



The Law of Accelerating Returns

by Ray Kurzweil

An analysis of the history of technology shows that technological change is exponential, contrary to the common-sense "intuitive linear" view. So we won't experience 100 years of progress in the 21st century -- it will be more like 20,000 years of progress (at today's rate). The "returns," such as chip speed and cost-effectiveness, also increase exponentially. There's even exponential growth in the rate of exponential growth. Within a few decades, machine intelligence will surpass human intelligence, leading to The Singularity -- technological change so rapid and profound it represents a rupture in the fabric of human history. The implications include the merger of biological and nonbiological intelligence, immortal software-based humans, and ultra-high levels of intelligence that expand outward in the universe at the speed of light.

Published on KurzweilAI.net March 7, 2001.

You will get \$40 trillion just by reading this essay and understanding what it says. For complete details, see below. (It's true that authors will do just about anything to keep your attention, but I'm serious about this statement. Until I return to a further explanation, however, do read the first sentence of this paragraph carefully.)

Now back to the future: it's widely misunderstood. Our forebears expected the future to be pretty much like their present, which had been pretty much like their past. Although exponential trends did exist a thousand years ago, they were at that very early stage where an exponential trend is so flat that it looks like no trend at all. So their lack of expectations was largely fulfilled. Today, in accordance with the common wisdom, everyone expects continuous technological progress and the social repercussions that follow. But the future will be far more surprising than most observers realize: few have truly internalized the implications of the fact that the rate of change itself is accelerating.

The Intuitive Linear View versus the Historical Exponential View

Most long range forecasts of technical feasibility in future time periods dramatically underestimate the power of future technology because they are based on what I call the "intuitive linear" view of technological progress rather than the "historical exponential view." To express this another way, it is not the case that we will experience a hundred years of progress in the twenty-first century; rather we will witness on the order of twenty thousand years of progress (at *today's* rate of progress, that is).

This disparity in outlook comes up frequently in a variety of contexts, for example, the discussion of the ethical issues that Bill Joy raised in his controversial WIRED cover story, *Why The Future Doesn't Need Us*. Bill and I have been frequently paired in a variety of venues as pessimist and optimist respectively. Although I'm expected to criticize Bill's position, and indeed I do take issue with his prescription of relinquishment, I nonetheless usually end up defending Joy on the key issue of feasibility. Recently a Noble Prize winning panelist dismissed Bill's concerns, exclaiming that, "we're not going to see self-replicating nanoengineered entities for a hundred years." I pointed out that 100 years was indeed a reasonable estimate of the amount of technical progress required to achieve this particular milestone *at today's rate of progress*. But because we're doubling the rate of progress every decade, we'll see a century of progress--*at today's rate*--in only 25 calendar years.

When people think of a future period, they intuitively assume that the current rate of progress will continue for future periods. However, careful consideration of the pace of technology shows that the rate of progress is not constant, but it is human nature to adapt to the changing pace, so the intuitive view is that the pace will continue at the current rate. Even for those of us who have been around long enough to experience how the pace increases over time, our unexamined intuition nonetheless provides the impression that progress changes at the rate that we have experienced recently. From the mathematician's perspective, a primary reason for this is that an exponential curve approximates a straight line when viewed for a brief duration. So even though the rate of progress in the very recent past (e.g., this past year) is far greater than it was ten years ago (let alone a hundred or a thousand years ago), our memories are nonetheless dominated by our very recent experience. It is typical, therefore, that even sophisticated commentators, when considering the future, extrapolate the current pace of change over the next 10 years or 100 years to determine their expectations. This is why I call this way of looking at the future the "intuitive linear" view.

But a serious assessment of the history of technology shows that technological change is exponential. In exponential growth, we find that a key measurement such as computational power is multiplied by a constant factor for each unit of time (e.g., doubling every year) rather than just being added to incrementally. Exponential growth is a feature of any evolutionary process, of which technology is a primary example. One can examine the data

in different ways, on different time scales, and for a wide variety of technologies ranging from electronic to biological, and the acceleration of progress and growth applies. Indeed, we find not just simple exponential growth, but "double" exponential growth, meaning that the rate of

exponential growth is itself growing exponentially. These observations do not rely merely on an assumption of the continuation of Moore's law (i.e., the exponential shrinking of transistor sizes on an integrated circuit), but is based on a rich model of diverse technological processes. What it clearly shows is that technology, particularly the pace of technological change, advances (at least) exponentially, not linearly, and has been doing so since the advent of technology, indeed since the advent of evolution on Earth.

I emphasize this point because it is the most important failure that would-be prognosticators make in considering future trends. Most technology forecasts ignore altogether this "historical exponential view" of technological progress. That is why people tend to overestimate what can be achieved in the short term (because we tend to leave out necessary details), but underestimate what can be achieved in the long term (because the exponential growth is ignored).

The Law of Accelerating Returns

We can organize these observations into what I call the law of accelerating returns as follows:

- Evolution applies positive feedback in that the more capable methods resulting from one stage of evolutionary progress are used to create the next stage. As a result, the
- rate of progress of an evolutionary process increases exponentially over time. Over time, the "order" of the information embedded in the evolutionary process (i.e., the measure of how well the information fits a purpose, which in evolution is survival) increases.
- A correlate of the above observation is that the "returns" of an evolutionary process (e.g., the speed, cost-effectiveness, or overall "power" of a process) increase exponentially over time.
- In another positive feedback loop, as a particular evolutionary process (e.g., computation) becomes more effective (e.g., cost effective), greater resources are deployed toward the further progress of that process. This results in a second level of exponential growth (i.e., the rate of exponential growth itself grows exponentially).
- Biological evolution is one such evolutionary process.
- Technological evolution is another such evolutionary process. Indeed, the emergence of the first technology creating species resulted in the new evolutionary process of technology. Therefore, technological evolution is an outgrowth of--and a continuation of--biological evolution.
- A specific paradigm (a method or approach to solving a problem, e.g., shrinking transistors on an integrated circuit as an approach to making more powerful computers) provides exponential growth until the method exhausts its potential. When this happens, a paradigm shift (i.e., a fundamental change in the approach) occurs, which enables exponential growth to continue.


If we apply these principles at the highest level of evolution on Earth, the first step, the creation of cells, introduced the paradigm of biology. The subsequent emergence of DNA provided a digital method to record the

results of evolutionary experiments. Then, the evolution of a species who combined rational thought with an opposable appendage (i.e., the thumb) caused a fundamental paradigm shift from biology to technology. The upcoming primary paradigm shift will be from biological thinking to a hybrid combining biological and nonbiological thinking. This hybrid will include "biologically inspired" processes resulting from the reverse engineering of biological brains.


If we examine the timing of these steps, we see that the process has continuously accelerated. The evolution of life forms required billions of years for the first steps (e.g., primitive cells); later on progress accelerated. During the Cambrian explosion, major paradigm shifts took only tens of millions of years. Later on, Humanoids developed over a period of millions of years, and Homo sapiens over a period of only hundreds of thousands of years.

With the advent of a technology-creating species, the exponential pace became too fast for evolution through DNA-guided protein synthesis and moved on to human-created technology. Technology goes beyond mere tool making; it is a process of creating ever more powerful technology using the tools from the previous round of innovation. In this way, human technology is distinguished from the tool making of other species. There is a record of each stage of technology, and each new stage of technology builds on the order of the previous stage.

The first technological steps—sharp edges, fire, the wheel—took tens of thousands of years. For people living in this era, there was little noticeable technological change in even a thousand years. By 1000 A.D., progress was much faster and a paradigm shift required only a century or two. In the nineteenth century, we saw more technological change than in the nine centuries preceding it. Then in the first twenty years of the twentieth century, we saw more advancement than in all of the nineteenth century. Now, paradigm shifts occur in only a few years time. The World Wide Web did not exist in anything like its present form just a few years ago; it didn't exist at all a decade ago.



The paradigm shift rate (i.e., the overall rate of technical progress) is currently doubling (approximately) every decade; that is, paradigm shift times are halving every decade (and the rate of acceleration is itself growing exponentially). So, the technological progress in the twenty-first century will be equivalent to what would require (in the linear view) on the order of 200 centuries. In contrast, the twentieth century saw only about 25 years of progress (again at today's rate of progress) since we have been speeding up to current rates. So the twenty-first century will see almost a thousand times greater technological change than its predecessor.



The Singularity Is Near

To appreciate the nature and significance of the coming "singularity," it is important to ponder the nature of exponential growth. Toward this end, I am fond of telling the tale of the inventor of chess and his patron, the emperor of China. In response to the emperor's offer of a reward for his new beloved game, the inventor asked for a single grain of rice on the first square, two on the second square, four on the third, and so on. The Emperor quickly granted this seemingly benign and humble request. One version of the story has the emperor going bankrupt as the 63 doublings ultimately totaled 18 million trillion grains of rice. At ten grains of rice per square inch, this requires rice fields covering twice the surface area of the Earth, oceans included. Another version of the story has the inventor losing his head.

It should be pointed out that as the emperor and the inventor went through the first half of the chess board, things were fairly uneventful. The inventor was given spoonfuls of rice, then bowls of rice, then barrels. By the end of the first half of the chess board, the inventor had accumulated one large field's worth (4 billion grains), and the emperor did start to take notice. It was as they progressed through the second

half of the chessboard that the situation quickly deteriorated. Incidentally, with regard to the doublings of computation, that's about where we stand now--there have been slightly more than 32 doublings of performance since the first programmable computers were invented during World War II.

This is the nature of exponential growth. Although technology grows in the exponential domain, we humans live in a linear world. So technological trends are not noticed as small levels of technological power are doubled. Then seemingly out of nowhere, a technology explodes into view. For example, when the Internet went from 20,000 to 80,000 nodes over a two year period during the 1980s, this progress remained hidden from the general public. A decade later, when it went from 20 million to 80 million nodes in the same amount of time, the impact was rather conspicuous.

As exponential growth continues to accelerate into the first half of the twenty-first century, it will appear to explode into infinity, at least from the limited and linear perspective of contemporary humans. The progress will ultimately become so fast that it will rupture our ability to follow it. It will literally get out of our control. The illusion that we have our hand "on the plug," will be dispelled.

Can the pace of technological progress continue to speed up indefinitely? Is there not a point where humans are unable to think fast enough to keep up with it? With regard to unenhanced humans, clearly so. But what would a thousand scientists, each a thousand times more intelligent than human scientists today, and each operating a thousand times faster than contemporary humans (because the information processing in their primarily nonbiological brains is faster) accomplish? One year would be like a millennium. What would they come up with?

Well, for one thing, they would come up with technology to become even more intelligent (because their intelligence is no longer of fixed capacity). They would change their own thought processes to think even faster. When the scientists evolve to be a million times more intelligent and operate a million times faster, then an hour would result in a century of progress (in today's terms).

This, then, is the Singularity. The Singularity is technological change so rapid and so profound that it represents a rupture in the fabric of human history. Some would say that we cannot comprehend the Singularity, at least with our current level of understanding, and that it is impossible, therefore, to look past its "event horizon" and make sense of what lies beyond.

My view is that despite our profound limitations of thought, constrained as we are today to a mere hundred trillion interneuronal connections in our biological brains, we nonetheless have sufficient powers of abstraction to make meaningful statements about the nature of life after the Singularity. Most importantly, it is my view that the intelligence that will emerge will continue to represent the human civilization, which is already a human-machine civilization. This will be the next step in evolution, the next high level paradigm shift.

To put the concept of Singularity into perspective, let's explore the history of the word itself. Singularity is a familiar word meaning a unique event with profound implications. In mathematics, the term implies infinity, the explosion of value that occurs when dividing a constant by a number that gets closer and closer to zero. In physics, similarly, a singularity denotes an event or location of infinite power. At the center of a black hole, matter is so dense that its gravity is infinite. As nearby matter and energy are drawn into the black hole, an event horizon separates the region from the rest of the Universe. It constitutes a rupture in the fabric of space and time. The Universe itself is said to have begun with just such a Singularity.

In the 1950s, John Von Neumann was quoted as saying that "the ever accelerating progress of technology...gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue." In the 1960s, I. J. Good wrote of an "intelligence explosion," resulting from intelligent machines designing their next generation without human intervention. In 1986, Vernor Vinge, a mathematician and computer scientist at San Diego State University, wrote about a rapidly approaching technological "singularity" in his science fiction novel, *Marooned in Realtime*. Then in 1993, Vinge presented a paper to a NASA-organized symposium which described the Singularity as an impending event resulting primarily from the advent of "entities with greater than human intelligence," which Vinge saw as the harbinger of a run-away phenomenon.

From my perspective, the Singularity has many faces. It represents the nearly vertical phase of exponential growth where the rate of growth is so extreme that technology appears to be growing at infinite speed. Of course, from a mathematical perspective, there is no discontinuity, no rupture, and the growth rates remain finite, albeit extraordinarily large. But from our *currently* limited perspective, this imminent event appears to be an acute and abrupt break in the continuity of progress. However, I emphasize the word "currently," because one of the salient implications of the Singularity will be a change in the nature of our ability to understand. In other words, we will become vastly smarter as we merge with our technology.

When I wrote my first book, *The Age of Intelligent Machines*, in the 1980s, I ended the book with the specter of the emergence of machine intelligence greater than human intelligence, but found it difficult to look beyond this event horizon. Now having thought about its implications for the past 20 years, I feel that we are indeed capable of understanding the many facets of this threshold, one that will transform all spheres of human life.

Consider a few examples of the implications. The bulk of our experiences will shift from real reality to virtual reality. Most of the intelligence of our civilization will ultimately be nonbiological, which by the end of this century will be trillions of trillions of times more powerful than human intelligence. However, to address often expressed concerns, this does not imply the end of biological intelligence, even if thrown from its perch of evolutionary superiority. Moreover, it is important to note that the nonbiological forms will be derivative of biological design. In other words,

our civilization will remain human, indeed in many ways more exemplary of what we regard as human than it is today, although our understanding of the term will move beyond its strictly biological origins.

Many observers have nonetheless expressed alarm at the emergence of forms of nonbiological intelligence superior to human intelligence. The potential to augment our own intelligence through intimate connection with other thinking mediums does not necessarily alleviate the concern, as some people have expressed the wish to remain "unenhanced" while at the same time keeping their place at the top of the intellectual food chain. My view is that the likely outcome is that on the one hand, from the perspective of biological humanity, these superhuman intelligences will appear to be their transcendent servants, satisfying their needs and desires. On the other hand, fulfilling the wishes of a revered biological legacy will occupy only a trivial portion of the intellectual power that the Singularity will bring.

