



Technology & Standards Watch (TechWatch)

www.jisc.ac.uk/techwatch

Horizon Scanning report 09_02

First published: Nov, 2009

Low Carbon Computing: a view to 2050 and beyond

by

Paul Anderson
Gaynor Backhouse
Daniel Curtis
Simon Redding
David Wallom

Reviewed by:

Paul Hobson
Associate Director - Deputy CTO
University of Cardiff

Zahl Limbuwala
Chairman
BCS Data Centre Specialist Group

Raza Rizvi
Director
activereach Ltd

To ensure you are reading the latest version of this report you should always download it from the original source.

Original source	http://www.jisc.ac.uk/techwatch
Version	1.1
This version published	November, 2009
Publisher	JISC: Bristol, UK
Copyright owner	Paul Anderson, Gaynor Backhouse, Daniel Curtis, Simon Redding, David Wallom

Table of Contents

1. Introduction	2
1.1 How low can 'low carbon' go?	3
1.2 Beyond targets	3
1.3 Summary and terms of reference	4
2. Policy state-of-play	5
2.1 European Action	6
2.2 Strategy and policy within the computer industry	7
2.3 Higher Education	8
3. Metrics and measurement	9
3.1 Life cycle Assessment	9
3.2 Carbon footprint	10
3.3 Labelling and Standards	10
3.4 What to do now with regard to metrics	15
4. Technologies and Standards	16
4.1 Data centres and energy-intensive computing	16
4.2 Specific HE/FE high-energy consumption applications	27
4.3 Storage	31
4.4 Networks	35
4.5 End User Devices (PCs etc)	41
4.6 Printers and printing-related technology	55
5. To 2050 and beyond	58
5.1 The future of the data centre	59
5.2 Energy security for ICT	63
5.3 Adapting to climate change	66
6. Conclusions and Recommendations	68
6.1 Low Carbon Roadmap	69
List of abbreviations	75
About the Authors	76
References	77

1. Introduction

The UK is the first country in the world to introduce a legally binding framework for tackling climate change, and the implications of this are likely to be far-reaching for the public sector. Over the coming years targets will be set, carbon emission budgets agreed and strategies and plans announced: all with the overall goal of cutting the UK's emissions by 80% by 2050.¹ What will this mean for the ICT sector, and in particular for those responsible for information services in HE and FE? This report sets out to explore how lowering carbon emissions from ICT will play out in practice and what the future may hold.

While issues with climate change are being communicated widely and energy use from other sectors is being tackled by a variety of financial instruments and measures, ICT has, until recently, been a very quiet problem. Computing has quietly followed Moore's Law since the 1960s and it has just been taken for granted that each new generation of equipment can provide more processing capacity without a corresponding increase in cost. What has received little attention until very recently is just how much energy ICT consumes and hence the greenhouse gas (GHG) emissions produced.

These emissions are rising rapidly. A recent report by Intellect, a prominent trade association for the technology sector, estimated the proportion of global emissions due to ICT as being approximately 2%, roughly the same as the aviation industry (Intellect, 2008). The report recognised the figure as being much higher for technology-heavy, knowledge-based economies such as Britain and the USA, where the estimate of the proportion of carbon emissions due to ICT ranges from 4% to 10% of the total. It also reported that DEFRA's Market Transformation Programme, which supports the development and implementation of UK government policy on sustainable products, states that ICT-related electricity use more than doubled between 2000 and 2005.

However, the impact of ICT on the environment has both positive and negative aspects. For example, a video-conferencing service for a college may use electricity, but the overall environmental impact will be small compared to the carbon emissions created by travel. Indeed, it is possible that we may need more ICT to develop a low carbon society. This can also be conceptualised as direct and indirect impacts. Direct environmental impact is caused by the equipment itself or the way in which it is configured, whereas indirect impact is measured by the effect produced through the business processes the technology changes or facilitates. For example, a PC that uses energy generates a direct negative impact. If the PC helps to run a car-sharing database it creates an indirect positive impact.

In-use electricity consumption is not the only concern. ICT impacts negatively on the environment because of its embodied energy, use of toxic or rare material, and the waste created at the end of a product's life. This is one part of a systemic view that, in a structured way, takes into account supply chain management and, ultimately, social and environmental factors such as the effect on humans of using hazardous waste materials in manufacture, the over-extraction of water sources in manufacturing plants or the social effects of dumping waste.

For the purpose of this report we class these wider concerns as part of the 'green' or 'sustainable' ICT agenda, for which JISC has already published work (James and Hopkinson, 2009a). This TechWatch report will therefore focus on so-called first- and second-order impacts: direct environmental impact of equipment and facilities and the impact of policies and processes for control/use of equipment. In line with much current thinking we use the term 'low carbon' to

¹ Using 1990 levels as the baseline.

describe this primary focus on reducing GHG emissions, whilst acknowledging the difficulties inherent in this term (see below). In addition, whilst TechWatch's remit is usually to scan the 5- to 10-year horizon, this report takes a perspective over 30+ years. This reflects the long-term nature of current, primarily European horizon scanning, which will increasingly drive climate change abatement measures across national boundaries, and which already takes a view to 2100 (EEA, 2005). There is, however, a caveat with respect to outlying risk scenarios as technology is notoriously difficult to predict over such long timescales, not least because of the pace of change in materials development and so-called 'Black Swan' events.

1.1 How low can 'low carbon' go?

Consideration of the term 'low carbon' can often result in confusion. There are different targets depending on whether you take a global, European or UK perspective, and whether you look at 2020, 2030, 2050 or even 2100 – all different milestones on the emissions reduction roadmap. If the UK's target of a 34% reduction below 1990 levels by 2020 is 'low carbon', then how do we describe the >80% reduction the UK will need to make by 2050? In addition, how helpful is it to think of 'targets' in this way? The European Commission's commitment to a reduction of emissions to 20% below 1990 levels may increase to 30% depending on the outcome of the UN COP15 Summit at Copenhagen in December 2009. In fact, the European Environment Agency (EEA), commenting on the United Nations Framework Convention on Climate Change, concluded that:

scientific uncertainties do not allow unambiguous determination of concentration levels [of GHG] below which this condition [of a stabilised atmosphere] could be fulfilled. Furthermore, determination of 'safe' or 'sustainable' concentration levels is not only a scientific issue but is also related to the observed and projected impacts and to perceptions, values and political negotiations (EEA, 2005, p. 8).

In addition, the report makes it clear that the substantial, low-cost emission reductions anticipated from cuts in nitrous oxide and methane from industry, waste management and agriculture will have been almost fully exploited by 2030. Bigger reductions will have to be made elsewhere.

What this means is that the figures we currently take for 'targets', whilst useful as basic communication tools to raise awareness and put the scale of the change needed into some kind of context, are actually shifting sands, based on a variety of factors that are likely to change. Whilst it is quite likely that they will have to be revised more severely, it is unrealistic to think that they will go in the opposite direction. In the run-up to 2020 universities must become ready to act beyond the quick hits afforded by interim technology developments and to have started to pull together across the sector to implement much bigger plans for a longer-term future – to 2050 and even beyond.

1.2 Beyond targets

In the short term (to around 2013) it is inevitable that much of the focus will be on meeting the targets of the Climate Change Act (CCA) and the UK government's 5-year carbon budgets. However, in order to promote long-term strategic thinking and avoid becoming too fixated on targets there needs to be a bigger-picture vision in mind. The UK government's Transition Plan (DECC, 2009) provides a context for this:

By 2050, virtually all electricity will need to come from renewable sources, nuclear or fossil fuels where emissions are captured and safely stored for the long term. By 2020, 40% will need to come from these sources, of which, 30% will be from renewables (DECC, 2009, p. 14).

This statement helps to conceptualise the scale of what needs to be done and by focusing on making electricity production 'GHG emission neutral' the impetus falls equally on reducing demand and generating supply. This is a theme that will be developed over the course of this report, leading to the exploration of a long-term strategy for cross-sector planning and action within HE and FE.

This report has therefore been structured around two main ideas: looking at the technologies and standards that will help to maximise the energy efficiency of ICT within the individual higher education institution (HEI) over the short to medium term² (Section 4) and looking at the cross-sector strategic planning and action that needs to take place from around 2013 in order to prepare for 2020 and beyond (Section 5).

1.3 Summary and terms of reference

In general, TechWatch reports serve a cross-section of the FE and HE community. Senior decision-makers ought to be able to gain from the reports a good feel for the essentials and capabilities of the technologies, their implications and where they might be heading. In addition, parts of the reports might push beyond the technical levels of expertise of such readers. Institutional IT managers and technical staff should find the reports to be useful reference resources.

In brief, the report looks at:

- best practice measures and standards for metrics
- short-term 'quick fixes' based on simple staff actions and/or low-cost investment
- longer-term solutions that either represent a more costly investment, or are based on more experimental technologies
- discussion of the issues that may affect how these technologies develop in the future
- a first attempt at a Low Carbon ICT Roadmap, which puts these issues into a framework that also takes into account what is currently known about the targets associated with the CCA.

A detailed examination of the main technology areas is confined to Section 4. While the level of technical discussion may be beyond the expertise of the lay reader, Section 6 provides a summary of the main trends that can be gleaned from the technology developments discussed elsewhere in the section. Section 6 also presents the Low Carbon ICT Roadmap – a first attempt at providing a structured view of the technology within the target framework that is emerging from the CCA. It is hoped that this will give a feel for the rate at which change needs to take place and the kinds of technologies and policies that are likely to form part of achieving that change.

Section 5 overlaps these short- to medium-term measures by examining the cross-sector planning and action that needs to take place from about 2013, in order to take the sector to 2050 and beyond.

In particular, it looks at:

- the future of the data centre: cloud computing, the G-Cloud, adaptive data centres and nano data centres
- Enterprise Architecture (EA)
- energy security, in particular the role of DC power
- adapting to climate change: the second strand of the European Commission's strategy for dealing with climate change.

The report format uses the Harvard system to provide links to the references section at the end of the report, and footnotes to provide supplementary or explanatory material, often in the form of links to websites. In this way the report both references its evidence base and provides a resource of follow-up material. All links to Web resources are correct as of 3rd October 2009.

² Up to around 2020.

2. Policy state of play

Like the Internet, saving carbon can become part of how business is done: every finance officer knowing their savings and liabilities from carbon.

Ed Miliband, Secretary of State of Energy and Climate Change (DECC, 2009, foreword)

The drive towards low carbon ICT is not happening in a policy or strategic vacuum. It is one part of a national, cross-sector strategy for tackling climate change and environmental degradation that is being driven from the highest levels of government, and interest in climate change policy can be traced back at least to the year 2000, when the old Department of Environment, Transport and Regions (DETR) published its *UK Climate Change Programme*.

Subsequent work in the mid-2000s was led by the then-new Department for Environment, Food and Rural Affairs (DEFRA) and culminated with the publication of its *Climate Change Strategic Framework* in 2007. This noted the need for and scale of urgent action to address climate change and outlined the vision for a low carbon economy. It discussed the need for an international dimension to any lasting solution and noted some of the key international institutions for action: the Kyoto protocol to the United Nations Framework Convention on Climate Change and the European Union (EU). It makes clear that the public sector will have an important role to play in developments towards a low carbon economy and that technology, energy efficiency and technology transfer will be of vital importance (DEFRA, 2007).

These government-level discussions over climate change strategy came almost in tandem with the commissioning and publication of the *Stern Review on the Economics of Climate Change* by HM Treasury (Stern, 2006). This reviewed the economic consequences of acting or not acting on climate change and argued that the economic costs of stabilising the climate are "significant but manageable" and delay would be "dangerous and much more costly" (p. xvi), with the costs of acting in the order of 1% of GDP. As Adair Turner points out in his book *Just Capital: The Liberal Economy* (2001), this kind of figure is equivalent to a nation delaying getting slightly richer by a handful of months.

The weight of the review's evidence, the depth of research, the importance of the Treasury in UK policy terms and the leadership of an establishment figure in the form of Sir Nicholas Stern, a former Chief Economist of the World Bank, have all led to the report's enormous impact on government and public sector discussions. The government's key strategic response to the Stern Review was the Climate Change Act 2008, which was supported by other government strategies and initiatives such as the 2007 Energy White Paper from the Department for Business, and DEFRA's Climate Change Programme, published in 2006, which sets out policies and priorities for action in the UK and internationally.

The Climate Change Act

The UK's Climate Change Act (CCA) became law in the late autumn of 2008. The Government makes the claim that this is the first legally binding climate change framework in the world and it is the key plank in the UK's national response to the need to decarbonise our economy. The full details can be found on the DECC's website, but in summary the CCA:

- sets legal targets to reduce CO₂ emissions through domestic and international action by at least 80% by 2050 and at least 34% by 2020, against a 1990 baseline
- creates a new, independent Committee on Climate Change to review targets and progress, and to report to Parliament
- allows for five-year carbon budgets, which will set binding limits on CO₂ emissions that ensure every year's emissions count. The first budget runs from 2008 to 2012
- provides some of the legal framework for the previously announced Carbon Reduction Commitment (CRC), a new emissions trading scheme for large, non-energy intensive commercial and public sector organisations (see below).

As an indication of how far these issues have risen up the political agenda a new department was created in October 2008, the Department of Energy and Climate Change (DECC), with the remit to bring about the transition to a low carbon Britain.

By way of follow-up to the passing of the CCA, the DECC has produced a white paper, *The UK Low Carbon Transition Plan* (DECC, 2009), which sets out strategic plans and targets that need to flow from the legal framework established in the Act. Its approach is that carbon will be 'budgeted' in the same manner as finance, with an overall aim to cut GHG emissions by 459 million tonnes. Of most relevance to this TechWatch report is the section on the workplace (Section 5) and, within that, plans for strategic carbon reduction within the public sector. Whilst business workplaces in general will be expected to provide an overall reduction of about 13%, the public sector is expected to reduce by 30% of 1999/2000 levels by 2020.

The CRC is the government's principal driver for carbon reduction within the public sector and it is likely to have the greatest short-term impact on the university and college sector. This scheme will begin in April 2010 and will involve larger public sector institutions and organisations purchasing carbon allowances, the total amount of which will be controlled from 2013. It is not clear at this stage what the situation will be with regard to embodied energy (the energy expended in the manufacture and distribution of a product), and whilst it is possible that embodied energy costs are to be 'paid for' and accounted by others further back in the supply chain, at least initially, this cannot be taken for granted in the long term.

Revenues raised will be recycled to participants according to how they perform. In addition, by 2018 there is an ambition (although not yet a definite commitment) to ensure all new public sector buildings are 'zero carbon'. Each department and its associated public sector institutions will be expected to produce their own carbon reduction plans by spring 2010.

The Higher Education Funding Council for England (HEFCE) is (at time of publication) consulting across the sector as to the level of carbon reduction that is feasible and the attendant strategy for achieving such targets.³

2.1 European action

Europe has committed itself to a strategy of working towards limiting global climate change to 2°C and announced a plan at a summit in Brussels in 2007 (Europa, 2007). Planned actions include:

- improving the EU's energy efficiency by 20% by 2020, in line with the Energy Efficiency Action Plan announced by the European Commission in October 2006
- increasing the share of renewable energy, across Europe, to 20% by 2020
- putting in place an environmentally safe strategy to promote the industrial use of carbon capture and storage technology
- strengthening and expanding the EU emissions trading scheme
- limiting emissions from transport, e.g. cars, civil aviation and transport fuels
- reducing CO₂ emissions from other sectors, e.g. residential and commercial buildings, and emissions of other GHGs from a range of different sources
- significantly increasing the EU budget for climate, energy and transport research again after 2013, as has been done for the European Community's Seventh Research Framework Programme.

The latter two points are of particular relevance to the education sector.

³ See http://www.hefce.ac.uk/pubs/hefce/2009/09_27/

2.1.1 EU baselines

One of the important aspects of work at the European level is to look at the trends for energy use within the ICT sector across the EU. Two main scenarios for the development of technology until 2020 have been produced: Business as Usual (BAU) and Ecological (Eco) (Bio Intelligence Service, 2008). The Eco scenario assumes that there is a significant 'push', including policy, market and technology improvements, for ICT-related energy efficiency measures to be implemented. Examples of the kind of measures that are assumed in the Eco scenario include: widespread uptake of and continual improvement to the EU's Energy Star scheme; '80 PLUS' energy-efficient power supplies to be mandatory from 2009; power management to be set to function as compulsory on delivery of a computer or device (measures that are discussed elsewhere in this report).

To create these scenarios, a baseline figure for energy use across the EU was calculated in various subcategories of ICT and other electronics (TVs, mobile phones, data centres, etc.) based on statistics for 2005 (Bio Intelligence Service, 2008). The report calculated that the total EU use of energy for this broad range of ICT and electronics in 2005 was 214.5 TWh per annum (with subcategory figures of 42.3 TWh for computers/monitors and 29.1 TWh for data centres/specialist servers – the two that this JISC report is most interested in). Their prediction for 2020 for the BAU scenario is a total of 409.7 TWh (the two subcategories are 59 and 96 TWh, respectively). The Eco scenario gives a prediction of 288.2 TWh total with 46 and 69.8 TWh in the two respective subcategories.⁴

2.2 Strategy and policy within the computer industry

Whilst policy and strategy work has been ongoing within government there has also been activity in a wide range of industrial sectors. Aware of both the likelihood of government legislation and the personal views and commitment of company staff, many industries have already started to act. This includes the industry most relevant to this report: the computer industry.

1. The British Computer Society (BCS) has set up a **Carbon Footprint Working Group**,⁵ which will write 'cookbooks' of best practice for particular technical areas. These will be publicly available to all sectors.
2. **Climate Savers Computing**⁶ is an initiative that was started by Google and Intel in 2007, and is a non-profit group of eco-conscious consumers, businesses and conservation organisations. It was started in the spirit of the WWF's Climate Savers programme to demonstrate that reducing emissions is good business. The goal is to promote development, deployment and adoption of 'smart' technologies that can both improve the efficiency of a computer's power consumption and reduce the energy consumed when the computer is in an inactive state.
3. **The Green Grid**⁷ is a global consortium dedicated to advancing energy efficiency in data centres and business computing ecosystems. It consists of many computer companies including major players such as Intel, IBM and Dell.

Within the UK public sector the CIO Council (an organisation composed of Chief Information Officers for large government departments) issued a strongly worded statement in November 2007

⁴ The report puts the dramatic increase in data centre energy use (more than three times under the BAU scenario) down to predications of a massive increase in the number of end-user devices making use of remote software and data sources.

⁵ See <http://www.bcs.org/server.php?show=ConWebDoc.23759>

⁶ See <http://www.climatesaverscomputing.org/>

⁷ See <http://www.thegreengrid.org/home>

promising reductions in its own carbon footprint. The strategy to achieve this was delegated to the Chief Technology Officers (CTO) Council and, in turn, a Green ICT working group was set up, which now contains 16 government departments. This is tasked with producing:

- a strategy for achieving reductions with associated targets;
- a set of top tips for change which can be drawn upon by government departments;
- a scorecard, co-authored with Gartner, who will provide a benchmarking capability based on that scorecard.

A strategy, Greening Government ICT, was produced in 2008 (along with a number of charters and roadmaps).⁸ Of particular note is the proposal that by 2009 "all departments are to address and consider the impact on carbon emissions of all new ICT purchases" (Cabinet Office, 2008, p. 5).

The Society of IT Managers (SocITM, a membership organisation that acts as a consultant to local government) is working towards similar targets for local government and a sustainable procurement unit will be set up within the Office of Government Commerce (OGC, the government lead on procurement) to provide a procurement framework for central and local government.

2.3 Higher education

A number of projects and initiatives are taking place in the HE/FE sector. Examples include:

- The Higher Education Environmental Performance Initiative (<http://www.heepi.org.uk>)
- Sustainable Information Technology In Tertiary Education (SUSTE-IT) (<http://www.susteit.org.uk>)
- HEFCE, Sustainable development in higher education (http://www.hefce.ac.uk/pubs/hefce/2008/08_18/)
- JISC's Getting Greener Programme (<http://www.jisc.ac.uk/whatwedo/programmes/greeningict.aspx>)

⁸ Further details at http://www.cabinetoffice.gov.uk/cio/greening_government_ict.aspx

3. Metrics and measurement

The European Commission's funded report into the environmental impact of ICT predicts that the waste produced by the computer industry and by the disposal of end-of-life equipment, much of which reaches landfill or is incinerated, will double by 2020 based on current projections (Bio Intelligence Service, 2008). The problem is not only the sheer scale of this waste but also the failure to recover materials that have value and could be reused (for example copper) and the release of hazardous substances (such as lead) into the environment. It should also be noted that the practice of shipping industrial waste (such as old ICT equipment) to developing countries for 'recycling', where it can cause considerable health and safety problems, is widely condemned.⁹

As one example, consider a desktop PC. Much of the energy used by a PC over its full life cycle is actually embodied, i.e. expended during the manufacturing and distribution processes. What proportion of the machine's total carbon footprint this represents will depend on how many years the PC is used for and how much energy it consumes. Because a PC may involve several thousand components and the ICT industry's supply chain is highly fragmented, there is little information available on the footprint of this 'embodied' stage of the life cycle.

In order for users and purchasers of equipment to make informed decisions about the negative impact of ICT there needs to be adequate ways of measuring this impact and labelling equipment accordingly. There are two important concepts related to this: carbon footprint and asset life cycle management. The following sections will consider each of these in turn before going on to examine a shortlist of the most important eco-labels.

3.1 Life Cycle Assessment

In general, measuring the environmental impact of the full product life cycle is an extremely complex area. This quantitative process is commonly referred to as life cycle assessment (LCA)¹⁰ and formal standards have been developed by the ISO and other agencies. LCA includes the carbon footprint (see below) for products and services but also takes into account a wide range of other environmental factors including those related to their manufacture, distribution and disposal. For the computer industry this is especially difficult. One reason for this is the diversity of potential impacts. As one example, the US Environmental Protection Agency (EPA) undertook research into the full life cycle impact (including production) of LCD and CRT computer displays. It categorised 20 different significant areas of environmental impact including global warming, solid waste landfill, photochemical smog, water supply acidification and ozone depletion (University of Tennessee Center for Clean Products and Clean Technologies, 2001). The question then arises as to the relative importance of these different factors: is climate change 'worse' than acidified water supplies?¹¹

A second difficulty for the computer industry is that the manufacturing processes for computer-related equipment are highly complex, constantly evolving, and involve long and wide supply chains. A perennial problem is the lack of reliable data detailing all the processes in the production chain, often compounded by the aura of strict commercial secrecy within which much of the industry operates. By way of illustration, the fabrication of microchips alone consumes considerable amounts of water and tens to hundreds of different chemicals and energy. Williams (2003) provides a crude calculation regarding the memory chips for a single computer, showing that 94 kg of

⁹ See, for example, <http://www.computerweekly.com/Articles/2009/09/23/237843/criminal-gangs-dump-toxic-it-in-developing-world.htm> and <http://www.jstor.org/pss/3650919>

¹⁰ The European Environment Agency has published a guide at <http://www.eea.europa.eu/publications/GH-07-97-595-EN-C>

¹¹ More detail on the background to these issues can be found in Williams (2003).

embodied fossil fuel are required. This may not seem much, but the author points out that this is 500 times the weight of the chip and compares it to a car, which only uses twice its weight. Indeed, James and Hopkinson (2009a) note that the production of the basic components for computer equipment – integrated circuits – has the most significant energy impact of the production life cycle.

3.2 Carbon footprint

Carbon footprint is the total emission of greenhouse gases (GHGs) from a product across its life cycle (Carbon Trust, 2007). Carbon footprint falls within a general LCA and unlike other environmental measures it focuses entirely on GHG emissions. Carbon footprints are typically expressed in tonnes of CO₂ equivalent (tCO_{2e}), reflecting the fact that other chemicals such as methane and nitrous oxide are also responsible for atmospheric warming. Of particular importance in this regard is the British Standard Institute's Publicly Available Specification (PAS) 2050, sponsored by the Carbon Trust and DEFRA, which sets out to provide a consistent method for assessing the life cycle GHG emissions of goods and services. When reviewing emissions that form the carbon footprint, the BSI PAS says that impact should be assessed over a 100-year period from the creation of the product (BSI, 2008). The BSI has also produced an accompanying guide, *How to assess the carbon footprint of goods and services*, which details the process that a company or institution needs to undertake to assess the carbon footprint of its goods and services.¹² There is anecdotal evidence that the BSI is currently working on a specific version of its PAS for the ICT industry.

3.2.1 Carbon accounting

At the moment a basic approach to carbon accounting seems to be sufficient. However, this is a rapidly moving area and there is a need to develop and use rigorous standards (Stocks, 2009). There is anecdotal evidence that organisations undertaking some form of carbon accounting are tending to use a mixture of the details of their electricity bills and a spreadsheet. However, more sophisticated application packages are starting to appear, partly as an extension of the existing corporate social responsibility market. The UK analysts Verdantix recently published a report looking at this emerging field that reviews 20 leading providers.¹³ The UK Carbon Counting Group¹⁴ organises a regular series of workshops in this area.

3.3 Labelling and standards

There are a number of initiatives that label a product with information about its environmental impact. These eco-labels, as they are known, are intended to provide the purchaser with two fundamental forms of information: product and process. The former relates to the environmental characteristics of the finished product (for example its energy consumption when in everyday use). Process labelling describes the manufacturing or production processes that result in the finished article. The labels can also be subdivided as follows:

- Type I: Provides for a certificate of approval granted by a recognised third party such as a government agency (i.e. self-certification is not allowed).
- Type II: A manufacturer declares that their product meets certain independent standards (i.e. self-certification).

¹² Available from the PAS 2050 website at <http://shop.bsigroup.com/en/Browse-by-Sector/Energy--Utilities/PAS-2050/>

¹³ See http://www.verdantix.com/index.cfm/papers/Products.Details/product_id/51/green-quadrant-carbon-management-software-forthcoming-/- (requires subscription).

¹⁴ See <http://www.carboncounting.org.uk/>

- Type III: Designed to provide a set of quantitative environmental information direct to the consumer (e.g. electricity consumption) so they can use this to make their own value judgements.

This categorisation forms part of the family of environmental management standards that makes up ISO 14000.¹⁵ A detailed discussion of eco-labelling and the use of these three types can be found in Johansson (1999) and Johnson (1998).

However, it is important to note that it is difficult to certify and label computer-related products. As already discussed, the computer industry produces extremely complex products, involving highly sophisticated manufacturing processes and a large range of materials.¹⁶ There is also a very fast turnaround in product development (a new PC model, for example, may only be on the market for 6 months) and it is a highly competitive, price-sensitive market. This makes it difficult, time-consuming and costly for independent third parties to both develop new metrics to test against and then certify new products.

These problems have led to various criticisms of eco-labelling and a general criticism of type II schemes, for example, is that there are no independent audits – manufacturers undertake their own assessment and simply report to the certifying authority (although they also test their competitors). Another criticism, particularly of the EU eco-label, is the slow process of certification (Saied and Velasquez, 2003) and at the time of writing not one PC or portable computer in the UK market has the label, according to the EU's database.¹⁷ Further detailed discussion of some of the criticisms of labelling schemes can be found in Johansson (1999).

Despite their limitations, there are a number of prominent labelling schemes that are gaining traction and anyone considering the environmental impact of ICT should be aware of them. There are a number that focus heavily on the energy use of equipment, and which are particularly relevant to this report, and these are described below.

3.3.1 Energy Star

Eco-labels focusing on energy use started in 1992 with the EPA's Energy Star label. This is a Type II label that focuses on the operational use and standby phases of a product. The scheme is voluntary and has been adopted within the EU following a joint agreement between the US government and the EU Commission.¹⁸ The Energy Star scheme covers a growing range of domestic, office and industrial products, but of most relevance to this report is the part of the scheme that covers computer-related products (PCs, laptops, games consoles, etc.). The latest version is known as Energy Star 5.0 and was adopted for use within the EU in June 2009 (Europa, 2009).

For a product to qualify as Energy Star 5.0 it must meet various energy use guidelines in three distinct operating modes: standby, sleep and running. The AC to DC power conversion must operate at a minimum of 85% efficiency in order to reduce the waste heat generated during this conversion. The kinds of figures expected for a typical PC include: less than or equal to 2 W when running in standby mode, less than or equal to 4 W in sleep and less than 50 W when idling (i.e. just running basic operating system activities and not undertaking a user's tasks). The exact details are complex, depending on the type of PC (size of memory, etc.), and whilst Energy Star 5.0 does not

¹⁵ See http://www.iso.org/iso/iso_catalogue/management_standards/iso_9000_iso_14000.htm

¹⁶ Williams (2003) notes that a typical PC includes perhaps 1,500 different components using several thousand different materials.

¹⁷ See <http://www.eco-label.com/>

¹⁸ See <http://www.eu-energystar.org/en/index.html>

cover data centre-type servers and blades,¹⁹ work is ongoing on a separate labelling standard for server equipment. The work has been split into two tiers with a final draft of Tier 1 being released for comment in May 2009.²⁰ Work on Tier 2 started in September 2009.²¹

3.3.2 EPEAT

Where Energy Star focuses entirely on energy use, other labels offer a wider set of environmental impact criteria. A prominent example is the Electronic Product Environmental Assessment Tool (EPEAT) eco-label. This looks at the whole life cycle of electronic equipment (currently just desktops, laptops and monitors), including use of hazardous materials, design for disassembly and recycling, energy use and disposal/take-back.

The criteria that the EPEAT seeks to measure are defined in the IEEE 1680 standard. This defines 23 required and 28 optional criteria including process-related information such as the use of cadmium in the manufacturing processes, and product-related issues such as energy use when running. The Energy Star specification for desktops, laptops and monitors is incorporated into the core criteria of EPEAT.

Strictly speaking, EPEAT is a Type II eco-label, although independent verification is undertaken periodically. The verification process may simply require the manufacturer to provide production reports, lab analysis or other data, but EPEAT reserves the right to independently obtain products and inspect them.

EPEAT criteria are organised into three bands:

- Bronze: basic compliance to all required criteria
- Silver: compliance to 50% of the optional criteria
- Gold: compliance to 75% of the optional criteria

The US Federal Government has mandated the Silver level of compliance for 95% of its purchasing and it is rapidly becoming an accepted standard for procurement in the UK, with the Environment Agency aiming for EPEAT Gold where suitable equipment is available. EPEAT and IEEE are now extending their work to update the existing monitor, desktop and laptop standards to form an IEEE 1680 family of standards and to extend coverage to include other devices such as servers, printers, mobile phones and televisions.²² These IEEE processes are open to all and there is a need for more input from the HE sector.

