The file below was compiled by www.inovasyon.org on 12th May, 2006.

****

Michael Dertouzos[1]

The late Michael Dertouzos was Director of the MIT Laboratory for Computer Science (LCS) since 1974. For more than a quarter century, the Lab has been at the forefront of the computer revolution. Its members and alumni have been instrumental in the invention of such innovations as time-shared computers, RSA encryption, the Spreadsheet, the NuBus, the X-Window system, the ARPAnet and the Internet. The Lab is currently home to the World Wide Web Consortium, an open forum of companies and organizations led by the Web's inventor.

Dertouzos spent much of his career studying and forecasting future technological shifts, and leading his lab toward making them a reality. In a 1976 *People* magazine interview, he successfully predicted the emergence of a PC in every 3-4 homes by the mid-1990s. In 1980, he first wrote about the Information Marketplace, with an ambitious vision of networked computers that has emerged as the trillion-dollar engine of commerce transforming our economy.

Most recently, Dertouzos was an advocate for what he calls "human-centric computing"-a radical transformation of the way we use computers. As part of this effort, LCS recently unveiled the $50 million Oxygen project, intended to make computers easier to use and as natural a part of our environment as the air we breathe.

Born in Athens, Greece, Dertouzos came to the U.S. as a Fulbright Scholar. Following a Ph.D. from MIT in 1964, he joined the MIT faculty, where he is now Professor of Computer Science and Electrical Engineering.

In 1968 Dertouzos founded Computek Inc. to manufacture and market one of the earliest graphical display terminals, based on one of his patents. He soon became the Chairman of the Board of Computek, where he introduced the first intelligent terminals in the early 1970's. He subsequently sold the company when he became Director of LCS. Since that time, Dertouzos has been involved in several high-tech start-ups, including Picture Tel

[1] http://www.kurzweilai.net/bios/frame.html?main=bio0018.html; 09th May, 2006

and RSA. In his consulting activities for companies such as Siemens Nixdorf, UPS, and BASF he has advanced business and Information Technology strategies.

During the Carter Administration, Dertouzos chaired a White House advisory group that redesigned the White House Information Systems. In February of 1995, he represented the U.S. as a member of the U.S. delegation led by Vice President Al Gore to the G7 Conference on the Information Society. In 1998 he was co-chairman of the World Economic Forum on the Network Society in Davos, Switzerland.

Dertouzos was a dual citizen of the U.S. and the E.U. He worked extensively with the European Commission, in particular as a frequent keynote speaker on ESPRIT and other EC technology programs. For several years he was an adviser to the Prime Minister of Greece, as well as to other governments.

Dertouzos was a member of the United States National Academy of Engineering and the Athens Academy of Arts and Sciences. He held an honorary doctorate from the University of Athens, and he received the B.J. Thompson Award (best paper) of the Institute of Electrical and Electronics Engineers (IEEE) and the Terman Award (best educator) of the American Society for Engineering Education. He was a member of the U.S. Council on Foreign Relations, and was honored by the Hellenic Republic as Commander of Greece's Legion of Honor.

Dertouzos is the author/co-author of seven books, including MADE IN AMERICA: Regaining the Productive Edge (MIT Press, 1989), with over 300,000 copies in print, and WHAT WILL BE: How the New World of Information Will Change Our Lives (HarperCollins, 1997), which has been translated into thirteen languages.

****

Finishing the Unfinished Revolution[2]

by Michael L. Dertouzos

In this manifesto, Dr. Dertouzos introduces a radical vision of human-centered computing intended to make computers more usable, based on natural interaction (such as speech recognition), automation, individualized information access, collaboration, and customization. MIT's Oxygen project, a prototype to test these concepts, is summarized in this excerpt from The Unfinished Revolution (HarperCollins, 2001),

Originally published 2001. Excerpt from the book *The Unfinished Revolution*. Published on KurzweilAI.net February 22, 2001.

Imagine it's the year 2020 and the radical change we are after has happened. Systems like Oxygen have finally risen above the machine level and have been serving human needs. How far have they gone toward helping us do more by doing less? Did they help us get rid of the many difficulties that surrounded computers back in the year 2000? Did they increase our productivity and make our systems easier to use?

