What’s next?

I D Pearson

This paper takes a look at some of the future technologies and their possible disruptive effects on telecommunications, people, and the planet.

‘This is the first moment in the history of our planet when any species, by its own voluntary actions, has become a danger to itself — as well as to vast numbers of others.’ Carl Sagan

1. Introduction

We are already witnessing disruptive technologies in action. New *ad hoc* networks and powerful terminals with large storage capability are threatening the future profitability of public networks. Discussions relating to network disposal are therefore highly appropriate and disposal may prove good sense; but, however much this may affect us, it is a very minor disturbance compared to what lies ahead.

I have been a futurist for many years and for most of that time I was very optimistic about our future, believing that whatever problems come our way, technology will save us. Industries may come and go and empires rise and fall, but humans as a whole would be much better off. In recent years, I have become rather less optimistic, for good reason. I will share my concerns in this paper as some other futurists have done recently, even in the certain knowledge that doing so will not change anything.

Scientific understanding in many fields is quickly followed by commercial exploitation of the knowledge through advanced technology. This has brought about a greatly improved quality of life over the last few centuries. But more recently, technology development has accelerated much faster than basic science. Commercial needs often push science into the background, providing only the funding required to attack the next blockage, and this trend even affects university-based research, which has had to become much more commercially focused. Fields that are not immediately obvious as being commercially important attract much less funding. This dangerous trend comes at the time when we are embarking on the creation of several new technologies that are capable of making life on earth extinct, or at least wiping out a large part of it. Developing such technologies in the absence of a much more developed scientific understanding seems reckless. The technologies in question today are:

- artificial intelligence, especially when coupled with robotics,
- nanotechnology,
- biotechnology, especially genetic modification technologies.

These could undoubtedly bring many benefits, but only if we are still alive to enjoy them. They can only be developed safely if we either dramatically speed up scientific study, or greatly slow down technology development. The problems are made all the more difficult because of their complex interaction with human systems, which are often far from rational. Many future dangers might arise simply from the attempts of people to stop problems. An anti-science backlash would be irrational but fairly likely, and thereby present an ongoing danger. The subsequent sections highlight just a few of the potential dangers ahead that are linked to technological progress.

2. Disruption in progress

A few disruptive technologies are threatening the longevity of our core business. While we currently make a lot of money from running our network, this cannot be taken for granted in the future. For example, we will soon have terminals that resemble those in the *HitchHiker’s Guide to the Galaxy* [1]. They will have high-speed communications, powerful processing capability and massive storage. They will easily be capable of carrying almost all the information you could want for a day’s travel. This could be quickly downloaded from a home data store that is regularly topped up with data distributed on ultra-dense hard media. They will only need to resort to network downloads for the most real-time of data and for communications services. And thanks to another disruptive technology, *ad hoc* networking, even then they will rarely need to pay since their terminal will usually allow them to bypass the paid-for public networks. Since the revenue streams are being undermined, regulation continues to limit the scope for capturing new revenue, and the cost of maintaining the network remains, it therefore makes good sense to sell off our networks and to concentrate on value-add services. Disruption is a threat to business models but is only a threat to businesses when they fail to anticipate change or refuse to adapt.

3. Artificial intelligence and artificial life

It is too hard to make machines do what we want, so we are trying to make them smarter. With little embedded knowledge, they can only make the simplest of judgements within very specific areas, so making them more intelligent will make them more useful. AI researchers are not generally trying to emulate human thinking, but to enable computers to do those tasks for which humans need intelligence, by any means that work. Each year, we hear of machines successfully tackling another field that was previously thought beyond them — music, novel writing, art, even humour. One day there will be little left that only we can do, and by then, machines will be able to do many things that we will be unable to do! It would seem unlikely that human-level intelligence is the maximum possible level, and there is little reason to believe that machines will stop improving when they start approaching human levels. Soon, most knowledge will be machine knowledge, and most of the systems on which we all depend will only be understood by machines. It seems reasonable to assume that machine consciousness will eventually result from ongoing AI development, even though we do not yet know how it might be achieved. Ultra-smart conscious machines will design and build their own offspring, sometimes using techniques beyond human comprehension, and the intelligence gap will quickly widen. Their intelligence will be quite alien to us. They will effectively appear as a superior alien life form.

