PAUL MOBBS

A practical guide to sustainable IT

Unit 7

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

Author
Paul Mobbs

Copy-editing
Alan Finlay

Layout proofreading
Lori Nordstrom

Publication production
Karen Banks and Flavia Fascendini

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

Commissioned by the Association for Progressive Communications (APC).

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

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

A practical guide to sustainable IT

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

South Africa
2012

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

LOCAL AND GLOBAL NETWORKS

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

While the use of electronic networks has developed computers into a significant new global technology, this has not come without an adverse impact. As communications have increased, so the amount of data flowing between users around the globe has increased, leading to a leap in the amount of data now routinely stored on personal computers; as more people connect to the network through mobile devices, more data is stored in large online systems to enable people to communicate on the move – and this is now maturing as a large cloud storage network; and while the efficiency of online technology has increased, the increase in data traffic[1] more than offsets that improvement, and so the energy and resources used to support the global communications network are still increasing.[2]

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

1. Minnesota Internet Traffic Studies (accessed June 2012), Internet Growth Trends & Moore's Law, www.dtc.umn.edu/mints/igrowth.html
2. Raghavan and Ma (2011). The Energy and Emergy of the Internet.www.cs.berkeley.edu/˜jtma/papers/emergy-hotnets2011.pdf

7.1. LOCAL NETWORKS

Local networks[3] connect computer users in a building, or people at home to their broadband router. Early networks used a number of different systems, but over the last two decades network connections have become standardised around different technologies that have been adapted to work in different spatial environments:

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

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

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

3. Wikipedia: 'Local area network'. en.wikipedia.org/wiki/Local_area_network
4. Wikipedia: 'Ethernet'. en.wikipedia.org/wiki/Ethernet
5. Wikipedia: 'Wi-Fi'. en.wikipedia.org/wiki/Wi-Fi
6. Wikipedia, 'Bluetooth'. en.wikipedia.org/wiki/Bluetooth
7. For example, see LessWatts, 'Tips & Tricks: Wi-Fi/Wireless'. www.lesswatts.org/tips/wireless.php

or for the connections between the ethernet hubs which link groups of computers and the network's data servers or high-speed internet connection. Depending upon the capability of the hardware and the operating system, it is possible to switch gigabit ethernet to work at the slower fast ethernet speed and conserve power[8].

Another major factor in running a network is the power used by the network electronics to provide connectivity. All devices use power, and when not in use they will still draw a minimum amount of energy while they wait to move packets of data between machines. For that reason it is better to have the smallest number of ethernet switches[9] or Wi-Fi base stations and use them to their maximum capacity – only adding new devices when the network experiences congestion. There is no easy advice to give on this point. How a network is physically laid out depends upon both the number of machines involved and the layout of the space they occupy. Theoretically the maximum length of an ethernet cable is 100 metres, although taking the cable through twists and turns around the edge of large rooms can quickly reduce the practical distance between the network switch and the computer. On small networks power can be saved by using a combined broadband/ISDN router[10] and multi-port ethernet hub rather than using separate devices.

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

7.2. USING NETWORKS

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

7.2.1. Email and attachments

Email is still the backbone of the way people communicate and network online. Email is simple to use, and has the benefit that files can be attached to the email and easily sent to one or many recipients. That ease of sending files is also one of the greatest problems with email. If the attachment has not been formatted for efficient delivery it will use an excessive amount of bandwidth in transmission, and if saved it will occupy a large amount of storage capacity.

The simplest way to reduce the size of the files sent via email is to use a data compression[12] program. This converts the data in text files and other raw data into a file which occupies a fraction of the space of the original file.

8. For example, see LessWatts, 'Tips & Tricks: Ethernet'. www.lesswatts.org/tips/ethernet.php
9. Wikipedia, 'Network switch'. en.wikipedia.org/wiki/Network_switch
10. Wikipedia, 'Residential gateway'. en.wikipedia.org/wiki/Residential_gateway
11. Wikipedia, 'Energy-Efficient Ethernet'. en.wikipedia.org/wiki/Energy-Efficient_Ethernet
12. Harris, Tom (accessed June 2012). How File Compression Works, How Stuff Works. www.howstuffworks.com/file-compression.htm

Box 7.1.

Reducing/scaling digital media files

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

Image files

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

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

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

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

Sound/audio files

Sound is recorded as raw audio data – this is what is found in WMV files, and so WMV is the least efficient way of sending audio data. Other file formats encode the data, and while this reduces the file size it progressively results in lower quality sound reproduction. There are three factors which have a great influence on file size: The sample rate; the number of channels; and whether the "codec" (the standard used to encode/decode the audio data) is lossy or lossless.