Back at the turn of the century, we had to read huge manuals to operate a word processor. Now, thanks to the natural interaction provided by human-centered systems,

[2] http://www.kurzweilai.net/meme/frame.html?main=/articles/art0128.html; 09th May, 2006

this "excessive learning fault" is largely gone. We talk to our systems and they understand enough to talk back and be useful. We still have to learn how to operate these machines, but the effort required on our part is much smaller.

In 2000, we typed and squinted a lot, doing all the electronic shoveling with our brains, eyeballs, and fingertips. Human-centric automation has freed us from this "manual labor fault," carrying out all sorts of tasks automatically. The "human servitude fault" is also largely behind us, since in the face of truly useful automation, service providers can no longer get away with those terrible automated phone operators that enslaved us through a maze of push-button choices.

The "overload fault," caused mostly by a dangerously expanding email habit, has also been brought under control. People have adopted human-centric attitudes-they no longer frantically send so much unsolicited e-mail, nor do they feel obligated to respond to every message they get. Most legislatures have passed laws obliging telemarketers to tag their messages with metadata that identifies the sender and the category of product or service being proffered, and filters used by essentially everyone let through only the ads that their masters wish to see.

Before human-centered systems we could barely find what we wanted through all the info-junk. Today, even though the info-junk has soared, we can find what we want with less work on our part, thanks to individualized information access and the ascent to meaning through the Semantic Web; the old "information access fault" has been largely circumvented. The "feature overload fault" is also out of the picture, because customization of our systems and applications to our individual needs have reduced the tendency of software developers to provide every conceivable feature in an attempt to please everybody. The old "crash fault" has been vanquished, too, because our human-centric software tracks the daily evolution of every program we touch, bringing us its most recent incarnation, and because when we run into trouble, the system takes us back to the most recent trouble-free state. Our machines do the backing up, not we. And no longer do we have to contend with the loss of time and peace of mind to port our software from one machine to another when we change machines. Nomadic software ensures that our info personality flows into whatever new piece of hardware we acquire or borrow, whenever and whenever this is necessary.

The "unintegrated systems fault" that made it impossible for me to use my calendar card during my plane ride to Taiwan is now a rare occurrence. The human-centric focus of technology has made the developers of operating systems and applications much more conscious of the need to serve people, and competition to supply consumers with this highly desired, higher level of operation has obliged them to do so.

Not all the computer faults have vanished. The "fake intelligence fault" continues to bother us, as software developers try to make systems more helpful by making them more "intelligent." And the collection of hundreds of automated procedures that we all have, while helpful in their individual tasks, conflict with one another at times. The "ratchet fault" where layers of old software pile up on top of one another is also present, because writing software continues to be more of a difficult arts-and-crafts proposition than a precise science, and we have not yet come up with any dramatic improvements to the software design process. We derive some comfort, however, from the fact that most of this ugliness resides well inside our systems, invisible to us.

In 2000, we all plied our trades and pursued our private info escapades with identical "personal" computers. Today, the machines adapt to our unique needs through customization. Back then, we could not easily reach people on the go, nor control our physical surroundings. Now we use the human-centered systems' ample reach to interact

with people in every place and time, and to control the devices and appliances we care about.

Human-centered systems also have made it possible for us to carry out new tasks. They help us work easily with one another across space and time, tracking our activities, helping us form secure collaborative regions, letting us annotate our conclusions, and generally helping us work much better than we could using only e-mail. Information work is now routine and occupies one-fourth of the world's economy, as people buy and sell human information skills across the world. Nearly 10 percent of that activity comes from India, which has doubled its GDP since 2000, mostly by selling clerical office work and software services. China accounts for 6 percent of total information work and Africa for 3 percent. The remaining 80 percent is within the industrial world.

Of the 1.5 billion people now using the Information Marketplace, some 300 million come from these three vast blocks of humanity--a feat made possible with a lot of good help from the people of the West, and partially from the progress in cellular Web access and in speech technology. Our goal of ensuring that many people become interconnected has been partially met, though we are by no means there. Compared with 2000, when fewer than 5 percent of the world's people were interconnected, the figure now approaches a respectable 20 percent, a quarter of which represents the developing world.