3.3.3 TCO

TCO is a confederation of Swedish trade unions and because of this the 'TCO Certified' label encompasses health and safety issues as well as environmental criteria.²³ It is a Type I scheme that covers a variety of office-related products including display monitors, computers, keyboards, printers and mobile phones, although it has particular prominence in the marketplace because of its information on displays and monitors. It has had far less impact within the market for desktop PCs (only Dell and Fujitsu Siemens computers) and notebooks (no manufacturers).²⁴

¹⁹ Although ES 5.0 does cover what it calls small-scale servers.

²⁰ See http://www.energystar.gov/index.cfm?c=archives.enterprise_servers

²¹ See http://www.energystar.gov/index.cfm?c=revisions.computer_servers

²² See <http://grouper.ieee.org/groups/1680/>

²³ See <http://www.tcodevelopment.com/>

²⁴ Figures from TCO website. See http://www.tcodevelopment.com/pls/nvp!/tco_search

3.3.4 EU Eco-label

The EU Eco-label (also often referred to as the EU 'flower') is a Type I scheme that covers a very wide range of products. Criteria are established for each particular product grouping (e.g. appliances) and the label indicates that a particular product has been assessed and performs well against the criteria when compared to the rest of the market. The award is therefore selective in that only those products with the lowest impact in their particular category are eligible, with this being defined, informally, as the best 20–30% (Johnson, 1998). There is a product category for PCs and the criteria include meeting the Energy Star criteria (see the EU website for further details²⁵). Each nation has what is known as a competency body that assesses manufacturers against criteria and for the UK this is DEFRA.²⁶ Other EU-related work of interest in this area falls under the EU Directive on energy end-use efficiency and energy services.²⁷

3.3.5 The 80 PLUS[®] Program

This is a US-based initiative from the electricity utilities to help the computer industry move towards what they call '80 PLUS' power supplies: computers, savers or monitors that have an AC to DC conversion which is at least 80% energy efficient at various loads. Suppliers meeting this and other related criteria are awarded one of three labels (Gold, Silver, Bronze). The programme has recently begun expanding its work to cover high-end power supplies used in dedicated data centres and has published results of a study proposing that a third-party clearing house should be established and that the efficiency targets it proposes should be implemented in the marketplace (Ecos and EPRI, 2008).

3.3.6 Climate Savers

Climate Savers is a non-profit group of consumers, producers and conservation organisations (including the WWF) that was started by Intel and Google in 2007. The group is focused on persuading computer manufacturers to commit to producing products that meet specified power efficiency targets. The scheme is more than an eco-labelling system as it also has elements of an activist campaign, for example by seeking to educate and involve individual consumers to commit to using products that meet its criteria and to use the advanced power management features on these products.

The labelling aspects of its work fall within Type II. Climate Savers starts with Energy Star 5.0 as a baseline and intends to gradually increase the target over the next few years. Like many other labelling schemes it awards Gold, Silver and Bronze levels of attainment. Full technical details are available on their website²⁸ but as an example, the Gold level currently requires the most recent Energy Star in addition to a minimum of 90% efficiency rating for the power supply at 50% of rated output. Climate Savers also has criteria for the data centre server market.

3.3.7 The Standard Performance Evaluation Corporation (SPEC)

SPEC, a non-profit organisation focused on benchmarking computer technical performance, has been around for 20 years and has recently augmented its work with the launch of SPECPower, to add power/energy benchmarks. This work provides an efficiency benchmark suite that manufacturers can run to gain a quantitative efficiency level for an individual computer or server. This is quite specific and difficult to run, however, and the output is too mathematical for most

²⁵ See http://ec.europa.eu/environment/ecolabel/index_en.htm

²⁶ Further details at DEFRA: <http://ecolabel.defra.gov.uk/index.htm>

²⁷ See <http://www.euractiv.com/en/energy-efficiency/energy-use-efficiency-energy-services/article-133534>

²⁸ See <http://www.climatesaverscomputing.org/about/faq/#4>

buyers. There are also intellectual property issues with regard to its wider use and it is thus likely to remain a marginal way of measuring utility computing servers for a while.

3.3.8 Data centre metrics

Much of what we have discussed so far in this section relates to individual items of equipment, but metrics are beginning to be developed to look at energy efficiency at the level of the data centre. This is an emerging area and can be complicated, as a recent report published by the BCS notes (Newcombe, 2008a). The Green Grid, an industrial consortium, has proposed several metrics to quantify data centre power efficiency (Rivoire, 2008). These metrics focus on efficiency as measured over time and as a function of physical space. The Data Center Density (DCD) metric is defined as the ratio of power consumed by all the equipment on the floor of the data centre to the area of the floor (giving kilowatts per square foot/metre). The Data Center Infrastructure Efficiency (DCiE) metric looks at the proportion of energy consumption that is due to overheads such as cooling and power supply infrastructure, taking a proportion of the total energy used to that actually received by the ICT equipment in question. Newcombe, in his report for the BCS (2008a), argues however that whilst these kinds of metric are very useful as a starting point, they focus on determining a rather utilitarian amount of IT processing per unit of energy rather than asking "what is the financial and environmental cost of each IT service that my data centre delivers to support a business process?" (p. 5). Such a focus moves away from thinking purely about processing power to considering the actual service that is being provided and the wider benefits associated with it. This is an ongoing debate and IT managers/relevant JISC staff should engage with it.

Of particular interest in the UK is the work of the BCS's Data Centre Special Interest Group. Its model of data centre efficiency, jointly developed with Romonet and the Carbon Trust, is the most advanced model of its type to date, although the model has not yet been extended to the whole ICT environment and HE is in a strong position to assist in that work. The BCS recently published a review of implications, for data centre managers, of the proposed CRC. This discusses some of the detailed issues that data centre managers will face with the introduction of energy reporting for CRC purposes.²⁹ JISC/HEFCE should consider working with the BCS's Data Centre Special Interest Group particularly on their data centre efficiency modelling work.

3.3.9 Green IT maturity assessments

The metrics and benchmarks discussed above should also be considered in the wider context of an organisation's strategic and planning response to environmental concerns. One emerging technique is that of green IT maturity assessments. These function in a similar manner to more traditional maturity modelling tools, such as CMMI, which help organisations to assess their business process maturity and enable comparison with best practice. Several of these have now been developed by companies such as IBM and Accenture, with the UK government collaborating with Gartner in the production of the HM Government Green ICT scorecard. The Environment Agency has been working on a metric for measuring the sustainability of an ICT outsourcing arrangement, but it is not yet ready for release. One notable example is the Greener ICT Maturity Model which is being assessed by a number of local authorities including Birmingham and the Corporation of London, and which rates an organisation's strategic response to the 'greening' of its ICT using an online tool (Street, 2008).

²⁹ CRC and data centre carbon reporting, v1.0.0. Available at http://dcsb.bcs.org/component/option=com_docman/task_cat_view/gid.22/Itemid.50/

3.4 What to do now with regard to metrics

As outlined in the introduction, this report argues that, with respect to reducing the environmental impact of ICT, the initial task facing managers in HE/FE is that of meeting the requirements of the CCA. A reduction in the carbon footprint of day-to-day use is therefore the most important imperative. HEFCE has issued draft guidance to institutions as to the steps that need to be taken in order to respond to the CCA.³⁰ It notes that around 80 "larger institutions" will be involved in the CRC when it launches in 2010. Institutions are recommended to develop a carbon management strategy, which will include an element of baseline calculation (i.e. measuring the current carbon footprint of the institution). HEFCE recommends using the Carbon Trust's HE Carbon Management Programme, although this is a fairly basic approach and institutions or individual departments interested in a more detailed process of carbon footprint calculation are referred to the work of BSI PAS 2050 (see above).

The Carbon Trust recommends that for most organisations a basic approach to carbon footprinting will suffice at this stage of tackling climate change. It recommends identifying a handful of major emission sources for an institution including the on-site electricity usage. This may also be broken down by department or even individual computer and server rooms. The Carbon Trust provides a calculator for converting electricity use to a footprint measure on its website. At the time of writing, 1 kWh of grid electricity is considered equivalent to 0.537 kg CO_{2e} and, as a monetary comparison, £1000 of electricity bill in HE/FE is roughly equivalent to 8.6 tonnes of CO_{2e} per year. Readers interested in further details of carbon footprint calculation are referred to DEFRA's environmental reporting website.³¹

In the rest of this report we separate the short- to medium-term technology 'fixes' that institutions can make internally from the more long-term, cross-sector approach that will eventually need to dominate. The former is dealt with in the next section and concentrates on the technologies and standards that are relevant to the discussion on carbon footprint reduction and, in particular, day-to-day energy-efficient computing. A key part of this will be to take account of the energy efficiency of equipment. Good benchmarking can help in this regard and well-defined benchmarks are needed to help with purchasing decisions and to provide a standardised means of comparison.

³⁰ See Annex B, *Carbon management strategies and plans: a guide to good practice*, of the **Consultation on a carbon reduction target and strategy for higher education in England** at http://www.hefce.ac.uk/pubs/hefce/2009/09_27/

³¹ See <http://www.defra.gov.uk/environment/business/reporting/index.htm>

4. Technologies and standards

In this section of the report we look at the ideas, technologies and standards associated with achieving lower carbon ICT. These have been broken down into six categories: data centres; HE-specific, energy-intensive computing; storage; networks; end-user devices; printers. Each category has its own subsection which begins with an introduction to the existing energy efficiency problems associated with the particular area of technology. We then go on to discuss the 'quick fixes' that can be implemented almost immediately, based either on mature technologies or staff effort, depending on resources and local circumstances. Finally within each subsection there is a discussion of longer-term solutions that either represent a more costly investment, or are based on more experimental technologies.

4.1 Data centres and energy-intensive computing

Within the HE/FE sector there has historically been a mix of centralised data centre set-ups and disparate, departmental server rooms. More recently there has been a move away from server rooms towards more centralised data centres, which are easier to manage. Bringing equipment together in this way creates opportunities for energy efficiency and the greater the centralisation, the more opportunities to manage environmental impact.

These data centres are a source of intensive energy use and therefore carry a heavy carbon footprint. Large amounts of processor-intensive equipment are packed closely together in a series of racks, in rooms containing, in some cases, thousands of machines. All this processing generates heat and in order to keep the temperature within the data centre down to a safe operating level considerable energy is spent in cooling. Given this it is not surprising that Gartner (2007) estimate that data centres use approximately a quarter of ICT-related energy. Figures provided by a recent JISC report indicate that HE has an estimated 215,000 servers, which will probably account for almost a quarter of the sector's estimated ICT-related CO₂ emissions (James and Hopkinson, 2009b), and which matches Gartner's estimate.

In tandem, the requirement for data centres has been growing rapidly over the last few years due to the growth in online services, changes in architecture (the rise of so-called 'cloud computing', which moves applications into the data centre) and the more general growth of information and data storage requirements. In 2007, the EPA predicted a 75% growth in data centre energy use in the USA over a 5-year period (from 2006) under the "current energy efficient trends" scenario³² (EPA, 2007, p. 48), but their estimates already look like they are being exceeded: the Uptime Institute reports that when surveying the top tier of data centres, they have recorded a 20–30% *annual* rise in energy consumption (Brill, 2008; Miller, 2008).

Indeed, there are some concerns that, notwithstanding the emissions issue, there may be problems with the basic provision of electric power to growing data centres. The Uptime Institute found that 42% of the 311 US data centre managers they recently surveyed expected to run out of electricity within 2 years, and 39% said they will exceed cooling capacity within the same period (Courtney, 2008). Its survey did not reveal any companies who felt that they were not running out of capacity – the differences reported in the survey were just in the amount of breathing time organisations felt they had before major works were required. In the UK there has been some debate about these kinds of supply issues, for example in the crowded south-east.³³

³² The EPA made a series of fairly conservative assumptions about trends in energy efficiency for this scenario, e.g. power management is enabled on 10% of servers.

³³ Next Generation Data (NGD), for example, used this justification in part for plans for an enormous data centre in Wales (see <http://news.techworld.com/operating-systems/105403/ngd-builds-uks-first-data-fortress/>).

4.1.1 The state of play: problems with data centres

Energy-related problems with data centres are of two fundamental types: **primary consumption** by the equipment that is installed in the facility and **secondary consumption** by the infrastructure services, including cooling of the IT equipment and distribution of power around the data centre. These latter issues are affected by the fact that data centres are often designed with a 10- to 20-year life cycle, which means that the pace of new technology adoption is consequently much slower than for ICT in general and that data centres inevitably lag behind the technological state-of-the-art.

The sheer scale of a data centre generates huge inefficiencies in the power delivery chain and as with all electricity distribution there is considerable power loss between the power station and the data centre building.³⁴ However, once the building has been reached there are particular problems relating to secondary consumption. Distributing the electricity around the centre as alternating current (AC) and converting to direct current (DC) at the equipment level consumes additional energy that is dissipated in the form of heat, which then has to be removed to keep the equipment at a safe operating level. Finally, providing infrastructure services such as cooling and uninterruptible power supply (UPS) consumes considerable amounts of electricity, which means that, overall, only about 7% of the power generated at the power station is actually consumed by the servers (Mingay, 2007; BT, 2009).

4.1.1.1 Uninterruptible power supply (UPS) systems

These are designed to plug the gap between a power failure and a standby generator kicking in so that equipment in the data centre does not fail. They also have the secondary function of providing some element of power conditioning (i.e. protecting equipment from over- and under-voltage and spikes on the grid supply). They are recognised as a source of energy inefficiency in the data centre set-up and it has been estimated by the US Department of Energy's Lawrence Berkeley National Laboratory (Berkeley Lab) that such efficiency losses in UPS represent about 5–12% of all the energy consumed in data centres.³⁵ Within HE, these systems usually rely on some form of battery back-up³⁶ and there are two main issues to consider:

- The inefficiency of charging and discharging batteries, and the environmental impact of those batteries (e.g. considerable amounts of high environmental impact material such as lead). Generally a 5-year replacement cycle is required, presenting a significant waste disposal issue.
- The inefficiency of converting AC power (as delivered to the data centre) to DC power (to charge the batteries), and then converting back to AC for distribution to the consuming IT hardware.

4.1.1.2 Cooling

The data centre's layout and the provision of cooling systems is often a source of considerable inefficiency. A few years ago a typical data centre rack would have equipment that drew 1–2 kW. Nowadays this is more likely to be around 12–20 kW and rapidly heading towards 35 kW. Even carefully designed racked blade systems using low-voltage components can consume up to 22 kW. These levels of power consumption produce significant heat that has to be disposed of in order for the servers to function within a safe temperature operating range. These are typical figures and it

³⁴ Around 7% according to BT (2009).

³⁵ See <http://hightech.lbl.gov/ups.html>

³⁶ Non-battery systems are generally only suitable for very large data centres that would normally not be found within the current to near-term HE set-ups.

should be noted that within the HE sector these levels are currently not reached,³⁷ although cooling is still a significant issue, with a report by SusteIT calculating that it is common for between a quarter and half the energy supplied to a data centre to be used in cooling processes (James and Hopkinson, 2009b). Indeed, in poorly designed data centres this ratio can be more than half.

4.1.2 Short-term fixes for data centre problems

There are a number of technologies and ways of working that data centre managers can begin to implement almost immediately, depending on the local situation. Such short-term solutions and best practices are elaborated on in a number of best practice guides, most notably the EU Code of Conduct on Data Centre Energy Efficiency (JRC Institute for Energy, 2008a). The EPA's report on data centres also provides detailed suggestions for best practice and state-of-the-art scenarios (EPA, 2007) as does the Green Grid consortium advice,³⁸ but these are both based on American experiences and some significant differences exist between European and US power and cooling systems.

To assist with the process of benchmarking and assigning efficiency there are a number of sources that may be of interest. Freely available tools are available from the BCS's Data Centre Special Interest Group (currently out in beta test³⁹), the result of a joint initiative between the BCS, Romonet and the Carbon Trust. In the USA, the National Data Center Energy Efficiency Information Program also includes a focus on certification and benchmarking⁴⁰ and the Department of Energy publishes case studies of best practice in energy saving for data centres.⁴¹ The University of Oxford recently commissioned a report on the design of a central data centre⁴² that may be of interest to readers as it includes a number of near-term recommendations drawn from the EU Code of Conduct and discussions between estates and ICT staff.

It is recommended that data centre managers study and implement the best practice guides provided by the EU and institutions may also decide to formally sign up as participants to the EU Code of Conduct, which involves a process of committing to various energy-saving best practice activities and can be used to assist with procurement processes.⁴³ However, due to the amount and complexity of advice that is available, the rest of this section provides a summary of short-term solutions (which either require reorganisation or investment in relatively mature technologies rather than representing technology challenges) and mid-term, experimental work that should be tracked as a source of potential future solutions. For data centres, these suggestions are based around the two fundamental sources of inefficiency in the data centre: primary and secondary consumption.

4.1.2.1 Primary consumption reduction: virtualisation and replacing hardware

For the first type of problem a simple hierarchy is helpful for taking action:

- decommission (i.e. if services or equipment are not being used, then they should be turned off and removed from the data centre)

³⁷ For example the Oxford Super Computer Centre has a machine room that was designed to have its cooling capacity limited to 11kW per rack. This is sufficient for current HPC needs, but as the centre strives for greater performance it is likely that there will be a move to greater energy density of the same order as that of the commercial sector (perhaps 20kW).

³⁸ See <http://www.thegreengrid.org/>

³⁹ See <http://dcs.org/content/view/full/45/59/> for more details

⁴⁰ A factsheet is available at http://www.energystar.gov/ia/partners/prod_development/downloads/NDCFactSheet.pdf

⁴¹ See http://www1.eere.energy.gov/industry/saveenergynow/case_studies.html

⁴² See <http://www.oucs.ox.ac.uk/greenit/oxford-central-machine-room-design.xml>

⁴³ The BCS is about to release an ISEB intermediate qualification in the EU Code of Conduct for Data Centres. The course and examination provides a practice guide over its 3-day duration on how to implement the code's best practices within your own data centre.

- consolidate (i.e. co-locate multiple services onto one server, for example host multiple websites or databases on a single piece of hardware)
- virtualise
- use more efficient hardware.

The idea behind the list is to make the smallest changes that realise the greatest benefits first. For example, there's little point in virtualising a server that should have been decommissioned 6 months ago.

a) Virtualisation

Over time the idea of a single machine running a single version of the operating system and supporting a small number of applications, more familiar as the client-server model used in the standard PC set-up, has been carried through to the data centre. It is now common to see situations where a very powerful but under-utilised server is running its own operating system and a single application.

However, with the significant increases in the processing capability of modern hardware has come the idea of being able to virtualise this kind of set-up: in essence to 'partition' a single piece of hardware to produce the effect of several individual machines running at once, which therefore automatically provides the required capacity. A single, powerful machine runs a piece of virtualisation software, often referred to as the hypervisor, which creates several virtual machines (VMs) each of which replicates, in software, a computer's hardware functions (and therefore the functionality of what was traditionally a separate user machine). The user interacts with a particular VM via some form of control panel on the user's screen.

This is not a particularly new idea, having been first used with the mainframe machines of the 1960s, but it was effectively abandoned following the widespread adoption of the client-server model. However in recent years the proliferation of powerful but under-used desktop PCs, along with departmental server rooms and central data centres, both with under-used server capacity – sometimes referred to as 'server sprawl' – has raised questions about the wisdom of this model and there is a growing interest in using the VM technique. This can target two different types of end use. Firstly, virtualisation of servers within a data centre, such that a single physical machine can replicate several servers. Secondly, virtual desktops that virtualise an individual user's desktop PC on a central, shared machine and either replace the desktop PC with a thin client or retain the existing PC for localised processing of a centralised application (through a process known as 'desktop streaming') (Gruman, 2006). However, whilst it is obvious that a process of deep consolidation of under-used equipment through virtualisation can contribute significantly to energy reduction, it is less clear with desktop virtualisation in which the PC is retained.⁴⁴

The key advantage of all this is that processing power, storage and networking requirements can be decoupled from the underlying physical equipment and provided as a 'virtual infrastructure service'. This means a set of services can be provided over the most efficient hardware and that individual VMs can be moved around between different hardware configurations as easily as moving around a software file. In this way, dozens of poorly utilised machines can be replaced by a few heavily loaded servers: a typical, virtualised server runs at 65–85% capacity compared with average server usages of 7–20% (IT Pro, 2008). Indeed, these figures may be even worse in some public sector organisations. The Environment Agency, for example, found these figures to actually be quite high when they undertook a 'real world' survey (unpublished) of their own server utilisation prior to a programme of virtualisation. They found most servers were utilised at less than 5%.

⁴⁴ In which case the benefits relate more to the management of multiple desktop PCs and efficiency of being able to provide centralised support.

However, there are up-front costs associated with virtualisation, which vary according to the type and scale being undertaken, but might include such things as the purchase of thin clients, new or extended server hardware, staff retraining and VM licensing. These costs can be substantial and there remains some controversy over the cost-benefit ratio of virtualisation,⁴⁵ with Forrester ranking server virtualisation and the introduction of thin clients as involving quite a high level of 'resource intensity' in a comparison of the cost-benefits of various green ICT techniques.⁴⁶

These concerns notwithstanding, virtualisation is recommended as a priority by the EU Code of Conduct for Data Centres and, indeed, Liam Newcombe of the BCS Data Centre Working Group points out that the Code goes as far as to say that senior management approval should be sought before any new business application is provided with a *non*-virtualised server facility in order to discourage server sprawl (Newcombe, 2008b).

Although virtualisation is a technical approach that is relatively mature it is likely to develop further in the next few years and so it is also covered in Section 5.1.4.

b) Use more efficient hardware

Once organisations know they are using the minimum of physical equipment, it is cost-effective to look at the types of equipment being used. There are a number of possible optimisations for server equipment:

- Prioritise energy-efficient equipment when undertaking procurement processes.
- Consider the operating temperature range of equipment and take account of the EU guidance for equipment to be purchased from 2012.⁴⁷
- Tackle power supply inefficiencies by, for example, using systems with lower-voltage processors (which can halve CPU energy use).
- Use systems with shared components such as power supplies, networking, fans, etc. These are commonly packaged as 'blade' systems.
- Use systems with active component management: for example only powering fans that are required and controlling their spin speed to optimise energy use.
- Use systems that can partially shut down in low-load situations (for example systems that shut down processor cores).

However, there is a note of caution to be struck with regard to the use of blades. While they may seem to provide solutions to problems such as energy inefficiency and scalability, proper comparisons need to be made between existing, discrete server systems and blade alternatives. If a blade system is chosen and operated at less than full load for the first 2 years of its life the user could find themselves actually drawing more power than the sum of the replaced discrete systems. While blade systems are improving they should not be treated as the default solution to all issues.

4.1.2.2 Secondary consumption reduction: cooling and power delivery

By tackling the primary consumption problems of the ICT equipment installed in the data centre the second type of energy utilisation problem, the power consumed by the infrastructure services, is automatically reduced. However, bearing in mind that most power consumption occurs during distribution and cooling, it is not enough just to maximise efficiency at the primary consumption level. In fact, it is always important to maintain a careful balance between the equipment producing heat (the servers) and the cooling infrastructure. It is therefore important, for example, to review the

⁴⁵ See http://www.pcworld.com/businesscenter/article/131683/gartner_says_virtualization_too_expensive.html

⁴⁶ See <http://www.avayaglobalconnect.com/greenintelligence/forresterreport2007.pdf>

⁴⁷ Starting in 2012 new servers should be able to withstand the extended air inlet temperature and relative humidity ranges of 5°C to 40°C (exceptional conditions up to +45°C) (JRC Institute for Energy, 2008a).

cooling equipment when undertaking a process of server consolidation that might result in decommissioning.

a) Cooling

There are currently two different types of cooling undertaken within a data centre, either **air volume cooling** or **liquid (normally water) rack cooling**. The primary differences between these two types of cooling are their cost and the capacity for system density, with water cooling able to run much higher energy density within individual racks.⁴⁸ The vast majority of HE sector data centres are air cooled, with the choice often being dictated by the physical space available for the data centre, because water cooling requires less space but is more expensive to fit out. However, even where space is not an issue, water-cooled systems are beginning to be installed in the higher end of UK HE data centres, two recent examples being Bristol⁴⁹ and the national supercomputer facility, HECToR (High End Computing Terascale Resources), based at the University of Edinburgh.⁵⁰

Data centre facilities are usually specified and laid out in accordance with the Telecommunications Industry Association's TIA-942 standard. This not only defines the level of operational resilience, tier and features, but also defines the optimal layout for an air-cooled data centre.⁵¹ The standard includes technical specifications for encouraging airflow and reducing the amount of heat, and recommends using a raised floor of tiles and arranging the equipment racks in an alternating pattern of hot and cold aisles. In the cold aisle, equipment racks are face-to-face and the floor tiles are perforated, and in the hot aisle they face away from each other and the floor tiles are whole. Cold air is drawn up through the holes in the floor in the cold aisles, passes through the equipment and then comes out in the hot aisle. A typical air-conditioned data centre, with underfloor supply of cold air from a computer room air conditioning (CRAC) unit and a layout in parallel lines, alternating hot aisles and cold aisles, is shown in Figure 1.

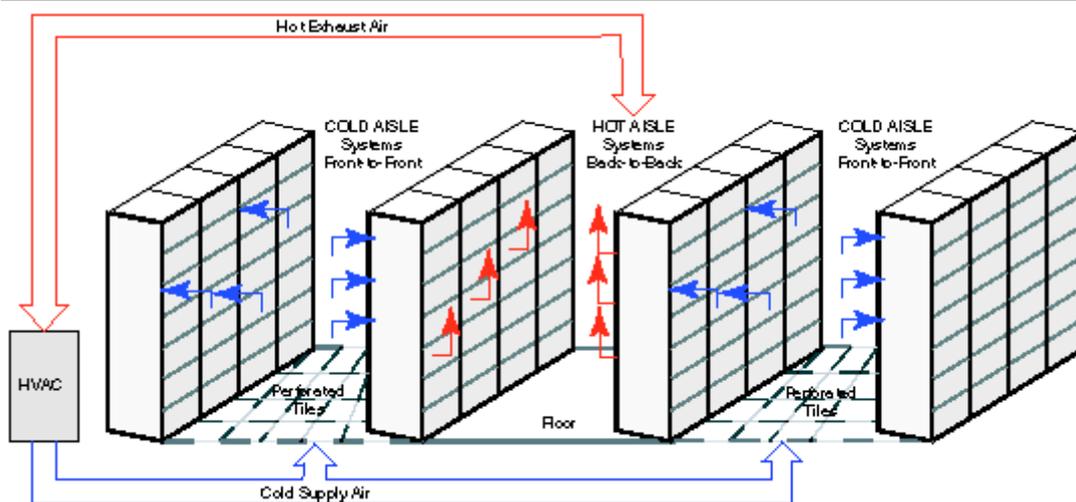


Figure 1 Typical TIA-942 air conditioned data centre

However, HE server rooms are often badly laid out, with hot air from one aisle of servers being drawn straight into the intakes of the next aisle. This is obviously inefficient, and often results in the need for additional air conditioning, which compounds the problem. It is also worth pointing out

⁴⁸ Because water carries away heat around 3,500 times better than air.

⁴⁹ See <http://www.bristol.ac.uk/news/2008/212017945293.html>

⁵⁰ See www.jisc.ac.uk/media/documents/publications/greenict-edinburgh.pdf

⁵¹ An overview is available from <http://www-wsp.adckrone.com/eu/en/webcontent/support/PDFs/enterprise/Generic/102264BE.pdf>

that the EU Code of Conduct on Data Centres notes that "facilities are often overcooled with air temperatures... colder than necessary resulting in an energy penalty" (JRC Institute for Energy, 2008a, p. 15). Institutions should make considerable effort to ensure their data centres are not wasting energy in this way.

Data centre managers looking to make immediate improvements should look to manage the airflow so as to optimise the cooling process. One simple measure is to seal the underfloor void in the aisles by moving floor tiles, using grommets around cables, etc. Once the floor is sealed, gaps in racks that let air flow freely between the hot and cold aisles can be filled with blanking plates.⁵² As a further measure, many data centres over the last 2 years have physically isolated the hot aisles from the cold aisles using plastic curtains at the end of aisles and transparent aisle roofs. This prevents turbulence and bleed of hot air back into cold aisles, and reduces the chance of air circulation not going right back to the CRAC plant. Within the HE institutional environment these types of simple solutions can be extremely cost-effective as well as being quick to install. A further simple solution, put forward in the EU Code of Conduct for Data Centres, is to deploy groups of equipment with substantially different environmental requirements in separate areas of the data centre with separate airflow and cooling provision.

More complex solutions make use of fluid dynamics to study the airflow within a data centre and then incorporate **active management technology** to control airflow. These systems will automatically adjust the cooling delivered to the data centre based on thermal requirements.⁵³

Related to managing airflow, **fresh air cooling** is extremely popular: taking the hot exhaust air, mixing it with fresh air from outside the building (if this is at the right temperature), and only firing up the CRAC plant when necessary. Newcombe (2008c) provides detail of this technique and its associate, air-side economisers. A case study, with an associated return on investment, can be found in Wilman (2007). This shows that appropriate use of computer-controlled outside air can save around 70% of the energy costs of running air conditioning. It is also worth noting that fresh air cooling has been applied extensively and successfully in 107 of BT's 21st Century Network data centres, saving over 85% of cooling costs (BT, 2009).

There are also alternatives to the TIA-942 style air-cooled data centre – other cooling models that can be more efficient, especially in environments with 'hot spots' of particular heat density (e.g. racks of blade servers). Several manufacturers have released racks that provide assisted cooling (e.g. cooling units and fans on the rear doors of equipment cabinets), and these offerings have now morphed into four distinct ranges of totally contained environment, requiring no external cooling input:

- hermetically sealed rack with integral cooling (e.g. HP's Freon-based Modular Cooling System);
- a self-contained aisle with integral cooling (e.g. APC's InfraStruXure);
- a portable data centre in a shipping container including fully integrated cooling (e.g. Sun's Modular Data Centre);
- systems that are water cooled within the blade chassis.

This equipment can produce some of the most efficient data centres available at present but they do have constraints, mainly that of having to introduce water plumbing in some cases. Another option for the smaller, departmental type of server configuration is to do away with the idea of a separate server room altogether and use cupboard-style data centres that make use of air from the

⁵² Subject to having equipment that cools by taking air from front to back. Some legacy equipment doesn't work like this.