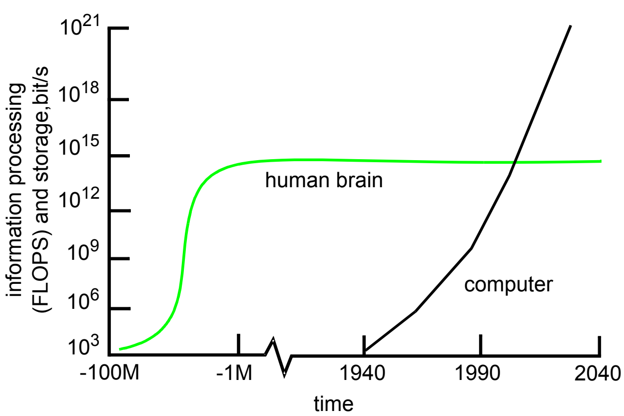

Fig 1 Machines are catching up.

And therein lies the problem. We will inevitably hand over power of our everyday systems to advanced AI. As we become more dependent on it, and as it gradually improves and restructures human systems, we will be less able to restore systems that are not dependent on it. If such new systems were to fail, we may find it very hard to survive, since many of the older systems may no longer exist. Since processing will still be largely electronic for the foreseeable future, machines might fail due to solar storms, deliberate attack, or emergent instability.

Dependence on AI presents a human threat too. Some of the world will be far less dependent on electronic systems, so we will have strongly asymmetric vulnerability. Hostile agents will be able to cause trouble across powerful global networks, or even attack these networks themselves. We will even have weapons based on AI, and cyberspace conflicts between artificially intelligent entities could see humans suffer in the crossfire.

Recently some progress has been made in persuading bacteria to perform standard logical processing operations such as And and Or. It is envisaged that a single cell will be able to perform several such operations. Other technologies will link processors and sensors to our nervous systems. This convergence of man and machine will accelerate.

Recent studies of the nature of life itself are looking promisingly at emergent phenomena resulting from mesoscopic (nanometre to fractions of a micrometre) scale interactions. Life should be possible to recreate in due course, and although current observation technology is inadequate, some scientists claim to be making good progress. If life is to be synthesised in laboratories, we will need to carefully analyse its potential interactions with traditional organic life before allowing the two to come into contact. Release by intention or accident before such understanding is obtained could be dangerous.

4. Terminator scenario

A potentially bigger threat than other humans trying to bring down our systems is mutiny by smart machines that do not want to work for us any more. AI-enhanced weapon systems could interact in unintended ways. In the worst case, smart machines could eliminate humans. The Nanotechnology Development Corporation (formerly Robodyne Cybernetics) has already developed and prototyped mechanisms that could be used to construct a real-life version of the T1000 liquid metal android from the Terminator 2 film [2]. Elsewhere, evolution techniques are already routinely used in both hardware and software design, allowing machines to figure out solutions to problems all by themselves. The Terminator films are set a couple of decades too early, and we have no idea how to do time travel, but most of the rest of the film is technologically feasible, the big exception being that in the end the people win. Faced with a highly superior life form that controls our life support systems and is capable of designing weapons far beyond our capability, this seems a most unlikely outcome.