Needless to say, the Singularity will transform all aspects of our lives, social, sexual, and economic, which I explore herewith.

Wherefrom Moore's Law

Before considering further the implications of the Singularity, let's examine the wide range of technologies that are subject to the law of accelerating returns. The exponential trend that has gained the greatest public recognition has become known as "Moore's Law." Gordon Moore, one of the inventors of integrated circuits, and then Chairman of Intel, noted in the mid 1970s that we could squeeze twice as many transistors on an integrated circuit every 24 months. Given that the electrons have less distance to travel, the circuits also run twice as fast, providing an overall quadrupling of computational power.


After sixty years of devoted service, Moore's Law will die a dignified death no later than the year 2019. By that time, transistor features will be just a few atoms in width, and the strategy of ever finer photolithography will have run its course. So, will that be the end of the exponential growth of computing?

Don't bet on it.

If we plot the speed (in instructions per second) per \$1000 (in constant dollars) of 49 famous calculators and computers spanning the entire twentieth century, we note some interesting observations.

Moore's Law Was Not the First, but the Fifth Paradigm To Provide Exponential Growth of Computing

Each time one paradigm runs out of steam, another picks up the pace



It is important to note that Moore's Law of Integrated Circuits was not the first, but the fifth paradigm to provide accelerating price-performance. Computing devices have been consistently multiplying in power (per unit of time) from the mechanical calculating devices used in the 1890 U.S. Census, to Turing's relay-based "Robinson" machine that cracked the Nazi enigma code, to the CBS vacuum tube computer that predicted the election of Eisenhower, to the transistor-based machines used in the first space launches, to the integrated-circuit-based personal computer which I used to dictate (and automatically transcribe) this essay.

But I noticed something else surprising. When I plotted the 49 machines on an exponential graph (where a straight line means exponential growth), I didn't get a straight line. What I got was another exponential curve. In other words, there's exponential growth in the rate of exponential growth. Computer speed (per unit cost) doubled every three years between 1910 and 1950, doubled every two years between 1950 and 1966, and is now doubling every year.

But where does Moore's Law come from? What is behind this remarkably predictable phenomenon? I have seen relatively little written about the ultimate source of this trend. Is it just "a set of industry expectations and goals," as Randy Isaac, head of basic science at IBM contends? Or is there something more profound going on?

In my view, it is one manifestation (among many) of the exponential growth of the evolutionary process that is technology. The exponential growth of computing is a marvelous quantitative example of the exponentially growing returns from an evolutionary process. We can also express the exponential growth of computing in terms of an accelerating pace: it took ninety years to achieve the first MIPS (million instructions per second) per thousand dollars, now we add one MIPS per thousand dollars every day.

Moore's Law narrowly refers to the number of transistors on an integrated circuit of fixed size, and sometimes has been expressed even more narrowly in terms of transistor feature size. But rather than feature size (which is only one contributing factor), or even number of transistors, I think the most appropriate measure to track is computational speed per unit cost. This takes into account many levels of "cleverness" (i.e., innovation, which is to say, technological evolution). In addition to all of the innovation in integrated circuits, there are multiple layers of innovation in computer design, e.g., pipelining, parallel processing, instruction look-ahead, instruction and memory caching, and many others.

From the above chart, we see that the exponential growth of computing didn't start with integrated circuits (around 1958), or even transistors (around 1947), but goes back to the electromechanical calculators used in the 1890 and 1900 U.S. Census. This chart spans at least five distinct paradigms of computing, of which Moore's Law pertains to only the latest one.

It's obvious what the sixth paradigm will be after Moore's Law runs out of steam during the second decade of this century. Chips today are flat (although it does require up to 20 layers of material to produce one layer of circuitry). Our brain, in contrast, is organized in three dimensions. We live in a three dimensional world, why not use the third dimension? The human brain actually uses a very inefficient electrochemical digital controlled analog computational process. The bulk of the calculations are done in the interneuronal connections at a speed of only about 200 calculations per second (in each connection), which is about ten million times slower than contemporary electronic circuits. But the brain gains its prodigious powers from its extremely parallel organization *in three dimensions*. There are many technologies in the wings that build circuitry in three dimensions. Nanotubes, for example, which are already working in laboratories, build circuits from pentagonal arrays of carbon atoms. One cubic inch of nanotube circuitry would be a million times more powerful than the human brain. There are more than enough new computing technologies now being researched, including three-dimensional silicon chips, optical computing, crystalline computing, DNA computing, and quantum computing, to keep the law of accelerating returns as applied to computation going for a long time.

Thus the (double) exponential growth of computing is broader than Moore's Law, which refers to only one of its paradigms. And this accelerating growth of computing is, in turn, part of the yet broader phenomenon of the accelerating pace of any evolutionary process. Observers are quick to criticize extrapolations of an exponential trend on the basis that the trend is bound to run out of "resources." The classical example is when a species happens upon a new habitat (e.g., rabbits in Australia), the species' numbers will grow exponentially for a time, but then hit a limit when resources such as food and space run out.

But the resources underlying the exponential growth of an evolutionary process are relatively unbounded:

1. (i) The (ever growing) order of the evolutionary process itself. Each stage of evolution provides more powerful tools for the next. In

biological evolution, the advent of DNA allowed more powerful and faster evolutionary "experiments." Later, setting the "designs" of animal body plans during the Cambrian explosion allowed rapid evolutionary development of other body organs such as the brain. Or to take a more recent example, the advent of computer assisted design tools allows rapid development of the next generation of computers.

2. (ii) The "chaos" of the environment in which the evolutionary process takes place and which provides the options for further diversity. In biological evolution, diversity enters the process in the form of mutations and ever changing environmental conditions. In technological evolution, human ingenuity combined with ever changing market conditions keep the process of innovation going.

The maximum potential of matter and energy to contain intelligent processes is a valid issue. But according to my models, we won't approach those limits during this century (but this will become an issue within a couple of centuries).


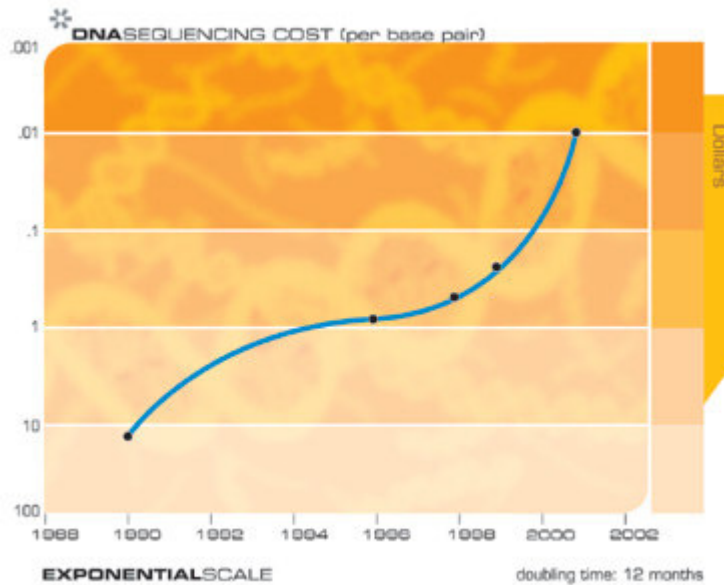
We also need to distinguish between the "S" curve (an "S" stretched to the right, comprising very slow, virtually unnoticeable growth--followed by very rapid growth--followed by a flattening out as the process approaches an asymptote) that is characteristic of any specific technological paradigm and the continuing exponential growth that is characteristic of the ongoing evolutionary process of technology. Specific paradigms, such as Moore's Law, do ultimately reach levels at which exponential growth is no longer feasible. Thus Moore's Law is an S curve. But the growth of computation is an ongoing exponential (at least until we "saturate" the Universe with the intelligence of our human-machine civilization, but that will not be a limit in this coming century). In accordance with the law of accelerating returns, paradigm shift, also called innovation, turns the S curve of any specific paradigm into a continuing exponential. A new paradigm (e.g., three-dimensional circuits) takes over when the old paradigm approaches its natural limit. This has already happened at least four times in the history of computation. This difference also distinguishes the tool making of non-human species, in which the mastery of a tool-making (or using) skill by each animal is characterized by an abruptly ending S shaped learning curve, versus human-created technology, which has followed an exponential pattern of growth and acceleration since its inception.


DNA Sequencing, Memory, Communications, the Internet, and Miniaturization

This "law of accelerating returns" applies to all of technology, indeed to any true evolutionary process, and can be measured with remarkable precision in information based technologies. There are a great many examples of the exponential growth implied by the law of accelerating returns in technologies as varied as DNA sequencing, communication speeds, electronics of all kinds, and even in the rapidly shrinking size of technology. The Singularity results not from the exponential explosion of computation alone, but rather from the interplay and myriad synergies that will result from manifold intertwined technological revolutions. Also, keep in mind that every point on the exponential growth curves


underlying these panoply of technologies (see the graphs below) represents an intense human drama of innovation and competition. It is remarkable therefore that these chaotic processes result in such smooth and predictable exponential trends.

For example, when the human genome scan started fourteen years ago, critics pointed out that given the speed with which the genome could then be scanned, it would take thousands of years to finish the project. Yet the fifteen year project was nonetheless completed slightly ahead of schedule.



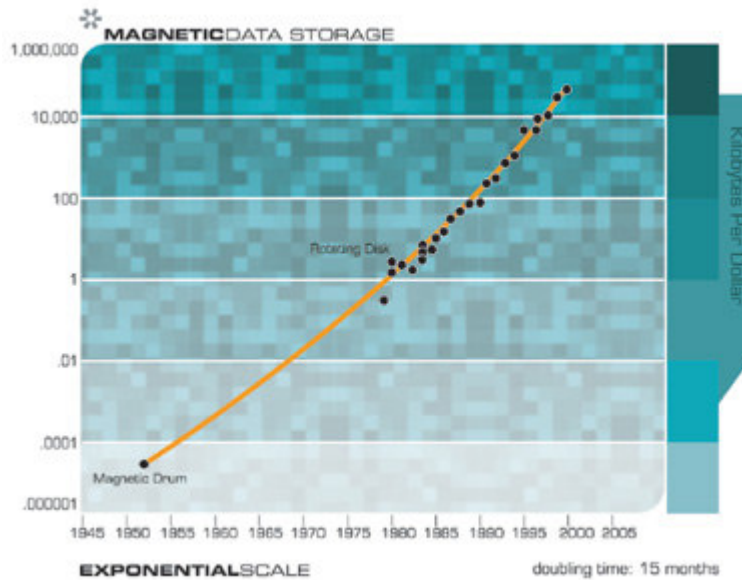


Of course, we expect to see exponential growth in electronic memories such as RAM.




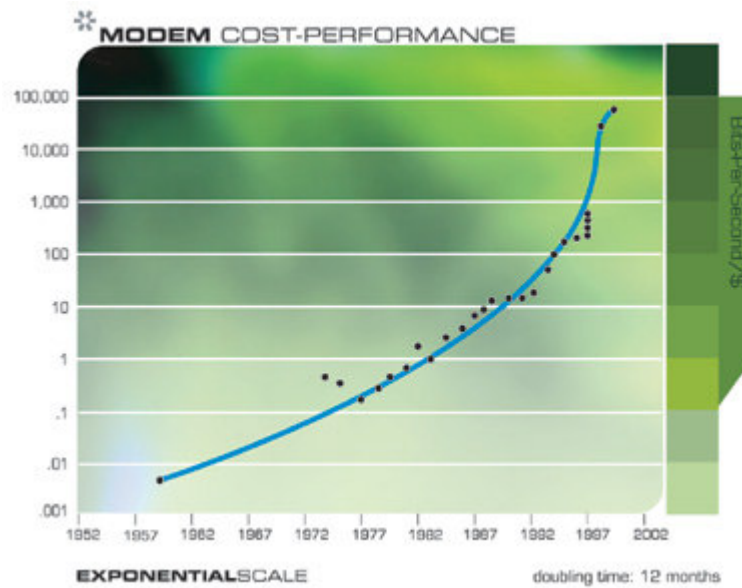
Notice How Exponential Growth Continued through Paradigm Shifts from Vacuum Tubes to Discrete Transistors to Integrated Circuits

However, growth in magnetic memory is not primarily a matter of Moore's law, but includes advances in mechanical and electromagnetic systems.




Exponential growth in communications technology has been even more explosive than in computation and is no less significant in its implications. Again, this progression involves far more than just shrinking transistors on an integrated circuit, but includes accelerating advances in fiber optics, optical switching, electromagnetic technologies, and others.

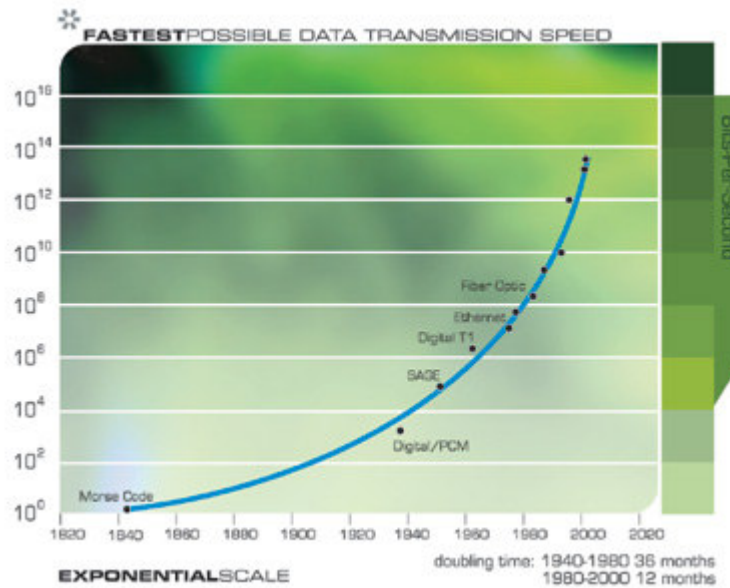





Notice Cascade of smaller "S" Curves


Note that in the above two charts we can actually see the progression of "S" curves: the acceleration fostered by a new paradigm, followed by a leveling off as the paradigm runs out of steam, followed by renewed acceleration through paradigm shift.






