# Socio-Technical Change: Computing, Bioscience, Maths

Dave Cliff 1[1]
Professor of Computer Science
Director, UK Large-Scale Complex IT Systems Initiative
University of Bristol
dc@cs.bris.ac.uk

## Introduction

In the 30 years since 1978 a number of technology trends and developments have had a significant effect upon society, and right now in early 2008 it seems entirely plausible that several of these technology trends will continue to maintain or increase in their societal importance and impact over the coming 30 years to 2038, and beyond. Technology advances will change society but, simultaneously, societal factors will influence the demands placed on technology research, development, commercialization, and application; and so society will shape and influence future technology advances.

Consider a couple of obvious examples. The past 30 years have witnessed a dramatic, explosive, increase in the socioeconomic importance of information technology (IT) systems. It is difficult to remember the extent to which the magnitudes of the Personal Computer (PC) and Internet revolutions were under-estimated until very shortly before their impact became undeniable. Right now, it seems very likely that the post-genomic biotechnology of "proteomics" and "synthetic life" are set (over the next 30 years or so) to have a similarly explosive and revolutionary effect on the world in general, and on advanced industrialized economies in particular. In these aspects of IT and of biotech we can think informally of the causation being from the lab to the world: researchers develop new technologies, those technologies change society.

But we know that causation runs in the opposite direction too. Events in the world, and changes in society, will determine research agendas and affect technology developments, or affect which technologies are commercially successful. That the world's climate is undergoing a period of significant change is no longer in dispute (although the extent to which that change is anthropogenic is still the topic of some debate). That the world's energy consumption is increasing at an alarming rate (regardless of any impact on climate) is, on a 25-50 year perspective, a very serious concern: many respected commentators argue that we are now past the oil peak[2]; some estimates of global uranium stocks imply that, if all major economies switched to traditional fissile nuclear reactors, the world's uranium would be entirely consumed within 30 years. Dealing with climate change, our rampant energy needs, our increasingly crowded planet, and our resultant pollution, are challenges not only for technology development but also for policy-makers and other agents of social change.

Sometimes the challenges lie entirely in the social domain: the deployment of proven technology, presenting no significant research challenges, might nevertheless be highly contentious. One result of the various terrorist attacks attributed to Al-Qaeda has been a greatly increased focus on "homeland security" monitoring and surveillance of people as they go about their everyday business. Here, technology can readily be brought to bear on the problem (for example, basic automated surveillance such as tracking cars by computerized CCTV license-plate recognition is, in principle, more effective and more practicable than having humans do the same job) but thorny social, legal, ethical, and political issues need also to be resolved, such as those concerning the trade-off between the individual's right to privacy, and the state's "duty of care" to detect and defeat terrorist activity.

Because of the interplay and interdependence between technology developments and societal change, making usefully accurate predictions concerning *specific* technology

---

[1] Copyright © D. Cliff, 2008. This document was commissioned in January 2008 by FutureLab as a briefing paper for the UK Government's Department of Children, Schools, and Families (DCSF). This version: March 2008.

[2] See, for example, R. Heinberg (2007) *The Party's Over: Oil, War, and the Fate of Industrial Societies*. 2nd Edition, Clairview Books

developments and how they might interact with social structures on a 15-to-50-year timescale is pretty-much impossible. Nevertheless, there are a number of fairly long-established *general* trends that can usefully be extrapolated out by one or more decades with some authority. That's the approach that we'll take here.

The invitation to write this Challenge Outline paper contained the following specification:

*This challenge would focus specifically upon attempting to explore the various ways in which cutting edge technological developments in computing, bioscience, mathematics, might interact with social structures and practices over the next 15-50 years and to understand how subsequent changes in social practices might have implications for education. This challenge will specifically aim to provide two perspectives on social and technological change. First, from a scientific perspective, this challenge would attempt to map out key trends likely to emerge in the field of automation/artificial intelligence, ubiquitous computing, and brain/world interfaces. Second, from a social scientific perspective, this challenge would explore how these developments might be mobilized, moderated, or transformed in social contexts.*

And so the rest of this paper aims to meet that specification. We concentrate here on computing, with particular reference to computerized automation, artificial intelligence, ubiquitous computing, and advanced interface issues. Bioscience makes an entry because within the next 50 years it seems quite likely that traditional computer technologies will be able to interface directly with the nervous system via implanted electrical connections, and also because novel computers will be built from living cells, instead of silicon chips. Mathematics is not discussed explicitly in any depth here, but *all* of the technology trends explored in this paper require continued effort by highly numerate researchers and practitioners and, in some cases, they also require the development of new mathematics tools and techniques, and so in that sense mathematics silently permeates this entire document.

Six specific major trends are discussed. Before we go on, it's worth pointing out that any discussion paper such as this is *partial* in both senses of the word: it is incomplete, and it reflects the biases of the author. Also, in the text that follows, bullet-points typeset in *italic* are quotes from the various Annex documents supplied by Futurelab/DCSF to define the brief for this paper.

## Six Major Trends

*What are the major trends in this area?*

**Trend 1: The continuation of "Moore's Law" and other curves.** Gordon Moore, a co-founder of Intel, observed in 1965 that the number of transistors that can reasonably be placed on a silicon chip doubles approximately every two years. This exponential growth in component-count is mirrored by corresponding increases in computing power per unit cost, hard disk space per unit cost, data-rate per optical fibre, and several other measures of digital technology capacity and capability. These exponential growth curves provide good fits to the data from digital technology for the last 40 years, and there are strong grounds to expect them to continue for the next few decades at least. Informally, these exponential growth curves yield two important observations: first, the most powerful computer (or other digital device) that you can afford today will cost half as much in two years time; second, in two years time for the same money you will be able to buy something twice as powerful. The compound effect over decades is somewhat more dramatic: a digital device bought today for £1000 will cost less than £31.50 in 2018, and less than £1 by 2028; the same £1000 will, in ten years time, buy a device 32 times more powerful in 2018; and by 2028 a device more than more than a thousand times as powerful as the 2008 version will be available for that same £1000.

**Trend 2: Once-Per-Decade Disruptions.** In 1982 Joel Birnbaum, at that time the Director of Hewlett-Packard's global research laboratories, noted that the dominant mode of commercial provision of computing had undergone a small number of dramatic changes, in each case involving a disruptive technology development that totally altered the industry, and that these disruptions occurred approximately once every ten years. History since 1982 has borne out Birnbaum's observation. The earliest commercially produced computers were (by today's standard) giant, slow, expensive, and cumbersome single-user "mainframe" machines, that were nevertheless unchallenged by any rival technology for over a decade. Eventually, a rival did emerge: advances in miniaturization and in software design led to the development of "minicomputers" which were much smaller and cheaper than mainframes; at the same time, advances in software enabled multi-user time-sharing operating systems so that one computer could be used simultaneously by several users. Approximately a decade later, the revolutionary idea of one-computer-per-desk was made possible, giving birth to the "personal computer" (PC), again driven by Moore's-law reductions in costs and corresponding advances in software and system design. A decade after that came networked computer systems for sharing resources, which led a bit less than a decade later to the sudden explosion of the internet and world-wide web. The net and the web are now well established, and the next transition (to utility-style provisioning where massively powerful computing data-centers are accessed remotely by users with charging via usage meters) is currently underway[3]. The next big shifts, in coming decades, are also starting to come into focus. It seems very likely that ubiquitous and pervasive computing (where almost all products – even cheap throw-away ones -- have significant digital sensing, processing, and communication capability built-in) will have a transformative effect, as the workplace and the home and the streets and even the body all become the "scaffolding" on which are placed tens or hundreds or thousands or millions of small computing devices that interact autonomously with one another, communicating over low-power wireless links such as Bluetooth[4]. Beyond that, within the next 20 years, Moore's law as applied to silicon chips starts to run into the problem that the transistors on the chips will be so small that they involve very few actual atoms of silicon, and so quantum effects become a major consideration: this has been recognized for some time, and research in quantum computing (also referred to as quantum information processing or QIP[5]) seems likely to mature, and to trigger a revolution in information processing at least as significant as those triggered by the invention of the transistor, the chip, and the net. There is one other disruption that seems now to be almost inevitable on a 20-40 year timescale, but that disruption is one which involves not just hugely more computing devices or microscopically smaller logic elements, it involves the engineering of information processing systems based on substrates other than silicon: substrates that are alive.

