymers in the disc).
- Keep the discs in an environment that has an even temperature, and protect them from extremes of heat and cold.
- Handle the disc by its edge and centre hole, and avoid touching the surface of the disc as the grease/dirt from fingertips encourages dust and fine grit to stick to the surface.

needs of large data centres rather than individual computers. The ecological advantages of the cloud¹⁷ are predicated on the basis that this new demand for computing power can be managed more efficiently in a large data centre than on a large number of small systems. The difficulty is that by keeping data in the cloud you are putting all your trust in the availability of the online services. If the service fails (e.g. the problems with the Blackberry service in 2011),¹⁸ or the user's account is compromised or blocked, or the company providing the service collapses, access to data in the cloud can be lost.

To have secure access to our most valued information at all times, it is necessary to keep back-ups where they can be accessed locally. The difficulty for the cloud computing model is that if users keep back-ups of their information on their own machine for the purposes of security, many of the arguments for the ecological efficiencies of the cloud system disappear.

6.2.6. Manually transporting data – the “Sneakernet”

In the early days of personal computing dial-up internet connections were very slow. While today's broadband internet connections run at a few hundred to a thousand kilo-bits per second, early dial-up modems ran at 0.3 kilo-bits per second. For this reason the fastest way to move data wasn't via a network, it was physically carrying floppy disks in a bag. As a metaphor for the idea of moving data on foot, this method of data transfer became known as the “Sneakernet”.¹⁹

While we may focus on the use of the internet, it is arguable that a person walking with a bag full of data can move information faster than many electronic networking technologies. That's because while the internet is fast for everyday small files and email transfers, when considering the movement of giga-bytes of data the throughput of the network can be

very slow. The different ecological impacts of moving data via electronic networks, or manually carrying or sending storage media via the postal system or a courier service, are also significant.

There are various figures for how much energy it takes to move a gigabyte of data across the internet.²⁰ A general figure is somewhere around seven kilowatt-hours per gigabyte,²¹ and, with power generation producing around 600g of carbon dioxide per kilowatt of power, that entails the emissions of around 4.2 kilos of CO₂ per gigabyte of data. An optical disc with a plastic jewel case takes around 16 kilowatt-hours of energy to produce, releasing about a kilo of CO₂ as a result of its production; and while there are no detailed figures, a DDS tape might take three times that because it is more complex product containing mixed materials.

Figure 6.3 shows a comparison of moving data via the internet (shown as hatched bars) and using storage media via Sneakernet (shaded dots). For the given storage capacity of each media type the impacts of moving that data electronically via the internet or creating the back-up and then transporting it are calculated. There are studies of the ecological impacts of the postal service,²² and conservatively these figures have been multiplied by a factor of five to take account of the higher package weight. The costs of moving data online have been assumed to be £1 per gigabyte, while the cost of sending an optical disc/DDS tape via post, and purchasing the blank media, is also calculated for comparison. What the results show is that:

- Producing an optical disc and a plastic jewel case to protect it takes around 16kW-h of energy and emits a kilo of CO₂; however, as moving data via the internet is also energy intensive, sending a DVD via the post has a slightly lower impact than sending the same data across the internet.

17.WSP/Accenture (2010). Cloud Computing and Sustainability: The Environmental Benefits of Moving to the Cloud. www.accenture.com/SiteCollectionDocuments/PDF/Accenture_Sustainability_Cloud_Computing_TheEnvironmentalBenefitsofMovingtotheCloud.pdf

18.ZDNet (October 2011). BlackBerry issues statement over downed services. www.zdnet.com/blog/btl/blackberry-issues-statement-over-downed-services/60450

19.Wikipedia, Sneakernet. en.wikipedia.org/wiki/Sneakernet

20.Koomey et. al. (August 2009). The energy and climate change impacts of different music delivery methods. download.intel.com/pressroom/pdf/CDsvsdownloadsrelease.pdf

21 .Taylor and Koomey (2008). Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising. imc2.com/Documents/CarbonEmissions.pdf

22.Pitney Bowes Inc. (2008). The Environmental Impact of Mail: A Baseline. www.pb.com/bv70/en_US/extranet/landingpages/Environ_Impact_Mail_Web.pdf


- Moving less than two or three gigabytes of data via the internet is more efficient in terms of energy and carbon emissions than posting an optical disc however, that's primarily because of the high energy and carbon density of the disc. If reusable media were used, such as a rewritable CD or DVD, that would significantly reduce the impacts. Even if only reused four times, a CD-RW disc would then be more efficient than the internet.
- Taking the cost of data at £1 per gigabyte, if using commercial postal services it is cheaper to send up to two to three gigabytes via a network than via an optical disc.
- While for everyday small-scale data movements (e.g. web pages download or email) this analysis doesn't have a great relevance, where these figures are very significant is the use of networks for backing up data. For example, Google has recently launched a large-capacity cloud storage service.²³ If used for backing up, routinely moving that quantity of data via the internet would take more power and emit more carbon than storing a DVD at a location where it could be held securely.

The implication of this is that backing up large quantities of data off-site has the lowest impact if carried out using storage media, such as optical discs, rather than using electronic networks. The difficulty is organising how the data will be transported to the other location, whether it is secure enough to hold the data, and accessible if it becomes necessary to retrieve the data. How these issues are dealt with, and solutions found, would have to be tackled as part of the process for planning and implementing a back-up procedure.

When data CDs first became widely available, but the use of computer networks was still in its infancy, the publication of digital content often used CD-based distribution as a means of moving large quantities of information in a structured format. Today, with the advent of high-speed networks, disc-based distribution of publications is going out of fashion in favour of network-based/web distribution of content – and more recently the movement of large blocks of data using peer-to-peer file

Figure 6.3.

Comparison of Internet versus Sneakernet impacts



sharing.²⁴ The relative ecological impacts of networked versus stored data distribution indicate that we may need to re-evaluate how information distribution is planned. As general rule, where the content is relatively static or is

23.BBC News (April 2012). Google Drive to offer free storage in the cloud. www.bbc.co.uk/news/technology-17831725

24.Wikipedia, Peer-to-peer file sharing. en.wikipedia.org/wiki/Peer-to-peer_file_sharing

not regularly updated, the impacts of distribution via optical discs may be lower than holding that information live online. For example, the data required to install an operating system today fills a DVD, and arguably the ecological impacts of distribution on DVD would be marginal-

ly better than via the internet and yet the trend today is for greater network-based software distribution to install and maintain operating systems, or to run software remotely within the a cloud system.²⁵

6.3. CONFIGURING THE DESKTOP TO IMPROVE DATA SECURITY AND USABILITY

Another significant area of everyday data loss is user error or system crashes which affect the work/program in current use. While many programs are getting better at recovering data following a system crash, it is possible to improve the way we use programs to reduce the likelihood of losing data.

Firstly, many programs – particularly word processor/office applications and text editors – can routinely create a back-up copy. This feature is enabled from the user configuration options for the program. While some will create a back-up file with a slightly modified file name, others will set up a specific directory to store back-up copies of the files/documents being worked on. Each time you save a copy of the file it stores the previous version as a back-up. If a crash corrupts the original file, or the contents of the file are accidentally lost or erased, or if saving the file erases some valuable data that had existed in the previous version of the file, it is possible to revert to the back-up copy and recover lost data.

Secondly, it's always a good idea to regularly save work. Some programs allow you to configure an auto-save option to save the document at a specified time interval. Rather than remembering to save, the program will take care of that for you. The difficulty is that sometimes, if you make a mistake, the auto-save option can over-write the back-up version of the file being worked on. The trick is to set the auto-save interval to be short

enough that you will not lose data if the system crashes, but not so short that it repeatedly stores the current document, which restricts the ability to reuse the previously saved copy of the document if data is lost during editing.

Thirdly, most programs allow you to “undo” the editing or processing of data. The more steps that it is possible to undo, the greater the chance that mistakes can be corrected. The main restriction on the undo facility is the amount of memory it takes up. Your options are limited if using a machine with limited memory space. If the machine has a lot of memory then it is possible, from the programs configuration options, to increase the number of undo steps. How many steps you set depends upon your preference – 30 to 40 steps is probably sufficient for most applications.

Finally, most operating systems use a “waste basket” to store files which have been removed from the hard drive. This means that you have to occasionally go in and empty the waste. While it is tempting to simply “empty trash”, or use a delete command to directly remove files from the hard drive, the waste basket feature serves a very useful purpose – *it stops the user deleting files accidentally*. Also, do not blindly use the “empty trash” command – always check what's in the waste bin before erasing the contents.

25.Wikipedia, Software as a service. en.wikipedia.org/wiki/Software_as_a_service

Box 6.4.

Information storage check-list

The more common and easily accessible a resource the less we need to back it up, the more unique or expensive a resource the greater the need to create a back-up copy – value the files on the computer using this general principle.

Separate the information held on the machine into different directories – one directory tree for current work, another for already archived files, and another for low value/superfluous files – to simplify the process of backing up and reduce the amounts of data requiring storage.

Identify roles and responsibilities for backing-up data, and agree policies and procedures for regularly backing-up data:

- System administrators/managers should be responsible for backing up at the system level, and co-ordinating the use of network-based software for backing up to a central server.
- Computer users should be responsible for backing up current work and other important data using removable storage.

Backing up with the least ecological impact requires that we match the type of data being stored with the characteristics of the storage media used:

- Data that is being archived for long periods, and other data which changes infrequently, should be backed up

on optical discs as these have the lowest ecological impact and a long storage life.

- Routine daily/weekly back-ups should be made using reusable media, such as USB sticks, memory cards or removable hard drives.

When creating archives on optical disc, create a text file containing a directory listing of the disc's contents – searching this disc using a text editor (or integrating it into a desktop search system) takes less time and energy than physically searching each disc.

Using the internet for backing up has a high impact when moving many gigabytes of data. While cloud computing is becoming popular, always consider the alternatives to overcome the data security and ecological drawbacks of using large amounts of online data storage.

When moving large quantities of data between two systems, posting optical discs/digital tapes can have a lower ecological impact and financial cost than transfer over a network.

Configure the features of desktop applications to preserve data: Always create back-up copies of files being worked on; auto-save at regular intervals; set/increase the number of “undo” operations; and always use the waste basket rather than directly deleting files.

LOCAL AND GLOBAL NETWORKS

Early personal computers had a number of advantages over previous information technologies, especially in fields such as word processing and the use of spreadsheets. However, it was when personal computers were first networked together, so that data could be shared between users in the same building or across the globe, that they began to reshape the world of information and communications. Today it is difficult to conceive of a desktop or laptop computer without a network connection, and over the last few years the general expectation is that this connection will be a high-speed broadband link rather than the slow dial-up connections which preceded them.

While the use of electronic networks has developed computers into a significant new global technology, this has not come without an adverse impact. As communications have increased, so the amount of data flowing between users around the globe has increased, leading to a leap in the amount of data now routinely stored on personal computers; as more people connect to the network through mobile devices, more data is stored in large online systems to enable people to communicate on the move – and this is now maturing as a large cloud storage network; and while the efficiency of online technology has increased, the increase in data traffic¹ more than offsets that improvement, and so the energy and resources used to support the global communications network are still increasing.²

As users of this global network there are various measures we can take to improve the ecological footprint of our network usage. Some of these steps are simple, while others are a matter of adopting different working methods. Some, such as the way we design websites, are more problematic because they entail redefining the way in which we present ourselves to the world online – breaking with the fashion for more elaborate web interfaces in order to slim down the amounts of data transacted.

1. Minnesota Internet Traffic Studies (accessed June 2012), Internet Growth Trends & Moore's Law, www.dtc.umn.edu/mints/igrowth.html

2. Raghavan and Ma (2011). The Energy and Emergy of the Internet. www.cs.berkeley.edu/~jtma/papers/emergy-hotnets2011.pdf

7.1. LOCAL NETWORKS

Local networks³ connect computer users in a building, or people at home to their broadband router. Early networks used a number of different systems, but over the last two decades network connections have become standardised around different technologies that have been adapted to work in different spatial environments:

- **Ethernet**⁴ is a wire-based network which physically links machines together. In terms of the number of machines connected it is the most widely used networking standard. ethernet uses multi-core cables to carry data between a local hub and the machines connected to it. In turn that hub can be connected to other hubs and machines, or using a router it can be connected to the global internet. Ethernet has historically used different speeds, the most common today being fast ethernet (100 megahertz, MHz) and gigabit ethernet (1,000MHz). Faster speeds are now in development, and while some use wire-based connections, increasingly fibre-optic cables are used to connect devices together.
- **Wi-Fi**⁵ is a radio-based network operating around 2.4 gigahertz. The frequency band is divided into a number of channels and machines switch between them in order to find an unused space to establish a wireless link. With the development of more portable equipment, and changes in the way workplaces are organised to encourage more mobile working, Wi-Fi has become a popular means of connecting to a local network because it is not necessary to find a cable or network socket to connect to. It has become popular in the home as it allows computers to be connected without the inconvenience of having to set-up a wired network.
- **Bluetooth**⁶ is also a radio-based network, although it only functions over very short distances. It became popular as a result of its adoption in mobile phones and wireless mobile headsets,

but is increasingly used to link computers and peripherals (such as the mouse, keyboard and printer). Recently it's become commonly used in games consoles to connect controllers and motion capture devices to the main console.

Each networking technology has a different ecological footprint. For example, Wi-Fi uses more energy to function than ethernet, sometimes twice as much, in order to communicate between the base station and the mobile device. That's because the energy of the radio signal falls away quite quickly the further the user is from the base station, and so both the base station and the device must use a high signal strength to maintain the communications link. On a typical laptop, while the video display and processor use the most power, the next most significant power drain is likely to be the Wi-Fi interface. This assumes that the Wi-Fi interface is used – if not then it should be disabled in the hardware/BIOS set-up options in order to conserve power (if allowed to routinely search for wireless connections the interface will create a large power drain as the device will search using full transmitter power). Most Wi-Fi routers and other hardware allow you to configure power saving options from their configuration interface. For example, where wireless devices are used in a small area it may be possible to switch the base station to transmit using its lowest power setting. Operating systems also give some scope to control the use of wireless links.⁷ In contrast to Wi-Fi, bluetooth uses less power because it involves short distance links.

The next significant factor in local network power consumption is the speed of communication. As a general rule of thumb, doubling the processing speed of digital electronics can increase power consumption by up to four times. This has relevance to the increasing use of gigabit ethernet. For most online browsing and video streaming fast (100MHz) Ethernet has enough capacity to handle routine data movement. Higher speeds become more relevant when very large files are routinely moved between computers (for example, backing-up machines across a local network),

3. Wikipedia: 'Local area network'. en.wikipedia.org/wiki/Local_area_network

4. Wikipedia: 'Ethernet'. en.wikipedia.org/wiki/Ethernet

5. Wikipedia: 'Wi-Fi'. en.wikipedia.org/wiki/Wi-Fi

6. Wikipedia, 'Bluetooth'. en.wikipedia.org/wiki/Bluetooth

7. For example, see LessWatts, 'Tips & Tricks: Wi-Fi/Wireless'. www.lesswatts.org/tips/wireless.php

or for the connections between the ethernet hubs which link groups of computers and the network's data servers or high-speed internet connection. Depending upon the capability of the hardware and the operating system, it is possible to switch gigabit ethernet to work at the slower fast ethernet speed and conserve power⁸.

Another major factor in running a network is the power used by the network electronics to provide connectivity. All devices use power, and when not in use they will still draw a minimum amount of energy while they wait to move packets of data between machines. For that reason it is better to have the smallest number of ethernet switches⁹ or Wi-Fi base stations and use them to their maximum capacity – only adding new devices when the network experiences congestion. There is no easy advice to give on this point. How a network is physically laid out depends upon both the number of machines involved and the layout of the space they occupy. Theoretically the maximum length of an ethernet cable is 100 metres, although taking

the cable through twists and turns around the edge of large rooms can quickly reduce the practical distance between the network switch and the computer. On small networks power can be saved by using a combined broadband/ISDN router¹⁰ and multi-port ethernet hub rather than using separate devices.

As part of the initiatives to control total and standby power consumption, ethernet/Wi-Fi hardware now incorporate power saving functions.¹¹ Unfortunately these only save significant amounts of power on larger ethernet switches. For smaller devices, for example in the home or a small office, the greatest power savings are achieved by physically turning off the power supply when the computers/network are not in use. As noted earlier in the section on hardware, it's a good idea to connect computers to a surge-protected socket with a build-in switch. In a small office/home the broadband router can also be connected to this type of central power socket, so that when the computer is shut down the router powers down too.

7.2. USING NETWORKS

Using networks efficiently has many benefits. The energy used to transmit data is directly proportional to the size of the files being transmitted – therefore using various means to minimise the amount of data you need to send will reduce the energy consumed by your use of electronic networks. And a result, by minimising the data moved, we can also reduce the amounts of data which clog the file systems of our computers.

7.2.1. Email and attachments

Email is still the backbone of the way people communicate and network online. Email is simple to use, and has the benefit that files can be

attached to the email and easily sent to one or many recipients. That ease of sending files is also one of the greatest problems with email. If the attachment has not been formatted for efficient delivery it will use an excessive amount of bandwidth in transmission, and if saved it will occupy a large amount of storage capacity.

The simplest way to reduce the size of the files sent via email is to use a data compression¹² program. This converts the data in text files and other raw data into a file which occupies a fraction of the space of the original file.

8. For example, see LessWatts, 'Tips & Tricks: Ethernet'. www.lesswatts.org/tips/ethernet.php

9. Wikipedia, 'Network switch'. en.wikipedia.org/wiki/Network_switch

10. Wikipedia, 'Residential gateway'. en.wikipedia.org/wiki/Residential_gateway

11. Wikipedia, 'Energy-Efficient Ethernet'. en.wikipedia.org/wiki/Energy-Efficient_Ethernet

12. Harris, Tom (accessed June 2012). How File Compression Works, How Stuff Works. www.howstuffworks.com/file-compression.htm

Box 7.1.

Reducing/scaling digital media files

How to edit and reformat digital media is a highly complex subject – certainly too complex to explain here. Even so, there are certain options in the editing of digital media which influence the size of files produced, reducing the energy used when they are sent over a network.

Image files

There are three factors which influence the quality of an image file: The size of the image, in pixels; the number of colours used to display the image; and the file format the image is stored as, specifically whether that format is a lossy or lossless format.

The size of the image file has a significant effect on the file size. The more pixels, the more data is required to store the image. Where the size of the file is greater than the size of the screen upon which it is displayed the image will automatically be scaled down to fit in which case it would require less data to store the image if it is scaled to a usable size before transmission. Older computer monitors had a resolution of 800 by 600 pixels, while newer ones can be anything from 1024 to 1920 by 768 to 1080. In deciding what size to scale the image to, consideration should be given to the likely resolution of the monitor it will be viewed upon.

The greater the number of colours in an image the more data has to be stored for each pixel. Image editing applications can reduce the number of colours in the image. The effect of this makes the image look progressively more grainy, although the use of filters to “dither” the colour distribution can help to reduce this slightly. Some file formats (e.g. JPEG) allow a large number of colours to be used while others only allow a specific number of colours (e.g. GIF).

The JPEG format is useful because it uses a “lossy” compression system to reduce the amount of data used to store colour information. This has the effect of making the image look progressively more blurry as the quality factor of the image is reduced and that loss of quality is permanent if you try and re-edit the image. Other formats, such as PNG, are lossless, and so produce the same quality no matter how many times they are edited – although the cost of that is a larger file size.

Sound/audio files

Sound is recorded as raw audio data – this is what is found in WMV files, and so WMV is the least efficient

way of sending audio data. Other file formats encode the data, and while this reduces the file size it progressively results in lower quality sound reproduction. There are three factors which have a great influence on file size: The sample rate; the number of channels; and whether the “codec” (the standard used to encode/decode the audio data) is lossy or lossless.

The sample rate, measured in kilohertz (kHz), is the number of sound measurements taken each second – the higher the sample rate, the clearer and more “deep” the audio reproduction will sound. Most audio codecs require specific sample rates to be selected. 8kHz sounds like a bad phone line conversation, 22kHz sounds like an FM radio, while 48kHz has the quality of a CD. For speech a low sample rate is possible; for music higher sample rates are required. The number of channels increases the file size; a stereo (two channels) file has almost twice the data as a mono file. To reduce the channels use an application to down sample from stereo to mono.

The file format chosen also has an effect on the file size. FLAC files are a lossless – format meaning that they do not lose their quality if edited. In contrast MP3 files lose data when the audio information is encoded, and so sound quality can be lost with each re-encoding. What influences file size the most is the bit rate of encoding (like sample rates it is measured in kilohertz). A 32kHz bit rate sounds like a poor phone line, while 320kHz sounds very clear and is often used to encode CD-quality music. Alternately data can be encoded using a variable bit rate (VBR) which varies the rate according to the complexity of the sound, producing a slightly more space-efficient encoding.

Video files

Video files are the most difficult media to edit as the standards used are more complex, and are often linked to proprietary codecs which require specialised application programs to edit them.

For video the frame rate is equivalent to the sample rate it represents the number of pictures displayed every second. The standard is 25, and while reducing this can make the video progressively more jumpy it has a large effect on the file size. Scaling down the size of the video image also has an effect on the file size. The audio track usually represents about a fifth to a tenth of the total file data, and while it can be edited it doesn't have a great effect on the overall file size.

At the simplest level this is done by replacing repeated sequences of numbers with a key which represents that sequence. For example, if we think of a large file which contains the text of a book, we could replace long words with numbers which represent those words. Some email applications will give the option of compressing attachments before they are sent. Alternately you can use a free or proprietary application to compress the file.

The problem with many file formats in common use today – such as *OpenOffice* or Adobe *Acrobat* – is that they are already compressed as part of their formatting. Therefore simply compressing the file will not reduce its size significantly. In these cases you have to change the way these files are created by the original application. Adobe Acrobat (and the “Export PDF” functions within office applications) allow you to reduce the image resolution of the document and compress any images which it contains (the text of Acrobat files is already compressed). The proprietary Adobe Acrobat application offers many more options to reduce the size of the file¹³ which can reduce the size of PDF files significantly. Many applications offer additional plug-ins or extensions to add extra functions which optimise file sizes. For example, OpenOffice Impress has a downloadable extension called *Presentation Minimizer*.¹⁴ This adjusts the formatting of fonts, images and animations within presentations to reduce the size of the file generated.