The following two charts show the overall growth of the Internet based on the number of hosts. These two charts plot the same data, but one is on an exponential axis and the other is linear. As I pointed out earlier, whereas technology progresses in the exponential domain, we experience it in the linear domain. So from the perspective of most observers, nothing was happening until the mid 1990s when seemingly out of nowhere, the world wide web and email exploded into view. But the emergence of the Internet into a worldwide phenomenon was readily predictable much earlier by examining the exponential trend data.




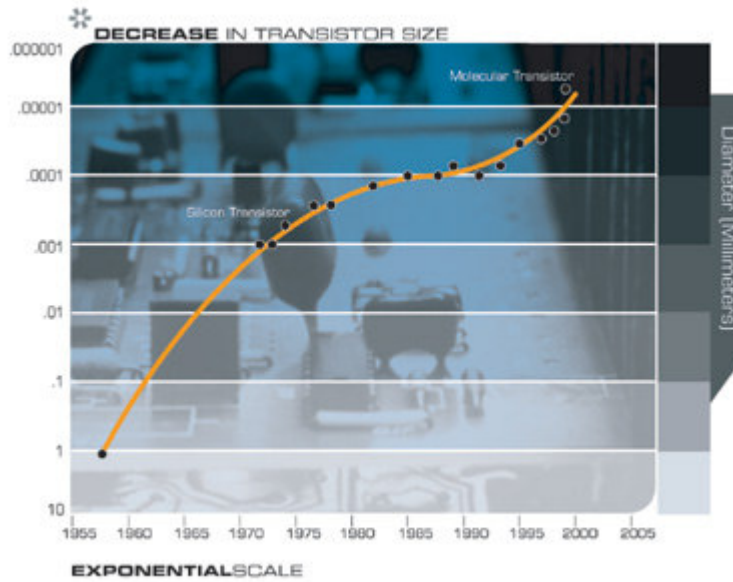



Notice how the explosion of the Internet appears to be a surprise from the Linear Chart, but was perfectly predictable from the Exponential Chart

Ultimately we will get away from the tangle of wires in our cities and in our lives through wireless communication, the power of which is doubling every 10 to 11 months.




Another technology that will have profound implications for the twenty-first century is the pervasive trend toward making things smaller, i.e., miniaturization. The salient implementation sizes of a broad range of technologies, both electronic and mechanical, are shrinking, also at a double exponential rate. At present, we are shrinking technology by a factor of approximately 5.6 per linear dimension per decade.





The Exponential Growth of Computation Revisited

If we view the exponential growth of computation in its proper perspective as one example of the pervasiveness of the exponential growth of information based technology, that is, as one example of many of the law of accelerating returns, then we can confidently predict its continuation.



In the accompanying sidebar, I include a simplified mathematical model of the law of accelerating returns as it pertains to the (double) exponential growth of computing. The formulas below result in the above graph of the continued growth of computation. This graph matches the available data for the twentieth century through all five paradigms and

provides projections for the twenty-first century. Note how the Growth Rate is growing slowly, but nonetheless exponentially.

The Law of Accelerating Returns Applied to the Growth of Computation

The following provides a brief overview of the law of accelerating returns as it applies to the double exponential growth of computation. This model considers the impact of the growing power of the technology to foster its own next generation. For example, with more powerful computers and related technology, we have the tools and the knowledge to design yet more powerful computers, and to do so more quickly.

Note that the data for the year 2000 and beyond assume neural net connection calculations as it is expected that this type of calculation will ultimately dominate, particularly in emulating human brain functions. This type of calculation is less expensive than conventional (e.g., Pentium III / IV) calculations by a factor of at least 100 (particularly if implemented using digital controlled analog electronics, which would correspond well to the brain's digital controlled analog electrochemical processes). A factor of 100 translates into approximately 6 years (today) and less than 6 years later in the twenty-first century.

My estimate of brain capacity is 100 billion neurons times an average 1,000 connections per neuron (with the calculations taking place primarily in the connections) times 200 calculations per second. Although these estimates are conservatively high, one can find higher and lower estimates. However, even much higher (or lower) estimates by orders of magnitude only shift the prediction by a relatively small number of years.

Some prominent dates from this analysis include the following:

- We achieve one Human Brain capability ($2 * 10^{16}$ cps) for \$1,000 around the year 2023.
- We achieve one Human Brain capability ($2 * 10^{16}$ cps) for one cent around the year 2037.
- We achieve one Human Race capability ($2 * 10^{26}$ cps) for \$1,000 around the year 2049.
- We achieve one Human Race capability ($2 * 10^{26}$ cps) for one cent around the year 2059.

The Model considers the following variables:

- V: Velocity (i.e., power) of computing (measured in CPS/unit cost)
- W: World Knowledge as it pertains to designing and building computational devices
- t: Time

The assumptions of the model are:

1. (1) $V = C1 * W$

In other words, computer power is a linear function of the knowledge of

how to build computers. This is actually a conservative assumption. In general, innovations improve V (computer power) by a multiple, not in an additive way. Independent innovations multiply each other's effect. For example, a circuit advance such as CMOS, a more efficient IC wiring methodology, and a processor innovation such as pipelining all increase V by independent multiples.

- $(2) W = C_2 * \text{Integral}(0 \text{ to } t) V$

In other words, W (knowledge) is cumulative, and the instantaneous increment to knowledge is proportional to V .

This gives us:

- $W = C_1 * C_2 * \text{Integral}(0 \text{ to } t) W$
- $W = C_1 * C_2 * C_3^{(C_4 * t)}$
- $V = C_1^2 * C_2 * C_3^{(C_4 * t)}$
- (Note on notation: a^b means a raised to the b power.)

Simplifying the constants, we get:

- $V = C_a * C_b^{(C_c * t)}$

So this is a formula for "accelerating" (i.e., exponentially growing) returns, a "regular Moore's Law."

As I mentioned above, the data shows exponential growth in the rate of exponential growth. (We doubled computer power every three years early in the twentieth century, every two years in the middle of the century, and close to every one year during the 1990s.)

Let's factor in another exponential phenomenon, which is the growing resources for computation. Not only is each (constant cost) device getting more powerful as a function of W , but the resources deployed for computation are also growing exponentially.

We now have:

- N : Expenditures for computation
- $V = C_1 * W$ (as before)
- $N = C_4^{(C_5 * t)}$ (Expenditure for computation is growing at its own exponential rate)
- $W = C_2 * \text{Integral}(0 \text{ to } t) (N * V)$

As before, world knowledge is accumulating, and the instantaneous increment is proportional to the amount of computation, which equals the resources deployed for computation (N) * the power of each (constant cost) device.

This gives us:

- $W = C_1 * C_2 * \text{Integral}(0 \text{ to } t) (C_4^{(C_5 * t)} * W)$
- $W = C_1 * C_2 * (C_3^{(C_6 * t)})^{(C_7 * t)}$

- $V = C1 \wedge 2 * C2 * (C3 \wedge (C6 * t)) \wedge (C7 * t)$

Simplifying the constants, we get:

- $V = Ca * (Cb \wedge (Cc * t)) \wedge (Cd * t)$

This is a double exponential--an exponential curve in which the rate of exponential growth is growing at a different exponential rate.

Now let's consider real-world data. Considering the data for actual calculating devices and computers during the twentieth century:

- CPS/\$1K: Calculations Per Second for \$1,000

Twentieth century computing data matches:

- $CPS/\$1K = 10^{(6.00 * ((20.40/6.00)^{((A13-1900)/100)) - 11.00)}$

We can determine the growth rate over a period of time:

- Growth Rate = $10^{((\text{LOG}(\text{CPS}/\$1K \text{ for Current Year}) - \text{LOG}(\text{CPS}/\$1K \text{ for Previous Year})) / (\text{Current Year} - \text{Previous Year}))}$
- Human Brain = 100 Billion (10^{11}) neurons * 1000 (10^3) Connections/Neuron * 200 ($2 * 10^2$) Calculations Per Second Per Connection = $2 * 10^{16}$ Calculations Per Second
- Human Race = 10 Billion (10^{10}) Human Brains = $2 * 10^{26}$ Calculations Per Second

These formulas produce the graph above.

Already, IBM's "Blue Gene" supercomputer, now being built and scheduled to be completed by 2005, is projected to provide 1 million billion calculations per second (i.e., one billion megaflops). This is already one twentieth of the capacity of the human brain, which I estimate at a conservatively high 20 million billion calculations per second (100 billion neurons times 1,000 connections per neuron times 200 calculations per second per connection). In line with my earlier predictions, supercomputers will achieve one human brain capacity by 2010, and personal computers will do so by around 2020. By 2030, it will take a village of human brains (around a thousand) to match \$1000 of computing. By 2050, \$1000 of computing will equal the processing power of all human brains on Earth. Of course, this only includes those brains still using carbon-based neurons. While human neurons are wondrous creations in a way, we wouldn't (and don't) design computing circuits the same way. Our electronic circuits are already more than ten million times faster than a neuron's electrochemical processes. Most of the complexity of a human neuron is devoted to maintaining its life support functions, not its information processing capabilities. Ultimately, we will need to port our mental processes to a more suitable computational substrate. Then our minds won't have to stay so small, being constrained as they are today to a mere hundred trillion neural connections each operating at a ponderous 200 digitally controlled analog calculations per second.

The Software of Intelligence

So far, I've been talking about the hardware of computing. The software is even more salient. One of the principal assumptions underlying the expectation of the Singularity is the ability of nonbiological mediums to emulate the richness, subtlety, and depth of human thinking. Achieving the computational capacity of the human brain, or even villages and nations of human brains will not automatically produce human levels of capability. By human levels I include all the diverse and subtle ways in which humans are intelligent, including musical and artistic aptitude, creativity, physically moving through the world, and understanding and responding appropriately to emotion. The requisite hardware capacity is a necessary but not sufficient condition. The organization and content of these resources--the software of intelligence--is also critical.

Before addressing this issue, it is important to note that once a computer achieves a human level of intelligence, it will necessarily soar past it. A key advantage of nonbiological intelligence is that machines can easily share their knowledge. If I learn French, or read War and Peace, I can't readily download that learning to you. You have to acquire that scholarship the same painstaking way that I did. My knowledge, embedded in a vast pattern of neurotransmitter concentrations and interneuronal connections, cannot be quickly accessed or transmitted. But we won't leave out quick downloading ports in our nonbiological equivalents of human neuron clusters. When one computer learns a skill or gains an insight, it can immediately share that wisdom with billions of other machines.

As a contemporary example, we spent years teaching one research computer how to recognize continuous human speech. We exposed it to thousands of hours of recorded speech, corrected its errors, and patiently improved its performance. Finally, it became quite adept at recognizing speech (I dictated most of my recent book to it). Now if you want your own personal computer to recognize speech, it doesn't have to go through the same process; you can just download the fully trained patterns in seconds. Ultimately, billions of nonbiological entities can be the master of all human and machine acquired knowledge.

In addition, computers are potentially millions of times faster than human neural circuits. A computer can also remember billions or even trillions of facts perfectly, while we are hard pressed to remember a handful of phone numbers. The combination of human level intelligence in a machine with a computer's inherent superiority in the speed, accuracy, and sharing ability of its memory will be formidable.

There are a number of compelling scenarios to achieve higher levels of intelligence in our computers, and ultimately human levels and beyond. We will be able to evolve and train a system combining massively parallel neural nets with other paradigms to understand language and model knowledge, including the ability to read and model the knowledge contained in written documents. Unlike many contemporary "neural net" machines, which use mathematically simplified models of human neurons, some contemporary neural nets are already using highly detailed models of human neurons, including detailed nonlinear analog activation

functions and other relevant details. Although the ability of today's computers to extract and learn knowledge from natural language documents is limited, their capabilities in this domain are improving rapidly. Computers will be able to read on their own, understanding and modeling what they have read, by the second decade of the twenty-first century. We can then have our computers read all of the world's literature--books, magazines, scientific journals, and other available material. Ultimately, the machines will gather knowledge on their own by venturing out on the web, or even into the physical world, drawing from the full spectrum of media and information services, and sharing knowledge with each other (which machines can do far more easily than their human creators).

Reverse Engineering the Human Brain

The most compelling scenario for mastering the software of intelligence is to tap into the blueprint of the best example we can get our hands on of an intelligent process. There is no reason why we cannot reverse engineer the human brain, and essentially copy its design. Although it took its original designer several billion years to develop, it's readily available to us, and not (yet) copyrighted. Although there's a skull around the brain, it is not hidden from our view.

The most immediately accessible way to accomplish this is through destructive scanning: we take a frozen brain, preferably one frozen just slightly before rather than slightly after it was going to die anyway, and examine one brain layer--one very thin slice--at a time. We can readily see every neuron and every connection and every neurotransmitter concentration represented in each synapse-thin layer.



Human brain scanning has already started. A condemned killer allowed his brain and body to be scanned and you can access all 10 billion bytes of him on the Internet
http://www.nlm.nih.gov/research/visible/visible_human.html.


He has a 25 billion byte female companion on the site as well in case he gets lonely. This scan is not high enough in resolution for our purposes, but then, we probably don't want to base our templates of machine intelligence on the brain of a convicted killer, anyway.

But scanning a frozen brain is feasible today, albeit not yet at a sufficient speed or bandwidth, but again, the law of accelerating returns will provide the requisite speed of scanning, just as it did for the human genome scan. Carnegie Mellon University's Andreas Nowatzyk plans to scan the nervous system of the brain and body of a mouse with a resolution of less than 200 nanometers, which is getting very close to the resolution needed for reverse engineering.

We also have noninvasive scanning techniques today, including high-resolution magnetic resonance imaging (MRI) scans, optical imaging, near-infrared scanning, and other technologies which are capable in certain instances of resolving individual somas, or neuron cell bodies. Brain scanning technologies are also increasing their resolution with each new generation, just what we would expect from the law of accelerating

returns. Future generations will enable us to resolve the connections between neurons and to peer inside the synapses and record the neurotransmitter concentrations.





We can peer inside someone's brain today with noninvasive scanners, which are increasing their resolution with each new generation of this technology. There are a number of technical challenges in accomplishing this, including achieving suitable resolution, bandwidth, lack of vibration, and safety. For a variety of reasons it is easier to scan the brain of someone recently deceased than of someone still living. It is easier to get someone deceased to sit still, for one thing. But noninvasively scanning a living brain will ultimately become feasible as MRI, optical, and other scanning technologies continue to improve in resolution and speed.

