

September 2009

TechNews is a technology, news and analysis service aimed at those in the education sector keen to stay informed about technology developments, trends and issues.

Please note that Becta does not accept any responsibility for or otherwise endorse any products contained in these pages including any sources cited.

Please navigate the newsletter by clicking on items within the table of contents, or using bookmarks:

Networking and wireless	3
Analysis: Grid computing	3
Networking and wireless news	6
Improving data centre scalability	6
EU invests in LTE Advanced research.....	8
Space routers to improve satellite internet	8
Internet, wireless data and online video trends	9
US researchers to develop 100Gbps network.....	10
802.11n wireless standard may be ratified in September [Updated]	10
ZigBee to offer 'power-free' wireless connections	11
M2M and the 'internet of things'	11
OneApp to run them all	12
Smartphone sales to outnumber PCs by 2011.....	13
Multimedia	13
Analysis: Augmented reality	13
Multimedia news	17
NVIDIA offers GPU support for ray tracing.....	17
Company aims to produce laser table-top touchscreens	18
Sky to launch 3D TV channel.....	18
New liquid crystal panels save power	19
Researchers improve LED production, create flexibility	19
Flash expected on smartphones in 2010	20
Video streaming over HTTP proposed.....	21
Print directly from the internet with web-connected printer.....	22
New information tag to compete with barcodes.....	23
Selected readers to see first video ads in print publication	23

Hardware.....	24
Analysis: Visualisation technologies.....	24
Hardware news.....	27
Server and mobile processor update, September 2009.....	27
Creating processors at nanoscales.....	29
Spasers could underpin optical computers.....	30
New transistor design cuts charger power loss.....	31
Manufacturers close to wireless charging standard.....	32
New approaches to improved touch screens.....	32
Darkfield tracking enables optical mice on glass.....	33
Assistive technology for the blind, deaf.....	34
The Digital Rosetta Stone to preserve data.....	35
New installation standards for audiovisual equipment.....	35
Software and internet.....	36
Analysis: Speech recognition.....	36
Software and internet news.....	39
Google announces Chrome OS.....	39
Blurring the line between browser and operating system.....	41
New Microsoft subscription licensing scheme launched.....	41
Nokia and Microsoft collaborate to improve Office on phones.....	42
'Best before' encryption key time-limits online data.....	42
W3C introduces new knowledge management standard.....	43
New algorithm sifts experts from spammers.....	44
'Sexting' creates concern for over a third of teens.....	44
Free-to-use tech training videos for teachers.....	45
Free e-textbooks linked with material purchases.....	45
TechNews Information.....	46
Disclaimer.....	46
Copyright and permitted use.....	46
To unsubscribe:.....	46
Feedback:.....	46

Networking and wireless

Analysis: Grid computing

At a glance

- Grid computing involves aggregating processing capacity, storage and other resources to achieve a task that would be beyond the resources of a given institution.
- The electricity grid is used by some as a model for grid computing, but processor power 'on tap' is a long way from reality for complex projects.
- Volunteer computing projects rely on individuals donating spare 'processing cycles' on their PCs to solve problems in areas such as medicine, astrophysics and climatology.
- Large academic research projects, such as those planned for the Large Hadron Collider (LHC), rely on grid computing to analyse their vast data sets.
- Issues around security and mutual trust when 'donating' capacity are balanced by advantages of higher resource utilisation and contributing to worthwhile projects.
- Educational networks could be used to advance volunteer projects or for in-house processing tasks, but security issues may limit the former while the latter may be more effectively achieved through 'cloud' services.

Getting on the grid

The broad definition of grid computing, otherwise known as utility computing, entails the notion of making computing power as available as the national grid - some strategists foresee a time when you will be able to plug a terminal into a 'wall socket' and get all the computing power you need. This view simplifies the current state of computing to 'pure' processor power, analogous to electricity, without reference to all the complexities of differing processor architectures, storage requirements, peripheral interactions and a host of other factors. In many respects cloud computing (see [TechNews 11/08](#)) offers these facilities by providing computing power and storage via the internet; the user does not know where those servers are located, but can lease the capacity required.

'Grid computing', in more common use and as discussed in this article, refers to a form of distributed computing whereby users can access spare capacity on other people's resources to deal with tasks that would take far too long on in-house hardware. Provision is enabled by a complex web of co-operative pacts and predefined service level agreements (SLAs) that are a far cry from the 'plug in, use now and get billed after' vision of utility computing. As defined on [Wikipedia](#):

Grid computing (or the use of computational grids) is the combination of computer resources from multiple administrative domains applied to a common task.

This definition indicates one of the key features of current computing grids: heterogeneity. There are many computing platforms and a whole mass of

communities, research projects and nascent standards, only some of which will be covered in this article.

Grid computing is most suited to scalable, massively parallel computing tasks. These applications can generally handle out-of-order processing, with algorithms that deal with late or missing results, and rely on 'message passing' protocols to control execution by allocating tasks, sharing progress and transferring completed data to the appropriate point. Such tasks include searching very large data sets, video rendering, climate simulations, genome analysis, processing particle physics data and drug research. Some projects involve 'volunteer computing' where people grant access for applications to run on spare processor capacity while their computer is idle. One of the most widely known examples is the [SETI@home](#) project, searching for signals attributable to intelligent sources among the radio background 'noise' of the universe. Some projects allow the greatest contributors to propose their own tasks to be run on the virtual, networked processor.

Many large academic research projects also use grid computing, taking advantage of facilities in partner organisations to process data during idle time, perhaps at night or between in-house applications.

Educause has a helpful article, [7 things you should know about Grid Computing](#) and the Worldwide LHC Computing Grid (WLCG) has published [Grid computing in five minutes](#). Grid computing was previously covered in TechNews 09/05.

The structure of the grid

The grid is inherently heterogeneous, a loose collection of processors, storage, specialised hardware (for example particle accelerators, electron microscopes and telescopes) and network infrastructure. For each task, appropriate hardware has to be discovered, processor time booked, network capacity scheduled (especially where large data sets are involved) and collation of results organised. Although this can be achieved on a peer-to-peer basis (in which no one machine has overall control), it is generally arranged as a client-server structure. 'Middleware', is often utilised to manage the applications and resources required to achieve a particular outcome, such as the [Globus toolkit](#) or Berkeley University's [BOINC](#) software (both of which are open source).

The complexities of managing grid applications are offset by significant advantages, including:

- access to resources beyond those available within a given institution
- optimisation of spare capacity
- flexibility to scale and reconfigure available resources
- avoidance of single points of failure in the computing infrastructure used
- data replication across a number of facilities
- provision of 'virtual' resources in-house, so that experienced researchers are less tempted to go to institutions elsewhere.

Academic institutions have created partnership groups for sharing resources, notably [GridPP](#) (for particle physics tasks in the UK), the EU's [EGEE](#) science network and the UK's [National Grid Service](#) (NGS); while international directories like the [GridGuide](#) provide international contacts. The [Open Grid Forum](#) (OGF) has been behind a number of substantive projects, especially developing standards for the protocols required to deliver and manage grid applications.

Volunteer computing

Volunteer projects are the simplest structure of grid computing: a server provides an application for users to download and a series of 'work units' to be processed during the processor's idle time; each work unit is independent, so results can be returned in any order. However, the researchers running the application do not know whether the user or client PC will produce accurate, authentic results, so tasks are generally randomly duplicated between users, with results compared to ensure validity. The owner of the client PC has to manage the installation and patching of the client application, while trusting that the application provider is doing the work purported, that no malware is being delivered and that the application will not interfere with the operation of the computer. Networks of PCs in schools and colleges could contribute huge numbers of spare processing cycles to these projects, but management overheads and security concerns often deter system managers from volunteering their resources.

Applications include research into disease, medicines, climate change, astronomy and particle physics. [GridRepublic](#) and the [World Community Grid](#) allow users to select the projects they wish to contribute to, while [Intel is promoting](#) its volunteer projects through Facebook. Many projects, such as the protein folding simulation [Folding@home](#), now support processing using games consoles and the parallel instruction pipelines found on graphics processors. (See 'GPU computing' in [TechNews 09/08](#).)

Research networks

Collaborative networks of academic researchers can assume that the infrastructure is trusted, diminishing the problems faced by public volunteer projects. However, the actual tasks are often far more complex, involving very large data sets and a much greater range of hardware, from desktop PCs through to supercomputers.

The Large Hadron Collider (LHC) [will be reliant](#) on massive grid computing capabilities to process the data that it is expected to produce. The WLCG has 11 Tier 1 and 140 Tier 2 data centres that will distribute the 15 million gigabytes (15 petabytes) of data created each year. The primary fibre optic network links will run at 10Gbps, allowing data transfers of several gigabytes per second through clustered channels.

Computing facilities at this scale represent a considerable investment, so the prioritisation, scheduling and job control are critical to effective use. A number of projects and protocols (in the widest sense) are focussed on this issue. For example:

- [GridFTP](#) is an extension to the standard internet file transfer protocol (FTP), allowing much larger blocks of data to be simultaneously and securely transmitted across multiple channels, as well as providing the facility for just part of a single, extremely large file to be downloaded.
- [GridCOMP](#) and [Phosphorus](#) are assembling frameworks and services to facilitate higher level project management.

Commercial opportunities

Large companies, including those involved in pharmaceuticals, aerospace simulations, data mining for market research and prospecting, also have immense processing requirements, but the data they handle can be commercially sensitive. SLAs must be legally watertight, covering issues like security, intellectual property and data protection (especially where personal information is held).

The function of [GridEcon](#) is to create an online auctioning, scheduling and management system for computational capacity. The EU's [SIMDAT project](#) had a wider remit, investigating all elements of grid computing, from protocol standardisation through to systems that allow companies to readily create virtual organisations through which they can define projects involving establishing, administering and securely taking down distributed processing and storage capacity.

The grid and the cloud

Many applications already run in the 'cloud', leasing facilities such as Amazon's Web Services (AWS) or Microsoft's soon to be launched Windows Azure Platform. Although these may use a distributed computing model to provide the services, they have a single point of accountability through the provider's SLA. The grid computing applications outlined in this article are far more complex, but they can provide computing power for 'free', or at a substantially reduced price, for academic researchers, while ensuring near full utilisation of expensive computing resources. This grid remains more informal in structure, collaborative in development and altruistic in nature, although it is becoming more formalised as the environment matures and the scale of individual projects increases, especially as commercial entities begin to adopt these approaches.

Educational establishments could consider donating spare computing cycles to advance areas of research considered to be for the good of humanity, although they need to factor in the management overheads that deployment is likely to incur and consider whether it will add significantly to energy consumption. Middleware, such as BOINC, could be deployed across a large institution to manage in-house processing tasks, or capacity could be leased from one of the cloud providers. However, access to massively scalable, grid computing resources is likely to remain the province of research organisations based in higher education and industry.

Networking and wireless news

Improving data centre scalability

Data centres hold growing masses of information, making management of the hardware and its interrelationships increasingly complex. This trend is being

exacerbated by the move to put data into the 'cloud' (see [TechNews 11/08](#)), which makes data more widely available, as well as gaining energy and other efficiencies from centralising and combining hardware. Such data may, in future, include information on learners and their educational progress. In order to support diverse applications, servers and data storage are 'virtualised', allowing them to function independently of the underlying hardware. This allows 'hot swapping' of logical servers between physical machines, which sounds simple enough on a small scale but becomes considerably more complex for large databases and multitudes of virtual machines.

Displaying a single Facebook page [is said to involve](#) 1,000 servers, each of which must be able to communicate the relevant information reliably. If any node or link is broken, the network protocols must be designed to find backup paths without creating closed loops that would flood the data centre network with traffic. Data centres have been structured on the same principles and protocols as enterprise networks, based on the seven-layer OSI (open system interconnection) reference model. Layer 3 uses IP addresses to route information, but these logical references must be mapped onto the MAC (media access control) address of a physical device. (The IP address structure for large numbers of devices needs to be partitioned into separate 'virtual' local networks, which requires human configuration, limiting the ease with which hardware and virtual machines can be migrated around the data centre's logical structure.) Layer 2 protocols (such as the widely used spanning tree protocol) use MAC addresses, so they are directly addressable, but generally do not scale well to large numbers of machines.

A new topology and set of protocols, dubbed PortLand by its creators at University of California, still broadcasts address resolution protocol (ARP) requests for servers and other resources that a device needs to connect with. However, instead of these requests being broadcast across the whole network, the request is immediately handled by a directory service that has a record of a pseudo-MAC (PMAC) address for devices on the network. By using MAC-based addressing, packets can be sent over standard switches and other Ethernet hardware without the complexities of introducing a full Layer 3 IP-based structure. On the final link, the PMAC is converted back into a true MAC, so that the intended device receiving the request will respond correctly. The system automatically creates the required PMACs for physical devices and logical servers, allowing easy migration of equipment and virtual machines around the whole structure. The research team believe that PortLand will readily scale to control over 100,000 servers in a data centre network.

Microsoft engineers [are proposing](#) a different approach, that they refer to as Virtual Layer 2 (VL2). This involves giving each service the impression that all the hardware needed is connected to a single Ethernet switch. Unlike PortLand, which still uses the traditional tree structure, VL2 implements a mesh network and enforces a random traffic routing around the structure to avoid overload on any given link. VL2 is briefly outlined on the [Microsoft research](#) website.

