The JOURNAL *of*
GEOETHICAL
NANOTECHNOLOGY

Terasem Movement, Inc.

Volume 1, Issue 3
August 2006

Ray Kurzweil, Ph.D.

Ray Kurzweil, a noted inventor and futurist, is author of The Singularity Is Near: When Humans Transcend Biology and co-author of Fantastic Voyage: Live Long Enough to Live Forever, and four other books; and has won numerous awards, including the 2001 Lemelson-MIT Prize, the world's largest in invention and innovation; and received the 1999 National Medal of Technology from President Clinton. He argues that while technology brings significant perils, we can't simply relinquish it. The key, he believes, is in understanding and learning to harness the accelerating progression of technology.

This article was adapted from a lecture given by Ray Kurzweil at the 1st Annual Workshop on Geoethical Nanotechnology on July 20, 2005 at the Terasem Retreat in Lincoln, VT.

Terasem Movement, Inc.
201 Oak Street
Melbourne Beach, Fl 32951

The JOURNAL of GEOETHICAL NANOTECHNOLOGY

Volume 1, Issue 3
August 2006

How We Can Manage Our Way Through the Intertwined Promise and Peril of Accelerating Change

Ray Kurzweil, Ph.D.

This article was adapted from a lecture given by Ray Kurzweil at the 1st Annual Workshop on Geoethical Nanotechnology on July 20, 2005 at the Terasem Retreat in Lincoln, VT.

Ray Kurzweil, a noted inventor and futurist, is author of The Singularity Is Near: When Humans Transcend Biology *and co-author of* Fantastic Voyage: Live Long Enough to Live Forever, *and four other books; and has won numerous awards, including the 2001 Lemelson-MIT Prize, the world's largest in invention and innovation; and received the 1999 National Medal of Technology from President Clinton. He argues that while technology brings significant perils, we can't simply relinquish it. The key, he believes, is in understanding and learning to harness the accelerating progression of technology*

A key question facing us to today is: How we can manage our way through the coming challenges, emphasizing the promise while avoiding the peril? We did not manage to avoid all the peril of technology in 20th Century. Fifty million people died in World War II, which is only one out of hundreds of wars in this century. Nevertheless, all of this conflict and thesee major events did not have any effect in inhibiting the pace of progress. If anything, it accelerated the ongoing progression of technology. Yet we can still clearly see that there is promise *and* peril.

I have some ideas about strategies on containing the perils. I believe this issue is the fundamental challenge facing human civilization. As powerful as 20th Century technologies were, 21st Century technologies are immensely more powerful. They will enable and multiply both our creative and destructive impulses.

One of the biggest issues I try to communicate is to distinguish the intuitive linear view of history from what I call the exponential view. It is remarkable how many otherwise very sophisticated people have a linear view of the future.

Frequently, I've been peered with Bill Joy as optimist and pessimist respectively, but I invariably end up defending Joy on the feasibility of the dangers. For example, in one recent dialogue, a Nobel Prize winning biologist said, "Oh, we're not going to see self replicating technology for 100 years." I said, "Well, where do you get that from?" He said, "It's hard to measure, but my intuition is that we've solved 1% of the problem over the last year." I said "That's actually my intuition also, and it will take 100 years at today's rate of progress. But the rate of progress is not a constant; it's accelerating."

This is not just a casual observation. I have been measuring this. I actually have a team of ten people that gathers key data, key measures of technology in many different areas, and we build mathematical models. I got into this because of my interest in being an inventor. I realized that my inventions had to make sense when I finished the project and the world is a very different place. Most projects fail, not because the R&D team cannot get it to work. Today, as I read business plans from people, 90% of those teams will do exactly what they say if they're given the resources. Yet 90 to 95% of those projects will still fail because the enabling factors needed for market success are not in place. Thus, I became an ardent student of technology trends.

This brings up a key issue, can you predict the future? The common refrain is that you cannot predict the future. It turns out that certain things are hard to predict. For example, will Google stock be higher or lower than it is today three years from now? That is hard to predict. What will the next wireless standard be: WiMax, 3G, CDMA?

Those things are hard to predict. But if you were to ask me, what would the cost of a MIPS of computing be in 2010, or the cost of sequencing a base pair of DNA in 2012, or the spatial and temporal resolution of brain scanning in 2014, I can give you a figure and it is likely to be correct. I have been doing this for 20 years and these trends have been tracking very accurately. There are very smooth exponential trends that go way beyond Moore's law.

"Information technology . . . ultimately will underlie everything of importance."

Moore's Law is one example of many of this basic exponential nature of the power that is measured in price-performance and bandwidth capacity of information technology. Information technology is not just electronic gadgets, but includes, for example, our understanding of biology and many other facets and ultimately will underlie everything of importance.

You might wonder, how could this be? If a particular project is unpredictable, how can the overall result of this unpredictable chaotic worldwide activity be predictable? We see that in other areas of science. Thermodynamics is a good example. It is impossible to predict the path of a single molecule in a gas, and yet if you take trillions of trillions of particles, all interacting unpredictably and chaotically, the overall properties are very predictable to a very high degree of precision according to the laws of thermodynamics.