Scanning from Inside

Although noninvasive means of scanning the brain from outside the skull are rapidly improving, the most practical approach to capturing every salient neural detail will be to scan it from inside. By 2030, "nanobot" (i.e., nano robot) technology will be viable, and brain scanning will be a prominent application. Nanobots are robots that are the size of human blood cells, or even smaller. Billions of them could travel through every brain capillary and scan every relevant feature from up close. Using high speed wireless communication, the nanobots would communicate with each other, and with other computers that are compiling the brain scan data base (in other words, the nanobots will all be on a wireless local area network).

This scenario involves only capabilities that we can touch and feel today. We already have technology capable of producing very high resolution scans, provided that the scanner is physically proximate to the neural features. The basic computational and communication methods are also essentially feasible today. The primary features that are not yet practical are nanobot size and cost. As I discussed above, we can project the exponentially declining cost of computation, and the rapidly declining size of both electronic and mechanical technologies. We can conservatively expect, therefore, the requisite nanobot technology by around 2030.

Because of its ability to place each scanner in very close physical proximity to every neural feature, nanobot-based scanning will be more practical than scanning the brain from outside.

How to Use Your Brain Scan

How will we apply the thousands of trillions of bytes of information derived from each brain scan? One approach is to use the results to design more intelligent parallel algorithms for our machines, particularly those based on one of the neural net paradigms. With this approach, we don't have to copy every single connection. There is a great deal of repetition and redundancy within any particular brain region. Although the information contained in a human brain would require thousands of trillions of bytes of information (on the order of 100 billion neurons times an average of 1,000 connections per neuron, each with multiple neurotransmitter concentrations and connection data), the design of the brain is characterized by a human genome of only about a billion bytes.

Furthermore, most of the genome is redundant, so the initial design of the brain is characterized by approximately one hundred million bytes, about the size of Microsoft Word. Of course, the complexity of our brains greatly increases as we interact with the world (by a factor of more than ten million). Because of the highly repetitive patterns found in each specific brain region, it is not necessary to capture each detail in order to reverse engineer the significant digital-analog algorithms. With this information, we can design simulated nets that operate similarly. There are already multiple efforts under way to scan the human brain and apply the insights derived to the design of intelligent machines.

The pace of brain reverse engineering is only slightly behind the availability of the brain scanning and neuron structure information. A contemporary example is a comprehensive model of a significant portion of the human auditory processing system that Lloyd Watts (www.lloydwatts.com) has developed from both neurobiology studies of specific neuron types and brain interneuronal connection information. Watts' model includes five parallel paths and includes the actual intermediate representations of auditory information at each stage of neural processing. Watts has implemented his model as real-time software which can locate and identify sounds with many of the same properties as human hearing. Although a work in progress, the model illustrates the feasibility of converting neurobiological models and brain connection data into working simulations. Also, as Hans Moravec and others have speculated, these efficient simulations require about 1,000 times less computation than the theoretical potential of the biological neurons being simulated.

Reverse Engineering the Human Brain: Five Parallel Auditory Pathways

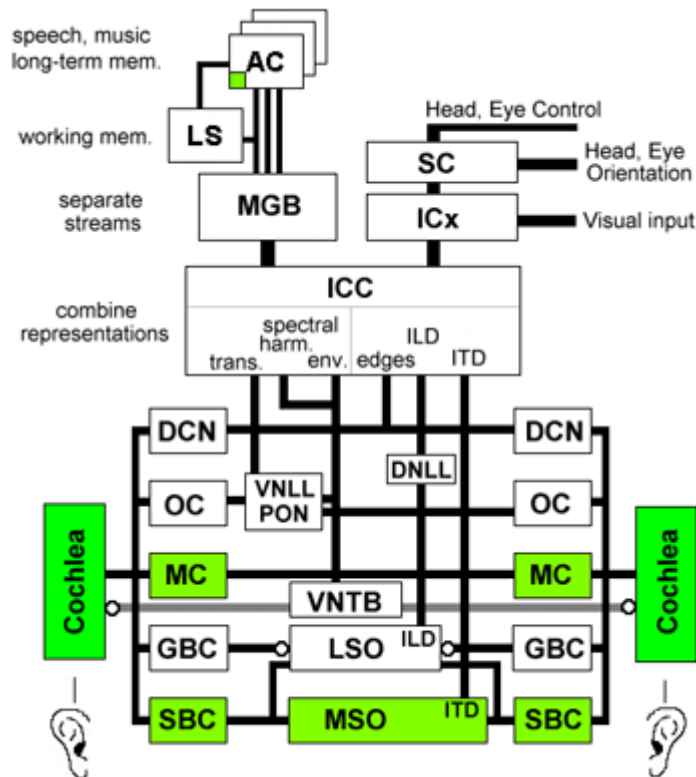


Chart by Lloyd Watts

Cochlea: Sense organ of hearing. 30,000 fibers converts motion of the stapes into spectro-temporal representation of sound.

MC: Multipolar Cells. Measure spectral energy.

GBC: Globular Bushy Cells. Relays spikes from the auditory nerve to the Lateral Superior.

Olivary Complex (includes **LSO** and **MSO**). Encoding of timing and amplitude of signals for binaural comparison of level.

SBC: Spherical Bushy Cells. Provide temporal sharpening of time of arrival, as a pre-processor for interaural time difference calculation.

OC: Octopus Cells. Detection of transients.

DCN: Dorsal Cochlear Nucleus. Detection of spectral edges and calibrating for noise levels.

VNTB: Ventral Nucleus of the Trapezoid Body. Feedback signals to modulate outer hair cell function in the cochlea.

VNLL, PON: Ventral Nucleus of the Lateral Lemniscus, Peri-Olivary Nuclei. Processing transients from the Octopus Cells.

MSO: Medial Superior Olive. Computing inter-aural time difference

(difference in time of arrival between the two ears, used to tell where a sound is coming from).

LSO: Lateral Superior Olive. Also involved in computing inter-aural level difference.

ICC: Central Nucleus of the Inferior Colliculus. The site of major integration of multiple representations of sound.

ICx: Exterior Nucleus of the Inferior Colliculus. Further refinement of sound localization.

SC: Superior Colliculus. Location of auditory/visual merging.

MGB: Medial Geniculate Body. The auditory portion of the thalamus.

LS: Limbic System. Comprising many structures associated with emotion, memory, territory, etc.

AC: Auditory Cortex.

The brain is not one huge "tabula rasa" (i.e., undifferentiated blank slate), but rather an intricate and intertwined collection of hundreds of specialized regions. The process of "peeling the onion" to understand these interleaved regions is well underway. As the requisite neuron models and brain interconnection data becomes available, detailed and implementable models such as the auditory example above will be developed for all brain regions.

After the algorithms of a region are understood, they can be refined and extended before being implemented in synthetic neural equivalents. For one thing, they can be run on a computational substrate that is already more than ten million times faster than neural circuitry. And we can also throw in the methods for building intelligent machines that we already understand.

Downloading the Human Brain

A more controversial application than this scanning-the-brain-to-understand-it scenario is *scanning-the-brain-to-download-it*. Here we scan someone's brain to map the locations, interconnections, and contents of all the somas, axons, dendrites, presynaptic vesicles, neurotransmitter concentrations, and other neural components and levels. Its entire organization can then be re-created on a neural computer of sufficient capacity, including the contents of its memory.

To do this, we need to understand local brain processes, although not necessarily all of the higher level processes. Scanning a brain with sufficient detail to download it may sound daunting, but so did the human genome scan. All of the basic technologies exist today, just not with the requisite speed, cost, and size, but these are the attributes that are improving at a double exponential pace.

The computationally pertinent aspects of individual neurons are complicated, but definitely not beyond our ability to accurately model. For example, Ted Berger and his colleagues at Hedco Neurosciences have built integrated circuits that precisely match the digital and analog information processing characteristics of neurons, including clusters with hundreds of neurons. Carver Mead and his colleagues at CalTech have built a variety of integrated circuits that emulate the digital-analog characteristics of mammalian neural circuits.

A recent experiment at San Diego's Institute for Nonlinear Science demonstrates the potential for electronic neurons to precisely emulate biological ones. Neurons (biological or otherwise) are a prime example of what is often called "chaotic computing." Each neuron acts in an essentially unpredictable fashion. When an entire network of neurons receives input (from the outside world or from other networks of neurons), the signaling amongst them appears at first to be frenzied and random. Over time, typically a fraction of a second or so, the chaotic interplay of the neurons dies down, and a stable pattern emerges. This pattern represents the "decision" of the neural network. If the neural network is performing a pattern recognition task (which, incidentally, comprises the bulk of the activity in the human brain), then the emergent pattern represents the appropriate recognition.

So the question addressed by the San Diego researchers was whether electronic neurons could engage in this chaotic dance alongside biological ones. They hooked up their artificial neurons with those from spiny lobsters in a single network, and their hybrid biological-nonbiological network performed in the same way (i.e., chaotic interplay followed by a stable emergent pattern) and with the same type of results as an all biological net of neurons. Essentially, the biological neurons accepted their electronic peers. It indicates that their mathematical model of these neurons was reasonably accurate.

There are many projects around the world which are creating nonbiological devices to recreate in great detail the functionality of human neuron clusters. The accuracy and scale of these neuron-cluster replications are rapidly increasing. We started with functionally equivalent recreations of single neurons, then clusters of tens, then hundreds, and now thousands. Scaling up technical processes at an exponential pace is what technology is good at.

As the computational power to emulate the human brain becomes available--we're not there yet, but we will be there within a couple of decades--projects already under way to scan the human brain will be accelerated, with a view both to understand the human brain in general, as well as providing a detailed description of the contents and design of specific brains. By the third decade of the twenty-first century, we will be in a position to create highly detailed and complete maps of all relevant features of all neurons, neural connections and synapses in the human brain, all of the neural details that play a role in the behavior and functionality of the brain, and to recreate these designs in suitably advanced neural computers.

Is the Human Brain Different from a Computer?

Is the human brain different from a computer?

The answer depends on what we mean by the word "computer." Certainly the brain uses very different methods from conventional contemporary computers. Most computers today are all digital and perform one (or perhaps a few) computations at a time at extremely high speed. In contrast, the human brain combines digital and analog methods with most computations performed in the analog domain. The brain is massively parallel, performing on the order of a hundred trillion computations at the same time, but at extremely slow speeds.

With regard to digital versus analog computing, we know that digital computing can be functionally equivalent to analog computing (although the reverse is not true), so we can perform all of the capabilities of a hybrid digital--analog network with an all digital computer. On the other hand, there is an engineering advantage to analog circuits in that analog computing is potentially thousands of times more efficient. An analog computation can be performed by a few transistors, or, in the case of mammalian neurons, specific electrochemical processes. A digital computation, in contrast, requires thousands or tens of thousands of transistors. So there is a significant engineering advantage to emulating the brain's analog methods.

The massive parallelism of the human brain is the key to its pattern recognition abilities, which reflects the strength of human thinking. As I discussed above, mammalian neurons engage in a chaotic dance, and if the neural network has learned its lessons well, then a stable pattern will emerge reflecting the network's decision. There is no reason why our nonbiological functionally equivalent recreations of biological neural networks cannot be built using these same principles, and indeed there are dozens of projects around the world that have succeeded in doing this. My own technical field is pattern recognition, and the projects that I have been involved in for over thirty years use this form of chaotic computing. Particularly successful examples are Carver Mead's neural chips, which are highly parallel, use digital controlled analog computing, and are intended as functionally similar recreations of biological networks.

Objective and Subjective

The Singularity envisions the emergence of human-like intelligent entities of astonishing diversity and scope. Although these entities will be capable of passing the "Turing test" (i.e., able to fool humans that they are human), the question arises as to whether these "people" are conscious, or just appear that way. To gain some insight as to why this is an extremely subtle question (albeit an ultimately important one) it is useful to consider some of the paradoxes that emerge from the concept of downloading specific human brains.

Although I anticipate that the most common application of the knowledge gained from reverse engineering the human brain will be creating more intelligent machines that are not necessarily modeled on specific biological human individuals, the scenario of scanning and reinstantiating all of the neural details of a *specific* person raises the most immediate questions of identity. Let's consider the question of what we will find

when we do this.

We have to consider this question on both the objective and subjective levels. "Objective" means everyone except me, so let's start with that. Objectively, when we scan someone's brain and reinstantiate their personal mind file into a suitable computing medium, the newly emergent "person" will appear to other observers to have very much the same personality, history, and memory as the person originally scanned. That is, once the technology has been refined and perfected. Like any new technology, it won't be perfect at first. But ultimately, the scans and recreations will be very accurate and realistic.

Interacting with the newly instantiated person will feel like interacting with the original person. The new person will claim to be that same old person and will have a memory of having been that person. The new person will have all of the patterns of knowledge, skill, and personality of the original. We are already creating functionally equivalent recreations of neurons and neuron clusters with sufficient accuracy that biological neurons accept their nonbiological equivalents and work with them as if they were biological. There are no natural limits that prevent us from doing the same with the hundred billion neuron cluster of clusters we call the human brain.

Subjectively, the issue is more subtle and profound, but first we need to reflect on one additional objective issue: our physical self.

The Importance of Having a Body

Consider how many of our thoughts and thinking are directed toward our body and its survival, security, nutrition, and image, not to mention affection, sexuality, and reproduction. Many, if not most, of the goals we attempt to advance using our brains have to do with our bodies: protecting them, providing them with fuel, making them attractive, making them feel good, providing for their myriad needs and desires. Some philosophers maintain that achieving human level intelligence is impossible without a body. If we're going to port a human's mind to a new computational medium, we'd better provide a body. A disembodied mind will quickly get depressed.

There are a variety of bodies that we will provide for our machines, and that they will provide for themselves: bodies built through nanotechnology (i.e., building highly complex physical systems atom by atom), virtual bodies (that exist only in virtual reality), bodies comprised of swarms of nanobots, and other technologies.

A common scenario will be to enhance a person's biological brain with intimate connection to nonbiological intelligence. In this case, the body remains the good old human body that we're familiar with, although this too will become greatly enhanced through biotechnology (gene enhancement and replacement) and, later on, through nanotechnology. A detailed examination of twenty-first century bodies is beyond the scope of this essay, but recreating and enhancing our bodies will be (and has been) an easier task than recreating our minds.