⁵³ For example, AdaptivCool provides a product that places dynamically controlled airflow regulators/boosters in the underfloor void.

surrounding general office (these are thus a variant of fresh air cooled systems). In the UK, Kell Systems produce such a range and, anecdotally, these may have been piloted in a number of universities.

b) Power delivery: voltage conditioning and phase balancing

The EU Code of Conduct for Data Centres (JRC Institute for Energy, 2008a) makes clear the importance of power delivery equipment and notes the substantial impact upon data centre efficiency of such equipment, which also tends to stay in operation for many years once installed. It argues that careful selection of the power equipment during the design process can deliver substantial savings through the lifetime of the facility. One area where efficiency savings may be achieved in the very near term in an existing facility is through conditioning the power supply.

Voltage conditioning and phase balancing products were first developed in Japan and the basic idea is that they optimise voltage, and therefore energy efficiency, by dealing with the discrepancy between the actual supply voltage you receive (207–253 V) and the optimum voltage your electrical equipment needs (220 V). They include products such as PowerPerfector and EMS PowerStar.

To date these types of products have not been widely used in HE data centres. This may be down to a lack of awareness or because in some cases of ICT power supply, known as switched-mode, they are not always suitable. However, in a case study of a PowerPerfector implementation at the University of Surrey the company claimed savings of around 11%.⁵⁴ A number of other companies also supply this form of product and most will assess the power supply beforehand to indicate what cost savings can be expected if a device is installed. The Environment Agency has recently installed such a product at six of its sites in the north-east.

4.1.3 Near-term developments

There are a number of technologies and techniques that are in the experimental stage, with a small number of early adopters undertaking work to explore their applicability. Should these experiments prove successful then it is likely that institutions may be looking to implement them in the next 3–5 years.

4.1.3.1 Uninterruptible power supply (UPS)

As already mentioned, UPS is a source of energy inefficiency in the data centre set-up and whilst differences in loading conditions and test procedures mean that manufacturer specifications can differ widely from measured results, Berkeley Lab estimates that efficiency losses in UPS represent about 5–12% of all the energy consumed in data centres. Near-term developments in UPS technology focus on reducing the inefficiency of charging and discharging batteries and the related environmental impact of those batteries, and there are two technologies worthy of mention:

a) Flywheel UPS systems spin up a large heavy wheel, in a vacuum, while the power is on, and then use very little power to maintain its spin speed. If the power goes off, the inertia of the flywheel keeps it spinning at high speed and this generates electricity to drive the servers, although it only maintains the power long enough for a reserve power system such as a diesel generator to take the strain. A slight variation on this uses the flywheel as back-up to a battery-based UPS, using the wheel's inertia to provide energy to iron out any slight glitches in the battery supply, which prolongs battery life (Fontecchio, 2006). The other significant disadvantage for both systems is their physical size, which would probably make them impractical in most HE server room settings,

⁵⁴ See <http://www.powerperfector.com/university-of-surrey.htm>

although it may be worth considering for larger data centres and high performance computing (HPC) set-ups.

b) Fuel cell power systems are generally considered to be a more long-term technology and are covered in more detail in Section 5, but one example that needs to be mentioned here is APC's hydrogen-based fuel cell for data centre use. It provides up to 30 kW, which the company argues is sufficient for a pair of server racks (Dubash, 2007). The company claims that the fuel cell takes about 30 to 40 seconds to come into operation once a power supply problem has been detected – about the same as a standard generator. It then keeps the data centre's UPS batteries charged until the mains can be restored.

4.1.3.2 Cooling

As already discussed, the process of cooling equipment consumes considerable energy within a data centre. When designing new data centres, or retrofitting old ones, there are two aspects to consider: firstly, to design so as to minimise the percentage of input energy that you actually use for the cooling and secondly, to use the heat by-product once it leaves the centre.

a) In-row and in-rack cooling

Traditionally, cooling has been designed at the room level, but newer data centre designs are starting to explore a more focused approach in which cooling is more closely integrated with the server equipment (Dunlap and Rasmussen, 2006). This type of system introduces cooling apparatus at the level of a row of racks or directly onto the front and rear of the rack systems. So far there has been only a limited number of trials within the HE sector.

Other solutions include the installation of specialised chassis that allow through-system cooling by removing the individual cooling fans from within the server chassis and putting them into the rack itself with innovative solutions such as vertical airflow. The disadvantage of these types of system is that they can only be used with specific server chassis.

b) Running the data centre at higher temperatures

Various companies and institutions are experimenting with a process of running data centres at temperatures that are theoretically not supported by the equipment inside. The key standard in this regard (even in the EU) is set by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), who provide the 2008 ASHRAE Environmental Guidelines for Datacom Equipment. This recommends an operating range of 18–27°C and various levels for humidity.⁵⁵ However, as noted earlier, EU best practice is to plan to purchase new equipment (from 2012) with a wider range of operating characteristics (5–40°C). It is worth noting that, under pressure, many equipment manufacturers have revised the acceptable operating temperatures of their equipment and institutions should check maximum possible temperatures with their equipment manufacturers.

Others are pushing further. As a consequence of ASHRAE and the generally accepted wisdom on what ICT equipment can handle, data centres have traditionally been built to support operating environments of between 20°C and 24°C and a prescribed humidity level. BT's experience is that equipment can run at higher levels than these, for example at up to 50°C (Morgan, 2007). Other companies such as Dell, Google and Intel have also been trying higher temperatures or decreased humidity control in an effort to determine ways to avoid unnecessary energy costs. However, not everyone is convinced of the wisdom of this approach.⁵⁶ It should be noted that pushing standard

⁵⁵ A certain level of 'dryness' is required in the data centre to avoid the build-up of static electricity.

⁵⁶ For a discussion of some of the issues see

<http://processor.com/editorial/article.asp?article=articles%2Fp3109%2F37p09%2F37p09.asp&guid>

commodity IT equipment over 35°C can in some cases lead to disproportionately higher energy consumption due to internal fans spinning up to maximum. This is particularly true in 1U 'pizza box' devices.

Related to this are a number of experiments to explore the practicality of a complete absence of air conditioning in a data centre environment. Intel, for example, took a data centre in the middle of the desert and set up one room as a standard (control group) data centre. The other was set up to have no air conditioning, but to use an air economiser to pump in unfiltered and unhumidified air from outside (complete with desert dust). The room with the economiser used 67% less energy than the control room. The reliability of the equipment in the rooms was not hugely different, suggesting that modern servers and switchgear are built for much harsher conditions than advertised.⁵⁷

4.1.3.3 Use of DC power

Getting power into and around the data centre is a source of considerable inefficiency. Electricity is distributed over long distances by the use of alternating current (AC), but much of the computer equipment within the data centre uses direct current (DC). Generally, data centres tend to distribute the power around the centre using AC and then convert to DC at the equipment level (although several varying and detailed power distribution methods are used⁵⁸). A typical PC or server will include a power supply box that converts AC from the mains into three DC voltages (3.3 V, 5 V and 12 V). The integrated circuit-related parts of a PC or server (e.g. the CPU motherboard) make use of the 3.3 V and 5 V, while the higher voltage is required by the disk drive motors (Mueller, 1999). The conversion of AC power to DC is one source of inefficiency because each conversion uses energy and this is dissipated in the form of heat, which then has to be removed.

In recent years there has been debate over the merits of AC versus DC power within the context of a data centre. Proponents of what is sometimes called the DC power architecture, argue that eliminating the AC to DC conversions can generate energy savings of between 10% and 20% (Chen, 2006), with the added bonus that DC can be 10% more efficient at transmission around the centre (Hughes, 2007). One example of the DC power architecture approach has been provided by the DC Power for Data Centres of the Future project which involved industry and the California Energy Commission.⁵⁹ Another involves BT, the largest corporate user of electricity in the UK, which is investigating the issue as part of their 21st Century Network programme. BT estimates that the overall grid-to-CPU loss (i.e. loss within the actual data centre after the electricity has arrived from the grid) is 60% for AC supply and 35% for DC (O'Donnell, 2007).

Proponents of DC point to the example of telecom data centres where 48 V DC is effectively a global standard. These centres tend to do the AC to DC conversion in one go and then feed to the equipment using DC around the centre to avoid the inefficiency of many smaller AC to DC conversions (Dubash, 2006). A number of leading server suppliers offer products that accept DC power in this manner and so it has been argued that it is not completely impractical for non-telecoms data centres to move over to DC to deliver power around the data centre (Mitchell, 2005; Harbaugh, 2009).

⁵⁷ Video and white paper of proof of concept are available from <http://ipip.intel.com/go/1652/intel-its-air-economizer-proof-of-concept/>

⁵⁸ There are potentially various power distribution methods and the interested reader may find the Green Grid report "Qualitative Analysis of Power Distribution Configurations for Data Centres" to be useful, although it focuses mainly on US and Canadian systems, which differ to the UK in that they make use of 480 and 600V AC at the mains. A summary of the report is available from http://re.jrc.ec.europa.eu/energyefficiency/pdf/IEECB08/IEECB08%20proceedings/096_Green%20Grid_final.pdf

⁵⁹ See <http://hightech.lbl.gov/dc-powering/>

However, many operations staff within the data centre industry are reluctant to make such a fundamental change without detailed evidence of its effectiveness and indeed safety (Chen, 2008). This is a source of considerable debate. A particular issue is that 48 V DC needs to be transferred around the data centre using a high current. This requires the use of large copper power cables that are expensive. For this reason the computer industry is considering the use of high-voltage DC (300–400 V) (Lidstrom, 2007), but this brings with it concerns about safety and requires the introduction of a new set of cable and connector standards that do not yet exist.

To help with these debates, researchers at Berkeley Lab set up experiments with Sun Microsystems, Intel, Cisco and others to demonstrate the DC concept and look at related technologies. The Green Grid provided a detailed critique of this work, which readers may find of interest.⁶⁰ The precise details of these kinds of experiments can get somewhat complicated because, as noted above, there are ongoing explorations as to precisely what voltage of DC is most efficient for this process. The complexity of all this is compounded by the North American slant of much of this work. Within the EU, AC to high-voltage conversion offers less scope for energy efficiency improvements than that provided by US systems. Despite this caveat it is recommended that data centre managers keep abreast of these developing debates. They may be particularly pertinent if widespread adoption of DC-based renewable energy generation takes off (see Section 5).

Google has been active in these power-related areas. One of its concerns has been that the motherboards in their equipment (which are based on pretty standard, off-the-shelf consumer PCs running Linux, rather than dedicated servers) need more than one DC voltage level – a source of inefficiency. They have worked directly with the hardware vendors to make sure that all the components on the motherboard use a single voltage level (Hölzle, 2005), although others in the industry point out that dedicated server manufacturers are going down this route anyway (Branscombe, 2008). Google has also extended this model to dynamically adjust power supply (Fan *et al.*, 2007), depending on what a device's components are doing at a particular moment. This is based on a process known as Dynamic Voltage Scaling, which was first introduced as a technique for mobile computing (Zhai *et al.*, 2004).

4.1.3.4 Shared data centres

As well as consolidating technology and services within an institution, some organisations are also starting to explore the potential for shared data centres. Under the auspices of HEFCE the universities of Derby, Salford and Sheffield Hallam recently undertook a study for the potential of such a centre⁶¹ and Yorkshire and Humberside Metro-area Network are looking into a virtual data centre under the same series of feasibility studies. This can be done via a commercial partner (using a partial or total outsourcing model) or by mutually agreed collaboration between like-minded organisations. However, there are significant hurdles that must be overcome for this to be satisfactory for the HE sector. Firstly, the Data Protection Act and other data management best practices impose a number of regulations on universities about who has control over sensitive personal records of students – information that must be kept for up to 10 years. Data centres housing this information have to be kept within UK legal jurisdiction.

The other problem inherent with a shared data centre is that while the implications may not be so critical for business systems (of which a single HE institution may have a largish number), for more specialist services such as research equipment control and supercomputing, then their localisation

⁶⁰ See The Green Grid Peer Review of DC Power for Improved Data Centre Efficiency at http://doe.thegreengrid.org/files/temp/E12EAE15-9518-E10E-7C99EA1A2325D105/White_Paper_12_-_LBNL_Peer_Review05.09.08.pdf

⁶¹ See <http://www.hefce.ac.uk/Finance/shared/feasibility/show.asp?id=12&cat=1>

can have significant benefits for the research that is enabled.⁶² These types of system are generally required to be resident within the physical infrastructure of the hosting institution, not least because of the need to employ specialist staff to operate and monitor the equipment, who would have to be employed by the institution.

4.2 Specific HE/FE high energy consumption applications

Within the HE/FE sector, standard business systems and teaching equipment are but one component within the typical university data centre set-up, with a significant number of other specialised systems that operate within this environment. These include equipment control systems and high performance computing (HPC). Each of these has its own operational challenges on top of those that are faced within a typical data centre. For example, they will be used in bursts, or be separated from the standard infrastructure for security and operational reasons, or be so fault-tolerant that turning them off is almost impossible. Within the operation of a research department they may also be core to the work of every member of staff.

4.2.1 High performance computing (HPC)

HPC is the application of powerful computational resources to a range of (generally) scientific and engineering problems. The systems used for these are called supercomputers, though this term may in fact be used to generalise a significant number of different types of systems. Early supercomputers were extremely large and complex bespoke systems, but recently this has changed. This is largely due to the improved performance of commodity systems, particularly Linux-based Beowulf systems. These generally use standard, off-the-shelf servers with a large amount of storage, which are connected by highly specialised networks.

Until relatively recently HPC has tended to push performance "at all costs" (Freeh *et al.*, 2007) but managers of HPC systems are beginning to realise the need for energy-efficient computing facilities. However, HPC environments have a number of unique features that offer energy efficiency challenges over and above those of a standard computing environment (Simon, 2008a, 2008b). For example, they are designed to operate on a single task and make it happen as quickly as possible. They are therefore, for example, less able to use lower clock-speed processors and are normally scaled for the peak load they are likely to expect from their user community. This inevitably means that their energy consumption is inefficient because they very rarely see this volume of work.

The requirement for speed is fundamental to the HPC operation, and this means that having the latest version of a processor and architecture is essential. To this end the individual computational nodes are more frequently replaced or relegated. This should be considered a problem because the embodied energy within the systems, due to their specialised nature, are extremely unlikely to be able to be reused within another sector.⁶³

The other major problem with these systems is their considerable cost which makes, in some instances, upgrade and replacement simply not possible. This means that the desire to have as up-to-date a system as possible is inevitably balanced by the need to have a system of some sort.

⁶² For example seismic modelling where close connectivity to the actual physical experiments is required.

⁶³ Despite making use of some commodity components these systems are essentially 'built-for-task' and difficult to reuse.

4.2.1.1 Short-term fixes for HPC problems

Many of the short-term fixes that could be brought to bear on HPC systems are similar to those for data centres in general. Mills (2008) provides a detailed round-up of some of these more general solutions, in the context of HPC, which include reviewing cooling systems and UPS power supplies. However, in addition, there are a number of specific solutions that can be considered, either to minimise energy consumption or to optimise usage:

- Managers must ensure that they operate their scheduling software to fill the system rather than according to priorities for specific users (irrespective of that user's work patterns). This will ensure that all nodes of the system are kept as fully loaded as possible.
- Each centre should employ software experts who are able to interact with users in order to show them how they can squeeze the optimal amount of performance out of the systems. An example of this is the national supercomputing facility, HECToR, which provides a Computational Science and Engineering service.
- Users should be educated so that they understand the importance of achieving the optimal number of users at all times and that appropriate types of work are being done. The significant cost and overhead for these types of systems means that they should be reserved only for the type of work for which they were designed. This means that they should be used for tasks that rely on their specialised hardware such as closely coupled parallel tasks. Other systems, such as high throughput computing, should be used where they are more appropriate.

In addition, procurement processes need to take account of the type of operation that the system is being procured for. For example, a generalised type of system may be able to cater for as many different types of user as possible but at the same time it will be unable to fully satisfy any single user's needs completely and will therefore be a compromise. To this end it is essential that a full user profile is carried out during the procurement, which may result in the system purchased having a lower headline figure performance but will mean that it will be more tailored to what the user community actually requires. To this end it is also worth considering a suite of different types of small machines rather than one larger system that will operate individually for a specific user community at very high utilisation. In summary, it is very important that everyone involved in the procurement of such systems has a clear view of what the hardware is intended for.

Specific hardware measures that can be taken include:

- When purchasing the cluster ensure that your worker nodes include no extra peripherals such as CD/DVD ROMs. These peripherals are unnecessary for modern HPC systems as they are able to install applications via the network.
- Other components within the cluster should be optimised for energy efficiency, for example by ensuring that a correctly specified network and appropriate power supply are included.
- When specifying and building an HPC cluster, energy utilisation should be taken into account for *all* of the services being procured. For example, questions should be asked as to whether UPS needs to be installed on all the various nodes within the system, as not every node needs it.
- Within a computationally intensive department it is possible to introduce an HPC cluster using desktop systems rather than dedicated servers. Installing a second network alongside the normal department network enables this. This type of system is limited by the specifications of the desktop systems used (for example the number of cores and memory) as well as performance of both networks.
- It is possible to build an HPC system using diskless technology within the worker nodes. This has several advantages from the point of view of energy utilisation. Firstly of course you are not running another peripheral within the system, which will consume energy, but this will

also allow for the system to be shut down and started up much more frequently because this is one of the last remaining moving components within a computer (other than fans) which can be built to withstand this type of operation. This will allow unused nodes within the system to be switched off much more quickly and efficiently. If local storage is still required, then solid-state drives can be used.

Where very large systems are being bought, which may be out of the reach of most HEIs, then serious consideration must be paid to the 'Green 500' list.⁶⁴ This list provides a ranking of the most efficient supercomputers in the world, based on performance per watt of consumed energy (but note that it doesn't take into account energy consumption by services).

4.2.1.2 Near-term developments

Within the HPC community, systems are being developed that move away from the traditional CPU-only system. These developments make use of co-processors that assist with a specific, computationally intense task and sit alongside more general-purpose processors from Intel and AMD. Unlike the multi-core set-up discussed later, these co-processors are physically separate from the central processor (i.e. not on the same IC die). Key amongst these are the development of systems that are able to make use of the massive computational power that Graphics Processing Units and other hardware systems developed for the computer gaming and graphics market provide. These offer significant improvements in the computational power per watt, though not without a significant increase in the complexity of the software that must be developed to make optimal use of the systems. In a similar manner, companies like ClearSpeed or nVidia offer 'accelerators', additional processors that handle floating point operations involved in science computation.⁶⁵ Essentially these co-processors take some of the burden from the CPU and allow it to work at a lower rate, thereby reducing the amount of energy it requires.

4.2.2 Equipment operating systems

There are a significant number of machines within an HEI that are used to control pieces of experimental equipment that must run continually. These systems are normally purchased with the equipment and because of this, users have a limited say in what is included. Inefficiency is caused by the need for these systems to operate 24 hours a day, although institutions should always check carefully whether this is really the case. An example of such a system is shown in Figure 2.

⁶⁴ See <http://www.green500.org/>

⁶⁵ See ClearSpeed for a list of articles and white papers: <http://www.clearspeed.com/newsevents/presskit/#WhitePapers>

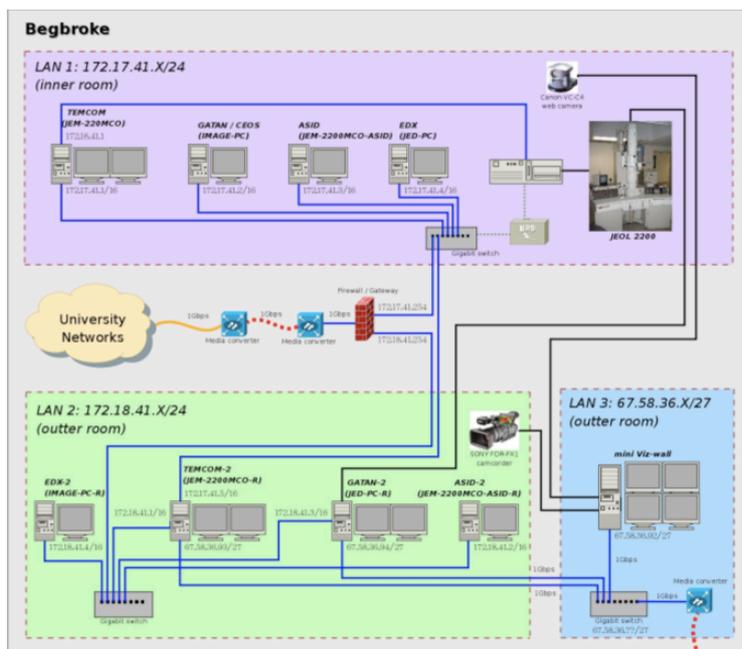


Figure 2 Example of duplicate system in experiment control

In this configuration there are two complete duplicate systems, one of which is directly connected to the microscope and the other is that used by the operator and experimenter. This isolates the equipment subnet from the users which, while beneficial for operational reasons, means that there are many physical systems each of which is fulfilling different roles but individually may be very lightly loaded.

4.2.2.1 Short-term fixes for equipment operating systems problems

The most obvious short-term fix is to ensure that any experimental equipment that comes with its own machine to control and manage the system should always be specified to be as small as is reasonably possible. Using shared file systems to receive the data as it is produced often allows a minimal specification to be used. A more complicated way to improve energy efficiency is to consolidate onto a single system using virtualisation. However, this requires confirmation from the equipment manufacturers that existing equipment support agreements will be upheld.

4.2.2.2 Near-term developments

A number of manufacturers of experimental equipment are now building their control systems such that they may be connected to the departmental network, thereby streaming experimental data directly to the users. This arrangement removes the need for a separate machine, therefore dramatically reducing the energy consumption from these individual units. It is essential that where such systems are offered by manufacturers, both the researcher who is specifying a solution and the procurement office are aware of them. They should be the preferred system unless there are specific functional requirements that cannot be met.

For those systems that cannot have inbuilt data streaming then physical positioning and planning of lab space to allow connection of multiple pieces of equipment to a single virtual host must be considered. This should be at the point of laboratory fit-out rather than a later addition.

4.3 Storage

The capacity requirements for long-term data storage have rocketed over the last few years and where it was once common to discuss storage in terms of gigabytes, it is now becoming common to discuss terabytes (TB). The drivers for this growth vary. In part there has been a tremendous upswing in the amount of data held by students/staff through the uptake in multimedia, video and music. There are also strong demands on institutions with the hosting of websites, staff blogs, social network facilities, etc. and the needs of disaster recovery planning are also driving storage increases. At an organisational level there are obviously considerable long-term records management and archiving issues, but there are also particular concerns for universities when it comes to the increasingly large experiment-related data sets that are being generated and made use of in research situations.

Coping with the demands of this growth as well as managing the requirement for energy reduction will be a major challenge for IS managers over the coming years.

4.3.1 The state of play: problems with storage

Storing information involves a process that has an environmental footprint, particularly the use of power to control and run the media in question (e.g. a hard disk system requires power in order to spin the magnetic disks). Otoo *et al.* (2009) cite research that 25–35% of the energy consumption at data centres is attributed to disk storage systems. This will continue to rise as faster and higher capacity disks and data-intensive applications demand reliable online access to data resources. As the capacity of media increases, the power required to run them also increases. Western Digital, for example, states that a typical 1 TB hard disk requires 13.5 W of power⁶⁶ and by comparison a 320 GB disk consumes around 8 W.⁶⁷ This is an active area, exemplified by ongoing work to devise new benchmarks for measuring the energy efficiency of database servers and storage subsystems. The Transaction Processing Performance Council, for example, has formed a working group to look into adding energy efficiency metrics to all its benchmarks (Poess *et al.*, 2008).

An obvious solution to the exponential growth in storage requirements is to be much more selective about what is stored and for how long. After that it gets more complex, and we have to look at the life cycle of each piece of information that could be stored – when it was created, how long we can wait to access it as it passes through different stages of its history, and when it can finally be archived or deleted. Such considerations may be supported by rules-based technology that assists efficiency by allowing multiple technologies, providing different service levels, to provide facilities for the different life cycle stages, without users of the information being aware of any change. Once information life cycle management is in place, then we can choose to place some of it on very fast response 'online' disks or, as it becomes less relevant to current time-bound transactions, to place it on slower media that use less energy.

However, managing the life cycle of information is, at least initially, a matter of policy, and the extent to which IT managers or technical staff will have the autonomy to implement an information life cycle policy will vary between institutions. Indeed, the strategic significance of information management is rising and JISC has recently commissioned a study into ways in which information management techniques could help in this respect.⁶⁸ There is also increasing interest in techniques

⁶⁶ See <http://www.wdc.com/en/products/GreenPower/index.asp>

⁶⁷ See Tom's Hardware comparison chart: <http://www.tomshardware.com/reviews/hardware-components,1685-14.html>

⁶⁸ See the work of the "Assessing the Role of Digital Information Management Practice in Reducing Environmental Impact of ICT Use in Higher Education" study, available at <http://www.jisc.ac.uk/whatwedo/programmes/infl1/greeningim.aspx>

such as EA,⁶⁹ which means that such policies are more and more likely to become part of corporate governance.

With this in mind, we focus this section on technology that reduces the energy demands of storage either by operating traditional storage systems with less power or through new techniques that allow sharing of storage capacities and thus offer economies of scale.

4.3.2 Short-term fixes storage problems

In general, data and storage management solutions tend to be implemented as piecemeal, partial solutions to an increasingly complex problem. There may be compatibility problems, for example between solutions from different vendors. In fact, there is little in the way of a true low-cost 'quick fix' for storage problems other than consolidating individual disks. While there are solutions based on mature technologies these tend to be expensive and would likely be implemented as part of an information life cycle policy. In this respect the main solutions focus on consolidation and virtualisation.

4.3.2.1 Consolidate individual disks

A relatively easy fix that can be widely adopted by individual staff and students is to review their current set-up with regard to hard disks on user machines. Many machines that have been around for a few years have had additional hard disk drive (HDD) storage added to increase storage space. As an example of what is being proposed here, Linux Magazine recently reported that replacing three older HDDs of low capacity with a state-of-the-art 500 GB hard disk could save 20 W (Casad, 2008).

4.3.2.2 Consolidation through SANs and NASs

As well as managing the information life cycle, it is important to look at where and how data is stored. Traditionally, storage has meant hard disks that are directly connected to a server or desktop PC (direct attached storage, or DAS). Many data centres still use DAS and, in addition, many institutions organise storage around departmental function (for example HR and Engineering have separately managed storage systems), thus multiplying the number of storage devices across the institution.

Newer ways of handling storage detach it from an individual PC or server and make it more of a function of the network. These techniques include storage area networks (SANs) and network attached storage (NAS), which both provide more efficient use of storage space by sharing capacity across the network. OnStor, a leading supplier of NAS, claims up to 90% energy saving through NAS.⁷⁰

Readers who are interested in further details of SANs, particularly within HE, are referred to JISC TechWatch's 2003 report on the subject (Chidlow, 2003).

4.3.2.3 Virtualising storage

Beyond consolidation, storage can be virtualised: instead of allocating a disk-sized segment for storage on a piece of physical media, a 'virtual disk' can be of any size and will be split into small

⁶⁹ Recent JISC reports include: Doing Enterprise Architecture: Enabling the agile institution (April 2009), and Unleashing EA: Institutional Architectures and the value of joined up thinking (July 2009), both available from <http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/earlyadopters.aspx>

⁷⁰ See http://www.onstor.com/bhive/t/5/go_green.jsp

chunks across a large number of physical disks (and perhaps more than one array on the SAN). The environmental benefits of this approach are realised through extensions to the basic technique such as data migration, 'lean provisioning'⁷¹ and 'thin provisioning'⁷² (Roussos, 2007; Graefe, 2008).

4.3.2.4 De-duplication

Much of the information stored by an organisation is in fact duplicate data. To cope with this, organisations are increasingly using data de-duplication technology to eliminate all but one copy. This is a technology with commercial solutions today often implemented at the SAN level. Some argue that as much as 20 times more data can be stored in a given storage capacity using this technology (Geer, 2008).

4.3.3 Near-term developments

It is highly likely that over the next few years research will intensify into ways to reduce the energy use of storage systems. Such work will be driven partly by the demands for ever-increasing storage capacity driven by peta-scale computing (e.g. data-intensive science experiments) and partly by changing legislative demands.

4.3.3.1 RAID

Redundant Arrays of Inexpensive (or Independent) Disks (RAID) is a technique that combines several hard disks into a single, secure storage mechanism. RAID arrangements are common in data centres and can consume as much as a third of the total energy. Until relatively recently developments in RAID technologies had focused on improvements in performance and reliability, but researchers have recently shown increasing interest in energy reduction. A variety of schemes and techniques have been proposed including Green RAID (Mao *et al.*, 2008), Diverted Access (Pinheiro *et al.*, 2006) and EERAID (Li and Wang, 2004). It is likely that such ideas will be incorporated into RAID technology in the coming years.

4.3.3.2 Archiving and MAIDs

Traditionally, tapes have been used for long-term storage and archiving – a process that decouples the storage medium from the access hardware (Storer *et al.*, 2008). There are a number of technical disadvantages to this way of working, including energy efficiency. As other forms of storage, such as hard disks, become capable of increased storage capacity, new possibilities will open up.

One energy-efficient approach to digital storage is Massive Arrays of Idle Disks (MAID). The technique was developed by researchers at the University of Colorado and is being commercialised by a company called Coplan (a number of other companies have followed suit, including Hitachi and EMC). MAID demonstrates "considerable energy-based cost savings" (Storer *et al.*, 2008, p. 2) mainly because these systems leave the majority of their hard disk platters 'spun down'. In essence it replicates the idea of only fetching, powering up and accessing the information on a particular tape when it is needed, but does this using disks. Various policies exist for how to distribute the data among the disks in order to minimise the number that are powered up (Colarelli and Grunwald, 2002) with Harnik *et al.* (2009) providing an overview of some of these methods. This of course delays the speed at which access to the data is obtained, but for archiving scenarios this is not seen as a major problem. Coplan staff claim a cost saving of 75–90% per energy unit of storage (Smith, 2008).