The greatest challenge in minimising the data transaction from email attachments relates to digital media files – video, sound and image data (see Box 7.1). When using free software this is a simpler task because all the programs required to reformat open media files are usually included with the operating system, or can be freely downloaded. Of the three, video is the most difficult media to work with. Often reducing the size of video files is a matter of trial and error in each case. Using the video editing application(s) available for your operating system the file is first imported, and then various options are tried to reduce the frame rate, size of the video frame and the quality of the audio track. In contrast, the process of re-

ducing the size of image and audio files is simpler and has more consistent results.

Audio and video files are processed using “codecs”¹⁵ – encoding and decoding standards which convert the compressed data into audio/video data ready for reproduction, and vice versa. The main consideration about processing digital media for transmission is quality; file size and the quality of the video, sound or image reproduction are directly related. Often it's not simply shrinking the size of the video/image which has the greatest impact on file size. Reducing the number of colours used to display the visual data, or reducing the sample rate of audio files or the frame rate of video files, will reduce the file size without necessarily affecting the quality of reproduction.

For example, one common problem is that the popularity of digital cameras/camera phones has led to people routinely emailing multi-megabyte images files. If they had taken a few minutes to scale the image to a fraction of its original size it could shrink the file significantly. As discussed earlier in Box 5.3, each one mega-pixel of a digital camera is roughly equivalent to 100 dots-per-inch (DPI) when displayed/printed. Therefore a six mega-pixel camera is roughly 600DPI, which is four to six times higher than the resolution of most video displays. As a result that image could be scaled to a third of its former size without significantly affecting the picture quality.

Last but not least, over the last few years the original standard of plain text email has been slowly replaced by HTML email.¹⁶ Rather like web pages, these use formatting codes to specify the font, size, and text decoration of the text, as well as adding features such as lists and tables. The difficulty is that all this additional formatting data adds to the size of the transmission – and when formatting is configured/used badly, or only a very small email is sent, there can be up to three or four times more formatting data within the body of the email than data (texts and images) that humans actually see. Although HTML formatting will not be significant for small volumes of email, when sending to popular email lists, or emailing to long distribution lists, the additional formatting can represent a large quantity of data. Many email applications now select HTML formatted email by default. Instead it would be more efficient to

13. Adobe (2004). Adobe Acrobat 7: Reduce PDF file size. www.adobe.com/designcenter-archive/acrobat/articles/acr7optimize/acr7optimize.pdf

14. OpenOffice (accessed June 2012). Extensions: Presentation Minimizer. extensions.openoffice.org/en/project/PresentationMinimizer

15. Wikipedia, 'Codec'. en.wikipedia.org/wiki/Codec

16. Wikipedia, 'HTML email'. en.wikipedia.org/wiki/HTML_email

use text-only email by default, and send HTML formatted email only when the use of formatting has a beneficial effect on the clarity/presentation of the data being communicated.

Reducing the impact of our email use is quite simple to achieve, provided that we are mindful of the way we construct the message before we send it. Controlling our use of email and email attachments, quite apart from the issue of their environmental impact, is also an issue of digital equity and inclusiveness. In less developed states, particularly in South America and Africa, which have lower regional connectivity, accessing the internet at any appreciable speed can be difficult. Even in many developed nations, residents of more remote rural areas cannot connect to high speed broadband. If those with high-speed broadband routinely email large data files across the web that can create problems for those who do not have the capacity to download those files at high speed. For those who are using metered access, having to pay for the amount of data they transact, downloading unnecessarily large files unfairly penalises them when the same *quality* of communication could have been achieved with a smaller file size.

7.2.2. Web browsing

Over recent years, especially since the widespread adoption of broadband connections, the size of web pages has grown significantly. This is in part the result of more graphically complex pages, and the increasing role that advertising and content tracking have in the design of many websites and blogs. Especially for sites with animated advertising, and which require large quantities of scripting code¹⁷ to control the behaviour of the page, the ecological impacts are not just related to the amount of data downloaded – executing local scripting can also require a large amount of the machine's processing power. For example, some newspaper sites now required a megabyte or more of data to be downloaded for each page; and on some web pages (especially where poorly designed or incompatible Javascript controls are used) viewing the page will increase the load on the processor, increasing the machine's power consumption.

The need to download data can to some extent be controlled through the configuration of the web browser. Some aspects of minimising

the downloading of data are part of the browser's configuration, while others require plug-ins or additional software:

- **Web cache** – Web browsers use a cache¹⁸ to hold copies of commonly downloaded files. When a page is requested if the cache already has a copy it checks if the copy on the site has been changed since it was last downloaded. If the content is unchanged the local copy is used and so it need not be repeatedly downloaded. If routinely browsing advertising/graphically rich websites it is a good idea to increase the size of the cache to take account of the large script and image/flash files used by these sites. For example, the Firefox browser defaults to a limit of 50 megabytes of cached files; extending this to around 100 to 120 megabytes will allow more of those repeatedly displayed adverts and the code/data files which go with them to be cached and will help to speed loading of the page and reduce the amount of data downloaded.
- **Pop-ups** – A common feature on advertising-driven sites is that certain actions by the user, such as clicking on buttons or lists, will cause another browser page to “pop-up.”¹⁹ This page invariably contains advertising or other unwanted information, and can launch additional code within the browser than can tie up the processing power of the system and even compromise security. Most browsers include an option to block pop-ups, which is enabled from the browser's configuration menu. Quite apart from avoiding additional data downloads, blocking pop-ups will reduce the load on the processor and avoid loading more data into the system memory each time a new browser window is opened.
- **Advertising filters** – Even with a pop-up blocking enabled, many sites now use dynamic HTML to display adverts or animated content which float²⁰ over the browser window. As these stay within the control of the existing window this circumvents control by pop-up blocking. Instead what is required is an advertising filter²¹ plug-

18. Wikipedia, 'Web cache'. en.wikipedia.org/wiki/Web_cache

19. Wikipedia, 'Pop-up ad'. en.wikipedia.org/wiki/Pop-up_ad

20. Wikipedia, 'Hover ads'. en.wikipedia.org/wiki/Hover_ad

21. Wikipedia, 'Advertising filtering'. en.wikipedia.org/wiki/Ad_filtering

17. Wikipedia, 'Client-side scripting'. en.wikipedia.org/wiki/Client-side_scripting

in for the browser, such as *Ad Block*,²² which monitors the use of dynamic HTML to prevent frames hovering over the browser window. As with pop-up blocking, this saves processor power and memory.

- **Disabling Java/Javascript and flash** – The dynamic content of web pages, especially that related to advertising, can be a drain on the power of older hardware. If this presents a problem the simplest measure is to disable Java and Javascript, and perhaps de-install Adobe flash player. This difficulty with this approach is that disabling Javascript will prevent many websites displaying their content properly. Few websites now maintain a static or “text only” access option which allows pages to be viewed without their active components – which can be a problem not only for older/lower power systems but also for those using Braille/text-to-speech screen readers.²³

7.2.3. Proxy servers

So far we have looked at the steps which can be taken to minimise the impacts of world-wide web use from the web browser. For larger organisations many of the above features can be enabled with a web proxy server.²⁴ Users of the network access the web through the proxy, and the rules defined

for serving pages will block and filter certain types of content as well as blocking selected websites. Proxy servers are becoming more popular on small and home networks too as they are a simple way to implement content filtering as part of “parental control” systems.²⁵

The greatest benefit of proxy servers is that they help to control the amount of data downloaded by users of the network. As popular content can be served from the local network they can also speed up access to information. Reducing the amount of data downloaded has a positive ecological impact, and where data download is metered it can also reduce costs. Recent research suggests that a proxy server might reduce the quantity of data downloaded by over 20%.²⁶

More generally, proxy servers are becoming an important part of network and internet security. Due to their capacity to filter connections made between machines and the outside world, the proxy can monitor use of the network and block connections to sites associated with malware and other security problem. For example, Microsoft's *Forefront Threat Management Gateway*²⁷ acts as a router, firewall and proxy server, integrating network security and control functions into one system. On Linux systems one of the most popular web proxy servers is Squid,²⁸ which can filter network traffic as well as caching regularly used files.

22.Adblock Plus (accessed June 2012). adblockplus.org

23.Wikipedia, 'Screen reader'. en.wikipedia.org/wiki/Screen_reader

24.Wikipedia, 'Proxy server'. en.wikipedia.org/wiki/Proxy_server

25.Wikipedia, 'Parental control'. en.wikipedia.org/wiki/Parental_control

26.Olatunde et. al. (2008). Proxy Server Experiment and the Changing Nature of the Web. www.ece.iit.edu/~tricha/papers/04554305.pdf

27.Microsoft (accessed June 2012).Forefront Threat Management Gateway 2010. www.microsoft.com/en-us/server-cloud/forefront/threat-management-gateway.aspx

28.Squid, www.squid-cache.org

Box 7.2.

A case study: The “efficiency” of web content

There is no standard definition of the “efficiency” of a website. Increasingly site design is about the visual appearance and usability of the user interface design, not the impact of the data streams generated when the site is used. With the wider introduction of broadband many websites have allowed the scale of data transmission to expand without any perceptible check on the impacts this has. Although there has been much discussion about the bloat of software, as yet there is little information available on the bloat of web content, and the impacts of this bloat on the increasing ecological footprint of information networks.

As part of efforts to produce a lower impact website, in 2011 the Free Range Network conducted research on the websites of UK campaign groups, political parties and new media. Sixteen websites, plus the Free Range Network's new site, were selected for study. On each site the ten most popular web pages were downloaded. In most cases a web page is not a single file; it is made up of a central HTML page, and associated with that are separate files to provide graphical images, control and formatting data. Together the 170 web pages downloaded amounted to over 10,000 individual files comprising a total of 160 megabytes of data. The content was sorted to produce a statistical snapshot of the 170 web pages, and while the sample is not large enough to produce statistically rigorous results, the analysis produced some startling facts:

- A fifth of the 170 web pages downloaded had 100 or more files associated with them – as a general trend those with a higher advertising content (primarily the new media sites) had the highest file counts;
- Measuring the size of the page and its associated files, a third of all the pages had a size of more than a mega-byte, and 2% were over three megabytes;
- Sorting the web pages by size, there is a general trend for larger pages to be dominated by Javascript code

files rather than text or graphical data – smaller pages were proportionately made up of more graphical data; and

- Isolating the types of information contained in each page, text and graphical data made up a third of the average page – the rest was formatting (style sheet) and control (flash/Javascript) data.

Using the knowledge gained from the study of other sites, the Free Range Network redesigned their site using a minimalist standard: Rather than use a database-driven dynamically created approach, a static page design was used; instead of using automated tools to generate page content, simple HTML editing tools allowed pages to be created with a minimal amount of HTML tags within them; control and formatting data were kept to a minimum; graphical data was processed to reduce its size, and a design scheme was chosen to minimise the need for overly complex or large graphics.

At the end of the site redesign, the average file size across the 10 pages from the Free Range Network's site was 169 kilo-bytes (about a sixth of a mega-byte), and on average each page had just eight graphic, style or Javascript files associated with it. As a result of the changes the amount of server space the website occupied decreased by 20%. When the new site design was evaluated alongside the results from the other sixteen, for most of the indicators used to assess “efficiency” the Free Range site had the smallest data transaction for an average web page.

The Free Range Network's efforts were essentially a scoping study to explore the issues related to the bloat of web content, and the type of steps which could be taken to address it as part of site design. This project is ongoing, and it is hoped that this snapshot of the influence of “data bloat” on web use, and its ecological footprint, will spur a more general debate and research on ecological web design strategies.

7.3. WEBSITES AND CONTENT

Although email and file sharing/download-ing make up a large part of internet traffic, a significant quantity of the impacts are related to everyday web browsing. While the impacts of the world-wide web can be controlled to a small extent through the configuration of the browser or use of a proxy server, the greatest savings on the impacts of the web can be made by the design choices which guide the development of sites. It's not so much a technological issue, or the type of content created, *it's all about design*. The operators of websites and services must deliberately set out to create a site that uses the least possible resources when it is accessed by its users.

The IT industry has been able to grow significantly over the last thirty years due, in a large part, to the increasing power and processing capacity of the equipment involved. It's really easy to grow your industry when the tools of the trade double their power and halve their utilisation costs every 18 months or so.²⁹ As a result the industry need not put a lot of effort into being more productive – *getting a greater output from your existing body of production resources* – when the power of IT systems is growing exponentially.

This is the root of the “IT productivity paradox” paradox.³⁰ It is not necessary to work hard to increase productivity if there is no incentive to create efficient resources; the rising power, capacity and functionality of technology is able to make up for the lack of attention to the efficiency of design. This, of course, is where the problem of bloat arises.³¹ If computers become more powerful there is no incentive to improve the quality of the code use to make them function; and for web bloat, if network speeds are rising there is no need to be more creative with site/information design. For example, when

mobile phone operators paid billions for their licences in the 1990s they had an incentive to get as much capacity out of their networks as possible – and through new transmission protocols they achieved a higher capacity than initially expected.³² In contrast, the designers of web systems have no such pressures on the quality and efficiency of their work.

Irrespective of whether we access the web or not, we're all paying for bloat. Not just in the extra money we have to pay to download and manipulate all that data – we're paying for it ecologically. Transferring greater quantities of data requires the internet industry to buy higher capacity hardware; that in turn uses up the finite stocks of rare metals, and generates toxic waste streams from manufacturing new machines and disposing of old equipment.³³ Making all that equipment also uses a large quantity of energy, as does running it – which also contributes to our depletion of finite energy resources, the production of pollution and climate change.³⁴

Unfortunately, there is no simple route to achieving greater efficiency in web design. In 2010, the Free Range Network undertook a small-scale study to evaluate the design-related impacts of web use³⁵ (see summary in Box 7.2). The purpose of the study was to highlight strategies to redesign their own website in order to make web pages more efficient when accessed – reducing data downloads and the amount of processing power required to display the page. A general problem is that

29. Wikipedia, 'Moore's Law'. en.wikipedia.org/wiki/Moore%27s_Law

30. Willcocks and Lester (1996). Beyond the IT Productivity Paradox, European Management Journal, vol.14 no.3 pp.279-290, June 1996. dx.doi.org/10.1016/0263-2373(96)00007-2

31. Wikipedia, 'Software bloat'. en.wikipedia.org/wiki/Software_bloat

32. Faist et. al. (2005). Life Cycle Assessment of the Mobile Communication System UMTS Towards Eco-efficient Systems. www.fraw.org.uk/files/economics/faist_2005.pdf

33. APC (2010). Global Information Society Watch 2010: ICTs and Environmental Sustainability. www.giswatch.org/en/2010

34. Pelley, Janet (2009). Consumer culture keeps carbon emissions high, Journal of Environmental Science and Technology, vol.43 no.16 p.6116. pubs.acs.org/doi/pdf/10.1021/es901774u

35. As yet unpublished, a general summary of the findings is available Mobbs, Paul (2011). Web design, environmental impact and the much ignored ecological efficiency of the Internet, Ecolonomics No.12, May 2011. www.fraw.org.uk/mei/ecolonomics/01/ecolonomics-012-20110523.pdf

the use of content creation programs encourages inefficient site design; those creating content do not understand the factors which cause bloat in online content, and so cannot format their information in a more efficient manner. The solution to this is to seek more direct means of content creation, which allow for the optimisation of the content generated. This would require a greater level of training and knowledge of web design, and the methods which can minimise the impacts of online content (such as the reformatting of files to reduce their size, outlined above in relation to email attachments). At present this is a novel area for research, and while there is a large effort devoted to search engine optimisation, or the inclusion of advertising and web analytics as part of site design, there is as yet little demand to optimise web content to reduce its download size and demand for power when displayed.

Many of the rules which apply to the sending of email attachments also apply to the creation of websites/web pages. The content created should be optimised to provide the level of design or reproduction quality for the

application it is intended to serve. Optimising PDF and images to reduce their size can save storage space, reducing the costs/impacts of web hosting, as well as reducing the scale of data downloads. Reducing the energy required to generate and serve content is directly related to the way the data is stored. For example, static web content does not require significant processing compared to dynamically generated content held in databases. Finally, managing the energy required to display a page is more complex as it requires optimisation of the scripting/code used to control the page. The difficulty is that the software industry does not apply consistent standards in the way browsers/readers interpret the content. This leads to an unnecessary duplication in scripting and style information in order to handle the display of content consistently across different operating system/browser platforms. It is especially difficult for sites which accept advertising as these blocks of code are loaded from external servers, and unless they are willing to look at the efficiency of their code it is not possible to address its impacts.

7.4. THIN CLIENT NETWORKS

The first computer time-sharing networks were designed around the client-server model.³⁶ Staff access the system from a number of simple "dumb terminals", sending commands to the central computer and receiving the results of the request on their terminal displays. As personal computers became more powerful, especially with the development of more user-friendly graphical user interfaces, having many stand-alone PCs on a network offered a more flexible and scalable alternative to traditional client-server networks.

Now the trend is swinging back towards client-server networks once more. As network speeds increase, and cheap and scalable rack servers become more powerful, the cost of developing networks of powerful stand-alone

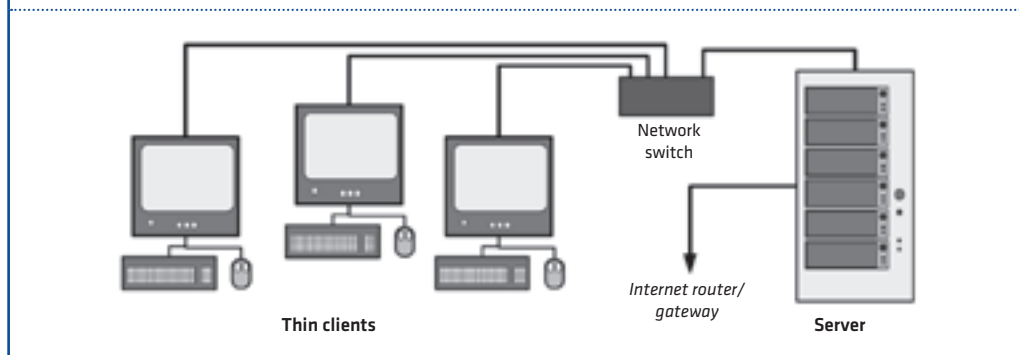
computers is becoming comparatively more expensive. Add to that the reduced costs of managing one large server rather than maintaining the operating systems of many stand-alone computers, and the costs of using client-server networks rather than many powerful desktop systems becomes an attractive option for larger organisations and education institutions.

In many ways a client-server network looks similar to a network of many powerful machines. The practical difference is that the power and primary cost of the hardware is based around the server. For example, client-server networks have become an attractive option for schools and educational institutions because of the lower upgrade costs. Rather than replacing all the machines on the network, the older/lower powered machines can be retained and their life extended by using

36. Wikipedia, 'Client-server model'. en.wikipedia.org/wiki/Client-server_model

Figure 7.1.

A client-server network



them as thin clients. Much of the expense of the upgrade is the cost of the powerful server which runs the system. Then, as the older clients fail or become unserviceable, they can be replaced with cheaper purpose-designed thin clients, further reducing the power consumption of the system.

The thin clients³⁷ run a minimal operating system: On old PC machines the system is installed like a conventional stand-alone operating system; on purpose-designed thin clients the system is often an embedded software system held in flash memory. Being able to do away with hard drives and other components of the stand-alone desktop system is one of the reasons why purpose-built thin clients are cheaper and use less power.

The central server runs a dedicated time-sharing system. This runs the programs requested by users and sends information to generate the desktop display back to the terminals. Time-share systems were originally developed for Unix machines. These ideas were taken on by the Linux community and developed into a variety of low-cost terminal server systems.³⁸ Their low cost, and ability to utilise cheaper/recycled equipment, has made these systems especially popular in education-

al establishments,³⁹ particularly in developing nations.⁴⁰ While initially sceptical,⁴¹ even Microsoft now offer a terminal server system⁴² compatible with Windows applications, and there are programs available to convert older Windows XP machines into thin clients.

The major difficulty with thin client networks is that the server represents a single point of failure. With stand-alone PCs, even when the network is down, it is still possible for people to carry out some work. With client-server systems any failure of the server or the network prevents everyone connected to the network from using the system. The other problem with thin clients is that they are not very flexible in their application. Although for routine applications they have few problems, if specialised or unconventional software is required for certain tasks it may not be easy to run this on the server. There is also an intellectual property issue as more expensive multi-

37. Wikipedia, 'Thin client'. en.wikipedia.org/wiki/Thin_client

38. Popular systems include: The Linux Terminal Server Project, www.ltsp.org; OpenThinClient, openthinclient.org; and ThinStation, thinstation.org

39. Edununtu, help.ubuntu.com/community/Edubuntu-Dokumentation/EdubuntuCookbook/ThinClient; see also Hargadon, Steve (accessed June 2012). Rescuing a School Technology Program: Linux Thin-client Overview.

40. Rais, Mark (accessed June 2012). Why Linux is the Logical Choice for the Classroom? www.reallylinux.com/docs/linuxclassroom.shtml

41. Lai, Eric (2007). Microsoft adjusts to add support for thin-client computing, ComputerWorld www.computerworld.com/s/article/9015280/Microsoft_adjusts_to_add_support_for_thin_client_computing_

42. Microsoft (accessed June 2012). Windows Embedded Thin Client. www.microsoft.com/windows/embedded/en-us/evaluate/windows-embedded-thin-client.aspx

user licences may have to be purchased to run certain software across a client-server network rather than a single licence for use on a stand-alone machine.

Overall, whether or not a client-server network creates cost and energy savings depends upon the types of tasks carried out by existing network us-

ers, and the extent to which any pre-existing hardware can be reused when creating the network. Given the current trends in computing, especially the shift towards more centralised/online services, it is likely that client-server networks will become more widely used in situations where many people share the same network system.