In related news, a team involving researchers at MIT, Carnegie Mellon University, and the networking company Akamai, [have proposed](#) a smart routing algorithm for

large, multi-site data centre networks that will direct processing to those locations where (for reasons such as the time of day) electricity is least expensive.

EU invests in LTE Advanced research

Current third generation (3G) wireless networks are adequate for voice calls and casual data traffic, but offer insufficient bandwidth for regular mobile browsing. New protocols, such as high speed packet access (HSPA) are providing incremental improvements to data rates and are often referred to as '3.5G' technologies. Vodafone [recently announced](#) that it has begun upgrading its network to give a maximum (theoretical) data rate of 14.4Mbps, although this is likely to be experienced as little better than 1-5Mbps due to distance from masts, interference and sharing bandwidth with other users.

Two technologies are competing to create next generation 4G data networks: WiMAX and LTE. Some operators ([for example Sprint](#)) have already begun to implement WiMAX in urban areas of the States, but it is seen as more likely to take hold in less developed regions of the world that have limited existing mobile infrastructure. Long Term Evolution (LTE) is seen as the more promising technology in Europe and most parts of the US, since it has been specifically designed by the Third Generation Partnership Project (3GPP) as the 'natural' upgrade path from existing 3G investments. Operators are starting to test equipment even though the underlying '3GPP Release 8' has yet to be formally ratified as an agreed standard. LTE promises theoretical download rates of up to 326Mbps, although 100Mbps is the more immediate target - again dependent on distance and shared bandwidth. Development of LTE was supported by the EU and the first roll-outs are expected next year, although widespread coverage will take another three to five years.

The European Commission has [now committed](#) €18 million (around £16 million) into developing the next set of standards, LTE Advanced. Although LTE is often called 4G, its bandwidth (in particular) does not meet the International Telecommunication Union's (ITU) [IMT-Advanced](#) requirement for a peak maximum speed of 1Gbps for links involving low mobility. LTE Advanced is being developed as 3GPP's [Release 10](#).

Proposals for LTE, which uses radio spectrum more efficiently, will be factored into the pan-European plans for the 'digital dividend'. (See [TechNews 07/09](#).) Because it can reach further, fewer cells will be required and better coverage should be achieved in many rural areas. The high data rates will allow bandwidth-hungry applications, such as full screen high definition television, to operate effectively over mobile networks.

Space routers to improve satellite internet

Cisco is working on an [Internet Routing In Space](#) (IRIS) project, part-funded by the US Defence Department. The terrestrial internet uses intelligent routers to identify a suitable path that avoids network congestion: as packets are transmitted, the router maintains a dynamic 'routing table' of the availability of potential paths to intended destinations. However, satellite connections tend to be bound to specific terrestrial

base stations to form point-to-point links, with no alternative paths for failed packets to reach their destination.

Flexibility can be provided by adding terrestrial routing to the satellite downlinks, but the distances involved create a delay (latency) of at least a quarter of a second for each satellite link used, so this is impractical for many purposes. Cisco's solution would be to combine a router into the communications hardware of the satellite itself, allowing some packets to be routed between packet streams on different antennae without having to return them to earth to do so. Hence, VoIP (voice over IP) calls could be connected directly between parties on the same satellite. Routing between satellites would also be possible if they were equipped with suitable antennae to communicate with each other.

Cisco is [planning to launch](#) its prototype 100Mbps router on Intelsat's IS-14 vehicle later this year. Due to the conditions in space, the hardware will have to be extremely robust and include high levels of redundancy. Power consumption and cooling are two other issues that the engineers must overcome.

Satellite connections are generally leased on a sole use model, meaning that bandwidth on one satellite may be idle while another is struggling to meet demand. Intra-satellite routing could enable a 'pay per use' model for bandwidth, as packets could be routed towards idle capacity on any satellite. This could make some forms of satellite communication cheaper, but would be unlikely to improve broadband services, which are linked to satellites at a specific location in geostationary orbit.

Satellite broadband was covered in [TechNews 05/09](#).

Internet, wireless data and online video trends

The majority of users in the US now access the internet through some sort of wireless connection, be that wireless LAN, Wi-Fi hotspot or mobile, [according to reports](#) of the latest Internet & American Life Project from the Pew Research Centre. The researchers found that 56 per cent were using a variety of wireless devices, including laptops, mobile phones and games consoles, for internet access. Both learners and staff are increasingly coming into schools with these personally owned devices, so expectations about their use are likely to increase. Bill Thompson, a columnist for the BBC, [writes about his experiences](#) of living with limited bandwidth and how frustrating he found it when attending a conference with inadequate wireless internet access.

[Further reports](#) emphasise the continued growth of watching online video, including a rise in the numbers watching complete television shows and films. According to Pew, more than three in five Americans watch video streams, which is nearly double the number in late 2006. The proportion was even higher among young adults, aged 18 to 29, at almost 90 per cent. Due to the quantity of information needed to create a video image, this has considerably increased network traffic, leading to reports of users experiencing throttling of video applications, such as BBC's iPlayer, during peak times. The current consensus in the UK is for effective 'net neutrality', where all data streams get equal treatment within the network, but exponential growth of video

may lead to companies seeking to add restrictions on intensive data use into their broadband contracts. The Pew researchers also found that users had been more ready to relinquish cable television than internet access during the current recession.

[Predictions by](#) Nokia Siemens Networks have suggested that data use will surpass voice on mobile networks by 2011, with volumes continuing to rise exponentially into 2013. This will continue to put pressure on network bandwidth, such that carriers will be looking to both upgrade core network infrastructure and to throttle back the most data intensive mobile applications, especially video.

Overall, Forrester Research [is reported](#) to have found that the number of people connected by any means to the internet will have risen to 2.2 billion by 2013, approaching half of whom (43 per cent) will be in Asia.

US researchers to develop 100Gbps network

Schools and businesses may use 1Gbps fibre Ethernet backbones for their networks, while individual computers are connected at 100Mbps. However, the rapid rise of video content will soon be putting pressure on some of these network links, so even faster core networks will be required. JANET is already testing 100Gbps network links (see [TechNews 06/09](#)), but Californian group is researching configurations that it hopes will eventually provide 1Tbps. (A terabit is 1000Gb.)

The Lawrence Berkeley National Laboratory has received funding to pilot 100Gbps links, which are ten times faster than the existing backbone, for ESnet (formerly the US Energy Sciences Network) that connects Department of Energy supercomputers. The [press release](#) notes that a climate modelling programme at the Lawrence Livermore National Laboratory is set to increase the data it stores nearly 20-fold, to 650TB, and a new distributed worldwide climate database would hold more than 6PB. (6PB is equivalent to 6,000TB, or 6 billion megabytes.) In addition to high bandwidth, data transfers at this scale require intelligent routing and bandwidth reservation systems to ensure successful international data transfers. These services are provided by ESnet's [OSCARS system](#).

Moves to higher bandwidths will provide higher capacity infrastructure for the organisations involved, but are likely to also help drive down costs of standard network components for lower capacity systems. High speed networks will also be required to ensure efficient operation of the major interconnects between large server farms operated by companies like Amazon, Google, IBM and Microsoft.

802.11n wireless standard may be ratified in September [Updated]

Many products on sale designed to work with wireless networks now use "draft-n" Wi-Fi protocols. The final 802.11n specification has yet to be finally approved by the IEEE, the body which is guardian of many technology standards. [PC Advisor reports](#) that a member of the working party for 802.11n has stated that a final draft specification would be sent to the main review committee in September. The Wi-Fi Alliance, which certifies wireless LAN products, has been using an earlier 2007 draft for 2 years and [does not expect](#) any major changes to the compliance certification process, while new products should remain backward compatible. However,

conflicting reports have emerged, with [PC Pro quoting](#) an engineer from wireless equipment manufacturer Buffalo, who suggested that there would instead be a third draft this autumn.

Standardisation should guarantee hardware interoperability, so some corporate purchasers have been awaiting final ratification of 802.11n to deliver certainty before committing to widespread deployment. The most recent approved Wi-Fi standard, 802.11g, provides a maximum data rate of 54Mbps. The new 802.11n standard is expected to be able to deliver 450Mbps (from the current draft) or 600Mbps when it is finalised, by employing MIMO (multiple-input multiple-output) technology. MIMO uses several antennae and advanced techniques to focus the signal towards receivers, to increase transmission range and support higher data rates. The network speed experienced by users will be considerably lower than 450Mbps due to network overheads, working in a less than ideal wireless environment and contention (competition) from others. Becta's [Technical specification: institutional infrastructure](#) contains further details of wireless networking standards and advice on creating wireless networks.

[Update 16/9/09: 802.11n [has been ratified](#), as most commentators expected.]

ZigBee to offer 'power-free' wireless connections

ZigBee is a low power wireless communications technology for mobile and embedded hardware, aimed particularly at remote monitoring and control systems. It is based on the IEEE 802.15.4 standard and governed by the [ZigBee Alliance](#). (ZigBee was more extensively covered in TechNews 02/05.)

The Alliance [recently announced](#) that a 'Green Power' standard will be available to participating companies at the end of this year. The new standard will mean that equipment connected to energy harvesting devices can be approved by the Alliance under its certification programme. Although some ZigBee devices can already 'run on regular batteries for years', energy harvesting technology could create a largely maintenance-free system. Such technologies include harvesting 'free' wireless power ([TechNews 06/09](#)), energy from vibrations and small-scale solar power. Environment control in buildings and remote monitoring of the 'health' of public infrastructure (for example bridges) are likely to be some of the first applications of these devices.

The Alliance [also announced](#) an agreement with the HomePlug Powerline Alliance, aimed at uniting systems to create the 'smart' electricity grid, such that thermostats and other devices could be adjusted using control information carried over the mains. (The HomePlug Powerline Alliance represents a range of companies that create data networks over the mains electricity infrastructure.)

M2M and the 'internet of things'

Machine to machine (M2M) technologies will be a key part of the 'internet of things' - an emerging trend in which embedded smart devices are networked wirelessly, allowing devices to talk to each other (M2M) as well as respond to instructions from a user. The internet of things could encompass tracking the location and use of items issued to students, environment control in buildings, smart energy grids that would

drive down electricity consumption, vehicle management, automated fault-reporting for plant and ICT infrastructure, fully automatic pilots and remote healthcare. A management layer sits between the user, who sets general targets, such as the required room temperature or a destination, and a complex control and monitoring system that interacts with the devices to achieve them.

Vodafone, by no means the only mobile operator in the M2M sector, [recently announced](#) a platform to support the roll out of new services, with facilities to control individual devices and collate large amounts of data. Mobile operators have obvious reasons to promote mobile integration of as many devices as possible, including hardware like cameras and e-book readers, as well as sensors and control systems. Quoting analyst Berg Insight, Vodafone says that M2M revenues for mobile operators are expected to grow from €3 billion last year to €8.9 billion in 2012 globally.

Qualcomm, one of the manufacturers of radio and processor circuits in this sector, recently held a 'Smart Services' conference at which a spokesman for the GSMA (the mobile industry forum) [stressed the need](#) for much greater interoperability in order to integrate diverse devices into a coherent structure. At the conference, Qualcomm [announced a joint venture](#) with Verizon Wireless to develop 'next-generation products' using Qualcomm's hardware expertise and Verizon's mobile technologies, including LTE wireless connectivity. (LTE is a 'fourth generation' mobile technology aimed at creating high bandwidth mobile data and voice networks.)

OneApp to run them all

Smartphones run all types of applications provided by the handset manufacturer, mobile service provider and (increasingly) downloaded from online application stores. These applications include games and interfaces for social networks, but also business and educational software. The phones that run these applications are expensive and often paid for using relatively expensive monthly contracts. Not only do many British learners not have access to such hardware, but it is beyond the means of most people in developing nations.

Microsoft [is introducing](#) OneApp, a 150KB software interface that runs on more basic 'feature phones', into emerging markets. The application collates icons for all the registered services onto a single display and then downloads and run only the appropriate part of the software for each feature of the service as it is required. Some of the processing is achieved in the 'cloud' (on remotely networked servers), so the network carrier has to be involved in implementing the system as a whole. This approach allows users to interact with sophisticated services without having to purchase advanced smartphones or pay for high volumes of data that would be accessed by loading the whole of a fully-featured application.

In addition to social networking, applications could be written for health monitoring, agricultural advice and 'online' learning. The press release suggests that applications would have to be specifically written using the OneApp development framework,

using tools like JavaScript and XML. The first implementation will be with Blue Label Telecoms in South Africa.

Smartphone sales to outnumber PCs by 2011

Based on the second quarter of this year (286 million), sales of mobile phones in all categories exceeded one billion. This is a reduction from the same quarter last year, but [reports say](#) that Gartner's analysis shows that 40 million of these handsets were smartphones - giving 160 million per annum - and that their market share is increasing. Another group of analysts, RBC, are saying that if this trend continues, sales of mobile phones will outnumber PCs by the end of 2011, with shipments of each surpassing 400 million per annum.