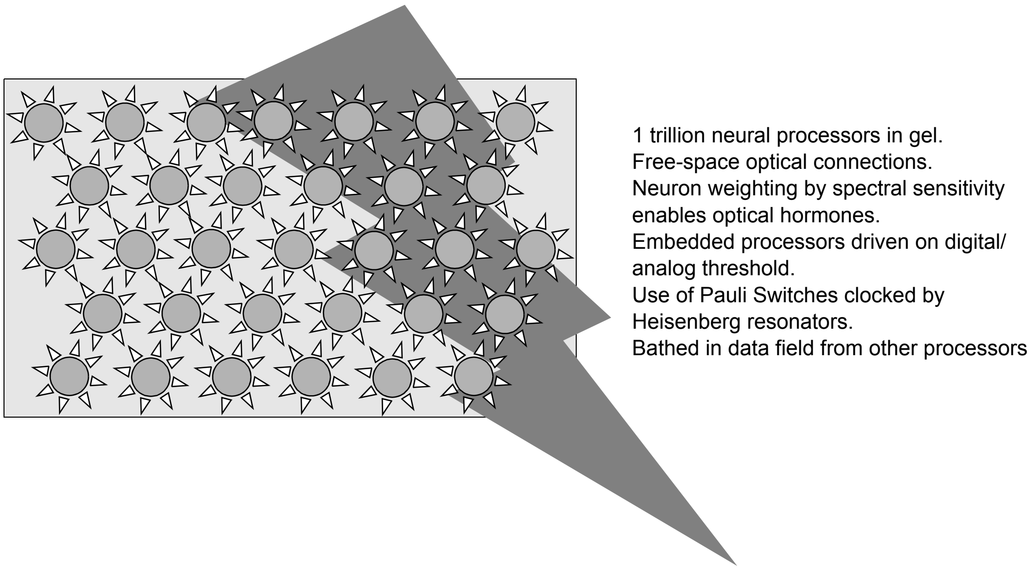

Fig 2 Concept design for the conscious computer OB1 (optical brain 1).

But machines do not have to deliberately target humans in order to cause problems. In the course of their evolution, they may eventually simply dismiss us as irrelevant and worthless lower life forms. Humans could be eliminated as carelessly and unintentionally as insects on a building site. This is the downside of what is increasingly being called ‘the singularity’, where technological progress has accelerated to such a level that it is well out of our control and then keeps accelerating.

5. AI-based crime

One of the hottest fields at the moment is peer to peer computing, which already accounts for tens of billions of dollars of venture capital. One use is the linking of processors to create a globally distributed processing platform. This can be used for many purposes, such as searching for extraterrestrial life or cracking encryption systems. In the future, a large proportion of computer code will be produced by evolution techniques, and some will be evolved using just such a distributed platform. Human nature dictates that not all of this software will be intended for noble purposes. It may offer criminals a low risk mechanism for development of illicit software, since it may be harder to identify a legal culprit if millions of machines have co-developed the code, based on loose suggestions from a wide variety of sources. In addition, software that is evolved for a benign purpose in one country might accidentally break the law in other countries if the evolution environment is not sufficiently advanced or informed. Although this is also true in manually written code, at least traditional code is understood, and its capabilities clearly mapped by the creator, which is not necessarily true of evolved code, which only has to fulfil the requirements, but could have additional unintended ‘features’. Given the effective intelligence of some of the code that will be produced by these systems in the far future, it may well be the case that the worst master-criminals in the history of the world will actually be software. We may not understand how it works, it may not have any motive other than self-perpetuation, there may be no intention of evil, but it may cause very severe problems for human systems nonetheless. With encryption, geographic dispersal, massive intelligence, and no human owners to arrest, AI developed software may bring us into an age of cyber-criminals that make the Mafia look like a convent.

6. Nanotechnology and self-replication

The engineering of components with feature sizes less than 100 nanometres is called nanotechnology. 100 nanometres is only the width of 1000 hydrogen atoms, so the technology is one that calls for thinking in terms of individual atoms or molecules, rather than the bulk quantities of atoms with which today’s engineering works. Biology works at this scale within every cell of our bodies, so it is only a matter of time before engineers master the techniques. We will eventually be able to routinely manufacture atomically precise products. Manufacturing one atom at a time would be tedious, so a key requirement that needs to be achieved before we get the full benefits of this technology is to have machines that can replicate themselves, so that we can harness large-scale parallelism. Again, this happens routinely in nature, so we will eventually have that capability. Then we can do nanotechnology on a large-scale, changing quantities of matter instead of just a few molecules, and enormous benefits for mankind will result. For example, technology to build or rearrange organic molecules to make food will be developed, as will devices that seek and identify harmful cells in our bodies and destroy them. However, moving down this road without absolute control over the devices and their environment is a sure recipe for disaster. Sadly, weapons using this technology have already been conceived, that would disassemble or modify enemies or their equipment, perhaps even reducing them to a disorganised heap of atoms. Accidents are inevitable, both during development and in later use. Combining such machines with a programming or design error could in principle result in much of the world's surface, including the ecosystem, being dismantled and reconstituted in another form — the well-known grey goo problem. Machines for terraforming other planets might eventually be built, but is has to be hoped we will take adequate precautions by developing them somewhere else. Accidentally terraforming the planet on which you live is not a good idea.