Box 7.3.

Networking check-list

Networking hardware:

- If possible use wired connections as this has a lower power consumption than using Wi-Fi.
- Organise the power supply to ethernet switches/Wi-Fi base stations so it can be easily switched off when the computer(s) they are connected to are switched off.
- Configure the network adapter to use the lowest practical speed for data transfer to reduce power consumption. This is primarily an issue with gigabit ethernet because its high capacity is rarely utilised in routine operations.
- Use the least amount of network hardware in order to reduce the overall demand for power – for example, rather than using two eight-port switches to connect machines use a single 16-port switch. Add new hardware only when network congestion becomes a problem.

Using networks:

- Take care when attaching files to emails, especially when sent to long distribution lists/email lists. Always consider the implications of sending large (multi-megabyte) files by email, and where possible scale down or process the attachment to reduce its size.
- Consider whether you need to use HTML formatted emails all the time – set the default on your email client to plain text and use HTML formatting only where it is necessary.
- To reduce the need to repeatedly download data, ensure that you have around 100 megabytes in your web browser cache to store the large files often used with advertising-driven websites.
- To further improve the demands made upon your system when using the web, block pop-ups from the browser and consider installing plug-ins to filter advertising and web bugs.
- For a more comprehensive approach to controlling the amount of data downloaded, and to filter connections

against certain types of content or to block access to certain sites, the most effective approach is to set up a proxy server.

Websites and content:

- Many of the rules which apply for the sending of email attachments also apply for the creation of websites/web pages. The impact of a website is proportional to its use reducing the size of individual pages and downloaded files will reduce the bandwidth required to serve the site.
- In general, static content uses less power to serve pages than database-driven websites, and static pages require less power to view than pages heavily dependent on client-side scripting.
- There has been little debate on the issue of web bloat and its impacts. Reducing the amount of data required to serve pages is a design issue, and to control that impact requires a positive effort to control the amount of data required to view content across the site.

Thin client networks:

- Thin clients utilise the power of a central computer to enable low-powered terminals to undertake common IT applications – consequently they are a way to use much older machines to achieve more powerful information processing tasks.
- As thin client networks become a cheaper and more popular way of providing access to information and networks, specialised low-power consumption terminal clients are being developed these consumed much less energy than traditional stand-alone “fat” computers used on large networks.
- Switching to a thin-client network is most commonly associated with the need to perform system upgrades, where the reduced cost and ability to recycle existing equipment create a more affordable alternative to replacing many stand-alone PC systems.

MAINTENANCE, SECURITY AND RESILIENCE

The term “sustainability” generally means the ability of a system to continue to function. In practice we must not only consider the ability to function, it's the ability to function even when external events are destabilising the system. Just like natural ecosystems, what a truly sustainable system must embody is not just the ability to work, but to continue to work after events have caused temporary interruptions or problems. This characteristic is called *resilience*,¹ and it's a feature that can be designed into the way we build and use computer technology.

This section looks at resilience of information systems in the broadest sense, considering how the equipment is used, how the hardware is maintained, and how the software is configured to maximise security. The objective of this process is to protect our information resources, and the machines that we use to access and manipulate them. That requires that we consider everything from how we define our needs and plan the system, the organisation of physical security, protection against malware and other system failures, right down to basic security measures such as passwords and access controls.

In many ways this section is all about combining all the actions covered in previous sections into a single, integrated method of managing our need for technology. Rather than separating different issues and looking at them in isolation we need to develop an integrated approach to system resilience and security, trying to minimise the ecological impacts of our demand for IT, and to refine the elements that make the system function to produce the optimal solution to our needs.

1. Wikipedia, 'Resilience (ecology)'. [en.wikipedia.org/wiki/Resilience_\(ecology\)](https://en.wikipedia.org/wiki/Resilience_(ecology))

8.1. SUSTAINABLE DESIGN

There is no good or bad way to practice sustainable design² – it either works to serve your particular needs or it doesn't. In order to adapt to the diverse way in which we use ICTs it is often necessary to unpack and adapt certain ideas and strategies to fit your unique circumstances. The true test of any system is that when you depend on it to work at its best, it functions as designed to protect the integrity and reliability of the system and the tasks it is designed to perform – and does so while internalising ecological principles which minimise the impacts of the system upon the environment.

8.1.1. Examining options

There is no one method to achieve a good sustainable design. When we build in an ecological dimension to planning our IT needs it quickly becomes apparent that we have to change our working practices, not just the tools we use to carry out our work. For example, while using more efficient machines can lower energy use, changing the culture or expectation of the people using the equipment offers many more options to reduce impacts (for example, turning it off/using power-saving options when equipment is not in use).

The best way to approach sustainable design is to think more strategically about our present and future needs. Thinking over a longer period allows time for cultural changes/changes in working practices to evolve alongside technical change. In practical terms, what many organisations seek to achieve is a compromise between “deep green” sustainability, where ecological principles completely determine how we work, and more practical sustainability, where we try and implement what is possible today and set longer-term goals for improvement. How we do this is a matter of setting the context within which we ask questions or find solutions to problems:

- Take a long-term strategy which defines the need for certain equipment over a number of

years – certainly longer than the lifetime of the equipment involved. By having an idea of what our needs might be in the future it is possible to invest in upgrades and replacements more strategically to reduce costs, and take advantage of new or more efficient systems.

- Always consider the effects of changing technical standards. Where certain technologies will obviously become obsolete we should plan for their replacement, although adopting new standards too early (for example, a new software package or operating system) might involve investment in equipment before its design has been perfected or its effectiveness proven.
- Technological standards can create inflexible restrictions – for example the use of proprietary designs which restrict compatibility with other similar systems (as a classic example, think of the many different types of USB or laptop power connectors in common use). Choosing hardware or software which uses generic standards avoids compatibility restrictions, and allows the switching of components/parts between different systems in the event of a fault in the equipment.
- It is important to question the cultural habits in the way people use ICT systems, just as we query the technical specifications of different technologies. Increasing efficiency and productivity often requires people to learn new skills, or adapt their methods of working to avoid habitual practices. Improving people's understanding and perceptions of the ecological impacts of their work, to be mindful of them, and to understand how they can work differently to address them, can be as cost-effective as finding technical solutions.
- Concentrate on what resources you already have. Optimising existing systems, using auditing or benchmarking tools to measure performance, can reduce the pre-existing ecological footprint – and might reduce the running costs or extend the working life of the equipment. More importantly, the process of optimising existing systems and working practices can highlight options or needs which

2. Wikipedia, 'Sustainable design'. en.wikipedia.org/wiki/Sustainable_design

were not previously understood, and this information and experience can become essential when planning a longer-term strategy.

- Measurement is critical – “what is not measured is not managed”.³ In order to compare the effect of two different ideas or options we have to be able to objectively measure their costs and impacts. Traditionally we use costs or prices to measure “the bottom line”. By building in values of both ecological and social impacts, in addition to traditional economic values, it enables decisions to be taken on a broader “triple bottom line”.⁴

Planning the transition towards a more sustainable way of working requires a balance between the ideal goals of sustainability, and the practical ability to find better methods of purchasing, operating and maintaining equipment. You must weigh up – as you perceive them – the risks, reliability and relative costs of different options to develop a sustainable solution to your IT needs, and then choose the set of options that can be reliably and easily maintained as part of your everyday routine.

8.1.2. Modular design

It can be difficult to consider the whole system design in terms of its sustainability. There are so many different aspects to sustainable IT that it is easy to become overwhelmed with the detail. To assist in the design process it helps to break down the system into its constituent parts, defining the larger system as a series of interconnected “modules”. For more complex systems, breaking down the problem into its constituent parts allows those with needs or expertise in different areas to evaluate the options that reflect their interest.

Creating a module design involves identifying the physical or logical boundaries which define each part of the whole system. The value of breaking down our needs or systems into smaller parts is that the problems related to each small part of the system are more easily quantified, and so can be more easily managed. Each part, independent of the rest, can be designed

to function in the best way possible – and it is by systematically working through each module one-by-one that we progressively produce a more sustainable system.

For example, an ethernet network operates according to a technological standard. The machines which connect to the network use the same standard irrespective of which manufacturer's equipment is used. For that reason we can consider the network and the machines which connect to it as representing different modules, and each module can be evaluated as a single unit in order to select the optimum design. For network hardware, this division enables us to measure the different performance of a wired network versus wifi, or a single large network versus a series of smaller subnets, and then combine those options to produce a system with the least impact in terms of economic and ecological costs. The software that computers use to connect to a local server on the network also functions according to a technological standard, distinct from the network hardware. This allows us to consider the desktop machines people use as distinct from the servers, routers and other parts of the network – enabling measurements or research to be carried out to find the most effective/efficient solution for each of these parts.

By adopting a modular plan as part of a longer-term strategy for improvement it is possible to progressively upgrade different parts of the system in isolation, without necessitating the upgrading or modification of other parts. Organisational complexity also plays a role here. Breaking down the organisation's infrastructure into manageable sections which mirror the structure of the organisation allows roles to be assigned, projects developed and targets set to implement changes, and the monitoring of longer-term objectives to be reported.

For large IT systems, developing a modular plan can be an elaborate process involving consultation, planning and finding options which meet the needs of many different system users. For small or stand-alone IT systems it can be a more informal process because in practice there are fewer alternative options to consider. What's important in either case is that, for each part of the system, you have a clear view of the purpose and function of the module, the options available to provide it, and the costs and impacts of each option. The overall cost or per-

3. Willcocks and Lester (1996). Beyond the IT Productivity Paradox, *European Management Journal*, vol.14 no.3 pp.279-290. [dx.doi.org/10.1016/0263-2373\(96\)00007-2](https://doi.org/10.1016/0263-2373(96)00007-2)

4. Wikipedia, "Triple bottom line". en.wikipedia.org/wiki/Triple_bottom_line

formance of the system as a whole is then an aggregate of the results for all the modules – and providing that each module represents the optimum design, the whole system will then provide the best solution.

8.1.3. Building-in resilience

While modular designs allows more flexibility in planning or upgrading, making it easier to optimise performance, when we consider security and resilience we think of “layers”. The concept of layering involves systematically building in security and reliability by spreading similar functions across different systems. This approach ensures that if one layer fails, then there are further layers to maintain system security and resilience. As with defining modules within a system, by breaking down the problem into individual layers, those responsible for that segment of the problem can take responsibility for it and work these tasks into their everyday routine – rather than having one person take on the whole task.

For example, the physical security in a building protects the computers inside the building; but if physical security fails and the computers are stolen then data back-up procedures should minimise the loss of data the machines contained. We can in turn break this example down into further layers to improve the system design. Building security can be broken down into separate zones, meaning that more important machines or information are protected by more levels of security. In a similar way, on a single computer non-sensitive information can be stored insecurely while secure data can be stored in encrypted formats.

Another common means of achieving better security and reliability is through redundancy.⁵ In the event of one piece of equipment/a module in the system failing there is always a spare available to replace it. For certain applications where high reliability is essential (for example the safety systems of nuclear power stations) the same piece of equipment might be replicated two or three times – which of course has a much higher ecological impact. In high-availability computer systems multiple machines work in parallel, enabling the whole system to keep working if one or two machines fail, and

even data storage might use an array of hard drives so that if one drive fails there is an exact copy of the data it contained on a second drive ready for use.

Unlike these highly specialised examples, improving the resilience of our everyday use of technology need not entail a large amount of expenditure and high ecological impacts. For example, on an average PC the only unique part of the machine is the data stored on the hard drive – the rest of the machine is expendable. By using a removable hard drive, taken from the machine at the end of each day and secured in a safe or locked cupboard, if the computer is stolen the data can be quickly reloaded from the old drive onto a new machine. Another option is to regularly back up the computer to an external hard drive so that, if the machine breaks down, the information on the external drive can be used on another machine.

Another way to develop resilience is to plan the purchase of equipment so that the likelihood of many failures occurring at the same time is minimised. All hardware has an expected working life – defined as a minimum by the guarantee period offered by the manufacturer. When a number of identical computers are purchased at the same time the chances are that they will begin to break down at roughly the same time too – especially if a certain batch of equipment contained a systemic flaw. Alternately, if you plan the procurement of equipment to a longer-term plan, then the purchasing or renewal of equipment can be staggered, and the chance that a large amount of equipment would break down over a short period is reduced.

At its simplest, creating resilience is all about having pre-planned alternatives to “business as usual”. It begins with everyday procedures such as backing up data, so that if a computer fails you can still have access to the data it contained. In addition you might choose to have email accounts with more than one service provider, so that if one provider's system fails you can still send and receive messages. As part of the design process, when evaluating options for different elements within the IT system it is wise to provide alternatives to essential parts or services to maintain the system in the event of unexpected occurrences. If you have the in-house skills, failures can be quickly fixed or worked around by repairing or

5. Wikipedia, 'Redundancy (engineering)

reconfiguring existing systems – perhaps utilising components from an expendable machine until longer-term repairs are made. For external services and support, it's important to

have a regularly updated list of service providers or trades people so that in the event of a problem occurring with the regular provider you can quickly switch to another.

8.2. PHYSICAL SECURITY

The physical security of the room or building where equipment is kept might not appear to be a critical component in sustainable IT. However, IT systems and information appliances (digital cameras, music players, etc.) are expensive pieces of equipment; replacing them entails the expenditure of a large amount of energy and resources too. Ensuring that they are physically secure, both from theft and casual damage, is an essential part of making the equipment function for as long as possible, and keeping the information resources on the equipment secure. Improving the physical security around computer equipment must therefore be considered an essential part of how these systems are used.

8.2.1. The building

Addressing the security of a building is entirely dependent upon those who look after the building. If you control the building, that's simple – *it's your problem*. If the building is shared between many tenants that's a different problem. Likewise, if you share a single space, be that a formal office or a space in a community centre, you are reliant on others to do things for you to assist your security. In that case improving security isn't just your problem, it's about convincing others of the benefits they might get from the process.

The physical security of a building is reliant upon the individual elements from which it is constructed:

- **Doors and walls.** Doors⁶ are a weak point because they are designed to open. A door can be strengthened by adding more hinges and locks, and reinforcing the door with ex-

tra wood or metal. If you use a deadlock, once locked the door can't be opened from the inside without a key – making it harder to remove objects if entry is gained by other means. However, there's no point strengthening a door to a point where it's much stronger than the walls around it (for example, a door set in a stud work and plasterboard wall) – otherwise it's just as easy to go in through the wall.

- **Windows.** Windows are a weak point, but often a last resort for illicit access due to the hazards of climbing over broken glass. Using key locks on window frames help because, once broken, the window frame can't be opened – which makes it much harder to climb through. The only effective way to secure a window is with internal bars or welded mesh securely fixed to the wall (external bars/fittings can be removed more easily).
- **Roof and floor spaces.** These are often overlooked. For adjoining buildings, if the roof space is shared then you'll need to secure any access point into the roof space from the inside. In offices which share a void above a suspended ceiling, you should also consider the likelihood of access from adjacent rooms. Likewise floor spaces can be vulnerable if their construction allows access from other rooms/spaces within the building.

While all the above are sensible measures, there is one very big note of caution. Yes, you can secure the building against various forms of external intrusion, either natural, accidental or deliberate. However, if the house/building is on fire and those inside can't get out, or the fire service can't get in, you've got a really big problem. The general point about security is that by working in layers you can spread the security risks through many different mechanisms rather than relying

6. Wikipedia, 'Door security'. en.wikipedia.org/wiki/Door_security

on just one or two options. By spreading the security measures to secure the most important items, it can become unnecessary to put intrusive or expensive general security measures in place around the whole building.

Traditionally the way to secure a building is with a lock. Big heavy locks might give the appearance of security, but many locks have basic flaws which allow the perception of security they give to be quickly bypassed. Both the building, rooms, and cabinets/cupboards in the room can be locked, but don't rely upon mechanical locks to guard against access by skilled operatives. All key-based locks can be picked⁷ if the person has the skills and the tools to do so. The standard front door lock, which clicks shut when pushed, is extremely insecure. By pushing a plastic card or a thin metal sheet (called a "shim") around the edge of the door it is possible to release the bolt and the door will open⁸ – unless the mechanism has a dead lock.⁹ Combination locks are more secure because they don't have a "key hole" that can be tampered with – but they still shouldn't be relied upon. Padlocks are also insecure, whether you use a key or combination version, because like front door locks they can be opened with a shim.¹⁰

8.2.2. The room

Working in one space makes it easier to secure the information and equipment that the space contains. If security measures can concentrate on that space, lesser physical security measures can be applied around the rest of the building. Work carried out in more than one location multiplies the security problem. General building security needs to be improved, and secure storage points – such as a filing cabinet or strong cupboards – need to be set up in each room/work location. You also have to give thought to how those measures are installed. For example, locked tamper-proof cupboards are not secure if they can be carried away – so if possible fix them to the wall or floor to prevent them being removed.

The greatest problem with securing the items in a room are those things which are physically difficult to secure – books, large desktop computers, DVD players, etc. Small high value items, such as digital cameras or external hard drives, can easily be locked in draws or cupboards. Laptops can also be locked away when not in use.

Which items we choose to secure in a room depends upon their value and/or whether they can easily be replaced. For those things which are replaceable – such as books, subscription publications and computer peripherals – the simplest option is to have insurance; if they're taken or destroyed you buy another with the insurance money (although some insurance policies might ask you to list all such items when obtaining the policy). If you want to keep other non-digital information safe – such as the original copies of important papers, certificates and other documents – then you'll need a fire-proof cupboard or safe to lock them inside when not in use. For irreplaceable items, such as film-based photographs or rare printed materials and books, the best option is to scan/digitise them, use the digital copies day-to-day and store the originals in a secure location away from the workplace.

The general problem is desktop computers, servers, laser printers and other expensive peripherals. If we look at where the "value" lies in computers, the hardware is expendable but the data that the computer contains is not – it's very valuable and often irreplaceable. For that reason it is easier to treat the hardware as expendable but organise the data it contains so it can be easily secured. On desktop computers use a removable drive caddy to hold the machine's hard drive so that, when not in use, the drive can be removed and securely locked away. Compared to the value of the information on the hard drive, a drive caddy is a minor expense, and can be easily installed in the machine. Laptops, mobile phones and other mobile devices should be backed up to some form of removable media – as outlined in unit 6.

8.2.3. Other building impacts

IT equipment uses electricity, but the environment within which these systems are used also consumes energy in the form of space heating, air conditioning, lighting, and fresh water for

7. Wikipedia, 'Lock picking'. en.wikipedia.org/wiki/Lock_picking


8. Wikihow, 'How to open a door with a credit card'. www.wikihow.com/Open-a-Door-with-a-Credit-Card

9. Wikipedia, 'Dead bolt'. en.wikipedia.org/wiki/Dead_bolt

10. Lock Pick Guide, 'Padlock shim'. www.lockpickguide.com/padlockshim.html

Figure 8.1.

Removable hard drive caddy



human consumption and flushing toilets. Creating a super-efficient computer system makes little sense if the environment within which it is used is not also optimised to minimise the impacts of the activities that the IT system supports. Again, this is the value of thinking in modules and layers – we can consider the needs of the whole system, and how best to address all the features involved.

When developing a sustainable IT system, it is possible to achieve similar, if not greater savings by attending to the design of the building and work areas.¹¹ This can be done through a formal building energy audit,¹² or less formal steps to tackle the main areas of energy and resource use:

- Space heating and/or cooling is the major consumer of energy in buildings,¹³ often making up over half the energy budget of the building –
 - Space heating can be controlled through better insulation, but it is far more effective to reduce the operating temperature of the environment.
 - The need for cooling/air conditioning can be managed by reducing the heat load

entering the workspace. CRT computer monitors use a lot of energy compared to flat screens, and switching to peripherals, printers and other devices which quickly switch to a low power standby mode when not in use will reduce the heat they produce.

- A major source of heating is sunlight entering through windows – especially in more formal office environments which have large windows – simple light-coloured screens and curtains are the simplest way to reduce the solar gain from windows.

- Lighting is a significant factor in commercial/office environments. While compact fluorescent lighting reduces energy consumption significantly compared to incandescent bulbs, the latest tubular fluorescent luminaires and LED lighting modules¹⁴ require even less power for the amount of light they can generate. Using timers on lighting systems can also ensure that lights are automatically switched off outside normal office hours.
- Water resources are an increasingly pressing ecological issue.¹⁵ Public water supply requires one to two kilowatt-hours per person per day to treat and supply drinking water

11. Whole Building Design Guide, 'Sustainable'. www.wbdg.org/design/sustainable.php

12. Wikipedia, 'Energy audit'. en.wikipedia.org/wiki/Energy_audit

13. Whole Building Design Guide, 'Optimize energy use'. www.wbdg.org/design/minimize_consumption.php

14. Whole Building Design Guide, 'Energy efficient lighting'. www.wbdg.org/resources/efficientlighting.php

15. Whole Building Design Guide, 'Protect and conserve water'. www.wbdg.org/design/conserves_water.php

- perhaps two or three times that figure in those areas which rely on desalinated water or water pumped for tens of miles from the nearest source. Minimising wastage and using low flush toilets are essential, and will become more so as the global climate warms.
- Indoor air quality is becoming a more significant health issue as research highlights the pollution created by modern construction materials and furnishings.¹⁶ This can be exacerbated by sealed air-conditioned building designs which reduce ventilation. The solution is to use products which do not contain solvents,

flame retardants, vinyl chloride and other volatile compounds which affect air quality.

- As well as using more sustainable IT systems, it's important to use products which have a lower environmental impact generally.¹⁷ Using recycled materials helps, but as a general rule we should seek to use fewer goods overall, more reusable and repairable goods, and try to extend the working life of products for as long as possible.

One solution for electricity supply is to buy from a renewable energy supplier, or seek to develop a renewable power installation. This issue is considered in section 11.

8.3. THE MACHINE

Many aspects of the installation and configuration of the computer are dealt with in units 3, 4 and 5. All machines need care and maintenance – as outlined in Box 8.1 – and by regularly attending to the well-being of the computer system you can improve its performance and extend its operating life. What is equally important in the day-to-day use of these systems are the security-specific aspects of computer use.