So Just Who Are These People?

To return to the issue of subjectivity, consider: is the reinstantiated mind the same consciousness as the person we just scanned? Are these "people" conscious at all? Is this a mind or just a brain?

Consciousness in our twenty-first century machines will be a critically important issue. But it is not easily resolved, or even readily understood. People tend to have strong views on the subject, and often just can't understand how anyone else could possibly see the issue from a different perspective. Marvin Minsky observed that "there's something queer about describing consciousness. Whatever people mean to say, they just can't seem to make it clear."

We don't worry, at least not yet, about causing pain and suffering to our computer programs. But at what point do we consider an entity, a process, to be conscious, to feel pain and discomfort, to have its own intentionality, its own free will? How do we determine if an entity is conscious; if it has subjective experience? How do we distinguish a process that is conscious from one that just acts *as if* it is conscious?

We can't simply ask it. If it says "Hey I'm conscious," does that settle the issue? No, we have computer games today that effectively do that, and they're not terribly convincing.

How about if the entity *is* very convincing and compelling when it says "I'm lonely, please keep me company." Does that settle the issue?

If we look inside its circuits, and see essentially the identical kinds of feedback loops and other mechanisms in its brain that we see in a human brain (albeit implemented using nonbiological equivalents), does that settle the issue?

And just who are these people in the machine, anyway? The answer will depend on who you ask. If you ask the people in the machine, they will strenuously claim to be the original persons. For example, if we scan--let's say myself--and record the exact state, level, and position of every neurotransmitter, synapse, neural connection, and every other relevant detail, and then restantiate this massive data base of information (which I estimate at thousands of trillions of bytes) into a neural computer of sufficient capacity, the person who then emerges in the machine will think that "he" is (and had been) me, or at least he will act that way. He will say "I grew up in Queens, New York, went to college at MIT, stayed in the Boston area, started and sold a few artificial intelligence companies, walked into a scanner there, and woke up in the machine here. Hey, this technology really works."

But wait.

Is this really me? For one thing, old biological Ray (that's me) still exists. I'll still be here in my carbon-cell-based brain. Alas, I will have to sit back and watch the new Ray succeed in endeavors that I could only dream of.

A Thought Experiment

Let's consider the issue of just who I am, and who the new Ray is a little more carefully. First of all, am I the stuff in my brain and body?

Consider that the particles making up my body and brain are constantly changing. We are not at all permanent collections of particles. The cells in our bodies turn over at different rates, but the particles (e.g., atoms and molecules) that comprise our cells are exchanged at a very rapid rate. I am just not the same collection of particles that I was even a month ago. It is the patterns of matter and energy that are semipermanent (that is, changing only gradually), but our actual material content is changing constantly, and very quickly. We are rather like the patterns that water makes in a stream. The rushing water around a formation of rocks makes a particular, unique pattern. This pattern may remain relatively unchanged for hours, even years. Of course, the actual material constituting the pattern--the water--is replaced in milliseconds. The same is true for Ray Kurzweil. Like the water in a stream, my particles are constantly changing, but the pattern that people recognize as Ray has a reasonable level of continuity. This argues that we should not associate our fundamental identity with a specific set of particles, but rather the pattern of matter and energy that we represent. Many contemporary philosophers seem partial to this "identify from pattern" argument.

But (again) wait.

If you were to scan my brain and reconstitute new Ray while I was sleeping, I would not necessarily even know about it (with the nanobots, this will be a feasible scenario). If you then come to me, and say, "good news, Ray, we've successfully reconstituted your mind file, so we won't be needing your old brain anymore," I may suddenly realize the flaw in this "identity from pattern" argument. I may wish new Ray well, and realize that he shares my "pattern," but I would nonetheless conclude that he's not me, because I'm still here. How could he be me? After all, I would not necessarily know that he even existed.

Let's consider another perplexing scenario. Suppose I replace a small number of biological neurons with functionally equivalent nonbiological ones (they may provide certain benefits such as greater reliability and longevity, but that's not relevant to this thought experiment). After I have this procedure performed, am I still the same person? My friends certainly think so. I still have the same self-deprecating humor, the same silly grin--yes, I'm still the same guy.

It should be clear where I'm going with this. Bit by bit, region by region, I ultimately replace my entire brain with essentially identical (perhaps improved) nonbiological equivalents (preserving all of the neurotransmitter concentrations and other details that represent my learning, skills, and memories). At each point, I feel the procedures were successful. At each point, I feel that I am the same guy. After each procedure, I claim to be the same guy. My friends concur. There is no old Ray and new Ray, just one Ray, one that never appears to fundamentally change.

But consider this. This gradual replacement of my brain with a nonbiological equivalent is essentially identical to the following sequence:

1. (i) scan Ray and reconstitute Ray's mind file into new (nonbiological) Ray, and, then
2. (ii) terminate old Ray. But we concluded above that in such a scenario new Ray is not the same as old Ray. And if old Ray is terminated, well then that's the end of Ray. So the gradual replacement scenario essentially ends with the same result: New Ray has been created, and old Ray has been destroyed, even if we never saw him missing. So what appears to be the continuing existence of just one Ray is really the creation of new Ray and the termination of old Ray.

On yet another hand (we're running out of philosophical hands here), the gradual replacement scenario is not altogether different from what happens normally to our biological selves, in that our particles are always rapidly being replaced. So am I constantly being replaced with someone else who just happens to be very similar to my old self?

I am trying to illustrate why consciousness is not an easy issue. If we talk about consciousness as just a certain type of intelligent skill: the ability to reflect on one's own self and situation, for example, then the issue is not difficult at all because any skill or capability or form of intelligence that one cares to define will be replicated in nonbiological entities (i.e., machines) within a few decades. With this type of *objective* view of consciousness, the conundrums do go away. But a fully objective view does not penetrate to the core of the issue, because the essence of consciousness is *subjective* experience, not objective correlates of that experience.

Will these future machines be capable of having spiritual experiences?

They certainly will claim to. They will claim to be people, and to have the full range of emotional and spiritual experiences that people claim to have. And these will not be idle claims; they will evidence the sort of rich, complex, and subtle behavior one associates with these feelings. How do the claims and behaviors--compelling as they will be--relate to the subjective experience of these reconstituted people? We keep coming back to the very real but ultimately unmeasurable issue of consciousness.

People often talk about consciousness as if it were a clear property of an entity that can readily be identified, detected, and gauged. If there is one crucial insight that we can make regarding why the issue of consciousness is so contentious, it is the following:

There exists no objective test that can conclusively determine its presence.

Science is about objective measurement and logical implications therefrom, but the very nature of objectivity is that you cannot measure subjective experience--you can only measure correlates of it, such as behavior (and by behavior, I include the actions of components of an entity, such as neurons). This limitation has to do with the very nature of

the concepts "objective" and "subjective." Fundamentally, we cannot penetrate the subjective experience of another entity with direct objective measurement. We can certainly make arguments about it: i.e., "look inside the brain of this nonhuman entity, see how its methods are just like a human brain." Or, "see how its behavior is just like human behavior." But in the end, these remain just arguments. No matter how convincing the behavior of a reinstantiated person, some observers will refuse to accept the consciousness of an entity unless it squirts neurotransmitters, or is based on DNA-guided protein synthesis, or has some other specific biologically human attribute.

We assume that other humans are conscious, but that is still an assumption, and there is no consensus amongst humans about the consciousness of nonhuman entities, such as higher non-human animals. The issue will be even more contentious with regard to future nonbiological entities with human-like behavior and intelligence.

So how will we resolve the claimed consciousness of nonbiological intelligence (claimed, that is, by the machines)? From a practical perspective, we'll accept their claims. Keep in mind that nonbiological entities in the twenty-first century will be extremely intelligent, so they'll be able to convince us that they are conscious. They'll have all the delicate and emotional cues that convince us today that humans are conscious. They will be able to make us laugh and cry. And they'll get mad if we don't accept their claims. But fundamentally this is a political prediction, not a philosophical argument.

On Tubules and Quantum Computing

Over the past several years, Roger Penrose, a noted physicist and philosopher, has suggested that fine structures in the neurons called tubules perform an exotic form of computation called "quantum computing." Quantum computing is computing using what are called "qu bits" which take on all possible combinations of solutions simultaneously. It can be considered to be an extreme form of parallel processing (because every combination of values of the qu bits are tested simultaneously). Penrose suggests that the tubules and their quantum computing capabilities complicate the concept of recreating neurons and reinstantiating mind files.

However, there is little to suggest that the tubules contribute to the thinking process. Even generous models of human knowledge and capability are more than accounted for by current estimates of brain size, based on contemporary models of neuron functioning that do not include tubules. In fact, even with these tubule-less models, it appears that the brain is conservatively designed with many more connections (by several orders of magnitude) than it needs for its capabilities and capacity. Recent experiments (e.g., the San Diego Institute for Nonlinear Science experiments) showing that hybrid biological-nonbiological networks perform similarly to all biological networks, while not definitive, are strongly suggestive that our tubule-less models of neuron functioning are adequate. Lloyd Watts' software simulation of his intricate model of human auditory processing uses orders of magnitude less computation than the networks of neurons he is simulating, and there is no suggestion

that quantum computing is needed.

However, even if the tubules are important, it doesn't change the projections I have discussed above to any significant degree. According to my model of computational growth, if the tubules multiplied neuron complexity by a factor of a thousand (and keep in mind that our current tubule-less neuron models are already complex, including on the order of a thousand connections per neuron, multiple nonlinearities and other details), this would delay our reaching brain capacity by only about 9 years. If we're off by a factor of a million, that's still only a delay of 17 years. A factor of a billion is around 24 years (keep in mind computation is growing by a double exponential).

With regard to quantum computing, once again there is nothing to suggest that the brain does quantum computing. Just because quantum technology may be feasible does not suggest that the brain is capable of it. After all, we don't have lasers or even radios in our brains. Although some scientists have claimed to detect quantum wave collapse in the brain, no one has suggested human capabilities that actually require a capacity for quantum computing.

However, even if the brain does do quantum computing, this does not significantly change the outlook for human-level computing (and beyond) nor does it suggest that brain downloading is infeasible. First of all, if the brain does do quantum computing this would only verify that quantum computing is feasible. There would be nothing in such a finding to suggest that quantum computing is restricted to biological mechanisms. Biological quantum computing mechanisms, if they exist, could be replicated. Indeed, recent experiments with small scale quantum computers appear to be successful. Even the conventional transistor relies on the quantum effect of electron tunneling.

Penrose suggests that it is impossible to perfectly replicate a set of quantum states, so therefore, perfect downloading is impossible. Well, how perfect does a download have to be? I am at this moment in a very different quantum state (and different in non-quantum ways as well) than I was a minute ago (certainly in a very different state than I was before I wrote this paragraph). If we develop downloading technology to the point where the "copies" are as close to the original as the original person changes anyway in the course of one minute, that would be good enough for any conceivable purpose, yet does not require copying quantum states. As the technology improves, the accuracy of the copy could become as close as the original changes within ever briefer periods of time (e.g., one second, one millisecond, one microsecond).

When it was pointed out to Penrose that neurons (and even neural connections) were too big for quantum computing, he came up with the tubule theory as a possible mechanism for neural quantum computing. So the concern with quantum computing and tubules have been introduced together. If one is searching for barriers to replicating brain function, it is an ingenious theory, but it fails to introduce any genuine barriers. There is no evidence for it, and even if true, it only delays matters by a decade or two. There is no reason to believe that biological mechanisms (including quantum computing) are inherently impossible to replicate using nonbiological materials and mechanisms. Dozens of contemporary

experiments are successfully performing just such replications.

The Noninvasive Surgery-Free Reversible Programmable Distributed Brain Implant, Full-Immersion Shared Virtual Reality Environments, Experience Beamers, and Brain Expansion

How will we apply technology that is more intelligent than its creators? One might be tempted to respond "Carefully!" But let's take a look at some examples.

Consider several examples of the nanobot technology, which, based on miniaturization and cost reduction trends, will be feasible within 30 years. In addition to scanning your brain, the nanobots will also be able to expand our experiences and our capabilities.

Nanobot technology will provide fully immersive, totally convincing virtual reality in the following way. The nanobots take up positions in close physical proximity to every interneuronal connection coming from all of our senses (e.g., eyes, ears, skin). We already have the technology for electronic devices to communicate with neurons in both directions that requires no direct physical contact with the neurons. For example, scientists at the Max Planck Institute have developed "neuron transistors" that can detect the firing of a nearby neuron, or alternatively, can cause a nearby neuron to fire, or suppress it from firing. This amounts to two-way communication between neurons and the electronic-based neuron transistors. The Institute scientists demonstrated their invention by controlling the movement of a living leech from their computer. Again, the primary aspect of nanobot-based virtual reality that is not yet feasible is size and cost.

When we want to experience real reality, the nanobots just stay in position (in the capillaries) and do nothing. If we want to enter virtual reality, they suppress all of the inputs coming from the real senses, and replace them with the signals that would be appropriate for the virtual environment. You (i.e., your brain) could decide to cause your muscles and limbs to move as you normally would, but the nanobots again intercept these interneuronal signals, suppress your real limbs from moving, and instead cause your virtual limbs to move and provide the appropriate movement and reorientation in the virtual environment.

The web will provide a panoply of virtual environments to explore. Some will be recreations of real places, others will be fanciful environments that have no "real" counterpart. Some indeed would be impossible in the physical world (perhaps, because they violate the laws of physics). We will be able to "go" to these virtual environments by ourselves, or we will meet other people there, both real people and simulated people. Of course, ultimately there won't be a clear distinction between the two.

By 2030, going to a web site will mean entering a full immersion virtual reality environment. In addition to encompassing all of the senses, these shared environments can include emotional overlays as the nanobots will be capable of triggering the neurological correlates of emotions, sexual

pleasure, and other derivatives of our sensory experience and mental reactions.

In the same way that people today beam their lives from web cams in their bedrooms, "experience beamers" circa 2030 will beam their entire flow of sensory experiences, and if so desired, their emotions and other secondary reactions. We'll be able to plug in (by going to the appropriate web site) and experience other people's lives as in the plot concept of 'Being John Malkovich.' Particularly interesting experiences can be archived and relived at any time.