Smartphones are much more capable of running web applications and browsing the internet than previous generations of phones, making them more able to interface with learning platforms. Many older students own smartphones and will increasingly expect to access educational content using them. This suggests that schools and colleges will need to plan to support a more diverse range of devices for home and remote learning.

Multimedia

Analysis: Augmented reality

At a glance

- Augmented reality (AR) combines virtual artefacts - information or 3D objects - to supplement the user's field of view. AR may also use audio and touch feedback.
- Printed tags can be used as 'fiduciary markers' to place virtual objects on the user's video display, providing a tangible link between the physical and real worlds.
- Applications on powerful mobile phones, equipped with GPS, a compass and accelerometers, can download and display location-based information over the handset's display.
- More powerful, semi-immersive systems use significant processing capability, requiring expensive, bulky hardware. These systems are mostly found in research and specialist environments, such as head-up displays for fighter pilots.
- The ability to 'annotate' and superimpose objects over real locations could have educational applications, while creating interactive environments using AR may help stimulate and involve younger learners.

The creative clash of the virtual and the real

Augmented reality (AR) spans the divide between the 'real' and virtual worlds, delivering data and computer generated imagery (CGI) to bear on the user's environment. The existing work has tended to display reference information or video content as an overlay on the user's field of view, but it could potentially add audio, touch feedback (haptics) and smells to create a more rounded sensory experience of the merged worlds.

Head-up displays (HUDs) have already been commercialised, especially in aeroplanes and some cars (for example, using [Microvision's display systems](#)), providing flight data, navigation prompts and weapons information. HUDs integrate a projector and some type of semitransparent display surface that reflects selected wavelengths of projected light, allowing the background view to pass through to the user unhindered. Newer systems, including those mounted on goggles or in a helmet (head-mounted displays or HMDs), may project the image directly onto the retina or use transparent display elements fused onto the visor or screen. These systems all come in the category of 'optical see-through', whereas systems that feed the background as well as the CGI image onto a display (and completely blank out the ordinary view) are known as 'video see-through'.

A wide range of applications has been suggested for augmented reality systems, from military, through medical diagnostics and training, to tourist information and entertainment. Although widely researched in the 1990s and earlier this century, the field has generally been constrained for all but niche uses by high prices, clumsy equipment and limited processing power. However, a new breed of AR is emerging as mobile phones gain sufficient capacity to capture data and create a compound image in real time, bringing data to the user rather than the user actively seeking it out.

ARTags and virtual objects

A simpler breed of augmented reality is based on basic two-dimensional barcodes, for example [ARTags](#) or CyberCode. These act as 'fiduciary markers', providing visual cues to the CGI processor about the intended location of virtual objects. Not only does the code translate into a specific image (dependent on the particular application), but the size, orientation and the visual shear (deformation) of the markers can be used to infer the position of 3D objects to be displayed. Some of these systems have been built using code that has been released under an open source licence, such as the [ARToolKit](#) from the Human Interface Technology Lab (at the University of Washington), which has also been commercialised through [ARToolworks](#).

One example of AR controlled by tags, which has been widely circulated, is GE's [Smart Grid application](#): the user prints the marker and holds it in the field of view of their webcam; a flash video application embedded in a web page generates an image of a windfarm or a solar energy plant, which the user can pan around as the printed marker is moved. Tags have been used with AR for a Nike marketing campaign in Hong Kong (as shown in [this video](#)) and in a much more sophisticated mobile [botanical field guide](#). A [Lego point of sale system](#) designed by [Metaio](#) uses the packaging as the marker, enabling the buyer to see what the Lego model will look like when it is completed, as demonstrated in [this YouTube video](#).

Gaming and AR

Inevitably, AR has links with entertainment, allowing users to explore virtual environments (such as the [Star Ship Enterprise](#)) or to play simple games (for example those from [CellaGames](#)). The Georgia Institute of Technology [Augmented Environments Lab](#) has created a zombie-shooter, called [ARhrrrr](#), as an experimental

environment for trialling AR techniques and systems. (See [YouTube video](#), which is also designed to market the graphics processing capabilities of NVIDIA's new Tegra platform.) [Novarama](#) has created a new bug-hunting game called Invizimals for the Sony Playstation Portable - children use the hardware's screen and camera to chase the virtual animals around their own homes and then lay down tags that are converted into traps on screen.

The popular 'first-person shooter' Quake was given an augmented reality makeover for playing the game across the University of South Australia campus some years ago. Although the AR environment was principally designed as a research project, the photos on the [ARQuake](#) website indicate the difficulty of commercialisation, revealing how the user was encumbered with hardware while playing the game.

Enhancing the real world

Much of the older augmented reality research was aimed at more 'serious' applications, including [construction](#), training vehicle technicians (see [this recent video](#)) and in [medical](#) fields. The Columbia University [Computer Graphics & User Interfaces Lab](#) has developed the [Opportunistic Controls](#) interface, which uses gestural controls to support motor mechanics. Zugara, a Japanese marketing agency, has demonstrated a simpler [online retail system](#) that allows shoppers to use a virtual dressing room to 'try on' clothes.

Nokia continues to develop its Mobile Augmented Reality Applications ([MARA](#)) project, with research into a broader field that has been termed '[mixed reality](#)', and Intel has used [AR to showcase](#) forthcoming mobile hardware.

The latest AR implementations rely on smartphones and other powerful mobile processing platforms to add information layers onto the hardware's video feed. [Layar](#) has used its system to deliver property information for house-hunters in the Netherlands, [SekaiCamera](#) creates a socially generated geo-tagged layer and [Wikitude](#) provides tourist information drawn from Wikipedia and other sources. New applications are appearing regularly, with a dynamic maps for the [London Underground](#) and the [New York Subway](#) system, and [TwittARound](#), an iPhone application that displays tags for people using Twitter locally. During the recent Wimbledon tennis tournament, visitors could use [IBM Seer Android](#) to discover real time tournament information and find out about facilities around the courts.

AR in education

Augmented reality has generated tantalising possibilities for education, but outcomes have been limited to date. There is considerable potential for training in manual and technical skills, showing students how to carry out procedures or which parts should be used and how they must be aligned. Tagged environments can be explored using AR applications on mobile phones - teachers and students could upload their own links and comments to be added to the virtual tags displayed on screen.

AR can create a tangible link between a marker in the hand and a 3D representation being manipulated on screen. At the lower end of the age range, early learners can use [tags to represent characters](#) from a story and see the action come alive on

screen as they move the tags around. Likewise, design students could create 3D models in a CAD package and use AR tools to overlay the results on a real scene. These design applications are coming close to market, with users able to examine [3D models of landmarks](#) built in Google's Sketchup using [ARSights](#), while Max Wainwright has demonstrated a simple AR system that is intended to work with 2Simple's online [2Design & Make](#). (The relevant part of [this video](#) starts at 10:25.)

The Human Interface Technology Laboratory New Zealand has been engaged in the [eyeMagic](#) books project, in which characters come 'alive' on the page as a child views the pages using a special stereoscopic video viewer. (An eyeMagic book is demonstrated on this [YouTube video](#).) [Futurelab documented](#) a more ambitious project, undertaken as part of the now cancelled BBC Jam service, which had children become players in an interactive story by acting out parts while carrying AR tags. Adrian Woolard, one of the BBC team, said, "You're only providing certain things like characters: they embellish it with their own stories, so they describe what the sun might look like, or whether it's a rainy day or a hot day." The same article describes a prototype system in which older learners can view an image of the skeleton and muscles overlaid on their own arm.

The Educause Learning Initiative has produced a helpful document (aimed more at higher education) entitled [7 things you should know about Augmented Reality](#).

The realities of AR

Augmented reality, despite years of research, remains a nascent technology in the consumer market. Generating information 'on the fly' is computationally expensive, producing on-going problems for high-end AR that include:

- lag time between the real world view and the superimposed virtual image
- aligning the CGI objects with the user's view
- requirements for substantial processing capability
- rapid power drain on portable systems
- heavy hardware due to large batteries
- developing a suitable user interface for each application
- potential distraction to pilots, car drivers and other users in critical situations
- total cost of systems.

A new generation of non-immersive AR that has lower hardware requirements has been growing up, such that the latest smartphones (equipped with a video camera, GPS, compass and accelerometers) are able to provide an augmented view of any landscape. These new applications draw in supplementary data over mobile networks, positioning information on screen in real time, but the markers displayed need not be as accurately placed as the navigational data on a fighter pilot's display. Developers are also exploring the potential for pointing camera-equipped mobile devices at objects to get further information, for example translating Chinese road signs or displaying show times for [films advertised on posters](#).

More ambitious AR systems, using physical markers as cues for the display of 3D objects in a real or virtual landscapes, largely remain in laboratories as experimental environments for testing human-computer interaction and generation of on-demand graphical imagery. Some educational applications (like the eyeMagic Books) are emerging and showing the potential of AR to enthuse and stimulate learners, but it remains unclear whether the learning objectives could not be equally achieved using other media or methods not involving technology.

Multimedia news

NVIDIA offers GPU support for ray tracing

Ray tracing is a set of alternative algorithms designed to add more accurate lighting and shadow to a scene rendered by a graphics processing unit (GPU). Traditionally, programs use image maps to calculate which objects from a three dimensional scene are viewable from the user's perspective. These objects are rendered as many thousands of individual triangles, determined by the brightness and colour of each surface in view, and then reprocessed by 'image shaders' and other graphics processing streams to add textures, shadows and highlights. These systems are now highly efficient, supported by thousands of processing 'cores' on more expensive graphics cards, generating real time imagery for 'first person shooter' games and virtual worlds. However, close inspection of still images from such applications reveals that reflections from curved surfaces and or other reflecting objects are often poorly rendered.

Ray tracing generates the image by beaming 'light' backwards from the eye of the viewer to surfaces in view and then following reflected beams on through the virtual image until they reach a light source. Depending on the type of surface, angle, curvature and so on, these reflected beams will take on the characteristics (especially colour and brightness) of each surface they encounter and the original light sources that would have originally generated them. However, this process becomes complex, as multiple beams must be cast 'back' from each surface to find possible primary light sources or secondary sources reflecting a primary source, introducing a level of complexity that can provide more photo-realistic results, but at added computational cost. Processing requirements are further compounded by translucent and diffusive surfaces, and projecting rays towards objects that may have no effect on the value of a particular pixel.

NVIDIA [announced](#) at the SIGGRAPH graphics conference earlier this month that its OptiX engine for real-time ray tracing would be widely available later this autumn. Like AMD's chips and Intel's forthcoming Larrabee processor, NVIDIA's high-end graphics cards contain multiple, highly parallelised graphics pipelines ideally suited to this type of repetitive task, but programming these can be an arduous assignment. The OptiX programmable pipeline is designed to simplify ray tracing problems on NVIDIA GPUs by providing a C programming interface that utilises NVIDIA's CUDA graphics environment.

NVIDIA sees the tools being used in many applications that need to model some type of radiative effect, such as optical design, acoustics and radiation research, as

well as rendering high resolution imagery. Using the GPU as a numerical co-processor potentially reduces the cost of the host system, as highly parallelised tasks can be offloaded to the GPU, while demanding less of the CPU. This type of system is more likely to be found in design and graphics departments, or in research labs. (GPU computing was covered in [TechNews 09/08](#).)

Company aims to produce laser table-top touchscreens

Light Blue Optics (LBO) [has received funding](#) to expand development on a number of projects involving holographic laser technology. Basic holograms are recorded as a set of diffraction patterns on a photographic plate produced by laser light scattered from an object interacting with a 'reference beam' that is produced by the same light source. When the plate is later illuminated by a similar source, the diffraction patterns create a three-dimensional image.

LBO's laser system does not create 3D images, but it uses a similar principle. Normal projectors selectively block a light source, using liquid crystal filters or micro-mirrors, to create an image. These systems effectively have a one-to-one correspondence between the imaging system and the pixels displayed on screen. The [holographic projector](#) uses diffraction patterns instead, creating an image through constructive and destructive interference effects as a coherent light source is reflected from a ferroelectric liquid crystal-on-silicon (LCoS) microdisplay unit. (The LCoS chip has a highly reflective layer which is revealed by controlling a liquid crystal matrix bonded on top of it. Instead of using normal liquid crystals that rely in being untwisted to control the passage of light, a ferroelectric system uses a material whose transmissive state can be 'flipped' using an electric field, which can take significantly less time.) Due to the way that each new image is generated from a unique diffraction pattern, reproducing new sub-frames from each normal video frame considerably reduces image noise and visual artefacts that might distract the viewer.

LBO claims that the system produces video with considerably greater brightness and contrast per watt of power used, compared with conventional LCD projectors. Due to the type of lens used and the construction of the image using diffraction patterns rather than a focussed light source, the image has a wide angle of projection, a very large depth of field and can be projected onto angled and curved surfaces without optical distortion. Without giving details, an [earlier press release](#) suggests that the projector could be used to convert any surface into a touch screen, presumably using sensors that respond to light scattered from the user's finger. [TechRadar reports](#) that very low power pico projector products (with a brightness of 10 lumens) with a limited 'wide' VGA (WVGA) resolution will be available from next year.