8.3.1. Securing the machine

Computer security begins with the design of the “the box” – the case containing the components of the computer system. The most important component in your computer is not the CPU, or the memory, or even the motherboard – it's the hard drive(s). That's because the hard drive holds all the data on the system. Protecting the box is therefore primarily centred around protecting the hard drives and the data they contain. The simplest option is to install a hard drive caddy into the machine, as discussed earlier. The other option with desktop boxes is to get a cage or a clamp. These fit over the case and then bolt to a table-top or the floor. Commercial ones are quite expensive, but they're fairly easy to make by anyone with basic metalworking skills.

Laptops are a different problem. It is possible to take the hard drive (or, on new disc-less machines, the flash card) out of a laptop. Usually they slot into a compartment inside the body, accessed through a flap in the case. The problem is the connectors on the laptop's hard drive are not designed for regular removal; if you did this repeatedly they will become damaged. However, if you were in an exceptional situation where you felt it necessary to secure the data on the laptop, you could remove the hard drive from the laptop in order to store it in a more secure location.

For more effective laptop security the most common option is a cable lock. One end of the steel cable is bolted to a wall or worktop, and the other slots into the body of the laptop and is locked with a key, preventing it from being taken away. The problem is that the hard disk is easily removable – and if the sensitivity of the data on the machine is a concern you would need to use hard-disc encryption as well to prevent theft of/access to the data on the drive.

16. Whole Building Design Guide, 'Enhance indoor environmental quality'. www.wbdg.org/design/ieq.php

17. Whole Building Design Guide, 'Use environmentally preferable products'. www.wbdg.org/design/env_preferable_products.php

Box 8.1.

Care and maintenance of electronic equipment

The human world is not always kind to ICT equipment: Dust can clog cooling ducts and reduce heat dissipation; food and drink can be spilled on keyboards and other equipment; and for mobile devices, careless handling can slowly degrade the internal components until, perhaps after a very heavy knock, they fail.

All computers, peripherals and other consumer electronics hardware should have a section in their user manual on care of the device. This will give you the basics of how to care for your equipment, how to carry out any routine cleaning or servicing, and what methods not to use to clean the device. For example some video displays, especially touch screens, can be sensitive to the solvents and detergents used in many household cleaners. In general all electronic devices can be easily cleaned using an antistatic or microfibre cloth. Stains and splashes from food or drink are best removed by gentle scrubbing with a non-abrasive cloth using warm water with no added cleaning agents. If you spill a drink on a laptop it's usually easier to buy a replacement keyboard for the machine – and that's certainly cheaper than buying a whole new machine.

For mobile gadgets the greatest risk is dropping them on a hard surface. The simplest way to guard against this is to buy a purpose-built soft case to hold the device. In the event the gadget is dropped the case prevents a large mechanical shock being transmitted from the hard surface through the case. Many mobile devices, such as iPods or mobile phones, can be fitted with a lanyard so the device can be worn around the neck or wrist. Some even incorporate headphones so that you do not stress the headphone connector of mobile phones and MP3 players by tugging on the cord.

A major problem for digital electronics is static electricity. This can be generated by human-made fibres in clothes and carpets, and is made worse in a hot, dry environment such as an office. Static discharges can damage electronic components when the inside of the machine is exposed, or when you touch the metal contacts on cables or connectors. Another way of generating static electricity is from air moving across dry plastic surface. For this reason you should never use an ordinary household vacuum cleaner to remove dust from electrical equipment. Quite apart from the static risk, the high force of the vacuum can damage internal fans or rip the keys from keyboards.

If you wish to remove dust there are small low-suction vacuum cleaners available for use with electrical de-

vices, which have conductive nozzles to prevent a build-up of static. You can also buy special anti-static brushes and cloths to remove dust without generating damaging static. Another option is to use small canisters of compressed air to blow the dust out of the case, fans and heatsinks. The cheapest option is to put your face close to the fan, keyboard or heatsink, close your eyes tightly to prevent damage from dust and grit, and then gently blow to remove the dust.

Dust is a particular problem with laptop computers. Laptops are often put down on top of soft furnishings and dusty table-tops – where the internal fan can suck up dirt and dust from beneath the machine and trap it inside against the heatsink. The internal heatsink of laptops use high density cooling fins with only a few millimetres gap between each one. Over a few months of inappropriate use these can easily clog with hair and dust, preventing the processor from being cooled adequately. You will notice this first when the area of the laptop where the processor is becomes unusually hot; then the processor will slow down during heavy operations because it can't lose heat at a sufficient rate. Finally the laptop may shutdown automatically to prevent damage to the system. Laptops usually have a panel in the base, sometimes secured by screws, which covers the fan and cooling fins – the user manual will usually show how to remove it. Then, using an anti-static brush or cloth, you can remove the wad of hair and dust which has built up between the fan and the heatsink.

Finally, as well as the physical hardware, the software systems of the device may need occasional maintenance. This is best carried out on a regular basis, such as regular calendar dates for following a routine back-up of information on the device. How this should be done is explained in the documentation written for the operating system, or on the producer's web site. An important part of system maintenance is installing software patches and upgrades. These not only fix system bugs, they also prevent malware infecting the machine from previously unknown flaws in the operating system. Some device manufacturers also offer firmware upgrades, to fix flaws within the programs which make the hardware function. Finally it is important to regularly “clean” the storage media inside the device of superfluous and unwanted files – for example, deleting the trash bin on the desktop. This is most appropriately done before performing a routine back-up.

8.3.2. Passwords and access controls

The purpose of access controls¹⁸ is to prevent anyone who shouldn't have access to a machine, or a place, having access. We've had complex machine-based access controls in common use for a few hundred years – *keys*. Computer-based access controls are more complex, but that doesn't mean they're any less fallible than their mechanical counterparts. While there are recent biometric¹⁹ or token access²⁰ systems available, the most common form of computer access control are passwords.²¹

Passwords are a means of access control, guarding against the unauthorised access to an information system. Some password systems are fairly weak, such as the PIN number²² used with mobile phones and bank cards. Some passwords, such as the passphrases²³ used with data encryption, are stronger because they are more complex. The principle of password access is that, as there are so many possible alternatives, it's unlikely anyone could guess the password. How “strong” a password is relies on how many characters it contains, how many different symbols each character can represent, and whether it has a unique or predictable format.

To illustrate how strong these access controls are let's look at how they work. A PIN number usually has four digits, each with 10 possible alternatives (0 to 9). The number of potential PIN numbers is $10 \times 10 \times 10 \times 10$, or 10^4 (ten to the power four), or 10,000 possible alternatives. In contrast an eight-character password with at least 62 possible alternatives per character (A to Z, a to z and 0 to 9) has 62^8 or about 220-thousand-billion (2.2×10^{14}) possible passwords. Passphrases are usually much longer, and can contain punctuation characters and spaces. For an 18 character passphrase, with around 80 possible alternatives per character, that's 80^{18} or eighteen million billion billion billion (or 1.8×10^{34}) possible alternatives.

Manually entering a PIN number every 10 seconds, assuming an unlimited number of tries, it would take nearly 28 hours ($10,000 \times 10$ seconds) to enter them all. This process is called a brute-force attack.²⁴ By using many processors in parallel, the specially designed “cracking” computers created for IT research and intelligence agencies can try *millions* of passwords *per second*. Against such professional cracking technology an eight-character password doesn't stand a chance and could be broken in a few weeks, but an 18-character passphrase might take months or years to crack.

In reality a brute-force attack doesn't start at “0” and step through each alternative. The process of password cracking²⁵ uses many steps to guess the most likely password/passphrase, based on the occurrence of certain words or letters in the language used. Another option is a dictionary attack²⁶ which tries words from a dictionary first, on the assumption that people use plain words for their password. By using people's habit of selecting whole words, names, birth dates, people's names and other such trends, the most common options can be eliminated first. However, an equally successful means of getting passwords is not the use of technology, but the use of human-to-human social engineering²⁷ techniques. Given the right scenarios, people will give away passwords, or can be tricked into doing so.²⁸

There are two ways to create strong security using passwords:

Firstly, by using more complex and random passwords and passphrases. While there's all sorts of recommendations on the length and format of passwords – such as how many upper/lower-case letters and numbers you should use – research on the use of passwords suggest

18. Wikipedia, 'Access control'. en.wikipedia.org/wiki/Access_control

19. Wikipedia, 'Biometrics'. en.wikipedia.org/wiki/Biometrics

20. Wikipedia, 'Security token'. en.wikipedia.org/wiki/Security_token

21. Wikipedia, 'Password'. en.wikipedia.org/wiki/Password

22. Wikipedia, 'Personal identification number'. en.wikipedia.org/wiki/Personal_identification_number

23. Wikipedia, 'Passphrase'. en.wikipedia.org/wiki/Passphrase

24. Wikipedia, 'Brute-force attack'. en.wikipedia.org/wiki/Brute-force_attack

25. Wikipedia, 'Password cracking'. en.wikipedia.org/wiki/Password_cracking

26. Wikipedia, 'Dictionary attack'. en.wikipedia.org/wiki/Dictionary_attack

27. Wikipedia, 'Social engineering (security)'. en.wikipedia.org/wiki/Social_engineering_(security)

28. Mitnick, Kevin (2003). *The Art of Deception: Controlling the Human Element of Security*. John Wiley & Sons, ISBN 9780-7645-4280-0 (paperback).

that this doesn't create much better security.²⁹ That's because people find it difficult to remember a truly random password, and so opt for easily guessable rules or formats which significantly reduce the variability and hence the security of their passwords. The solution is to use a mnemonic password,³⁰ created by a simple process which takes a more easily remembered phrase and reduces it to a string of letters and numbers (see Box 8.2).

Secondly, use multiple passwords for different functions on a machine. As more people use passwords on remote internet systems, the risks of using the same or similar passwords for all access controls is that if the passwords you use over a network are disclosed, it is possible to launch remote attacks on your computer system using similar style passwords to attempt a breach of security. Therefore when using multiple passwords, use different passwords, and a different style of password, on your local machine and on internet services. The difficulty is that remembering multiple passwords is difficult – but again, using mnemonic passwords based upon a more easily remembered phrase can solve this problem.

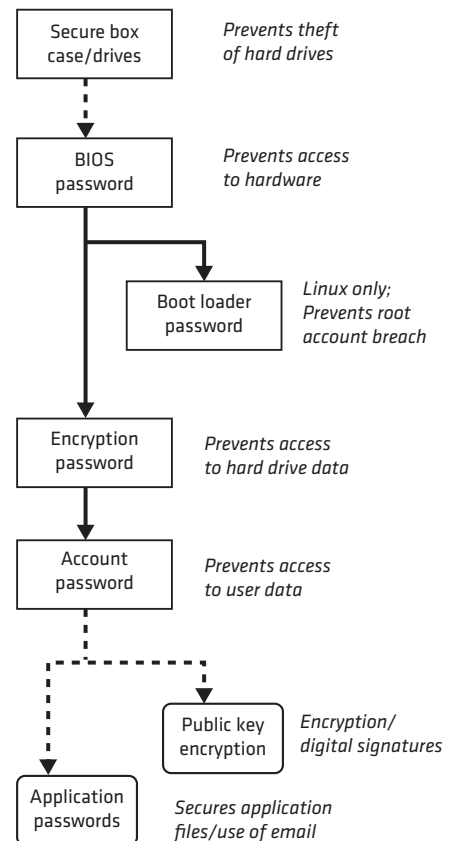
With most computer systems it is possible to set multiple levels of access control/passwords to secure access. First and foremost, it should be remembered that if the machine itself isn't secure – in particular the hard drives containing the data – then no amount of access controls will make the system secure. If an unencrypted hard drive is removed from one machine it can be plugged into another computer and read quite easily. Therefore, as outlined earlier, it is necessary to improve the security of all the layers in the system in order to make sure that security functions as an holistic process.

Assuming you've physically secured the system hardware, then there are various methods of using passwords to control system access:

- **BIOS password.** BIOS³¹ is a firmware program held on the motherboard. If you set a user

Figure 8.2.

Use of multiple passwords to secure desktop systems



or system password, when the machine first starts up this has to be entered before the machine will boot an operating system. Unfortunately the BIOS password is easily circumvented by clearing the contents of the BIOS chip or removing the internal battery from the machine – although if that were the case you'd notice when you next used the computer because you wouldn't be asked for your usual password.

- **Boot loader password.** This is only of relevance on Linux-based operating systems. Ordinarily the boot loader program would load the operating system. When the machine boots it is possible to interrupt the boot loader to give it instructions – which can with a little knowledge of the Linux system be used to circum-

29. Yan, Jianxin (2000). The memorability and security of passwords some empirical results. Computer Laboratory Technical Report 500, University of Cambridge. www.cl.cam.ac.uk/techreports/UCAM-CL-TR-500.pdf

30. Computer Academic Underground (2007). Mnemonic Password Formulas: Remembering Secure Passwords. www.uninformed.org/?v=76a=3&t=pdf

31. Wikipedia, 'BIOS'. en.wikipedia.org/wiki/BIOS

Box 8.2.

Mnemonic passwords

Almost all user-based security features are based on the use of passwords, and understanding how passwords should be constructed and used is an important aspect of user security. For example, if the password “password” were used for many years, for every password on a computer system, that would be incredibly insecure; likewise if the same password was used with a sequential number added for different programs, that’s insecure because the pattern can be easily guessed. In contrast, if each of the different passwords required to start and login to the machine were different, was only used for a few months or preferably less, and looked something like “cv6Td2Qb”, that presents a far greater security challenge.

When creating passwords people routinely substitute “1” for “i”, “0” for “o”, “5” for “s”, or “3” for “e”, in order to add numbers to a dictionary word – creating something like “Pa5sw0rd” instead of “password”, or “acc3s5” instead of “access”. Such variations can easily be deduced and tested by password cracking programs, using a dictionary or word list to guess words. The greater problem is that because we need many different passwords, users might use just a few across the whole system, or they might add easily guessable changes – such as consecutive numbering of a root word (“password01”, “password02”, etc.).

To produce easily remembered and secure passwords you must find an approach that suits you. If you’re the sort of person who can remember long strings of numbers and digits then you could use truly random passwords – using a random password generator to create them if necessary. If that’s not the way you think

then the strongest approach isn’t to use a *mnemonic password* that reduces an easily remembered phrase into a string of characters. To make a mnemonic password begin with a favourite phrase, such as a line from a song or a poem. Then take the initial letter of each word in the phrase to make the password, while substituting easily remembered words or numbers to deal with repetition and significant meaning – such as substituting the word “space” with a space (‘ ’) or underscore (‘_’) character, or the word “up” with the caret (“^”) character. For example:

- Take the phrase, “It is a far, far better thing that I do” – this reduces to “iaf,btld” (initial letters, ignoring the repetition of characters, and including the punctuation);
- The phrase, “may the force be with you” – this reduces to “mT4BwU” (alternates character case with each word, but uses the sound of “force” to indicate a number); and
- The word “encyclopaedia” – stringing out the word by its syllables this can reduce to “Ns1cl0PdA”.

By turning the characters of a password into “mnemonics” – small units that represent an idea – you can construct a seemingly random password in a way that’s far easier to remember. We each have our own meanings and preferred ways of associating and breaking down words and phrases – which introduces some randomness in the setting of the mnemonic. As a result, even if the phrase is the same, two people may not reduce it to the same set of mnemonics

vent the passwords of the system's "root" user account. Setting a password on the boot loader prevents this.

- *Hard-drive encryption password.* If the file system has been configured to encrypt the contents of one or more hard drives, you will be prompted for this password before the machine can finish loading the operating system. Hard-drive encryption is the best line of defence against the theft of the computer or its hard drives. The disadvantage is that it uses a lot of processor power to run the encryption process (a problem on older machines); and if you experience a data corruption problem you'll potentially lose access to all the data on that partition, or the entire hard drive.
- *User account passwords.* These are set when the computer is installed, and users can select their own passwords afterwards. The purpose of a user password is to protect data on a shared machine, or prevent opportunistic access to the user's files after they have logged in – for example, through the use of a password lock on a screen saver, and through the access permissions controlling access to files within the system. Although the system administrator can still change and override the user's account password, users can implement their own passwords outside the control of the system administrator using file locks and file encryption.

If you use a wireless or infra-red keyboard, that can be even more insecure because it's designed to transmit a weakly encrypted signal³² which can be easily cracked by those with the required expertise and computing power.³³ In turn, any "strong" encryption codes entered and intercepted from a wireless keyboard can be disclosed by cracking the "weak" encryption its signals are encoded with. Basically, never use any wireless gadgets (including networks) for secure or sensitive work as they're a weak link in system security. Another potential problem is the local network. Any other computer connected to the network, due to deliberate use or infection with malware, can monitor data pass-

ing between local machines in order to detect passwords and other security information exchanged between machines.

8.3.3. File wrappers and locks

If you are the only user of a machine security is less of a problem; but if a machine is shared between a number of people that creates security issues that must be addressed. That not only includes desktop machines, but also access to a shared file or network server, and of course online services. As more data is lodged within "the cloud" using online services, how we secure the files we use from unauthorised access is becoming a more important issue.

There are two ways of securing the data inside files:

Firstly, using *file wrappers*. As the name implies, a "wrapper" is something the file is enclosed in in order to protect it. The two commonest forms of wrapper are file encryption and digital signatures. Wrappers can be used with files stored on the hard drive of the computer, files stored on a back-up storage media, and to protect files in transit on a network or files stored on remote systems. While the mathematical encryption of data is the basis of both file encryption and digital signatures, they serve very different purposes:

- *File encryption*³⁴ is intended to prevent access to the data in a file. Using mathematical functions the digital content of the file is scrambled to prevent access except by use of a digital key. In symmetric file encryption³⁵ both the person encoding the file and those decoding the file has to have a copy of the key – which presents a problem because sharing the key, by any means of communication, might lead to its disclosure. To avoid this, public key encryption³⁶ was developed. Anyone can encrypt data using the public half of the key and send it to the recipient. The recipient then decodes the data using their own private key – and only the holder of the private key can decrypt the data.

32. Leyden, John (2007). Microsoft wireless keyboards crypto cracked. www.theregister.co.uk/2007/12/03/wireless_keyboard_crypto_cracked/

33. Securiteam (2007). 27Mhz Wireless Keyboard Analysis Report. www.securiteam.com/securityreviews/6G0030KKKI.html

34. Wikipedia, 'Encryption software'. en.wikipedia.org/wiki/Encryption_software

35. Wikipedia, 'Symmetric-key algorithm'. en.wikipedia.org/wiki/Symmetric-key_algorithm

36. Wikipedia, 'Public-key cryptography'. en.wikipedia.org/wiki/Public-key_cryptography

- *Digital signatures*³⁷ are not intended to prevent access to the content of the file, they are a means of protecting the contents of the file from being modified. As with encryption, a person uses a program to generate a digital signature, and this produces a block of data which can be attached to the file or held separately. Anyone can then check the authenticity of the data by checking the digital signature against the file using another program. With public key encryption this has become simpler as the public key can be used to verify the authenticity of any file that has been signed with the user's private key.

Many email programs now include features which allow the use of both digital signatures and encryption. This allows the easy encryption of emails or attachments to prevent disclosure; or by loading the public keys of email correspondents into your system, every time an email is received from them the program can verify that the email comes from that person (email is a very insecure system, and without digital signatures emails can easily be forged).

As shown in figure 8.2, each user can configure their own public- and symmetric-key encryption to protect their data on the machine, and use the passwords lock feature many applications use to secure files and the data they contain. These work separately from the security of the operating system – meaning that each user can implement their own data security irrespective of that applied at the system level. The principle security flaw in this process is the system administrator. They have the ability to log

the user's activity using key-logging software which monitors the keystrokes entered from the keyboard. Therefore, if you don't know or trust the system administrator of the machine you have to work on, it's best to assume that the system is not secure – even if you have the ability to configure high-grade encryption of the data on the system.

The second type of password protection is the built-in passwords used for many popular office-based applications. This allows the use of passwords to control access to the content of files – without the password the file cannot be opened and read. Some applications, such as the programs used to create PDF files, allow you to set conditions for how the document can be accessed – preventing those accessing the file from printing, using cut-and-paste or modifying the contents unless the file is unlocked with the correct password.

The general problem with the file locks used by application programs is that they're very weak. Most applications – such as PDF files, word processor files and spreadsheets, or ZIP archive files – have downloadable “cracking” programs available to find the password and unlock the file. These are brute-force programs, and with a reasonable amount of parallel computing power it is possible to crack the passwords on applications within a matter of hours or days. In contrast, if you wish to protect the integrity of the document's contents rather than just restrict access to it, cryptographic digital signatures are far more complex and can't be easily circumvented.

37. Wikipedia, 'Digital signature'. en.wikipedia.org/wiki/Digital_signature

8.4. MALWARE AND PHISHING

Malware³⁸ – or to use the more popular label, “viruses” – have been a problem with computers for at least the past 20 years. Originally they were carried in files and on floppy disks swapped between machines. Today, with the advent of high-speed broadband, the malware problem has become much larger and more complex, and can be transferred from the active code used in websites, email, as well as software programs. The motivation for producing malware has also changed; 20 years ago writing malware was a past-time for computer geeks and pranksters, whereas today the production of malware is increasingly related to organised crime and the work of intelligence agencies.

Most of the viruses, trojans and related malware in circulation are designed specifically for use with the Windows operating system. That is partly a reflection of the widespread use of Windows – if you want to write malware for nefarious purposes then it makes sense to target the most popular operating system. Both Linux³⁹ and Mac OS have far fewer problems with malware. That's partly because they are used less than Windows, but it is also related to the design of the operating system itself. Unix-like operating systems, such as Linux and Mac OS, are designed to be more secure and to enforce strict control over how users can use the operating system, unlike the Windows system which has traditionally been designed for convenience and ease of use. There *are* examples of Linux and Mac viruses. In addition there are also cross-platform viruses that can use Java, or the scripting language used in office-based programs such as Microsoft Office and OpenOffice, which can affect all major operating systems. Even the Android mobile phone system now has malware developed for it.