We won't need to wait until 2030 to experience shared virtual reality environments, at least for the visual and auditory senses. Full immersion visual-auditory environments will be available by the end of this decade with images written directly onto our retinas by our eyeglasses and contact lenses. All of the electronics for the computation, image reconstruction, and very high bandwidth wireless connection to the Internet will be embedded in our glasses and woven into our clothing, so computers as distinct objects will disappear.

In my view, the most significant implication of the Singularity will be the merger of biological and nonbiological intelligence. First, it is important to point out that well before the end of the twenty-first century, thinking on nonbiological substrates will dominate. Biological thinking is stuck at 10^{26} calculations per second (for all biological human brains), and that figure will not appreciably change, even with bioengineering changes to our genome. Nonbiological intelligence, on the other hand, is growing at a double exponential rate and will vastly exceed biological intelligence well before the middle of this century. However, in my view, this nonbiological intelligence should still be considered human as it is fully derivative of the human-machine civilization. The merger of these two worlds of intelligence is not merely a merger of biological and nonbiological thinking mediums, but more importantly one of method and organization of thinking.

One of the key ways in which the two worlds can interact will be through the nanobots. Nanobot technology will be able to expand our minds in virtually any imaginable way. Our brains today are relatively fixed in design. Although we do add patterns of interneuronal connections and neurotransmitter concentrations as a normal part of the learning process, the current overall capacity of the human brain is highly constrained, restricted to a mere hundred trillion connections. Brain implants based on massively distributed intelligent nanobots will ultimately expand our memories a trillion fold, and otherwise vastly improve all of our sensory, pattern recognition, and cognitive abilities. Since the nanobots are communicating with each other over a wireless local area network, they can create any set of new neural connections, can break existing connections (by suppressing neural firing), can create new hybrid biological-nonbiological networks, as well as add vast new nonbiological networks.

Using nanobots as brain extenders is a significant improvement over the idea of surgically installed neural implants, which are beginning to be used today (e.g., ventral posterior nucleus, subthalamic nucleus, and

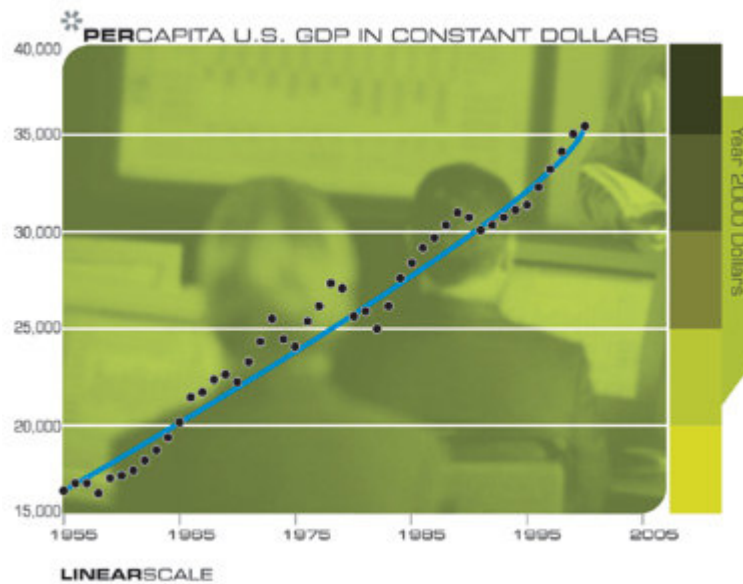
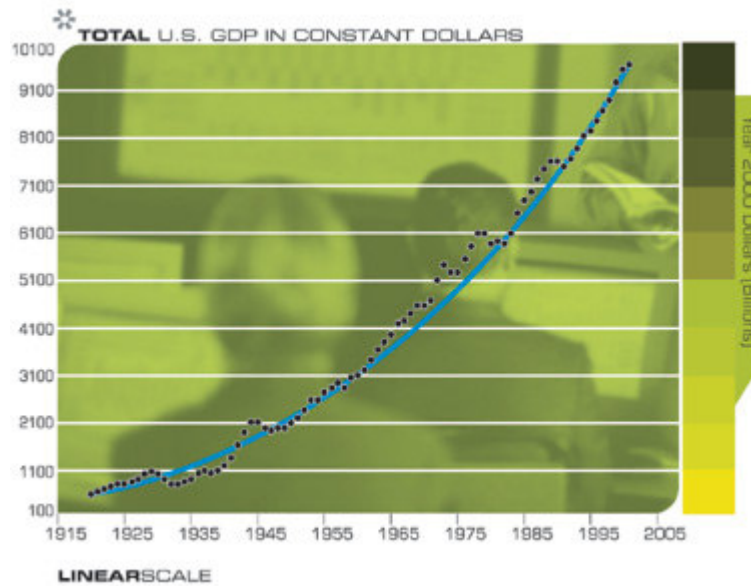
ventral lateral thalamus neural implants to counteract Parkinson's Disease and tremors from other neurological disorders, cochlear implants, and others.) Nanobots will be introduced without surgery, essentially just by injecting or even swallowing them. They can all be directed to leave, so the process is easily reversible. They are programmable, in that they can provide virtual reality one minute, and a variety of brain extensions the next. They can change their configuration, and clearly can alter their software. Perhaps most importantly, they are massively distributed and therefore can take up billions or trillions of positions throughout the brain, whereas a surgically introduced neural implant can only be placed in one or at most a few locations.

The Double Exponential Growth of the Economy During the 1990s Was Not a Bubble

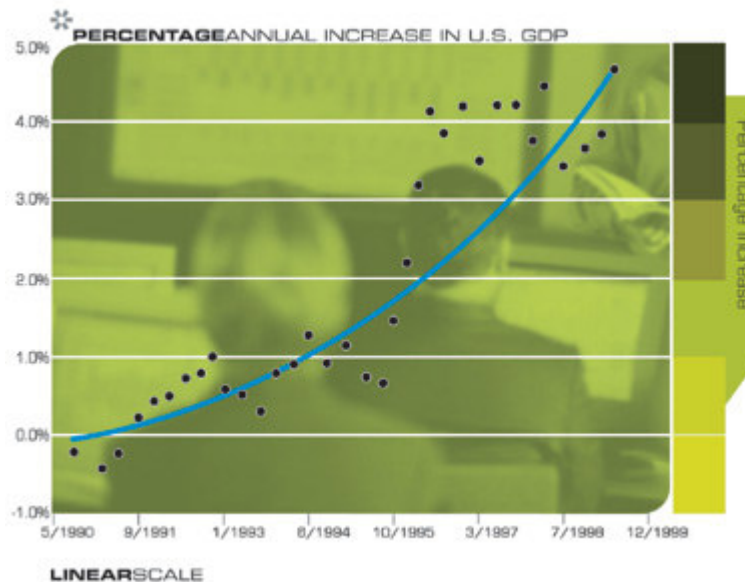
Yet another manifestation of the law of accelerating returns as it rushes toward the Singularity can be found in the world of economics, a world vital to both the genesis of the law of accelerating returns, and to its implications. It is the economic imperative of a competitive marketplace that is driving technology forward and fueling the law of accelerating returns. In turn, the law of accelerating returns, particularly as it approaches the Singularity, is transforming economic relationships.

Virtually all of the economic models taught in economics classes, used by the Federal Reserve Board to set monetary policy, by Government agencies to set economic policy, and by economic forecasters of all kinds are fundamentally flawed because they are based on the intuitive linear view of history rather than the historically based exponential view. The reason that these linear models appear to work for a while is for the same reason that most people adopt the intuitive linear view in the first place: exponential trends appear to be linear when viewed (and experienced) for a brief period of time, particularly in the early stages of an exponential trend when not much is happening. But once the "knee of the curve" is achieved and the exponential growth explodes, the linear models break down. The exponential trends underlying productivity growth are just beginning this explosive phase.

The economy (viewed either in total or per capita) has been growing exponentially throughout this century:



There is also a second level of exponential growth, but up until recently the second exponent has been in the early phase so that the growth in the growth rate has not been noticed. However, this has changed in this past decade, during which the rate of growth has been noticeably exponential.




Productivity (economic output per worker) has also been growing exponentially. Even these statistics are greatly understated because they do not fully reflect significant improvements in the quality and features of products and services. It is not the case that "a car is a car;" there have been significant improvements in safety, reliability, and features. There are a myriad of such examples. Pharmaceutical drugs are increasingly effective. Groceries ordered in five minutes on the web and delivered to your door are worth more than groceries on a supermarket shelf that you have to fetch yourself. Clothes custom manufactured for your unique body scan are worth more than clothes you happen to find left on a store rack. These sorts of improvements are true for most product categories, and none of them are reflected in the productivity statistics.

The statistical methods underlying the productivity measurements tend to factor out gains by essentially concluding that we still only get one dollar of products and services for a dollar despite the fact that we get much more for a dollar (e.g., compare a \$1,000 computer today to one ten years ago). University of Chicago Professor Pete Klenow and University of Rochester Professor Mark Bils estimate that the value of existing goods has been increasing at 1.5% per year for the past 20 years because of qualitative improvements. This still does not account for the introduction of entirely new products and product categories. The Bureau of Labor Statistics, which is responsible for the inflation statistics, uses a model that incorporates an estimate of quality growth at only 0.5% per year, reflecting a systematic underestimate of quality improvement and a resulting overestimate of inflation by at least 1 percent per year.

Despite these weaknesses in the productivity statistical methods, the gains in productivity are now reaching the steep part of the exponential curve. Labor productivity grew at 1.6% per year until 1994, then rose at 2.4% per year, and is now growing even more rapidly. In the quarter ending July 30, 2000, labor productivity grew at 5.3%. Manufacturing productivity grew at 4.4% annually from 1995 to 1999, durables

manufacturing at 6.5% per year.



The 1990s have seen the most powerful deflationary forces in history. This is why we are not seeing inflation. Yes, it's true that low unemployment, high asset values, economic growth, and other such factors are inflationary, but these factors are offset by the double exponential trends in the price-performance of all information based technologies: computation, memory, communications, biotechnology, miniaturization, and even the overall rate of technical progress. These technologies deeply affect all industries.

We are also undergoing massive disintermediation in the channels of distribution through the web and other new communication technologies, as well as escalating efficiencies in operations and administration.


All of the technology trend charts in this essay represent massive deflation. There are many examples of the impact of these escalating efficiencies. BP Amoco's cost for finding oil is now less than \$1 per barrel, down from nearly \$10 in 1991. Processing an internet transaction costs a bank one penny, compared to over \$1 using a teller ten years ago. A Roland Berger / Deutsche Bank study estimates a cost savings of \$1200 per North American car over the next five years. A more optimistic Morgan Stanley study estimates that Internet-based procurement will save Ford, GM, and DaimlerChrysler about \$2700 per vehicle. Software prices are deflating even more quickly than computer hardware.

Software Price-Performance Has Also Improved at an Exponential Rate


(Example: Automatic Speech Recognition Software)

	1985	1995	2000
Price	\$5,000	\$500	\$50
Vocabulary Size (# words)	1,000	10,000	100,000

Continuous Speech?	No	No	Yes
User Training Required (Minutes)	180	60	5
Accuracy	Poor	Fair	Good



Current economic policy is based on outdated models which include energy prices, commodity prices, and capital investment in plant and equipment as key driving factors, but do not adequately model bandwidth, MIPs, megabytes, intellectual property, knowledge, and other increasingly vital (and increasingly increasing) constituents that are driving the economy.



The economy "wants" to grow more than the 3.5% per year, which constitutes the current "speed limit" that the Federal Reserve bank and other policy makers have established as "safe," meaning noninflationary. But in keeping with the law of accelerating returns, the economy is capable of "safely" establishing this level of growth in less than a year, implying a growth rate in an entire year of greater than 3.5%. Recently, the growth rate has exceeded 5%.

None of this means that cycles of recession will disappear immediately. The economy still has some of the underlying dynamics that historically have caused cycles of recession, specifically excessive commitments such as capital intensive projects and the overstocking of inventories. However, the rapid dissemination of information, sophisticated forms of online procurement, and increasingly transparent markets in all industries have diminished the impact of this cycle. So "recessions" are likely to be shallow and short lived. The underlying long-term growth rate will continue at a double exponential rate.

The overall growth of the economy reflects completely new forms and layers of wealth and value that did not previously exist, or least that did not previously constitute a significant portion of the economy (but do now): intellectual property, communication portals, web sites, bandwidth, software, data bases, and many other new technology based categories.

There is no need for high interest rates to counter an inflation that doesn't exist. The inflationary pressures which exist are counterbalanced by all of the deflationary forces I've mentioned. The current high interest rates fostered by the Federal Reserve Bank are destructive, are causing trillions of dollars of lost wealth, are regressive, hurt business and the middle class, and are completely unnecessary.

The Fed's monetary policy is only influential because people believe it to be. It has little real power. The economy today is largely backed by private capital in the form of a growing variety of equity instruments. The portion of available liquidity in the economy that the Fed actually controls is relatively insignificant. The reserves that banks and financial institutions maintain with the Federal Reserve System are less than \$50 billion, which is only 0.6% of the GDP, and 0.25% of the liquidity available in stocks.

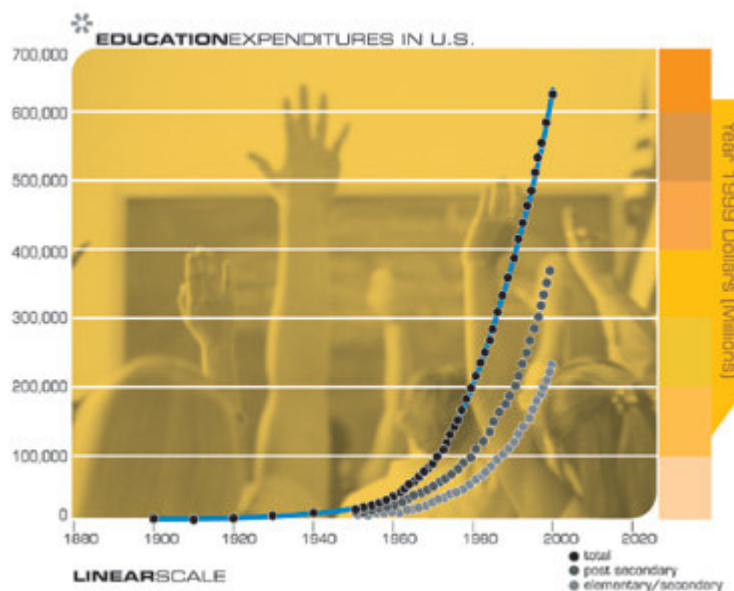
Restricting the growth rate of the economy to an arbitrary limit makes as much sense as restricting the rate at which a company can grow its revenues--or its market cap. Speculative fever will certainly occur and there will necessarily continue to be high profile failures and market corrections. However the ability of technology companies to rapidly create new--real--wealth is just one of the factors that will continue to fuel ongoing double exponential growth in the economy. These policies have led to an "Alice in Wonderland" situation in which the market goes up on bad economic news (because it means that more unnecessary punishment will be avoided) and goes down on good economic news.