The evolution of technology, which is a continuation of the process that gave rise to the technology-creating species, is itself a chaotic activity with a vast number of unpredictable projects, all of which give rise to a predictable outcome. I am going to quickly show some samples of that in order to demonstrate how pervasive this is.

Image 1 shows that the basic paradigm shift rate, the rate at which we introduce new ways of doing things and adopt new technologies, is accelerating.

Image 1: Mass Use of Inventions (click to enlarge in new window)

It took half a century to adopt the telephone, which is the first virtual reality technology that allows me to be with someone else despite being hundreds of miles apart. That never happened before, a century ago. That took half a century to be adopted by a quarter of the U.S. population. More recent technologies – the PC, cell phone, the Web - were measured in a few years time.

These are all logarithmic graphs - meaning as you go up, the graph it represents multiplying generally by a factor of 10. So a straight line on the logarithmic graph is exponential growth. This is better than exponential growth: the Web was adopted in seven years time, according to this. We have had exponential progression in the adoption of new technologies.

In the first few chapters of The Singularity Is Near, I articulate a theory of evolution, starting with biological evolution, leading to technological evolution. Image 2 shows the key events on both biological and technological evolution on this double logarithmic graph.

Image 2: Countdown to Singularity, Logarithmic Plot (click to enlarge in new window)

This shows how long ago in powers of 10 the event took place and how long it took until the next paradigm shift. The first paradigm shift - basically the evolution of biology itself, cells, in particular DNA/RNA, where evolution created a little information processing system, a computer system to keep track of its experiments - took billions of years.

Evolution works through indirection, it creates a capability and then uses that capability to evolve the next stage. That is why the next stage goes more quickly and why the fruits and products of an evolutionary process grow exponentially in power.

So the next stage, the Cambrian explosion, where all the animal body plans evolved, took only ten million years and was one hundred times faster. The biological environment kept accelerating. Homo sapiens evolved in only a few hundred thousand years. Then again, working through indirection, evolution used one of its products: a species that combined a higher cognitive function and an opposable appendage to bring in the next stage, which is technology, which went a little bit faster.

It took only tens of thousands of years for the first stage - fire, wheel, stone tools, and so on. There is only a very small genetic change between us and our still unidentified primate ancestor. A few gene changes allowed a larger cerebral cortex to give us more analytical skills. A very small genetic change moved the pivot point of the thumb up about one inch. Although a chimpanzee's hand looks similar to ours, chimps do not have a power grip or fine motor coordination. This enabled us to manipulate the environment to reflect our mental models. We always use the latest technology to bring the next stage. A half millennium ago, the printing press took a century to be adopted.

Interestingly, this makes a straight line, with technological evolution continuing this evolutionary process. If we look on a linear graph (Image 3), it looks like everything just happened.

Image 3: Countdown to Singularity, Linear Plot (click to enlarge in new window)

Some people say that I only put points on this graph that fit on the straight line. So I took fourteen different lists; these were not thinkers who were trying to prove or disprove my point. Most of them did not even talk about acceleration. They include Carl Sagan's Cosmic Calendar, the American Museum of Natural History, the Encyclopedia Britannica, different reference works as to what were the key events in biological evolution and technological evolution. We do see some spreading of the points in the results. There is disagreement about a few things, such as when did agriculture start, when did language start, how long did the Cambrian explosion take? Nonetheless, the result is a very pervasive straight line. When economist Ted Modis grouped these into what he calls canonical milestones, he comes up with a very similar graph.

Image 4: Paradigm Shifts (click to enlarge in new window)

Image 5 shows that there is a key difference when you look over a significant period of time between the linear and exponential view.

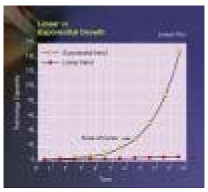

Image 5: Linear vs. Exponential Growth (click to enlarge in new window)

Most public policy in government is based on linear models, which workvery well for short periods of time. You can see that the linear and exponential views are very similar. You can take an exponential progress at any stage and if you take a very small piece of it, it looks like a straight line, more or less. Yet if you go over a significant period of time, there is a great divergence.

The social security debate is unusual in that they are actually talking about 2042. That is the date I have for the Singularity. They are saying that there might be a three or four year increase in longevity and 1.7% increase per year in economic growth and so on. But this linear view will not match reality.

Moore's Law (Image 6) says that information technologies double their power every year.

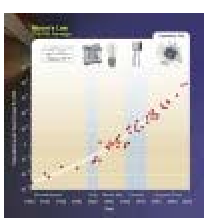

Image 6: Moore's Law (click to enlarge in new window)

It is actually double exponential growth, but right now it iss doubling price performance capacity every year and it is very pervasive. Consider my own personal experience. When I came to MIT, a computer took up space larger than this room, yet was less powerful than your cell phone today. There have been 24 doublings of price performance just in terms of MIPS. This does not even take into consideration all the ways in which computers today are more powerful.