**Trend 3: Computing as Bioscience.** Computers are *artificial* information processing device, but they are constructed very differently from the *natural* information processing systems found in the nervous systems of animals. Although still very much at the experimental stage, with correspondingly primitive results, researchers are starting to understand how it might be possible to engineer computing systems built from living nerve cells (*neurons*). As advances in genetic engineering and stem-cell research increase, it is feasible (but only on a very long timescale) that computing devices built from networks of living neurons can be created to serve particular needs and uses. The motivation here is not the science-fiction scare-mongering of creating an artificial brain with super-human intelligence: the astonishing capabilities of the human brain are highly unlikely to be surpassed by any artificial device within the next 50 years, whatever that device is made of. Rather, although the human brain is indeed notable for its embodiment of amazing faculties of cognition and creativity, there are more mundane engineering reasons why it is attractive to engineers: it is very compact (approx 1300cc), very light (approx 1.4kg), runs cool ($37_o$ C), uses cheap sources of energy (glucose, oxygen) and is a very low-power device (approx 20W). A conventional silicon-based computer system with the processing power of the human brain could, feasibly, be constructed (or planned, at least). It would certainly be several orders of magnitude larger, heavier, hotter, and more power-hungry

---

[3] See, Nicholas Carr (2008) *The Big Switch: Rewiring the World, from Edison to Google*. Norton

[4] See Hewlett-Packard's *Memory Spot:* http://www.hpl.hp.com/news/2006/jul-sept/memoryspot.html

[5] 5 See, for example: http://www.cs.bris.ac.uk/Research/QuantumComputing/

than a biological brain. As it happens, many of the engineering attractions of biological brains also hold for other types of cell that are more easily manipulated for engineering purposes than neurons are. Researchers at Southampton University[6] have successfully controlled a small hexapod walking robot using a simple "brain" made from cells of the slime-mould *Physarum polycephalum.* Other researchers are working on computing using interactions among enzymes and other biomolecules: currently such experiments are proof-of-concept studies conducted in test-tubes. Nevertheless, the idea of subsequently genetically-engineering bacteria or other simple life-forms to embody these computational "circuits" of interacting molecules is entirely realistic. It seems likely that these non-neural bio-computers will mature more rapidly than neuron-based technologies. Some researchers question the need to work with neurons at all (after all, aeronautical engineers give aeroplanes wings, but not feathers: and so it might be argued that slavishly copying Mother Nature's details may not be as useful as extracting her general principles). Nevertheless, advances in neuroscience do offer increasing credibility to the prospect of *directly* interfacing artificial computing devices (even current silicon-based ones) with the human nervous system. There are manifest opportunities in restoring diseased or damaged functionality. Such technology is already familiar from the 100,000 or more profoundly deaf patients worldwide who have had a partial sense of hearing restored by the addition of a cochlear implant device. Recent experimental studies of direct implants into the visual cortex of the brain have given blind patients a primitive sense of sight[7]. Experimental studies of using digital technology to restore motor and sensory capabilities to patients with spinal-cord injuries are also making advances, as are more general-purpose neural "microarray" implants intended for direct interfacing in enhanced/augmented human-computer interaction purposes, rather than for overcoming damage or disability. It seems plausible that, as such technologies mature in coming decades, much higher-resolution restoration of lost sensory capability will be possible, and advances in "intelligent" robotic prosthetics mean that motor functionality can be restored to patients with damaged/missing limbs or with loss of motor control to those limbs. More mundanely, many thousands of domestic pets (and, subsequently, one professor of cybernetics[8]) have had radio-frequency identification (RFID) "chips" implanted under the skin, to aid in identifying them via a scanning device. It's not impossible that citizens will accept RFID implants as a means of semi-permanently carrying a digitally readable identification code, although the use to which such an implant can be used is likely to be the topic of intense debate. Giving paramedics attending accident scenes the ability to identify an unconscious/injured patient by scanning the victim's RFID chip is plausibly benign, but it's not hard to imagine more Orwellian applications for state monitoring of the populace. Finally, it is worth emphasizing here that the prospect of "downloading" one person's memories or sensations into some digital medium and then "uploading" them into another person's head (or even back into the original donor's head) will remain firmly in the realms of science fiction and ill-informed journalism for the next 50 years at least. (I sincerely doubt that it will ever be possible, but it's beyond the scope of this paper to discuss why).

Much more likely, in the medium term at least, is the prospect of "brain-content" being *externalized* into discrete devices and information appliances. Many people are already reliant on their mobile phone's SIM-card to remember, on their behalf, the phone numbers of their friends and family: I don't know my sister's mobile phone number, because my phone knows it for me. As Moore's Law and related technology growth-curves deliver ever more powerful information devices in ever smaller packages, we can already envisage the prospect of an entire three-score-and-ten years of a human life being recorded in digital video and audio and then being stored on a single iPod[9]. If *everything* you have ever seen

---

[6] See http://www.newscientist.com/article.ns?id=dn8718

[7] See http://www.bmj.com/cgi/content/full/330/7481/30

[8] See http://www.wired.com/culture/lifestyle/news/2000/09/38467

[9] Computer Scientist Alan Dix was the first to note this, though he did so in 2004 and hence his calculations involved iPods that, by today's standards, have much smaller maximum amounts of storage. Here I've updated Dix's calculations to give 2008 figures. Dix observed that 70 years is $2.2*10_9$ seconds, which at 100kbits per second for the compressed audio & video requires approximately 27.5 terabytes of storage. In 2008, that much storage is provided by 170 top-of-the-range 160Gb iPods, altogether costing a total of less than £40k, and

or heard can be digitized and held on one device, then that device can come to serve as your ultimate *aide memoire*. The research questions that then arise center on how we develop new tools to index, manipulate, and access such huge quantities of data[10]. This is not a bioscience issue. Rather, it is a new challenge for computer science. In principle, artificial intelligence (AI) techniques could help. We turn to AI next.