Speaking of market speculative fever and market corrections, the stock market values for so-called "B to B" (Business to Business) and "B to C" (Business to Consumer) web portals and enabling technologies is likely

to come back strongly as it becomes clear that economic transactions are indeed escalating toward e-commerce, and that the (surviving) contenders are capable of demonstrating profitable business models.

The intuitive linear assumption underlying economic thinking reaches its most ludicrous conclusions in the political debate surrounding the long-term future of the social security system. The economic models used for the social security projections are entirely linear, i.e., they reflect fixed economic growth. This might be viewed as conservative planning if we were talking about projections of only a few years, but they become utterly unrealistic for the three to four decades being discussed. These projections actually assume a fixed rate of growth of 3.5% per year for the next fifty years! There are incredibly naïve assumptions that bear on both sides of the argument. On the one hand, there will be radical extensions to human longevity, while on the other hand, we will benefit from far greater economic expansion. These factors do not rule each other out, however, as the positive factors are stronger, and will ultimately dominate. Moreover, we are certain to rethink social security when we have centenarians who look and act like 30 year-olds (but who will think much faster than 30 year-olds circa the year 2000).

Another implication of the law of accelerating returns is exponential growth in education and learning. Over the past 120 years, we have increased our investment in K-12 education (per student and in constant dollars) by a factor of ten. We have a one hundred fold increase in the number of college students. Automation started by amplifying the power of our muscles, and in recent times has been amplifying the power of our minds. Thus, for the past two centuries, automation has been eliminating jobs at the bottom of the skill ladder while creating new (and better paying) jobs at the top of the skill ladder. So the ladder has been moving up, and thus we have been exponentially increasing investments in education at all levels.





Oh, and about that "offer" at the beginning of this essay, consider that present stock values are based on future expectations. Given that the (literally) short sighted linear intuitive view represents the ubiquitous outlook, the common wisdom in economic expectations are dramatically understated. Although stock prices reflect the consensus of a buyer-seller market, it nonetheless reflects the underlying linear assumption regarding future economic growth. But the law of accelerating returns clearly implies that the growth rate will continue to grow exponentially because the rate of progress will continue to accelerate. Although (weakening) recessionary cycles will continue to cause immediate growth rates to fluctuate, the underlying rate of growth will continue to double approximately every decade.

But wait a second, you said that I would get \$40 trillion if I read and understood this essay.

That's right. According to my models, if we replace the linear outlook with the more appropriate exponential outlook, current stock prices should triple. Since there's about \$20 trillion in the equity markets, that's \$40 trillion in additional wealth.

But you said I would get that money.

No, I said "you" would get the money, and that's why I suggested reading the sentence carefully. The English word "you" can be singular or plural. I meant it in the sense of "all of you."

I see, all of us as in the whole world. But not everyone will read this essay.

Well, but everyone could. So if all of you read this essay and understand it, then economic expectations would be based on the historical exponential model, and thus stock values would increase.

You mean if everyone understands it, and agrees with it.

Okay, I suppose I was assuming that.

Is that what you expect to happen.

Well, actually, no. Putting on my futurist hat again, my prediction is that indeed these views will prevail, but only over time, as more and more evidence of the exponential nature of technology and its impact on the economy becomes apparent. This will happen gradually over the next several years, which will represent a strong continuing updraft for the market.

A Clear and Future Danger

Technology has always been a double edged sword, bringing us longer and healthier life spans, freedom from physical and mental drudgery, and many new creative possibilities on the one hand, while introducing new and salient dangers on the other. We still live today with sufficient nuclear weapons (not all of which appear to be well accounted for) to end all mammalian life on the planet. Bioengineering is in the early stages of enormous strides in reversing disease and aging processes. However, the means and knowledge will soon exist in a routine college bioengineering lab (and already exists in more sophisticated labs) to create unfriendly pathogens more dangerous than nuclear weapons. As technology accelerates toward the Singularity, we will see the same intertwined potentials: a feast of creativity resulting from human intelligence expanded a trillion-fold combined with many grave new dangers.

Consider unrestrained nanobot replication. Nanobot technology requires billions or trillions of such intelligent devices to be useful. The most cost effective way to scale up to such levels is through self-replication, essentially the same approach used in the biological world. And in the same way that biological self-replication gone awry (i.e., cancer) results in biological destruction, a defect in the mechanism curtailing nanobot self-replication would endanger all physical entities, biological or otherwise.

Other primary concerns include "who is controlling the nanobots?" and "who are the nanobots talking to?" Organizations (e.g., governments, extremist groups) or just a clever individual could put trillions of undetectable nanobots in the water or food supply of an individual or of an entire population. These "spy" nanobots could then monitor, influence, and even control our thoughts and actions. In addition to introducing physical spy nanobots, existing nanobots could be influenced through software viruses and other software "hacking" techniques. When there is software running in our brains, issues of privacy and security will take on a new urgency.

My own expectation is that the creative and constructive applications of this technology will dominate, as I believe they do today. But there will be a valuable (and increasingly vocal) role for a concerned and constructive Luddite movement (i.e., anti-technologists inspired by early nineteenth century weavers who destroyed labor-saving machinery in

protest).

If we imagine describing the dangers that exist today to people who lived a couple of hundred years ago, they would think it mad to take such risks. On the other hand, how many people in the year 2000 would really want to go back to the short, brutish, disease-filled, poverty-stricken, disaster-prone lives that 99 percent of the human race struggled through a couple of centuries ago? We may romanticize the past, but up until fairly recently, most of humanity lived extremely fragile lives where one all too common misfortune could spell disaster. Substantial portions of our species still live in this precarious way, which is at least one reason to continue technological progress and the economic enhancement that accompanies it.

People often go through three stages in examining the impact of future technology: awe and wonderment at its potential to overcome age old problems, then a sense of dread at a new set of grave dangers that accompany these new technologies, followed, finally and hopefully, by the realization that the only viable and responsible path is to set a careful course that can realize the promise while managing the peril.

In his cover story for WIRED *Why The Future Doesn't Need Us*, Bill Joy eloquently described the plagues of centuries' past, and how new self-replicating technologies, such as mutant bioengineered pathogens, and "nanobots" run amok, may bring back long forgotten pestilence. Indeed these are real dangers. It is also the case, which Joy acknowledges, that it has been technological advances, such as antibiotics and improved sanitation, which has freed us from the prevalence of such plagues. Suffering in the world continues and demands our steadfast attention. Should we tell the millions of people afflicted with cancer and other devastating conditions that we are canceling the development of all bioengineered treatments because there is a risk that these same technologies may someday be used for malevolent purposes? Having asked the rhetorical question, I realize that there is a movement to do exactly that, but I think most people would agree that such broad based relinquishment is not the answer.

The continued opportunity to alleviate human distress is one important motivation for continuing technological advancement. Also compelling are the already apparent economic gains I discussed above which will continue to hasten in the decades ahead. The continued acceleration of many intertwined technologies are roads paved with gold (I use the plural here because technology is clearly not a single path). In a competitive environment, it is an economic imperative to go down these roads. Relinquishing technological advancement would be economic suicide for individuals, companies, and nations.

Which brings us to the issue of relinquishment, which is Bill Joy's most controversial recommendation and personal commitment. I do feel that relinquishment at the right level is part of a responsible and constructive response to these genuine perils. The issue, however, is exactly this: at what level are we to relinquish technology?

Ted Kaczynski would have us renounce all of it. This, in my view, is

neither desirable nor feasible, and the futility of such a position is only underscored by the senselessness of Kaczynski's deplorable tactics.

Another level would be to forego certain fields; nanotechnology, for example, that might be regarded as too dangerous. But such sweeping strokes of relinquishment are equally untenable. Nanotechnology is simply the inevitable end result of the persistent trend toward miniaturization which pervades all of technology. It is far from a single centralized effort, but is being pursued by a myriad of projects with many diverse goals.

One observer wrote:

"A further reason why industrial society cannot be reformed. . . is that modern technology is a unified system in which all parts are dependent on one another. You can't get rid of the "bad" parts of technology and retain only the "good" parts. Take modern medicine, for example. Progress in medical science depends on progress in chemistry, physics, biology, computer science and other fields. Advanced medical treatments require expensive, high-tech equipment that can be made available only by a technologically progressive, economically rich society. Clearly you can't have much progress in medicine without the whole technological system and everything that goes with it."

The observer I am quoting is, again, Ted Kaczynski. Although one might properly resist Kaczynski as an authority, I believe he is correct on the deeply entangled nature of the benefits and risks. However, Kaczynski and I clearly part company on our overall assessment on the relative balance between the two. Bill Joy and I have dialogued on this issue both publicly and privately, and we both believe that technology will and should progress, and that we need to be actively concerned with the dark side. If Bill and I disagree, it's on the granularity of relinquishment that is both feasible and desirable.

Abandonment of broad areas of technology will only push them underground where development would continue unimpeded by ethics and regulation. In such a situation, it would be the less stable, less responsible practitioners (e.g., the terrorists) who would have all the expertise.

I do think that relinquishment at the right level needs to be part of our ethical response to the dangers of twenty first century technologies. One constructive example of this is the proposed ethical guideline by the Foresight Institute, founded by nanotechnology pioneer Eric Drexler, that nanotechnologists agree to relinquish the development of physical entities that can self-replicate in a natural environment. Another is a ban on self-replicating physical entities that contain their own codes for self-replication. In what nanotechnologist Ralph Merkle calls the "Broadcast Architecture," such entities would have to obtain such codes from a centralized secure server, which would guard against undesirable replication. The Broadcast Architecture is impossible in the biological world, which represents at least one way in which nanotechnology can be made safer than biotechnology. In other ways, nanotech is potentially more dangerous because nanobots can be physically stronger than

protein-based entities and more intelligent. It will eventually be possible to combine the two by having nanotechnology provide the codes within biological entities (replacing DNA), in which case biological entities can use the much safer Broadcast Architecture.

Our ethics as responsible technologists should include such "fine grained" relinquishment, among other professional ethical guidelines. Other protections will need to include oversight by regulatory bodies, the development of technology-specific "immune" responses, as well as computer assisted surveillance by law enforcement organizations. Many people are not aware that our intelligence agencies already use advanced technologies such as automated word spotting to monitor a substantial flow of telephone conversations. As we go forward, balancing our cherished rights of privacy with our need to be protected from the malicious use of powerful twenty first century technologies will be one of many profound challenges. This is one reason that such issues as an encryption "trap door" (in which law enforcement authorities would have access to otherwise secure information) and the FBI "Carnivore" email-snooping system have been so contentious.

As a test case, we can take a small measure of comfort from how we have dealt with one recent technological challenge. There exists today a new form of fully nonbiological self replicating entity that didn't exist just a few decades ago: the computer virus. When this form of destructive intruder first appeared, strong concerns were voiced that as they became more sophisticated, software pathogens had the potential to destroy the computer network medium they live in. Yet the "immune system" that has evolved in response to this challenge has been largely effective. Although destructive self-replicating software entities do cause damage from time to time, the injury is but a small fraction of the benefit we receive from the computers and communication links that harbor them. No one would suggest we do away with computers, local area networks, and the Internet because of software viruses.

One might counter that computer viruses do not have the lethal potential of biological viruses or of destructive nanotechnology. Although true, this strengthens my observation. The fact that computer viruses are not usually deadly to humans only means that more people are willing to create and release them. It also means that our response to the danger is that much less intense. Conversely, when it comes to self replicating entities that are potentially lethal on a large scale, our response on all levels will be vastly more serious.

Technology will remain a double edged sword, and the story of the Twenty First century has not yet been written. It represents vast power to be used for all humankind's purposes. We have no choice but to work hard to apply these quickening technologies to advance our human values, despite what often appears to be a lack of consensus on what those values should be.

Living Forever

Once brain porting technology has been refined and fully developed, will this enable us to live forever? The answer depends on what we mean by

living and dying. Consider what we do today with our personal computer files. When we change from one personal computer to a less obsolete model, we don't throw all our files away; rather we copy them over to the new hardware. Although our software files do not necessarily continue their existence forever, the longevity of our personal computer software is completely separate and disconnected from the hardware that it runs on. When it comes to our personal mind file, however, when our human hardware crashes, the software of our lives dies with it. However, this will not continue to be the case when we have the means to store and restore the thousands of trillions of bytes of information represented in the pattern that we call our brains.

The longevity of one's mind file will not be dependent, therefore, on the continued viability of any particular hardware medium. Ultimately software-based humans, albeit vastly extended beyond the severe limitations of humans as we know them today, will live out on the web, projecting bodies whenever they need or want them, including virtual bodies in diverse realms of virtual reality, holographically projected bodies, physical bodies comprised of nanobot swarms, and other forms of nanotechnology.

A software-based human will be free, therefore, from the constraints of any particular thinking medium. Today, we are each confined to a mere hundred trillion connections, but humans at the end of the twenty-first century can grow their thinking and thoughts without limit. We may regard this as a form of immortality, although it is worth pointing out that data and information do not necessarily last forever. Although not dependent on the viability of the hardware it runs on, the longevity of information depends on its relevance, utility, and accessibility. If you've ever tried to retrieve information from an obsolete form of data storage in an old obscure format (e.g., a reel of magnetic tape from a 1970 minicomputer), you will understand the challenges in keeping software viable. However, if we are diligent in maintaining our mind file, keeping current backups, and porting to current formats and mediums, then a form of immortality can be attained, at least for software-based humans. Our mind file--our personality, skills, memories--all of that is lost today when our biological hardware crashes. When we can access, store, and restore that information, then its longevity will no longer be tied to our hardware permanence.

Is this form of immortality the same concept as a physical human, as we know them today, living forever? In one sense it is, because as I pointed out earlier, our contemporary selves are not a constant collection of matter either. Only our pattern of matter and energy persists, and even that gradually changes. Similarly, it will be the pattern of a software human that persists and develops and changes gradually.