Malware isn't just a security issue. Machines with malware running in the background use more processor power, and if part

of botnets⁴⁰ they will use the broadband connection to shift large amounts of data as part of spam or denial of service attacks. This results in higher energy consumption, and for those whose broadband connection is metered it can also result in very large bills for the data sent over the network. However, it is also true that running anti-virus software, which routinely scans the operating system and incoming/outgoing data for malware, will add to the load on the processor and use more energy, and on older machines it might reduce system performance slightly.

For Windows machines anti-virus software is an essential part of running an internet-connected system. There are some free anti-virus packages available for Windows – most notably Microsoft's own *Windows Security Essentials*.⁴¹ Programs which are paid by subscription usually provide a higher level of protection, providing not only regularly updated anti-virus functions but also the filtering of web traffic and email. A particular problem with Windows are trojan programs which infect the machine and monitor its operation, potentially sending secure personal data to online fraudsters who can use it to commit crime.

For Mac and Linux machines anti-virus software is still an optional extra as it's not essential to everyday operation – although it might be if these systems grow in popularity. Where some form of anti-virus software is required is on Linux servers; not to protect the server, but to prevent malware being transmitted by the server which might affect Windows users. While Linux users often talk of the system's immunity to malware, it is just as susceptible to unknown security flaws as other systems – although one feature of open source software is that programmers are free to study the code in order to find such flaws and fix them. Even if someone were to knowingly download a program or script containing

38. Wikipedia, 'Malware'. en.wikipedia.org/wiki/Malware

39. Wikipedia, 'Linux malware'. en.wikipedia.org/wiki/Linux_malware

40. Wikipedia, 'Botnet'. en.wikipedia.org/wiki/Botnet

41. Microsoft (accessed June 2012). Microsoft Windows Security Essentials. windows.microsoft.com/en-GB/windows/products/security-essentials

malware and execute it, depending upon the permissions used to configure the system, it is likely to affect only the user account it is run from. The use of the suite of security hardening tools which are now available for use with Linux, such as *Security-Enhanced Linux*,⁴² reduce the likelihood that malware would cause significant damage to the system.

Lastly, an increasing problem on the internet is *phishing*.⁴³ Phishing is the use of the internet as a means to make people disclose sensitive financial and personal information. That information can then in turn be used to perform various forms of fraud and identify theft.⁴⁴ Often such frauds ask people to pay small sums of money as a “handling charge” in order to arrange the transfer, or to receive the goods promised, which ultimately will never turn up.

Email users will regularly receive emails telling them that they have won money, or someone wants to give them money, in return for their bank details. The best response to emails such as this is to hit the delete key. A

significant problem is the use of HTML emails as this allows the display of web addresses to indicate one web location, such as a bank, when the underlying link takes the user to a web site where they will be persuaded to release sensitive information, or load software which might install malware on their system. Before clicking on any link in an HTML email it is always essential to look at the web address. Often this will give an indication of fraudulent use as the address does not match the site/organisation it claims to represent. While a more technical issue for novice users, if you inspect the header data contained in the email it is possible to look at which address the email originated, and then check if that address matches the real address of the organisation's claimed identity.

Another recent phishing phenomena are hoax phone calls where an operator tells you that they have “detected a problem with your Windows computer”⁴⁵ (at present this is a problem specific to Windows – Mac/Linux users usually have the call terminated when they say that they do not use Windows). The caller asks

42. Wikipedia, 'Security-Enhanced Linux'. en.wikipedia.org/wiki/Security-Enhanced_Linux

43. Wikipedia, 'Phishing'. en.wikipedia.org/wiki/Phishing

44. Wikipedia, 'Identity theft'. en.wikipedia.org/wiki/Identity_theft

45. Microsoft (accessed June 2012). Avoid tech support phone scams.
www.microsoft.com/en-gb/security/online-privacy/avoid-phone-scams.aspx

you to perform some commands on your computer, and then pay to download some software to solve the problem. In fact, apart from the scam of making you pay for software that is not required, downloading that software is likely to introduce real malware onto the system.

As a general anti-phishing/online fraud avoidance rule, any organisation wishing to offer money, products or IT support, either on the phone or online, should not object to giving you

a company name, land-line telephone number and postal address at which you can contact them. In order to ensure that you have redress for any consumer fraud, it's important that the organisation is based within your resident legal jurisdiction. Any organisation which is not willing to give a telephone number or postal address to verify their identity, or which insists on conducting business only through websites or email, should not be trusted.

SECTION 9

END OF LIFE AND DISPOSAL

All types of tools and machines have a limited functional life. How they are used, cared for and maintained has an important role to play in maximising that working life, but in the end they will fail or, compared to a new device, become inefficient to use. Electrical devices, due to the toxic materials they contain, represent a hazard when disposed of inappropriately. For this reason ensuring that all electrical items are collected and reprocessed to maximise the recovery of materials, and minimise the impact to the environment of waste disposal, is an important part of how we manage this equipment at the end of its life.

Assessing when a device has reached the end of its life, and then disposing of it in an ecologically safe way, are the last two stages in the life cycle of ICT equipment. In this unit we'll examine the environmental impacts of electronics and electronic waste disposal. In the next section we'll examine ways in which devices might be reused to prolong their service life.

9.1. THE PROBLEM OF TECHNOLOGICAL OBSCOLESCENCE

When we examine our use of computers the issue of “obsolescence” not only applies to failed hardware; we also need to consider the logical as well as physical methods used to store data. A desktop computer or any other information storage device, such as a smart phone, is only as sustainable as the software and hardware used to store the data it holds. The reasons behind that statement are not obvious so let's consider a practical example.

One of the earliest mass market personal computers was the *Apple II*.¹ It used 5¼-inch floppy disks, and ran a bespoke disk storage system, and could store around 100 kilobytes of data. In the early 1980s the first versions of the IBM Personal Computer² also used 5¼-inch floppy disks, and ran the PC-DOS disk operating system, which wasn't compatible with Apple's system. At the time there was another competitor disk operating system, *CP/M*,³ which was also incompatible. Jump forward 25 years to the world of computing today. Both the Apple, early IBM and CP/M disk operating systems are obsolete. If you had any 5¼-inch floppy disks containing some important information it would be very difficult to access it. 5¼-inch floppy disk drives became obsolete long ago, and even the 3½-inch floppy drive is now almost history – abandoned in favour of CDs/DVDs and USB memory sticks. There are commercial archive services which read old computer media and convert old file formats to modern equivalents, but they are expensive. Therefore much of the data generated on these early computer systems has effectively been lost.

Unlike books, which are directly accessible to human interaction, access to digital information is intermediated through technological standards – and as those standards change, important or historically valuable data can be lost. For example, if an early Apple or IBM computer user had written a literary or technical work of great significance, but it was never published in print, it is now inaccessible as a result of the obsolescence of the tech-

nology used to store the information. One notable example of this problem was the BBC's *Doomsday Project*,⁴ an update to the English Doomsday Book carried out in the mid-1980s, the data from which was almost lost due to the obsolescence of the technology used to store it.

If data being created today has value for the future then you have to guard against the obsolescence of the hardware and software used to create it. Thinking about how we create information, and how we store it for the future, is an essential part of how we should plan our use of computers. The physical media which we store the data on, the file formats which we use to hold the data, and the operating systems which we use to run the programs which read those files, are all likely to become obsolete one day. While operating systems evolve, many of the programs we run on them stay largely the same. Even so, while the name or purpose of the program may not change, the way the data is physically stored by those programs changes from generation to generation of technology. In the proprietary software world, newer applications retain some backwards compatibility⁵ with the older version, but only for two or three generations. As a result old files can, after a time, become inaccessible to more recent programs. If we are creating valuable work and we want to ensure that the information is available for “future generations” – an idea which forms the core of the concept of sustainability – then we have to think carefully how we decide to create and store our data.

The basic rule to guard against future obsolescence is to keep copies of data in different formats, preferably formats which do not use proprietary encoding systems, and do not lock up the data in ways which might block access to it in the future. There are a number of commonly used formats which, due to their history of use to exchange data between different operating systems or software applications, are more suited to long-term data storage:

1. Wikipedia, 'Apple II'. en.wikipedia.org/wiki/Apple_II

2. Wikipedia, 'IBM Personal Computer'. en.wikipedia.org/wiki/IBM_Personal_Computer

3. Wikipedia, 'CP/M'. en.wikipedia.org/wiki/CP/M

4. Wikipedia, 'BBC Doomsday Project'. en.wikipedia.org/wiki/BBC_Domesday_Project

5. Wikipedia, 'Backward compatibility'. en.wikipedia.org/wiki/Backward_compatibility

- If you are using a proprietary program, such as *Microsoft Word*, then save copies of the most important files in other file formats – such as plain text, RTF, PDF, or an open formatting standard such as XML or HTML.
- When creating graphical data, or using proprietary CAD or publishing programs, export a copy of the final work to a common open image format (such as TIF, JPG, PNG or GIF) or if the program permits export it as a PDF.
- When using databases or spreadsheets, export a copy of the data into flat text-based files (such as tab or comma-separated tables) or as an XML data file.
- Avoid compressing or collecting data inside archive files – for example ZIP, RAR, TAR, GZ, etc. Not only are compressed files more likely to suffer a greater amount of data loss as a result of later file corruption, certain proprietary compression formats may fall out of use and become inaccessible in the future.
- As a general rule when storing data for long-term retrieval, unless there are security concerns, do not lock files using the encryption or password locks of applications – it's likely the password will be lost. If you need to en-

sure the integrity of files use detached digital signatures to verify their authenticity, for example by using a cryptographic hash function⁶ such as MD5 or SHA2.

- A popular file format, or a format conforming to an open standard (not a patented software format), is preferable to storing data using a file format tied to a single program or application.

What's important is that the conversion of data files is carried out when you are finishing and archiving a large project. That's because at a later date, when the system/application is about to become obsolete, to suddenly convert all the data that you have created over many months or years will be an onerous chore. By converting the formats when you finish a project/writing a significant piece of work, you also guard against hardware obsolescence. If all the files are openly readable on their current storage media (e.g. CD, DVD, memory card, etc.) then they can be easily copied to any new standard of storage media in the future. By collating our data in a future-proof way today we remove the obstacle of converting large amounts of data in a short space of time in the future, and make it easier to migrate those files to new types of storage media.

9.2. PLANNED/PERCEIVED OBsolescence AND SERVICE LIFE

If we look at graphs of how fast computers have become – for example the graphs which illustrate Moore's Law⁷ – we might believe that we are working many times faster than we were five or ten years ago. In reality that's not true. As the power of computers has developed, and the speed of computer networks has increased, so the amounts of data being moved around have grown too. This raises an interesting ecological paradox for the entire IT industry; as its capacity has never been constrained, the IT industry has never had to try and make more efficient software or data standards. The result of this is that while the pro-

cessing of data has increased for system users, the perceived increase in performance is nothing like the actual increase in system speed. A large part of those speed/capacity increases have been expended moving more and more complex data.

One factor in the increasing bloat of software is that older hardware can appear to become slow and inefficient. This usually happens when operating systems are significantly upgraded, for example the transition from *Windows XP* to *Windows Vista*. As a result, perfectly serviceable hardware may be scrapped due to changes in software, not as a result of the hardware's inability to function. This is an

7. Wikipedia, 'Moore's Law'. en.wikipedia.org/wiki/Moore's_Law

6. Wikipedia, 'Cryptographic hash function'. en.wikipedia.org/wiki/Cryptographic_hash_function

example of *perceived obsolescence*.⁸ Whether it is because the older hardware cannot run new software to the users' expectations, or because older hardware is rejected because of the appeal of new/more fashionable products, the user scraps the system and buys a new one. The idea of perceived obsolescence has been at the heart of the marketing of new products since the 1950s,⁹ and has been criticised since its development because of the costs to the consumer and the environment.¹⁰

Another process related to the use of technology is *planned obsolescence*.¹¹ This involves the developers or vendors of a particular system or product deciding not to support or service older equipment in order to promote the adoption of a newer model. Sometimes the operating life of a device can be hard wired into the logic of its electronics – and without having access to the detailed designs of the system it can be difficult to circumvent these restrictions. More commonly the producers of hardware or software will upgrade systems without backwards compatibility,¹² restricting the ability of the latest systems to read files from or export data to older versions. This tends to be more of a problem with proprietary systems, as open source/free software systems allow developers and enthusiasts to continue support for older versions long after the equipment has ceased being supported by its manufacturers.

For example, according to the environmental report produced for the first series of Apple iPad, it has a design life of three years.¹³ However, after just two years, the latest upgrades to the iPad's software cannot be used with the first series iPad, and so users of those devices are being forced to upgrade to the latest model

of hardware if they wish to continue receiving software updates.¹⁴

The problem with planned obsolescence is that there is little that the user can do to avoid the cost and environmental impacts of being forced to upgrade – that's why many consider planned obsolescence to represent a greater problem than perceived obsolescence because it takes away consumer choice from the decision to upgrade.

In the final analysis, the end-of-life of any device must be a balance between the service it gives, the cost of using it in its current form, and the costs or benefits of upgrading it. Often that balance is reshaped by external forces rather than being motivated by a change in the way we organise our use of the system. For those who perform largely office-related and internet/communication activities, the use of the same computer system for a significant length of time should not affect the way they carry out that work. What tends to create problems are changes to the applications and related software required to view web content, to read documents or files imported onto the system, or problems finding compatible hardware when existing devices cease to function.

As noted elsewhere in this guide, one of the best ways to minimise our ecological impacts is to extend the service life of the equipment we use. In order to achieve that goal we must find ways to manage these external incompatibilities, finding alternatives which avoid the need to upgrade until it is absolutely necessary. Just because a computer or other device ceases to have a viable function in one role does not mean it would not have a viable future in another. As outlined in unit 4, older PC hardware can perform a variety of functions – from a small file server to backing up a laptop or other machine via a network cable, to a machine which provides local services on the network, or just a machine to “play” with, practising your skills manipulating hardware or installing software. Provided that the application serves a useful purpose, it is a valid use of that technology. Only when a machine no longer has a useful application should it be disposed of – and even then other uses may be found for it if the machine is recycled via a local computer refurbishing or training scheme.

8. Leonard, Annie (2008). The Story of Stuff: Planned and Perceived Obsolescence.

www.youtube.com/watch?v=N2KLYYKJGkO

9. Dannoritzer, Cosima (2010). The Lightbulb Conspiracy. www.imdb.com/title/tt1825163/

10. Packard, Vance (1970). The Waste Makers. Reissued by IG Publishing, 2011. ISBN 9781-9354-3937-0.

11. Wikipedia, 'Planned obsolescence'. en.wikipedia.org/wiki/Planned_obsolescence

12. Wikipedia, 'Backwards compatibility'. en.wikipedia.org/wiki/Backwards_compatibility

13. Apple Computer (2010). iPad Environmental Report. images.apple.com/environment/reports/docs/iPad_Environmental_Report.pdf

14. Bevan, Kate (2012). You mean my two-year-old iPad can't take this year's software? The Guardian, 4th July 2012. www.guardian.co.uk/technology/2012/jul/04/apple-ipad-software-update

9.3. IMPACTS OF E-WASTE

When hardware does reach the end of its life it has to be disposed of. How that is carried out is intimately linked to the types of materials these devices contain, and how the ecological impacts of those materials can be controlled to prevent harm. It is difficult to control what substances our electrical devices are made from, as many manufacturers provide little information on the environmental impacts of production with the goods we buy – although some manufacturers now produce goods free of toxins such as brominated flame retardants and PVC.¹⁵

Ultimately, no matter how much we can reduce the toxic load, the physical volume of electrical waste will always represent a large expenditure of energy, resources and pollution to create it – which is why concentrating on maximising the use of these systems for as long as possible is so important. Perhaps due to the separation of our use of modern technology from wider ecological debate, there seems to be a sense that people have the right to access digital devices¹⁶ irrespective of the impacts that might have. For example, we may argue about the carbon footprint of air travel or power generation at great length, but the ecological footprint of consumer electronics is equally problematic.¹⁷

As discussed in unit 2, our dependence upon scarce mineral resources poses some difficult questions for our technological society. If digital and consumer electronics require these resources then just how much are we willing to compromise to own them? This is not an abstract question – it's one that's already being played out in the global market for resources, and the ways in which the high prices for these materials drives

their production in areas stricken by conflict, or using forced labour (so called “blood metals”). Similarly poor conditions exist in relation to the breaking up and reclamation of waste electronics in many parts of the developing world. The highly complex mixture of materials in e-waste, and the problems of recovering the substances they contain without causing any toxic pollution, make their reclamation in the developed world expensive. Even when the best reclamation technology is used, and high recycling rates are achieved, toxic waste streams requiring safe long-term disposal are still produced.

As a result, when many electronic devices reach the end of their life, they are exported from western states to poorer countries with less demanding environmental laws for “recycling”. In many west African states,¹⁸ India¹⁹ and east Asia,²⁰ this is creating a highly toxic legacy for future generations. As these schemes tend to target only the easy-to-extract metals (gold or steel) it means that the metals valuable to digital devices are scrapped or lost in the system. Even when old computers and mobile phones are exported to Africa for reuse they will, after a short period of use, be discarded – and in many African and east Asian states the lack of any formal collection systems and advanced processing of e-waste means that they are unlikely to be responsibly recycled.

As pressure groups highlight the issue, and consumers begin to question the environmental costs of the products they buy, government and policy organisations are beginning to address the problem of e-waste.²¹ The United Nations

15.Greenpeace International (2010). Why BFRs and PVC should be phased out of electronic devices. www.greenpeace.org/international/en/campaigns/toxics/electronics/the-e-waste-problem/what-s-in-electronic-devices/bfr-pvc-toxic/

16.De Decker, Kris (2008). The right to 35 mobiles, Low Tech Magazine, February 13th 2008. www.lowtechmagazine.com/2008/02/the-right-to-35.html

17.Greenpeace International (2010). Toxic Transformers Briefing. www.greenpeace.org/international/en/publications/reports/toxic-transformers-briefing

18.Basel Action Network (October 2005). The Digital Dump Exporting Reuse and Abuse to Africa www.ban.org/films/TheDigitalDump.html

19.Greenpeace International (August 2005). Recycling of Electronic Wastes in China and India. www.greenpeace.org/international/PageFiles/25502/recyclingelectronic-wasteindiachinafull.pdf

20.Basel Action Network/Silicon Valley Toxics Coalition (2002). Exporting Harm: The High-Tech Trashing of Asia www.ban.org/E-waste/technotrashfinalcomp.pdf

21.United Nations (2010). As e-waste mountains soar, UN urges smart technologies to protect health. www.un.org/apps/news/story.asp?NewsID=33845&Cr=waste&Cr1

Environment Programme²² has been working to develop a framework to control the impacts of e-waste under the Basel Convention²³ (see box 9.1), which controls the global trade in hazardous wastes. In March 2012, the Basel Convention adopted new rules to address the movement of e-waste to African states.²⁴

Slowly the regulatory process is catching up with the problems of e-waste. To make this process successful, a more general change in the culture of use surrounding our consumption of electrical goods is required – both reducing their toxic content and maximising their useful life in

order to minimise their impact on the global environment. In the end this is a design issue; it is a matter of how we choose to build technological systems. It also requires the users of electrical goods to be mindful about the impacts of the goods they demand, and how they should care for and responsibly dispose of those goods in order to minimise the global impacts of their use. If we respect the physical boundaries to the natural world then we can make a truly sustainable culture. The difficulty is that recognising these limits inevitably means applying limits to ourselves.

9.4. E-WASTE DISPOSAL SCHEMES

As legislation has only recently begun to catch up with the increasing use of electrical goods, the system of controls for e-waste management is, at the global level, still fragmented (see box 9.1). While global legislation is in development, both the electronics industry and many national governments are waiting to see what procedures current negotiations will enact.

At present the world leader in the control of e-waste is the European Union. The directive on e-waste (the WEEE directive) was agreed a decade ago. This enacted controls over the disposal of e-waste. Since then the EU has also agreed on the *Directive on the restriction of the use of certain hazardous substances in electrical and electronic equipment (RoHS)*. This seeks to control the range and concentration of toxic materials in industrial and consumer products, and complements the objectives of the WEEE Directive by reducing the toxic contaminants present in the waste stream.

While recent revisions to the WEEE Directive have sought to increase the amounts of electrical waste collected, there has been criticism of the way the problem is being tackled.²⁵ At present the focus of the directive is on waste collection, and without any other legislation which seeks to control the life cycle and service life of goods. This means that functional/usable electrical goods can be disposed of – which under the current system is likely to lead to them being crushed, fragmented and incinerated. For those who seek to reuse computers and other electrical items, this takes away items which they might have been able to reuse as the centralised collection of e-waste inevitably leads to usable computers and other goods being scrapped.

At present the greatest difficulty with the regulation of e-waste, and to some extent the laws on the restriction of hazardous substances, is that they deal with waste as an end-point of the consumer process. We are still looking at resource use as a linear process – involving production and disposal – rather than a cyclical process²⁶ which focuses on reuse, recycling and zero waste production.²⁷

22. UNEP (2010). Urgent Need to Prepare Developing Countries for Surge in E-Wastes www.unep.org/Documents/Multilingual/Default.asp?DocumentID=6126&ArticleID=6471

23. Wikipedia, 'Basel convention'. en.wikipedia.org/wiki/Basel_convention

24. United Nations (2012). UN-backed initiative to address electronic waste problem in Africa adopted. www.un.org/apps/news/story.asp?NewsID=41570&Cr=Electronic+Waste&Cr1

25. Guardian Environment Network (2012). EU beefs up electronic waste recycling. www.guardian.co.uk/environment/2012/jan/24/eu-electronic-waste-recycling

26. Leonard, Annie (2010). The Story of Electronics. www.storyofelectronics.org

27. Wikipedia, 'Zero waste'. en.wikipedia.org/wiki/Zero_waste

Box 9.1.