**Trend 4: Artificial Intelligence Remains Hard.** Research in Artificial Intelligence (AI) over the past 50 years has yet to deliver the machines with super-human general intelligence beloved of science fiction writers, and there are little or no grounds to expect such a development in forthcoming decades either. All the indications from the history of AI research to date indicate that *general* human-level cognitive faculties are just extraordinarily hard to recreate via standard engineering practices. Even perceptual processes that we humans think of as essentially effortless, such as recognizing visual scenes or following a conversation in a noisy environment, are astonishingly hard to reliably engineer into computer systems. Nevertheless, for some *specific* tasks or problems usually associated with or requiring intelligence if performed by a human, computers can now routinely outperform people. This is especially true of tasks or problem-domains that involving integrating large amounts of data in small amounts of time (such as financial trading), or searching vast spaces of scenarios and possible actions (such as playing chess). As processors increase in speed, and as data-stores increase in size (both following Moore's Law style growth curves) so in some cases the sophistication of AI algorithms counts for less: the fact that no human could combine so much data or search so many possibilities at such speed is sufficient to give the impression of intelligence, despite the fact that the algorithm is actually really rather dumb.

This point is illustrated most starkly in automated optimization. There are several algorithmic approaches for attempting to identify an optimal solution in a (usually vast) space of possible solutions. Very often this is a needle-in-a-haystack situation, one where human ingenuity and creativity has long held the upper hand. Some of the most productive optimization algorithms are also the most dumb. In fact, the dumbest algorithm of all is also the best, in the sense that it will *certainly* find the optimum: *exhaustive search* generates, and tests, every possible solution; once all possible solutions have been tested, the best will be known. Let's say that having a computer generate and test one possible solution takes one second, and let's say that there a billion possible solutions. On a single computer, such an exhaustive search would take over 31 years to complete. If the search can instead be distributed over a networked cluster or grid of, say, 1000 computers (not a large number by current standards), then the same search will terminate in less than 12 days. Moore's Law indicates that, within a decade, such a search can be done in less than a day on the same-sized cluster for the same cost. The algorithm is the same, idiotic, exhaustive search: but the wait-time falls from 31 years to less than a day because of ongoing trends in the underlying computer technology. Now if it takes a human more than a day to come up with a solution, then the computer-based solution wins. And, crucially, if we treat this AI system as a black box (i.e., don't ask how it actually works) then it might *appear* to be "intelligently" or even "creatively" generating new solutions.

Prof. Sir Tim Berners-Lee, the scientist widely credited as the inventor of the world-wide-web, has in recent years turned his efforts to rallying support for a new vision of the next-generation web, one in which the computers responsible for housing and shifting the "content" of the web actually have the ability to reason about the *meaning* of that content. This is an initiative known as the Semantic Web, and it has attracted significant attention from AI researchers in academia and industry alike[11]. If the Berners-Lee vision of the Semantic Web can be delivered, it could have a transformational impact at least as great as

---

collectively fitting in a small suitcase. Assuming that maximum iPod storage capacity continues to double roughly once every two years while the price stays roughly constant, then the number of iPods required halves every two years, and so in 16 years time the number of top-end iPods required will be $170/2_{(16/2)}$ which equals 0.664. That is, by 2024, the digitized video and audio capturing an entire 70-year human life will fit on a single iPod, and that iPod will only be 66% full, and that iPod will still cost £250.

[10] http://www.memoriesforlife.org/

[11] T. Berners-Lee, J. Hendler, & O. Lassila "The Semantic Web" *Scientific American*, May 17, 2001.

that of the current world-wide web. And, as with exhaustive search, it could give the impression of (the web of) computers being intelligent, even if the inner workings of the semantic web do not actually involve advanced AI algorithms modeling human thought processes.

Despite this, it is somewhat sobering to observe that, after 50 years of worldwide AI research effort, the online retailer Amazon.com recently set up a successful new subsidiary online business called *Amazon Mechanical Turk*,[12] subtitled *Artificial Artificial Intelligence,* where members of the general public can sign up to make themselves available as the online "brains" for simple "AI" tasks that computers *still* can't do, such as recognizing or classifying images. The human participants behind the "MTurk" are paid on a piecework basis, where each piece of work may only take a few seconds to complete – and generates a commensurately minimal payment. For as long as a human brain is available for hire (even if only for a few seconds) at a cost less than, and with accuracy more than, an advanced AI computer system, human brains are likely to remain the first choice.

**Trend 5: 3D Printing and Plastic Electronics** In the past 20 years, so-called "3D printers" have moved quite firmly from proof-of-concept to successful (if expensive) commercial products, commonly used for rapid prototyping in manufacturing engineering. A conventional laser or inkjet printer allows one to go from a virtual, digital, two-dimensional (2D) document or image (specified in a format such as PDF and previewable on a computer display) to a real 2D print of that document or image. A 3D printer does much the same, starting with a digital specification of a 3D object – e.g. from a computer-aided design (CAD) software package, but the end result is exactly that three-dimensional physical object made from a plastic – rather than a 2D pattern of ink on paper. It now looks likely that 3D printers will within a few decades be as common as their 2D counterparts are now. The impact of this trend has the potential to be highly significant.

There are a variety of techniques employed in 3D printers, and it is beyond the scope of this paper to review them all. As illustration, in one technique, microscopic granules of plastic are scattered on a flat level "bed" to a depth of less than a millimetre, and a computer-controlled laser then traces over the exposed top-surface of those granules. Where the laser hits the granules, they fuse to form a solid. If the laser skips a spot, the granules remain in powder form. On the first pass, the laser draws some kind of a 2D pattern on the exposed upper surface of granules, creating a gossamer-thin solid "print" of that 2D pattern on the top of the plastic-granule dust. Then, more granules are scattered on the surface, building it up by a small additional depth, and the laser is then made to scan over the new surface of "virgin" granules, again fusing some of them. This time though, the beam will not only fuse adjacent granules in the top layer; it can also fuse granules to those solids already fused in the previous pass(es). This scatter-fuse-scatter-fuse layering process is repeated for a large number of layers, and then the remaining un-fused granules are removed (e.g., by blowing them away with compressed air). What remains is a 3D plastic object, built by successive accretion of layers of fused plastic granules.

Let's take a simple example. To construct a *solid* upright cone in a 3D printer, on the first pass the laser would scan a solid (filled) disc onto the surface of granules, and in successive layers the laser would scan a succession of discs each concentric with the previous layer's, but with a reduced radius, until in the final layer the laser fuses a single "point" marking the apex of the cone. For comparison, to create a *hollow* upright cone, the first layer would again be a disc (possibly with a small hole somewhere), but successive scans of the laser would draw concentric empty circles rather than filled discs, until the final apex point is reached. When the apex is complete, the surrounding granules can be blown away, leaving a cone filled with granules. A few vigorous shakes should then empty the granules from the interior of the cone, through the hole "printed" into its base at the start, leaving the desired hollow plastic cone.

This simple example is worth pursing in a little more detail. For the sake of argument, imagine that there is a need for a hollow plastic cone containing a small solid plastic sphere

---

[12] http://www.mturk.com/

and a small solid plastic cube. Under traditional manufacturing processes, the cube and the sphere would probably be manufactured separately, and then inserted into the cone before the cone's base is sealed. In a 3D printer, all three objects can be generated simultaneously – the two smaller objects "printed" directly in the hollow void of the cone. Or, imagine that there is a need for an object with a specific honeycomb interior structure: giving very high strength with very low weight – possibly with the honeycomb varying within the structure to support specific lines of stress or load[13]. Using traditional techniques, such honeycombed objects are very hard to make. Using 3D printers, they are effectively no more difficult than a sphere.