⁷¹ See <http://www.networkworld.com/community/node/31166>

⁷² See <http://www.information-age.com/channels/storage/features/313741/thin-provisioning.thtml>

4.3.3.3 NAND Flash

Another technology that is becoming mainstream is the NAND Flash memory and in particular its use in solid state disk drives (SSDs). This technology was initially used in small portable devices such as MP3 players, but increasing capacity has opened up the possibility of introducing it into more mainstream uses. Binstock (2008) notes that these kinds of drives are beginning to appear in the workplace, especially in situations where speed of access is the dominant criteria. Further information on this can be found in Leventhal (2008) who argues that SSDs could form a new tier in the storage hierarchy. Others are looking into incorporating a small flash memory into a traditional disk drive (a hybrid disk) to explore the impact of the flash memory on the performance and energy consumption of the disk (Deng *et al.*, 2008; Wu *et al.*, 2008).

Some have claimed that, because SSDs involve no moving parts (as opposed to a spinning hard disk) they represent a good, long-term prospect for energy-efficient computing. Leventhal (2008) quotes figures of 12.6 W for a 750 GB HD, 5 W for a 1 G DRAM memory module and only 2 W for a 128 GB SSD (although warns of the difficulty of being completely accurate). Binstock (2008) argues that this is debatable, saying that: "On a watt/GB basis, HDDs outperform SSDs by a larger margin. In other words, SSDs are not greener from a power-savings point of view than HDDs" (p. 1). It is likely this debate will continue as the technology improves. For example, Samsung announced in January 2009 a 100 GB SSD with a claimed energy use of less than 2 W in operating mode compared with more than 8 W for an HDD. One of the problems, however, is the comparison with a moving target as manufacturers of HDD continually focus their efforts on improving the energy efficiency of their technology.⁷³ Another, more technical factor relates to SSDs and what is known as the flash translation layers that are tuned to different tasks, with some not suited to the task they are sold for, resulting in power consumption that turns out to be larger than a standard disk drive (Choudhuri and Givargis, 2007). Nevertheless, they are a technology that must be tracked over the next 2 years; they are initially being introduced as an acceleration technology, but have the potential to completely replace hard disks over the next 5 years. It is worth noting that a number of blade server systems are now being shipped with such storage.

4.3.3.4 Compression

Another area of research likely to deliver further developments is that of data compression. By using clever mathematics to develop ever more efficient compression algorithms, storage growth can at least be kept to a lower level. Obviously this is not without some computational cost and there are trade-offs between system performance and the use of compression, and/or its sister technique de-duplication.⁷⁴ It also worth noting that Kothiyal *et al.* (2009) investigated the potential for data file compression to save energy in GNU/Linux server systems and found they could make no generalised conclusions one way or the other.

4.3.3.5 Dynamic power management of disks

Recently there has been academic interest in the idea of dynamically managing the power of disk storage facilities (Otoo *et al.*, 2009) in a similar manner to power management used in end-user devices such as a PC (see Section 4.5). Much of the work to date has considered a single disk only and attempts to find an optimal idle waiting period after which a disk should be moved to an operating state which consumes less power.

⁷³ For example Western Digital's GreenPower series, which promises to reduce power consumption by over a third and uses 13.5 W to spin a 1 TB disc drive.

⁷⁴ Some of these trade-offs are discussed at <http://searchstorage.techtarget.com.au/articles/29948-Data-compression-vs-data-deduplication-technologies-for-SMBs>

4.4 Networking

Historically, very little has been done to optimise networking efficiency (Ethernet, local area networks [LANs], routers, Internet backbone, etc.), apart from individual manufacturer action to increase the throughput of data per watt of their devices. However, as the Internet grows there is a steady proliferation of more sophisticated networking equipment needed to cope with the increased volume of data throughput, with a subsequent increase in direct energy consumption and, indirectly, electricity to provide cooling. In addition, as desktop computing facilities, servers and storage are progressively optimised, the spotlight has inevitably fallen on network inefficiencies.

While there are no HE-specific figures for the UK, work undertaken for JISC (James and Hopkinson, 2009a) estimated that of an institution's total ICT-related electricity usage, the proportion used by network-related equipment was in the range of 8–15%, depending on the kind of institution. These HE figures contrast with Gartner (2007), which estimates that networking of various different types forms approximately 31% of the ICT-related energy footprint, and with the Climate Group (2008) which estimates 37% of 2007 levels. The discrepancy may be due to differing ways of determining what is classed as networking, but either way this is a substantial and, more importantly, rapidly growing area.

4.4.1 State of play: problems with networking

Across a wide range of communications scenarios (LANs, metro area networks, the Internet backbone, etc.) there is a steady move to the adoption of Ethernet even though it is particularly energy inefficient. Various devices now receive their electric power via the network, for example Voice over IP (VoIP) telephones, which are often left permanently powered-up, even when office blocks are empty at weekends or holidays. In fact there are many reasons for Ethernet's increasing ubiquity, and interested readers are referred to a recent TechWatch report: *100G Ethernet and beyond* (Rizvi, 2009).

Ethernet networks have three main sources of inefficiency: they are generally designed for the busiest load scenario (which by definition exceeds their average utilisation), and the links and topologies within these networks are rarely optimised, which means that capacity is often poorly managed.⁷⁵ Thirdly, because Ethernet is a carrier sense protocol it maintains a carrier signal at full power regardless of the distance to the recipient or the amount of data being transmitted: even when a network link is being used at a fraction of its capacity it uses nearly the same energy as when fully utilised (Gupta and Singh, 2003). Nedeveschi *et al.* (2008) argue that "The implication of these factors is that most of the energy consumed in networks is wasted" (p. 1).

Related to this is the difficulty involved in measuring energy efficiency in networking scenarios. The complex mix of technologies, applications and performance variables means that there is little by way of unambiguous energy efficiency metrics standards (Ceuppens *et al.*, 2008). This is backed up by the Smart2020 report (Climate Group, 2008), which noted the complex environments which pertain in Internet backbones, fixed line and mobile telecommunications infrastructure and the difficulty for operators in obtaining a clear view of energy usage. Many of the issues discussed are likely to also pertain to LANs. It is generally accepted that such fixed network operators have not considered energy consumption a major factor in their cost bases. While this is starting to change (Berl *et al.*, 2009), most attention has focused on meeting the technical requirements of developing low-energy wireless networks such as ZigBee because of operational (e.g. use of batteries with laptop-based WiFi) and regulatory requirements.

⁷⁵ Nedeveschi *et al.* (2008) quote backbone utilisation of under 30% as typical.

The rest of this section provides a summary of short-term solutions (which either require reorganisation or investment in relatively mature technologies rather than representing technology challenges) and mid-term, experimental work that should be tracked as a source of potential future solutions.

4.4.2 Short-term fixes for networking problems

Just as we have seen with data centres, a process of equipment consolidation and replacement can be beneficial for improving the overall energy efficiency of an institution's network. As networks are over-specified, to allow for all data bandwidth eventualities, Chabarek *et al.* (2008) argue that looking at network topology and in particular how power demands of switchers and routers work under different loads and configurations can generate benefits in energy use.

Immediate solutions to the energy usage of network equipment centre on reviewing current arrangements and undertaking a process of consolidation, perhaps with the increased use of virtual LANs.

4.4.2.1 Equipment consolidation

New equipment can offer advantages with regard to energy efficiency and it is noticeable that in recent months network manufacturers have begun to introduce new 'green' ranges. However, a particular issue for IT managers who wish to consolidate their networks and refit with new equipment is that the provision of information on the energy efficiency of networking equipment is somewhat behind other sectors of the ICT industry. Two recent developments are of interest here: the Energy Consumption Rating and the EU Code of Conduct on Energy Consumption of Broadband Communication Equipment (JRC Institute for Energy, 2008b).

The **Energy Consumption Rating (ECR)** is an open initiative that is developing a method to measure the energy efficiency of networking equipment. It was set up in October 2008 by two leading US network product vendors, Juniper Networks and IXIA, in association with Berkeley Lab.⁷⁶ Their proposal is that a normalised metric be developed and applied to different classes of equipment. For example, the metric would calculate energy efficiency for a range of equipment from different vendors falling into the class of edge routers. A comparison could then be made across the specific class. The metric involves a formula that includes the energy consumption in watts and effective throughput (in bits per second). Equipment with a lower ECR will expend less energy to move the same amount of data payload around the network.

The specification is currently in draft form (v1.04),⁷⁷ and although the group has developed a metric and a methodology for testing equipment, the initiative is seen, at this stage at least, as more a call to action for other vendors to get involved, but to date there has been little sign of this. Others, such as Cisco, have indicated to the press that they would rather work with existing standards bodies such as the ITU.⁷⁸

The **EU Code of Conduct on Energy Consumption of Broadband Communication Equipment** provides guidance and target power levels for various classes of equipment (e.g. ADSL modems, WLAN access points, etc.) (Bertoldi, 2008).

⁷⁶ See <http://www.networkworld.com/news/2008/102808-juniper-ixia-ecr-initiative.html>

⁷⁷ See <http://www.ecrinitiative.org/>

⁷⁸ See http://www.pcworld.com/businesscenter/article/153029/efficiency_drive_moves_to_networks.html

4.4.2.2 Virtualisation

Related to physical consolidation is the virtualisation of networks – going from physical local area networks (LANs) to virtual ones (VLANs). In the virtual model, the organisation may have multiple 'separate' networks (e.g. human resources, finance, faculty/department) but these are not physically separated – instead, they share network switches and inter-switch links, and the switches are configured by software to separate the traffic that runs on them. This model is beneficial in terms of agility and maintainability, but a key benefit is also that it uses less energy than multiple parallel, physical networks. Most equipment manufacturers now offer VLAN support, at least at enterprise or service provider level.⁷⁹

4.4.2.3 Energy-efficient equipment

In recent months equipment manufacturers have started to introduce more energy-efficient network, wifi and router equipment. This has been particularly noticeable at the lower, consumer end of the network chain, for example, individual wifi boxes.⁸⁰ It is likely that across a campus some of these devices are being used in individual departments, halls of residence and research areas and that replacement of equipment may be of benefit.

4.4.3 Near-term developments

There are a number of developments that are likely to come to fruition over the next 2–3 years, designed to ameliorate the main sources of networking inefficiency.

4.4.3.1 Speed rationalisation

Designing for the busiest load is a mindset that has dominated network design in order to cope with worst-case scenarios. However, some people are starting to look at the benefits of speed rationalisation across an organisation's total network. For example nowadays, equipment is often fitted out with Gigabit Ethernet as standard. However, the power difference between a medium-loaded 100 Mb/s link (100BaseT) and the same load on a Gigabit (1000BaseT) link is perhaps 3 or 4 W – not so much per link, but very substantial when extrapolated across an enterprise. Gunaratne *et al.* (2005) give a figure of 3 W between an Ethernet line card running at 10 Mb/s and one at 1 Gb/s (when no traffic). Cooper (2007) provides similar figures of around 3.2 W.⁸¹ Gunaratne *et al.* (2005) of University South Florida argue that their experiments and results show that "even for the busiest user on the USF campus it is possible to operate his or her link at a lower data rate (in this case 10 MB/s instead of the maximum possible 100 MB/s) with no significant increase in delays" (p. 306). Their work used a form of Ethernet known as Auto-Negotiation and required adaptation of the network card, but the authors argue that the costs of this would be offset by the energy saved. Such work is closely related to Adaptive Link Rate.⁸²

4.4.3.2 Port and switch shutdown

In the same way as desktop PCs and servers should be shut down after hours, data network managers should consider shutting down segments of corporate or institutional networks. After office users have gone home, data ports and VoIP phone ports should be powered down, and a number of manufacturers have started providing this functionality. This will save energy in any

⁷⁹ Held (2003) provides an introduction and Decisys have produced *The Virtual LAN Technology Report*, a detailed technical introduction, available at <http://www.3com.com/nsc/200374.html>

⁸⁰ Examples include Netgear Wireless-N Router (WNR2000) and D-Link Xtreme Gigabit routers.

⁸¹ See <http://www.slideshare.net/us056444/saving-energy-with-smart-cabling>

⁸² See <http://efficientnetworks.lbl.gov/enet-adaptive.html>

event, but is particularly relevant when the ports in question supply Power over Ethernet to phones or other devices such as wireless access point aerials. In this case, the end devices will also be powered down.⁸³ Ports in multi-port LAN switches should also be shut down when they have no active equipment connected and a number of manufacturers have started to produce switches that can do this automatically.

Manufacturers have also begun to incorporate power management software into their systems so that equipment can be switched off remotely. This is a rapidly changing area of development and a number of announcements have been made recently. Cisco has introduced Energy Wise,⁸⁴ and Go Green⁸⁵ has been introduced as the result of a partnership between Extreme Networks (data switch manufacturer) and Avaya (VoIP equipment manufacture). Unfortunately, neither offering is generic to other types of telecoms and data equipment, so a situation like desktop PC power management (where central software will power off devices from many different manufacturers) is yet to come. However, Cisco has announced that it is working to develop a Software Development Kit to allow independent software to access the same power control facilities.

4.4.3.3 Energy-efficient equipment

In the longer term, networking equipment manufacturers will need to invest more resources into developing equipment that tackles power management, as manufacturers of end user devices such as PCs have done. Nedeveschi *et al.* (2008) argue that this needs to happen at two levels. Firstly, network equipment ranging from routers to switches and network interface cards (NICs) will need power management primitives added at the basic, hardware level. Gupta and Singh (2003) describe this as "putting to sleep some of the sub-components" (p. 20) of the router equipment – for example line cards, processors, memory – when not in actual use.⁸⁶

This 'sleeping' will be enabled by software which must interact with the network's protocols, and this leads us to the second point: network protocols will need to make use of these hardware primitives to best effect. This relates to the previous subsection (port and switch shutdown), but also relates to the design of routing protocols and the network topology, for example by allowing coordination to facilitate changes of routes during low periods so as to aggregate traffic along a few routes only, while allowing others to sleep. There are a number of detailed technical issues with regard to changing protocols and these are discussed in more detail in Gupta and Singh (2003).

The Energy Efficient Internet project based at the University of Florida is one example of the research work being undertaken to improve this situation.⁸⁷ It focuses on problems at 'the edge' of the network where office and consumer PCs and other networked devices are often idle but need to remain powered up to handle potential network transactions. Other projects, such as the Energy Efficient Digital Networks⁸⁸ at Berkeley Lab, are looking at the deeper core of the Internet and the energy efficiency of router and backbone equipment. The pace of these developments has picked up recently, and a new international conference, the Workshop on Green Communications, or GreenComm, took place in June 2009.⁸⁹ A second workshop will take place in late 2009. Protocols will be redeveloped over the next few years and new equipment introduced in order to rectify the

⁸³ Although it is well to be aware that some institutions use the network at night to update machines and back up user data while the network is quiet. The use of Power On LAN also needs to be considered.

⁸⁴ See <http://www.eetimes.com/news/design/showArticle.jhtml?articleID=212902697>

⁸⁵ See <http://www.avaya.com/rt/default.aspx?CurrentPath=emea/en-us/corporate/pressroom/pressreleases/2007/pr-121107b.htm>

⁸⁶ An alternative for saving energy is to clock the hardware components at a slower speed.

⁸⁷ See <http://www.csee.usf.edu/~christen/energy/main.html>

⁸⁸ See <http://efficientnetworks.lbl.gov/enet.html>

⁸⁹ See <http://www.greencomm09.org/home.html>

energy profligacy of network equipment. Of particular interest is the work currently being undertaken into energy efficient Ethernet.

4.4.3.4 Proxy support for sleep modes and split TCP

The Internet (and therefore Ethernet) was designed with the idea that devices at the edges are either fully on or off, and large amounts of energy are consumed by devices maintaining a fully on network connection: there is no 'sleep' mode. It is estimated by the Ethernet Alliance that "billions of dollars" worth of electricity is used to keep Ethernet (and connected devices) fully on at all times and that a solution could save perhaps half the energy used by desktop PCs (Nordman and Christensen, 2007, 2009). They go on to propose two potential solutions: firstly, to redesign Internet architecture and protocols to take account of power states; secondly, to encapsulate the intelligence required to maintain the network presence in another 'separate' entity.

According to Nordman and Christensen (2009), the second solution makes use of a "network connectivity proxy" (NCP), which keeps the sleeping PC's network connection 'alive' and wakes it only if required. In effect the NCP 'covers' for a sleeping PC and will be designed to use far less power. The proxy may be an external device (a wiring cupboard first-level router) or a specially adapted NIC, known as a smart-NIC. Provision for such a mechanism has been included in the recent Energy Star 5.0 requirements and an ECMA standards group (TC32-TG21: Proxying Support for Sleep Modes⁹⁰) is working on the technical details of standardisation that will encourage the development of hardware to support Energy Star 5.0 requirements. Nordman and Christensen (2009) expect PCs with the necessary hardware and software to be available in 2010 and argue that for an average PC the technique will lead to savings in the order of 400 kWh per annum. They recommend that in the meantime network managers identify any usage models or applications that proxying cannot support so the users of these can be separated from the rest who can.

A related problem is that of persistent network TCP connections in client-server scenarios. Many applications maintain a permanent TCP protocol connection between a client and server so that, if there is no other traffic, both sides can generate and respond to 'keep alive' messages at least every 2 hours.⁹¹ The problem is how to respond to these messages without requiring the full resources (and hence energy) of the client or server. One proposed solution is for split TCP connections (Gunaratne *et al.*, 2005). The TCP connection is split with the addition of a shim layer between the socket's interface and the application. This shim layer fakes a persistent connection to any applications that require this even when a client powers down. Related work is being undertaken by another group who are looking at what they call Somniloquy (Agarwal *et al.*, 2009).

4.4.3.5 Energy efficient Ethernet

If Ethernet continues its evolution as the primary mechanism for transmitting multimedia, the power demands for high-speed circuits will need to be reduced. With this in mind, manufacturers and researchers have begun work on a new standard for energy efficient Ethernet (EEE), which will be known as IEEE 802.3az. The IEEE development work is defined by a number of criteria (Energy Efficient Ethernet Task Force, 2009):

- distinct identity (i.e. branded differently to standard Ethernet)
- market potential and commercial feasibility
- technical feasibility

⁹⁰ See <http://www.ecma-international.org/memento/TC32-TG21.htm>

⁹¹ Strictly speaking the TCP standard does not mandate any particular time for this, but TCP software often includes an 'unofficial' keep alive message every so often. See http://www.tcpiipguide.com/free/t_TCPConnectionManagementandProblemHandlingtheConnec-3.htm

- backwards compatibility

Work currently consists of a number of streams:

- Low Power Idle (LPI): roughly equivalent to power management on a laptop, LPI enables the circuit to go into a low power 'doze' mode when there is no traffic being sent. It also provides a method to optionally negotiate the wake or resume-time from LPI, allowing systems to safely enter deeper sleep states to save additional power.⁹²
- Variable physical transmission speeds (at the PHY level of Ethernet): switching the physical media transmission speed down when there is less traffic but turning it up when full speed is required (Adaptive Rate).
- Turning off channels: above 1 Gb/s, Ethernet transmits through multiplexed channels and the power requirement gets larger as the speed increases. This can result in a 24 W power requirement for each end of a 10 Gb/s link. However, it may be possible to turn off some channels at times.
- Investigating new signalling mechanisms.
- Reducing transmit amplitude: this is akin to turning the volume down. A specific focus here is providing a backward-compatible, lower power 10 Mb/s link standard.

Backwards compatibility is a very important aspect of this standards work and is one of the reasons for Ethernet's increasing popularity. For example, a specific focus of the IEEE work is the creation of a lower power and fully backward-compatible 10 Gb/s standard.

The IEEE 802.3az Working Group has been active since 2007, and plans to deliver a published standard by September 2010. The main participants in the standardisation process are large semiconductor companies such as Broadcom, but there is a part to be played in this process by universities. The potential to reduce the in-use carbon footprint of ICT is large – this could halve the power requirement for wide and local area networking in future, and should be tracked.

⁹² See <http://www.ethernetalliance.org/press-room/press-releases/185-ethernet-alliancer-announces-updates-from-ieee-8023-task-force-interim-meetings-.html>

4.5 End user devices

We define end user devices as those systems and devices that the end user interacts with in order to undertake information processing and communication activities. This category includes desktop PCs, monitors, laptops, netbooks, e-readers, PDAs and mobile phones. There are obviously many millions of such devices in operation already and therefore the combined energy consumption and related environmental footprint are significant. Matthews and Matthews (2003) quote figures of over 130 million PC computers being sold per year around the world and over a billion mobile phones were sold in 2007.⁹³ The UK public sector alone spends £1 billion a year on new PCs (Doyle, 2008). More importantly, the rate of annual growth is expected to rise dramatically, not least through emerging economies such as China. It is expected that there will be more than 4 billion PCs and laptops by 2020 and 50% of the world's population will own a mobile phone (Climate Group, 2008).

Around 40% of users in UK HE supplement their desktop with a laptop (James and Hopkinson, 2009c). While laptops often perform better than desktops in terms of energy economy, they are only suitable for short periods of use⁹⁴ so their role is limited at an institutional level. Laptops per se are therefore not considered here as a separate category of end user device, although there is some discussion of laptop technology. Also within this section, while there is some discussion of mobile device technology and its potential contribution to energy efficiency, there is no detailed investigation of different devices. An institution's ICT manager can expect to have little control over individual devices and it is unlikely that this will form a part of energy reduction measures under the CCA/CRC other than in a general sense of needing to reduce energy consumption.

This section will therefore focus on the environmental implications of the standard end-user set-up: the desktop PC and monitor. In offices and teaching rooms, such a set-up will usually be networked and the separate implications of this have been discussed elsewhere. It is worth noting in passing the scale of the challenge when we consider the energy requirements of the stock of PC systems, with James and Hopkinson (2009c) estimating a computer stock of 760,000 PCs for UK HE, accounting for almost half of the sector's estimated ICT-related CO₂ emissions of 275,000 tonnes and an ICT-related electricity bill of £61 million in 2009. Similar figures are quoted for FE.

4.5.1 State of play: problems with end user equipment

The power demands of the average (UK 2007 non-domestic stock modelled) desktop computer, LCD (liquid crystal display) monitor and CRT (cathode ray tube) monitor is shown in Table 1.

	Power demand (W)		
	On	Sleep	Standby (Off)
Desktop	74.16	4.60	2.80
LCD	35.51	1.15	0.96
CRT	61.39	3.68	2.03

Table 1 Power demand for end user desktops and monitors⁹⁵

⁹³ See <http://www.itpro.co.uk/166446/mwc-2008-billion-mobile-phones-sold-each-year>

⁹⁴ For longer periods, the ergonomics of the laptop mean that it is likely to fall foul of the Health and Safety Executive's Display Screen Equipment (DSE) Regulations. In order to comply, the machine would need to be combined with: (1) a separate monitor capable of being tilted on a vertical and horizontal azimuth; (2) a separate keyboard and mouse. A docking station would most often be used to facilitate this.

⁹⁵ The Market Transformation Programme (MTP) provides an online tool at <http://whatif.mtprog.com>, which provides estimates for power demand for various types of devices at set dates. The figures for this table were taken from the

It should be noted that these are average figures for the stock. From such average figures we can quickly calculate the carbon footprint, over the course of a year, of a PC left on at all times (approximately 75 W for the PC + 35 W for the monitor = 110 W). Over a year, 8,760 hours × 110 W = 964 kWh. Where 1 kWh of grid electricity is considered equivalent to 0.537 kg CO₂ (based on the Carbon Trust's Greenhouse Gas Conversion Factors⁹⁶) then the total, annual, CO₂ equivalent is approx. 517 kg).^{97,98}

The difference in the power demand of the two main monitor technologies is apparent: LCD monitors require around half that of their CRT counterparts – which goes part of the way to explaining why the latter are no longer available to buy. However, the issues are not necessarily as straightforward as this. CRT monitors, for example, have much less embodied energy than LCD and this makes the total life cycle impact more difficult to calculate. Also, while the low power draw of all units in their respective sleep/standby/off modes is testimony to the effectiveness of initiatives such as Energy Star, the time that the equipment spends in such low power states is often far less than it could be. For example, the use of screensavers means that neither the monitor nor the desktop is going into a low power state. In fact, the desktop will be consuming more energy than it otherwise would have been because of the slightly higher processing power required to run the screensaver.

There are three key problems with respect to making computers and monitors more energy efficient: equipment usage patterns, embodied energy, and the implications of Wirth's Law.

4.5.1.1 Equipment usage patterns

Equipment usage patterns are a function of: (1) the amount of time a user is making productive use of their equipment, (2) the effectiveness with which the computer and monitor enter their low power states when not in use. When plugged into the mains electricity a desktop computer may be in one of six power states, described in Table 2 (over page).

figures quoted in the tool for the year 2008, except for the CRT figures, which were taken from 2007, the last year available.

⁹⁶ See http://www.carbontrust.co.uk/resource/conversion_factors/default.htm. The figures used by the Carbon Trust are based on the figures for DEFRA's greenhouse gas conversion factors, full details for which are at <http://www.defra.gov.uk/environment/business/reporting/conversion-factors.htm>

⁹⁷ These conversion factors are for the current electricity generation mix of the UK; for other countries they may be higher or, for example in hydro- and nuclear-rich Sweden, far lower.

⁹⁸ To fully appreciate the consequence of this figure it is worth noting that the environmental writer George Monbiot (2006) has calculated that in order for the world to be seriously tackling climate change by 2012 then the per person limit of carbon emissions per annum should be around 800 kg.

ACPI	Description	Approx. wake-up time	Typical PC power demand
S0	working state or on-idle	none	highest (70–110 W)
S1	soft standby – hard drive powered down	2–3 s	high (65 W)
S2	as above with power to CPU cut	3–4 s	high (60 W)
S3	only the RAM receives power	Max. 10 s	low / v. low (4 W)
S4	hibernate – save to disk and power down	30–60 s+ (with proviso ⁹⁹)	very low (3 W)
S5	power down – minimal standby state	60 s minimum – sometimes minutes	very low (3 W)

Table 2 ACPI power states, typical wake-up times and power demand for an average UK non-domestic desktop in 2007 (see footnotes 95 and 99)

As can be seen from Table 2, the difference between a desktop computer being 'on-idle' and it being in standby state S3 (from which it takes no more than 10 s to wake up) is significant. The latter state draws only 5% of the power of the former. Powering the machine down completely to a minimal standby state means only 2.5% of the power is being used. Despite this, several studies (Canon, 2006; Fujitsu, 2005; Nordman *et al.*, 2000; Webber *et al.*, 2006) have shown that in networked computing environments between 25% and 60% of desktop computers are left in their 'on-idle' state when not in use. Indeed, Nordman and Christensen (2009) note the increasing propensity for people to almost never turn off their PC, even leaving it running in on-idle over the weekends. The reasons for this centre around:

- risk aversion among departmental IT staff who feel that it may cause them more work if power management, or the capability for users to configure power management, is enabled;
- lack of incentive for ICT department staff (who are rarely if ever aware of, let alone responsible for, their equipment's associated running costs);
- lack of awareness or choosing not to configure power management settings for energy efficiency where users have the capability.

4.5.1.2 Embodied energy

A key question with respect to low carbon ICT is whether or not to replace existing equipment with newer machines. Although there is a considerable embodied energy/environmental cost in replacing a machine, new equipment may use less power and operate more efficiently, if chosen appropriately. Although this question applies to the whole range of ICT-related equipment it is particularly pertinent with regard to end user devices such as PCs, in part because individual users are more likely to have some say over when they replace their own machines and replacement costs are relatively low compared to other ICT equipment (<£1000), but also because many public sector institutions have established 'refresh' cycles for user machines. In addition, manufacturers are increasingly using their 'green' and energy credentials to stimulate sales and this has now become a key marketing instrument.

However, the answer is not always as straightforward as it might first appear and, indeed, is the source of considerable debate amongst environmental computing experts. An example will help illustrate the dilemma. Consider a comparison between monitors:¹⁰⁰

⁹⁹ With regard to Table 2, the time taken by a computer to recover from states S4 and S5 will depend very much on the processing capabilities and installed RAM on the machine (and also how well it has been maintained (e.g. defragmented). It is not uncommon to see older machines running XP with 0.5 GB RAM take up to – and in some cases over – 10 minutes to boot up. This serves as a major disincentive to getting users to power-down their machines.

- Monitor (1) is an old CRT device and has a total annual in-use electrical energy consumption of 158 kWh (83kg CO_{2e}).
- Monitor (2) is a modern LCD device and has a total annual in-use electrical energy consumption of 84 kWh (44kg CO_{2e}).

Looked at in terms of energy-in-use, the LCD monitor consumes half the energy and emits half the CO_{2e} of the CRT monitor. However, in terms of *embodied* emissions, the CRT device contains 236 kg of CO_{2e}, and the LCD unit, 423 kg of CO_{2e} (Socolof *et al.*, 2005).

The total life cycle emissions of the monitors (in kg of CO_{2e}) will therefore be:

- CRT: 236 + 83y
- LCD: 423 + 44y (where y equals years of use)

While the LCD leads to lower emissions in use, its greater embodied GHG emissions mean that it would take nearly 5 years of use for the LCD to overtake the CRT. Faced with the question: 'Suppose we have a functional, CRT monitor and are considering replacing it with an LCD unit on the grounds of reducing GHG emissions alone – would this make sense?' The answer would clearly be: 'No. Retain the CRT until it is no longer fit for purpose.'

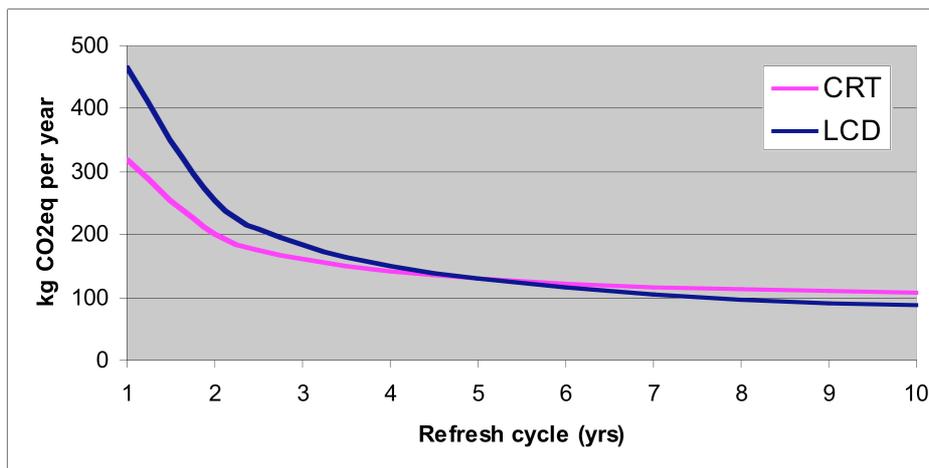


Figure 3 Variation of GHG emissions with refresh rate in CRT and LCD monitors

Figure 3 shows that unless the monitor is to be kept for at least 5 years, the GHG emissions will be higher for an LCD monitor than its CRT equivalent. Clearly, this example is based solely on average usage patterns; similar models could be drawn to reflect exceptional usage patterns or future changes in embodied GHG intensity (which may arise from better recycling facilities, etc.). It should also be borne in mind that there is some debate over the exact embodied energy proportion. Most studies seem to indicate that over the lifetime of equipment, the 'use' proportion of CO₂ is higher than the manufacture and James and Hopkinson (2009c) quote a number of studies in support of this view. However, there is by no means a consensus and in part this seems to depend on what type of PC is used in the model, the kind of use it is put to, whether a number of machines are modelled, etc.