The control and recycling of electronic waste

There are various schemes around the world which seek to control the production and disposal of electrical waste. Some are run by industry organisations whilst others are mandated by national and regional law. The notable schemes which exist at present are:

- *The Basel Convention*

The *Basel Convention on the Control of Transboundary Movements of Hazardous Wastes and Their Disposal* provides a basic minimum standard for the international movement and processing of hazardous substances, including waste electronics. In 2006 the parties to the Treaty agreed the *Nairobi Declaration on the Environmentally Sound Management of Electrical and Electronic Waste*. Under this agreement regional systems to control the shipment and processing of electronic waste are now being created. The first of these regional agreements, covering African nations, should be completed during 2012. for more information see – www.basel.int

- *Sustainable Electronics Initiative*

The USA has no federal law on the control of electronic waste – waste electrical goods are dealt with as another part of the general municipal waste system. Certain parts of the e-waste stream, such as batteries or television monitors, are covered by other legislation, and some states implement controls over the disposal of articles containing toxic substances. In 2009 the Sustainable Electronics Initiative was set up by the Institute of Natural Resource Sustainability to encourage a more proactive effort on electronics waste, both managing the disposal of waste but also trying to minimise the production of it. It is hoped that this

and similar schemes will eventually lead to a national/federal initiative on e-waste. Some states, such as Washington and California, are already enacting their own stricter controls over electronic waste. For more information see – www.sustainelectronics.illinois.edu

- *Waste Electrical and Electronic Equipment Directive*

Globally the European *Waste Electrical and Electronic Equipment Directive* (the “WEEE” Directive) is the flagship scheme for the control and reclamation of electronic waste. It covers all electrical goods, both computers and every electronic appliances. Consumer electronic goods sold after 2005 are required to be collected by the companies producing or selling them – although most companies contribute towards industry-based schemes which organise the collection on their behalf through retailers and local authorities. For more information see – ec.europa.eu/environment/waste/weee/index_en.htm (contacts for each EU member can be accessed from this site).

- *Chinese RoHS/electronic waste law*

In 2007, the Chinese government enacted the *Administrative Measure on the Control of Pollution Caused by Electronic Information Products* law. This law has many similarities to the European Union's Restriction on Hazardous Substances (RoHS) laws, and is intended to encourage more responsible waste disposal, provide better information on consumers of the hazards of certain goods, as well as providing an incentive for cleaner production.

For more information see – www.chinarohs.com

REUSE AND RECLAMATION

Computers and the internet have enabled a capacity for human interaction and creativity that has not previously existed in our history. However, while many enjoy the products of the digital electronics revolution, comparatively few understand the principles behind how these tools function, and fewer still have the skills required to reconfigure these systems to perform different functions. In order to ensure that these technologies support the needs of the broadest spectrum of civil society, more people need to engage with the technical mechanics of information systems – both to improve their everyday use and extend and adapt their function to meet new challenges.

One of the simplest ways to engage more people in the techniques and skills required to maintain and adapt information systems is to set up computer reuse and reclamation projects. Such projects offer a wide scope of opportunities for those who wish to take part:

- They can learn more about how these systems can be used to create more interactive media through the development of online content and audio/video media recording/editing;
- They can learn about the functioning of information systems through the installation and configuration of operating systems and software; and
- They can develop electrical and mechanical skills through the maintenance, disassembly, and construction of machine hardware.

In this unit we'll examine the last of these options – developing practical computer reuse and recycling projects. Such projects are the foundation on which the community use of information technology is based. Creating the independent capacity to reclaim and re-purpose IT hardware allows a whole range of other projects which rely on the use of these technologies to take place.

10.1. ICT “RE-CYCLING” AND “DOWNCYCLING”

The term “recycling” has been widely used since its promotion by the green movement in the 1970s. What is less familiar is the term “downcycling”.¹ Recycling implies that goods or materials are reused for similar purposes; downcycling implies that the waste is processed to produce materials of a lower quality.

The quality of the material produced by reclamation is important. When raw materials are processed into finished goods a large quantity of energy is required, and this entails the production of pollution. The investment of energy and resources in manufactured goods is often described as embodied energy.² The benefit of recycling is that much of the energy and resources investment in the original production of the goods is retained, avoiding the need to expend resources again to produce those materials. Downcycling often involves breaking down or disassembling the goods to produce materials of a lower quality which, while they may be useful, require that the value of energy embodied within the goods is diminished. It is technically possible to re-manufacture goods from the downcycled materials, but doing so requires that we invest a greater quantity of energy and resources to achieve this.

When we look at the options for reusing digital electronics and information systems, the principles of downcycling and recycling can help us find the best options for the use of these goods. Many commercial electronic waste companies crush and fragment electrical goods,³ even if they are serviceable, because that is the cheapest way to extract value from the collected waste. Community-based schemes, provided they have the skilled technicians available, can look instead at the options for recycling – finding new uses for serviceable goods, and only discarding those items which cannot be reused because they have reached the end of their serviceable life.

This is why the level of technical skill possessed by the community is so important to how we can maintain the value invested in the production of technology. Community-based organisations can find ways of reusing electrical goods which may not be open to commercial reclamation companies – for example repairing electrical goods to supply at low cost to low income families. Community organisations can also respond to the needs of a community by using these materials to develop low-cost community projects – such as community ICT access and training projects. Such projects help to retain the value of the energy and resources originally invested in the production of these goods, and can make them available at low cost to those who may otherwise not be able to access them if they had to pay the market value for newly manufactured goods.

10.1.1. Equipment triage

The first important step in the reclamation of electrical goods is “triage”. Like medical triage, this involves separating the waste goods received by the project into different groups depending upon the serviceability of the items. This needs to be carried out by someone with experience of reclaiming electronic goods; so that they can diagnose any faults or problems which exist with them, and classify them according to the likelihood of repairing/reconditioning them.

As the electrical goods come into the project they must be sorted into three categories:

- *Serviceable* – the goods have no faults, and are in working order. In order to comply with electrical safety regulations in many states⁴ this may require a qualified electrician to carry out earthing and electrical safety tests in order to certify them ready for reuse. Once certified these goods can immediately be reused by the local community without the need for further work or modification.
- *Unserviceable* – the goods may have minor faults which need rectifying. In contrast to

1. Wikipedia, 'Downcycling'. en.wikipedia.org/wiki/Downcycling

2. Wikipedia, 'Embodied energy'. en.wikipedia.org/wiki/Embodied_energy

3. BBC News (2008). Throw-away Britain: How to recycle a television. news.bbc.co.uk/1/hi/uk/7747975.stm

4. Wikipedia, 'Portable appliance testing'. en.wikipedia.org/wiki/Portable_appliance_testing

the financial restrictions which operate within commercial reclamation organisations, where the most practical economic option is to recover the scrap value of the materials, this is the type of repair and reclamation work which could be carried out at a low cost if a project could be developed to organise the required training and safety certification. Once repairs have been carried out the goods can be passed as “serviceable” and made available for use.

- *Unusable* – for a variety of reasons goods can be unusable if it is not practical or economic to refurbish/repair them. That does not mean the whole unit has no use and must be scrapped. Parts of the mechanism could be stripped down and components reused in order to provide the parts to repair unserviceable goods. Stripping down old equipment can also act as a valuable teaching aid for those learning the skills of electrical repair, and being able to separate waste materials into separate metals, plastics and waste electronics allows a high price to be made from the sale of these materials to commercial reclamation companies.

Repairing computers and information devices is a very much simpler task than trying to repair the previous generation of analogue electronics. That's because digital electronics have a more standard format – although the electronics are broadly similar it's the software which makes them function in different ways. Desktop PCs are a good example of this. As outlined in unit 3, the circuit boards, cards, hard drives and other components are manufactured according to a standard design which makes individual parts interchangeable. An unserviceable PC, or a PC with an outdated motherboard, could have its case, hard drives and other components reclaimed and rebuilt with a more powerful motherboard – only the old motherboard need be scrapped.

How the reclamation process is carried out is often determined by the structure and objectives of the project. There are many examples of computer reclamation projects in operation around the world: International organisations, such as Computeraid,⁵ specifically work with aid organisations to provide computers to developing countries for education and community projects; some groups are more focussed on

recycling, such as Free Geek⁶ (an urban project in the US), who work on a semi-commercial basis providing a community service, training and volunteering opportunities related to the reuse of ICTs; others, such as The Container Project⁷ (based in rural Jamaica), are wholly community based, using IT reclamation as part of wider cultural and educational projects within the local community.

10.1.2. Finding a purpose for the equipment

Projects can run into trouble if their objectives are poorly designed, or because they lack the capacity to provide the service they desire. In the worst case they can rapidly become a dumping ground for waste electronics. Especially in those states where the disposal of electrical goods is regulated, donating equipment to community groups can be seen as the easy option for disposing of old equipment. As a result the groups can become quickly submerged in old equipment with no immediate outlet for its reuse, and no funds to pay for the disposal of it.

An important element within community-based recycling projects is having a community or body of people who are ready and able to put the reconditioned equipment to good use. Irrespective of how many technicians the group has available to carry out work, without a group of people ready and able to use/take-away the equipment the project will again become a storage area – albeit a store of repaired rather than scrap electrical equipment.

Setting up a recycling project depends upon finding people with the skills and motivation to “seed” their skills into the project. Even if the project has sufficient people to undertake the work, it is important that those administering the project take a cold, hard attitude to the acceptance of any and all waste electronics. There has to be a balance between the demand for the service by the wider community the project supports and the acceptance of excessive volumes of waste, or unusable types of equipment. Although having too little waste to recondition can be a problem for projects, having too much is equally problematic because it gives rise to storage problems and potential

5. Computeraid UK. www.computeraid.org

6. Free Geek. www.freegeek.org

7. The Container Project. www.container-project.net

regulatory liabilities for storing e-waste – all of which increase the costs and legal liabilities of the project.

Ideally any community-based recycling scheme should have its purpose defined by the users of the goods to be recycled. In that way the supply of waste goods doesn't drive the project: it's the needs of the users. By having a clear purpose the triage of equipment can proceed in a

more efficient way. If those involved know the standard or purpose of equipment they require, waste can be refused before delivery; and when waste goods are accepted the triage process can focus more clearly on accepting only that equipment which meets their specifications. As a result of a more focussed approach, the project is less likely to become a dumping ground for unusable/superfluous types of waste electronics.

10.2. RECLAMATION PROJECTS

If we look at the examples of successful computer/IT recycling projects around the world we can find that there are certain characteristics which are common to all. Those wishing to develop ICT education and computer recycling projects should consider these factors, and look at the information available to guide their own thoughts and planning.

10.2.1. Types of project

As outlined earlier, projects which recycle computer and other electrical equipment can have a varied purpose:

- For some groups the recycling of equipment is the only reason they exist, and those joining the project would not look at social or educational applications of their work. An example would be local Linux user groups⁸ who undertake recycling work as part of their promotion of free and open source software. This is not always the case, and for some technology centred groups reclaimed and re-purposed technology is a critical part of how they promote their social and political agenda – for example, the Institute for Applied Autonomy.⁹
- Others may use the equipment they recycle to support community-based education and training initiatives. Some charities and education agencies give financial support to IT re-

cycling projects as a means of providing skills training in economically disadvantaged areas, such as supporting practical training for the long-term unemployed. Such funding can form an important part of the project's budgets, as well as providing connections to the local business community in order to source discarded equipment. There are a number of examples of this around the globe, such as Hole-in-the-Wall education in India.¹⁰

- For some reclamation work is incidental to the purposes of the project, but the work carried out can provide a means to provide cheap IT equipment to support the educational or cultural activities of the group. Some groups have developed operating system support in native languages, usually ignored by mainstream systems developers, as part of their work to support the local indigenous culture – for example the NepaLinux¹¹ system in Nepal; others use recycled technology to offer new opportunities to express local culture, such as the iStreetLab project¹² or the Sheffield Access Space.¹³

Irrespective of the purposes, reclamation projects are usually based around the vision and drive of a handful of people. Most computer reclamation groups are motivated by people who have in-depth practical ICT skills – although they may not necessarily be employed by the in-

8. Wikipedia, 'Linux user group'. en.wikipedia.org/wiki/Linux_user_group

9. Vimeo (accessed June 2012). Institute for Applied Autonomy. vimeo.com/channels/iaa/6075432

10. Hole-in-the-Wall. www.hole-in-the-wall.com

11. NepaLinux. www.nepalinux.org

12. iStreetLab. istreetbhae.ning.com/video/mervin-s-bin

13. Sheffield Access Space. access-space.org

dusty. Others within the group may represent specific community or educational stakeholders who work alongside the more IT-literate members, bringing their community organising and networking skills to the project.

While groups may come and go, what marks out the most successful projects is that they are able to transfer the skills held by the founders to successive members who have joined at a later date. Although organisational and logistical problems often dominate the activities of the organisers, particular in more commercial projects, the longest-running projects are often the most successful in transferring skills to a large number of people within their community of users. Without such skills transfer the whole project might be jeopardised by one or two members leaving, taking with them the essential skills required to undertake practical recycling work.

As it is the sharing of skills which identify some of the most successful and longest-running projects, the groups which have had the greatest success have been those which are allied to the free and open source software movement.¹⁴ It may be that because the free and open source philosophy puts a great emphasis on skills sharing and collaborative development, they are more likely to create an environment which is more supportive and inclusive when training new members of the group.

10.2.2. Space

Most successful recycling projects, in addition to having the motivated people with the right mix of skills, usually possess another important commodity – *space*. It takes very little space to recycle a computer; a small desktop or workbench and a few hand tools is all that is required. Even so, to be successful a project needs sufficient space:

- To provide for the storage of computers, both waste computers which have been through the triage process, and also for the secure storage of the tools and test equipment used in the reclamation process;
- To provide sufficient space for education and

training to take place, with sufficient seating to accommodate the groups who take part; and

- To provide, if possible, a space for a wider community of users to use the reclaimed equipment informally, perhaps with internet access – note that public participation is an important part of promoting the work of the project, and a means to gain new active members.
- For the most successful projects, an exhibition space, or small theatre space for showing training/educational films or multimedia presentations created by the project.

Many projects fold when the space they originally used is lost. Some get around this by using vans or buses, converted into travelling workshops, to act as a base for the project. Others use low-cost or reclaimed buildings, such as shipping containers.¹⁵ Projects which do not have a secure base will often founder because of the problems of continually relocating, and because not having a fixed base prevents them from building up a critical mass of membership in the area. Often the most successful projects have an association with educational establishments, particularly universities, as this provides a secure base and a regular turnover of enthusiastic members.

10.2.3. Tools

The simplest recycling projects can be based around nothing more than a workbench and some hand tools. For more ambitious projects, particularly tied to education and training, a larger and more secure space is required in order to store the test equipment required for servicing IT equipment.

As outlined earlier in unit 3 computers are assembled and serviced on a workbench. A bench, with sufficient lighting to allow for the inspection and repair of equipment, is also an essential part of carrying out the initial triage on the waste goods taken in by the project. To ensure safe working, a bench is required for each person undertaking recycling work.

In addition to the basic hand tools required for servicing, projects can also benefit from having access to:

14. Wikipedia, 'History of free and open source software'.
en.wikipedia.org/wiki/History_of_free_and_open_source_software

15. Paul Mobbs (accessed June 2012). The Container Project.
www.fraw.org.uk/mei/container_project

- Sufficient storage space, both racks/shelving and drawers, to safely store IT equipment, computer components and tools;
- A properly installed “signal earth” to connect anti-static mats and wrist straps to for servicing computer circuit boards (connecting to the mains earth can be highly dangerous);
- A voltmeter/electronic multimeter for circuit testing;
- A bench space for soldering/de-soldering and electrical assembly work, with good ventilation to extract the fumes created by this process (this is an important part of electronics/multimedia based projects as they often involve the assembly and repair of many different types of audio/hifi equipment and cables);
- Test units for identifying/testing memory modules and other components;
- Network cable test unit for checking/manufacturing ethernet cables;
- A bench power supply;
- Speakers, audio amplifiers, data projectors and other peripherals to aid in the practical use of the equipment refurbished by the project.

Ordinarily projects would have a mains electricity supply. However, projects which work in less developed states with a poor quality power grid, or which work in a mobile environment such as

a bus, may incorporate renewable energy supplies. They may also teach the skills to create and maintain off-grid power supplies as part of the work they do. Diversifying from an IT-centric project to a more broad-based approach to education and technology is commonly found in successful recycling projects, in part because it reflects the “do it yourself” culture¹⁶ which motivates many of those taking part.

In unit 3 (box 3.3) and unit 4 (box 4.2) we looked at the range of applications it is possible to create with reclaimed equipment. Many of these are related to the networking of computers to perform more complex functions – such as running websites or file servers. While this can be difficult for the average person to organise at home, IT recycling projects offer the ideal space to experiment with wired and wireless networks and the configuration of services to run over networks. Having such a capability can be extremely valuable for training as fully functional networks are usually only found in universities and corporate IT training centres. For projects which include arts or multimedia components, this also allows the testing and simulation of more complex multimedia content before publishing it online. If the site doesn't have a live internet connection, it's also possible to simulate online services over a local network¹⁷ to provide experience and training in the use of email, the web and other internet services.

10.3. WASTE DISPOSAL

As noted above, space is also required for storing the waste generated by the project. In some states this may be required to meet a certain standard in order to comply with environmental and safety laws. For projects which accept large volumes of e-waste, the sorting and separation of different types of waste for recycling can represent a valuable source of income. How much space is required, and how it is organised, is usually determined by the contacts/opportunities the group has for disposing of its sorted waste to merchant recyclers.

Ideally, no electronics recycling project should operate without first identifying a vi-

able disposal route for the waste it generates. In those states which regulate electronic waste disposal, the waste produced by community groups, even not-for-profit groups, may be classed as commercial waste and disposal will cost. That need not be a barrier if the groups

16. Wikipedia, 'DIY ethic'. en.wikipedia.org/wiki/DIY_ethic

17. For an example of “offline” Internet services training see Free Range Network (accessed June 2012). The history of the Community-Linux Training Centre Project. www.fraw.org.uk/projects/community_linux/cltc_history.shtml

charges a gate fee¹⁸ for the waste they accept, and the fees collected are sufficient to fund the disposal costs of the project. Before setting up any project it is important to identify any regulatory hurdles that need to be resolved before the project opens. Often such restrictions only involve a one-off fee for registration, although there may be a burden of paperwork to complete afterwards in order to demonstrate compliance with the law.

Another important factor is the availability of disposal contractors. These might be e-waste reclamation companies, but if the project can sort different waste streams to the required standard then they could sell the materials to

merchant recyclers. This creates a positive environmental benefit because separated waste materials can be more easily and cheaply recycled than bulk disposed waste – reducing both the pollution created and the resources wasted. Bulk electrical waste is usually crushed, fragmented and then incinerated to recover only the metal components; waste separation allows more metals to be recovered from the e-waste, and the plastics can be recycled rather than incinerated. Separation may also create a stream of usable components which can be sold to others engaging in similar work – for example people who have trained with the project and who are renovating/maintaining their own equipment.

18. Wikipedia, 'Gate fee'. en.wikipedia.org/wiki/Gate_fee

RENEWABLE POWER

When you have taken all practical measures to adapt your equipment and procedures to more efficient computing, there remains one external factor which can be improved: the power supply. Around two-fifths of the carbon emissions produced each year are the result of electricity generation.¹ By comparison, transportation produces just over a fifth, and industrial emissions are another fifth. Of the emissions from power generation the majority are the result of coal burning – and coal burning also creates problems due to the emission of acid gases, which damage wildlife and crops, and polluting heavy metals. Two-thirds of the world's power is generated from fossil fuels,² and although de-carbonising power production will not, on its own, solve the problem of climate change, it is an essential step in tackling the problem.

For most ICT users their source of power is most likely to be the electricity grid. Just as recent operating systems have become inextricably linked to the use of broadband data connections, much of our modern electronics is tied to the 24-hour-a-day availability of a mains power supply – and in many states there are few other options. In that respect, the modern electricity grid mirrors the global economic process generally; it is built upon an underlying assumption that there will be a never-ending supply of energy and resources in order to make society function. In contrast, users of off-grid power systems know that there are finite limits to their power supply; and the scale and seasonal variation of off-grid power systems requires that electricity use must be monitored and adapted if they are to have power available when they need it.

How we use electrical power is influenced by the equipment we use, but the source of power production has a significant effect on our ecological footprint. How you go about addressing this issue is dependent

1. International Energy Agency (2011). CO2 Emissions from Fuel Combustion. www.iea.org/co2highlights/co2highlights.pdf

2. International Energy Agency (2011). Key World Energy Statistics. www.iea.org/publications/freepublications/publication/key_world_energy_stats-1.pdf

upon the budget you have available to purchase alternatives to fossil-fuelled grid power, and your technical capabilities to purchase and operate these alternatives.

In this section we'll look at three potential options for improving the ecological impact of your power supply:

- Changing the contract or tariff paid to your electricity supplier in order to support lower carbon or renewable energy technologies – this is the simplest option as it requires little change on the part of the consumer;
- Installing a grid-connected renewable power generation system – this is a more complex and expensive option than simply changing electricity supplier; and
- Developing an off-grid power supply system – this a more technically challenging option, although in some parts of the world off-grid power supplies are the only option to run ICT equipment.

11.1. IMPROVING THE SOURCE OF YOUR ELECTRICITY SUPPLY

The ability to vary the source of your electricity supply is dependent upon the level of liberalisation of the power grid in your area. In more developed states the supply of electricity is carried out by private companies, either partly or wholly regulated by government. As part of this process consumers may have a choice of different power suppliers, and each supplier will have a range of electricity tariffs for different types of electricity generating technologies. In less developed states power generation comes in a variety of forms, from wholly state controlled to wholly privatised. The general problem here is that there is often a restricted choice of power sources available to purchase through the grid.

Different power generation technologies can be grouped according to how “green” they are (see Box 11.1). While some forms of power are obviously renewable other power sources are less beneficial for the environment. For example, municipal waste incineration can create electrical power, and it is often promoted as an environmentally advantageous technology, although research suggests that waste incineration produces less power than the energy which could be saved if that waste had been recycled,³ especially paper⁴ and plastics. In the same way, the burning of biomass (wood and plant matter) can be less damaging to the environment, but if the land the fuel was grown upon had previously been forest, or agricultural land producing food crops, the impacts are not much better than using fossil fuels. For this reason it is essential to consider what

power sources are used to create the electricity that support our ICT needs. Rather than considering just the direct carbon emissions, it's important to use life-cycle analysis studies of power production which take land use change and other indirect effects into account⁵.