This is just one example of many. Moore's Law is just the vertical stripe on the right in Image 6, shrinking the size of transistors on an integrated circuit, but there has been exponential growth for a hundred years. These are the 49 "famous" computers, going back to the first data processing equipment, used in the 1890 American census, which was old punch card machines. Around 1940, we see the relay-based computer that broke the German enigma code, and then vacuum tube based computers predicted the election of Eisenhower in 1952. They were shrinking vacuum tubes, making them smaller and smaller to keep the exponential growth going. That hit a wall, but it did not stop the overall progression. When one paradigm comes to an end, it actually creates research pressure to create the next paradigm. Thus, transistors, which had a niche application in radio, were brought over to computers.

So Moore's Laws is not the first, but the fifth paradigm to provide this exponential growth. We have been talking for some time now that that will come to an end. The first predictions were 2002. Intel now predicts that by 2022, the key features of transistors will be a few atoms in width and we will not be able to shrink them further. Will that be the end of Moore's Law? Yes. Will that will not be the end of the exponential growth of computing?

We will then go to the next paradigm. Because we have been talking about the end of the Moore’s law paradigm for some time now, there has been increasing research on the sixth paradigm, which is three-dimensional molecular computing. We live in a three-dimensional world. Our brain, although it uses a very slow chemical switching, is organized in three dimensions. We might as well use the third dimension. When I talked in *The Age of Spiritual Machines* in 1999 that the next wave would be three-dimensional molecular computing, nanotubes were a very powerful, likely candidate. I pointed out that DNA computing also would be interesting. That was very controversial then, but it is now a mainstream view because there has been so much progress in three-dimensional molecular computing.

Image 7 shows Hans Moravec's chart. Supercomputers are marching right along now up to 10 to the 14th.

Image 7: Evolution of Computer Power/Cost (click to enlarge in new window)

Moravec's estimate of the computation capacity needed to functionally emulate the human brain is 10 to the 14th; mine is 10 to the 16th, which is a little more conservative. Many different ways of looking at this process of performance, such as dynamic RAM, are also going through different paradigms.

Image 8 shows the smooth curve of average transistor price.

Image 8: Average Transister Price (click to enlarge in new window)

You could buy one transistor for a dollar in 1968. When I was a high school student in New York, I would hang out at the surplus electronics shops on Canal Street, which are still there, and buy a telephone relay with support circuitry for $40. This was a million times slower than a transistor, but otherwise equivalent to one transistor. You could buy 10 million transistors in 2002; it is now about 50 million transistors for a dollar. It is a very smooth curve.

This is not the output of some table top experiment. This is the result of measuring worldwide activity that is very chaotic. There have been bankruptcies, IPO's, accusations of one country dumping products in another, and wars and we nonetheless have this very smooth progression. Unlike Gertrude Stein's roses, it is not the case that a transistor is a transistor. As we have made them cheaper, they have actually gotten better because they are smaller. The electrons have less distance to travel, so we have exponential growth in the speed. The cost of a transistor cycle has been coming down by half every 1.1 year. If you add in other levels of innovation, it is about one year now to double the price-performance of electronics. That is 50% deflation.

This actually affects every aspect of information technology, including biology. It took us 15 years to sequence HIV; we sequenced SARS in 31 days. The economists say that that is a danger. They have been worrying as much about deflation recently as inflation. We had deflation during the Depression - a completely different phenomenon, that was the collapse of confidence, collapse of the money supply. This deflation is due to improved productivity, but the economists say that's all very good, but if you can get the same capability for half the money a year later, you may increase your purchasing somewhat, but you are not going to keep up with doubling consumption every year. You are not going to buy twice as much every year. You will therefore have a contraction of the economy at least as measured in dollars, which would be a bad thing.

Actually, what we find is that we more than keep up with it. There has been 18% growth for the last 50 years, 18% per year, compounded, in electronics and information technologies in general in dollars, despite the fact that you can get twice as much per dollar each year. The reason is as new capabilities come to be cost-effective, it opens up whole new applications. People did not buy iPods for $10,000 five years ago. Again, we have very smooth exponential growth; information technology is already 8% of the economy. It will be the majority of the economy by the 2020s.

It is very pervasive. Image 9 shows Magnetic Data Storage.

Image 9: Magnetic Data Storage (click to enlarge in new window)

This is not Moore's Law; this is not shrinking transistors on an integrated circuit. It i's shrinking magnetic spots on a substrate: a different technological problem, with different engineers, different companies, different countries, yet the same exponential progression.

The biotechnology revolution, now underway, is very profound. The old paradigm was not to understand how biology worked, it was called drug discovery, which meant we would just find something that happens to work. For example, we discover something that lowers blood pressure, yet we have no idea why it works or how it works. Invariably, it would work to some extent, but have lots of side effects taht we did not understand. We did not have effective models of the information processes underlying biology. We still do nt have them in full, but we are making exponential progress in gaining them. We are gaining the means of actually reprogramming our biology. We have these little software programs inside us: 23,000 of them called genes. They were written, that is to say evolved, tens of thousands of years ago.