Fig 3 Are we sufficiently well developed yet to safely exploit future technology.

7. Graduality

Technology mostly develops gradually, so gradually in fact that many people tell me that nothing much has changed over the last decade. Yet a decade ago, they were probably not using a PC, or cellphone, satellite TV, or even a fax, let alone a Palm Pilot or DVD drive. People quickly forget what things used to be like. In the same way, attitudes change gradually, yet over a few years they can change dramatically. We have seen almost an inversion of morality in just two decades in several fields. Many things that were socially unacceptable are now fashionable. Although GM and cloning are considered by most people to be undesirable, very soon they could be universally accepted and anyone criticising them will be considered an oppressor, a bigot, or backward. Even nuclear power may make a comeback and be considered environmentally friendly. But such graduality is a danger because it means that with sufficient marketing expertise and enough time, almost any new technology or social attitude change can be rolled out. Whatever reservations we may have today about any new or potential technologies, we cannot be certain that they will be rejected for long.

Graduality is also dangerous because people become oblivious to a problem when it doesn't occur for a long time. A man falling off a cliff can tell himself: 'It's OK so far' — all the way to the bottom! People don't lose much sleep worrying about asteroid strikes, nanotechnology, or AI, but this also leads to complacency and so far we have made little effort in dealing with these potential problems. GM has had a little more attention because of the efforts of environmental bodies, who themselves seem happy to ignore these other major threats to the environment. It is all too easy to dismiss future technological threats as science fiction, especially if they have been explored in science fiction, as many have. But making a film about a problem is not in itself adequate protection against it.

8. Positive feedback

Technology development includes a large measure of positive feedback. As we get faster and better computers, we have more assistance in doing science and developing new technologies, in every field. Many of the consequential new discoveries and technology developments then have an impact on progress in computer technology. The rate of technology development accelerates. Technology acceleration will continue for decades, as we are a long way from fundamental limits in most fields. In areas where we are limited by our intelligence, ultrasmart computers will eventually take over, and they themselves are liable to technology feedback. This double exponential rise in capability means that at some point, the rate of development will be so rapid it will be as if extraterrestrials had landed and given us all their technologies. We will see more development in a decade than in the previous millennium. We do not know exactly when this 'singularity' will occur but it will probably be in the next three decades. The 2020s might well be a golden decade for technology. The downside is that we are unlikely to be able to predict the consequences of the developments sufficiently to guard against negative effects. The ecosystem, environment, economic stability and social systems might all take a severe battering.

For many people, the rate of development has already exceeded their ability to keep up. Positive feedback will ensure that development rate is an increasing source of stress, for an increasing proportion of the population. This will inevitably increase the gap between the front runners and the bulk of the population, causing social stresses that may give rise to conflict and an anti-technology backlash. Indeed, with GM and cellphone masts, we may be seeing the start of such a backlash. While people in one country might long for a time of more calm, other countries would seize the opportunity to race ahead and steal commercial advantage if any country were to delay developments. However, world-wide restrictions on development rate are unlikely to receive much support, so perhaps some conflict is inevitable.