In states with a liberalised energy supply many different producers supply the grid with power. The power produced from these sources is then balanced by the amount of power individual users buy from the grid. Therefore with a “green” energy tariff, while you may not be physically using renewable-generated electricity, the amount you consume will be balanced by the amount of renewable power entering the grid. By contracting with a provider of renewable electricity you can purchase some or all of your electricity needs from renewable sources. While the cost of a renewable supply tariff is usually higher than the average grid price, how “green” the sources are is often reflected in that price. Large-scale hydro and waste incineration are often priced around the same rate as fossil-fuelled electricity. The most sustainable sources, such as a geothermal, wind and solar power are usually more expensive.

While not a solution for all the ills of the modern world, buying renewable electricity is an important step in moving society towards more sustainable operation. Unless people are willing to invest in non-fossil fuel electricity the alternatives required to address climate change will not be created. If it is affordable, buying renewable electricity is a means to encourage investment in those alternative sources of energy.

3. GAIA (2007). Incinerators vs Zero Waste: Energy and the Climate. www.no-burn.org/downloads/GAIA_Incinertors_vs_ZeroWaste.pdf

4. European Environment Agency (2006). Paper and cardboard - recovery or disposal?, EEA Technical report No 5/2006. www.eea.europa.eu/publications/technical_report_2006_5/at_download/file

5. Benjamin Sovacool (2008). Valuing the greenhouse gas emissions from nuclear power: A critical survey, Energy Policy, Vol.36 pp.2940-2953. www.nirs.org/climate/background/sovacool_nuclear_ghg.pdf The figures from this paper are quoted for each energy source listed in Box 11.1.

Box 11.1.

Impacts of electrical power generation technologies

The impacts of our electricity supply depend upon the sources used to generate it. At present the global power system is dominated by the use of fossil fuels renewable power sources make-up less than a fifth of supply. The list below outlines the impact of different energy sources, ordered from the highest to the lowest level of carbon emissions. The figures are the life-cycle emission of carbon dioxide (in grams of carbon dioxide per kilowatt-hour, $\text{gCO}_2/\text{kW-h}_e$) for a unit of electricity produced from each source:

- **Coal** ($960\text{gCO}_2/\text{kW-h}_e$) Coal is primarily used for power generation around the world. There are different grades of coal, and while high quality bituminous coal produces less carbon emissions, the use of lower quality brown coal and lignite, or even peat, will produce more.
- **Heavy oil** ($778\text{gCO}_2/\text{kW-h}_e$) Heavy oil is a low quality, sticky, tarry form of oil and for that reason it is cheaper than the price of oil quoted in the media. While its low price makes it an alternative to coal for power production, it tends to produce more soot, acid gases and heavy metal emissions than higher quality diesel fuel.
- **Diesel** ($778\text{gCO}_2/\text{kW-h}_e$) Diesel is often used for power generation as a back-up for the large power plants which supply the grid. In states with a poor quality power supply, diesel generators are often used as an alternative during blackouts.
- **Natural gas** ($443\text{gCO}_2/\text{kW-h}_e$) Natural gas is used primarily in more developed nations. As it is a higher quality fuel it produces less emissions than other fossil fuels.
- **Nuclear** ($66\text{gCO}_2/\text{kW-h}_e$) While there is much controversy over the use of nuclear power, it only makes up 13% of global power generation less than is produced from large hydroelectric dams.
- **Geothermal** ($38\text{gCO}_2/\text{kW-h}_e$) Geothermal power is produced in volcanically active areas, such as Iceland, the US or Kenya. It uses hot rocks to create steam to generate power.
- **Solar photovoltaic** ($32\text{gCO}_2/\text{kW-h}_e$) Photovoltaic (PV) cells turn sunlight into electrical power. While the greatest solar resource is in the tropics, even at higher latitudes photovoltaic cells can still produce a viable amount of power. Some manufacturers now produce solar PV kits to power laptops and mobile phones.
- **Biomass** ($14\text{--}35\text{gCO}_2/\text{kW-h}_e$) Biomass is plant matter. It can be burnt in power stations in the place of coal, or turned into liquid fuels such as biodiesel for use in generating equipment.
- **Solar thermal** ($13\text{gCO}_2/\text{kW-h}_e$) Solar thermal generation is used in desert regions, using mirrors to focus solar heat and create high-pressure gas to turn power-generating turbines.
- **Small-scale hydroelectric** ($13\text{gCO}_2/\text{kW-h}_e$) Small-scale hydro uses small flows of water in streams, sometimes without the use of a dam to trap water. As water is relatively heavy, micro-hydro is a good source of power for off-grid electrical systems.
- **Biogas/anaerobic digestion** ($11\text{gCO}_2/\text{kW-h}_e$) Biogas is created by the digestion of animal manure and plant matter by bacteria. It produces methane which can be burnt in modified generators or gas engines to produce electricity.
- **Onshore wind** ($10\text{gCO}_2/\text{kW-h}_e$) Onshore wind is one of the fastest growing sources of renewable electricity. While the largest turbines now produce up to five-million watts of power, small-scale turbines producing 75 to 150 watts can be used to power a laptop computer.
- **Large-scale hydroelectric dams** ($10\text{gCO}_2/\text{kW-h}_e$) Large hydroelectric dams which tap the power of the world's largest river basins produce 16% of the world's power. While they produce low-carbon electricity, they are highly damaging to build and often flood valuable agricultural land and wildlife habitats.
- **Offshore wind** ($9\text{gCO}_2/\text{kW-h}_e$) Offshore wind is slightly more efficient than onshore wind because of the higher and more consistent wind resource available out at sea. Even so, it is more expensive because of the problems of building and developing turbines at sea.

11.2. GRID-CONNECTED RENEWABLE POWER SYSTEMS

Another way of using renewable electricity is to generate your own from on-site renewable systems.⁶ This option depends upon the suitability of the location to construct an efficient renewable power system, and whether the electricity supplier/power grid operator allows power to be fed back into the grid. Ideally the amount of generating capacity should match the average amount of power consumed, although the cost, size and ability to dump excess power back into the grid are all factors in the planning and installation of grid-connected systems.

An important consideration in creating a grid-connected power system is cost. While developing a grid-connected renewable system is likely to be competitive with buying renewable electricity from the grid, it will never be a means to reduce the cost of electricity overall. That's because, even in those nations where electricity is very expensive, the cost of installing and maintaining small-scale grid-connected systems will be greater than the large economies of scale offered by large power grids. But, as noted above, if we wish to develop an alternative to the fossil-fuelled power grid then it is small, local and largely self-contained systems such as this (collectively called distributed generation)⁷ which will be required to address our current demand for electricity.

Figure 11.1 shows a typical grid-connected power system. Electricity produced by one or more renewable generating technologies is supplied to a synchronising inverter. This converts the low voltage direct current (DC) of small renewable technologies into high voltage alternating current (AC) used by the power grid. This allows the locally generated power to flow into the system without generating interference with the power from the grid.

If the renewable power sources do not generate enough current the grid supplies the difference between what is being generated and

what is being used. Where the local system produces more than required it is possible, if the grid operator allows it, to feed that excess power back into the grid. This is usually done by having two electricity meters. One measures the power used from the grid, while the other measures the power flowing back to the grid. When the utility company sends the bill they calculate the price of the power supplied to the grid and subtract that from the price of the power consumed (note, grid operators may not pay the same amount for the power you supply them as they charge for the power they supply to you).

Over the last two decades grid-connected renewable power systems have become popular in many states because they allow people to produce their own energy⁸ without the problematic restrictions of being wholly cut off from the grid. In some European states governments pay a premium for the power produced,⁹ and so the public can earn money from operating these systems in homes and businesses (they have become especially popular on rural farms where the large space available allows a much larger scale of installation). One difficulty that has arisen with these systems is that the synchronising inverter requires electricity to function. If the generating system produces insufficient power, or it's sited in a very poor location, the system can actually consume more electricity than it creates (this was a particular problem identified with small wind turbines designed for installation on rooftops in the UK).¹⁰

The problem with grid-connected renewable systems is that many of them do not function during a power cut. That's because the synchronising inverter requires the signal from the power grid to function – even if you are generating power, it will not be supplied to

6. For example, see Wikipedia: 'Grid-connected photovoltaic power system'. en.wikipedia.org/wiki/Grid-connected_photovoltaic_power_system

7. Wikipedia: 'Distributed generation'. en.wikipedia.org/wiki/Distributed_generation

8. Wikipedia: 'Grid-connected photovoltaic power system'. en.wikipedia.org/wiki/Grid-connected_photovoltaic_power_system

9. Wikipedia: 'Feed-in tariff'. en.wikipedia.org/wiki/Feed-in_tariff

10. Encraft (2009). Warwick Wind Trials. www.warwickwindtrials.org.uk

Box 11.2.

Uninterruptible power supplies

While not directly related to renewable energy, this is a topic of relevance to the provision of power to computers. Uninterruptible power supplies (UPS) are a means of protecting equipment against power cuts. For servers and small networks, especially when using a client-server network, a UPS protects against the data loss caused by small fluctuations or temporary interruption of mains power. In the most developed states, where the electricity grid has a very high reliability, UPSs are only used in large data centres and corporate networks – where high reliability is an essential part of the services provided. In less developed states, where brown-outs and temporary interruptions are more commonplace, the use of UPSs to prevent data loss can be more common – although often the costs of these units deters their wider use.

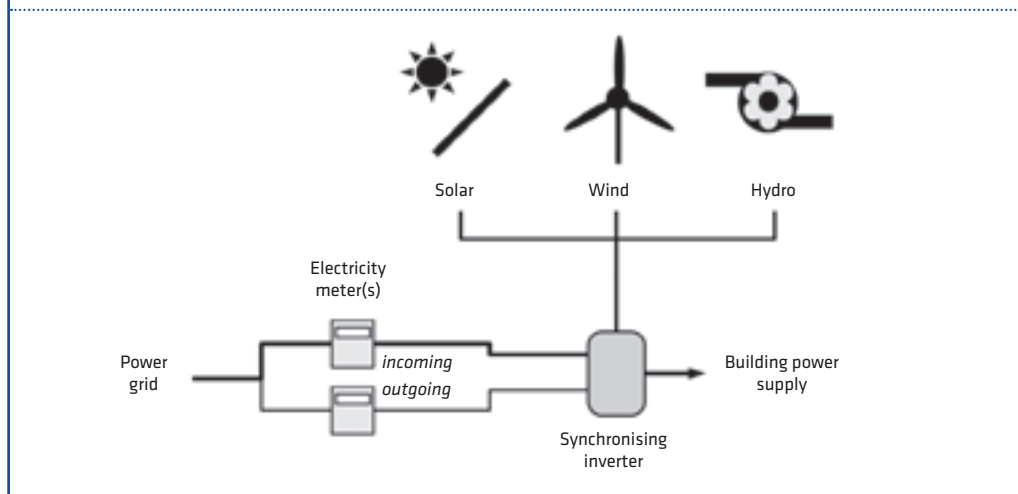
How expensive the UPS is depends upon how long you want it to function. The cheapest only provide a few minutes of power – just enough to enable the server/desktop to close down the running programs

and prevent data loss, or to prevent small fluctuations in the supply interrupting the computer. More expensive systems use a large bank of batteries to provide power for longer periods, or have a small battery bank to keep the equipment running while a generator starts up to supply power.

If data loss from an unreliable power supply is a problem, using a UPS is often a good investment because of the time and data saved. However, for most small computer users having a UPS for a single desktop machine can be expensive it is more practical to use a laptop computer instead. The internal battery of the laptop will cover temporary losses of power just like a UPS, although you will have a problem if using other mains-powered printers and peripherals. Using a laptop is also better suited to off-grid power supplies, not only because of their greater efficiency but also because laptops run at low voltages which can be matched by most small, renewable-power systems.

Figure 11.1.

Grid-connected power system



the building. To get around this problem some systems incorporate battery storage, allowing them to keep the inverter running without a grid supply. In effect, they function like a large uninterruptible power supply for the whole building. Due to their greater complexity, these systems

are more expensive to install and require more maintenance. Even so, if the reliability of the grid supply is an issue, grid-connected renewable generation can be a means of securing the power supply to the building in the event of the grid going down.

11.3. OFF-GRID RENEWABLE POWER SYSTEMS

An advantage of renewable power technology is that it allows the development of power systems well beyond the reach of the power grid.¹¹ This has obvious benefits for developing nations, and also for the use of mobile systems (for example, built into vehicles or vans) which allow ICTs to go on tour into rural areas. The critical factor in the design of an off-grid system is the average amount of power that needs to be supplied, and how long it must be supplied for. These two measures determine the amount of equipment required and thus the costs of the system.

Unlike the power grid, where more power can be supplied instantaneously on demand,

the generating and storage capacity of an off-grid system has absolute limits to its use. In order to make the costs and scale of the off-grid power supply manageable, the equipment used must function as efficiently as possible. The cost of supplying each additional kilowatt-hour of power is greater than for mains-powered equipment, and adding more power generation or storage capacity to supply inefficient equipment represents an unnecessary expense. As a result, the cost of more expensive but efficient computers and other equipment can often be justified by the cost savings from the power saved.

Figure 11.2 shows a schematic of a typical off-grid power system. By examining this you should be able to understand how these

11. Wikipedia: 'Off-the-grid'. en.wikipedia.org/wiki/Off-the-grid

systems work. There are many companies producing systems such as this, and they will construct/install it in order to ensure its safe operation in accordance with national building and electrical codes. With a basic understanding of electronics, it's relatively simple to construct one from individual parts. Either way, in order to specify the size and performance of the installation it is necessary to have a basic grasp of the principles involved.

At the heart of the system is the battery storage.¹² Its capacity determines how much power the system can supply, and for how long. It's also the most significant hazard in the system. A battery is a reserve of potential energy. When supplied at the required rate that's quite safe, but if a major fault occurs that potential energy can be released almost instantaneously – creating both a fire and flash burn hazard. Guarding against this is a matter of system design, protecting against short circuits, power surges, etc; and mechanical good design, ensuring that the battery is housed in a suitable enclosure to protect it against physical damage, rain, frost, and heat.

Off-grid systems¹³ operate at low voltages using direct current. A commercial system might use 24 volts or 48 volts, whereas most self-built systems will use 12 volts as this is the standard used in leisure/off-grid consumer systems. The voltage is a factor because it has an effect on efficiency – the higher the voltage, the more efficient the system. It is also important to consider the ease of maintaining and repairing the system – 24 and 48-volt components are relatively harder to source than the 12-volt units available through many outdoors and mobile home dealers.

There are a number of different battery technologies available, each with different costs and characteristics. The batteries used in cars and lorries are of low quality, and while they can be used to store power it's very inefficient to do so as they lose so much during charging, and can only be discharged by a small amount before cell damage occurs. Most batteries for power systems, while similar to automotive batteries, are more advanced “sealed”

lead-acid batteries. These have a longer life, are more efficient to charge, and can use up to half of their rated capacity before battery damage occurs. There are a number of different types of deep-cycle battery, from the basic leisure batteries used in mobile homes, to more advanced industrial batteries used in uninterruptible power supplies, to the highly specialised gel batteries designed for use in solar PV systems. How well the system performs over its lifetime depends to a large extent on the type and quality of the battery technology used.

Which power source is used to charge the battery will depend upon the feasibility of each technology for the application chosen. Some technologies are relatively mobile while others are only viable on a fixed site. Another factor is power density – how much power can be produced with a given amount of space and equipment:

- Solar photovoltaic¹⁴ (PV) panels are the simplest option. They're essentially a passive technology – you put them in the sunshine and they produce power.
- Wind power¹⁵ is the next most dense source of energy. This is more complex to construct as it requires a tower to be erected.
- Hydro power¹⁶ is the next most dense. The difficulty is that this requires the installation of pipework to tap a source of falling water – the greater the height the water drops, the greater the pressure in the pipe and the higher the power output.
- It is possible to use a generator powered by biofuel¹⁷ or biogas. This is a heavier and more complex operation, but it represents a very dense source of power as gas and biofuel contain a lot of energy in a small volume of fuel.

The battery store has a fixed capacity. If the battery was continually charged past that point it would slowly degrade the battery, and in the worst case could lead to a fire or release of flammable gases. To protect against this

12. Wikipedia: 'Lead-acid battery'. en.wikipedia.org/wiki/Lead-acid_battery

13. For a general introduction see the Homepower Magazine website. homepower.com/basics/started/

14. Wikipedia: 'Photovoltaic system'. en.wikipedia.org/wiki/Photovoltaic_system

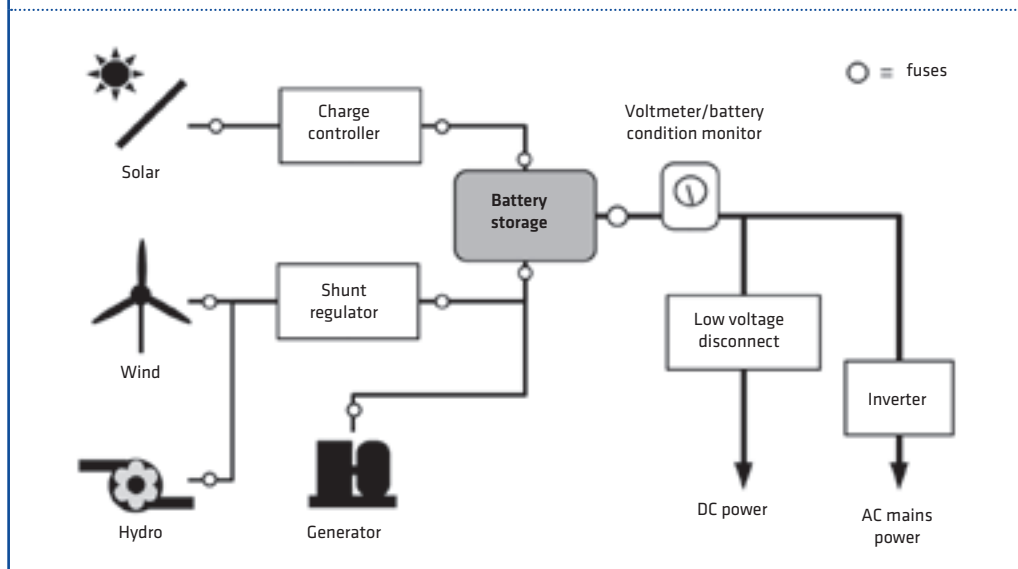
15. Wikipedia: 'Wind power'. en.wikipedia.org/wiki/Wind_power

16. Wikipedia: 'Microhydro'. en.wikipedia.org/wiki/Microhydro

17. Wikipedia: 'Biofuel'. en.wikipedia.org/wiki/Biofuel

Figure 11.2.

A typical off-grid power system



the battery must have a cut-out device which prevents overcharging:

- Certain power sources, such as solar PV, can be automatically disconnected to prevent overcharging. This is achieved with a voltage-controlled switch called a *charge controller*.¹⁸ When the battery reaches full capacity its voltage begins to rise exponentially. As it rises above a set point the controller disconnects the panels. More advanced controllers for larger PV systems (half a kilowatt or greater) - called *maximum power point tracking controllers* - sense the optimum operating voltage of the solar panels and adjust according. This increases the efficiency of the system by 15% in summer and up to 30% in winter.
- Other power sources, especially wind and hydro, cannot be disconnected to prevent overcharging. For example, if you disconnected a wind turbine the resistance to the wind created by the battery is removed and the turbine would spin faster and faster until it ripped itself apart. In these cases a shunt regulator

is used. Like a charge controller it senses the battery voltage, but instead of disconnecting it switches the current to a bank of high capacity resistors which dump the excess current as heat (in the most ecological designs, the excess power might even be used to heat water).

Note that many of the diesel/petrol generators¹⁹ designed for use with batteries include over-charge controls as part of their design, and will gradually slow the engine to an idle tick-over once the battery is charged (more advanced models will automatically turn off and on in response to the change in battery voltage). If using a very basic generator without these advanced monitoring systems it should be connected to the battery using a charge controller. This will disconnect the load and the generator engine should automatically slow down to an idling tick-over.

At the simplest level, using the power stored in the battery involves connecting a load across the terminals. In practice it is more complex as you also need to monitor the battery condition to prevent it being over-discharged.

18. Wikipedia: 'Charge controller'. en.wikipedia.org/wiki/Charge_controller

19. Wikipedia: 'Engine generator'. en.wikipedia.org/wiki/Engine-generator

Just like over-charging, regularly exceeding the battery's discharge limit will cause damage. Most commercial off-grid systems have a single computerised controller. This monitors both the charging and discharge of the battery store, and gives a read out of how much energy is stored inside the battery. Self-built systems assembled from individual parts usually have a separate battery monitor – often no more than a voltmeter – and a low voltage disconnect unit. Just like a charge controller, this monitors the battery voltage and in the event it falls too far it disconnects the load to prevent damage to the battery.

The power from the battery can be used directly if the equipment functions at that voltage. As noted above, it's essential to use a low-voltage disconnect unit to protect the battery if you use the power directly. For devices which do not operate at the battery voltage you will need either a power regulator to drop the voltage down to the required level, or a power converter (sometimes called a “DC to DC” or “buck converter”) to step-up the voltage to the required level. For example, most small self-built systems operate at 12 volts but most laptops use 18 to 20 volts. There are a number of power converters available on the market. Often these are designed for use in cars, and can be adapted to work with small battery-powered supplies.

Finally the battery's DC voltage can be converted into mains AC using a *power inverter*.²⁰ Some very expensive inverters can take any input voltage, but most inverters are designed to be used with a specific battery voltage either – 12, 24 or 48 volts. There are two general types of inverter:

- *Modified sine wave inverters* produce a very rough approximation of mains voltage. This means they are more efficient, but the modified sine wave can affect the operation of voltage sensitive equipment such as TVs and video recorders, data projectors and desktop

computers. Mains lighting and many types of motor-driven equipment are usually unaffected.