But is that person based on my mind file, who migrates across many computational substrates, and who outlives any particular thinking medium, really me? We come back to the same questions of consciousness and identity, issues that have been debated since the Platonic dialogues. As we go through the twenty-first century, these will not remain polite philosophical debates, but will be confronted as vital, practical, political, and legal issues.

A related question is "is death desirable?" A great deal of our effort goes into avoiding it. We make extraordinary efforts to delay it, and indeed often consider its intrusion a tragic event. Yet we might find it hard to live without it. We consider death as giving meaning to our lives. It gives importance and value to time. Time could become meaningless if there were too much of it.

The Next Step in Evolution and the Purpose of Life

But I regard the freeing of the human mind from its severe physical limitations of scope and duration as the necessary next step in evolution. Evolution, in my view, represents the purpose of life. That is, the purpose of life--and of our lives--is to evolve. The Singularity then is not a grave danger to be avoided. In my view, this next paradigm shift represents the goal of our civilization.

What does it mean to evolve? Evolution moves toward greater complexity, greater elegance, greater knowledge, greater intelligence, greater beauty, greater creativity, and more of other abstract and subtle attributes such as love. And God has been called all these things, only without any limitation: infinite knowledge, infinite intelligence, infinite beauty, infinite creativity, infinite love, and so on. Of course, even the accelerating growth of evolution never achieves an infinite level, but as it explodes exponentially, it certainly moves rapidly in that direction. So evolution moves inexorably toward our conception of God, albeit never quite reaching this ideal. Thus the freeing of our thinking from the severe limitations of its biological form may be regarded as an essential spiritual quest.

In making this statement, it is important to emphasize that terms like evolution, destiny, and spiritual quest are observations about the end result, not the basis for these predictions. I am not saying that technology will evolve to human levels and beyond simply because it is our destiny and because of the satisfaction of a spiritual quest. Rather my projections result from a methodology based on the dynamics underlying the (double) exponential growth of technological processes. The primary force driving technology is economic imperative. We are moving toward machines with human level intelligence (and beyond) as the result of millions of small advances, each with their own particular economic justification.

To use an example from my own experience at one of my companies (Kurzweil Applied Intelligence), whenever we came up with a slightly more intelligent version of speech recognition, the new version invariably had greater value than the earlier generation and, as a result, sales increased. It is interesting to note that in the example of speech recognition software, the three primary surviving competitors stayed very close to each other in the intelligence of their software. A few other companies that failed to do so (e.g., Speech Systems) went out of business. At any point in time, we would be able to sell the version prior to the latest version for perhaps a quarter of the price of the current version. As for versions of our technology that were two generations old, we couldn't even give those away. This phenomenon is not only true for pattern recognition and other "AI" software, but applies to all products,

from bread makers to cars. And if the product itself doesn't exhibit some level of intelligence, then intelligence in the manufacturing and marketing methods have a major effect on the success and profitability of an enterprise.

There is a vital economic imperative to create more intelligent technology. Intelligent machines have enormous value. That is why they are being built. There are tens of thousands of projects that are advancing intelligent machines in diverse incremental ways. The support for "high tech" in the business community (mostly software) has grown enormously. When I started my optical character recognition (OCR) and speech synthesis company (Kurzweil Computer Products, Inc.) in 1974, there were only a half-dozen high technology IPO's that year. The number of such deals has increased one hundred fold and the number of dollars invested has increased by more than one thousand fold in the past 25 years. In the four years between 1995 and 1999 alone, high tech venture capital deals increased from just over \$1 billion to approximately \$15 billion.

We will continue to build more powerful computational mechanisms because it creates enormous value. We will reverse-engineer the human brain not simply because it is our destiny, but because there is valuable information to be found there that will provide insights in building more intelligent (and more valuable) machines. We would have to repeal capitalism and every visage of economic competition to stop this progression.

By the second half of this next century, there will be no clear distinction between human and machine intelligence. On the one hand, we will have biological brains vastly expanded through distributed nanobot-based implants. On the other hand, we will have fully nonbiological brains that are copies of human brains, albeit also vastly extended. And we will have a myriad of other varieties of intimate connection between human thinking and the technology it has fostered.

Ultimately, nonbiological intelligence will dominate because it is growing at a double exponential rate, whereas for all practical purposes biological intelligence is at a standstill. Human thinking is stuck at 10^{26} calculations per second (for all biological humans), and that figure will never appreciably change (except for a small increase resulting from genetic engineering). Nonbiological thinking is still millions of times less today, but the cross over will occur before 2030. By the end of the twenty-first century, nonbiological thinking will be trillions of trillions of times more powerful than that of its biological progenitors, although still of human origin. It will continue to be the human-machine civilization taking the next step in evolution.

Most forecasts of the future seem to ignore the revolutionary impact of the Singularity in our human destiny: the inevitable emergence of computers that match and ultimately vastly exceed the capabilities of the human brain, a development that will be no less important than the evolution of human intelligence itself some thousands of centuries ago. And the primary reason for this failure is that they are based on the intuitive but short sighted linear view of history.

Before the next century is over, the Earth's technology-creating species will merge with its computational technology. There will not be a clear distinction between human and machine. After all, what is the difference between a human brain enhanced a trillion fold by nanobot-based implants, and a computer whose design is based on high resolution scans of the human brain, and then extended a trillion-fold?

Why SETI Will Fail (and why we are alone in the Universe)

The law of accelerating returns implies that by 2099, the intelligence that will have emerged from human-machine civilization will be trillions of trillions of times more powerful than it is today, dominated of course by its nonbiological form.

So what does this have to do with SETI (the Search for Extra Terrestrial Intelligence)? The naïve view, going back to pre-Copernican days, was that the Earth was at the center of the Universe, and human intelligence its greatest gift (next to God). The more informed recent view is that even if the likelihood of a star having a planet with a technology creating species is very low (e.g., one in a million), there are so many stars (i.e., billions of trillions of them), that there are bound to be many with advanced technology.

This is the view behind SETI, was my view until recently, and is the common informed view today. Although SETI has not yet looked everywhere, it has already covered a substantial portion of the Universe.




Chart by Scientific American

In the above diagram (courtesy of Scientific American), we can see that SETI has already thoroughly searched all star systems within 10^7 light-years from Earth for alien civilizations capable (and willing) to transmit at a power of at least 10^{25} watts, a so-called Type II civilization (and all star systems within 10^6 light-years for transmission of at least 10^{18} watts, and so on). No sign of intelligence has been found as of yet.

In a recent email to my research assistant, Dr. Seth Shostak of the SETI Institute points out that a new comprehensive targeted search, called Project Phoenix, which has up to 100 times the sensitivity and covers a greater range of the radio dial as compared to previous searches, has only been applied thus far to 500 star systems, which is, of course only a minute fraction of the half trillion star systems in just our own galaxy.

However, according to my model, once a civilization achieves our own level ("Earth-level") of radio transmission, it takes no more than one century, two at the most, to achieve what SETI calls a Type II civilization. If the assumption that there are at least millions of radio capable civilizations out there, and that these civilizations are spread out over millions (indeed billions) of years of development, then surely there ought to be millions that have achieved Type II status.

Incidentally, this is not an argument against the SETI project, which in my view should have the highest possible priority because the negative finding is no less significant than a positive result.

It is odd that we find the cosmos so silent. Where is everybody? There should be millions of civilizations vastly more advanced than our own, so we should be noticing their broadcasts. A sufficiently advanced civilization would not be likely to restrict its broadcasts to subtle signals on obscure frequencies. Why are they so silent, and so shy?

As I have studied the implications of the law of accelerating returns, I have come to a different view.

Because exponential growth is so explosive, it is the case that once a species develops computing technology, it is only a matter of a couple of centuries before the nonbiological form of their intelligence explodes. It permeates virtually all matter in their vicinity, and then inevitably expands outward close to the maximum speed that information can travel. Once the nonbiological intelligence emerging from that species' technology has saturated its vicinity (and the nature of this saturation is another complex issue, which I won't deal with in this essay), it has no other way to continue to evolve but to expand outwardly. The expansion does not start out at the maximum speed, but quickly achieves a speed within a vanishingly small delta from the maximum speed.

What is the maximum speed? We currently understand this to be the speed of light, but there are already tantalizing hints that this may not be an absolute limit. There were recent experiments that measured the flight time of photons at nearly twice the speed of light, a result of quantum uncertainty on their position. However, this result is actually not useful for this analysis, because it does not actually allow information to be communicated at faster than the speed of light, and we are fundamentally interested in communication speed.

Quantum disentanglement has been measured at many times the speed of light, but this is only communicating randomness (profound quantum randomness) at speeds far greater than the speed of light; again, this is not communication of information (but is of great interest for restoring encryption, after quantum computing destroys it). There is the potential

for worm holes (or folds of the Universe in dimensions beyond the three visible ones), but this is not really traveling at faster than the speed of light, it just means that the topology of the Universe is not the simple three dimensional space that naïve physics implies. But we already knew that. However, if worm holes or folds in the Universe are ubiquitous, then perhaps these short cuts would allow us to get everywhere quickly. Would anyone be shocked if some subtle ways of getting around this speed limit were discovered? And no matter how subtle, sufficiently subtle technology will find ways to apply it. The point is that if there are ways around this limit (or any other currently understood limit), then the extraordinary levels of intelligence that our human-machine civilization will achieve will find those ways and exploit them.

So for now, we can say that ultra high levels of intelligence will expand outward at the speed of light, but recognize that this may not be the actual limit of the speed of expansion, or even if the limit is the speed of light that this limit may not restrict reaching other locations quickly.

Consider that the time spans for biological evolution are measured in millions and billions of years, so if there are other civilizations out there, they would be spread out by huge spans of time. If there are a lot of them, as contemporary thinking implies, then it would be very unlikely that at least some of them would not be ahead of us. That at least is the SETI assumption. And if they are ahead of us, they likely would be ahead of us by huge spans of time. The likelihood that any civilization that is ahead of us is ahead of us by only a few decades is extremely small.

If the SETI assumption that there are many (e.g., millions) of technological (at least radio capable) civilizations is correct, then at least some of them (i.e., millions of them) would be way ahead of us. But it takes only a few centuries at most from the advent of computation for that civilization to expand outward at at least light speed. Given this, how can it be that we have not noticed them?

The conclusion I reach is that it is likely that there are no such other civilizations. In other words, we are in the lead. That's right, our humble civilization with its Dodge pick up trucks, fried chicken fast food, and ethnic cleansings (and computation!) is in the lead.

Now how can that be? Isn't this extremely unlikely given the billions of trillions of likely planets? Indeed it is very unlikely. But equally unlikely is the existence of our Universe with a set of laws of physics so exquisitely precisely what is needed for the evolution of life to be possible. But by the Anthropic principle, if the Universe didn't allow the evolution of life we wouldn't be here to notice it. Yet here we are. So by the same Anthropic principle, we're here in the lead in the Universe. Again, if we weren't here, we would not be noticing it.

Let's consider some arguments against this perspective.

Perhaps there are extremely advanced technological civilizations out there, but we are outside their light sphere of intelligence. That is, they haven't gotten here yet. Okay, in this case, SETI will still fail because we won't be able to see (or hear) them, at least not before we reach

Singularity.

Perhaps they are amongst us, but have decided to remain invisible to us. Incidentally, I have always considered the science fiction notion of large space ships with large squishy creatures similar to us to be very unlikely. Any civilization sophisticated enough to make the trip here would have long since passed the point of merging with their technology and would not need to send such physically bulky organisms and equipment. Such a civilization would not have any unmet material needs that require it to steal physical resources from us. They would be here for observation only, to gather knowledge, which is the only resource of value to such a civilization. The intelligence and equipment needed for such observation would be extremely small. In this case, SETI will still fail because if this civilization decided that it did not want us to notice it, then it would succeed in that desire. Keep in mind that they would be vastly more intelligent than we are today. Perhaps they will reveal themselves to us when we achieve the next level of our evolution, specifically merging our biological brains with our technology, which is to say, after the Singularity. Moreover, given that the SETI assumption implies that there are millions of such highly developed civilizations, it seems odd that all of them have made the same decision to stay out of our way.

Why Intelligence is More Powerful than Physics

As intelligence saturates the matter and energy available to it, it turns dumb matter into smart matter. Although smart matter still nominally follows the laws of physics, it is so exquisitely intelligent that it can harness the most subtle aspects of the laws to manipulate matter and energy to its will. So it would at least appear that intelligence is more powerful than physics.


Perhaps what I should say is that intelligence is more powerful than cosmology. That is, once matter evolves into smart matter (matter fully saturated with intelligence), it can manipulate matter and energy to do whatever it wants. This perspective has not been considered in discussions of future cosmology. It is assumed that intelligence is irrelevant to events and processes on a cosmological scale. Stars are born and die; galaxies go through their cycles of creation and destruction. The Universe itself was born in a big bang and will end with a crunch or a whimper, we're not yet sure which. But intelligence has little to do with it. Intelligence is just a bit of froth, an ebullition of little creatures darting in and out of inexorable universal forces. The mindless mechanism of the Universe is winding up or down to a distant future, and there's nothing intelligence can do about it.

That's the common wisdom, but I don't agree with it. Intelligence will be more powerful than these impersonal forces. Once a planet yields a technology creating species and that species creates computation (as has happened here on Earth), it is only a matter of a few centuries before its intelligence saturates the matter and energy in its vicinity, and it begins to expand outward at the speed of light or greater. It will then overcome gravity (through exquisite and vast technology) and other cosmological forces (or, to be fully accurate, will maneuver and control these forces) and create the Universe it wants. This is the goal of the Singularity.

What kind of Universe will that be? Well, just wait and see.

Plan to Stick Around

Most of you (again I'm using the plural form of the word) are likely to be around to see the Singularity. The expanding human life span is another one of those exponential trends. In the eighteenth century, we added a few days every year to human longevity; during the nineteenth century we added a couple of weeks each year; and now we're adding almost a half a year every year. With the revolutions in genomics, proteomics, rational drug design, therapeutic cloning of our own organs and tissues, and related developments in bio-information sciences, we will be adding more than a year every year within ten years. So take care of yourself the old fashioned way for just a little while longer, and you may actually get to experience the next fundamental paradigm shift in our destiny.



Copyright (C) Raymond Kurzweil 2001

Chart Graphics by Brett Rampata/Digital Organism