3D printers were originally intended for rapidly producing prototypes of designs that, if they went into full-scale manufacturing, would then be produced from injection-molded plastic or die-cast metal. But 3D printing technology has already matured to the point that researchers are exploring their use in *rapid manufacturing* rather than in rapid prototyping. That is, where the output of the 3D printer is not an evaluation-model of a planned final product. Rather, it *is* the final product.

So far, so mechanical. Using new technology to create novel plastic objects will have many uses, but the relevance to advanced computing and IT is that recent developments have made it clear that it is very likely that, very soon, it will be possible not only to print 3D plastic objects, but also to *embed electronics in the 3D print at the time of printing.* There are two ways this can be done. The first method, achievable right now, is to precisely drop small self-contained processor/memory/receiver/transmitter units (roughly the size of a grain of rice)[14] onto the granule-bed, and fuse them into or onto the 3D object, where they communicate with each other via low-power and low-distance radio or infra-red links. The second method, already achievable in principle but not yet realistic at commercial scales, is to "print" so-called "plastic electronics" onto (or into) the 3D physical object as it is built. Plastic electronics involves using 2D or 3D printing technologies to build up circuits from printed layers of insulators, conductors, and semiconductors.

Thus, this gives strong grounds for talking of a future, two or three decades hence, in which 3D physical artefacts can be "printed" from a computer, where not only are the physical appearance and mechanical capabilities of the artefact determined by the printed 3D plastic, but also the artefact's electronic control and information-processing capabilities are determined by the embedded printed electronics.

Although 3D printing and plastic-electronics technologies do not have as long a track-record as silicon chips, the available data for these new technologies already points to Moore's-Law-style exponential halvings in cost (or doublings in power) over coming decades. Say, for the sake of argument, that currently a 3D printer costs £500,000. Such a price places it well beyond the pocket of pretty-much all UK schools and homes, given its specialist capabilities. Wait twenty years, and if these technologies are on a Moore's Law curve, then the same functionality will be available for £500 (in today's money).

What happens, in 30 years say, when anyone can design and print a new mechatronic device in their schoolroom or home office? "Mechatronic" is a standard term for systems involving integrated mechanical, electronic, and software components. In recent years, mechatronics has increasingly been associated with research in autonomous mobile robotics. The 30-year future prospect of people designing (or, at least, customising an existing design for) their own robots in software and then printing them for immediate use seems wholly alien today, much like the current world-wide-web, or cheap photo-quality digital printers, say, would have seemed 30 years ago in 1978. Remember that, in 1978, many people in the UK had not yet used an electronic pocket calculator.

---

[13] See, for example, D. W. Thompson (1992) *On Growth and Form.* Canto/Cambridge University Press

[14] See Hewlett-Packard's *Memory Spot:* http://www.hpl.hp.com/news/2006/jul-sept/memoryspot.html

**Trend 6: Large-Scale Socio-Technical Systems of Systems** Increasingly, a standard desktop PC is no longer a single computer. So-called "multi-core" processor chips are, in essence, multiple computers (cores) on the same piece of silicon. Just as many hands make light work, so multiple cores can dramatically improve the performance of a PC. "Dual-core" is now a standard phrase in the lexicon of high-street PC retailers, and top-of-the-range PCs with quad-core processors are already starting to appear. The main processor chip in the Sony PlayStation 3 games console is a special-purpose 8-core device, called a "Cell B.E.", designed in collaboration with computer giant IBM[15]. So, a single PS3 is actually eight computers in one box. This trend toward multi-core processors with more and more cores looks set to continue, if only for the pragmatic reason that it is currently much easier to produce ten copies of a 2.5Ghz-speed processoon a single chip than it is to create a single 25Ghz processor c

But the proliferation of cores inside desktop boxes is not the only way in which people can gain ready access to the combined results of multiple computers. The past decade has seen a growth in so-called "clustered" computing where the same many-hands-make-light-work approach is applied at the level of multiple computer boxes in one big room rather than multiple computer cores on one big chip. A little over a decade ago, various research teams realized that, by connecting together a few tens or hundreds of standard desktop PCs and having them synchronized by some management software, it is possible to create a cluster of machines offering high-performance computing (HPC) power that would otherwise only have been available in a multi-million-dollar specialist supercomputer. In the past couple of years, the next of Birnbaum's once-per-decade transitions in the provisioning of commercial IT has clearly started to occur, triggered by this. Huge "data centres" housing thousands or tens of thousands of commodity PCs (each primarily consisting just of the motherboard in a thin box – referred to as a "blade server" – because inessential items such as DVD drives, keyboards, and screens need not be replicated) have been built to offer computing facilities that are vast in terms both of processing power and storage capacity, and which are accessible by remote users over high-speed internet connections[16]. Every time someone uses the Google search-engine, they benefit from the truly vast data-centres that Google has built to search and index the world-wide-web. Google's business model makes the use of their data-centres free to the casual search-engine user, because Google generates its revenues from lucrative sales of advertising space on the search-results web-pages presented to such users. But, increasingly, companies are starting to offer access to their data-centres for HPC and data-processing on a "metered" pay-per-use basis. One phrase commonly used to describe this is "utility" computing, because compute-power and data-storage start to be generated and sold much like the longer-established utilities of electricity, gas, and water. The cost of building the data-centres is huge, but the cost to a user of "renting" a few of the blades for a few minutes or hours is tiny in comparison. The need for business, or homes, or schools, to own a powerful computer facility is greatly diminished because remote large-scale computing facilities of immense power can be connected to commonplace "access devices", such as PCs or mobile phones, via telecoms internet links, at very low cost to the end-user.

At the same time, the IT systems supporting major companies, and major public-sector institutions and organizations, can also be viewed as distributed large-scale complex IT systems. Here, thousands or tens of thousands of PCs are networked together, within and between buildings around the country or around the globe, to support the overall business of that company or organization. Increasingly, organizations find themselves critically reliant on the IT systems that they use to support their business, even when their business itself has very little to do with IT.

So, at the level of the computer chip, and the data-centre, and the organization-wide IT system, we currently see rapid increases in the number of computing elements that interact with one another to form the IT "fabric" on which business processes and activities ever more critically rely. And at the same time we also see rapid increases in the social or

---

[15] See http://domino.watson.ibm.com/comm/research.nsf/pages/r.arch.innovation.html

[16] See Nicholas Carr (2008) *The Big Switch: Rewiring the World, from Edison to Google*. Norton

economic severity of the consequences of failures in that fabric. Ten years ago, if the IT servers in a hospital broke down, the knock-on effects in terms of provision of healthcare services were not great. In ten years time (when we can expect most patient health data to be digitised at the point of capture), the consequences of an IT failure in a hospital could be really very much more severe. Similar stories can be told for many other public services and for most private-sector corporations and enterprises. Worse still, all these IT systems are increasingly becoming linked together and capable of interacting as so-called *systems-of-systems (SoS)*. There is very little established engineering practice for SoS.