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

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

Video files

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

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

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

The problem with many file formats in common use today – such as *OpenOffice* or Adobe *Acrobat* – is that they are already compressed as part of their formatting. Therefore simply compressing the file will not reduce its size significantly. In these cases you have to change the way these files are created by the original application. Adobe Acrobat (and the "Export PDF" functions within office applications) allow you to reduce the image resolution of the document and compress any images which it contains (the text of Acrobat files is already compressed). The proprietary Adobe Acrobat application offers many more options to reduce the size of the file[13] which can reduce the size of PDF files significantly. Many applications offer additional plug-ins or extensions to add extra functions which optimise file sizes. For example, OpenOffice Impress has a downloadable extension called *Presentation Minimizer*.[14] This adjusts the formatting of fonts, images and animations within presentations to reduce the size of the file generated.

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

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

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

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

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

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

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

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

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

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

7.2.2. Web browsing

Over recent years, especially since the widespread adoption of broadband connections, the size of web pages has grown significantly. This is in part the result of more graphically complex pages, and the increasing role that advertising and content tracking have in the design of many websites and blogs. Especially for sites with animated advertising, and which require large quantities of scripting code[17] to control the behaviour of the page, the ecological impacts are not just related to the amount of data downloaded – executing local scripting can also require a large amount of the machine's processing power. For example, some newspaper sites now required a megabyte or more of data to be downloaded for each page; and on some web pages (especially where poorly designed or incompatible Javascript controls are used) viewing the page will increase the load on the processor, increasing the machine's power consumption.

The need to download data can to some extent be controlled through the configuration of the web browser. Some aspects of minimising the downloading of data are part of the browser's configuration, while others require plug-ins or additional software:

- ***Web cache*** – Web browsers use a cache[18] to hold copies of commonly downloaded files. When a page is requested if the cache already has a copy it checks if the copy on the site has been changed since it was last downloaded. If the content is unchanged the local copy is used and so it need not be repeatedly downloaded. If routinely browsing advertising/graphically rich websites it is a good idea to increase the size of the cache to take account of the large script and image/flash files used by these sites. For example, the Firefox browser defaults to a limit of 50 megabytes of cached files; extending this to around 100 to 120 megabytes will allow more of those repeatedly displayed adverts and the code/data files which go with them to be cached and will help to speed loading of the page and reduce the amount of data downloaded.
- ***Pop-ups*** – A common feature on advertising-driven sites is that certain actions by the user, such as clicking on buttons or lists, will cause another browser page to "pop-up.[19] This page invariably contains advertising or other unwanted information, and can launch additional code within the browser than can tie up the processing power of the system and even compromise security. Most browsers include an option to block pop-ups, which is enabled from the browser's configuration menu. Quite apart from avoiding additional data downloads, blocking pop-ups will reduce the load on the processor and avoid loading more data into the system memory each time a new browser window is opened.
- ***Advertising filters*** – Even with a pop-up blocking enabled, many sites now use dynamic HTML to display adverts or animated content which float[20] over the browser window. As these stay within the control of the existing window this circumvents control by pop-up blocking. Instead what is required is an advertising filter[21] plug-

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

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

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

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

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

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

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

7.2.3. Proxy servers

So far we have looked at the steps which can be taken to minimise the impacts of world-wide web use from the web browser. For larger organisations many of the above features can be enabled with a web proxy server.[24] Users of the network access the web through the proxy, and the rules defined for serving pages will block and filter certain types of content as well as blocking selected websites. Proxy servers are becoming more popular on small and home networks too as they are a simple way to implement content filtering as part of "parental control" systems.[25]

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

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

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

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

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

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

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

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

28.Squid, www.squid-cache.org

Box 7.2.

A case study: The "efficiency" of web content

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

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

- A fifth of the 170 web pages downloaded had 100 or more files associated with them – as a general trend those with a higher advertising content (primarily the new media sites) had the highest file counts;
- Measuring the size of the page and its associated files, a third of all the pages had a size of more than a mega-byte, and 2% were over three megabytes;
- Sorting the web pages by size, there is a general trend for larger pages to be dominated by Javascript code files rather than text or graphical data – smaller pages were proportionately made up of more graphical data; and
- Isolating the types of information contained in each page, text and graphical data made up a third of the average page – the rest was formatting (style sheet) and control (flash/Javascript) data.

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

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

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

7.3. WEBSITES AND CONTENT

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

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

This is the root of the "IT productivity paradox" paradox.[30] It is not necessary to work hard to increase productivity if there is no incentive to create efficient resources; the rising power, capacity and functionality of technology is able to make up for the lack of attention to the efficiency of design. This, of course, is where the problem of bloat arises.[31] If computers become more powerful there is no incentive to improve the quality of the code use to make them function; and for web bloat, if network speeds are rising there is no need to be more creative with site/information design. For example, when mobile phone operators paid billions for their licences in the 1990s they had an incentive to get as much capacity out of their networks as possible – and through new transmission protocols they achieved a higher capacity than initially expected.[32] In contrast, the designers of web systems have no such pressures on the quality and efficiency of their work.