For example, the work of Masanet and Horvath (2006) concluded that manufacturing and use impacts were roughly comparable for a hypothetical model office environment with 5,000 PCs and a refresh rate of 3–4 years. On the other hand, Moshnyaga (2009) quotes figures from 2002

¹⁰⁰ Based on typical consumption figures from the Market Transformation Programme, as before.

indicating PC manufacture uses three times the energy of 3 years' worth of PC use. He then updates the figures for a typical machine in 2007, indicating that there has been a 7% increase in the energy required to produce a machine in the 5 years (due to the use of new technology such as a CPU with far more integrated circuit transistors). The energy used when the machine is 'in-use' has remained about the same. He concludes that: "the total energy associated with modern PC is dominated by the production energy. Reducing computer lifespan leads to heavy loss of unused energy" (p. 91). Based on his research he recommends that the most effective way to save energy is to reuse PCs, perhaps with upgraded components, arguing that "an old PC has to be reused until the production energy for a new PC is offset by the low cumulative energy of usage [and] upgrade of the old PC".

An additional point to make with regard to upgrading is that usually the most important upgrade is to add additional memory and that a careful choice of timing for this can prolong the useful life of a PC by at least a year or two. Indeed, upgrading any 32-bit operating system PC to the maximum addressable memory of 4 GB might be a sensible future-proofing tactic.

4.5.1.3 Wirth's Law and the evolution of energy-hungry processing

The evolution of desktop PC sales has been driven, to a degree, by built-in obsolescence, both in software and hardware. This has been neatly encapsulated in 'Wirth's Law' (Wirth, 1995), which describes the 'Wintel' symbiosis where progressive versions of the Windows operating system require ever more powerful processing capability in order to function efficiently, resulting in functional desktop machines being replaced, before their time, just to keep up.¹⁰¹

A substantial proportion of the energy consumed by computers goes to the CPU's integrated circuitry (IC), which is at the heart of the computational process. The more transistors that can be crammed onto the silicon, and the faster they can be made to work, the more processing power can be squeezed out of the CPU. However, with each increase in CPU power comes a corresponding decrease in energy efficiency. IC engineers talk about the 'power wall': the way that performance increases are limited by the ability of the IC to take in electric power and dissipate heat.

When thinking about IC and CPU issues, there are five closely interrelated issues to take into consideration: clock speed (frequency), voltage, the number of transistors, the materials they are made from, and the space they are crammed into. Each of these has a bearing on what is called 'power leakage' or 'leakage power', i.e. wasted current flowing in a circuit that is not being used. Such waste generates heat, which has to be dissipated and is an inefficient use of the electric energy provided to the IC. There are two sources of power leakage in the circuitry: gate oxide leakage (from the materials used in the construction of the circuitry) and sub-threshold leakage, a physical property of IC transistors which is caused by an 'imperfect' switch.

The race to cram more, ever-smaller transistors onto the CPU means that these leakages become progressively greater with each new generation of computer, and devices must either have some form of cooling added (for example an integrated fan), or have their computational power limited (Borkar, 2003). In an attempt to control the amount of heat generated, chip designers gradually decreased the voltage needed by the transistors in order to help deal with the leakage problem.¹⁰² However, there are limits to this, and today's processors exhibit high levels of leakage and thus heat dissipation. Andy Grove of Intel says: "The main worry for the future is leakage power. As opposed

¹⁰¹ As an example, on its release in the UK, the Microsoft Vista operating system, in its most basic form, was not able to run properly on around half of the existing computer stock. In its interim report on Microsoft Vista and Office 2007, BECTA found that "over half of the current infrastructure could not run Vista even with the Aero graphics engine turned off (estimate 55%) and virtually none of the current infrastructure could run Vista with Aero switched on (estimate 0–6%)" (p. 2). See <http://publications.becta.org.uk/display.cfm?resID=28199>

¹⁰² Physics dictates that power consumption by the chip is proportional to the square of the voltage.

to the power that is used productively when transistors do work, leakage power is consumed just by the circuit receiving power. The smaller the transistor, the more leakage there is. It is now growing to be a considerable percentage of the total CPU power, and will only get worse" (Grover, 2003, para. 34).

4.5.2 Short-term fixes for end user equipment problems

There are a number of technologies and ways of working that could be implemented in an educational setting almost immediately, depending on the local situation. James and Hopkinson (2009c) recommend a strategic approach that takes account of the needs of users (through the use of an audit) and that plans for appropriate energy-efficient equipment procurement. As these are primarily managerial processes we don't discuss them here; instead we focus on specific technical issues.

4.5.2.1 Remove screensavers

Screensavers were originally deployed to prevent 'burn-in' on older CRT monitors. As newer CRTs are much less prone to this and it is not something that affects LCD monitors, the original need for screensavers no longer applies. As already mentioned, screensavers often prevent both the PC and monitor going into a low power state so for this reason users should be encouraged to remove screensavers. Computer Weekly recently reported that Barclays Bank has banned screensavers and estimates annual energy cost savings of £1 million.¹⁰³

4.5.2.2 Enable power management

There are likely to be many times during the working day, and certainly at 'off-peak' periods in the evenings and weekends, when the user is not using the machine at all – and yet it is left on. To combat this, manufacturers have come up with ways to allow a computer to enter various 'sleep' or 'idle' states in which energy is conserved. This was formalised in 1996 by the Advanced Configuration Power Interface (ACPI), a cross-platform standard to be adopted by hardware and software manufacturers in order to enable computers and monitors to enter low power modes when not in use.¹⁰⁴

The concept is straightforward and on Windows machines power settings are set through the power management interface. However, its implementation is not always simple, especially with desktop machines (laptops, owing to the need to conserve battery energy when on the move, appear to implement ACPI far more effectively, at least when running on battery power¹⁰⁵). How users actually make use of these facilities has an enormous bearing on energy use over the operational life of a PC, as can be seen from Table 3, which describes a variety of scenarios for desktop and LCD monitor energy consumption patterns over a year based on the average power demands given in Table 2.

¹⁰³ See www.computerweekly.com/Articles/2007/12/04/228429/discovering-the-financial-value-of-green-technology.htm

¹⁰⁴ It currently stands at Version 4.0.

¹⁰⁵ Manufacturers have put considerable thought and effort into how to exploit ACPI as a means of extending battery life between charges. However, as soon as the machine is plugged into a mains power source, such configuration is invariably dropped. This is an unfortunate scenario and one that suggests that battery life between charges is far more important to the consumer than overall energy savings. Energy Star make a note of this on their website: 'On laptops, be sure to activate these [power management] settings in the AC power profile – not just the DC (battery power) profile.' http://www.energystar.gov/index.cfm?c=power_mgt.pr_power_management

	Description	Hardware	Hours per working day			Annual combined kWh	Separated kWh
			On	Sleep	Off		
1	Both always-on (365 days per year)	Desktop	24	0	0	981	688
		LCD monitor	24	0	0		293
2	Desktop always-on (365 day pa), <i>monitor</i> power managed	Desktop	24	0	0	754	688
		LCD monitor	7.5	1.5	15		66
3	Both always on except weekends and holidays	Desktop	24	0	0	644	453
		LCD monitor	24	0	0		192
4	As Scheme 3 but with automatic <i>monitor</i> power management	Desktop	24	0	0	518	453
		LCD monitor	7.5	1.5	15		66
5	As Scheme 3 but with automatic <i>desktop</i> power management	Desktop	7.5	1.5	15	352	160
		LCD monitor	24	0	0		192
6	Fully enabled power management	Desktop	7.5	1.5	15	226	160
		LCD monitor	7.5	1.5	15		66

Table 3 Comparison of desktop and monitor power consumption under a variety of power management regimes¹⁰⁶

4.5.2.3 Power supply efficiency

Internal power supplies on PCs are generally from 50% to 80% efficient, whereas external power supplies that sit outside the PC casing (sometimes referred to as power 'bricks') can be as low as 20% efficient. The 80 PLUS[®] program, created by Climate Savers Computing,¹⁰⁷ sets out to remedy this by putting forward a roadmap for 80% and 85% efficient power supplies. However, efficiency is increasing all the time: for example Cisco is now shipping 90% efficient power supplies. Many of the more efficient power supplies are still offered as options, so institutions are advised to ask their suppliers for the maximum efficiency available.

4.5.2.4 More efficient 'fat clients'

For a marginal increase in capital cost it is possible to buy desktop PCs that are far more energy-efficient than units typically in use. However, up until recently there has been little incentive for these to be purchased in educational institutions because IT procurement budgets do not usually need to take into account the cost of electricity. However, this is likely to change under the CCA and more efficient fat clients will probably become more attractive as a way to reduce energy consumption.

More efficient fat clients make use of low-energy components (some derived from laptops), which means that machines consume around half the energy-in-use of typical Energy Star 4.0 rated desktops (and are around 20% better than Energy Star 5.0). As an example, the UK company VeryPC manufactures a range of desktop PCs which, they claim, use between 26 and 39 W in average use. Such performance is achieved by using low-power disk drives and dual-core

¹⁰⁶ Figures compiled from the MTP 'What if' website (see footnote 95).

¹⁰⁷ A vendor initiative supported by the WWF (see Section 3.3.5).

processors, and future machines could take these efficiency levels further still by using more cores per processor and replacing spinning disk drives with solid-state memory.

A factor in the short-term development of fat clients is the role of 32-bit and 64-bit operating system software. To date, most software and operating systems have been written to function on 32-bit hardware architecture. New versions are being released that take advantage of the development of 64-bit hardware (e.g. can address 2^{64} of computer memory – around 17 million terabytes). Apple's new Snow Leopard operating system is one example, as is the forthcoming Windows 7. The introduction of machines and software for 64-bit is likely to cause a significant increase in average power consumption (and embodied energy in production) of a typical desktop PC.

4.5.2.5 'Slim' clients

Another alternative, just coming to market, are known as ultra-low power PCs, and offer a close approximation to the abilities of a full-power PC. Conceptually these devices sit between more efficient fat clients and thin clients and might be termed 'slim clients'. An example is the Fit-PC 2 which is based on a low-power Intel Atom processor, runs Windows and provides a great deal of the functionality of a standard PC for around 9 W.¹⁰⁸

4.5.3 Near-term developments

There are a number of technologies and techniques that are either expensive or in the experimental stage, with a small number of early adopters undertaking work to explore their applicability. These technologies need to be tracked and evaluated as they become more mature, with an eye to future implementation.

4.5.3.1 Advanced display technologies

LCD monitors are now the norm and are becoming increasingly efficient over time with gradual improvements in the luminous efficacy of their backlighting systems and with reductions in standby consumption. However, innovations in monitor technology have the potential to increase efficiency further, with organic light emitting diode (OLED) technology reported to have the potential to supersede LCD technology by 2020.

The trouble with LCD is that it requires a relatively large amount of energy to backlight the screen, whereas OLED displays don't need to be backlit because OLEDs *are* the light source. Approximate figures for the energy consumption comparison between OLEDs and LCD, by way of a simple illustration, are:¹⁰⁹

1. A 17-inch monitor requires approximately 400 lumens to achieve sufficient brightness.
2. A 17-inch LCD monitor expends around 30 W in achieving sufficient backlighting to achieve this.
3. A 17-inch OLED monitor should require only 4 W to achieve the same brightness.

There also seems to be a general feeling that OLED displays will be easier to manufacture and therefore, perhaps, less resource intensive and with lower embodied carbon emissions. However, detailed figures on this are not yet in the public domain. Samsung are currently shipping OLED televisions and LG are launching a 15-inch monitor in December 2009, although these are likely to

¹⁰⁸ The company provides a comparison sheet that shows what aspects of a standard PC it can and cannot replicate. See <http://www.fit-pc.co.uk/index.html>

¹⁰⁹ These figures are based on rough workings of MTP figures (see footnote 95) and assume OLED display requirements of about 10% of the energy needed to replicate the same brightness of an LCD producing 400 lumens.

(initially) be cost-prohibitive for most institutions. Costs should come down as production ramps up and the technology is widely diffused.

A number of developments in flat screen display technologies were documented in detail in a JISC Technology Watch report (Anderson, 2006) including OLEDs, field emission displays, electronic paper and surface conduction electron-emitter displays (SED). There are also a number of interesting developments in laser-based projection technology. Readers interested in the longer-term development of display technologies are referred to this report.

4.5.3.2 Thin clients

Thin clients are end user devices that may be used to replace desktop computers in client-server networked environments. The thin client has minimal processing power of its own – typically just enough to run the onboard software required to link it with the servers. The thin client is the interface between the monitor, keyboard and mouse of the user and the processing power of the servers.

The processing power of desktop computers tends to be largely redundant for most of the time under typical usage patterns. Average CPU utilisation rates hover around the 5% mark and this is demonstrated by Figure 4, based on a report prepared for Energy Star (Calwell, 2005).

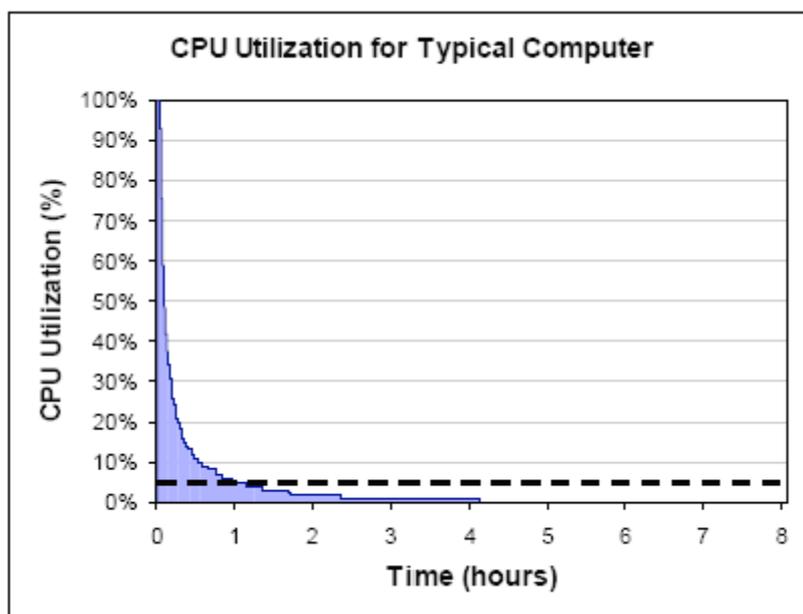


Figure 4 Typical desktop CPU utilisation over a working day (from Calwell, 2005)

Typically, the power demand of a desktop computer may be around 80 W at 0% CPU utilisation, with that demand rising to a maximum of 120 W at 100% utilisation (as an example, the Dell Optiplex 745 Mini Tower – a popular business machine – ranges from 76 W at idle to 114 W at full power). At, say, 10% utilisation, 84 W of power may be drawn: the additional processing utility being provided by only 4 W, with the remaining 80 W being required to keep the machine running.

With thin clients, the processing is provided by external servers. In a typical environment, a ratio of around 20 thin clients per server may be used. The average CPU utilisation of the server would therefore generally be far higher than that of the average desktop. As a simplified example, if, for instance, an average utilisation of 60% were achieved on a server with, say, double the power draw of a typical desktop (and assuming the server has a similar demand profile to that of the desktop), then overall power demand would be 208 W ($160 + 60\%$ of 80). This equates to a little over 10 W

of back-office processing power per thin client. If this figure is doubled to account for air-conditioning, etc. in the server room, this leaves 20 W per thin client. One must then add to this the power draw of the thin client itself; assuming this is another 10 W, the overall power demand rises to 30 W.

In addition, owing to their small form factor, both in terms of bulk materials and lower overall processing power requirements, a thin client environment should have far lower embodied GHG emissions than its equivalent networked desktop environment. Furthermore, owing to the fact that software is stored on the server rather than the client, there is a high likelihood of significantly extended life cycles: the machines, unlike desktops, do not become redundant after a couple of years when the latest new operating system comes along.

Great claims are often made for the potential of thin client technology with both BECTA¹¹⁰ and the Fraunhofer Institute¹¹¹ providing detailed discussion of their findings. In addition, public sector agencies such as the Rural Payments Agency are beginning to experiment with thin client technology, aiming to replace 5,000 PCs with Sun Ray thin clients.¹¹²

However, thin client solutions are not without their problems. James and Hopkinson (2009c) list a number of caveats with regard to their use within the HE/FE sector. Of particular note is their view that there is, to date, no clear evidence that thin client devices will in practice have significantly longer lives than desktops in educational settings. Added to this are concerns that some thin clients are increasing rather than decreasing their energy consumption and, at least within HE, may be used to supplement rather than replace desktops. This is a complex area and still open to considerable debate.

4.5.3.3 Netbooks and mobile Internet access

One trend is the rise of low cost and portable computing devices whose primary function is to provide access to the Internet and thus to a variety of online services. There is a wide variety of netbooks, smartphones and other hand-held devices and the boundaries increasingly blur. In a sense they are a form of thin client that interacts with the 'cloud' (see Section 5.1.1) rather than a local client/server system. Indeed, many of the caveats regarding thin clients are likely to apply. Because they are designed to be highly portable and thus rely on battery power they have been designed to be energy efficient, using much the same technology as laptops. However, as the overall 'energy TCO' of connecting to the cloud goes beyond individual devices (see Section 5.1.1), it remains to be seen whether these devices can really be considered an improvement on the desktop PC. As with PCs, researchers and technologists are beginning to explore the longer-term energy issues with these devices. Belaramani and Dahlin (2009), for example, explore the energy costs of the synchronisation of data between these kinds of devices and the Internet cloud.

4.5.3.4 Shared PCs

One new architectural model for delivering applications and office tools to end users is the shared PC. The concept is simple: take a PC and connect multiple monitors to it, with a keyboard and

¹¹⁰ The BECTA study into thin client use in schools found that these systems used less energy, produced less waste heat (and noise of fans) and allowed older PCs to remain in use (through reconfiguration to a thin client device). They did however note a number of problems; see <http://partners.becta.org.uk/index.php?section=rh&rid=13802>

¹¹¹ The Fraunhofer work found that replacement of desktop computers with thin clients could reduce overall emissions per terminal by at least 54%. Their work considered both operational and embodied GHG emissions. See http://it.umsicht.fraunhofer.de/TCecology/index_en.html

¹¹² See <http://www.sun.com/customers/software/rpa.xml>

mouse per monitor. Up to eight monitors (although six is a practical limit) can be connected to a single PC, and all the users can connect independently and simultaneously.^{113,114}

As with thin client environments, the model relies on the fact that the processors in modern PCs, where used in normal office environments, are heavily under-utilised. As mentioned previously, average CPU utilisation rates rarely average more than 5%. With six users sharing the same CPU, there should be ample capacity but at a fraction of power demand (Figure 5).

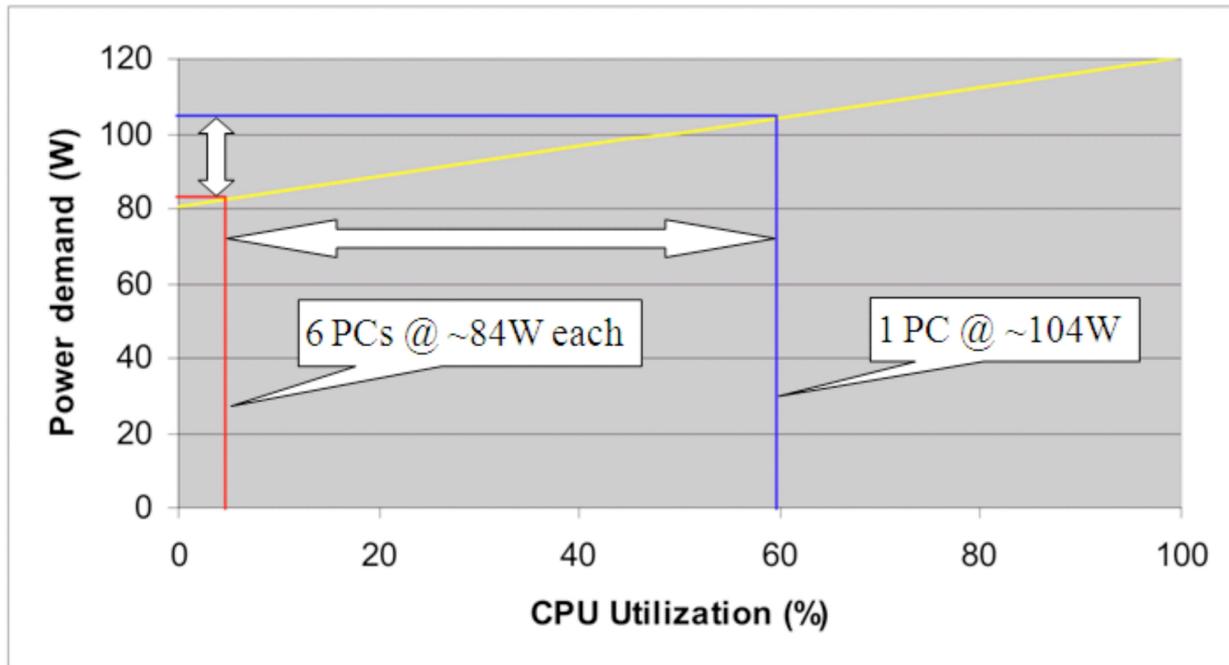


Figure 5 CPU utilisation rates

The concept is essentially a lighter version of the thin client environment but a single desktop is used as the server. The environmental benefits in terms of both embodied and operational GHG emissions will be similar. The set-up may be best suited to small office spaces.

4.5.3.5 Motherboard design

Many motherboard designs require three different DC voltages and the microprocessor itself works at yet another voltage. Power saving could come from designs that allow for a single DC voltage across the whole board and a further reduction in this single voltage level itself could have "dramatic effects" on power consumption (Forge, 2007). Google has been active in this area, working directly with hardware vendors to make sure that all the components on the motherboard use a single voltage level (Hölzle, 2005), and introducing Dynamic Voltage Scaling, a power management technique first developed for laptops and other mobile devices (Zhai *et al.*, 2004).¹¹⁵

¹¹³ NComputing is one company offering this concept (<http://www.ncomputing.com/>).

¹¹⁴ See <http://emergingtechnologies.becta.org.uk/index.php?section=etn&rid=14501>

¹¹⁵ The use of this in HPC environments is discussed at length in Kappiah, N. 2005. **Just-in-Time Dynamic Voltage Scaling: Exploiting Inter-Node Slack to Save Energy in MPI Programs**. See http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1559985

4.5.3.6 CPU developments

What is interesting about the discussion of the future direction for CPUs is the way that traditional 'rules of thumb' of hard core computer and electronic science are being upgraded to take account of the new imperative for energy efficiency. In order to deliver improved processor speeds and move towards what is called gigascale computing there are several developments to watch out for: power consumption, multi-core processors and photonics. These developments, some of which are still in the laboratory phase, point to the direction of travel for energy-efficient computing devices.

a) Power consumption

Development work on the core electronics design and physical properties of IC components continues with a focus on reducing the power input requirements. Intel, for example, has introduced various improvements in the design of their core processors that assist with energy efficiency and has also undertaken research and development in the materials science of the silicon used for the processor's IC gates, resulting in a new hafnium-based 45 nm gate (known as the Nehalem and introduced in 2008), which the company argues reduces electricity leakage from the transistors.¹¹⁶ AMD have also been working on a 45 nm design, codenamed Shanghai, and are making similar claims as to improved energy efficiency. These new 45 nm IC designs are currently making their way into computer products, with plans for 32 nm announced in September 2009.

A related area is that of the IC industry's pursuit of zero-standby-dissipation technology (Kimura, 2009). The details are complex but in essence the idea is that using non-volatile chip technology will allow equipment to be placed into very low power standby modes (at a few percent of the energy that is currently used) when the user is not making use of the device. The device can instantly turn itself on again when required. This will allow the development of PCs and other devices that operate in what is being called 'normally-off' mode.

b) Multi-core

There has been a recent trend amongst CPU manufacturers towards multi-core processors, with several CPUs on a single integrated circuit, sometimes referred to as Chip MultiProcessors (CMP). Typically, this has been a dual-core or quad-core, although there are plans for oct-core and researchers are actively looking at many-core processors with hundreds or even thousands of cores on a single IC die (Woo and Lee, 2008). Well-known examples of multi-core include Intel's Celeron dual-core and PowerPC 970MP (which is used in high-end Apple PCs). Chip manufacturers like Intel argue that a move to multi-core not only allows increases in computational processing power but can offer advantages in energy efficiency.

Intel continues to pursue such developments and is looking ahead to what it describes as the Tera-scale computing era, in which processors will be expected to deliver a trillion floating point operations per second. Their R&D programme, the Tera-Scale Computing Research Programme, is looking at the challenges of developing systems with dozens or even hundreds of energy-efficient cores. Intel argues that such parallel systems offer considerable scope for what they call Energy Management, allowing control of the thermal load by judiciously spreading computational tasks around the tera-processor's individual component cores.¹¹⁷ These developments will also combine with work taking place at the circuit level in the physics of IC circuits, an area in which Intel has set a goal of '10 times' improvement in performance per watt over the next 10 years.

Keeping down the power requirements of these individual cores will be of paramount importance according to Woo and Lee (2008), who argue that power is becoming the critical issue in scaling up the number of cores that an IC die can accommodate and they propose an appropriate update to

¹¹⁶ See <http://www.intel.com/technology/architecture-silicon/next-gen/whitepaper.pdf>

¹¹⁷ See <http://techresearch.intel.com/articles/Tera-Scale/1421.htm>

Amdahl's law of scaling.¹¹⁸ In part this is because of the problems inherent in dissipating the heat generated from modern processors, but as new pressure on reducing input energy takes hold this will also become important. They propose that including power in the Amdahl equation leads to the conclusion that the optimal solution is to have many small, energy-efficient processors integrated (on the same die) with a single, 'full-blown' processor. They point to the examples of Sony-Toshiba-IBM's Cell Broadband Engine¹¹⁹ and the Parallel-On-Die (POD)¹²⁰ project.

c) Photonics

Interest is growing in the potential for photonics to be used for handling on-chip communication between artefacts in the CPU (for example, handling communication between different cores within a multi-core system). As the roadmap for future CPU development is increasingly focused on multi-core, the requirements for on-chip communication will become more and more important. Shacham *et al.* (2007) offers an introduction to these issues.

This integration of optical comms into CPUs and ICs is considered to be a "Holy Grail because it could deliver the superior performance of optics at the same costs structure as electronics (or better)" (Gunn, 2006, p. 59). It will also help to reduce energy use because developments in conventional electrical interconnections currently consume more and more power.

4.5.3.7 Operating system design

Related to the design of computer electronics is the issue of the operating system that controls the basic functionality of the machine. The well-known operating systems expert Andrew Tanenbaum (2002) documents two main ways in which the operating system's kernel can play a part in energy reduction: to turn off parts of the computer (mostly input/output devices and the hard disk) when they are not in use, and to tell the application programs to use less energy (which may involve some degradation in an application's performance, although this may not be detectable by the user). Thanks to the levels of interest in lower carbon computing, energy efficiency is likely to feature heavily in operating system design of the future. Many of the ideas and research areas for this have emerged from the ubiquitous computing and wireless sensor research domains where extremely low power and long-life batteries are the order of the day and this has dictated the development of specialist operating systems.

There are two points to make in general. Firstly, a new or revised operating system's capacity to undertake energy efficiency-related tasks and techniques is likely to become a focus of attention for purchasers (and hence marketers) in the coming years. As one recent example, Microsoft's Vista operating system is claimed in press releases and other company materials to be its most energy efficient to date, with a number of sleep and standby settings. There have also been discussions in news items and blogs about the relative energy-efficiency merits of Windows and Linux.¹²¹ Secondly, there may well be some element of configurability in these operating system techniques, and thus becomes an area which institutional IS technical staff may well need a deeper level of engagement than the end users.

Another point to consider with operating systems is that a particular operating system and its newer release may demand an increased specification for the hardware. There is much talk in the industry

¹¹⁸ A formula for working out the maximum speed-up (performance) of a processor architecture through use of many processors and parallelisation. The original formula, by chip design Gene Amdahl, dates back to the late 1960s and does not take account of power or energy use.

¹¹⁹ See <http://www.ibm.com/developerworks/power/cell/>

¹²⁰ See <http://arch.ece.gatech.edu/pod.html>

¹²¹ See, for example, *Linux uses 12 percent less power than Windows 2008, study finds* (http://news.cnet.com/8301-13505_3-9964563-16.html).

of 'bloat-ware' and needing to 'beef-up' hardware on every software upgrade. Microsoft engineers, for example, will admit that they build a new version of the operating system with a view to what the typical off-the-shelf, new hardware will be in perhaps 6 years' time (Anderson, 2007). The launch of Vista, in particular, caused considerable controversy in this regard with a number of activist organisations such as Greenpeace protesting and a recent report by the UK Office of Government Commerce, concerning trials of open source software, came to the conclusion that a move to open source systems can be a way of avoiding this cycle of forced upgrade (OGC, 2007).

4.5.3.8 Wild cards

Finally, there remain a few wild cards – ideas that are just that: interesting ideas that represent a different way of thinking that diverges from the status quo. They are presented here as insight that may provide a 'left-field' answer to energy efficiency in the future.

a) Clock speed

Until recently, chip design focused mainly on increasing the frequency at which the processor operated: the higher the clock speed, the faster the transistors work and the more processing power the CPU can deliver. But this came at a cost in terms of increased power consumption and heat generation from the processors. According to figures presented at an academic conference on the future of processors, the actual clock itself can use substantial amounts of energy, with opinions ranging from 25% to as much as 60% of the energy used in the processor. The author, Gustafson (2007), points out that the human brain runs at approximately 1 Hz compared to upwards of 2.5 GHz in a computer and questions the logic of trying to chase ever-faster clock speeds. He argues that thought should be put into how to increase the work done within longer clock cycles.