As the number of interacting components that determine the overall behavior of such systems increase, so the prospect of accurately predicting the system-level behavior becomes ever more difficult. Technically, these are *complex systems*, which is a shorthand way of saying that there are mathematical nonlinearities ("tipping points") in the responses of the components, and in the interaction between the components, which compound across the system in such a way that even if you know everything about all the components and all their interactions, it may not be possible *even in principle* to accurately predict the system-level behavior. This is particularly worrying because it implies that there may be types of failure which can occur that we simply do not know about until an instance of that type of failure *actually does* occur. And, it is in the nature of complex systems that in certain circumstances the effects of the failure or malfunction of one or two components can interact and ripple out, causing disruption across the entire system/SoS, or even collapse of the entire system/SoS.

A new scientific field, known as *complexity science*, has emerged in the last 20 years, aiming to advance our understanding of how nonlinear interactions between components relate to system-level behavior, in a variety of diverse systems including economies of interacting traders, evolving populations of interacting animals, and brains composed of interacting neurons. Issues in large-scale complex IT systems are also studied in complexity science, and increasingly the people who design and maintain large-scale complex IT systems are starting to experiment with ideas drawn from economics and from biology to design better IT systems. Within the IT research community, there is talk of trying to create large-scale IT systems that are self-organizing, self-managing, self-regulating, and self-repairing. (These are known as "self-star" IT systems, because the asterisk character "*", commonly referred to as the "star" character, is used by computer scientists as a "wildcard" symbol that matches all words: hence "self-*"). Increasingly, engineers aiming to create self-star IT systems are drawing inspiration from naturally occurring self-star entities such as individual organisms, entire populations of animals, and even economic markets. For example, researchers are experimenting with building self-star IT systems where supply and demand of resources within the system are regulated by an internal artificial market[17], where the arrangement of hard-disk storage is based on findings of how ant-colonies organize food-foraging[18], and where computer viruses and other "malware" code are detected and dealt with by the IT system's artificial "immune system"[19].

The trends identified and discussed above point, for better or for worse, to a future critically reliant on engineered systems of interacting adaptive entities, where each of those interacting adaptive entities *are themselves* engineered systems of interacting adaptive entities or components. And (it gets worse) the systems that find themselves interacting with each other and effecting one another may not have been designed and built with a view to them subsequently being linked up to those other systems. The technical name for this is *Systems-of-Systems*, or "SoS". In many cases of significant interest, at least one of the adaptive interacting entities in a SoS is a *social entity*: possibly a lone human operator, perhaps a team or a firm, maybe an entire company or government department. That is to say: many crucial SoS are inherently *socio-technical systems*. We in the UK (and in many other advanced economies) have over recent years become ever more reliant on national-

---

[17] See, for example, http://www.marketbasedcontrol.com/

[18] See http://www.hpl.hp.com/techreports/2003/HPL-2003-221.pdf

[19] See, for example, http://www.cs.unm.edu/~immsec/begin.html

scale and international-scale SoS, that have been assembled "organically" from pre-existing systems that were never engineered (i.e., rather than designed from scratch as a single unified whole). Yet there is almost no well-established engineering practice for socio-technical SoS, and, relatedly, there are some significant known gaps in our scientific understanding of socio-technical SoS too.

The lack of a comprehensive understanding of socio-technical SoS, and the relative immaturity of complexity science analyses and large-scale engineered self-star IT systems, points to a worrying prospect for the future. Put bluntly, the complexity of the large-scale complex IT systems, and of the socio-technical SoS that they combine to create, is going up very fast indeed. Our socio-economic dependency on these systems is also rapidly increasing. The problem is that our ability to understand and manage these SoS is also going up, but there is now genuine concern in academia and industry that it may just not be going up fast enough. We face the worrying prospect that, sometime soon, society finds itself critically dependent on interacting and interlocking large-scale IT systems forming various broader socio-technical SoS that no-one really understands, and that are capable of failing in ways that no-one predicts, until those failures actually occur.

## Discussion: Uncertainties, Discontinuities, and Stasis

> *What are the major uncertainties?*
> *What profound discontinuities or breaks with the existing trends might emerge?*
> *What is likely to stay the same and why? What are the forces acting against radical change?*
> *What are the problems or challenges that are likely to emerge and which need to be solved to move the field forward?*

The first five trends identified above are, in essence, extrapolations from happenings in the more-or-less recent past. And, in each, the observation is essentially optimistic: the claim is that new developments are likely to improve on the state of the art: things will get desirably faster or cheaper or smaller. However, Trend 6 is somewhat more pessimistic: here the issue is a problem that is believed to be rapidly growing and which may possibly be heading, if not for some kind of crisis, then for some major difficulties at least. For this reason, the discussion that follows deals first with Trends 1 to 5 taken together, and then with Trend 6 separately.

For Trends 1 to 5, the primary uncertainties and discontinuities will most likely be, in the phrasing of Donald Rumsfeld, a consequence of "unknown unknowns". That is, underlying the projected continuation of these trends, there are "known knowns" (things we already know how to do, that will underpin the continuation of the trends) and there are also "known unknowns" (things that we either don't yet know how to do at all, or don't know how to do sufficiently cheaply/reliably). But these known unknowns are not a cause for undue concern: in any field of inquiry the list of "known unknowns" forms the challenges that define the current research agenda(s). All of the first five trends look likely to continue so long as progress along the research agendas in the relevant fields does not uncover any major unpleasant surprises. And any such major surprise would most likely come in the form of an "unknown unknown": something nasty that we simply just don't yet know to identify as a significant problem. Of course, it's also possible that the major surprise is a *nice* surprise: if a brilliant researcher suddenly unlocks the secret of making truly intelligent robots, then there could be a *positive* discontinuity as human-like intelligent robots become available for mass production (or bespoke home production via Trend 5) much sooner than predicted above.

In fact, the primary cause of any uncertainty or discontinuity in the first five trends is more likely to be exogenous than endogenous. By an *endogenous cause*, I mean something specific to the above sets of trends, internal to those fields of science and engineering that drive the technology development. *Exogenous causes* of uncertainty or discontinuity would be external to the activities driving these trends. Examples might include major shocks to the price of energy, availability of raw materials, or availability of skilled researchers. But such exogenous causes are so major that they would have some impact on *all* forms of socio-technical change, not just those trends listed above.

For Trend 6, the possibility of significant problems lurking in the future is already recognized. To be clear, the claim in Trend 6 is not that society is facing imminent catastrophic melt-down. Rather, as society comes to rely ever more on large-scale complex IT systems (LSCITS) – as those LSCITS become core to the critical paths of delivery of healthcare, education, law-enforcement, defence, government, and very many major business and enterprise sectors, and as the LSCITS for each of those become interconnected to form ultra-large-scale systems of systems (SoS) – so the need to be well prepared for the effects of the growth in complexity of those LSCITS is greatly emphasized, as is the potential for sudden wide-ranging failures. Socio-technical Systems-of-Systems are sufficiently novel that significant research effort is required to understand how traditional engineering techniques can best be amended or extended to deal with them, both in designing and deploying future systems, and in managing and maintaining existing ones[20].

## Social Justice, Social Mobility

*What are the implications for social justice or social mobility?*

Primary social justice and mobility implications come from adoption profiles and usage patterns, rather than from the technology itself. If we take "social justice" to mean equality of opportunity, rights, and benefits, then it is possible to envisage ways in which the technology trends discussed above increase social justice, and equally possible to dream up scenarios in which social justice is decreased. Similarly, if we measure "social mobility" by monitoring the extent to which children born to poor families fulfill their potential, relative to those born to wealthy families, then again it's not hard to come up with stories of how these technology trends might increase mobility, or how they could reduce it, but the differences between those stories will be in how the technology is deployed and used, not in the technology itself. The devil lies in the detail.