- *Sine wave inverters* create a fully compatible mains supply, although doing this can use 30% to 50% more energy than using a modified sine wave.

If you are only using an inverter, a low-voltage disconnect is not usually required as most inverters include an automatic disconnect. However, when you buy the inverter you should always check the voltage at which the inverter's disconnect functions to ensure that it doesn't over-discharge the type of battery you are using.

For those without experience of electronics or mechanics all this may seem rather daunting. Even so, if you were to buy a commercially produced system the considerations and specifications that you need to answer to make a purchase would cover much of this same ground – although perhaps not in the detail explored here. Before moving on to build or use a much larger system, you might find it helpful to buy a small educational solar power kit. These use exactly the same system components, albeit with a fraction of the power capacity. This enables you to learn more about the design and construction of these systems, and to get a basic grounding in the principles of their operation, before you move on to constructing more large-scale systems.

It should be noted that, even with the best-designed off-grid system, there may be times when it simply runs out of power. That's the nature of renewable energy; it is variable, and occasional natural variation will challenge the assumptions made in the design of most systems. On these occasions we just have to accept that we do no work – *nature has given us a holiday and we should do something else which does not involve the consumption of electricity!*

20. Wikipedia: 'Power inverter'. en.wikipedia.org/wiki/Power_inverter

SECTION 12

FURTHER INFORMATION AND INDEX

FURTHER INFORMATION

This publication cannot be a solution in itself; it is a starting point from which you can hopefully gain the inspiration to change and adapt your use of ICTs to minimise your impacts on the environment. Below, sorted by theme, is a list of reports, books, videos and websites which can help to fill in the detail that is beyond the scope of this guide:

Green & sustainable IT

- Global Action Plan (2007). *An Inefficient Truth*. www.it-energy.co.uk/pdf/GAP_An_Inefficient_Truth_Dec_2007.pdf
- Unhelkar, Bhuvan (2011). *Green IT Strategies and Applications: Using Environmental Intelligence*. CRC Press, ISBN 9781-4398-3780-1.
- Willcocks and Lester (1996). *Beyond the IT Productivity Paradox*, European Management Journal, vol.14 no.3 pp.279-290, June 1996. [dx.doi.org/10.1016/0263-2373\(96\)00007-2](https://doi.org/10.1016/0263-2373(96)00007-2)
- Murugesan, San. *Harnessing Green IT: Principles and Practices*, IT Pro (Journal of the IEEE Computer Society), January/February 2008. www.sis.pitt.edu/~dtipper/2011/GreenPaper.pdf
- International Energy Agency. (2009) *Gadgets and Gigawatts – Summary*, OECD/IEA, Paris. www.iea.org/Textbase/npsum/Gigawatts2009SUM.pdf
- Velte and Elsenpeter (2008). *Green IT: Reduce Your Information System's Environmental Impact While Adding to the Bottom Line*. McGraw-Hill Osborne, ISBN 9780-0715-9923-8.
- *Whole Building Design Guide: Sustainability* – www.wbdg.org/design/sustainable.php
- Randers, Jorgen (2012). *2052: A Global Forecast for the Next Forty Years*. Chelsea Green Publishing, ISBN 9781-6035-8467-8.
- APC (2010). *Global Information Society Watch 2010: ICTs and Environmental Sustainability*. www.giswatch.org/en/2010
- Bardi, Ugo (2011). *The Limits to Growth Revisited*. Springer Publishing, ISBN 9781-4419-9415-8.
- Turner, Graham (June 2008). *A Comparison of the Limits to Growth with Thirty Years of Reality*, Commonwealth Scientific and Industrial Research Organisation (CSIRO). www.csiro.au/files/files/plje.pdf
- Raghavan and Ma (2011). *The Energy and Emergy of the Internet*. www.cs.berkeley.edu/~jtma/papers/emergy-hotnets2011.pdf
- European Commission (2010). *Critical raw materials for the EU*, CEC. ec.europa.eu/enterprise/policies/raw-materials/files/docs/report_en.pdf
- WWF-UK (2010). *Common Cause: The Case for Working with our Cultural Values*. assets.wwf.org.uk/downloads/common_cause_report.pdf

Ecological limits, sustainability and resource depletion

- Leonard, Annie (2007). *The Story of Stuff*. www.storyofstuff.org/movies-all/story-of-stuff/
- Smithsonian Institute/MIT (March 2012). *Perspectives on Limits to Growth: Challenges to Building a Sustainable Planet*. si.edu/consortia/limitstogrowth2012
- Computeraid International – www.computeraid.org
- iFixit. *Self-Repair Manifesto*. www.ifixit.com/Manifesto
- Sheffield Access Space – access-space.org
- Mueller, Scott (2011). *Upgrading and Repairing PCs* (20th Edition). QUE Publishing, ISBN 9780-7897-4710-5.

Computer reclamation

Free and open source software (FOSS)

- Free Software Foundation – www.fsf.org
- Free Range Network (2009). *Thinking Outside the Cage: An introduction to Free/Open Source Philosophy*. www.fraw.org.uk/fwd?j1
- Lesswatts: *Saving Power with Linux* – lesswatts.org
- Open Source Ecology – opensourceecology.org
- Open Hardware – www.openhardware.org
- Open Source Windows – opensourcewindows.org

Electronic waste management and cleaner production

- *Waste Electrical and Electronic Equipment Directive* – ec.europa.eu/environment/waste/weee/index_en.htm
- Leonard, Annie (2010). *The Story of Electronics*. www.storyofelectronics.org
- U.S. Geological Survey (July 2006). *Recycled Cell Phones – A Treasure Trove of Valuable Metals*, USGS Fact Sheet 2006-3097. pubs.usgs.gov/fs/2006/3097/fs2006-3097.pdf
- Greenpeace International (2010). *Why BFRs and PVC should be phased out of electronic devices*. www.greenpeace.org/international/en/campaigns/toxics/electronics/the-e-waste-problem/what-s-in-electronicdevices/bfr-pvc-toxic/
- Greenpeace International (2010). *Toxic Transformers Briefing*. www.greenpeace.org/international/en/publications/reports/toxic-transformers-briefing
- Basel Action Network (October 2005). *The Digital Dump – Exporting Reuse and Abuse to Africa* www.ban.org/films/TheDigitalDump.html
- Basel Action Network/Silicon Valley Toxics Coalition (2002). *Exporting Harm: The High-Tech Trashing of Asia*. www.ban.org/E-waste/technotrashfinalcomp.pdf
- Greenpeace International (August 2005). *Recycling of Electronic Wastes in China and India*. www.greenpeace.org/international/PageFiles/25502/recyclingelectronicwasteindiachinafull.pdf
- Global Witness (2009) *Faced with a Gun, What Can You Do?* www.globalwitness.org/media_library_get.php/980/1277197135/report_en_final.pdf

INDEX

Rather than use page numbers, the index below uses section numbers/box numbers to direct the reader to the most relevant information across a number topics.

Check lists

backing-up/removable media box 6.3
backing-up mobile devices. box 6.1
desktop systems box 5.5
e-waste. box 9.1
hardware buying box 3.1
information storage box 6.4
networks. box 7.3
operating systems box 4.5
sustainability check list box 2.2
upgrading box 3.6
Cloud computing 6.2.5

Computers

BIOS 3.3.8
chipsets 3.3.3
hard drive/drive caddy 3.3.7, box 3.4, 6.2.2, 8.2.2
memory 3.3.5
motherboards 3.3.2, fig. 3.3
PC cards 3.3.8
ports and connectors table 3.3
power management 4.4, box 4.4
power supplies. 3.3.9
Processor/CPU 3.3.3, box 3.2
organising files/data 5.5
resource footprint 2.3, box 2.1
security. 8.3
speed/power 3.3.4
sustainability 2.2, 8.1
types of hardware 3.1.1
upgrading box 3.6
video display 3.3.6, box 5.1
Consumer electronics box 8.2
digital cameras box 5.3
e-readers 5.4
obsolescence 9.2
resource footprint box 2.1
waste electronics. 9.4
Data storage & backing-up 6.1-6.2

developing procedures 6.2
digital signatures 8.3.3
file encryption and locks 8.8.3
organising files/data 5.5
off-site 6.2.4
removable media. Box 6.2, table 6.1
reducing/scaling digital media files . . box 7.1
technological obsolescence 9.1
Ecological design/impacts. 2.3, fig. 2.3
backing-up data 6.1.2
e-waste. 9.3-9.4
impacts of electricity generation . . . box 11.1
modular design 8.1.2
Restriction of Hazardous Substances box 2.2
technological obsolescence 9.1
toxic materials. box 2.1, 9.3
Economic growth 2.2
Email. 7.2.1
Embodied energy 2.3.2
Energy use 5.2, 11.1-11.3
buildings. 8.2.3
Energy Star. box 2.2
EPEAT. box 2.2
scale of impact 5.1
measuring consumption box 3.5
off-grid electricity 11.3
one watt initiative. 5.2.1
power factor. box 3.5
power management 4.4
power supplies. 3.3.9
sources of electricity. 11.1, box 11.1
uninterruptible power supplies box 11.2
user feedback 4.4.3, 5.1
Free and Open Technology. 4.1
hardware 3.2.2
unbundling. 4.1.2
Hardware 3.1-3.4
applications and specifications 4.2.2, box 4.2
assembling. 3.4
buying. box 3.1
open hardware. 3.2
upgrading box 3.6
Intellectual property rights 4.1.1, 6.1.1

Life-cycle analysis	2.3	assembling PCs	3.4
Apple Computer	3.2.1, 9.2	care and maintenance	box 8.2
computers	box 2.1	closed hardware	3.2.1
mobile phone	box 2.1	reclamation projects	10.2
Limits to growth	2.2.2	service life	9.1-9.2
Linux	4.1, box 4.1	thin clients	7.4
Malware and phishing	8.4	triage of used equipment	10.1.1
Metadata	5.5, box 5.4	uses for recycled machines	4.2.2, box 4.2
Mobile phone		Resource depletion/footprint	2.2
backing-up mobile devices	box 6.1	blood metals	2.3.2, 9.2
growth in use	2.2.1, fig.2.1	copper	2.2.1
resource footprint	box 2.1	scarcity	2.3.2
Networks	7.1-7.4	Security	8.2-8.4
local networks	7.1	buildings/workplaces	8.2
off-site back-ups	6.2.4	configuring programs/applications	6.3
proxy servers	7.2.3	passwords	8.3.2, box 8.2
sneakernets	6.2.6	Sustainability	1.1, 2.1
thin clients networks	7.4	definitions	1.1
wake-on-LAN	4.4.2	information technology	1.2, 2.4
wired/wireless networks	7.1	options	8.1
Operating systems	4.1-4.4, box 3.3	recycling and downcycling	10.1
documenting systems	4.3, box 4.3	renewable electricity	11.1-11.3
fitting machine to hardware	4.2.1	resilience	8.1.3
free vs. proprietary software	4.1	technological obsolescence	9.2
system requirements	box 3.3	Technological convergence	2.1.1
thin clients	7.4	Waste	9.3-9.4, 10.3
Peripherals	3.4, table 3.3, 5.2.1	Web browsing	7.2.2
Printing	5.3	Web sites/online content	7.3
printing versus screen/e-readers	box 5.2	Viruses	<i>see malware</i>
Repair and reuse	3.2, table 4.1, 10.1-10.3		

ANNEX 1

Guidance for internet service providers

During the late 1980s, when the public first began accessing dial-up bulletin boards and internet systems, the role of internet service providers (ISPs) was highly specialised and served a limited audience. With the coming of the world-wide web in the mid-90s, the opportunities for public participation in computer-mediated communication and expression grew, and the role of service providers grew with this trend. In turn, as the social and political importance of electronic media grew, civil rights campaigners and ISPs sought to tackle the legal, economic and technological restrictions hampering wider public access to the internet – ensuring that the online community was open to as wide a range of social and ethnic groups as possible.

That objective became easier to secure in the first decade of the new millennium with the emergence of new ways to access electronics networks, such as mobile/smart phones, as well as the development of many different interactive and social media services which allowed easier communication. Almost 30 years after it left the closed world of academia and industrial research, the internet and electronic services are a vital part of many people's everyday lives – and even those who do not directly use them benefit from their use in logistics, manufacturing and government administration.

Now we have to deal with the impacts of that success; not only in terms of direct energy consumption, but also the indirect use of energy and resources in the manufacture of digital electronics and the waste and pollution this generates. This creates a number of problems with many people's expectations of information systems because, just like the problems of the human system within the physical world, balancing the demand of information systems for materials and resources, with the ability of the environment to sustain those activities indefinitely, will require redesigning our use of certain technologies.

At the same time the present economic crisis, which arguably has a relationship to energy and

resource use,¹ is generating calls for much wider social and economic change – from the online organisation of the Occupy movement, to the use of electronic media in the organising of the Arab Spring. Electronic communications and internet services have a role to play in enabling these changes, giving society a means to understand, debate and act to solve its problems more directly.

In the short-term, addressing the historic legacy of our development of IT systems is likely to involve slowing the pace of development. That's not just an implication of the current economic downturn and the financial pressure it is creating; in order to improve the quality of the systems we use, and reduce their ecological impacts, we have to extend service life and that inevitably means slowing the pace of upgrading. If we look at the roots of human sustainability, it is the scale of growth in the use of information systems which has driven the ecological issues surrounding their use today. Slowing that physical scale of growth by optimising the productivity of existing infrastructure can reduce the capital costs related to the turnover of equipment.

The primary role of ISPs is to provide network connectivity, technical services and training to those who use its electronic networks. In carrying out these functions ISPs have to opportunity to take a lead in reforming the way we acquire, use and dispose of technology. That begins by changing the way the service provider meets its own needs. As their own efforts to address these problems develop, it is likely they will need to engage their users/customers in that process to make the changes viable and more effective – for example optimising software to reduce system load or bandwidth use, or scheduling tasks to spread the load across a longer time period, utilising the space capacity available at off-peak hours. This requires change in the ways ISPs function, although it might also create new areas to provide support

1. Hall and Klitgaard (2011). *Energy and the Wealth of Nations: Understanding the Biophysical Economy*. Springer Publishing, ISBN 9781-4419-9397-7.

and services as part of their business model. In a world where economic pressures are causing more businesses and individuals to “make do” with the resources they already have, the skills ISPs develop to manage their own systems offer the potential to offer new support services to solve their customers problems.

Changing the way society uses online services is a complex problem – like many other ecological issues.² Encouraging people to change their lifestyles requires that they understand the links between their own small actions, the global impacts of all those similar actions by society, and understand why changes to our established ways of working can have both personal and global benefits. Education and raising awareness is the most direct method to allow people to understand their everyday impact, and make the types of changes which benefit them the most. This is a role which the internet, and ISPs as the creators or hosts of content, are well suited.

Alongside their traditional role in providing internet services, ISPs are in an ideal position to provide advice and training on minimising the impacts of ICTs. Whether they draw on their own knowledge and abilities to produce their own resources tailored to their users needs; or work within the projects and programmes developed

by the industry; or collaborate more proactively with researchers, educators and campaigners, ISPs are in a position to help promote change in the way society uses ICTs through the contact they have with their users/customer base.

Box A1 lists 10 priority areas for change: some are found in all green IT guides, such as improving energy efficiency; others are rarely included in guides, such as highlighting the role of ISPs in defending the public's right to expression and communication. If you review the content of the whole guide, the points in Box A1 are not exhaustive. What they are meant to do is act as general headings or signposts; they are there to point a direction from which you can carry out your own research, measurement and analysis.

Unit 12.1.1 ('Green & sustainable IT') lists a number of publications and reports which examine the issues. These are usually intended for use by medium and large system operators – smaller systems have not been well served by green ICT research as the costs and savings are often only significant when applied across a large IT systems. To redress this balance, much of this guide is tailored to the needs of those using small networks and stand-alone computer systems. The publications listed in unit 12.1.1 are more appropriate to the highly specific needs of large IT system users.

2. WWF-UK (2010). Common Cause: The Case for Working with our Cultural Values. assets.wwf.org.uk/downloads/common_cause_report.pdf

ANNEX 2

Guidance for policy makers & campaigners

Modern democratic principles are founded upon the work of 16th and 18th Century Enlightenment philosophers.¹ Exploiting the revolutionary information technology of the 16th Century – the metal movable type printing press – they were able to share and then collaboratively develop ideas which formed the basis of the political, economic and technological sciences which make society function today. Today the technological diversity and ease of use of information systems is driving a similar revolutionary process. However, this time the process is broadening the ability for groups across the social spectrum to access and share information. The result – from the crowd funding² of social projects, to the production of computer software which uses first/native languages, to the recent upheavals in the Middle East – is a trend towards greater social debate and diversity within societies where ICTs are widely available.

The downside of this is that as more people access and use the network, so the ecological impacts of this the system multiply. At the same time, while computing and network speeds increase, the movement of more data at higher speeds creates a rebound effect,³ erasing many of the technological efficiencies and energy savings that these new system create. Finding solutions to these ecological impacts does not require prohibiting or restricting the use of ICTs in the future. By redesigning the technologies involved to reduce the impacts of their use – both the physical hardware, software, and technological standards the system is built upon, and the manufacturing and waste disposal processes related to the device – we can create a significant reduction in the global impact of these systems.

To illustrate the subtle implications of this approach, let's look at a highly significant develop-

ment in the field of ICT policy. On 5 July 2012, the UN Human Rights Council approved a resolution stating that the right to freedom of expression on the internet should be protected by state parties. To quote the most significant part of the text,⁴

...the General Assembly at its sixty-sixth session on freedom of expression on the Internet, affirms that the same rights that people have offline must also be protected online, in particular freedom of expression, which is applicable regardless of frontiers and through any media of one's choice, in accordance with articles 19 of the Universal Declaration of Human Rights and the International Covenant on Civil and Political Rights.

This is a significant breakthrough in the long campaign for digital rights, and should be celebrated. Unfortunately, while this applies to the policies created by national and international governmental agencies, it does not have direct effect upon the handful of hardware and software companies who produce the bulk of the world's information systems. As a result new programs or operating systems could introduce restrictions on communication and there would be nothing the public could do to remedy that situation. Consumers could change their operating system or hardware, but as the production of equipment is concentrated amongst so few companies there is not a lot of choice to be had – and often these choices lead to other problems or arbitrary incompatibilities between different hardware and software technologies.

There are a number of solutions to this, and many of them have a positive effect on the sustainability of ICT equipment. In unit 4 we examined the idea of “unbundling”⁵ – removing the restrictions which close the design of products to a limited set of technologies, preventing the use of components from other manufacturers. On this

1. Wikipedia, 'Age of Enlightenment'. en.wikipedia.org/wiki/Age_of_Enlightenment

2. Wikipedia, 'Crowd funding'. en.wikipedia.org/wiki/Crowd_funding

3. Wikipedia, 'Rebound effect (conservation)'. [en.wikipedia.org/wiki/Rebound_effect_\(conservation\)](http://en.wikipedia.org/wiki/Rebound_effect_(conservation))

4. APC (July 2012). The UN recognises freedom of expression on the internet as a human right. www.apc.org/en/node/14772/

5. Wikipedia, 'Product bundling'. en.wikipedia.org/wiki/Product_bundling

issue we're approaching an interesting moment in time. One of the successes of the IBM-compatible PC was that, while the design was copyrighted, the functional components of which it was made were the property of a range of hardware manufacturers. That meant once PCs became widely used, consumers, and the PC marketing companies they bought their systems from, has a wide choice of parts with different qualities in order to construct the type of system that suited them – but, due to the “IBM-compatible PC” standard, the components could work together.

Now that era is slowly coming to an end as the world moves away from desktop PCs to other kinds of information appliances. As we move from the open design of the IBM PC to mobile phones, tablets, notepads and palm-top devices, we are losing the freedom to maintain, repair, reconfigure and reuse computing technologies. Unlike the advantages of the IBM PC's open design, these new appliances are mostly closed technology, unable to be modified by the user or retailer, and which vest control over the technology they contain to a handful of companies involved in their production.

In sustainability terms, this is taking us in the wrong direction; it encourages obsolescence and short-service lives because of the lack of competition in the market for parts, and because there is very little chance that we could repair the technology if it went wrong (proprietary components are often linked to one manufacturer, so buying one small part can often cost a large proportion of the cost of buying a new device). To make this clearer let's look at this another way. Even without a coherent policy on sustainable design, consumer pressure for greater service and reliability of household white goods – such as washing machines, fridges or cookers – mean that they can last five to fifteen years before they are physically worn out. Why then should a laptop or mobile phone only be used for two or three years, and a mobile phone even less? We

need the same consumer pressure to push the design standards and service life of consumer electronics in a more sustainable direction.

Box A2 contains a list of points we should be working towards in order to make ICTs more sustainable. Note that the list does not include many of the practical ideas outlined in the text of the guide. That is because, from the design of our hardware to the use of power management or security features, many of those things are already achievable – the fact they are not implemented is because people do not understand how to implement them, or the benefits of doing so.

Box A2 focuses on those essential elements of sustainable ICTs which we currently *do not* have the ability to enact – and for which change in national and international environmental and industrial policy is required to encourage change in the world's manufacturing systems. The list represents a kind of sustainability manifesto for the types of political, economic and technical changes that we need to enact to improve the ecological performance of information systems.

Finally, no matter how clean and green our information technologies become, it makes little difference if the rest of our everyday lifestyle is not sustainable. ICTs are only one small aspect of our personal ecological footprint; food, transport, buildings and the use of other consumer goods far outweigh the effect of using computers and electronic networks. Adapting our use of computers to significantly reduce energy and resource use is very important; ICTs are reliant of certain types of minerals, as well as fossil fuels and chemicals, the production of which will become more constrained as we move through the 21st Century. However, in terms of the relative scale, we cannot solve the ecological problems of the human system with ICTs alone – we must also adapt those other aspects of our lives to take control of the unique set of impacts that our demand for those goods has.

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