A principal objective of human-centric computing was to develop the gas pedal, steering wheel, and brakes of the Information Age. We have done so in the form of the five human-centric technologies, which became the applications interface of our new systems, were adopted by a new breed of applications, and have sent our productivity soaring. And by infusing these technologies into the Internet and Web, we have transformed these old media from their "voyeurism and exhibitionism" state into a full-fledged Information Marketplace.

Information technology has come well into our lives, and, as expected, we notice it less. Human-centered systems have liberated us from thinking about technology to thinking about what we really want to do. We can rejoice in the knowledge that our beastly computer menagerie of old has been almost fully domesticated!

But we are not quite where the Industrial Revolution was in 2000, because our information systems have not vanished as completely as the motors had back then. More work will be needed as human-centric information technologies continue to penetrate new areas of our personal and professional lives. When our information systems finally vanish in another decade or two, that will be the signal that the Information Revolution is done.

Let's suppose that these estimates for the year 2020 or so are correct, and the Information Revolution has been finished in the same sense that the Industrial Revolution is now over. Will we then be better off? Or will we have become efficiency freaks, bent on being productive every moment of our lives, in the process losing our peace of mind, our humanity, and our heart and soul? What will we do with all the time we save-work more? Will computers that operate at a more human level help us be more human? Or will our increased preoccupation with information drown us? Will increased collaboration across the planet lead to a uniform global culture? Will automation and superior information access make us lazy and excessively dependent on our machines? Or will the new capabilities encourage us to follow the high road? Will simpler systems reach beyond the fraction of the globe they now serve, to the billions of still unconnected and unengaged people? If so, will the systems help poor people become wealthier? Ultimately, how far might we go with human-centered computers toward enhancing our humanity?

Let's get some answers.

Info Royalty

We begin our search for the big picture with a small step: What might human-centric computing do to our rational, utilitarian selves? The answer is straightforward, especially in comparison with the Industrial Revolution. If you like what cars, airplanes, electricity, and chemicals have done for you, then you will like what the new information tools will give you.

You will be able to do more work, especially of the office variety, in less time. You won't be as frustrated, because your systems will be easier to use and more responsive to your needs. Your health will be improved through less expensive but faster, more accurate, and higher-quality medical systems. And you will have faster access to more of the world's products and services, tailored to your special desires. All the services you normally use, from getting an appliance fixed to finding the right lawyer or a comfortable future home, will be faster and better. You will have more options on receiving instruction, and even becoming educated. New entertainment will surround you, rich in content and interactions with other people. And you will have fun in new ways, as you play with it. You will also interact more easily and reliably with your family members, wherever they and you may be. Your thoughts and ideas will touch more people, and you will have the option to visit more of the thoughts and ideas of your fellow human beings. Organizations will function more efficiently, too, including governments, which will be able to better reach and interact with their constituencies.

These utilitarian benefits are qualitatively similar to the benefits we and our ancestors received from the plow and the motor. These earlier tools helped increase human productivity dramatically. Nowhere else is their combined effect more visible than in the generation of food, which went from absorbing all people in ancient times, to occupying a mere 2 percent of the industrial world population today--a whopping 5,000 percent productivity increase. These industrial innovations also helped us live better and have fun in new ways, with bright lights, automobiles, aircraft travel, consumer electronics, useful medicines, and so much more. As with human-centric technologies, these industrial developments made it easier for people to carry out their professional and personal lives. Just compare all the personal and professional things we can do today using the automobile with what people could do in earlier times using their feet and an occasional horse.

Of course, you might argue that cars did bad things for the family, the environment, and our soul, or that factory automation displaced jobs and led to the atrophy of our muscles. The same scenarios will be repeated in the Information Age.

How about leadership, responsibility, honesty, and those other human qualities we treasure? The answer is that you'll be able to use the new tools to either further or diminish these qualities. Any change will be up to you. As for the new ills that human-centric computing may bring--theft at a distance of our money, sexual advances toward our children, misinformation about us, cross-border crimes--the same answer applies: The new tools, like all technology, can and will be used for good and for evil. The angels and the devils are not in the machines, but in you and me. Since the ratio of angels to devils stems from human nature, this proportion is not likely to change. The balance between good and evil in the world won't be affected by the onset of human-centric systems.

Almost all the arguments you can fashion today about what the world of information will do to us were raised during the Industrial Revolution. So ask yourself if, considering everything you care about, you are better off with that socioeconomic movement behind you. Or would you be happier if it never happened? With almost no exception, the people of the industrial world have elected to live in it rather than in a cave, foraging to feed

their families. This suggests that despite protestations here and there, people overwhelmingly prefer the industrial to the preindustrial way of life.