For instance, in the UK of 2008, stating that someone owns a mobile phone is much less informative than stating that they do not. In the last 15-20 years, owning a mobile phone has gone from being a rare and status-laden thing (indicating high income and/or high job responsibility) to being a commonplace event, with penetration levels well in excess of 80%. For this reason, the prospect of allowing people to vote in local and general elections via their mobile phones has mass appeal in 2008 that it did not have in 1988, and the idea has been actively explored by the UK government. But it could be argued that introduction of mobile-phone voting is significantly more socially divisive now than it would have been in 1988, in the sense that back then such mobile-phone voting would have been an option available only to a small socio-economic elite with manifestly elevated levels of disposable income – and so the opportunity available to the vast majority of voters would have been unchanged. In 1988 the argument would be that folk rich enough to own a mobile phone would just have one more opportunity to spend their money to save their time (the cost of operating a mobile phone being more than the cost of walking to a polling station), and the only advantage it gives them is one of convenience: their vote is still just one vote. Furthermore, and crucially, in 1988 the most disadvantaged members of society were still members of that majority of non-mobile-owners who would be unchanged by the introduction of mobile-voting. By 2008 the vast majority now have ownership of a mobile phone and hence would have access to the new voting opportunity, yet the minority without a mobile phone is likely to include disproportionately high numbers of the most disadvantaged members of society: those too poor to own a phone, those incapable of operating one, and so on. So in 2008 the introduction of mobile-phone voting increases voting opportunities for the majority, but thereby selectively reinforces relative inequality,

---

[20] A major new UK research and training initiative is aimed at addressing this: see www.lscits.org

because the least advantaged members of society are no longer members of the majority being offered that additional opportunity.

In this example, the social justice implication lies not with the technology *per se*, but with the government's policy decisions on how mobile telephony technology might be used in the electoral process, and on the demographics of mobile-phone ownership at the time the decision comes into force. Technology developments create new opportunities for governments, and sometimes new threats too. The best way of identifying the implications for social justice and social mobility is to attempt to predict patterns of take-up (i.e., how the technology ownership/access "penetration" grows – and perhaps decays – over time) and usage (i.e., what people do with the technology once they have access to it).

Alas, the take-up and usage patterns for mass-market application of any technology is notoriously difficult to predict with great accuracy. Two examples from the development and commercialization of mobile phone technology in recent years illustrate this. First, the huge enthusiasm for SMS text-messaging famously came as a great surprise to pretty-much everyone in the UK mobile phone industry at the time. Second, the comparative lack of enthusiasm for third-generation (3G) mobile functionality came as an unwelcome surprise to those telecoms companies that won in the UK 3G spectrum auctions: their massive investments were rewarded with painfully slow uptake from customers of the new facilities that 3G offers. In the first we have a cheap, simple, but primitive technology becoming a roaring success, in the second it seems that an entire industry failed to guess that major (but expensive) increases in technology sophistication offered little actual interest to the consumers.

Despite the difficulty of predicting these things, there are people whose profession it is to do so, and taking a number of opinions is generally better than taking one. So, to evaluate the social justice and social mobility implications, we need to consult widely, with experts and with the wider public, and to take time and care in doing so.

Thankfully, the lead-time from proof-of-concept to full-scale commercialization (which is most usually the step before any major societal impact) in the technologies discussed here is sufficiently long that it is possible to discuss and evaluate in advance likely major effects on society to a fair extent, but this requires responsible, reasoned, and well-informed debate from experts and practitioners, policy-makers, and media commentators. Unfortunately, reasoned, responsible, and well-informed debate on scientific and technology developments are ever more a rarity in the UK. Headlines such as "new technology poses no obvious threat to health/society" do not sell nearly so many newspapers as apocalyptic scare-stories do. Recent public "debate" (i.e. arguments and counter-arguments from parties holding different views, and media coverage of those arguments) over possible/presumed dangers in radiation from high-voltage power lines, mobile telephone handsets and antennae, the MMR combination vaccine, environmental impact of GM foodstuffs, anthropogenic global warming (and so on: the list is disappointingly long) can all too often be readily characterized as ill-informed and poorly argued, to greater or lesser extents, on *both* sides of the argument. For example, people worried by MMR were, unsurprisingly, not best calmed by authority figures telling them that there was simply "no risk" – most people intuitively understand that any non-placebo medical intervention carries *some* risk, even if it is not the actual risk that concerns the worried people. Simplistic attempts to calm people just worried them more, and also reduced their trust in the authority figures that were trying to calm them.

In closing this section I observe without comment that in 2005 the Cabinet Office of the UK Government published a strategic briefing document "*Transformational Government: Enabled by Technology*"[21]. In this 25-page document discussing how IT and telecoms technology advances could affect and change government practice and delivery, the phrase "social justice" occurs only once, in the first paragraph of the Introduction, where it is listed as one of the three major challenges that face modern government. The phrase "social mobility" occurs nowhere.

---

[21] http://www.cio.gov.uk/documents/pdf/transgov/transgov-strategy.pdf

## Implications for Education

*What is the extent of educational "influence" in this area – to what extent can educational systems impact on these trends?*
*How might educational goals be affected?*
*How might educational processes be affected?*

The primary influence of education in the areas discussed in this paper lies mainly in the extent to which the education system prepares citizens of the country to deal with the effects of these technology trends, and enables citizens to make informed decisions about choices involving those effects. The broad qualitative nature of all six trends are unlikely to be affected much by the specific schools-level education policy of the UK, in the sense that the technology trends discussed here are consequences of global R&D efforts and their associated global socioeconomic driving factors. However, the UK education system has a definite quantitative influence, insofar as UK-educated scientists and engineers have a long history of making major contributions to the various fields of inquiry relevant to these trends. To take an absurdly extreme example: if the UK ceased to allow any education whatsoever in science or engineering, it's reasonable to expect that long-term progress along all of the trends outlined above would be somewhat slower in one respect or another because of the consequent loss of input from UK-educated researchers and practitioners. In that sense, the only influence our education system is likely to have on these trends is to have some long-term effect on the speed at which they occur. UK education policy cannot reverse or stop them. From a perspective of national education goals and processes, the real issue is how we adjust those goals and processes to cope with and accommodate the societal effects of these trends, not how we influence the trends themselves.

With regard to how goals and processes might be affected, it's worth exploring two questions. The first is: how might alterations in technology-enhanced learning driven by these technology trends affect education goals and processes? And the second is: how might education goals and processes be altered to take account of the effects of these technology trends in society? These questions form the basis of the two core challenges recommended here for further exploration.