Sky to launch 3D TV channel

A number of new films have recently been released in 3D, so that users wearing polarised glasses can perceive depth in the image, bringing the scene more 'alive'. To produce a smooth 3D image, twice the number of frames must be projected with alternate frames showing the appropriate image for each eye; some form of filter placed across the screen or in glasses worn by viewers prevents each eye from

seeing details intended for the other eye. (See [TechNews 11/08](#) for further information.)

Sky has [announced](#) that users of its Sky+HD service will be able to view some broadcast content in 3D on one of its channels next year. The system will use the subscribers' existing set-top boxes, but users will need to have a '3D Ready' television. These televisions may have an array of micro-lenses across the pixels, ensuring the correct image is delivered to each eye, or synchronised shutter glasses that alternately blank the view of each eye, or a system that uses polarised glasses like in cinemas. There are no agreed standards for recording, transmitting and replaying 3D content, so users will need to ensure that hardware is directly compatible with Sky or other content delivery systems.

Sky already has some content from live events that it has recorded using 3D techniques, while current and future releases of films (both new and re-processed for 3D) will be available. However, the range of 3D resources that might be directly relevant to education - rather than entertainment - remains limited.

New liquid crystal panels save power

Screens generally use power to refresh an image, often 60 times per second (60Hz) in modern liquid crystal displays (LCDs). Normal LCDs use polarising filters and the alignment of liquid crystals to control the passage of light from the backlight, which reduces the brightness of the screen and creates difficulties viewing the image in bright daylight. LCDs also rely on capacitors to release charge to control the alignment of the liquid crystals, but, because the power is being released, regular refreshes are needed to maintain even a static image.

[Sharp has developed](#) a 'Memory-LCD' panel technology based on a different display structure that integrates memory into each pixel. The polymer network liquid crystal (PNLC) display relies on ambient light and a reflective back plate. A PNLC is similar to a polymer dispersed liquid crystal display which, without power, is milky white, as the liquid crystals are randomly arranged in the polymer network, scattering light in all directions; when a voltage is applied across a pixel, the resulting field causes the crystals to line up and allow light to pass. Power is saved by dispensing with the backlight and integrating memory into each pixel, which reduces the refresh rate to 1Hz for a static image. The effect of these changes is to reduce power consumption by a factor of 65 compared to that used by a standard LCD panel.

The current design has a one inch square (24x24mm) viewable area and only emits a single colour, but the company intends developing other sized panels and full colour versions. Sharp says that it will make designs for mobile phones, watches and other applications that need to be light-weight and work well in bright daylight.

Researchers improve LED production, create flexibility

Organic LEDs (OLEDs) use carbon-based compounds, which are often designed to be flexible, and can be readily manufactured using processes more akin to printing, making it possible to create them cheaply. However, the current generation are not nearly as bright and do not last as long as 'inorganic' LEDs (sometimes called

ILEDs) which are found in many common objects, from lighting, through instrument panels, to large public screens at sports events. Such LEDs are more expensive to produce and must be individually placed (using robotic manufacturing technology) to create dynamic displays, further increasing the cost of the device. However, with brightness around 400 times as great as OLEDs, they are more suited to many applications, especially in daylight.

A diverse team of engineers [has developed](#) a way to mass-produce LEDs and then utilise an automated transfer mechanism, which can pick them up and 'print' them onto a flexible plastic sheet, along with epoxy glue and the necessary grid of connectors used to control the display. The LEDs are arranged as a loosely connected, sheet lattice, so that the pickup 'stamp' can be configured to select those LEDs needed for a particular job. Thus, by using (say) every fifth LED for a particular application, the pickup can be used five times on the same sheet before the LEDs are exhausted. (Due to their brightness, many proposed applications, for example transparent displays, need relatively few LEDs to create an image, so they must be spaced to allow through the background image.)

A significant part of the LED manufacturing process is that the necessary layers are deposited on to a 'sacrificial' substrate. The LED layers are extremely thin (about 2.5 micrometres) and they can be made in sizes as small as 10 micrometres, as they are not cut using a saw. Instead, a square grid of grooves is etched across the wafer, polymer 'posts' are used to connect two corners of the squares and then the underlying substrate is dissolved away to leave the lattice of LEDs.

The team has only released results for production of red LEDs so far, but there is no fundamental reason why blue and green LEDs could not be produced in the same way. If the predicted substantial reduction in manufacturing costs is proven when the process is scaled up, the new technique could lead to displays 'printed' on curved and flexible surfaces, such as visors and glasses for use in augmented reality applications, or wearable health monitors. In the shorter term, it could lead to cheaper LED lighting.

Flash expected on smartphones in 2010

Streaming video and interactivity on web sites are often delivered using Adobe's Flash suite. (See TechNews articles on ['Live streaming' 05/09](#) and ['Web applications' 01/09](#).) Adobe estimates that nearly 80 per cent of web video is delivered using Flash. Although smartphones may have custom applications that directly handle video for popular sites, such as YouTube or Facebook, and third party Flash 'player' applications can be downloaded for some phones, Adobe's own Flash Player is unsupported on the main mobile phone platforms. (Adobe has developed a more limited Flash Lite 'run time' player for Symbian, Windows XP and Linux on mobile devices, while a version of Flash Player 7 is available for Windows Mobile. The specifications are [available here](#).)

Adobe [recently announced](#) that Flash Player 9 will be installed at manufacture on the up-coming HTC Hero mobile phone, which uses the Android operating system, although [this means that](#) interactive applications will only be able to use ActionScript

2. It has also been reported that a beta (test) programme for Flash Player 10 on Android, Symbian, Windows Mobile and webOS phones will begin in October. The notable exceptions are phones from Apple (iPhone) and RIM (Blackberry), [possibly due to](#) technical difficulties or security concerns. Flash Player 10 is the version currently available for most PC operating systems and web browsers.

[TechNews 07/09](#) reported that Flash would be coming to Symbian, Windows Mobile and Palm's webOS smartphones from October this year. Reporting on an Adobe analyst event, [ReadWriteWeb](#) has said that Adobe's platform for creating rich, interactive internet experiences will recognise touch screen controls, GPS and accelerometers in mobile devices in the near future. This will allow augmented reality applications to run through Flash Player, using the position and viewing direction of the phone to overlay information tags and images across the phone's live camera display. (See [TechNews 07/09](#) for an explanation of Augmented Reality.) Providing programmers with access to these features on completely different phones would allow developers to create Flash applications that run using common code on all the supported hardware platforms, considerably reducing development costs.

Applications written in Flash (or Microsoft's competing Silverlight platform) and other web technologies can be used to support rich learning experiences that include video and custom interaction. Since such applications run in the browser, they can be used by learners at home or in other locations. If these experience can be ported to mobile devices, then the vision of 'anywhere, any time, any device' learning will be much closer to reality. The future HTML 5 standard will also provide much enhanced media and interactive capabilities. Microsoft Silverlight was recently upgraded to version 3 ([TechNews 07/09](#)).

In related news, Adobe [has released](#) an open source media framework (OSMF) for its Flash video technology. While this does not mean that the company has 'open sourced' the main Flash Player browser plug-in, which will still be required for playback, it will make it simpler for developers to control flash media assets using custom players. This would allow companies to extensively 'skin' the playback environment with their own controls and branding.

Video streaming over HTTP proposed

Video embedded in web pages is normally handled by third-party plug-ins which take control of that part of the display and stream video using separate protocols, such as the real time protocol. (See the 'Live streaming' article in [TechNews 05/09](#) for further details of some of the available protocols.) Although RTP is efficient, it uses client ports requests that may be blocked by firewalls. However, HTTP (hypertext transfer protocol) has to be permitted for any normal internet communication to occur, so using it would overcome these problems. This could open up large amounts of video to educational establishments, although filters would still be able to block video streams on a site-by-site basis.

[Ars Technica reports](#) that Apple has submitted video streaming protocols, based on HTTP, to the relevant standards body, the Internet Engineering Task Force (IETF). A server using HTTP Live Streaming would take a video file created using the MPEG-2

codec and break it into a large number of smaller temporary files defined by a '.m3u8' playlist. The client computer would then download each file using HTTP and plays them back in the predetermined order. To create a live stream, these playlists would be chained together, with a new location tagged to the end of the active playlist. The system would involve delays - Apple illustrates the schema with a latency of ten seconds - due to capturing and 'chunking' up the stream into the required files. If the client is unable to access the streamed content sufficiently quickly, the protocol allows for pointers that would provide alternate versions of the stream that use lower resolution or frame rate.

Apple has indicated that these streaming techniques would be used for web-based audio and video playback using its forthcoming QuickTime X software in the new 'Snow Leopard' version of its MAC OS X operating system, as well as on the iPhone's update to OS 3.0.

In related news, [PC Pro reports](#) that the forthcoming HTML 5 standard will not embed an agreed set of video standards. Most streamed video now comes through Adobe's Flash Player or Microsoft's Silverlight plug-in, but many in the working group for the next version of HTML wanted to define the video protocols that should be associated with a new <video> tag. (Hypertext markup language, HTML, uses tags to control layout and define non-textual elements on web pages.) Agreement would have mandated future browsers to support the codecs directly. However, lack of consensus means that the new tag will still link to video, but browsers will have to continue to use plug-ins, unless their developers decide to provide native support for a particular codec.

Print directly from the internet with web-connected printer

HP [has released](#) details of a new printer that features an integrated internet connection and a touch screen which controls applications that access content from the web. For example, users can select content streams from USA TODAY, which will be updated in real time, and print items of interest whenever they wish. Likewise, support for Google Maps and Calendar allows consumers to request details using the touch screen and print directly without needing a PC. New applications linked to other content can be downloaded from the HP Apps Studio. External devices, including desktops, notebooks and some models of smartphone, can be connected to print to it using cabled, Bluetooth or Wi-Fi connections.

This printer exemplifies a continuing trend towards converged, internet-connected hardware. Just as smartphones integrate many services, so other manufacturers are looking at ways to draw new consumers to their devices by adding functionality available elsewhere. Most households that might buy such a printer are likely to already own at least one PC, so it remains to be seen whether consumers will want to use a 4.3 inch (11cm) touch screen to carry out tasks that would otherwise be handled using the computer connected to the printer.

The HP Photosmart Premium with TouchSmart Web is an all-in-one PSC (printer-copier-scanner) that also has fax capability and has been given a US target retail price of \$399 (about £250) for its launch in the autumn.

New information tag to compete with barcodes

Barcodes represent information in machine-readable patterns. The type of pattern familiar to most people is 'linear' because there is one set of information that can (effectively) be scanned in a horizontal direction. This limits the amount of information stored in the barcode, unless the pattern of vertical lines is greatly extended. Newer '2D barcodes', including Data Matrix, High Capacity Colour Barcodes, QR Codes and Shotcodes, have multiple rows of information and are processed as an image, rather than using the common laser scanners seen in supermarkets. (See 'QR Codes and other 2D barcodes' article in [TechNews 06/09](#).)

Researchers at the MIT Media Lab have created another new form of data tag, which they are calling the [Bokode](#). The bokeh effect in photography results in bright objects not in focus creating a blurry circle in the viewfinder, which resolves into a clear image as the camera is focussed at the correct distance. Bokodes use this effect to create a small (3mm), bright data 'spot' on an item that is out of focus to the human eye (which focuses on the surface on which the Bokode is mounted) but which digital cameras resolve and magnify because their focus is practically at infinity (unless an object is very close). The Bokode uses a low-power LED light source, covered by a data mask and a microlens that creates the necessary optical effect. The mask is based on a 10x10 array of Data Matrix codes that store the actual information.

Bokodes can store data in a very dense array - the authors say they have detected 2.5 micrometre patterns from over 4m (13 feet) using an ordinary digital camera. Due to the optical nature of the system, it is possible to process the camera's image to calculate its distance and angle from the Bokode, and to read more information as the camera approaches the tag. This would permit developers to build augmented reality applications that produced a 3D object on screen that would be seen in the appropriate zoom and perspective as a user moved towards and around the tagged item. (See [TechNews 07/09](#) on Augmented Reality.) Bokodes could be used to overlay information on a mobile phone's display as learners moved around a historical site or museum, as a way to identify books in a library or as markers when playing live simulations, where the tags would give information about the players.

Bokodes are unlikely to displace barcodes in the short term due to the requirement for a powered LED light source and a good quality lenslet - a [BBC report](#) says that a single code costs around \$5 (£3) to produce. The researchers also need to overcome physical limitations in the system, such as the 20 degree 'field of view' of the current prototype.

Selected readers to see first video ads in print publication

Entertainment Weekly, a show business magazine from the US, [is to circulate](#) the first ever 'video-in-print-media' insert advertisement to selected subscribers. The video will be displayed on a 320x200 pixel screen that is just 2.7mm (0.1 inches) thick. In a similar way to greeting cards that play tunes, when the page is turned the display will start to show clips from the forthcoming CBS autumn schedule and advertisements from Pepsi. Buttons provide control for the presentation and a rechargeable battery powers the 40 minutes (maximum) of playback that can be

stored on the embedded chip. The reinforced screen is based on the standard thin film transistor (TFT) liquid crystal technology used in many laptops and the battery can be recharged using a mini-USB connector integrated in the device.