4.6 Printers and printing-related technology

Colleges and universities make considerable use of printer technology and related digital image systems such as photocopiers. James and Hopkinson (2009d) estimate that perhaps 2–3% of the UK tertiary education sector's total electricity consumption, and at least 10% of its ICT-related consumption, is related to such digital printing.

4.6.1 The state of play: problems with printing

Many of the problems produced by printing are in fact problems with changing user behaviour rather than technology: users are resistant to sharing printers, there is a tendency to print out documents 'just in case', and some people still print emails. In general, changing user behaviour is outside the scope of this report but we include an overview of the relevant technology and environmental impact, in particular:

- the environmental footprint due to the creation of the paper
- the impact through the life cycle of the ink/toner
- the embodied footprint of the printer device itself
- the energy use of the printer.

Whilst a range of detailed figures for different technologies is available we present an indicative example based on energy costs entailed in the life cycle of printing/copying one sheet of paper as outlined in James and Hopkinson (2009d).

Firstly, there is the creation of the actual piece of paper: 17 Wh is consumed in making a sheet of paper from virgin pulp, and 12 Wh when made from recycled material (the latter is also, in most cases, a cleaner manufacturing process). Secondly, there is the energy consumption of the printer, which, based on a laserjet printer when standing idle (switched on, ready, but not printing) is around 285 W.¹²² When a user prints a sheet of paper there is an additional cost of typically 2–4 Wh, which gives a rough total for printing a sheet of paper of around 300 W.¹²³

4.6.2 Short-term fixes for printers and copiers

There are a number of relatively easy fixes that could improve the environmental impact of printer technology and it hardly goes without saying that purchasing lower weight, recycled paper is a quick win, along with setting printers to default black/white and duplex printing rather than single-sided colour (which can be over-ridden manually if need be). Many organisations produce guidance on quick fix measures and it is now common for universities and colleges to have advice in place on how to reduce the amount of printing and on techniques for saving paper and ink.¹²⁴ The prevalence of information in this area obviates the need for a detailed review of quick fix measures in this report, and institutions that have not yet begun to implement printing-related savings are advised to review existing advice and target simple changes first, such as putting printers onto timers or setting up the sleep mode to ensure that they power down out of office hours.

¹²² Compare this with a standard desktop PC, which uses perhaps around 60 W when idling (i.e. switched on and ready, but not processing a complex application).

¹²³ James and Hopkinson report that the production of toner and inks is fairly negligible when compared to the energy and other environmental costs of paper and the physical equipment.

¹²⁴ London Met, for example, has produced a handout: www.londonmet.ac.uk/londonmet/library/r98029_3.pdf and the University of Aberdeen has an energy management plan which includes ways to reduce printing costs: <http://www.abdn.ac.uk/environment/energy/management.php>

Longer-term solutions, albeit based on mature technologies that are currently available, require a print management strategy, which involves some form of centralised control and even auditing of the printing process. Such methods can reduce print volumes significantly, with 20% being quoted by James and Hopkinson (2009d) but again, they are predominantly about changing behaviour rather than technology per se, so we provide only a summary here. The interested reader is referred to the James and Hopkinson report.

A simple print strategy has around four main components:

- print and paper auditing
- print substitution
- consolidation of print devices
- print management (including use of centralised print control and management¹²⁵).

The first involves a process of auditing what the current situation is with regard to printing in an institution, followed by some element of monitoring and improvement. Substitution involves looking at alternatives to printing (e.g. viewing a long PDF on a display screen or e-reader device rather than printing). This may involve some quite detailed and far-reaching strategic review and decision-making on the information life cycle in an organisation.

If a refresh or office move is due, office printers can be consolidated to larger multi-functional devices (MFDs). These devices bring greater numbers of features to those who really need to print, but are by their nature further away from the average user. In theory this should reduce the number of print requests as the walk down the corridor to collect the printed material often makes people think twice about whether they really need it (a form of behavioural 'nudge'). MFDs also have a lower embodied footprint and a lower operational energy use than the separate devices they replace. Xerox has calculated that, on average, a single multifunctional printer, copier and fax replaces one copier, four smaller printers and a fax, and halves the annual energy requirement from 1400 kWh to 700 kWh (Intellect, 2008).

However, it is very important that in situations where MFDs are deployed, aggressive power management (e.g. fast sleep times, device selected on standby power use) is used otherwise it is quite likely that overall power consumption will increase. One further change that can be made where networked MFDs are available is print on collect. This technology does not print the document when the user requests it, but when they go to pick it up from the printer (by typing in a PIN). This reduces both wasted printing and user frustration at someone else taking their urgent printout.

Finally, as is the case with other ICT equipment, careful procurement and a planned refresh cycle can help reduce the energy costs of printing.¹²⁶ As new technology is introduced and as schemes like the Energy Star are improved, newer machines are likely to use less energy.

4.6.3 Near-term developments

4.6.3.1 Reusable paper

Many projects have investigated ways to take an ordinary printed sheet of paper and prepare it for reuse by removing the existing ink (i.e. within the office, rather than sending it off for recycling at a

¹²⁵ These make use of rules-based software to route different types of print jobs to different printers. For example, they can be set so that when something relatively unimportant such as an email is printed out it is sent to a printer that uses low-quality, recycled paper.

¹²⁶ As we have seen with discussions over replacing office PCs, embodied energy and waste disposal have to be taken into account.

plant). Counsell and Allwood (2006) reviewed 104 patents filed in this technology area and methods include decolouring the toner or ink, obscuring existing print, and various chemical methods for print removal.

Of these, one of the more viable may be Xerox's reusable paper, a concept that has at least made it to the demonstration phase by researchers at the Palo Alto Research Center. It involves the use of special paper that has been coated with photosensitive chemicals and which fades after around 24 hours, allowing the piece of paper to be reused the next day. Xerox claims that 25% of the printing undertaken in a typical office results in a piece of paper that is read and then put in the bin on the same day.¹²⁷

4.6.3.2 Electronic paper

A development that is likely to have significant benefits for low carbon computing is that of electrophoretic electronic paper, or 'e-paper'. This is essentially a form of display technology that aims to emulate the extremely high level of readability and flexibility of ordinary paper, by developing displays that are very portable, lightweight and flexible. E-paper consists of a sheet of thin plastic made up of thousands of tiny bubbles, each of which contains 'ink' which is activated on application of an electric charge. Each bubble acts as a single pixel in the 'paper' display, being either black or white depending on the charge that is applied. Ideally, these systems should be bistable (the visible effect of which is that once the ink particles are in the required position they will stay there even after the charge is removed), thereby requiring very little power. A printed page could be created on the e-paper, read by the user, and then 'wiped' to be reused for another page as required.

¹²⁷ See http://news.cnet.com/8301-11128_3-9930674-54.html

5. To 2050 and beyond

Forty years since the Eagle landed on the moon, the idea of a new Apollo project has become shorthand for how we should tackle climate change: politics forcing through the technological limits, a decade-long push, and a nation unified for a shared goal.

Ed Miliband, The Guardian (Comment is Free blog), 20 July 2009

The challenge has been set. Science has shown us the limits of our world's atmosphere and tackling climate change will be a major goal for the rest of our lives. How, then, might ICT in education and other parts of the public sector respond to this over the long term? What are the goals for our sector's 'Apollo mission'?

In Section 4 we considered individual technical areas and, for each, considered some of the short- and medium-term technology developments that might be employed in order to implement the kind of energy reduction required by individual HEIs in the run-up to 2020. In the longer term, however, there is a need to rise above the level of the institution and engage in sector-wide strategic thinking. As the sociologist Anthony Giddens (2009) argues, in order to deal with climate change there is a society-wide need for a return to more planning. With this in mind, this section of the report looks at some 'big picture' technical areas for the sector, including the future of the data centre and cloud computing, how we might share ICT resources, and how we might generate energy for ICT. First of all, though, we need to get a feel for the kinds of pressures ICT will face in 2020.

It is anticipated that the measures outlined in the EU Eco scenario (see Section 2.1.1) will, for example, reduce average electricity requirements from desktop computers by an estimated 40%, laptops by 25% and monitors by 38% by 2020, compared with a 2005 baseline. However, because the actual numbers of stock are set to grow, the overall percentage reduction for desktops is pared back to only 13% by 2020.¹²⁸ This stock growth factor is worse for laptops and monitors, where the total energy use is almost trebled and doubled respectively by 2020.¹²⁹

This tells us that due to the growing numbers of ICT devices, even if we adopt the measures outlined in the EU's Eco scenario, public sector institutions will not meet the UK's target of a 30% reduction in energy use by 2020.¹³⁰ This means that public sector ICT will need to innovate at an even faster pace than that envisaged by the EU's Eco scenario. This is further complicated by the knowledge that ICT can contribute to a reduction in GHG emissions from other sectors through dematerialisation – the move from the physical to the virtual – and adding 'intelligence'. As the Centre for Alternative Technology, as just one example, notes in their recent report on moving to a 'zero carbon' Britain, ICT will have a big role to play in energy reduction through so-called smart grids, intelligent buildings and in changes to the way people work and travel (CAT, 2009). Great claims are made for this transformation potential of ICT and it is likely that HE will face a situation where – thanks to the CRC and similar measures – there will be considerable pressure on individual institutions to both reduce emissions produced by ICT, and to make more use of it in order to assist other areas. Added to this is the inevitable growth in the education sector in general as more people opt for tertiary education.

How might this circle be squared? Before we plunge into some technology areas that could help, we should note in passing the potential importance of Enterprise Architecture (EA) to this kind of big picture planning. EA is a high-level, strategic technique designed to help senior managers achieve

¹²⁸ Total desktop estimated electricity drops from 23 TWh/a to 20 TWh/a (see table in Bio Intelligence Service, 2008, p. 47).

¹²⁹ Total laptop estimated electricity goes from 5 TWh/a to 16 TWh/a by 2020.

¹³⁰ As outlined in the UK Low Carbon Transition Plan (DECC, 2009).

business and organisational change. It provides techniques for aligning organisational vision and strategy with business operations and supporting ICT provision. In recent months JISC, UCISA and others have been working together to explore the potential for this technique to help transform the sector and, in particular, lay the groundwork for common ways of working and shared services. Environmental considerations are starting to be fed into discussions about architecture and The Open Group, a leading proponent of EA, is starting to explore the connections between sustainability and institutional information architectures.¹³¹

5.1 The future of the data centre

Universities and colleges are particularly data-rich organisations. Whether we consider the future role of libraries as electronic repositories or the ever-increasing volume of experimental data that is being generated and stored, we can't escape the fact that data is a big part of an institution's *raison d'être*. It is therefore vitally important when considering the long-term future of institutional ICT infrastructure to think about the future of the data centre.

The Rocky Mountain Institute, which is well known in the field of energy reduction, has proposed that the data centres of the future could use 95% less energy. They argue that of 100 W going into today's data centre only 2.5 W gets to the useful computational part (Palmitier and Newman, 2008) and they map out some strategies for the future design of data centres (using today's technologies), some of which have already been discussed in Section 4. There is no doubt that individual institutions will need to upgrade and refit their existing data centres in more energy-efficient ways. However, new challenges are emerging to the basic idea of the departmental or institutional data centre that need to be considered in any long-term planning process.

5.1.1 Cloud computing

Many in the computer industry argue that we are in the midst of a paradigm shift as computing moves away from desktop PCs and migrates to the Internet 'cloud'. The development of broadband has led to reliable online access to third-party software and data storage services accessed through a browser, a process that is becoming known under various terms including service clouds, cloud computing or Software as a Service (SaaS). A detailed discussion of the different terminology and services is beyond the remit of this report but interested readers are referred to recent work by CETIS in this area.¹³² For the purpose of this report we will refer to these techniques generally as cloud computing. Although, strictly speaking, cloud computing is not a future technology (in that it is available now), migrating to the cloud is not without its problems, some of which (data security, legal jurisdiction, etc.) have been introduced elsewhere in this report.

Planning for lower carbon ICT will inevitably involve examining the potential of a move to the 'cloud', whether this is provided for as a commercial arrangement or as part of some sort of sector-wide shared service. Indeed, some universities are already starting to explore the potential for shared data centres. The Universities of Derby, Salford and Sheffield Hallam, for example, recently undertook a study to look at the potential for such a centre under the auspices of HEFCE¹³³ and Yorkshire and Humberside Metro-area Network (YHMAN) are looking into a shared virtual data centre under the same series of feasibility studies.

¹³¹ JISC has published two substantial reports in this area: *Doing Enterprise Architecture: Enabling the agile institution* (April 2009) and *Unleashing EA: Institutional Architectures and the value of joined up thinking* (July 2009). Both are available from <http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/earlyadopters.aspx>

¹³² See <http://blogs.cetis.ac.uk/cetisli/2009/07/07/cloud-computing-in-institutions/>

¹³³ See <http://www.hefce.ac.uk/Finance/shared/feasibility/show.asp?id=12&cat=1>

JISC is currently preparing a series of reports into academic use of cloud computing that will consider many of the issues in more detail, but for now it is sufficient to say that the environmental benefits of cloud computing are dependent upon two key factors:

1. The GHG emissions associated with the client: much has been made of the suitability of dedicated, low-power devices. However, the optimal situation is that cloud computing clients should consume very little energy¹³⁴ and this does not exclude desktop devices per se.
2. The GHG emissions associated with the cloud.

This last is starting to be seen as increasingly important. Berl *et al.* (2009) note that: "In principle, cloud computing can be an inherently energy-efficient technology for ICT provided that its potential for significant energy savings that have so far focused on hardware aspects, can be fully explored with respect to system operations and networking aspects" (p. 1). The authors argue that in order to develop energy-efficient cloud computing a comprehensive approach will need to be taken including energy efficiency at all system levels: server hardware, cooling systems, networks, communications protocols and the service or applications themselves. They argue that while much work has been done at the hardware (server) end, an area of work that needs further research is the optimisation of software to take account of working in a cloud environment (i.e. potentially spread over more than one server) in an energy-efficient manner. Techniques are likely to include new forms of workload scheduling that take account of energy use and quality of service trade-offs between performance/availability and energy use. Harizopoulos *et al.* (2009) agree, arguing that more effort needs to be put into how software can manage data in an energy-efficient manner. Liu *et al.* (2009) also agree, arguing that the most important work required to develop what they term as 'GreenCloud' relates to virtualisation, and the energy considerations of migrating and consolidating VMs onto physical servers.

Another aspect of the GHG emissions associated with the cloud is that of networks. As we have seen in Section 4, network equipment and associated protocols are not energy efficient. One implication of the rise of cloud computing will be a consequent rise in network traffic as users seek to access their remote data and make use of software services. The increase in network use will therefore be particularly inefficient.

As we have already seen, a number of techniques have been proposed for short-term fixes, including adapting bandwidth based on need and putting modems and network cards into 'sleep' states when not required. However, these are only local measures, applying to individual pieces of equipment, and new research is needed into making use of these energy-saving techniques at a network-wide level. Such techniques are referred to as 'traffic engineering' or 'shaping' and although controversial¹³⁵ are a growing area of interest. Vasic and Kotic (2008) propose such energy-aware traffic engineering for the network as a whole; Bathula *et al.* (2009) propose traffic shaping for optical networks.

In addition, network designers and researchers are starting to take into account the need to design energy efficiency into their systems and Bolla *et al.* (2009) argue that power consumption should be considered as a metric for network performance evaluation comparable to more traditional measures such as throughput and delay.

¹³⁴ See, for example, the CherryPal (www.cherrypal.com), which is claimed to consume just 2 W of power in use and have low embodied GHG emissions, or the 'slim client' concept discussed in Section 4.5.2.5.

¹³⁵ Similar techniques can be used to give some IP traffic priority over others for commercial rather than environmental reasons, resulting in what some argue would be a 'two-tier' standard of service.

5.1.1.1 G-Cloud

The recently published Digital Britain report (BIS/DCMS, 2009) included the first public discussion of internal government plans for the use of cloud computing by the public sector. These plans are being fleshed out by the CIO Council and Intellect. The vision is for a private cloud for government – the 'G-Cloud' – and a roadmap has been established for moving towards its creation as part of the rationalisation of government and public sector data centres. John Suffolk, Her Majesty's Government Chief Information Officer (CIO), outlined the context of these plans as part of an information strategy with ten strands in a blog posting. This included the statement: "Rationalise the data centre estate...Reduce from the central government 130+ to c9–12. Design a data centre eco system that is scalable, secure, green and economical" (Suffolk, 2009, point 3). He also discusses the creation of a Government Application Store (G-AS).

5.1.2 Adaptive data centres

If cloud computing is to be a key tool for the move to lower carbon computing then, as we have seen, the GHG emissions associated with the actual cloud will need to be much lower. In general, there is great interest in the ideas surrounding the concept of the smartgrid, in which we distribute and use energy in a more intelligent manner thanks to a convergence of electricity distribution technologies and the Internet. It is likely that cloud computing and shared data centres will play their part in this work. Andy Hopper and colleagues at the University of Cambridge propose a new approach to data centres – making them what they term **adaptive** (Hopper and Rice, 2008).

Existing data centres are built to maximise service availability and to provide considerable redundancy. A typical tier 4 data centre (the highest level of fault tolerance) requires an additional watt of energy for the fault-tolerance support infrastructure for every watt of energy that actually reaches a server. Such centres are designed to provide a constant service to their users. Adaptive data centres would make use of the heterogeneity that actually exists between the required service and fault tolerance levels of different applications in order to reduce the support infrastructure energy overhead.

How far could such work go? Researchers at MIT announced over the summer of 2009 that they had been working on algorithms for controlling data centres that routed data processing around the world based on the prevailing electricity price in a particular region (Knight, 2009). Although these ideas have a potential fit with the idea of acting in a more intelligent fashion with regard to energy demand and use, they raise the possibility that the driver for such techniques will be economic rather than environmental. This so-called 'carbon leakage' raises questions about how responsible a worldwide cloud service will actually be in practice. For example, will there be a tendency to offload data processing to countries that can provide cheaper electricity because they have developed a low carbon energy system or because they are cheap for other reasons? How will issues of legal jurisdiction be resolved in these scenarios? These are big questions that will test the world's economists and politicians in coming years, but it goes without saying that if our overall aim is to reduce *global* emissions, then badly conceived quick fixes that seek to offload the problem elsewhere will eventually come unstuck. Cloud computing should not provide a mechanism for offloading our own sector's energy responsibilities to commercial third parties or even other countries.

5.1.3 Nano data centres: rethinking the Internet and data centre

Whilst attention within the computer industry is focused on the idea of cloud computing and centralised data processing, others are thinking about the situation in a different way. Researchers at Telefonica and Thomson Research recently proposed what they call a "radical solution to data

hosting and delivery for the Internet of the future" (Laoutaris *et al.*, 2008). They propose hosting data at the network's edge using a combination of the ideas behind data centres and peer-to-peer networking. Their scheme, ECHOS, would make use of the 'boxes' that provide access to the Internet at its edge, e.g. cable modems, ADSL boxes and set-top boxes to create what they call 'nano data centres'. By making use of the storage capacity of these devices through a mass peer-to-peer process they envisage a scenario where much of the Internet's data load (particularly video-related material) would be handled by nano data centres, leaving major, centralised data centres to cope with demand spikes.

The authors give no figures for the potential to reduce energy but point out that over-capacity in centralised data centres is absorbing huge amounts of unnecessary energy and that set-top boxes and cable modems are already drawing energy when their users are at home. The authors note that much research is needed to take this from an idea to a working model but others have started to explore the idea further (Hea *et al.*, 2009). Blackburn and Christensen (2009) take the idea even further, proposing that network operators install file-sharing devices on the side of residential dwellings as part of the broadband service. These boxes would make use of peer-to-peer technology through a process they refer to as Green BitTorrent. This extends the well-known BitTorrent algorithm to allow for 'sleep' modes amongst the network nodes.

5.1.4 Virtualisation

As we have already seen, virtualisation is seen as an important technique for reducing energy in the data centre through consolidation of equipment and it is a key technique in cloud computing. Commentators on low carbon computing regularly cite virtualisation as a technique that is likely to become even more important (see O'Connor, 2008, for example). A presentation by senior VMWare staff at the 33rd International Symposium on Computer Architecture details some thoughts on future directions (Herrod, 2006), noting that computer hardware manufacturers are increasingly incorporating technologies that offer explicit support for virtualisation and that operating systems are being built to be able to handle VMs.

A more recent development, essentially a throwback to the 1960s, is IBM's experimentation with virtualising large numbers of individual data centre servers onto a single, back-office mainframe, which they argue can save as much as 80% of the power consumption (Doyle, 2008). The company has a project called BigGreen, which has virtualised nearly 4,000 servers onto 30 System Z mainframes.

This is a complex area, and while it is beyond the scope of this report to go into it in depth, it is recommended that university IS staff continue to track developments in virtualisation. In particular, it is virtualisation's place in the wider IS context that is of interest and this relates to the growing interest in EA.¹³⁶ For example, Francis and Richardson (2009) recently presented a Green Maturity Model for Virtualisation, aimed at information and enterprise architects. They point out that good architects consider many factors when designing applications and information systems and increasingly one of those factors is environmental. They argue for 'green' information architectural design and present virtualisation as a key technique. They also present a formula for defining the return on the embodied energy costs of buying a new virtualisation server. We are likely to hear a lot more over the next few years about EA, virtualisation, cloud computing and their relationship with lower carbon ICT.

¹³⁶ See <http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/earlyadopters.aspx>

5.2 Energy security for ICT

Season policy with a dash of utopian thinking. Why? Because, however it happens, we are working our way towards a form of society that eventually will be quite different from the one in which we live today. We have to chance our arm.

Anthony Giddens, *The Politics of Climate Change*, 2009, p. 13

When planning for the medium to long term (2020 to 2050 and beyond), universities may well have to give consideration not only to reducing their consumption of electricity, but also to generating their own supplies. This should partly be considered a defensive move: it is entirely possible that we will be entering an age when relying on the national grid to supply the amount of energy required by data-intensive HEIs might become limited. On the other hand this could be considered to be a unique opportunity for universities and colleges to plan ahead, think 'outside the box' and provide leadership in new innovation. To take one example of this kind of thinking, UC San Diego in the USA is in the process of building a zero carbon emission energy system for its data centre needs. The system is a combination of solar technologies and fuel cells. An electrochemical process will be used to generate electricity from waste methane (from a nearby household sewage plant) and the electricity will be stored in a variety of batteries and air-compression systems.¹³⁷ 2 MW of solar panelling will provide additional energy.

5.2.1 Renewable micro-generation

It is almost certain that there will be a complete redevelopment of the electricity supply process over the next 20 years (see, for example, Maloney *et al.*, 2008). This is known as the smartgrid and the vision is of a convergence between the Internet and the energy supply. The Internet will be used to exchange telemetry data that will optimise how the electricity supply system is used and also allow individuals and small-scale generators to 'feed in' electricity from their own renewable supply operations. In this vision, each of the big energy utilities will be just one of many nodes on a Europe-wide electricity supply network. Individual consumers and business users will make far more use of energy consumption data and will switch to appliances that can 'interact' with the energy grid, e.g. to monitor and control their own energy use to match information about the current state of electricity. One of the implications of all this, for the individual HE/FE institution or perhaps even faculty, will be that of micro-generation.

In the future there will be an increased reliance on the generation of local power, for local needs, via various forms of renewable micro-generation and the Stern Review indicates that micro-generation could provide between 30% and 40% of the UK's electricity demand by 2050. It is likely that we are entering an age of decentralised power supply in which small, local sources of energy through solar, wind and hydro will be located in buildings and workplaces and have increasing importance. A recent Foresight report (Hinnells, 2007) summarises the technologies and issues with regard to micro-generation, and both the Micropower Council¹³⁸ and the University of Southampton's Sustainable Energy Research Group¹³⁹ provide useful general information. A number of government and regional strategies are starting to emerge to help develop micro-generation in the public sector. Examples include the Scottish¹⁴⁰ and Welsh¹⁴¹ Energy Efficiency and Micro-generation Strategies.

¹³⁷ See <http://ucsdnews.ucsd.edu/newsrel/science/12-08FuelCellProject.asp>

¹³⁸ See <http://www.micropower.co.uk/welcome.html>

¹³⁹ See <http://www.energy.soton.ac.uk/microgen/microgen.html>

¹⁴⁰ See <http://cci.scot.nhs.uk/Publications/2007/03/09144516/0>

¹⁴¹ See <http://wales.gov.uk/about/departments/dein/publications/microgenplan?lang=en>

One example of micro-generation within the HE sector is Warwick University, which has installed small wind turbines¹⁴² and in the future different operating units are expected to contribute to the generation of the power they use. It is not impossible to envisage a future in which any new proposal to introduce substantial amounts of new equipment such as ICT will be expected to be accompanied by a parallel proposal to demonstrate where the electricity is to come from.

Such considerations are likely to feature in forthcoming debates over the sector's move toward shared services. There is already interest in the creation of a cross-sector data centre infrastructure, with JANET currently developing a sector-wide strategy for data centres that will look at bringing its shared services model to the provision of highly efficient, scalable solutions. At a JISC strategy workshop in April 2009 there was a proposition for such a shared national data centre infrastructure to be located near renewal energy sources such as tidal power (Backhouse, 2009). How much further could this idea be taken? In Section 4 we discussed the use of DC power for distribution around the data centre – but this assumes an AC supply. What if the supply was already DC? Discussions over the local provision of electricity lead into the debate about the form of power being generated for use by ICT.

5.2.2 DC power

Weizsacker *et al.* (1997), well-known energy efficiency engineers from the Rocky Mountain Institute, estimate that switching, in general, from AC power input to DC would result in between six and ten times increased efficiency (based on the US power system). They also cite research from Germany showing that using photovoltaic (PV) solar panels to generate sufficient AC power for a family of four would require an area of panels of at least 30 m². The problem is that PV generates DC electricity which has to be (inefficiently) converted to AC in order to power today's appliances and computers. If household appliances were all converted to DC then the panels needed by the aforementioned family home would only need to be some 8 m².

Of course, AC power comes into its own when electricity needs to be transported over long distances,¹⁴³ but in a localised, renewable micro-generation set-up such as that which might be feasible in some HE and FE institutions, the logic of moving to a DC power supply becomes more obvious. It should be noted, when looking to the future, that the potential for rewiring domestic dwellings for DC is being considered by some researchers and energy commentators (Friedeman *et al.*, 2002; Watson *et al.*, 2006) and this is likely to extend to office and institutional buildings. Other work worthy of note is that of Tom DeFanti and Greg Hindley of the GreenLight project at the California Institute for Telecommunications and Information Technology, who are researching the applicability of DC power for data centre and computational use.¹⁴⁴

5.2.3 Hydrogen and fuel cells

When considering renewable micro-generation over the long term one should not forget the potential for hydrogen, an energy source that has been widely tipped as having a long-term future because the waste emission produced in the process of releasing the energy is simply water. There is talk of a 'hydrogen economy', particularly in transport (Stern, 2006), although there are many

¹⁴² See http://www2.warwick.ac.uk/newsandevents/news/wind_energy_at/

¹⁴³ Although it should be noted that high voltage DC (HVDC, in hundreds of kV) is being explored as an option for long-distance transmission of electric power, particularly for the exchange of power at national boundaries and under water. Its proponents argue that at such high voltages, power loss is lower than with AC and that a very wide, pan-European 'supergrid' of HVDC could be used to smooth out the intermittent nature of solar and wind farms.

¹⁴⁴ See <http://www.docstoc.com/docs/11258670/Project-GreenLight---DC-Modular-Direct-Current-Datacenters-for-Improved-Energy-Efficiency>

practical and technological barriers to overcome before it can be used on a large scale.¹⁴⁵ It is particularly relevant to this discussion as it is a potential source of energy for fuel cells, which have a role in the provision of back-up power to data centres.

At the moment hydrogen is not a low carbon fuel source, as it must be produced in the first place from an industrial process that requires a fossil fuel input. Either a pure source of hydrogen is generated industrially in advance of its use and then provided to the fuel cell in the form of a storage tank attached to the device, or it is produced (or 'reformed') within the device from another hydrocarbon-rich fuel such as an alcohol like ethanol.¹⁴⁶ In the future it is expected that the hydrogen will be generated as required from electricity and water, where the electricity is from a local, renewable source. One proposed idea is that of 'solar-hydrogen': hydrogen produced from water and solar energy (Nowotny, 2005).¹⁴⁷

5.2.4 Energy harvesting for pervasive computing

Over the time frame that we are considering, computing will become more 'pervasive' as ever smaller computational devices become more embedded into our urban fabric, our physical workplaces and even the clothes we wear. We are likely to see far more sophisticated mobile devices, hand-held systems and new physical interfaces to information resources. Work on energy in the world of wireless sensor networks and RFID may have long-term implications for such pervasive computer-related devices and, indeed, for ICT in general. Wireless sensors are expected to operate in the physical environment (or perhaps embedded in the body), relatively far from any power supply, and thus there has been a great deal of research into long-life and miniature batteries for these devices.

Recently this work has started to look at ways of drawing electric power from the local environment (Mingo, 2008), and has resulted in techniques sometimes referred to as 'energy harvesting' (Weddell *et al.*, 2008). This could be from thermal heat in the immediate surrounds, small amounts of energy from artificial light sources, vibrations created as a device is moved, or from radio-frequency energy.

This technology differs from micro-generation in that it works on a very local scale. One example illustrates this – making use of energy from the immediate surroundings through what are known as heat conversion chips (Bio Intelligence Service, 2008, p. 52). It is early days in this research work, but these chips make use of a physical process known as thermionic energy conversion to produce electricity from heat. It is envisaged that such chips could be located alongside a CPU in order to convert the heat produced by the CPU into electricity, which would then be either reused by the processor or used to help with cooling via a fan.¹⁴⁸

Energy harvesting techniques like these are currently only sufficient to power tiny wireless sensors and RFID devices, which have very low energy requirements (1 V with a power density of 1 $\mu\text{W}/\text{mm}^3$). As materials technology improves and miniaturisation increases over the next 10+ years, we are likely to see some of these developments influencing the development of the components in pervasive computing and perhaps in more traditional systems such as the desktop

¹⁴⁵ See <http://www3.imperial.ac.uk/icept/ourresearchactivities/hydrogenandfuelcell> and the International Partnership for the Hydrogen Economy: <http://www.iphe.net/>

¹⁴⁶ A detailed discussion of the mechanism and types of fuel cell can be found at the Rocky Mountain Institute's website <http://www.rmi.org/sitepages/pid184.php> and at the Fuel Cell Today portal <http://www.fuelcelltoday.com/reference/education-kit>

¹⁴⁷ The London Hydrogen Partnership, which is supported by the LDA and the Mayor of London, provides further information on their website: <http://www.london.gov.uk/lhp/index.jsp>

¹⁴⁸ The EU cite Eneco as one company working in this area. See http://pesn.com/Radio/Free_Energy_Now/shows/2007/01/06/9700221_Eneco_thermal_electric/#NIST_Report

PC. Research in this area – known as Zero Power ICT – forms part of the EU's Future and Emerging Technology (FET) agenda.