How much software do you use that you have not updated in the last 20 months let alone 20,000 years? One such "software program" called the fat insulin receptor gene says hold onto every calorie, because the next hunting season may not work out so well. That was a great program 20,000 years ago. Calories were few and far between. It was very good to hold onto them. We would like to reprogram that now. We have a new technology called RNA interference, which can turn genes off. Little RNA fragments in a RNAi medication go into the cell, latch onto the messenger RNA expressing a gene, and deactivate it; thereby turning off that gene, and this works very well.

What would happen if you turned off the fat insulin receptor gene in the fat cells? This was done at the Joslin Diabetes Center in mice. These mice ate ravenously and remained slim, and got the health benefits of being slim. They did not get diabetes or heart disease. They lived 20% longer. They got the health benefits of caloric restriction without the restriction. There are five pharmaceutical companies rushing to bring that to the human market.

At the recent Time Life Conference on the Future of Life from the 50th Anniversary of the Discovery of DNA, all of us speakers were asked, what would the next 50 years bring. Every speaker except Bill Joy and myself used the last 50 years as a model for the next 50 years, which is not valid. Even Watson himself said that in 50 years, we will have drugs that allow you to eat as much as you want and remain slim. I said Jim, "We have already done that in animals. Of course the FDA will slow it down, but it's not going to be 50 years; it will be in the next 5 to 10 years."

Image 10 shows that we went from $10 to 2 cents in the cost of sequencing a base pair of DNA from 1990 to 2004. This graph represents a doubling every year of the amount of genomic data we've been collecting.

Image 10: DNA Sequencing Cost (click to enlarge in new window)

The Genome Project was not a mainstream project when it was announced in 1990.

Mainstream scientists said we just had our best PhD students, the most advanced equipment, and around the world, we have collected one ten-thousandth of the human genome. There is no way that we are going to do this in 15 years; it will be a hundred years at least. Yes, they'll speed it up. Ten years later, the skeptics were still going strong saying "I told you this wasn't going to work." I mean, here we are 10 years in a 15 year project and you've finished 2% of the project. Yet it is the last several doublings that go from 1% to 100% and the project was indeed done on time. This has continued now through the proteome. We have not yet reversed engineered biology, but we are making progress at an exponential rate.

We also see exponential growth in communications technology. There are many ways to measure communications technology, such as Internet data, traffic, the Internet backbone, the size of the Internet, and so on.

Image 11: Internet Hosts, Logarithmic Plot (click to enlarge in new window)

I had a little piece of this curve on the number of nodes on the Internet when I wrote my first book, The Age of Intelligent Machines, in 1985. We went from 10 thousand nodes serving 2,000 scientists, to 20,000, and then to 40,000 at yearly intervals. Nobody had heard of it; it was ARPANET. It was clear to me that this doubling trend was going to continue. Ten years later, it would be 10 million going to 20 million to 40 million and then it would be on everybody's radar screen. I put a prediction in about that and that is what happened.

Image 12: Internet Hosts, Linear Plot (click to enlarge in new window)

Image 12 is a linear graph of the same data. From the linear trend, it looked like the Internet came out of nowhere in the mid-1999's. But you could see these trends emerging if you look at the province in which they really reside, which is exponential progression.

Image 13: Decrease in Size of Mechanical Devices (click to enlarge in new window)

Another exponential term is miniaturization, with technology shrinking in size, not just the electronics, but mechanical at a rate of about five per linear dimension per decade, which is over 100 in 3D volume.

We actually do now little machines for the first time that can do very complex tasks at the molecular level, at least in experiments. There is a little robot that walks with a convincing humanlike gate, built at the molecular level. These are experiments, but they do show the ability to manipulate matter and create machines that are reliable at that level.

We have reverse-engineered red blood cells. They are fairly simple devices. This brings up an issue regarding biology. Although biology is quite remarkable and intricate, it is actually very suboptimal, because biological evolution made certain assumptions, like building everything out of proteins, which is a very limited class of materials that you can roll up from a linear sequences of amino acids or doing signaling in interneuronal connections at a few hundred feet per second versus electronic, which is a million times faster.

Conservative analyses of Freitas' design indicate that if you replace 10% of your red blood cells with these respirocytes, you could do an

"For the first time we can actually see not only our brain creating our thoughts but our thoughts creating our brain."

Olympic sprint for 15 minutes without taking a breath or sit at the bottom of your pool for four hours. He also has a white blood cell design. These are circa 2020 scenarios. I actually watched one of my red blood cells on a microscope. It took an hour and a half to destroy a bacteria; this robotic system would just do something similar in a few seconds.

If this sounds futuristic, I would point out a couple of things: there are already four major conferences on BioMEMS, biological microelectronic mechanical systems. I would not call them nanobots quite yet, but they are capsules, devices that are blood-cell size, several microns in size, with nanoengineered features, that are performing therapeutic functions in animals.

One scientist actually has a very sophisticated device, with nanoengineered seven-nanometer pores, that cures Type 1 Diabetes in rats. It lets out insulin in a controlled fashion and blocks the antibodies, because Type 1 is an autoimmune disease. We can see that the idea of blood cell-sized devices with nanoengineered features in the bloodstream performing therapeutic health functions is not quite as futuristic as it may sound.