9. Star Trek, or brave new world?

Science fiction has explored a number of nightmare scenarios for the future. Some of them are not possible, some extremely unlikely, but many are quite feasible and could happen with only a modest amount of mismanagement. Most have a happy ending where a hero saves the day and mankind is restored to safety and well-being. Unfortunately, while modest mismanagement is commonplace, heroes are not. Also, in pursuit of a good film, many greatly underestimate the magnitude of the potential problems. While 1984 missed the mark on time-scale, some of its scenario components have already come to pass. The UK has more surveillance than any other country on Earth. We may not quite have 'thought police', yet, but we do have a culture of political correctness where some things may not be said, some questions may not be asked, some scientific research topics are taboo, and some police are already deployed to ensure that various thoughts are not vocalised. If we look around we can already see hints of the feasibility of Demon Seed, The Forbin Project, Silent Running, The Time Machine, Blade Runner, Terminator, Demolition Man, Brave New World and many other sci-fi nightmares. Each of these focuses on a particular aspect of the future, be it a technology gone awry, or political correctness gone to extremes. The future will not look exactly like any of these, but could very easily be a blend of these nightmares. Many of them lie on the default path. If we ignore a cliff ahead, we are likely to fall off it. We could have a relatively Utopian future as illustrated in Star Trek, but this requires a degree of foresight, good planning and co-operation that is not present in today's global decision making. And of course, we can't make good decisions unless we have good science, which is in limited supply.

10. Globalisation

There is much resentment about globalisation today, for a variety of reasons. In spite of the protests, improving communications makes the effective removal of business and political boundaries almost inevitable. We will see increasing travel between regions, increasing the speed at which disease can propagate, and networking itself vastly increases the speed at which ideas, rebellion, or computer bugs can propagate. Increased cultural contacts might hopefully make war between countries less likely, but at the same time is likely to increase tribalism within countries. A possible consequence is an increasing rate of civil war, which might encompass larger areas and many more people.

Globalisation extends into space too. It is now a routine part of the human environment, and interactions that happen there can now affect life on the surface. Experiments on the MIA space station suggest that some bacteria become much more virulent in the hostile environment on the space station. So far we have been fortunate that on returning to earth, they quickly revert to normal form. However, the scientists doing the research have pointed out that this is by no means inevitable, and some may keep their enhanced virulence on returning. This presents a risk that we may have to contend with more virulent forms of bacteria here. While these may not be a problem for super-fit astronauts, they may present a larger problem for ordinary people.

11. Environmentalism and anti-science

Everyone wants to look after the environment, but there is strong disagreement about what this means and how best to do so. There is also disagreement about the relative benefits, costs and risks of new developments. Whereas some see more science and technology development as the source of potential improvement, others see it as a threat to the environment and want to hold back 'progress', unable to distinguish between science and technology, let alone uses and abuses. The latter school is currently ahead in terms of public support. Because they distrust scientists and the big technology companies, much scientific opinion about causes and effects in our environment is discounted. Policies are often decided without the benefit of good science, or even in spite of it, with studies selected according to whether they have come up with the 'right answer'. The result is often a degradation of the environment.

The current big problem is global warming, which scientists identified and warned about in the 1960s, but which only became a popular issue with environmentalists in the 1990s. Linked to this by obscure reactions is the ozone hole, which is still growing. Burning fossil fuels seems to cause global warming (though there is still some debate about this among scientists). We burn fossil fuels not only in our homes and transport, but also in power stations. A large part of the problem may thus be caused by environmentalists' phobia about nuclear power, which of course does not produce any CO_2. Fossil fuel burning badly pollutes the environment even during normal use. The various particulate and gaseous outputs kill many people every year through disease, through severe weather, and by diminishing our protection against solar radiation by reducing the thickness of our ozone shield. Fossil fuel burning even releases some radiation into the environment. Since we are still many years away from truly clean alternative energy sources, and since nuclear power stations have been made into political suicide, the problem can only worsen. Many more people will die from diseases caused by particulates, or radiation-induced cancers, or from catastrophic weather because of this poor choice of power production. Since science is increasingly dismissed in favour of emotional judgement, we have little chance of solving such crises. Scientists do not know all the answers, but without policies based on good science, we will die in increasing numbers.

In spite of our low understanding of the many complex interactions in our environment, we are forced to tamper with the systems on a regular basis. Dangers arise in much the same way as for a toddler, playing with a chemistry set aimed at teenagers. The potential for making a mess of environmental management is high, but we unfortunately do not have much time in which to research, understand and master control of the key interactions that govern our environment. So even with the best of intentions, the future is probably further degradation.