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

Unfortunately, there is no simple route to achieving greater efficiency in web design. In 2010, the Free Range Network undertook a small-scale study to evaluate the design-related impacts of web use[35] (see summary in Box 7.2). The purpose of the study was to highlight strategies to redesign their own website in order to make web pages more efficient when accessed – reducing data downloads and the amount of processing power required to display the page. A general problem is that

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

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

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

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

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

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

35.As yet unpublished, a general summary of the findings is available Mobbs, Paul (2011). Web design, environmental impact and the much ignored ecological efficiency of the Internet, Ecolonomics No.12, May 2011. www.fraw.org.uk/mei/ecolonomics/01/ecolonomics-012-20110523.pdf the use of content creation programs encourages inefficient site design; those creating content do not understand the factors which cause bloat in online content, and so cannot format their information in a more efficient manner. The solution to this is to seek more direct means of content creation, which allow for the optimisation of the content generated. This would require a greater level of training and knowledge of web design, and the methods which can minimise the impacts of online content (such as the reformatting of files to reduce their size, outlined above in relation to email attachments). At present this is a novel area for research, and while there is a large effort devoted to search engine optimisation, or the inclusion of advertising and web analytics as part of site design, there is as yet little demand to optimise web content to reduce its download size and demand for power when displayed.

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

7.4. THIN CLIENT NETWORKS

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

Now the trend is swinging back towards client-server networks once more. As network speeds increase, and cheap and scalable rack servers become more powerful, the cost of developing networks of powerful stand-alone computers is becoming comparatively more expensive. Add to that the reduced costs of managing one large server rather than maintaining the operating systems of many stand-alone computers, and the costs of using client-server networks rather than many powerful desktop systems becomes an attractive option for larger organisations and education institutions.

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

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

Figure 7.1.

A client-server network

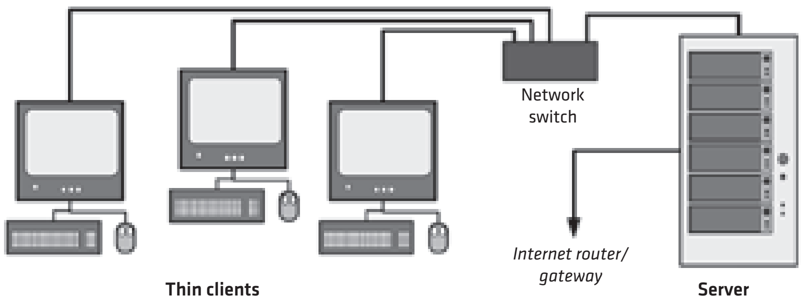

them as thin clients. Much of the expense of the upgrade is the cost of the powerful server which runs the system. Then, as the older clients fail or become unserviceable, they can be replaced with cheaper purpose-designed thin clients, further reducing the power consumption of the system.

The thin clients[37] run a minimal operating system: On old PC machines the system is installed like a conventional stand-alone operating system; on purpose-designed thin clients the system is often an embedded software system held in flash memory. Being able to do away with hard drives and other components of the stand-alone desktop system is one of the reasons why purpose-built thin clients are cheaper and use less power.

The central server runs a dedicated time-sharing system. This runs the programs requested by users and sends information to generate the desktop display back to the terminals. Time-share systems were originally developed for Unix machines. These ideas were taken on by the Linux community and developed into a variety of low-cost terminal server systems.[38] Their low cost, and ability to utilise cheaper/recycled equipment, has made these systems especially popular in educational establishments,[39] particularly in developing nations.[40] While initially sceptical,[41] even Microsoft now offer a terminal server system[42] compatible with Windows applications, and there are programs available to convert older Windows XP machines into thin clients.

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

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

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

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

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

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

42.Microsoft (accessed June 2012). Windows Embedded Thin Client. www.microsoft.com/windowsembedded/en-us/evaluate/windows-embedded-thin-client.aspx user licences may have to be purchased to run certain software across a client-server network rather than a single licence for use on a stand-alone machine.

Overall, whether or not a client-server network creates cost and energy savings depends upon the types of tasks carried out by existing network users, and the extent to which any pre-existing hardware can be reused when creating the network. Given the current trends in computing, especially the shift towards more centralised/online services, it is likely that client-server networks will become more widely used in situations where many people share the same network system.

Box 7.3.

Networking check-list

Networking hardware:

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

Using networks:

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

Websites and content:

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

Thin client networks:

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

A practical guide to sustainable IT

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

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

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

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

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

APC

IDRC | CRDI