## Core Challenge 1:
## Impact of technology trends on technology-enhanced learning.

Currently the use of PCs and data-projectors with touch-sensitive "interactive whiteboards" routinely provides technology-enhanced learning for children at all stages in UK schools. There is a wealth of web-content freely available and specifically intended as teaching material, plus an even larger volume of more free general-purpose reference content (such as Wikipedia, Google Earth, Google Maps, etc). Exponential growth in the amount of such data available, the density of storage media for holding that data, the amount of CPU-power available to process that data, and the speed of the networks over which the data travels all looks set to continue in coming decades. In some cases the web material offers a straight substitution and saving over traditional paper-based publications (why buy an expensive set of *Encyclopedia Britannica* when Wikipedia is available free?). In other cases, the very low (and falling) marginal costs of digital media production make cost-effective the teaching of creative arts that twenty years ago would have been prohibitively expensive: for instance, children can experiment with composing scores for multiple instruments – from a duo to a full orchestra –via a cheap polyphonic synthesizer "virtual orchestra" software package that will play the score back to them at the press of a button, or record the audio to a CD. Similarly, the marginal cost of taking a single shot in digital photography, or of recording an hour of digital video, are so low that wide-scale teaching of photography or movie-making are both now much more readily affordable prospects. One consequence of the rise of digital technologies applicable to teaching is that the financial barriers to entry are lowered, and so the sourcing of teaching material for home-schooling is much less expensive than it was previously: on that basis, we could see a significant rise in home-schooling in coming decades.

Nevertheless, economies of scale and sharing of resources still mean that investment in more expensive technology assets and facilities may still be the domain of traditionally-structured schools. Just as, 30 years ago, a typical UK secondary school might have only a small number of pre-PC-style microcomputers so, in 10-20 years time, a typical UK secondary school might have only a handful of 3D printers producing 3D artifacts with integrated plastic electronics. Pupils in 2028 might routinely engage in the design of autonomous mobile robots with integrated electronics that are "instantly" produced by a 3D printer, in much the same way that today's students can engage in creating digital music, photos, or videos that are then "instantly" preserved on CD or DVD burners, or photo-quality inkjet printers. Today, with 3D printers costing tens or hundreds of thousand of pounds, such a proposition may seem outlandish. But this proposal is no more outlandish than the proposition in 1963 that, by 1983, secondary schools would routinely house multiple interactive "personal" computers with (by the standards of 1963) huge RAM capacities and astonishing processor-speeds and graphics capabilities[22].

The human teacher's role seems unlikely to be usurped by a machine anytime in the foreseeable future. We humans are social animals, not data-banks. The education process and the role of the human teacher in that process is as much about stewardship of the child's emotional, social, psychological, and physical development as it is about filling young heads with declarative facts and procedural knowledge. Artificial Intelligence techniques may create more sophisticated intelligent tutoring systems (ITS), better able to form a model of the student's current level of learning, or learning difficulties, and thereby able to tailor the tuition to that student, but there are no credible claims that such adaptive ITS could (or should) entirely replace the human teacher. It is plausible to envisage adaptive intelligent physical play equipment (e.g., a multi-gym or adventure-playground that recognizes the individual student and sets physical challenges accordingly, or an autonomous robot baby that helps teach basic childcare) but again the most likely use-cases within the next fifty years for these technologies seems to be within traditional schooling structures. In that sense then, it seems likely that schools of 2058 will be as similar to those of today as our current schools are to those of 1958, or even of 1908.

But, in typing these words, I'm painfully aware that **I am writing here very much as a non-expert.** For an expert view, professional judgment should be sought from researchers and practitioners whose core expertise lies in the domain of technology-enhanced learning. Four leading UK-based researchers with strong expertise in the intersection between education and the technology trends discussed here are identified below who could respond from positions of considerable authority. But, before specific names are mentioned, we should discuss the second core challenge that these technology trends present.

## Core Challenge 2:
## How should education change to meet the effects of these trends?

As the six technology trends continue, one of their primary combined effects will be to alter the knowledge and skills demanded or expected of workers by society. This is most obviously true from the discussion of Trend 6. Although potential crises should be avoidable by the development of new tools and techniques, it seems increasingly likely that those new tools and techniques will lead to engineered systems that are significantly different from what we are familiar with today. So, regardless of the use of technology to enhance or enable the education process, there is the question of how we should change education to best fit the future in which the new types of engineered system play so many key roles.

---

[22] As it happens, 1983 marked a brief and truly significant turning-point in ICT education in the UK. In 1982-83, the *BBC Computer Literacy Project* was entering its second year, and sales of the *BBC Micro* computer were exploding. Only a couple of years earlier, c.1980-81, an average secondary school might have had one or two early microcomputers such as the popular *Research Machines 380Z*, purpose-built for the UK education market from 1978-1985. By 1984-85, uptake of the BBC Micro, and investment by education authorities, had exceeded most reasonable expectations and dedicated suites housing a sizeable number of BBC Micros could be found in very many schools.

The last 25 years have seen "computer skills" taught at ever earlier stages in the UK curriculum, where now even reception-class kids are encouraged to gain comfortable levels of familiarity with the standard PC input devices (keyboards and mice), and Key Stage 1 specifies particular programmes of ICT study. Thirty years ago, such a prospect could have seemed highly unlikely. What then, are the prospects for education 25 or 30 years from now?

The indications from Trends 1-6 are that advanced computing systems will, via one route or another, shift from mechanical-scale systems to biological-scale systems. "Biological-scale" here refers to the number of interacting components making, the adaptivity of those components, the complex nature of the interactions between them, and the way in which those interactions combine to generate the overall behavior of the system. So students of computer science will need a familiarity with biological concepts, but (as was mentioned in the discussion of Trend 6) they seem likely to also need to understand aspects of economics and of mathematics and of social and organizational psychology. Suggestions that "computer science" as an academic subject might be subsumed by a new science of generic information processing systems are starting to look entirely plausible[23].

The proposition that in less than three decade's time there will be commonplace and widespread acceptance and usage of information-processing machines unlike anything that we are familiar with today might seem like an incautious long-shot. But, with hindsight, it is clear that exactly the same proposition could have been made at pretty-much any point in the past of 30 or more years ago and would (30 years later) subsequently have proven to be true. In 1978 the current-day prominence of ICT in the schools curriculum of 2008 would have been a prediction likely to attract derision. What then, will be the new components of the educational curriculum, necessary in 30 years time to accommodate the technology changes outlined here? Once again, the prospect of unknown unknowns cannot be ruled out; but in the apparent absence of such, there are a couple of issues that are worth mentioning to start the discussion.

First, so-called *systems thinking,* closely allied to the notion of *complexity science,* concentrates on those systems where the overall system-level behavior is difficult to predict even when given detailed knowledge of the behavior of the system's components and their interactions. Systems thinking and complexity science both lean away from reductionism in ways that certainly make their integration into traditional schools science curricula a challenging prospect. Yet, to people lacking any education in systems thinking, the technological future will likely prove very difficult to understand or comprehend.

Second, and more fundamentally, there seems little point in attempting to educate people in systems thinking and complex systems unless they have already been educated to a minimally sufficient standard of scientific literacy and associated mathematical competence. The concern here is not with a lack of people knowing particular facts about specific sciences such as physics, chemistry, or biology. Rather, there are good reasons to worry about the increasing proportion of people seemingly incapable of making well-informed judgements about scientific (or pseudo-scientific) arguments presented to them. Basic, foundational, subject-independent, abilities to design (or critique the design of) an experiment, to know how to deal with the quantities of data that might result from such an experiment, to visualize that data, analyze it, and draw statistically rigorous conclusions from such analyses, are a set of skills that should be given a much higher priority in future revisions of UK educational curricula.