We already have devices in the human body that can download software from outside the body. Consider the Parkinson's implant, the latest generation of which allows you to download new software to your neural implant from outside the patient. This is not a nanobot, but if you apply these exponential trends of doubling price-performance of computation and communication, the shrinking technology at an exponential rate, it is conservative to expect that these devices, which are all ready working at some level in animal experiments, will be quite sophisticated by the 2020's.

If we continue that exponential trend of computation through this century, $1,000 of computation will equal even my more conservative estimate: ten to the sixteenth calculations per second for functional emulation of the human brain by 2020. That was a controversial notion in 1999, but it is pretty much a mainstream view today that we will have plenty of hardware computation to emulate human thinking by around 2020.

Image 14: Exponential Growth of Computing (click to enlarge in new window)

Now the controversy is, will we have the software? The ultimate source of the software of human-level intelligence is really to understand the best example we have of a complex, intelligent system, which is human intelligence. We are making much more rapid progress in doing this than people realize. Chapter 4 in The Singularity is Near is about this issue; why we can be confident in reverse-engineering the human brain. We have made more progress than people realize. Brain scanning is growing exponentially in spatial and temporal resolution. The latest generation can see non-invasively individual interneuronal connections. For example, there is an exciting new technology from the University of Pennsylvania that can see individual interneuronal connections signaling in real time. For the first time, we can actually see not only our brain creating our thoughts but our thoughts creating our brain. As we think about a subject, we are creating new spines and new synapses; we can actually watch that now. We are getting experientially more data about the brain.

Image 15: Noninvasive Brain Scanning (click to enlarge in new window)

The next question is, can we understand this information? Maybe it is just inherently too complex for our brains to understand. Doug Hofstadter muses that maybe our brains are just below that threshold

> *"Although the brain is not simple, the apparent complexity is much greater than the actual complexity."*

necessary to understand our own intelligence. Maybe any system is inherently below the threshold needed to understand itself. If we were more intelligent and able to understand it, then we would necessarily be that much more complex and so would never catch up with it. We are finding that we are able to accurately model in mathematical terms specific regions of the brain as we get the data. Although the brain is not simple, the apparent complexity is much greater than the actual complexity. Consider that the design of the brain is in the genome. You can show the genome has about 30 to 100 million bytes of information in it, compressed, and we are able to understand the methods that it encodes.

There are about two dozen regions in the brain where we have very detailed models and simulations. Image 16 is a block diagram of 15 regions of the auditory cortex, where scientists on the West Coast have created detailed models and computer simulations of those regions: how these regions code auditory information and transform it.

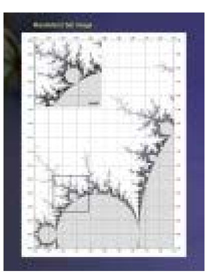

Image 16: Mandelbrot Set Image (click to enlarge in new window)

Applying psychoacoustics tests to this simulation gets very similar results to applying these same tests to human auditory perception. There is a similar system for the cerebellum, which comprises more than half the neurons in the brain. Again applying skill formation tasks, which is what the cerebellum does, to the simulation gets very similar results to experiments on human skill formation. It does not prove that these models are perfect, but it does show that as we are getting the data, we are actually able fairly rapidly to express them in the language of mathematics. If we can do that, we can simulate them.

This Mandelbrot set appears to be a very complex-looking formula. As we look deeper into the image, we see complexity within complexity. Yet the design – the formula – for this image is only six letters long.

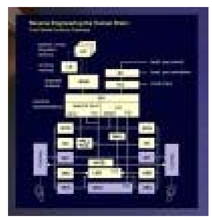

Image 17: Reverse Engineering the Human Brain (click to enlarge in new window)

Similarly, the way the genome actually creates the brain is that there is a lot of stochastic randomness within constraints. For example, there are only a few genes, a few thousand bytes of information that describe how the cerebellum is wired. It says the following, that there are four neuron types that are organized like this. You wire them in this fashion, and now repeat 10 billion times and add a little bit of randomness within the following constraints each time. Then you have this essentially randomly wired cerebellum that over time interacts with a complex environment and the child gathers skills, learns to walk, and talk and catch a fly ball. The child's cerebellum gets filled up with a lot of complex information, but there is actually very little information in the genome that describes the design of this system. Models often get simpler, not more complex, as we go up to a higher level.

All of this is driving economic growth, even on a per capita bases. Underlying this is exponential growth in the value of a human hour of labor, which went from $30 to $130 in 45 years. And the adoption of these technologies is exponential. Here is the adoption of e-commerce; it's now a trillion dollars, which is already meaningful on the world stage. You might say, wasn't there a boom and a bust in e-commerce? There's a similar graph for telecommunications; there was a boom and a bust there also.