5.2.5 The emerging science of power management

Computer science has undertaken a great deal of theoretical work over the decades to analyse the complexity and performance of computation. The focus tends to be on formally working out how efficient algorithms are and what the mathematical bounds to problems might be (for example, readers who studied computing formally may remember lectures about the NP-complete problem). However, very little of this work takes any account of energy usage. As the authors of a recent paper in this area note: "Just as the formal notations of computational complexity have immensely benefited hardware and software design, a formal treatment of energy power and thermal issues should also lead to significant advances in how such systems are designed in the future" (Kant, 2009, p. 99). The current understanding of the relationship between computation, algorithms and hardware performance is rudimentary.

The solution is research work into an abstract model of computational machinery that takes account of energy and power management. This is the emerging theoretical field of the science of power management (SciPM) and researchers involved in this work envisage a situation where such modelling can lead to energy-aware computer operation that will make use of load balancing, processor speed, sleep states and other parameters that can be tuned to allocate different resources within a system (Cameron *et al.*, 2009).¹⁴⁹

5.3 Adapting to climate change

We need to avoid making the problem worse, so cutting carbon emissions is a priority. But all of us – individuals, businesses, Government and public authorities – will also need to adapt our behaviour to respond to the challenges of climate change.

DEFRA website, 24 July 2009¹⁵⁰

The Stern Review argued that the economic costs of adapting to climate change (e.g. making buildings more resilient to storm damage, improving flood defences) would be higher than the costs of avoiding climate change and, moreover, the longer any such efforts were delayed, the more expensive it would be. This logic has resulted in a primary focus on *mitigating* climate change, so as to reduce its overall impact. However, as global temperatures rise, they will bring changes in weather patterns, rising sea levels and increased frequency and intensity of extreme weather events such as storms, floods, droughts and heatwaves. There are likely to be profound effects on food supply, health, industry, transport and local ecosystems. Even if we were able to convert to a low carbon economy overnight there would still be impacts from the climate change we've already incurred, and some adaptation will be needed.

This two-pronged 'mitigate and adapt' policy is at the heart of the EU's strategy for dealing with climate change and in April 2009 the European Commission presented a framework for adaptation measures and policies to reduce the EU's vulnerability to the impacts of climate change (European Commission, 2009). Although the EU has a major role in coordinating action, developing models, spreading best practice and collecting information (through a Clearing House), it is envisaged that individual member states will take up the bulk of the work.

¹⁴⁹ The NSF recently sponsored a workshop on the SciPM in Arlington, Virginia; see <http://scipm.cs.vt.edu/>

¹⁵⁰ See <http://www.defra.gov.uk/environment/climate/index.htm>

In the UK, government departments and devolved authorities such as the Scottish government are starting to roll out plans for adaptation in areas such as planning, law and transport and an advisory service, the UK Climate Impacts Programme (UKCIP), has been created. It is highly likely that plans and strategies for adaptation will begin to filter down to individual agencies and public sector institutions in the near future. DEFRA argues that adaptation needs to be built in to the normal planning and risk management processes of organisations.

How does this relate to the future role for ICT in universities and colleges? Increasingly we are seeing ICT as fundamental infrastructure that underpins much of the day-to-day work of an institution. As such its role is being increasingly included in high-level strategic planning through techniques such as 'LEAN' thinking and EA. It is early days, but as these high-level plans will soon be incorporating adaptation planning it is clear that ICT has a role to play. Indeed, a barrier to adaptation, according to UKCIP, is the short-term nature of many planning horizons,¹⁵¹ something that was echoed in a recent JISC-funded report on strategic management of information technology.¹⁵²

More prosaic measures include:

- ICT will have a role to play as a communication tool for exchanging data, best practice, etc. about adaptation.¹⁵³
- There is considerable scope for incorporation of ICT into our built environment, e.g. heating control and monitoring through building sensors.¹⁵⁴
- Training programmes for facilities managers, caretakers, ICT technicians and technical support staff to improve day-to-day control over buildings.¹⁵⁵
- Training for information managers in business change and work practices that could reduce carbon footprint.

But perhaps the key role that ICT will play in adaptation has yet to be really defined. Predicting the future accurately is notoriously difficult and this is especially true for low carbon computing. Due to the urgency of climate change there is a huge level of interest in reducing energy across all sectors of the economy, with money and research time being poured into developing energy-efficient technologies and researching new energy sources. This is an extremely fluid area and we can expect great strides to be made in the next few years.

As the International Institute of Sustainable Development notes about information technologies: "They are transformative technologies because they put intelligence at the edges of networks, thereby maximising users' capacity to create and adapt. They are transformative technologies because they enable this creativity to be widely shared at every level, from local to global" (MacLean, 2008, p. 1). There is clearly a role here for the 'think communities' that are universities and colleges.

¹⁵¹ See http://www.ukcip.org.uk/index.php?id=56&option=com_content&task=view

¹⁵² See http://www.jisc.ac.uk/whatwedo/programmes/programme_jos/project_lfhe.aspx

¹⁵³ For example, the WeAdapt wiki (www.weadapt.com) provides an informal collaborative environment for organisations involved in adaptation.

¹⁵⁴ The Foresight report, *Powering our lives*

(<http://www.foresight.gov.uk/OurWork/ActiveProjects/SustainableEnergy/ProjectHome.asp>) provides horizon scanning material on this subject.

¹⁵⁵ For example, The Nottingham Declaration, a local government initiative for dealing with climate change that recently published advice for schools: [http://www.energysavingtrust.org.uk/nottingham/Nottingham-Declaration/Events-resources/News-Items/Consultation-on-carbon-management-for-schools/\(energysavingtrust\)/261311](http://www.energysavingtrust.org.uk/nottingham/Nottingham-Declaration/Events-resources/News-Items/Consultation-on-carbon-management-for-schools/(energysavingtrust)/261311)

6. Conclusions and recommendations

As carbon budgeting is rolled out across the public sector, individual departmental managers and technical staff will be expected to work much more closely with information services, procurement teams and estates to help build a lower carbon computing infrastructure. This report will therefore be read by different readers with varying levels of technical knowledge and the conclusions and recommendations we make here should provide a more general overview of the key points. However, it should be noted that this is not comprehensive and is in no way a substitute for reading the full report.

In brief, Section 4 reviewed the main technologies and standards that ICT managers in HE and FE are likely to consider when preparing to reduce institutional energy consumption with respect to servers and data centres, networks, PCs and monitors, printers, and data storage. It covered some of the short-term solutions that could be implemented immediately and which have few or no technological or managerial barriers, and it also reviewed some of the technological developments that are likely to come to fruition over the next few years. The short-term solutions are largely common sense, although their implementation may cause inconvenience to already overburdened technical staff. The latter developments represent technologies that are currently at an experimental or prototype level, or are likely to be too expensive for widespread public sector adoption. Based on this there are certain observable trends:

- Information services, estates and procurement will need to work together under a new regime – one in which energy efficiency has become paramount. Electricity costs will begin to figure as part of the institutional computing TCO. This will have an impact on how carbon accounting is handled and at what level decisions will be made.
- Historically, it has been difficult to certify and label computer-related products with information about their environmental impact. Despite their limitations, there are a number of prominent labelling schemes that are gaining traction and an area that has seen a particular concentration of activity is data centre metrics.
- There has been a rapid growth in data centre requirements, with vastly increasing amounts of data processing taking place, leading to higher densities of equipment and a correlating growth in the amount of energy consumed. This increases the level of heat generated within data centre racks and more traditional methods of cooling are coming under scrutiny as they are seen as energy inefficient.
- There is a trend towards a more centralised data centre infrastructure, both within institutions, and, more recently, between institutions through the shared services agenda and the rise of 'cloud' computing.
- There is likely to be a debate over coming months about the practicality of moving towards the widespread use of thin clients as a way of reducing energy consumption for information services. This will be complicated by a number of factors including: the increasing energy efficiency of new PC technology ('fat clients' and 'slim clients'); the realities of moving to either 'shared' computing or 'cloud' computing (see Section 5).
- Energy efficiency will become an increasingly important characteristic in future releases of operating systems and applications. Avoiding 'bloat'-ware and enforced upgrades will become increasingly important and the opportunities presented by open source software, in particular, should not be overlooked in this respect. There may also be opportunities to make use of

'LEAN' concepts and modular services to reduce 'waste' produced by the overlap of functionality across different applications.¹⁵⁶

- Display monitor technology is about to experience a step change, similar to the change from CRT to LCD, with the widespread, affordable introduction of OLED technology. Display technologies is a fast-moving area and for detailed information readers are referred to the JISC TechWatch report on Advanced Display Technologies (Anderson, 2006).
- Historically, very little has been done to optimise networking capacity for energy efficiency. A particular difficulty for IT managers who wish to consolidate their networks and refit with new equipment is that the provision of information on the energy efficiency of networking equipment is somewhat behind other sectors of the ICT industry. Network managers need to be prepared for a new generation of Ethernet equipment which takes account of energy efficiency, particularly with the ratification of IEEE 802.3az in late 2010.
- Institutions are seeing exponential growth in data storage requirements and coping with the demands of this growth, as well as managing the requirement for energy reduction, will be a major challenge for IS managers over the coming years. In addition, data is increasingly being seen as a valuable resource and its strategic significance in terms of the competitiveness of individual HEIs is rising to the top of the institutional agenda. In the future, the management and storage of data is likely to become a matter for corporate governance through the adoption of techniques such as EA. Reducing the environmental impact caused by storage technologies is therefore likely to be caught up with bigger strategic questions about data and institutions that have already begun to address these issues are likely to garner a long-term advantage.
- Printing and photocopying use around 10% of HE/FE ICT-related energy consumption. A developing trend is the move towards MFDs that incorporate fax, copier and printer in one device and as long as these devices are used to consolidate a number of users on to one such device, rather than supplement existing devices, they can provide energy efficiencies.

In addition to these trends, we have looked at some of the key technologies and standards within an overall context of planning and action and have organised them into a structure that we call the Low Carbon ICT Roadmap.

6.1 Low Carbon ICT Roadmap

In Section 2.1.1 we looked at the EU's 'BAU' and 'Eco' ICT futures scenarios. The Eco scenario uses a number of assumptions about policy, the development of technology and the market for more energy-efficient solutions. The kinds of general assumptions that it makes include: increased information to procurers; adoption of an EU-wide 'green' procurement scheme; increased adoption and use of Energy Star and other metrics and increased finance and support for technology innovation. Although this is useful for the kind of macro-level modelling of ICT energy use that the EU is undertaking, it does not really provide much in the way of specific guidance for ICT managers and technical staff in HE and FE.

In addition, the UK has already agreed targets that go beyond the Eco scenario so it is less relevant, generally, to a UK audience. In its favour, the Eco scenario is ICT-specific, and this is something that is currently lacking in HE and more generally across the public sector, although some areas

¹⁵⁶ The work of the JISC Flexible Services Delivery programme may have bearing on these discussions. See <http://www.jisc.ac.uk/whatwedo/programmes/flexibleservicedelivery.aspx>

(e.g. the MoD) have said that they plan to create such a scenario (Cook, 2008). With these things in mind we propose what we call a Low Carbon ICT Roadmap. This is an attempt to flesh out some of the generalities that the EU's Eco scenario points towards, to add in additional ones, and to adjust the timescale in keeping with the CCA and the CRC. It is far from complete and is offered as a first attempt at devising a model for the education sector – more of a thought experiment than a fully developed model – and is focused almost entirely on in-use energy reduction in line with current UK targets.

	Data centres	End user devices	Storage	Networks	Printers
2009	Climate Change Act (CCA) comes into effect. Low Carbon Transition Plan published.				
	HEFCE consulting on HE's contribution to national carbon reduction plan and strategy for achieving this.				
	<ul style="list-style-type: none"> * Basic benchmark of existing data centre performance * Review with respect to EU Code of Conduct/BCS tools suite 	<ul style="list-style-type: none"> * Use device's power management software * Implement switch off policy for evenings and weekends * Disable screen savers 	<ul style="list-style-type: none"> * Undertake storage life cycle review * Consolidate individual user disks 	<ul style="list-style-type: none"> * Review equipment in light of EU Broadband Code of Conduct 	<ul style="list-style-type: none"> * Make use of 'quick' wins such as recycled paper, lower weight paper, black/white printing, etc.
2010	April: introduction of Carbon Reduction Commitment (CRC) in public sector with "larger institutions" expected to participate in a test of the carbon allowance market mechanism. In this phase, carbon allowances will not be controlled. In tandem, by spring, all government departments to have produced their own carbon reduction plan.				
	<ul style="list-style-type: none"> * Incorporate simple airflow 'fixes' such as sealing floor voids * Consolidate servers and/or purchase more efficient equipment (e.g. blades, multi-core) * Advanced power management * Decommission unused servers * Virtualise at server level * Explore use of voltage conditioning equipment * Explore potential for fresh air/liquid cooling * Track shared services/JANET shared data centre developments * Publication of JISC reports into academic uses of cloud computing (May) 	<ul style="list-style-type: none"> * Instigate centrally controlled power management * Buy Energy Star 5.0 rated PCs * Buy most efficient power supplies (>80% efficient) * Explore local advantages and disadvantages of thin clients, shared PCs, energy-efficient 'fat clients' and 'slim clients' 	<ul style="list-style-type: none"> * De-duplicate data * Virtualise storage disks * Instigate information systems architecture with a view to expanding this to an EA. Include analysis of virtualisation, SAN, NAS, MAID and SSD 	<ul style="list-style-type: none"> * Explore network consolidation, virtual LANs and speed rationalisation * Proxying Support for Sleep Modes (ECMA) and automatic port and network shutdown equipment start to appear 	<ul style="list-style-type: none"> * Develop and implement campus-wide print strategy to reduce overall printing levels * Consolidation of print devices/use of MFDs

Cross-sector planning and knowledge sharing: technology testing and data sharing; low carbon shared services; development of enterprise architectures to ensure that business processes and ICT are optimised for energy efficiency

2011	First sale of carbon allowances takes place in April. Participants will have to purchase allowances, at a fixed price, to cover their forecast emissions for 2011/12.				
	* Airflow modelling and redesign of data centres * Investigate operating equipment at higher temperatures than recommended/optimum (preparing for new equipment guidelines from EU in 2012)		* Use of hybrid disks (SSDs and hard disks) * Green RAID and similar technologies	* Purchase equipment based on new standard: for energy efficient Ethernet (IEEE 802.3az)	
2013	Total amount of carbon allowances to be capped under the CRC and allowances are auctioned to determine a market price. Second phase of carbon budgets begins.				
	* Fuel cells for UPS (hydrogen?) * Explore possible DC power re-equip * Possible introduction of 'adaptive' data centres	* OLED display monitors * Power-efficient operating systems * Many-core CPUs (with thermal load balancing) * Thermionic co-power * Zero-standby-dissipation technology	* Widespread use of dynamic power management	* Widespread use of centralised power management software	* Possible introduction of reusable paper
2018	All new public sector buildings to be 'zero carbon'. Third phase of carbon budgets begins.				
2020	Public sector to have cut greenhouse gas (GHG) emissions by 30% on 1999/2000 levels. 30% of all electricity to come from renewables.				
	* Zero or very low power ICT through energy harvesting becomes available for data centres * Data centres increasingly powered by renewable micro-generation	* Low power optical computing systems * Zero or very low power ICT through energy harvesting for end user devices * New display technologies such as field emission displays.			
2030	Low cost emission reductions from cuts in nitrous oxide and methane almost fully exploited. Bigger reductions to be made elsewhere.				
2050	>80% GHG reduction from 1990 levels. Virtually all electricity to come from a mix of renewable sources, nuclear or fossil fuels where emissions are captured and stored safely for the long term.				

Key

UK Low Carbon Transition Plan		European Environment Agency		HEFCE	
-------------------------------	--	-----------------------------	--	-------	--

6.2 Recommendations

6.2.1 Metrics and measurement recommendation

1. HEFCE has issued guidance to institutions as to the steps that need to be taken in order to respond to the needs of the CCA using the Carbon Trust's HE Carbon Management Programme. However, this is a fairly basic approach and institutions or individual departments interested in a more detailed process of carbon footprint calculation are referred to the work of BSI PAS 2050.

6.2.2 Data centre recommendations

2. Data centre managers should implement the latest best practice and, in particular, the EU code of conduct. It is strongly recommended that JISC and HEIs join in work with the BCS Data Centre Special Interest Group, possibly to learn how the BCS data centre simulator can help optimise their data centre infrastructure.¹⁵⁷
3. In the Roadmap, individual data centre managers are advised to undertake a technical survey on the use of voltage conditioning to assess its appropriateness for local operating conditions. JISC/HEFCE should consider funding activity that pools the data generated during these surveys.
4. More consideration needs to be given to the issues involved in a move towards shared services and third-party 'cloud' services. Connections need to be made with the emerging government plans for G-Cloud and further detailed work needs to be commissioned into these areas.
5. Investigative work should be funded into the practical realities of a wholly DC-supplied data centre built near a source of renewable energy (see Section 5).

6.2.3 Network recommendation

6. The IEEE 802.3az working group plans to deliver a published standard by September 2010. The main participants in the standardisation process are large semiconductor companies but there is an important role to be played in this process by universities. JISC should consider expanding its work on standards to incorporate 802.3az.

6.2.4 End user devices recommendation

7. There is a lack of authoritative advice on best practice with respect to the refresh cycle of PCs and monitors. Current guidance varies depending on what type of PC is used in the model, equipment usage patterns, whether a number of machines are modelled, etc. JISC should consider developing sector-specific models that can be used in different scenarios and these should include investigating the feasibility and effect of extending refresh cycles by upgrading components.

6.2.5 Storage recommendations

8. Data on the performance of storage technology tends to be anecdotal and varies depending on factors such as interoperability. Reliable data on the environmental impact of data policies and life cycle and storage technologies will become increasingly valuable as constraints and even legislative controls on the environmental impact of data growth come into force. As a sector-wide organisation, JISC has the oversight to collect and facilitate the generation of

¹⁵⁷ For further details see BCS: DCSG Simulator Software Project <http://dcs.g.bcs.org//content/view/45/59/>

such data and consideration should be given to incorporating such work into programmes such as EA and Flexible Services Delivery.

9. There is currently a dearth of clear information on the energy profile of SSDs. JISC should facilitate the generation and sharing of information about performance and energy use, for example, when HEIs buy equipment which incorporates SSD (e.g. newer blade servers) and use this to provide a mechanism for comparing and standardising measurement.
10. The use of MAID for archival and repository activities should be explored further in a HE/FE context.

6.2.6 Printing recommendation

11. Institutions should develop and implement a print strategy with a view to reducing the amount of unnecessary printing. This is a long-term process involving considerable amounts of behaviour change rather than technological fixes. Office printers can be consolidated to larger MFDs and this should be considered as part of the wider print strategy.

List of abbreviations

BCS	British Computer Society
CCA	Climate Change Act
CPU	central processing unit
CRC	Carbon Reduction Commitment
ECR	Energy Consumption Rating
EA	Enterprise Architecture
EEA	European Environment Agency
EEE	energy efficient Ethernet
EPA	Environmental Protection Agency (USA)
GHG	greenhouse gas
HEFCE	Higher Education Funding Council for England
HEI	higher education institution
HPC	high performance computing
LAN	local area network
UPS	uninterruptible power supply
VM	virtual machine

About the Authors

Paul Anderson

Paul is a technology writer and is the Technical Editor of JISC TechWatch. A graduate of Leeds University (computer science), he has over twenty years' experience of working in computing in both academia and industry. In 2007 he was awarded the EPSRC's Computer Science Writer of the Year prize (professional category).

During his lunch break he blogs about technology at <http://techlun.ch> and can be contacted by email at: p.anderson [AT] intelligentcontent.co.uk.

Gaynor Backhouse

As the Project Manager of JISC TechWatch, Gaynor commissions and develops reports on the future of technology. She also experiments with technology and in 2002 she won an innovation award for the design of a new electronic publishing system. Along with Paul she co-authored the Green Party's general election manifesto in 2005.

Daniel Curtis

Daniel is a senior researcher in energy conservation at the University of Oxford's Environmental Change Institute. He is also the Evaluation Lead on the JISC-funded Low Carbon ICT project, which aims to reduce the university's carbon footprint resulting from the operation of its ICT infrastructure. His background is in energy conservation, especially in the fields of efficiency and renewable energy technologies. Over a varied career he has worked on a range of energy projects in several countries and with many organizations, including the Centre for Alternative Technology in Wales, and the International Energy Agency in Paris.

Contact: <http://www.eci.ox.ac.uk/people/curtisdaniel.php>

Simon Redding

Simon is an innovation consultant for the Environment Agency, which he represents on government and international panels relating to sustainable ICT.

David Wallom

David is the Technical manager of the Oxford University e-Research Centre, which involves him in a wide range of technology infrastructure developments including the development of a university-wide computational grid. He is also the technical architect on the JISC-funded Low Carbon ICT project, which aims to promote the efficient use of existing ICT resources. David is also currently chair of the UK e-Science Engineering Task Force.

Contact: <http://www.oerc.ox.ac.uk/people/david-wallom>

References

- Agarwal, Y., Hodges, S., Chandra, R., Scott, J., Bahl, P., Gupta, R. 2009. *Somniloquy: Augmenting Network Interfaces to Reduce PC Energy Usage*. 6th USENIX Symposium on Networked Systems Design and Implementation, 22–24 April 2009, Boston, Massachusetts. Available online at: http://www.usenix.org/event/nsdi09/tech/full_papers/agarwal/agarwal_html/index.html
- Anderson, P. 2006. **Advanced Display Technologies**. JISC Technology and Standards Watch. JISC: Bristol, UK. Available at: <http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/horizonscanning/hs0503.aspx>
- Anderson, P. 2007. Personal notes from a lecture given by Dr Andrew Herbert, Director of Microsoft Research, Cambridge UK, given at Leeds University Computing Dept's 50th anniversary celebrations. See: <http://techlun.ch/2007/04/04/what-cost-software-innovation/>
- Backhouse, G. 2009. JISC TechWatch strategy position paper (unpublished) presented to JISC strategy workshop, April 2009.
- Bathula, B., Alresheedi, M., Elmirghani, J. 2009. **Energy Efficient Architectures for Optical Networks**. Presentation to the London Communications Symposium, 3 September 2009. Available at: http://www.ee.ucl.ac.uk/lcs/previous/LCS2009/LCS/lcs09_33.pdf
- Belaramani, N., Dahlin, M. 2009. **Achieving Energy Efficiency while Synchronizing Personal Data**. Technical Report TR-09-21, The University of Texas at Austin, August 2009. Available at: <ftp://128.83.120.155/pub/techreports/tr09-21.pdf>
- Berl, A., Gelenbe, E., di Girolamo, M., Giuliani, G., de Meer, H., Dang, M. Q., Pentikousis, K. 2009. *Energy-Efficient Cloud Computing*. **The Computer Journal**. Oxford University Press/British Computer Society: UK. Advance access copy available at: <http://comjnl.oxfordjournals.org/cgi/content/abstract/bxp080>
- Bertoldi, P. **European Code of Conduct for Broadband Equipment**. See: http://www.itu.int/dms_pub/itu-t/oth/09/05/T09050000010004PDFE.pdf
- Binstock, A. 2008. Solid State Drives – How Green Are They for IT? Greener Computing blog, 5 Aug 2008. Available at: <http://www.greenercomputing.com/blog/2008/08/05/solid-state-drives-how-green-are-they-it>
- Bio Intelligence Service. 2008. **Impacts of Information and Communication Technologies on Energy Efficiency**. Executive summary of the final report to the European Commission DG INFSO (tender no. CPP 16A-2007/2007/S 68-082361), September 2008. Available at: http://cordis.europa.eu/fp7/ict/sustainable-growth/studies_en.html
- BIS/DCMS [Department for Business, Innovation and Skills and Department for Culture Media and Sport]. 2009. **Digital Britain**. Final report CM7650, June 2009. TSO: UK. Available at: http://www.culture.gov.uk/what_we_do/broadcasting/6216.aspx
- Blackburn, J., Christensen K. 2009. **A Simulation Study of a New Green BitTorrent**. First International Workshop on Green Communications (GreenComm'09), 18 June 2009, Dresden, Germany. Available at: <http://www.csee.usf.edu/~christen/class3/greenComm09a.pdf>
- Bolla, R., Bruschi, R., and Ranieri, A. 2009. *Energy-aware equipment for next-generation networks*. In: Proceedings of the 4th international Conference on Future Internet Technologies (CFI

'09), 17–19 June 2009, Seoul, Korea. ACM: New York, NY. DOI=
<http://doi.acm.org/10.1145/1555697.1555707>

Borkar, S. 2003. *Getting Gigascale Chips: Challenges and Opportunities in Continuing Moore's Law*. **Queue magazine**. October 1, 2003. Available at: <http://queue.acm.org/detail.cfm?id=957757>

Branscombe, M. 2008. *How Google is redesigning your datacenter*. **IT Pro** website, 17th April 2008. Available at: <http://www.itpro.co.uk/features/188775/how-google-is-redesigning-your-data-centre.html> (last accessed 29/9/08)

Brill, K. 2008. *Findings on Data Center Energy Consumption Growth May Already Exceed EPA's Prediction Through 2010!* Vodcast, **Uptime Institute website**. Available at: <http://uptimeinstitute.org/content/view/155/147/>

BSI. 2008. **PAS 2050:2008, Specification for the assessment of the life cycle greenhouse gas emissions of goods and services**. Available at: <http://www.bsigroup.com/en/Standards-and-Publications/Industry-Sectors/Energy/PAS-2050/>

BT, 2009. *Green IT comes of age: data centres that use 60% less power*. BT case study, 21 CN. Available at: http://globalservices.bt.com/InsightsDetailContentAction.do?Record=efficient_data_centres_article_all_en-gb

Cabinet Office. 2008. **Greening Government ICT: Efficient, Sustainable, Responsible**. Available at: http://www.cabinetoffice.gov.uk/media/66177/greening_government_ict.pdf

Calwell, C. 2005. **Results of Field and Lab Testing of PCs**. Presentation at Energy Star stakeholder meeting, 21 June 2005. Available at: http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/computer/ResultsTestingPCs.pdf

Cameron, K., Pruhs, K., Irani, S., Ranganathan, P., Brooks, D. 2009. Report of the Science of Power Management Workshop. Final report (version 31st Aug 2009) of the Workshop on the Science of Power Management (SciPM '09), April 9–10 2009, Arlington, VA, USA. Available at: <http://scipm.cs.vt.edu/SciPM-ReportToNSF-Web.pdf>

Canon 2006. *On the tenth day of Christmas... ..UK offices waste enough energy to roast 4.4m turkeys*. Press Release, 13 Dec 2006. [no longer available online]

Carbon Trust. 2007. **Carbon Footprinting: an introduction for organisations**. Available at (requires registration): <http://www.carbontrust.co.uk/publications/publicationdetail?productid=CTV033>

Casad, J. 2008. *Green IT*. **Linux Magazine**, Oct 2008. Linux New Media Ltd: Manchester, England.

CAT. 2009. **Zero Carbon Britain: an alternative strategy**. Available at: <http://www.zerocarbonbritain.com/>

Ceuppens, L., Sardella, A., Kharitonov, D. 2008. *Power Saving Strategies and Technologies in Network Equipment Opportunities and Challenges, Risk and Rewards*. **International Symposium**

on Applications and the Internet (SAINT '08), 28 July–1 Aug 2008, Turku, Finland.

<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4604616&isnumber=4604514>

Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., Wright, S., 2008. *Power Awareness in Network Design and Routing*. **Proceedings of IEEE INFOCOM 2008**, the 27th Conference on Computer Communications, 13–18 April 2008, Phoenix, Arizona, USA. IEEE Communications Society (ComSoc): USA, pp.457–465.

Chen, A. 2006. *Energy-Efficient Direct-Current Powering Technology*. Lawrence Berkeley National Laboratory Environmental Energy Technologies Division Newsletter. Vol. 7, No. 1, Fall 2006. Available at: <http://eetd.lbl.gov/newsletter/nl24/pdf/eetd-nl24.pdf>

Chen, A. 2008. *A Vision for Energy Efficiency at Data Centers* (part 3). **Conservation Update**, September/October 2008, US Dept of Energy. Available at:

http://apps1.eere.energy.gov/state_energy_program/update/feature_detail.cfm/fid=85/start=3

Chidlow, S. 2003. **Storage Area Networks**. JISC Technology and Standards Watch, November 2003. JISC: Bristol, UK. Available at:

<http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/horizonsscanning/hs0307.aspx>

Choudhuri, S. and Givargis, T. 2007. *Performance improvement of block based NAND flash translation layer*. In: **Proceedings of the 5th IEEE/ACM international Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS '07)**, 30 Sept–3 Oct 2007, Salzburg, Austria. ACM: New York, NY, pp.257–262. DOI=

<http://doi.acm.org/10.1145/1289816.1289878>

Climate Group. 2008. **Smart 2020: enabling the low carbon economy in the information age**.

Global e-Sustainability Initiative: Brussels, Belgium. Available at: <http://www.smart2020.org/>

Colarelli, D., Grunwald, D. 2002. *Massive arrays of idle disks for storage archives*. In:

Proceedings of the 2002 ACM/IEEE Conference on Supercomputing. IEEE Computer Society Press: California, USA.

Cook, T. 2008. UK MoD Climate Change Strategy. Presentation to the Institute for Governance and Sustainable Governance, Importance of Military Organisations in protecting the climate conference 2008, 3–5 November, Paris, France. Available at:

<http://www.igsd.org/conferences/Paris2008/2d%20Cook.pdf>

Counsell T.A.M., Allwood J.M. 2006. *Desktop paper recycling: A survey of novel technologies that might recycle office paper within the office*. **Journal of Materials Processing Technology**, Vol. 173, Iss. 1 (March 2006), pp.111-123. Available at: doi:10.1016/j.jmatprotec.2005.11.017.