[TechNews 09/08](#) reported on a simpler moving display advertisement, due to be published in Esquire magazine, based on electrophoretic 'e-paper' technology from E Ink. (See the 'E-paper and e-books' article.) The new hardware will support a much higher frame rate - as TFT screens can be updated far faster than e-paper - allowing true video playback. Like rotating billboards in public spaces and the large video displays now found commonly on the London Underground and elsewhere, the current application is really an advertising 'gimmick' designed to catch readers' eyes and draw their attention to specific content of interest to the advertiser. Nevertheless, it demonstrates the degree to which the technologies used have been optimised in terms of both size and power consumption. Similar developments are producing better screens for mobile devices, but it would be unlikely for learners to find this technology used for video in 'traditional' textbooks due to the cost and the availability of more suitable display devices.

Hardware

Analysis: Visualisation technologies

At a glance

- Teachers often need to demonstrate skills and show images to classes, but technologies like OHPs cannot project opaque objects
- Visualisation technologies enhance traditional demonstrations through capturing video, providing fine zoom facilities and reducing key physical risks
- A visualiser (a video camera mounted on a flexible arm) connected to a projector puts live video of any object onto a screen for large groups to see
- Other specialist technologies, such as digital microscopes, and video connectors to standard technologies, such as webcams, smartphones and calculators, can be used as alternatives for some tasks

Visualisation

'A picture is worth a thousand words' underpins the educator's desire to present relevant imagery to learners, but the practicalities of providing such pictures without large numbers of text books or considerable (and too often illegal) photocopying is another issue. Visualisation technologies allow a whole class to see what is on the teacher's desk without having to gather at the front. Applications need not be limited to simply displaying static images, but provide opportunities for learners to see the detail of scientific experiments, to observe practical skills, to examine real objects and see exemplars of work being undertaken.

Gathering learners around the teacher's desk introduces health and safety risks, especially where dangerous chemical reactions are being demonstrated; yet pushing students back to their seats limits their ability to see the details of the experiment. Visualisation technologies throw the image up onto a screen or whiteboard,

potentially providing a better view than assembling around the apparatus. Visualisation technologies can reduce disturbance caused by learners moving around the class, prevent crowding at the front, give enhanced views through zooming or image colour adjustments, and provide a permanent record of the work undertaken if connected to a storage device.

Visualisation technologies

The overhead projector (OHP) became an essential tool to many teachers, but it can only deal with transmissive images and not opaque objects. OHPs have generally been superseded by interactive whiteboards (IWBs) and other projection technologies, but IWBs do not come with video input as a standard part of the equipment. A webcam can be connected to the host PC, with any image displayed dependent on the quality of the camera, the functionality of the supplied software and the stability of the camera mount.

Earlier technologies, especially the epidiascope (or 'opaque projector') put live images onto a screen, but hot bulbs could damage delicate specimens, they were expensive and there was no means to store live demonstrations.

Digital microscopes and certain mobile phones can be connected directly to a PC to provide some visual imagery, but one of the most flexible solutions that has been emerging over recent years has been the visualiser. This hardware is, in essence, 'a camera on a stick', but that description does not do the technology justice.

Educational applications

Visualisation technologies can be used in many ways, some of which are written up as case studies on the [Visualiser Forum](#), which is run by a consortium of educationalists and vendors. Applications include:

- demonstrating scientific experiments
- examining objects and documents, such as historical artefacts, illustrations in books, natural specimens and maps
- illustrating practical skills, for example fine brushwork, sewing, assembling electronic components or correct use of a protractor
- modelling writing by developing a collaborative structure, inserting grammar and punctuation or adding in detail like adjectives
- displaying other devices, such as a calculator, games console or phone
- facilitating peer evaluation and assessment, by putting learners' work on screen for all to see
- assisting vision impaired pupils
- supporting meetings by displaying official documents, curriculum plans or policies for discussion.

Teachers TV provides videos that discuss the use of visualisers, including [Hard To Teach - Secondary English Using ICT](#) and [Resource Review - Primary ICT Kit](#).

One issue teachers must consider before displaying items is copyright, which is not an area that the Visualiser Forum can provide guidance on at this time.

Visualisers

A visualiser is a combination of a video camera on an 'arm', a light source, a base and (optionally) a platform on which to place objects and documents. To be used effectively, it needs to be connected to a projector system, although some models now come with storage cards so that captured video can be saved for later viewing. The hardware features and software functionality vary between models, but the following could be considered during procurement:

- Image sensor and lens quality, including (as for a digital camera) the number of megapixels that make up a single image and the depth of field (focus limitations)
- Hardware and software controls for camera zoom, fine positioning, lighting, image capture, white balance and similar
- Flexibility of the arm for positioning the camera - will it be important to capture images (such as posters) on walls or for videoconferencing?
- Connectivity for output to a display, either directly or through a PC
- The need for a platform, which may help give scale to objects but can reduce flexibility
- Software features, such as live presentation; image capture and editing; split screen viewing (to compare to a previous version or to a recorded demonstration); time lapse capture; and stop-motion animation capabilities
- Weight, if the device is to be moved around.

The price of a reasonable quality starter model should be around £275, rising to over £1,000 for more advanced hardware. Tom Cooper, Strategic Leader ICT for the London Borough of Lewisham, compares the cost of a visualiser (with sound system and projector) at £1,500 to £2,000, with that of an IWB at over £3,000 (including installation), and considers that learners will derive far more benefit from the former compared with the passive use commonly observed for IWBs. Visualiser vendors include [AVerMedia](#), [ELMO](#), [Genee Vision](#), [Lumens](#), [Samsung](#) and [WolfVision](#).

The basic concept of a visualiser is undergoing little development, but manufacturers are improving hardware to provide HD resolutions, data storage on USB memory sticks and SD cards, and wireless connectivity. There have also been improvements in software integration through increased use of TWAIN (a set of protocols that allows a wide range of Windows and Linux software to access features of imaging hardware) and introduction of visualiser 'buttons' on IWB interfaces, allowing direct access to the device.

Alternative technologies

A number of specialist visualisation technologies exist, although visualiser cameras are generally suited to capturing the screen of any handheld device. Intel Play QX3 microscopes were given to all maintained primary and secondary schools in England as part of Science Year in 2001/2 and the upgraded QX5 is now produced by [Digital](#)

[Blue](#). Various companies, including [Veho](#), produce alternative USB microscopes that range in price from £30 to £100, or considerably more for 'professional' devices. Enthusiasts point out that many visualisers are adequate for general close-up work and that it is often possible to focus the camera on the eyepiece of a standard optical microscope.

It is possible to connect some smartphones to a television, to a projector, or to a computer which has suitable inputs on its graphics card. This setup can offer live video or recorded images on an *ad hoc* basis, but may not provide the ease of use of a visualiser, and is limited by the capabilities and battery life of the phone. Calculator manufacturers, for example Casio and Texas Instruments, provide connectivity or emulators for their more complex scientific models, but these are specialist applications.

Uptake and embedding

Dave Smith, ICT Consultant for the London Borough of Havering and Chair of the Visualiser Forum, believes that visualisers will become widespread now that prices have become more affordable. The hardware has seen greater adoption in Japan and the US, where IWBs are less prevalent, compared with Europe. [Futuresource Consulting predicts](#) that the number of UK classrooms equipped with a visualiser will rise from one in thirty last year to one in ten by the end of 2010. Many local authorities have focussed on procuring IWBs, but some London Boroughs, including Barking and Dagenham, Havering, and Lewisham, have directed investment towards visualisers.

Mr Smith believes that many teachers who are wary of technology find visualisers more intuitive to use than IWBs and that they have a greater impact on learning. As such, they assist in a personal as well as a technological transition from traditional to digital tools. Tom Cooper sees the technology as an interim between analogue projectors, like OHPs, and fully digital systems in which learners and teachers will collaborate, create digital artefacts and share them electronically. However, Dave Smith believes that paper will remain prevalent in schools for another fifteen to twenty years, providing a continuing, important role for visualiser technologies.

Hardware news

Server and mobile processor update, September 2009

The Hot Chips conference at the end of August saw launches and updates on many of the server processor platforms. IBM gave [further details](#) of its forthcoming 8-core POWER7 processor package, which supports four instruction threads on each core, giving 32 concurrent execution threads available in each socket on a server's motherboard. To ensure that the cores have plenty of data to work with, IBM is providing 32MB of cache on each chip and a pair of 4-channel DDR3 memory controllers. (DDR3 is a faster memory implementation than DDR2, which is still used in many PCs, because twice as much data can be fetched on each clock cycle.) The cache memory is formed from embedded dynamic RAM (eDRAM), which has a lower processor count per megabyte than the SRAM used by competitors, making the whole chip package smaller and (presumably) more power efficient. The

POWER7 processor, which will also have 4 and 6-core versions, is [expected to debut](#) in the middle of next year.

Sun, which is being purchased by Oracle, [gave few details](#) of its 16-core 'Rainbow Falls' processor, which may be released as the UltraSparc T3. One notable feature is that the package has cryptographic algorithms embedded in hardware, making secure communications and storage much faster than doing the same calculations with software. Fujitsu also described its 8-core, 2GHz Sparc64 VIIIfx processor.

These manufacturers are producing servers that run Unix variants or Sun's Solaris operating system; Windows and Linux products use the x86 instructions originally developed by Intel. Due to the size of the latter market, which includes around 90 per cent of desktops and laptops, x86 processors cost less but are not as highly optimised for server applications. AMD provided [further information](#) on its 'Magny-Cours' processor, which will feature 12 cores by combining two 6-core 'Istanbul' Opteron server processors in the same package. These chips will run somewhat slower in order to fit the same power envelope as the Istanbul designs, but faster cache memory and the additional cores should increase overall performance. AMD expects to release the new processors in the first quarter of next year.

Intel's processor strategy relies on its Nehalem architecture in most of its latest product lines. (See [TechNews 04/09](#).) 8-core Nehalem-EX processors - aimed at the server market - [will appear](#) later this year. Compared with the 12MB of Level 3 cache on Magny-Cours, Nehalem-EX will feature 24MB of L3 cache. (Cache memory holds data 'nearer' to the processing cores, avoiding time-wasting fetches from the computer's main memory across congested memory buses.) The current architecture can support up to eight processors (in other words 64 cores) that communicate using Intel's QuickPath Interconnect technology. However, IBM's POWER7 can be installed in systems with as many as 32 sockets, giving a total of 256 cores on a single server 'blade'.

These powerful processor systems will be used in corporate data centres and at the heart of some of the larger educational networks. As many physical servers are 'virtualised' into fewer large systems, the processing load of each server will increase but the range of tasks will remain largely the same. (In this context, virtualisation describes the partitioning of resources on a single physical server into separate 'spaces', each apparently running their own operating system and services.) Hence multi-core systems are expected to give the same performance, or better, running many processes in parallel, but saving power by reducing the amount of physical hardware and running processors that are far more optimised for power consumption than previous generations.

The trend towards virtualisation, as well as the credit crisis, may explain some of the [30 per cent fall](#), year-on-year, in server sales. [Reports](#) say that TheInfoPro has estimated that 30 per cent of all new servers were virtual installations in 2008, predominantly in the x86 environment, and expected to rise to 80 per cent by 2012.

Intel also gave further details of its processors for mobile computing at the Hot Chips conference in August. The main forthcoming mobile platform is codenamed 'Moorestown', which will be powered by the 'Lincroft' line of SoC Atom processors. (System on a chip, or SoC, puts the graphics processor, memory controller and other key functions on the same chip as the main processor core.) Due to fine-grained power control, Intel [claimed](#) that Moorestown will reduce idle power consumption by 50 times compared to previous 'Menlow' designs, although the company gave no details of peak power drain. Depending on workload, each sector of the chip can have its own power settings and clock speed, although transition between power states can introduce latency (delay) as the processor adjusts.

Intel is aiming Moorestown at the 'high-end' smartphone market, as well as MIDs (mobile internet devices). The Atom processor is largely competing with ARM derivatives in the mobile space, for example those from Freescale, Texas Instruments and VIA Technologies. Atom chips may be more powerful and based on the same 'x86' instruction set used on desktop machines, but this comes at the expense of power drain and, with poor battery life a regular gripe from smartphone owners, this may deter manufacturers from using the platform in new handsets. MIDs, which are small, multimedia, internet-enabled handheld devices, have only recently come to market, so their popularity has yet to be proven.

The first of Intel's new dual-core 32nm designs, 'Arrandale' (mobile) and 'Clarkdale' (for laptops and desktops), will be shown at the Intel Developer Forum (IDF) later this month. Some details of Arrandale [were given](#) at Hot Chips, including the fact that a 45nm-based combined memory controller and graphics processor will be integrated into the same package, although not on the same chip as the processor. Reduction in transistor size to 32nm for the main processor should increase speed while reducing power consumption, as the electrical signals have shorter distances to travel. Likewise, putting other key components in the same package will take that traffic off the main data bus out of the processor socket while improving speeds, due to these components being in close proximity to the main cores.