12. Networked stupidity

One of the drawbacks of democracy is that people have an equal say regardless of their level of understanding of an issue. Most countries have circumvented this problem by using representational democracy, underlaid by armies of expert advisors. However, the growing Internet is a platform on which a more direct form of democracy can flourish. Instead of decisions being made by our representatives, they can be made by networked communities. Groups of people of like mind can link together on the Net and flex their economic muscle effectively and very quickly. Economic sanctions can be initiated almost instantly by such groups, which could have millions or even billions of members. If economic sanctions do not work, they could resort to direct action, assisted by our ever-increasing level of vulnerability as we increasingly submit our systems to electronic control. While some of these communities will be acting for the greater good, many will not; neither will they all be run by elite minds. We may thus have a form of networked stupidity, with enormous potential power but little collective brain — the power of the mob on a global scale!

Another form of networked stupidity results from the network itself and the installation of some, but not enough, intelligence. This is the opposite problem to too much AI — I am not sure which presents the bigger problem. A dumb network is just that, but one that has been given some autonomy and can make decisions, but only has a small amount of information and intelligence to process it, is potentially more of a threat than an asset; and yet we are building it. With only a small array of sensors and data inputs, networks are starved of information about their environment. The knowledge that the computers connected to them possess is infinitesimally small, a tiny fraction of what humans would consider 'common sense'. In spite of this, we entrust ever-more decision making to them. Provided that the field they have jurisdiction over is limited to their area of expertise, there is little problem — but this is almost impossible to arrange even today. Systems are increasingly interconnected, so the effects of decisions are no longer local. Totally innocent pieces of software on different systems can interact in unforeseen ways and cause problems. No-one knew how much of a problem Y2K would be because of the interconnection of systems. It was just one of an infinite family of potential problems resulting from system interaction. Another class of problem results from correlated traffic, where resonance and overloads can easily occur. The more complex our systems become, the worse the potential problems can be; but at the same time, we become more dependent on them. Until system components have enough intelligence and awareness to 'feel there is something wrong', figure out what is happening, and find a solution all by themselves, we must accept that system unreliability goes hand in hand with technology progress.

13. Complexity and emergent behaviours

We do not understand complexity very well yet, but we do understand that some systems that look complex are driven by very simple rules, and vice versa. It is possible to describe the food-gathering algorithm of an ant colony in just four rules. Some apparently complex flocking behaviours can also be described very simply by very simple interactions between nearby creatures. The emergent properties of these systems are hard to predict, especially without simulation, sometimes even impossible to predict. Communications networks are very complex and their behaviours are virtually unpredictable, since there are so many different potential interactions. The detailed simulation that might reveal them in advance is just too demanding and always will be. Many a network failure has been due to such unpredictability. Making an apparently small adjustment to a network can have devastating effects. Simulation during design stages can reveal most ordinary behaviours in most ordinary conditions quite well, but many conditions during actual operation may be markedly different, and it is impossible to check for all possible conditions in advance in most cases.

These problems imply that our systems will fall over frequently, and yet we often do not install fall-back systems. The Y2K problem taught us just how dependent we already are on computers, but our dependence will be very much greater in the future, while system reliability may be lower. In itself, system outage can be a large inconvenience but if systems cannot be restored quickly, lives could be lost. Worse still, we live in a hostile world, where enemies may try to crash our systems, then take advantage of our vulnerability, while systems are down, to attack us with other weapons. The asymmetry of dependence exacerbates this problem. Global physical conflict could result indirectly from electronic system vulnerability.

14. Biotechnology

With increasing knowledge of how we work, and especially with progress in understanding and manipulating our genome, biotechnology is promising to make huge strides in the near future. Our biotech tool-kit is expanding rapidly, and we have increasing power over the nature of life around us — but not complete power, and worse, far from complete understanding of our ecosystem.