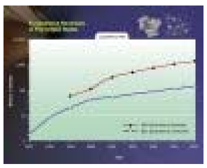

Image 18: E-commerce Revenues in the United States (click to enlarge in new window)

But the boom and bust was strictly a capital market phenomenon, a Wall Street phenomenon. Wall Street looked at the Internet and said, "Wow, this is going to transform every business model." Thus all the values went off the charts and then a year later, when every business model had not been turned on its head, they said, "I guess that was wrong," and everything went the other way. The actual adoption is exponential, but that does not mean instantaneous. It is now getting some real traction, with a trillion dollars of e-commerce revenue. We do have companies that are basically e-commerce companies, like Google for example, with close to a hundred billion dollar market cap, and Ebay, which actually harnesses the value of the net. Yet the adoption is very smooth exponential growth. In fact, when you see this boom and bust phenomenon, it is generally a harbinger of a real revolution. There was a little one for AI in the 1980's. In the 19th century, there was a boom and bust for the railroads, and so on. Information technology, narrowly defined, will be a majority of the economy by 2020. It is already deeply influential in every other aspect of the economy.

One of my companies does speech synthesis, which we also developed in another one of my companies, and commercial language text translation. This speech to speech language translation system, which is basically a translating telephone, will be a routine feature of your cell phones early in the next decade.

The speech sounds natural but this is the latest generation of speech synthesis, which is synthesized with concatenated diphones, and has pretty good inflection. We introduced a product recently, which is a pocket-size reading machine for the blind (http://www.knfbreader.com/).

Image 19: Reader for the Blind

It uses this kind of speech synthesis and optical character recognition and some intelligent image processing. We introduced it at the National Federation of the Blind convention. It fits in your shirt pocket, so a blind person if they are at this meeting they can take the handouts and read them. Just by snapping the picture and it reads it out loud. It does the OCR and the image cleanup.

Let me quickly mention some scenarios, but then get to the issue of promise and peril. Computers are getting smaller. They are already under our arms, in our pockets, and they will soon be in our clothing. We are developing some intelligent clothing in United Therapeutics in a joint venture with Kurzweil Technologies. It will be like an undershirt that will actually monitor your health. If you collapse on the golf course, it will call 911 and direct the ambulance by GPS.

Early in the next decade, images will be written directly to your retina from your eyeglasses, so you won't have to carry around displays. It will create high-resolution, full-immersion, virtual reality environments. I actually have an early version of a virtual-reality three-dimensional technology: Teleportec. I give about a third of my presentations using that. It looks as though I am there in three dimensions; you can see the local background behind me as I move around. It is fairly elaborate, but early in the next decade, this will be fairly routine types of technology.

> *"[Humans] are getting closer to technology and ultimately, we'll merge with it."*

It is really when we go to the end of the third decade of this century that we will have enough turns of the technology doubling screw that these technologies will be very profound because of the double-exponential growth of computation, communication, and our understanding of biology. As powerful as they are today, these technologies will be a billion times more capable by the end of 2020s. We will have completed the reverse-engineering of the human brain. We will have models and simulations that express the power of human intelligence that will add to the AI tool kit.

We already have hundreds of applications of narrow AI, programs that perform functions at human levels that used to require human intelligence. Every time you send an e-mail or place a cell phone call, intelligent algorithms route the information. If you get an electrocardiogram, it comes back with an automated diagnosis. Computers are flying and landing airplanes, guiding intelligent weapons, and are responsible for billions of dollars of daily investments in the stock market. We will gain the knowledge of the full range of human intelligence. The power of human intelligence is reflected in our pattern-recognition capabilities, which is still a unique advantage of human intelligence.

We will combine the powers of human pattern recognition with the natural advantages of machine intelligence, which are speed and repeatability. Machines can remember billions of things; we humans are hard-pressed to remember a handful of phone numbers. Machines can share their knowledge at electronic speeds. Humans are limited to the bandwidth of language, which is a millions times slower.

In my view, this will not be an alien invasion of intelligent machines. We are getting closer to technology and ultimately, we will merge with it. The killer app of nanotechnology is nanobots, which will be in the environment producing inexpensive energy, cleaning up the results of 19th century industrial era environmental degradation, and most importantly, going inside our bodies and brains.

My other current book, Fantastic Voyage: Live Long Enough to Live Forever talks about three bridges to radical life extension. Bridge one is applying today's knowledge aggressively so that us baby boomers can be in good shape when we have the full blossoming of the biotechnology revolution. At that point, we will master the information processes underlying biology. That in turn will be a bridge to the full blossoming of the nanotechnology revolution, where we can send nanobots inside our bodies and brains to keep us healthy, to reverse DNA errors, to remove debris, and to destroy pathogens and cancer cells. The nanobots will also go inside our brains to provide full-immersion virtual reality from within the nervous system and, most importantly, to extend human intelligence.

One of the major objections to radical life extension is the claim that life would be boring if we lived for hundreds of years and that is actually true. If we had radical life extension without radical life expansion, life would get boring. It would be an "endless do loop," to use Vernor Vinge's phrase. By merging intimately with our technology, to which we are already getting closer, we will be able to expand our horizons.

> *"[T]he genetic revolution is with us today and we are gaining the means . . . to reprogram biology."*

That is the promise side of the equation. The peril side is also daunting. I think we can take a measure of comfort from how well we have done with one new form of self replicating pathogen, which is software viruses. When they were first introduced, alarmists said, "This is going to destroy the effectiveness of networks. These first viruses are not very sophisticated, but they are going to get more sophisticated. They will become stealthy, people will place them in various places, and trigger them with different messages." All of that was true; they have become more and more sophisticated, but they have not shut down networks. Nobody has taken the Internet down for even a few seconds. Nobody is saying, "Well let's get rid of the Internet, because the problem with software viruses is so terrible."