Creating processors at nanoscales

Moore's 'law', which suggests that the number of transistors packed into a given area of a chip will double every two years, has proven remarkably robust over the last forty years. (With other advances, this leads to an approximate doubling of processor power every 18 months.) However, increased transistor density requires etching techniques that can produce increasingly fine lines at nanoscales, in turn necessitating highly controlled light sources that use ever shorter wavelengths to alter the materials that are to remain as a result of etching. Current systems use deep ultraviolet light, which has a shorter wavelength (under 300 nanometres) compared with the visible spectrum, but producing and controlling shorter wavelengths is increasingly difficult. Eventually, photolithography will reach fundamental limits due to absorption of light by the optical elements and other parts of the equipment; X-ray lithography and electron beam lithography are among proposed alternatives.

Researchers at Purdue University in the US [are investigating](#) extreme ultraviolet lithography, which uses a wavelength 13.5nm. Their approach involves creation of a plasma (a partially ionised gas), which contains free electrons, that is heated by 'pinching' using magnetic fields, so that the electrons produce high energy photons of extreme ultraviolet light. However, this process is difficult to control, as the mirrors used to manipulate the resulting beam become coated with condensing plasma and the energy yield is currently only one to two per cent of the 100kW (or more) that would be needed to power the system. The team is using computer simulations to try to improve the yield.

Scientists from [IBM](#) and the [California Institute of Technology](#) are developing techniques for etching patterns smaller than 22nm by using DNA 'scaffolding' to organise the deposition of carbon nanotubes. Long strands of viral DNA are manipulated into the required patterns using short strands of synthetic DNA, but the process is carried out in a saltwater solution that deposits the resulting folded 'origami' shapes randomly across the substrate. The IBM team was able to create negatively charged nano-scale patterns on the silicon wafer using electron beam lithography, which, with free positively charged magnesium ions added to the solution, caused the DNA patterns to be precisely deposited. Once in place, the DNA structures can be used to organise the carbon nanotubes and other components required, with spacing down to just 6nm. (The most advanced production technologies use 32nm or more, with 22nm on the horizon.) According to [reports](#), actual processors manufactured on this basis are eight to ten years away, while Intel engineers are hopeful of breaking the 22nm barrier using conventional lithographic techniques. Nevertheless, production equipment for self-organising DNA-based manufacturing could be considerably less expensive.

Processors based on smaller components will have greater processing power but produce less heat and use less electrical power. Exchanging silicon for carbon-based nano-materials could also improve speed and reduce current leakage, especially at these scales. Such processors would significantly increase the capabilities and battery life of smartphones and all kinds of other computing devices.

Spasers could underpin optical computers

Optical (or photonic) computing promises to be extremely fast and energy-efficient, since information will travel at the speed of light without the heat loss caused by electrical resistance. However, besides the problems of converting and processing data using optical methods, the size of normal optical components is dictated by the wavelength of light - around 500nm for the visible spectrum - compared with silicon components which will soon be in production at scales of 32nm. For some time theory has predicted that photons could be produced (without standard optical components) using surface plasmon lasers, or 'spasers'.

Lasers create spatially coherent light over a very narrow range of wavelengths through electrical or optical excitation ('pumping') of a source that emits photons. These photons are then emitted through a carefully constructed optical system that ensures that the beam is in phase (with the wave patterns in 'step') and spatially aligned. Surface plasmons result from optical excitation of the electrons on the

surface of a metal, producing clouds of electrons that create mobile field effects that can further excite a photon source close to the metal's surface. [Scientists](#) from Cornell, Norfolk and Purdue Universities in the US created 44nm gold spheres coated with two ultra-fine glass layers. As described by [Ars Technica](#), when excited by a pumping laser, dye in the outer layer emits photons with a green wavelength that then interact with surface plasmons on the gold sphere to amplify the production of photons, some of which escape as pulses of light.

It may be possible to make spasers as small as 1nm, creating considerable potential for integrating them with silicon circuitry for ultrafast computing devices. (While current processors can only work in the gigahertz range, optical computers have the potential for operating well in excess of a terahertz, or 1,000 times as fast.) However, a great deal of work needs to be done on other aspects of storing and processing photons before optical computing can become a reality, as well as reducing the power requirements of the pumping laser. In the meantime, spasers could be used as photon sources for extremely high resolution microscopy and other applications.

A second design for spasers has been [described by researchers](#) from UC Berkeley in the US, using cadmium sulphide nanowires.

New transistor design cuts charger power loss

Power supplies and chargers for computers, mobile phones and many other devices dissipate 30 per cent or more of their energy as waste heat, while using power even after the main hardware is switched off. (Unless the power is cut to a basic transformer, it will continue to use power even while apparently idle.) Where computers are clustered in a limited space, the wasted heat often leads to further power being consumed by cooling equipment.

[Fujitsu in Japan claims](#) that a new approach to constructing power supplies could reduce energy loss by a third. A large inefficiency in power supplies is the introduced by the silicon-based transistors that are used to reduce the voltage of the alternating current (AC) supply before it is converted into direct-current (DC). Fujitsu is developing hardware based on gallium-nitride high electron-mobility transistors (GaN HEMTs) which can switch more rapidly while using as much as a fifth less power. GaN HEMTs have been used before in other applications, but previous designs have only produced a switching voltage of 0.5V, rather than the 2V to 4V required in a power adaptor. Using new deposition and etching techniques, Fujitsu has achieved the required voltages in transistors that can switch more rapidly, producing higher switching frequencies, which could also lead to power adaptors that are 'one-tenth current sizes'. The new design also enables 'complete current-interruption in stand-by mode'.

In related news, [it is reported](#) that most of the main mobile phone manufacturers have agreed to support the EU's 'universal charger' proposals. (See [TechNews 03/09](#).) From 2010, chargers will start to appear that use a micro-USB connector which will be common to all compliant phones.

Manufacturers close to wireless charging standard

[TechNews 01/09](#) carried details of electrical chargers that transfer power through a coil embedded in a mat, using the principle of induction. A device (or a charger dock) - also containing a coil - is placed on the mat; as the alternating current generates and then rapidly destroys an electromagnetic field, the receiving coil gathers this as energy and uses it to charge the device's batteries.

The Wireless Power Consortium is [now close](#) to an agreed standard for this type of technology, intending to test the interoperability later this month. The group will use the Qi (pronounced 'chi') trademark to badge compatible devices, with reference to the Chinese symbol's meaning of 'vital power'. The standard covers equipment with a power rating up to 5W, such as mobile phones and MP3 players. The consortium includes household names such as Duracell, Olympus, Philips, Samsung, Sanyo, Texas Instruments and other brands.

Duracell has [announced plans](#) for myGrid charging mats to be released in the US later this autumn. The range is [said to be](#) a rebranded version of [WildCharge](#)'s products, with no indication that they fall within the Qi scheme. The suggested retail price for a pack with the charger and connectors for many Blackberry, Motorola and Nokia models will be nearly \$80 (around £50), while additional connectors are recommended to be about \$35 (£22). No UK release has been announced. Competing products come from other companies, such as [PowerMat](#).

In related news, Intel has [further demonstrated](#) power transfer using its Wireless Resonant Energy Link (WREL) technology - a 60W bulb was lit at a range of several feet with 70 per cent power efficiency. A 10MHz radio signal is broadcast between two coils, with the slave coil resonating to the frequency of the transmitted signal. However, as Intel admit, this high frequency 10MHz radio signal (based on alternating current, AC) will need to be converted to a stable direct current (DC) supply - suited to most mobile hardware - without significant loss of energy during the process. Intel is also considering the regulatory and safety issues involved. An [Intel video](#) shows the product in action, transferring an audio signal.

New approaches to improved touch screens

Touch screens are now expected on many types of handheld device, although they have been commonly available on PDAs and tablet PCs for over ten years. Learners may find touch screens easier to use on handheld devices, especially as more complex web applications are developed that require types of input other than through a keypad. The main types of screen available fall under one of two technologies:

- Resistive screens, which sense a point of contact that is formed as two surfaces are pushed together, allowing a current to flow. Many types of object can be used to make contact, although precision will depend on the area of contact made and the spacing of the conducting elements (such as a grid of ultra-fine wires). Styluses are generally provided such devices to allow for fine interaction, although there is a risk that they will be lost.

- Capacitive screens sense contact from a finger, which changes the surface charge where touched. Capacitive devices require some form of conductive object to drain the charge at the point of contact - fingers work well, but most styluses fail as they are plastic tipped to avoid damage to the transparent electronics printed on the surface of the screen. Capacitive systems can respond more rapidly than resistive devices and can detect multiple simultaneous touches.

Fingers can be much easier to use - and do not get lost - but they are limited when it comes to fine detail, such as handwriting recognition or manipulating small icons. HTC [is reported](#) to have applied for a patent covering embedding a small magnet in the head of the stylus or manufacturing it from magnetic material. As the pointer is moved across a capacitive surface, a minute current is induced, which can be detected using the screen's standard electronics. As well as reducing the point of contact, the use of a stylus means far less of the screen is obscured, compared with using a finger. Although this system would be cheap to incorporate with a capacitive screen, it would still be vulnerable to loss of the stylus, allowing continued finger input but limiting work requiring fine detail.

ZDnet [also reports](#) on a patent application by RIM (manufacturers of Blackberry phones) for hybrid devices: a normal resistive screen is manufactured to which the transparent electronics required for a capacitive screen are added. The latter can detect a person's immediate proximity, as charge will leak across a small air gap, enabling the phone to turn off touch functions when it is put near the user's face. A single controller is used for both systems, reducing the overall cost.

In related news: Touch screens on most internet tablets and mobile devices can only track one or, at best, two fingers, limiting the range of gestures that can be made and the way that the user interacts with the device. Multi-touch appears to give more intuitive control over devices, but programmers have still to provide many examples of how it can be used in a meaningful way, especially in an educational context. (The [TechNews 06/09](#) article on 'Gesture control' gives some more detail on touch devices.)

Synaptics [has released](#) details of a new ClearPad 3000 Series of touch sensors for handheld consumer electronics devices, which can recognise up to ten simultaneous touches. The sensors have a touch resolution of under a millimetre, are less than half a millimetre thick and can be a maximum of 20cm (8 inches) across the diagonal. Synaptics envisages that the main initial attraction of these capacitive sensors will be for games designers. Touch 'tables', such as Microsoft's Surface, already support multi-touch input, but there are few alternatives for handheld devices. Engineering samples will be available towards the end of this year.

Darkfield tracking enables optical mice on glass

Optical mice work by taking pictures of the surface they are placed on and tracking movement of textural features. The surface is lit by an LED or laser diode and image processing software embedded in the mouse monitors changes to pass back to the computer as movement information. However, glass and highly polished surfaces

give few movement clues to most optical mice, so they are not able to pick up relevant tracking information.

Logitech [has launched](#) a new series of mice that use 'dark field' tracking when on difficult surfaces. Scientists employ dark fields in microscopy by using slightly oblique lighting to show up salient features against a dark background. The Logitech mice use similar techniques, looking for minute scratches and flaws in the glass surface that reflect light back. The Darkfield Laser Tracking mice [have two lasers](#) to illuminate the surface, using standard tracking algorithms when on a textured surface and dark field methods on glass. Logitech [suggest](#) that the mouse can pick up details down to a fifth of the width of a human hair, so long as there are at least 44 such details per square millimetre. Such mice would be an advantage to many users, 40 per cent of whom the company found have glass tables in the home, although the mice will not be as responsive for gaming applications.

Assistive technology for the blind, deaf

Researchers [are reported](#) to have developed the technology for an active, full page Braille display. Most of the commonly available Braille displays use a single line of 40, 65 or 80 Braille cells (3x2 patterns of dots) that work by mechanically raising a pin through a hole to represent each dot. Blind readers scan the text by swiping their fingers across the row of cells, sensing the individual patterns as the letters that form words, and then refreshing the row as each line has been read. These machines are expensive and have little direct relationship with modern computer displays, although they function well for services similar to instant messaging.

Electromechanical systems for full page Braille displays (40 cells by 25 lines) would be very expensive, as the technology must be extremely robust and reliable. The new system uses electroactive polymers (EAPs) to create the dots. These polymers deform when a voltage is applied across them, leading some to refer to them as 'artificial muscles'. In the Braille application, a sheet made from a dielectric elastomer is printed on both sides with a pattern of electrodes; when a voltage is applied between a selected pair of electrodes, the EAP sheet deforms, causing a bump to be raised, which is then released when the voltage is removed. Due to basing production on sheet technology, rather than individual mechanical actuators, the new display has potential to be much cheaper, as well as being able to represent more graphical information by using the whole page. Systems based on EAPs could also be used to underpin haptic (tactile) feedback mechanisms for other applications. (See brief [TechNews 11/08](#) article on haptics.)

In related news: Deaf-blind people experience considerable limitations in communicating in the daily contexts that most of us take for granted - unable to see gestures or to hear questions, it is difficult to know that you are being addressed or that a response is required. 'Print-on-palm communication' techniques can help, where someone spells out a message on the palm of the deaf-blind person's palm, but this process can be awkward and slow.

HumanWare has launched the [DeafBlind Communicator](#) (DBC), a handheld Braille device for the deaf-blind user paired over a Bluetooth link with a specially adapted

mobile communicator that can be handed to bus drivers, shop assistants or whoever the owner wishes to speak to. The respondent both sees and hears a message that informs them of the owner's disability and invites them to type responses via the communicator unit, which the owner then reads using the Braille unit. DBC units can be linked to each other (or to any 'teletype' device) via landline, and the DBC can also be used to send standard SMS messages using the phone capabilities of the communicator unit. The Braille unit also contains a word processor, a browser and other applications, allowing users to write documents and communicate over instant messaging services. A [BBC News report](#) suggests that the system costs £4,400 and £5,400 depending on the model selected.