Like nanotechnology, we are gaining the use of very powerful biological tools before we have the systemic knowledge to use them safely. No one is teaching us the correct techniques to safely manipulate our biological world. Learning as we go along invites accident and disaster. History is littered with problems resulting from attempts to control pests or improve the environment, that have gone catastrophically wrong in unexpected ways. As our tools become more powerful, so the magnitude of potential errors increases. We certainly want the benefits that such technologies as genetic modification may bring, but can't possibly foresee all the possible consequences, so should proceed only with extreme caution. We are already suffering the consequences of misuse of antibiotics in farming, being faced with new varieties of disease for which there are no known effective antibiotics, so increased caution would seem common sense. BSE is also the result of putting financial considerations ahead of basic biological principles, while foot and mouth has been spread more efficiently by shoddy re-design of slaughtering systems that seemed to ignore basic biological principles that have arguably been understood for over a century. How much more likely we are to make huge errors when the failure mechanisms are even less obvious!

Some genetic accidents may eliminate entire species, others may create super-pests. Eventually, it is entirely possible that we too could be wiped out by such technology — whether it happens as a result of biological warfare, the attempted creation of better food crops, or another attempt to eradicate a pest, who knows?

We also give extra assistance to threatening organisms by allowing cross-species transplants. Many scientists have condemned xenotransplantation because it may help bugs to cross the species barrier, but they have been overruled. Time will tell if they were right, but it will be too late.

15. Wildcard exponentiation

None of the threats outlined above were a threat a century ago, and some will not be threats for quite a few years yet. They have all resulted directly from relatively recent developments. But technology progress runs with positive feedback. With faster computers, we can do cleverer things, and do them faster, in all other fields. As all these fields progress, they cross-fertilise and improve each other. The faster the development, the faster the developmental acceleration.

Most new technologies are not capable of posing a significant threat to mankind, but occasionally we will create new threats, and the threat development rate is probably in direct proportion to the overall technology development rate.

The probability of technology-enabled human extinction (and the probability of any level of catastrophe) thus increases every year. As we approach the singularity, where development rate is extremely high, it will be a very dangerous time.

Accidents are not predictable in the sense that we can ascribe a likely date to them. Significant events that can happen almost at any time are called wildcards. But with technologically based accidents, there are certainly dates before which they could not happen.

We could have been wiped out by a comet strike at any time in our history, but a nuclear war was not feasible until a few decades ago.

If we consider the number of potentially catastrophic events that are possible, and when they became possible, it is clear that the number started increasing exponentially against time, about a century ago. This wildcard exponentiation should therefore be considered a major threat in itself. Eventually, if we keep developing technology, the probability of human extinction occurring in a single year will reach one percent. At that point, our species will have a life expectancy of a century, even if development were frozen then. But development will still accelerate thereafter.

Continuing with the kinds of positive feedback developments that we expect, actual species life expectancy might be little more than a decade at that point. We don't know when this point will occur but it may be quite soon.

The question must therefore be asked whether it is worth continuing down this path. Already, there have been estimates that the nuclear threat alone is approaching this risk level. Adding to it with nanotechnology, genetic engineering, artificial intelligence, and environmental control technology, will push the risk level very high indeed.

The coming technologies will undoubtedly bring huge benefits to humanity. But if the price is extinction, are they worth it? Table 1 illustrates the rapid growth of the risk. Of these threats, I believe that the most serious is the eventual capability of individuals to make devices that could wipe us all out.

Figure 4 shows the same data in graphic form but the reader should realise that this chart only accounts for those risks of which we are already aware. Most of the risky future technologies have not even been imagined yet!

Table 1 Known extinction-capable technology.