They do cause billions of dollars of damage, but the benefit from electronic decentralized communication has been hundreds of times greater than the problems caused by these software pathogens. We have an evolving immune system, a technological system that responds to new threats and that responds very quickly within hours, or days. If there is some very clever new type of software pathogen, there is a response in place very quickly. It does not work perfectly, there is a lot of damage, and it is a very chaotic system. We cannot cross software pathogens off our list of concerns, but this evolving immune system actually has worked very well.

One of the reasons that it has worked well is that this is an area where we have no regulation. There is no certification of practitioners despite the deep influence that software programmers and creators of software and information technology have. There is no certification of products. We put out new aspects of this technological immune system without certification. It is a self-regulating system and the pathogen writers have equal access to the tools of creation, as do the scientists and engineers we rely on to protect us. The system works extremely well.

I will now jump to the biological world, because the concerns come in this order: G,N,R. The revolution that we are in the early stages of now is really the genetic revolution (G). True nanotechnology (N) is not here yet. There are early adoption technology in terms of nanoparticles, but those are not really nanoengineered machines, although there are early experiments with nanoengineered systems. I think nanotechnology is something that we will see in terms of the way with which we intend it in the late teen years and 2020s. But the genetic revolution is with us today and we are gaining the means, as I mentioned, to reprogram biology.

> *"The dangerous technologies are the same ones that are beneficial and you can't relinquish them without a broad totalitarian system."*

The same tools that can empower us to overcome cancer can also empower bioterrorists to create a bio-engineered virus that would combine three bad characteristics: be deadly, spread easily through the air, and be stealthy, that is, have a long incubation period. New viruses come along but they do not happen to be at the worst part of the spectrum on all of those characteristics. For example, SARS spreads pretty easily, but not through the air. It was pretty deadly, with about a 30% death rate. It is not very stealthy, because it has a fairly short incubation period. With software viruses, we do have regulation, which slows down the responsible practitioners. A bioterrorist does not have to put his or her invention through the FDA. This is an issue about which I have given testimony to Congress.

We are close at hand to some effective new technologies that have a broad-spectrum effect against biological viruses and we need to accelerate those. We need to consciously, as a society, put more stones on the defensive side of the equation. I have advocated a Manhattan style project to develop defensive biological virus technologies. A good example is RNA interference. I have described a rapid response system based on RNAi. We obtain a new virus; we sequence it in a few days, which we can do now. We develop an RNA interference medication which works against viruses, because viruses are genetic information. We have shown that we can stop viral diseases with

RNA interference. We would then rapidly gear up production of this medication. We could do that with today's technology, but nobody is organizing that. There are other methods that should also be pursued. For example, we should greatly accelerate vaccine development and develop new means of production that do not rely on eggs.

But the meta lesson here is we really need to address the regulatory issue. I am on the Army Science Advisory Group, which advises the Army on science and technology issues. The Army is responsible for bioterrorism response in this country. They are very concerned about the FDA, because it is not going to be feasible to test responses to bioterrorism agents using the normal regulatory model. We need to put a few stones on the defensive side of the scale by spending money specifically on the defensive side of the equation. As opposed to the calls for relinquishment, which basically will not work. The dangerous technologies are the same ones that are beneficial and you cannot relinquish them without a broad totalitarian system. That was the lesson in the novel Brave New World. It does not work; it just drives the technologies underground, where they continue in a less stable fashion. The responsible scientists then would not have easy access to the knowledge needed to defend civilization.

We do need ethical guidelines. The Asilomar guidelines in the biotechnology field have worked reasonably well. Yet they are not foolproof and obviously, irresponsible practitioners such as terrorists are not going to follow those guidelines. We are tantalizingly close to having broad-spectrum anti-biological virus tools. It is going to be a race. Someone could right now put out an existential-threat pathogen. We want to make sure we have the defenses ready when we need them.

The knowledge and tools to do that are much more widespread than the knowledge and tools to create an atomic bomb. It is not easy to create an atomic bomb, to get the knowledge, let alone the materials. Yet you can go to a routine college laboratory and all the tools and knowledge are there to create a bio-engineered pathogen.

That is the threshold we are on now. The next major challenge will be nanotechnology. There is a lot of effort now to say that nanotechnology manufacturing does not require self-replication. That is basically true, but the lesson from that is not that self-replication is no longer a concern, that grey goo was not a real specter. Grey goo *is* a possibility. In fact, manufacturing that does not use self-replication still does have self-replication hidden within the system. Techniques that use things like the broadcast architecture, where an entity does not have all the codes needed to self-replicate itself are a good idea. That inherently will make nanotechnology safer than biology. In biology, as a cell replicates, it has all of the replication codes within it. It does not have to go to some centralized server. But there are ways of defeating this broadcast architecture if you are determined to defeat them.