This device provides deaf-blind people with a communications system that uses component that are readily understood by the people they wish to talk to and could greatly assist learners in talking to teachers and other students.

The Digital Rosetta Stone to preserve data

Data is already stored on a multitude of different media using a considerable number of formats - few new PCs are now equipped with any kind of floppy disk drive, while the old 12 inch laser disk readers needed for the BBC Domesday Project (see [Wikipedia](#)) are now all but unobtainable.

A [BBC News report](#) gives details of a project that aims to create a digital archive that would remain readable for at least 1,000 years. The reports suggests that the life of data on a CD would be about 15 years and a DVD is even worse at 10 years. While this may make current optical disks adequate for everyday back up, these time scales are totally inadequate for long-term archives recording the knowledge, culture and history of our present civilisation. Although vulnerable to intense magnetic fields, data on hard drives stored in controlled conditions may take 40 years to degrade to a state where the content is unreadable. Silicon is much more stable, such that data maintained in an archive kept at less than two per cent humidity could provide the solution.

A research team from Keio University in Japan are testing masked ROM (MROM) chips in a stacked array to create the [Digital Rosetta Stone](#). MROM uses a 'write once' process that employs an electron beam to etch the necessary patterns into silicon through a mask that represents the data. Data is written to a stack of wafers (comprising a grid of MROM chips) and sealed into a package that contains controllers, capacitors to store power and data registers. The data is then read using an induction process that transfers energy by generating electromagnetic fields - a reader device is put alongside the data 'Stone' to transfer power and control commands, and then detect the required output following a predefined sequence. The stack of four 15 inch (38cm) wafers can hold 2.5TB (2,500GB) of data that can be read at 150Mbps.

New installation standards for audiovisual equipment

Becta has worked with partners to produce a new standard for the installation of audiovisual equipment. Approved by the BSI as [Publicly Available Specification \(PAS\) 122:2009](#), it covers the process of installation (pre-installation advice, site

surveys, actual installation, commissioning and signing off) and the relevant health and safety issues. As the preamble says, PAS 122 has been 'written to give end users and customers a level of confidence and assurance when purchasing audiovisual (AV) installation services... It has been noted over the past number of years that resellers of interactive whiteboards and AV equipment use third party sub-contractors to install their equipment. It has become apparent following completion of a number of audits that there are no strict guidelines relating to how the installations are carried out and that the processes employed by suppliers and subcontract staff differed greatly.' Appendices are provided with sample documents and further advice to installers and users.

The standard is designed to cover a wide range of AV installations, but it would be particularly relevant to institutions about to install whiteboards or large display screens. PAS 122 is available from BSI, priced at £85 for members.

Software and internet

Analysis: Speech recognition

At a glance

- Speech recognition enables all kinds of voice interaction with computers, but the technology has proven more difficult to implement than many expected.
- Learners could dictate essays, control machines and (where needed) use assistive technology, if voice input were more prevalent. However, students would need guidance in effective use of such systems.
- Speech recognition is hindered by regional accents, background noise, poor telecommunications and other factors, in addition to the complexities of deciphering speech itself.
- Voice 'tags' on mobile phones and voice commands for machinery and industrial applications can work well, as they use a restricted set of inputs.
- Transcription software and online services are gaining some ground, although questions remain over accuracy and the need for human intervention.

The promise of speech

The idea of speech recognition - to replace dictation or used to control devices - has been something of a computing 'holy grail' for the last forty years. Television programmes have appeared to promise comprehensive systems that are 'just around the corner', but the reality has been more disappointing, with a number of high profile buyouts and business failures around the turn of the millennium. A task that we do 'naturally' is extremely difficult to automate, although success could render a keyboard and mouse obsolete for many tasks.

Educational establishments would benefit greatly if speech recognition could be used to document meetings, to transcribe interviews, for learners to write essays and for a whole host of similar tasks. Further, speech recognition could provide captioning on live or recorded video for hearing impaired learners, or as an aid for students with poor motor skills who may find use of a pen or keyboard particularly difficult. Voice

commands could be used to operate machinery in technology labs, while the user's hands were engaged in operating other controls, or for physically disabled users to interact with hardware.

Speech recognition is distinct from voice recognition, in that the latter seeks only to determine the identity of the speaker rather than understanding what they are saying.

The problem of speech

Speech recognition has been hindered by a complex interaction of factors, including:

- recognising the target language that is being spoken
- discerning individual words - humans tend to blur the boundaries of words
- homophones - words that sound alike, for example bow and bough
- the problem of context, in which the meaning of words can change (such as 'wicked') and the introduction of new or specialist vocabulary
- regional accents and dialects
- minor speech impediments, such as stuttering, mispronunciation of particular consonants or undue sibilance
- inconsistency in speech, brought on by emotion, stress, tiredness or illness
- background noise
- microphone quality
- connection or call quality when mediated by telecommunications
- processor power and software capability.

One problem is rarely addressed: we tend to think in different ways when we speak compared to when we write. Imagine a verbatim record of a meeting, even without all the 'huhs' and 'errs', it would create a very unhelpful set of minutes. Likewise, creating an essay tends to involve significant pauses while structuring arguments or searching for the correct phrase - these can all be picked up by a dictation system as we 'think out loud' and the software will then try to interpret them according to its speech model.

Speech applications

Voice commands work well in restricted circumstances, as a limited vocabulary is required. Contacts lists can be accessed using 'voice tags' on mobile phones and some control or navigation systems will respond to specific commands. In the case of the former, the software only has to compare the spoken 'tag' against pre-recorded samples of the user's own voice, while in the latter a small set of commands is chosen to be relatively distinct and unambiguous. Ford Synch 3.0 (covered in this [TechRadar article](#)) uses voice commands to interact (via a Bluetooth connected phone) with navigation and other data stored on remote servers. The CEO of [Sensory Inc](#) was [recently commented](#) on a future where speech-controlled Internet devices (SCIDs) would be prevalent, such as alarm clocks that could be commanded to give you a weather report.

Direct voice input (DVI) is used to interact with military avionics systems, for example [Qinetiq's Vocs](#) application, as part of a voice user interface (VUI). Computing has [a report](#) on the use of a system from [Voiteq](#) that uses speech for warehouse management, ensuring that items are correctly stored and picked without operators in a dusty environment having to physically interact with terminals. IBM and Nuance [recently agreed](#) to collaborate on implementing voice systems across a wide range of industries. Call centres often use basic speech recognition as part of their computer telephony integration (CTI) systems, which process calls, allocate operators and record outcomes from conversations. However, automated voice systems have met with mixed success and have revealed the danger of aggravating a frustrated user before they are routed to a human assistant.

A number of companies are developing speech systems that interact with the content of phone calls, or which combine voice control with web applications. Among these are [Ditech Networks](#), who offer in-call monitoring for key words that bring up menus to (for example) search for a pizza outlet or create calendar entries, and [Ribbit](#) (now owned by BT), who offer a variety of application programming interfaces (APIs) that interact with Flash and other web programming technologies. These applications have great potential, but raise both privacy concerns and the issue of call quality interfering with accuracy.

[Google](#) and [Yahoo](#) are among companies that offer search applications powered by voice commands that interact with a remote search engine to provide results on a user's mobile phone. These services can make it easier to input searches, as users do not have to type using fiddly numeric keypads, as well as safer for drivers who want to access information in a 'hands-free' mode. Once the search term has been spoken, it is processed remotely and the results are returned as a standard search listing. However, BBC technology correspondent Rory Cellan-Jones found back in [November 2008](#) that Google's system had a distinct bias to US pronunciations, although a Google spokesperson [told the BBC](#) in April this year that these problems had been largely resolved.

A Guardian [video report](#) features a Microsoft system that is being developed which searches speech to index video streams. The context is a home media centre operated by a standard remote control pad, which is used to find a specific video and jump to the correct location where the search term is being discussed. The main speech processing is almost certainly carried out in a remote data centre.

Speech recognition and transcription

Dictation and speech transcription have been subject to intense research, with a recent desire to transcribe voicemail to forward as email or SMS adding impetus. Following reorganisations and takeovers early this decade, Nuance's [Dragon NaturallySpeaking](#) and Windows Speech Recognition in [Windows Vista](#) and Windows 7 are two of the main systems available to consumers. (Nuance suggests that Dragon users can achieve 'up to 99%' accuracy while dictating as fast as 120 words per minute.) Earlier versions relied on 'discrete speech', where users had to leave gaps between words so that the software could correctly identify them; modern systems can interpret 'continuous speech' with no gaps.

Speech recognition is a subset of natural language processing (NLP). This uses sophisticated statistical models to identify speech segments and compare these to dictionaries of phonemes (distinct units of sound), which are then built up into words to be matched to language-specific and specialist application dictionaries.

Recognition is generally improved through 'training', whereby users record sections of prose and make corrections to the program's output, so that the software can define a set of mathematical transformations used to 'deform' the standard speech model to the user's actual voice.

Online services, such as [SpinVox](#) and [VoxSciences](#), aim to transcribe voicemail messages as either emails or texts, to make it easier for mobile users to keep in touch. [Jott](#) seeks to transcribe dictated voice memos and new VoIP services, such as [Google Voice](#), also advertise voicemail transcription. (See [TechNews 03/09](#) for details of Voice over IP.)

SpinVox was involved recently in [public debate](#) with the [BBC](#) concerning the level of human intervention compared to machine transcription involved in its service. The company has not revealed details that it considers commercially sensitive, but the exchange exposes the particular difficulties of transcribing fragmentary messages, which lack context and are sent over 'noisy' telecommunications links.

A voice future?

Speech recognition remains a niche application to most users, even though they may encounter it regularly in call centre systems, but for a small group of users it could be a significant enabling technology. Despite huge increases in processor power since the late 1990s, the issues outlined at the start of this article remain far from immaterial, such that effective use tends to rely on personalisation through creating voice tags or 'training' the software. Moving processing into the 'cloud' centralises capacity and can aggregate 'machine expertise', suggesting that much more effective systems will be available before long, although that may remain another unfulfilled promise for several years.

Speech transcription is unlikely to see widespread use in education over the short to medium term; the few using the current generation of relatively expensive, complex, specialist software will require access to powerful, personally-owned devices. If it became more common, learners would need coaching in the thinking skills required to use such systems effectively. However, learners with disabilities have found some of these systems very helpful, while specific speech input for controlling hardware or for inputting search terms on mobile devices are emerging as real alternatives.

Software and internet news

Google announces Chrome OS

Google [plans to launch](#) Chrome OS, a new, free, Linux-based operating system (OS) for mobile devices, later next year. (This is not to be confused with the already available Chrome browser.) Google has its Android operating system for mobile phones (which has a substantial Linux core) and there are any number of versions of

Linux already in the mobile devices sphere, so there has been considerable discussion about Google's intentions. One of the clearest factors is that Google has brand profile that the mainstream Linux competitors lack, although Google is not immediately positioning itself as a competitor to Microsoft or Apple.

Google has already lined up [a number of hardware vendors](#) in the mobile devices space and stated, in the original announcement, that 'Chrome OS is an open source, lightweight operating system that will initially be targeted at netbooks'. At present Microsoft has managed to almost monopolise that segment, albeit with the aging Windows XP operating system. Most netbooks were originally supplied with Linux, but users showed a marked preference for the familiar XP environment. Although netbooks have become more powerful, there has been some concern that the forthcoming Windows 7 will not perform very well on low power systems with limited hardware. Experience to date seems to indicate the Windows 7 may perform as well as XP, but that would still leave users waiting too long (in their perception) for the machine to boot. Further, because Windows loads up drivers for most of the hardware at boot time, netbooks have tended not to show the battery life that many had anticipated.

Some manufacturers have been supplying dual boot systems that use an alternative, streamlined version of Linux that boots quickly to provide access to a browser, email and other networked tasks. (See 'Instant-on hardware' in [TechNews 07/09](#).) This approach gives users rapid access to the applications they most need while on the move without taking much power from the battery. Chrome OS would probably deliver the greatest threat to these systems initially, as many believe that manufacturers would install it in parallel with some version of Windows, to give users the flexibility of both options.

This may be how Google OS will roll out, but it does not really give a proper indication of Google's strategy, which is to get everything on the web. Google makes a large part of its money from advertising on web pages, so it has every incentive to get users online as fast as possible. Once there, they can access all kinds of web applications that will perform most of the tasks that user's currently use on the desktop, including its own Google Apps suite. While these applications may be deficient compared to Microsoft Office and other desktop products, this gap may have narrowed by the time the operating system launches. Further, if Google 'owns' the complete stack from hardware through to browser, it can make it operate as a coherent whole, ensuring web applications perform as efficiently and as securely as those outside the browser. Although that is not Google's absolute stated aim, since any Linux-based browser could be used on Chrome OS, it must surely be their ambition.