Evidence for the dramatically negative effects of scientific illiteracy in the UK are not hard to find. In recent years we have seen the UK high courts quash convictions of innocent women who had been wrongly convicted of the murder of two or more of their own infant children. The deaths of these infants, it is now acknowledged, are likely to be due at least in part to the siblings' shared genetic heritage and/or home environment. That is, crucially, the deaths are not statistically independent events. In the trials that led to the wrongful convictions, the testimony of a medical expert was subsequently demonstrated to be based on flawed statistical arguments. The flaws in the statistical argument advanced by the

---

[23] Aaron Sloman writes wonderfully on this at: www.cs.bham.ac.uk/research/projects/cogaff/misc/cs-future.html

expert witness were shamefully simplistic, and yet the real shame of this story is not that a single witness gave unreliable evidence (that happens all the time). The real shame of this is that not a single person in the courtroom had sufficient knowledge of basic scientific method to recognize the errors in the witness's testimony. Not the judge; no-one on the defence team; and no-one on the jury.

There are, worryingly, many other such stories to be told involving ignorance of basic scientific method coloring what should be rational debate and decisions. Health concerns over power lines, mobile phones, GM foods, and the MMR combination vaccine have each been the topic of significant debate. If people's general school-age education had left them with a better understanding of scientific method and the associated tools of probability and statistical analysis, the debates over these issues might have been somewhat more rational and better-informed. Well-meaning politicians might realize the folly of endorsing something as having "no risk". Equipping our citizens with the intellectual tools that will help them make informed decisions about future technology trends seems like a decent thing to do.

## A Draft Plan for Addressing the Core Challenges.

Involvement from leading UK-based researchers with strong expertise in the intersection between education and the technology trends discussed here is crucial. Four[24] of the most appropriate individuals to engage with are:

Prof Lydia Plowman, Institute of Education, University of Stirling.
Prof Rose Luckin, London Knowledge Lab.
Prof Michael Hammond, Warwick Institute of Education, University of Warwick.
Mr Paul Shabajee, Hewlett-Packard Research Labs Europe.

As the primary expertise of these individuals is in technology-enabled or technology-enhanced teaching and learning, it is very likely that they can see opportunities, or threats, in the technology trends identified here that I am simply unaware of.

But in addition to seeking the views of those individuals on the content of this document, expert "second opinions" should be sought on the arguments of this document itself. Other UK computer science and complexity science academics whose views and comments might usefully be sought include:

Prof John Hogan, University of Bristol.
Prof Tom Rodden, University of Nottingham.
Dr Seth Bullock, University of Southampton.
Prof Noel Sharkey, University of Sheffield.

It would also be valuable to involve philosophers of cognitive science and artificial intelligence, for their breadth of exposure to ideas and research at the interface between advanced computing technology and current understanding in neuroscience, psychology, and philosophy of mind. Individuals who might be involved here include:

Prof Aaron Sloman, University of Birmingham
Prof Andy Clark, University of Edinburgh.
Prof Maggie Boden, University of Sussex.
Dr Michael Wheeler, University of Stirling.

To generate robust evidence, and to capture fresh and challenging thinking, but on limited budgets and with no opportunity for new empirical research, good value should be gained from holding two small-group workshops, each lasting one or two days, and each involving, for the sake of argument, the twelve individuals named above, plus a few more, plus, say, four representatives from DCSF and FutureLab. So, perhaps twenty people. The first workshop would bring these twenty individuals together around a table for the first time,

---

[24] I've arbitrarily limited these three lists of possible experts to four names for each field of expertise. The lists are not intended to be exhaustive, and for each list there are several other UK experts who could be invited to contribute if required.

and so the emphasis would be on building familiarity and common understanding among those present. This Challenge Outline paper should be used as a “strawman” document, i.e. as an initial catalyst for the discussions at the first workshop. For instance, each attendee could be asked to spend some time critiquing the document, or revising or extending some of its arguments, or introducing new topics for consideration. The intention of the first workshop, then, is to initiate an informed discussion among an appropriate mix of experts focused on the two core challenges: how technology-enabled education is likely to be affected by future trends in advanced computing trends such as those discussed here, and how education can best serve the society of ten or twenty years time that will be shaped by these technology trends. The intent of holding a second workshop, with the same attendees meeting for approximately the same duration, is to conclude that discussion and finalize the capture of its outcomes. For that reason, the second workshop would be roughly eight to twelve weeks after the first, with the interim period being spent by the participants working, individually, or in small groups, or *en masse*, on various sections of the new document. Communication by email and phone should be sufficient to support these interactions (that is, there would not be a need for a large travel budget to support visits among the participants during the period between the two workshops). If that process actually results in a total re-write of this document, then so be it. The primary outcome of the second workshop would be a set of recommendations about the future nature, role, and organization of education in the light of the identified technology trends. There would not be a pre-commitment to arriving at a consensus view: if there are strong differences of opinion, then those differences should be reflected in the final document. The final agreed document should be made publically available by publishing it on appropriate websites (e.g. the websites of DCSF, and/or of FutureLab), and possibly also a web discussion forum could be set up to capture and disseminate subsequent responses and contributions from a wider audience.

To ensure that the work is credible with the wider academic community, the intention is that the final document would form the basis for one or more papers submitted for peer-review to appropriate leading international academic journals. Significant engagement from other key project stakeholders – such as ministers and policy-makers – would be best secured by short summary briefing documents listing key outcomes of the discussion workshops, and also by securing positive media coverage. For this reason, at least one of the experts participating in the discussion workshops should have expertise in public engagement.

None of the people listed above have been approached about their proposed involvement. If funding is approved for these two workshops, then planning for the first to be in May 2008 and the second to be in September 2008 would give participants sufficient time to prepare for each workshop, and would be in line with the *Beyond Current Horizons* overall timetable.

## Author Biography

Dave Cliff is a Professor of Computer Science at the University of Bristol. He has a BSc in Computer Science and an MA and PhD in Cognitive Science. He previously served in academic faculty jobs at the University of Sussex (UK), at the MIT Artificial Intelligence Lab (USA), and at the University of Southampton (UK). From 1998-2005 Cliff worked as an industrial research scientist: formerly as a Department Scientist at the Hewlett-Packard Labs European Research Centre in Bristol, where he founded and led HP’s Complex Adaptive Systems research group; and latterly as a Director in Deutsche Bank’s Foreign Exchange (FX) Complex Risk Group, on Deutsche’s City of London FX trading floor. In October 2005, Cliff was appointed Director of the UK national research and training initiative in the science and engineering of Large-Scale Complex IT Systems (LSCITS). Phase One of the LSCITS Initiative commenced in October 2007, is funded by almost £10m of UK public funds, and will involve more than 250 person-years of research effort: full details are available at www.lscits.org. Cliff is author or co-author on over 70 academic publications, inventor or co-inventor on 15 patents; and he has undertaken advisory and consultancy work for a number of major companies and for the UK Government. He has given well over 100 invited keynote lectures and seminars; and he and his work have frequently been featured both in the press and on TV and radio. He is a visiting professor at the University of Leeds, a Chartered Fellow of the British Computer Society, a member of EPSRC’s ICT Strategic Advisory Team, and a member of the CPHC.