> *"Ultimately, when you have achieved a certain level of intelligence in the nanosystem itself, a static non-replicating immune system is not going to work. That is the lesson that biology 'learned.'"*

I do not think the main danger is actually from accidents. I think it is very hard to just accidentally put together something that is deadly at a massive scale. Certainly, accidents can happen on a small scale. For example, it is not easy to just sort of accidentally put together an atomic bomb. Any of these real dangers require very exquisite engineering directed at a destructive goal. But we cannot assume that people will not in fact do that. Recent history shows that people will do that for whatever reason. We ultimately will need a nanotechnology immune system.

Does a nanotechnology immune system need to have self-replication for the blue goo, that is to say the good nanobots to keep up the with grey goo or the bad nanobots? I had an interesting debate in writing this book with Rob Freitas about this very issue, who has a proposal for a nanotechnology immune system that does not require self replication, where you detect the nanobots, you quickly decide on a design of an antidote, some nano system that would destroy it, kind of like RNA interference in the biological world, and then use nanomanufacturing to put out large numbers of them.

I think, without going through the details, a scenario like that will work in the early stages. Yet any particular system that you put in place is not going to last indefinitely. The technology is going to get more sophisticated. The nanobots themselves will get more intelligent. Ultimately, when you have achieved a certain level of intelligence in the nanosystem itself, a static non-replicating immune system is not going to work. That is the lesson that biology "learned." It evolved an immune system that does have self-replication in it. This immune system itself can represent a risk. Bill Joy points out that the immune system itself could turn on us.

And that, of course, is true in the biological world as well. We have auto-immune diseases, but that is not a reason to not have an immune system. We would not last very long without one. We will need an immune system. Ultimately, the early ones may not need self-replication or the self-replication can be hidden or it can use the broadcast architecture. Yet ultimately, we will need an immune system that does have self-replication.

Then one could say, wait a second, that sounds just like the first stage of a stealth destroyer, because the ultimate nightmare scenario of nanotechnology is not just grey goo – a nanobot that self-replicates, because the front of destruction would move very slowly. Nanobots cannot move very quickly. Therefore, you would see it happening and you could deal with it. The real nightmare scenario is: The nanobots self-replicate, using up one in thousand trillion of the carbon atoms. It is very stealthy; nobody notices it. Then they plant themselves throughout the biosphere. At some trigger, they start self-replicating in place. Then the front of destruction does not have to move. It has already seeded the entire biosphere and it would take about 90 minutes to self-replicate and destroy the biosphere.

"Intelligence is the most powerful force in the universe."

Here, you cannot at that time start trying to design a self-replicating immune system. One could point out that putting an immune system in place that has self-replicating capabilities looks just like the first stage of the stealth destroyer scenario. Of course, the immune system is there to protect us, not to destroy us. But how can you tell? These are all interesting questions and biology has dealt with them in a chaotic manner. Ultimately, we will need an immune system. Yet we must take steps to make sure it is friendly.

Finally, we come to the third challenge, R, which stands for robotics, which really means strong AI, AI at a human level, AI at a level that has a copy of its own design, and is able to actually go and improve its own design in a closed-loop self improvement cycle that could ultimately be very rapid. This is the sort of runaway concept of how once you get strong AI, it can rapidly improve its own capabilities.

What if you have unfriendly AI? This would be AI that does not have our values, such as the value of maintaining human civilization. That is the most daunting challenge. If we come back to the nanotechnology challenge, you can see that if you describe any level of nanotechnology, we can design a system, whether it is the broadcast architecture or some more sophisticated version of it or some sophisticated technological immune system that deals with it. Basically the solution has to be more intelligent than the challenge. If you have some level of potentially destructive nanotechnology, you need a response that is more intelligent.

The problem with artificial intelligence is if a system is more intelligent than you and if it is bent on your destruction, there is really not much you can do about it. Intelligence is the most powerful force in the universe. History has shown that the civilization with the more sophisticated technology that has harnessed its intelligence in its technology to a more effective degree generally has prevailed in history, for good or bad.

That is really the most daunting challenge, because when we have nanotechnology, it can protect us from the dangers of biotechnology. The nanobots can protect us from rogue biological viruses. Intelligent technology can protect us from rogue nanotechnology. Yet what is going to protect us from rogue intelligence that is greater than our own intelligence?

Ultimately, the solution I come up with is that we have to embody values of openness, free exchange of information, a democratic society, and all the values that that we embody in our civilization. Because this is not an alien invasion of intelligent machines coming from over the horizon. It is emerging from within our civilization. It is not going to be distinct from us, even within the room. We are going to have intelligent processes running inside our biological bodies and brains. We will have non-biological systems that are derived from the reverse engineering of biological systems. It is going to be one civilization. That civilization needs to embody the best of human values. That is the best we can do to assure that this future intelligence that we are creating, which will be us, will embody the best of human values. If you say, that does not sound foolproof; it is not, but that is really the best we can do.

For additional information, please visit www.kurzweilai.net and www.singularity.com.

Ray Kurzweil, Ph.D. is a best-selling author, inventor and entrepreneur. He is a recipient of the National Medal of Technology.