Google has said that Chrome OS, which will be developed on a protected open source model in a similar way to Android, will be free. At present, [Google has stated](#) that Android will continue to be its OS of choice for mobile phones, but it remains to be seen whether Google might attempt to push Chrome OS 'down' into this space as smartphones become more powerful. However, even though applications would not run on both platforms using the same code, with Android already in the mobile

sector, it seems more likely that Google would push Chrome OS 'up' into Microsoft's space, if the operating system takes off in the way the company hopes. This could turn PCs into browser-based thin clients, with most processing being handled in data centres in the cloud. Unless this happened, Chrome OS would not be much threat to Apple, as the latter has so far distained the whole notebook market segment, but it could become a major competitor in Microsoft's normal territory.

Blurring the line between browser and operating system

Web browsers are constrained in their access to operating system resources, including memory and file storage. This, in part, protects users from some forms of malicious software, while limiting the damage that could be done by a rogue application crashing the browser. Web pages were originally designed around HTML, which is a static page description language, but modern web applications are demanding more access to the computer's resources, including webcams, printers and other devices.

Microsoft has been running a [research project](#) that has developed a prototype Gazelle platform, which aims to run some of the operating system's functions, such as application protection and resource allocation, in a way that secures the browser. Not only can separate web applications currently conflict within the same browser instance, but even on a single web page there may be active content derived from several sources (or 'principals') which may hold up the rendering of that page. Thus, poorly written advertising code might not only prevent a particular page from loading, but also temporarily freeze all other open browser tabs. To solve these problems, Gazelle puts each principal in its own operating environment, although recent Internet Explorer versions and the Google Chrome browser already isolate processes within the browser itself and Mozilla's Firefox also uses a degree of application 'sandboxing'. (See [Ars Technica](#) for more detail on the differences between these systems and Gazelle.)

Gazelle also manages access to system resources, such as webcams and printers, which must otherwise be operated through a browser plug-in. (For example, the many video capture and live streaming websites that have recently appeared use Adobe's Flash plug-in to manage a webcam.) Microsoft has stressed that the project is purely research at present and that the company has no specific plans to create a Gazelle Browser.

On a related note, Google [has announced](#) that its [Native Client](#) add-on is moving out of the initial research phase. Native Client (NaCl) will be an extension to the Chrome browser that will allow web applications to directly run x86 code. (See [TechNews 01/09](#).)

New Microsoft subscription licensing scheme launched

Microsoft has added a new pilot software licence to the schemes currently available to UK schools. The Subscription Enrolment for Schools UK Pilot (SESP) allows schools to count the number of machines to be licensed once per year and multiply that by the number of products to be bought, giving the number of 'units' to be licensed, with a minimum total of 50 units. The license is a subscription model that

must be renewed annually - although figures can be recalculated each time - entitling institutions to all relevant upgrades while the licence remains valid. Both the Windows and Apple Mac platforms are covered and schools can opt to 'downgrade' in order to ensure that computers are running the same version of a particular product or to cover legacy hardware. A range of re-licensing options is available, with terms of one or three years, and schools can also opt to pay for licenses for students to use on personally owned hardware.

The new licence should prove more cost-effective for many schools, compared to the School Agreement, as cover is not required for the entire 'eligible' ICT estate (including those using alternative software or unable to run the licensed products). Ray Fleming from Microsoft UK has a [useful blog post](#) comparing the new licences to the existing Select and School Agreement licenses. The new agreement has resulted from negotiations with Becta and other stakeholders.

Nokia and Microsoft collaborate to improve Office on phones

Microsoft and Nokia [have announced](#) an alliance based on developing productivity software for the Symbian mobile phone platform. The agreement covers production of mobile versions of Microsoft's Office, communications and other business applications for Nokia's E series and subsequent smartphones. The first product, to be released next year, will be Office Communicator Mobile, Microsoft's application that integrates messaging, presence and voice over IP (VoIP) services.

This announcement brings native Microsoft Office compatibility to a wider range of mobile devices, enabling learners, teachers and others to both view and edit documents while on the move. Support for SharePoint Server will provide better access to resources hosted on learning platforms that are based on Microsoft's technologies. Inevitably, Office Mobile versions have reduced functionality due to operating on lower power processors and interacting through small screens using either a keypad or touch interface.

Microsoft, naturally, has mobile applications for its own Windows mobile environment, installed by a number of hardware vendors, while many other mobile operating systems have third party office-compatible applications available. Even though Nokia [is reported as having](#) 45 per cent of the mobile phone market, a Nokia executive [has publicly admitted](#) that the latest move is a direct response to the success of RIM's Blackberry devices. The greatest future challenge may emerge from web-based productivity suites, although mobile browsers have some way to go before they will be able to support a complete set of applications.

'Best before' encryption key time-limits online data

Web page owners do not necessarily want data to be permanently open to view, perhaps due to sensitivity, copyright arrangements or simply because the contents are out of date. Researchers at the University of Washington have created a system that uses data encryption techniques to ensure that such information is no longer accessible once it has passed the appropriate date.

Data security is usually guaranteed using encoding algorithms that are difficult to attack and shared security 'keys' that provide access to the data. These keys are created for a single browser session or are held permanently somewhere on the client's hard disk. So long as the user holds the key, the data can be accessed. The University of Washington team have created a system that requires users to have access to an encryption key in order to read, for example, the contents of an email or a document shared on a web application. A browser plug-in (working with a pre-installed application) transparently accesses the time-limited key in order for the users to see the required content.

The novel part of the [Vanish](#) system is the way in which it fragments the key into many pieces and distributes it around a peer-to-peer (P2P) network, instead of the user holding the complete key (which could be copied and later used to retrieve sensitive data). The distribution principle is based on the same 'torrent' technology used by many P2P file-sharing networks. The originator feeds text through the Vanish software, which performs the encryption and then distributes the key around the P2P network, without keeping a permanent copy. Anyone with access to the encrypted text and the Vanish plug-in can read the data, so users need to make sure that it is held in systems that have other appropriate protection measures. The advantage here is that, although the encrypted message may be permanently stored, the key effectively self-destructs after (in the prototype system) about eight hours, rendering the file unreadable. Thus, although held at some unknown location 'in the cloud' or backed up to other media, sensitive content is protected once the key has expired. An [Ars Technica article](#) gives a helpful description of the way the key is distributed, protected from loss and then expires. Moving the 'chunked' key between P2P servers would preserve content for more than eight hours.

W3C introduces new knowledge management standard

Classifying objects and data requires people to create taxonomies that describe and categorise the 'objects' under consideration. These taxonomies may be centrally defined, such as the Dewey Decimal system used in libraries, or derive informally from user input (for example tags added to content on video sharing sites), often described as 'folksonomies'. The same or similar data are often described using disparate classifications in different contexts or in frameworks created by different organisations; creating a larger, searchable data set that can be searched online requires these descriptions to be cross-referenced and linked, allowing students, staff and all kinds of users to access the widest range of available information.

The World Wide Web Consortium (W3C) [has published](#) the Simple Knowledge Organization System ([SKOS](#)) as a means to bridge such knowledge structures and make them accessible to web-based discovery engines. SKOS is part of W3C's programme to support the development of the 'semantic web' (which some have dubbed 'Web 3.0'), in which data is linked more by meaning than it's location on the internet. SKOS is a language that describes the structure of new and existing dictionaries, thesauri, taxonomies and the like, enabling knowledge management systems to link these structures and to offer better, targeted information from a more diverse range of sources.

The W3C has [also approved](#) POWDER, the [Protocol for Web Description Resources](#), which guarantees the compatibility of web pages against standards for basic browsers on mobile devices. W3C suggest that one spin-off benefit from POWDER is that pages are attributed to a 'publisher', so this could form the basis for automated validation systems that check that page content derives from the source claimed.

New algorithm sifts experts from spammers

Researchers at Cornell University, New York, [are investigating](#) automated methods to identify online information that has been posted by experts and other reliable parties, rather than casual users or spammers. Many systems on social networking sites rank updates and images by some form of 'popularity', emphasising the number of hits a post receives, or the numbers of items a user has posted. Automated systems can also pick up on user-generated tags that describe the content of an update, such as 'hashtags' on the popular Twitter micro-blogging service. However, spambots (automated advertising services) often find these tags and use them to promote their products, so the tag #swineflu could soon be flooded with links to (potentially) disreputable online pharmacies.

The Spamming-resistant Expertise Analysis and Ranking (SPEAR) algorithm uses a self-reinforcing system to validate posts - reputable users are those who post the most useful information and the best information is posted by reputable users. Of course the system has to first identify the reputable users, which it does on the basis that popular users who find new information soonest are most likely to be 'in the know', especially if other users then pass on that same link to the information to their friends. 'Users' that later use the same tags to link to content after a topic has become popular are much more likely to be spammers.

The researchers tested their algorithm on the social bookmarking site Delicious and found, when searching for JavaScript experts, that the top two were professional developers and that none of the top 200 were spammers.

'Sexting' creates concern for over a third of teens

'Sexting' is the practice of sending moderate or explicit sexual messages or images to friends using texts (SMS), multimedia messages (MMS) and online services. This practice can be risky when carried out between friends, as messages may later be used out of context, the mobile device may be picked up by others or the 'friend' may not be the person you thought they were. 'Sexting' may result from undue coercion or may be used as a form of cyber-bullying - the content of messages can be used for extortion or direct abuse of former amorous companions, or new messages may be sent to shock and intimidate unrelated victims.

A survey by the charity [Beatbullying](#) is [reported by BBC News](#) to show that over a third of teens had been subject to such bullying and that by far the majority of such messages came from other teenagers. Taking, sharing or holding 'indecent' pictures of persons under 18 is actually illegal, although the Police may take contextual factors into account [according to the BBC](#). Another report by [BBC's Newsbeat](#)

illustrates the pressure that can be exerted on a young person to engage in 'sexting' and the consequences it had for her life.

Becta's e-safety advice and links can be found at [here](#).

Free-to-use tech training videos for teachers

JISC [has launched](#) a set of new '[Web2practice](#)' videos through its Netskills programme to support educators' use of Web 2.0 tools. Topics covered include collaborative writing, podcasting and micro-blogging. (The last of these, including tools such as Twitter, was covered in [TechNews 03/09](#).) Each video is supported by a PDF guide that can be printed on an A4 sheet.

Russell Stannard, a principal lecturer in Multimedia/ICT at the University of Westminster, has added further resources to his free [Teacher Training Videos](#) collection, which helps staff and students learn how to use PC-based software, as well as online tools, and which seeks to promote good practice. New content is added on a regular basis.

Free e-textbooks linked with material purchases

A textbook publisher, Flat World Knowledge, [is to release](#) more of its titles on the internet under a Creative Commons (CC) licence. A trial using some of its business and economics titles in 30 US colleges earlier this year revealed savings of 82 per cent compared with paper versions, amounting to \$82 (£52) per student per class. The new titles include texts on psychology, sociology and genetics, with 50 titles due to be in development by summer 2010.

The CC licence allows users to access the complete work online and to re-edit and use relevant sections as required. The publisher expects to profit from making the same texts available as PDFs for \$19.95 (which will work with many types of e-book reader) and as audio texts for \$39.95 (approximately £25 and £12.50 respectively). A number of print options and supporting materials will also be available on a paid-for basis. Ars Technica [reports](#) that around two thirds of students purchased such additional materials, with an average value of \$30 (£19), suggesting that the business model has some prospect of success.

In May 2009 the Governor of California, Arnold Schwarzenegger, announced a plan to provide text books to all the state's schools in an open source, digital format. However, [commentators have remarked](#) on the difficulty of launching such projects in the past, especially in view of specific details of the curricula and other requirements defined by the state's education legislation.

TechNews Information

Disclaimer

While every care has been taken in the compilation of this information to ensure that it is accurate at the time of publication, Becta cannot be held responsible for any loss, damage or inconvenience caused as a result of any error or inaccuracy within these pages. Although all references to external sources (including any sites linked to the Becta site) are checked at the time of compilation, Becta does not accept any responsibility for or otherwise endorse any information or products contained in these pages including any sources cited.

We cannot be aware of the uses to which you may put this information nor of the environment in which you are working. Consequently you should take care to obtain professional advice relating to your circumstances before making use of this information.

Copyright and permitted use

The material featured is subject to Becta copyright protection unless otherwise stated. You may reproduce the Becta copyright-protected content, free of charge, in any format or medium without specific permission, provided you are not reproducing it for profit, material or financial gain.

You must reproduce the material accurately and not use it in a misleading context. If you are republishing the material or issuing it to others, you must acknowledge its source, copyright status and date of publication.

The permission to reproduce Becta copyright protected material does not extend to any material that is identified as being the copyright of a third party. You must obtain authorisation to reproduce such material from the copyright holder concerned.

Copyright in the typographical arrangements (including template design and graphics), logos and trademarks, all software compilations, underlying source code and software on this website are copyright Becta or copyright of third parties as identified. All rights reserved.

To unsubscribe:

<http://lists.becta.org.uk/mailman/listinfo/technews>

Feedback:

We welcome your feedback. Email comments to:
technews@lists.becta.org.uk

Publisher details

Becta
Millburn Hill Road, Science Park, Coventry, CV4 7JJ.
Tel: 024 7641 6994
Fax: 024 7641 1418
Email: becta@becta.org.uk