Courtney, M. 2008. *Data centres running out of electricity*. **Computing.co.uk**. 26 March 2008.

Available at: <http://www.computing.co.uk/computing/news/2212745/centres-running-dry>

DECC (Department of Energy and Climate Change). 2009. **The UK Low Carbon Transition Plan**, 20 July 2009. The Stationery Office: UK. Available online at:

http://www.decc.gov.uk/en/content/cms/publications/lc_trans_plan/lc_trans_plan.aspx

DEFRA. 2007. **Climate Change Strategic Framework**. Available at:

http://www.doeni.gov.uk/climate_change_bill_strategy_document.pdf

- Deng, Y., Wang, F., Helian, N. 2008. *EED: Energy Efficient Disk drive architecture*. **Information Sciences**, Vol. 178, Iss. 22 (15 November 2008), pp.4403–4417.
- Doyle, E. 2008. *The Green Jungle*. **Computer Weekly**. 4 Dec 2007.
- Dubash, M. 2006. *Is DC power what every data centre needs?* **ComputerWorldUK.com**, 12th October 2006. Available at: <http://www.computerworlduk.com/technology/servers-data-centre/infrastructure-management/in-depth/index.cfm?articleid=248>
- Dubash, M. 2007. *APC shows off fuel cells for the data center*. **Network World**. 6th June 2007. Available at: <http://www.networkworld.com/news/2007/060607-apc-shows-off-fuel-cells.html>
- Dunlap, K., Rasmussen, N. 2006. **The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers**. APC White Paper, No. 130. See: <http://www.datacentres.com/papers/papers/apcWP130.pdf>
- Ecos and EPRI. 2008. **Server Research Project: Efficient Power Supplies for Data Centre and Enterprise Servers**. Final Report, Feb 2008. 80 Plus/Ecos: USA. Available at: <http://www.80plus.org/research.htm>
- EEA (European Environment Agency). 2005. **Climate Change and a European low-carbon energy system**. EEA Report No 1/2005. Office for Official Publications of the European Communities: Luxemburg. Available online at: http://www.eea.europa.eu/publications/eea_report_2005_1
- Energy Efficient Ethernet Task Force. 2009. **Task Force document repository** (8th June 2009 update). Available at: <http://www.ieee802.org/3/az/>
- EPA. 2007. **Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431**. 2nd August 2007. Available at: http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf
- Europa. 2007. **Limiting Global Climate Change to 2 degrees Celsius**. European Commission press release, 10 January 2007 (MEMO/07/16). See: http://ec.europa.eu/environment/climat/future_action.htm
- Europa. 2009. *Energy Star: Commission and the US reinforce the energy efficiency standards for office equipment*. European Commission press release, 17 June, 2009. Available at: <http://europa.eu/rapid/pressReleasesAction.do?reference=IP/09/943&format=HTML>
- European Commission. 2009. **Adapting to climate change: Towards a European framework for action**. White paper COM(2009) 147 final. Available at: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2009:0147:FIN:EN:PDF>
- Fan, X., Weber, W., Barroso, L.A. 2007. *Power provisioning for a warehouse-sized computer*. In: **Proceedings of the 34th Annual international Symposium on Computer Architecture (ISCA '07)**, 9–13 June 2007, San Diego, California, USA. ACM: New York, NY. DOI=<http://doi.acm.org/10.1145/1250662.1250665>

Fontecchio, M. 2006. *Data center power backup isn't just about batteries anymore.*

SearchDataCentre.com website, 6th July 2006. Available at:

http://searchdatacenter.techtarget.com/tip/0,289483,sid80_gci1197541,00.html

Forge, S. 2007. *Powering down: remedies for unsustainable ICT.* **Foresight**. Vol. 9, Iss. 4. Emerald Group Publishing Ltd. DOI: 10.1108/14636680710773795

Francis, K., Richardson, P. 2009. *Green Maturity Model for Virtualization.* **Microsoft**

Architecture Journal. Jan 2009. Available at: [http://msdn.microsoft.com/en-](http://msdn.microsoft.com/en-us/library/dd393310.aspx)

[us/library/dd393310.aspx](http://msdn.microsoft.com/en-us/library/dd393310.aspx)

Freeh, V., Lowenthal, D., Pan, F., Kappiah, N., Springer, R., Rountree, B. 2007. *Analyzing the Energy-Time Tradeoff in High-Performance Computing Applications.* In: **IEEE Transactions on Parallel and Distributed Systems**. Vol. 8, Iss. 4 (June) 2007. IEEE Press: Piscataway, NJ, USA.

Also available from North Carolina State University website:

<http://www4.ncsu.edu/~vwfreeh/tpds.pdf>

Friedeman, M., Elisa C. Boelman, E., van Timmeren, A., Joop Schoonman, J. 2002. **Concept for a DC low voltage house.** Sustainable Building 2002 conference. Available at:

http://web.bk.tudelft.nl/bt/installaties/onderzoek/paper_sb02_217.pdf

Fujitsu 2005. *Failing to switch off computers could cost British business up to £123.2 million per year.* Press Release, 24 Oct 2005. [no longer available online]

Gartner. 2007. *Gartner Says Data Centres Account for 23 Per Cent of Global ICT CO2 Emissions.*

Press release, 11 October 2007. Available online at: <http://www.gartner.com/it/page.jsp?id=530912>

Geer, 2008. *Reducing the storage burden via data duplication.* **IEEE Computer**, Dec 2008. IEEE Computer Society: New York, NY, USA.

Giddens, A. 2009. **The Politics of climate change.** Polity: UK.

Graefe, G. 2008. *Database servers tailored to improve energy efficiency.* In: Apel, S., Rosenmüller, M., Saake, G., Spinczyk, O. (eds) **Proceedings of the 2008 EDBT Workshop on Software Engineering For Tailor-Made Data Management (SEDMDM '08)**, 29th March, Nantes, France.

ACM: New York, NY. Available at: <http://doi.acm.org/10.1145/1385486.1385494>

Grover, A. 2003. *Modern System Power Management.* **Queue magazine**. Vol. 1, Iss. 7 (October 2003). Also available at:

<http://delivery.acm.org/10.1145/960000/957774/grover.html?key1=957774&key2=1350226521&coll=GUIDE&dl=GUIDE&CFID=59061253&CFTOKEN=27941786>

Gruman, G. 2006. *Desktop virtualization: Making PCs manageable.* **InfoWorld.com**, 11th

September 2006. Available at:

http://www.infoworld.com/archives/emailPrint.jsp?R=printThis&A=/article/06/09/11/37FEvirtcase_desk_1.html

Gunaratne, G., Christensen, K., Nordman, B. 2005. *Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections and scaling of link speed.*

International Journal of Network Management, Vol. 15 Iss. 5, pp.297-310. John Wiley and Sons: USA.

- Gunn, C. 2006. *CMOS Photonics for High-Speed Interconnects*. **IEEE Micro**. March-April 2006. Available at: http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1624327
- Gupta, M., Singh, S., 2003. *Greening the Internet*. In: Feldmann, A., Zitterbart, M., Crowcroft, J., Wetherall, D. (Eds). 2003. **Proceedings of SIGCOMM 2003** Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, 25–29 August, 2003, Karlsruhe, Germany. ACM: Germany, pp.19–26.
- Gustafson, J. 2007. **Strategies for solving the heat/power problem**. Presentation to the Frontiers of Extreme Computing 2007/Zettaflops workshop, 21 – 25 October 2007, Santa Cruz, CA, USA. Available at: <http://www.zettaflops.org/fec07/index.html>
- Harbaugh, L. *Cut power costs with DC power*. **InfoWorld**, 14th January 2009. Available at: <http://www.infoworld.com/t/platforms/cut-power-costs-dc-power-479>
- Harizopoulos, S., Meza, J., Shah, M., Ranganathan, P. 2009. *Energy Efficiency: The New Holy Grail of Data Management Systems Research*. Fourth biennial Conference on Innovative Data Systems Research (CIDR '09), 4–7 January 2009, Asilomar, California, USA. Electronic proceedings available from: <http://www.cidrdb.org/cidr2009/program.html>
- Harnik, D., Naor, D., Segall, I. 2009. *Low power mode in cloud storage systems*. **International Parallel and Distributed Processing Symposium 2009 (IPDPS '09)**, Rome, Italy. IEEE Computer Society.
- Hea, J., Chaintreub, A., Diotb, D. 2009. *A performance evaluation of scalable live video streaming with nano data centers*. **Computer Networks**, Vol. 53, Iss. 2 (13 February 2009) pp.153-167.
- Held, G. 2003. 4th Edition. **Ethernet Networks: Design, Implementation, Operation and Management**. Wiley: Chichester, UK.
- Herrod, S. 2006. **The Future of Virtualisation Technology**. Presentation to 33rd International Symposium on Computer Architecture, 17–21 June 2006, Boston, MA, USA.
- Hinnells, M. 2007. *Enabling technologies for demand reduction and microgeneration in buildings*. **Foresight** website. Available at: http://www.foresight.gov.uk/Energy/Enabling_technologies.pdf
- Hölzle, U. 2005. *Viewpoint: The dark side of computing power*. **CERN Courier**. Available at: <http://cerncourier.com/cws/article/cern/29476>
- Hopper, A., Rice, A. 2008. *Computing for the Future of the Planet*. **Philosophical Transactions of the Royal Society**. A 366(1881):3685–3697, 2008. Available at: <http://rsta.royalsocietypublishing.org/>
- Hughes, D. 2007. **Data Center Design**. Presentation to the 21st Large Installation System Administration Conference, 11–16 November 2007, Dallas USA. Available online at: http://www.usenix.org/event/lisa07/htgr_files/hughes.pdf
- Intellect. 2008. **High Tech, Low Carbon: the role of technology in tackling climate change**. Intellect: London, UK. Available at: <http://www.intellectuk.org/hightechlowcarbon>
- IT Pro, 2008. *Virtualisation offers IT Departments the way to a greener path*. **IT Professional (News Briefs)** vol. 10, no. 1, Jan/Feb 2008, pp.7–11. Available at:

<http://www.computer.org/portal/web/csdl/abs/html/mags/it/2008/01/mit2008010007.htm>,
DOI:10.1109/MITP.2008.1

James, P., Hopkinson, L. 2009a. **Sustainable ICT in Further and Higher Education**. Final version 13 January 2009. JISC: Bristol, UK. Available online at:
<http://www.jisc.ac.uk/publications/documents/sustainableictfinalreport.aspx>

James, P., Hopkinson, L. 2009b. **Energy Efficient Data Centres in Further and Higher Education**. 27th May 2009. Available at: <http://www.susteit.org.uk/publications/index.php> (item 4)

James, P., Hopkinson, L. 2009c. **Energy and Environmental Impacts of Personal Computing: A Best Practice Review**. 27 May 2009. SusteIT: Bradford, UK. Available from:
<http://www.susteit.org.uk/publications/index.php>

James, P., Hopkinson, L. 2009d. **Energy Efficient Printing and Imaging in Further and Higher Education**. 13th Jan 2009. Available at: <http://www.susteit.org.uk/publications/index.php> (item 5)

Johansson, D. 1999. **The influence of eco-labelling on producers of personal computers**. Thesis, MSc Environmental Management and Policy. IIIIEE, Lund University: Sweden. Available at:
[http://www.iiiee.lu.se/Publication.nsf/e36f5f3f8fa200a6c1256b4200480181/74eb2507df88a130c1256c37002fff33/\\$FILE/comm2000_3.pdf](http://www.iiiee.lu.se/Publication.nsf/e36f5f3f8fa200a6c1256b4200480181/74eb2507df88a130c1256c37002fff33/$FILE/comm2000_3.pdf)

Johnson, E. P. 1998. *Developing European ecolabels for PCs*. In: **Proceedings of the 1998 IEEE International Symposium on Electronics and the Environment (ISEE '98)** 4-6 May, Illinois, USA. IEEE: USA, pp.9–13.

JRC Institute for Energy. 2008a. **Code of Conduct on Data Centres Energy Efficiency**. Version 1.0, 30th October 2008. Institute for Energy, Renewable Energy Unit: Ispra, Italy. Available at:
http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative_data_centers.htm

JRC Institute for Energy. 2008b. **Code of Conduct on Energy Consumption of Broadband Communication Equipment**. Version 3, 18th November 2008. Institute for Energy, Renewable Energy Unit: Ispra, Italy. Available at:
http://re.jrc.ec.europa.eu/energyefficiency/html/standby_initiative_broadband%20communication.htm

Kant, K. 2009. *Towards a Science of Power Management*. **IEEE Computer**, Vol. 42, No. 9 (Sept 2009), pp. 99–101. Available at:
<http://www.computer.org/portal/web/csdl/doi/10.1109/MC.2009.303>

Kimura, M. 2009. *Zero-Standby-Power ICs to Enter Practical Use*. **Tech On!** 20 July 2009. Available at: <http://techon.nikkeibp.co.jp/article/HONSHI/20090629/172348/?P=1>

Knight, W. 2009. *Energy-Aware Internet Routing*. **MIT Review**. Aug 2009. Available at:
<http://www.technologyreview.com/business/23248/page1/>

Kothiyal, R., Tarasov, V., Sehgal, P., and Zadok, E. 2009. *Energy and performance evaluation of lossless file data compression on server systems*. In: **Proceedings of the Israeli Experimental Systems Conference 2009 (SYSTOR 2009)**, 4–6 May 2009, Haifa, Israel. SYSTOR '09. ACM: New York, NY, pp.1–12. DOI= <http://doi.acm.org/10.1145/1534530.1534536>

Laoutaris, N., Rodriguez, P., Massoulie, L. 2008. *ECHOS: Edge Capacity Hosting Overlays of*

Nano Data Centres. Editorial note to ACM SIGCOMM **Computer Communication Review**, Vol. 38, No. 1 (Jan 2008).

Leventhal, A. 2008. *Flash Storage Today*. **ACM Queue magazine**, 1st July 2008. Available at: <http://queue.acm.org/issuedetail.cfm?issue=1413254> (item 3)

Li, D. and Wang, J. 2004. *EERAID: energy efficient redundant and inexpensive disk array*. In: **Proceedings of the 11th European Workshop on ACM SIGOPS**, 19–22 September 2004, Leuven, Belgium. ACM: New York, NY, 29. DOI= <http://doi.acm.org/10.1145/1133572.1133577>

Lidstrom, S. 2007. **Data Center High Voltage Connectors**. Presentation to the DC for Data Centers Workshop, 12 July 2007, Santa Clara, CA, USA. Available at: <http://hightech.lbl.gov/presentations/dc-powering/dc-stakeholders/4-Connectors.pdf>

Liu, L., Wang, H., Liu, X., Jin, X., He, W. B., Wang, Q. B., and Chen, Y. 2009. *GreenCloud: a new architecture for green data center*. In: **Proceedings of the 6th International Conference industry session on Autonomic Computing** and Communications Industry Session (ICAC-INDST '09), 15 June 2009, Barcelona, Spain. ACM: New York, NY, pp.29-38. DOI= <http://doi.acm.org/10.1145/1555312.1555319>

MacLean, D. 2008. **ICTs, Adaptation to Climate Change and Sustainable Development at the Edges**. International Institute for Sustainable Development (IISD) commentary document, August 2008. Available at: www.iisd.org/pdf/2008/com_ict_climate.pdf

Maloney, C., de las Heras Bueno, J., Grier I Fisa, M., Sanchez Jimenez, M. 2008. **Novel ICT solutions for Smart Electricity Distribution Networks**. Presentation, ICT 2008, 25–27 November 2008, Lyon, France. Available at: http://ec.europa.eu/information_society/events/cf/item-display.cfm?id=873

Mao, B., Feng, D., Jiang, H., Wu, S., Chen, J. 2008. *GRAID: A Green RAID Storage Architecture with Improved Energy Efficiency and Reliability*. In: **Proceedings of the IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems (MASCOTS '08)**, 8–10 Sept. 2008, Baltimore, MD, USA. IEEE Computer Society: New York, NY, USA. Available at: <http://ieeexplore.ieee.org/Xplore/login.jsp?url=http%3A%2F%2Fieeexplore.ieee.org%2Fstamp%2Fstamp.jsp%3Farnumber%3D4770574%26isnumber%3D4770548&authDecision=-203>

Masanet, E., Horvath, A. 2006. *Enterprise Strategies for Reducing the Life-Cycle Energy Use and Greenhouse Gas Emissions of Personal Computers*. **Proceedings of the IEEE International Symposium on Electronics and the Environment**, 8–11 May 2006, Scottsdale, AZ, USA, pp.21-26. DOI: 10.1109/ISEE.2006.1650025

Matthews, H., Matthews, D. 2003. *Information Technology Products and the Environment*. In: Kuehr, R. and Williams, E. (eds). 2003. **Computers and the Environment**. Kluwer Academic Publishers: Dordrecht, Netherlands.

Miliband, E. 2009. *One giant leap for a greener Britain*. **The Guardian, Comment is free** (invitation-only blog), 20th July 2009. Available at: <http://www.guardian.co.uk/commentisfree/2009/jul/20/climate-change-low-carbon-economy>

Miller, R. 2008. *Uptime: Energy Use Soars at Top Data Centers*. **DataCenterKnowledge.com**, 12th March 2008. Available at: <http://www.datacenterknowledge.com/archives/2008/03/12/uptime->

[energy-use-soars-at-top-data-centers/](#)

Mills, E. 2008. **Supercomputers: super polluters?** Lawrence Berkeley National Laboratory, Paper LBNL-359E. Available at: <http://repositories.cdlib.org/lbnl/LBNL-359E/>

Mingay, S. 2007. *Green IT – a new industry shockwave*. Gartner report, 7th Dec 2007. Summary available at: <http://mediaproducts.gartner.com/reprints/avaya/vol2/article10/article10.html> and presentation slides available at: http://www.ictliteracy.info/rtf.pdf/Gartner_on_Green_IT.pdf

Mingo, N. 2008. Towards zero power ICT. Presentation, ICT 2008, 25–27 November 2008, Lyon, France. Available at: ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/fet-proactive/20081125-27-lyon-mingo_en.pdf

Mitchell, R. 2005. *Chilling out with DC Power*. **Computerworld**, 9th May 2005. Available at: http://www.computerworld.com/s/article/101542/Chilling_Out_With_DC_Power

Monbiot, G. 2006. **Heat: how to stop the planet burning**. Allen Lane: London, UK.

Morgan, G. 2007. *Crossed Wires*. **Information Age**. 10th December 2007. Available at: <http://www.information-age.com/channels/data-centre-and-it-infrastructure/features/270271/crossed-wires.shtml>

Moshnyaga, V. 2009. *Untraditional approach to Computer Energy Reduction*. In: **Integrated Circuit and System Design**. Revised selected papers of the Power and Timing Modeling, Optimization and Simulation (PATMOS '09) Workshop. Springer: Berlin/Heidelberg. See: <http://www.springerlink.com/content/u7232755w176/?p=a6291f6617f14dc38249a876977594ce&pi=0>

Mueller, S. 1999. 11th Edition. *Chapter 21 Power Supply and Chassis/Case*. **Upgrading and repairing PCs**. QUE: Indiana, USA.

Nedevschi, S., Popa, L., Iannacone, G., Ratnasamy, S., Wetherall, D. 2008. *Reducing Network Energy Consumption via Sleeping and Rate-Adaptation*. In: **Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)**, 16–18 April, 2008, San Francisco, USA (online). Available at: <http://www.usenix.org/events/nsdi08/tech/nedevschi.html>

Newcombe, L. 2008a. **Data centre energy efficiency metrics**. British Computer Society: London, UK. Available at: <http://www.bcs.org/upload/pdf/data-centre-energy.pdf>

Newcombe, L. 2008b. *Energy Efficient Data Centres*. **Presentation** at University of Oxford ICT Forum, July 2008. Available at: <http://www.ictf.ox.ac.uk/conference/2008/programme.html#plenary2> (actually labelled as Low Carbon IT)

Newcombe, L. 2008c. **Data Centre Cooling**. October 2008. Available at: <http://www.susteit.org.uk/publications/index.php> (item 9)

Nordman, B. Christensen, K. 2007. **Improving the Energy Efficiency of Ethernet-Connected Devices: A proposal for proxying**. Industry white paper. Ethernet Alliance: Texas, USA. http://www.ethernetalliance.org/attachments/122_PROPOSAL_FOR_PROXYING_EDIT.PDF

- Nordman, B. Christensen, K. 2009. *Greener PCs for the Enterprise*. **IT Pro**, July/Aug 2009. IEEE: California, USA. Available online at:
<http://www2.computer.org/portal/web/computingnow/0809/whatsnew/itpro>
- Nordman, B., Meier, A., Pietter, M.A. 2000. **PC and Monitor Night Status: Power Management Enabling and Manual Turn-off**. US Dept of Energy. Available at:
<http://www.osti.gov/bridge/purl.cover.jsp;jsessionid=21D82FBD7216729DED7A9A1814B810CB?purl=/795948-j3RR18/native/>
- Nowotny, J. 2005. *Solar-hydrogen: Environmentally safe fuel for the future*. **International Journal of Hydrogen Energy**, Vol. 30, Iss. 5 (April 2005), pp.521-544.
- O'Connor, M. 2008. *Reuse and Recycle. But First, Virtualize*. **Greener Computing**, website, 19 Aug 2008. Available at: <http://www.greenercomputing.com/news/2008/08/19/reuse-and-recycle-first-virtualize>
- O'Donnell, S. 2007. *The 21st Century Data Centre*. Presentation, **Eco Responsibility in IT Conference**, 8th November 2007, London.
- OGC. 2007. **Open Source Software Trials in Government**. Final report. Office of Government Commerce. Available at:
<http://www.ogc.gov.uk/documents/CP0041OpenSourceSoftwareTrialReport.pdf>
- Otoo, E., Rotem, D., Tsao, S. 2009. *Energy Smart Management of Scientific Data*. In: **Proceedings of 21st International Conference Scientific and Statistical Database Management (SSDBM 2009)**, 2–4 June 2009, New Orleans, LA, USA. Springer: Berlin/Heidelberg, Germany, pp. 92–109. Available at: <http://www.springerlink.com/index/0545326q7r124218.pdf>
- Palmintier, B., Newman, S. 2008. **Systems thinking for radically efficient and profitable data centre design**. Presentation to the Next Generation Data Center Conference, 4th–7th August 2008, San Francisco, USA. Available at:
http://www.rmi.org/images/PDFs/Energy/Systems_thinking_for_radically_efficient_and_profitable_datacenters.zip (5th August)
- Pinheiro, E., Bianchini, R., and Dubnicki, C. 2006. *Exploiting redundancy to conserve energy in storage systems*. In: **Proceedings of the Joint international Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '06/Performance '06)**, 26 – 30 June 2006, Saint Malo, France. ACM: New York, NY pp.15-26. DOI= <http://doi.acm.org/10.1145/1140277.1140281>
- Poess, M., Nambiar, R. 2008. *Energy Cost, the key challenge of today's data centre: a power consumption analysis of TPC-C results*. In: **Proceedings of the 34th International Conference on Very Large DataBases 2008 (VLDB '08)**, 23–28 August 2008, Auckland, NZ. Available at:
https://www.cs.auckland.ac.nz/research/conferences/vldb08/index.php/VLDB_08
- Rivoire, S. 2008. **Models and Metrics for Energy-Efficient Computer Systems**. Ph.D. dissertation, Dept of Electrical Engineering, Stanford University, June 2008. Available at:
http://csl.stanford.edu/~christos/publications/2008.suzy_rivoire.phd_thesis.pdf
- Rizvi, R. 2009. **100G Ethernet and beyond: preparing for the Exabyte Internet**. JISC: Bristol, UK. Available online at:
<http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/horizonsscanning/hs0901.aspx>

Roussos, K. 2007. *Storage Virtualization Gets Smart*. **Queue**, Vol. 5, Iss. 6, September/October 2007. ACM: New York, USA, pp.38–44.

Saied M., Velasquez, G.T. 2003. PCs and Consumers – A Look at Green Demand, Use and Disposal. In: Kuehr, R. and Williams, E. (eds). 2003. **Computers and the Environment**. Kluwer Academic Publishers: Dordrecht, Netherlands.

Shacham, A., Bergman, K., Carloni, L. 2007. *The Case for Low-Power Photonic Networks on Chip*. In: 44th ACM/IEEE Design Automation Conference 2007 (DAC '07), 4–8 June 2007, San Diego, CA, USA. IEEE, pp.132-135. Available at: <http://www2.dac.com/data2/44th/44acceptedpapers.nsf/browse>

Simon, H. 2008a. *Is HPC going green?* **Scientific Computing**. May/June 2008. Available at: <http://www.scientificcomputing.com/HPC-Going-Green.aspx>

Simon, H. 2008b. **The Greening of HPC - Will Power Consumption Become the Limiting Factor for Future Growth in HPC?** Presentation to the HPC User Forum, 13–14 October 2008, Stuttgart, Germany. Available at: http://www.hpcuserforum.com/presentations/Germany/EnergyandComputing_Stgt.pdf

Smith, R. 2008. *Green Data Storage With MAID*. **Information Week**, 5th November 2008. Available at: http://www.informationweek.com/blog/main/archives/2008/11/green_data_stor.html

Socolof, M., Overly, J., Geibig, R. 2005. *Environmental life-cycle impacts of CRT and LCD desktop computer displays*. **Journal of Cleaner Production**. Volume 13, Issues 13-14 (Nov/Dec 2005). Elsevier. pp.1281–1294. DOI:10.1016/j.jclepro.2005.05.014

Stern, N. 2006. **The Stern Review on the Economics of Climate Change**. Cambridge University Press/HM Treasury: Cambridge, UK.

Stocks, J., 2009. *Developing a Rigorous and Rransparent Methodology for Carbon Counting*. **Presentation to Carbon Accounting Conference**, Heriot-Watt University, Edinburgh, 11 March 2009. See: <http://www.carboncounting.org.uk/>

Storer, M., Greenan, K., Miller, E., Voruganti, K. 2008. *Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk-Based Archival Storage*. In: **Proceedings of 6th USENIX Conference on File and Storage Technologies (FAST '08)**, 26–29 February 2008, San Jose, California, USA. Available at: <http://www.usenix.org/events/fast08/tech/> (10.30 a.m. session).

Street, T. 2008. **Greener ICT – saving the planet**. Presentation at SocITM conference, 13–14 October, 2008. Available at: <http://www.socitm.gov.uk/NR/rdonlyres/3D910500-C2D5-4648-BBF1-77187C49A86C/0/Soc08TerryStreet.ppt>

Suffolk, J. 2009. *Government Cloud*. **John Suffolk** (as HM Government Chief Information Officer) **blog**. Available at: <http://johnsuffolk.typepad.com/john-suffolk---government-cio/2009/06/government-cloud.html>

Tanenbaum, A. 2002. *Power Management in Operating Systems*. **InformIT**. 22 February 2002. Available at: <http://www.informit.com/articles/article.aspx?p=25451>

Turner, A. 2001. **Just Capital: the liberal economy**. Macmillan: London, UK.

University of Tennessee Center for Clean Products and Clean Technologies. 2001. **Desktop Computer Displays: A Life-Cycle Assessment**. U.S. Environmental Protection Agency: Washington DC, USA. Available at: <http://www.epa.gov/dfe/pubs/comp-dic/lca/>

Vasic, N., Kostic, D. 2008. **Energy-Aware Traffic Engineering**. EPFL Technical Report. Available at: <http://infoscience.epfl.ch/getfile.py?docid=24397&name=eate-tr-jun09&format=pdf&version=1>

Watson, J., Sauter, R., Bahaj, B., James, P., Myers, L., Wing, R. 2006. **Unlocking the power house: Policy and system change for domestic micro-generation in the UK**. Report of the ESRC Sustainable Technologies Programme. Available at: http://www.energy.soton.ac.uk/publications/unlocking_the_power_house_report.pdf

Webber, C.A., Roberson J.A., McWhinney, M.C., Brown, R.E., Pinckard, M.J., Busch, J.F. 2006. *After-hours power status of office equipment in the USA*. **Energy**. Vol. 31, Iss. 14 (November 2006), pp.2823–2838. Available at: http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6V2S-4HSY4JM-1&user=10&rdoc=1&fmt=&orig=search&sort=d&docanchor=&view=c&acct=C000050221&_version=1&_urlVersion=0&_userid=10&md5=889fd65795dfe8238a44f07319290d93

Weddell, A., Harris, N., White, N. 2008. **Alternative Energy Sources for Sensor Nodes: Rationalized Design for Long-Term Deployment**. IEEE International Instrumentation and Measurement Technology Conference (I²MTC 2008), 12–15 May 2008, Vancouver Island, Canada.

Weizsacker, E., Lovins, A., Lovins, L. H. 1997. **Factor Four: Doubling wealth, halving resource use**. Earthscan Publications: London, UK.

Williams, E. 2003. *Environmental impacts in the production of personal computers*. In: Kuehr, R. and Williams, E. (eds). 2003. **Computers and the Environment**. Kluwer Academic Publishers: Dordrecht, Netherlands.

Wilman, F. 2007. **Fresh Air Cooling in Data Centres**. FutureTech white paper, Feb 2007. Available at: <http://datacentredynamics.fr/ME2/dirmod.asp?sid=&type=contact&mod=Contact%20Form&cfid=5153C6313ADC42E493ABCA07173413CA> (registration required)

Wirth, N. 1995. A Plea for Lean Software. *Computer*. Vol. 28, Iss. 2 (February 1995), pp.64-68. DOI= <http://dx.doi.org/10.1109/2.348001>

Woo D., Lee, H. 2008. *Extending Amdahl's law for energy efficient computing in a many core era*. **IEEE Computer**. Vol. 41, Iss. 12 (Dec 2008). <http://doi.ieeecomputersociety.org/10.1109/MC.2008.530>

Wu, F., Hu, H., Liu, K. 2008. *A new energy saving storage system: SERAID for disk array*. In: **Proceedings of the Eighth International Symposium on Optical Storage and 2008 International Workshop on Information Data Storage**. SPIE, Vol. 7125, pp.71251I-71251I-6. Available at: <http://adsabs.harvard.edu/abs/2008SPIE.7125E..53W>

Zhai, B., Blaauw, D., Sylvester, D., Flautner, K. 2004. *Theoretical and practical limits of dynamic voltage scaling*. In: **Proceedings of the 41st Annual Design Automation Conference (DAC '04)**, 7–11 June 2004, San Diego, CA, USA. ACM: New York, NY. pp.868-873. DOI= <http://doi.acm.org/10.1145/996566.996798>