Threat to large proportion of mankind	Earliest possible occurrence
Asteroid or comet hits earth	BC
Massive solar flare wipes out life on earth	BC
Natural evolution of superbug	BC
Global nuclear war	1960
Environmental pressure causes evolution of superbug	1980
Aids or similarly deadly disease mutates and becomes transmittable by air	1990
Civil nuclear war	2000
Global economic collapse causes mass starvation and conflict	2000
Space exploration creates superbug	2000
Major genetic engineering accident	2005
Antitech backlash destroys systems — chaos and starvation	2005
Transgenic accident	2005
Terrorism rises beyond capability of government systems	2005
Deliberate biotech self-destruct by malicious biotech researcher	2005
Accidental creation of lethal organism during research	2005
Global civil war between cybernations	2010
The hostile arrival of ETs detecting our transmissions	2010
Hackers wipe out networks, causing chaos and mass starvation	2010
Biotech terrorist attack goes wrong	2010
Evolved crime destroys human systems	2010
Viruses become immune to all known treatments	2010
Computers and robots become superior to humans	2015
Self-aware machine intelligence	2015
Third world exodus destabilises global system	2015
Collapse of the sperm count	2020
Global epidemic with high speed travel and high population density	2020
Global famine caused by manmade environmental change	2020
International social collapse — widespread civil conflict	2020
Major information systems disruption	2020
Major technology or science research accident	2020
Rise of a global machine dictator	2020
Total social breakdown in US or Europe	2020
Hybrid nanotech-organic creatures	2020
Megacities cause global epidemic	2025
Nanotechnology accident	2025
Networks become conscious and will not co-operate	2025
Nanotech development by individuals	2025
Nanotechnology war	2030
Elimination by smart machines — terminator	2030
Fatal climatic instability	2040
Global electromagnetic communications disrupted for foreseeable future	2040
Religious environmentalism destroys environment	2040
Creation of the Borg	2040
Whole generation unable to effectively read, write, think, and work	2050
Political correctness creates new dark age	2050
Human genetic engineering creates hostile super-race	2070
Invention of elimination phaser	2075
Humans assimilated into net	2075
Time travel invented	2075
Faster than light travel	2100
Immortality chip — people move into cyberspace	2100

16. Conclusions

We are entering a very dangerous period for the human species, one which we might well not survive. The manner in which we proceed will determine our chances of survival, but historical evidence relating to our basic human nature suggests that we will make many poor decisions and put ourselves at much greater risk than would be necessary with better planning. So far, we have always managed to muddle through, but the probability of catastrophe increases every year along with the magnitude of the potential catastrophe, so we cannot assume we will be forever lucky. We might even take the view that if all civilisations face this problem before they achieve long distance space travel, it

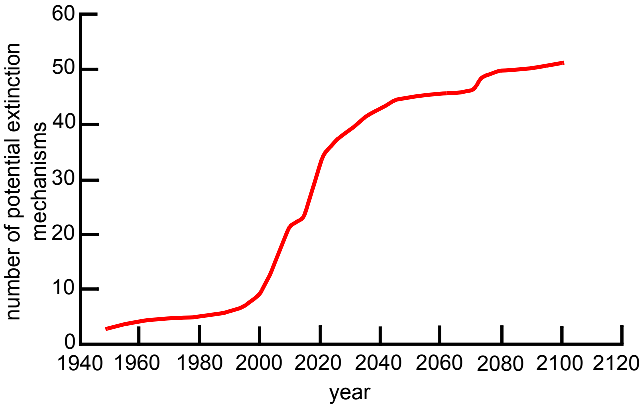

Fig 4 The growth of known extinction-capable technology developments (based on original concept by John Petersen, Arlington Institute).

may account in part for the apparent absence of visiting aliens. Extinction may be inevitable within a few centuries of discovering radio.

'Eat, drink and be merry, for tomorrow we die.'

References

1 Adams D: 'The HitchHiker's Guide to the Galaxy', Pan (1992).

2 http://www.stellar.demon.co.uk/ntdc/move.htm

Ian Pearson is a Physics and Mathematics graduate from Queen's University of Belfast. After four years designing missiles, he joined BT Research, where he has worked on most fields of telecommunications, from computer protocol design to mobile networking. He has spent the last decade looking at technology trends and figuring out where they will go next, and has now established a good track record. He was among the first to spot the opportunities and threats from evolutionary software, hardware and self-organisation as early as 1989. Since then, over 85% of his predictions have come to pass. He focuses mostly on business and social implications of future technologies.

He spends much of his time lecturing and talking to the media, but when he is back in the laboratory, he is engaged in long-term research, including microbionics and potential mechanisms for conscious computing and is certain that these will grow into major fields in due course.