I can already hear the dissonant chorus: "People can't help it." "They think they are better off but they aren't." "This is a utilitarian society that has lost its compass heading. No wonder they like it. They have lost their sense of direction." I'll address these deeper questions about technology's ultimate impact upon humanity in a moment. Meanwhile, it is safe to conclude that from a utilitarian perspective, we will be better off with our new information tools, for the same reasons that we continue to be satisfied with the greater utility made possible by the industrial advances of the previous two centuries.

Does all this mean that human-centered computers will simply continue the same sorts of benefits? Not quite. The gains will be sufficiently different to induce a qualitatively new social change--something akin to a new social order. In a strange way, we'll be able to do many of the things that were the province of wealthy people, past and present. Kings and rich folk have always had servants that catered to their every wish. With human-centered computers, we, too, will end up surrounded by many automated servants--scripts and specialized procedures ready to cater to our needs. Rich people have always had better access than the rest of us to the information they need, because they have the right connections and can afford the expense of finding and obtaining what they need. So will we with the human-centric force of individualized information access. Rich people have always had products and services customized to their desires. So will we through customization. Rich people don't need to work, because their wealth breeds more wealth. This won't happen to us completely, or overnight, but the expected threefold increase in human productivity, made possible by human-centric systems, could free up two-thirds of the time we now spend working. . . if we elect to realize the savings in this way. The collective benefits of human-centered machines will give us enough of the capabilities now reserved for the rich to make us feel like royalty.

Just as the Industrial Revolution produced a new middle class, the Information Revolution, through its human-centric technologies, will create a new "info royalty" class. Who knows? A few decades from now, human-centered machines may return human beings to the princely benefits of earlier feudal times, when the rich had servants, and the master reigned supreme . . . except that almost everyone will have a chance of being the master!

Will we then be better off? That will be up to us. The history of kings and princes shows that they have gone in every conceivable direction during their spare time. If we follow in the Information Era what we did in the Industrial Era, we'll work harder with the time saved by our new royal status, so that we may acquire even greater wealth. On the other hand, we may elect to devote the time we save to other endeavors that please or uplift us. We'll have the luxury of choosing our course.

Such a societal shift would be more profound than an incremental utilitarian improvement in human productivity and ease of use. Doing more with less effort would then have the added meaning that we would be able to act more like kings than serfs.

Global Reach

Who would have believed 15 years ago that poor programmers in Bangalore, India, would sell their software services to the West, putting together companies like Infosys (which in July 2000 was valued at nearly $25 billion), which collectively employed 60,000 programmers, whose standard of living is now pulling their region's economy upward at 25 percent per year?

That ray of sunshine is particularly important for the hope it brings to the developing world. For, if the new royalty class is limited to the people now interconnected via the Web, humanity won't be doing more by doing less. The new "royalty" would stand for a tiny fraction of the world population. And that would be just as bad as the real royalty of old, reverting us to an era of a privileged few, likely to be followed by bloody revolutions, as was feudalism. This is why I insist that a primary imperative of finishing the Information Revolution is that the new technologies of information reach as many people as possible.

Fortunately, there are many ways to improve the global reach of information technology. Communications could be provided by low-earth-orbiting satellites operated by such companies as McCaw Communications and Globalstar that whip around the earth. When these birds are over the industrial nations they are very busy, but when they are over the developing world they are doing nothing. Let's pay the low marginal cost to leave them on. In addition, hardware and software makers, training outfits, and communication service providers could offer their wares to the poor at deep discounts. We citizens could help cover the cost by instructing our governments to offer attractive tax breaks to these suppliers. Individuals could also donate money or time. Organizations like the World Bank, which spends over $30 billion annually in structural loans to the developing world, could put a good part of these funds into worthy information technology projects.

Armed with the excitement of these prospects, a few of us techies got together with a colleague from Nepal, fully expecting to boost his nation's economy by 20 percent through clever use of the Internet. Unfortunately, we quickly found that even if we got him the communications, hardware, software, and training for free, we would still fall short of our goal. That's because only 27 percent of the Nepalese are literate, and of those, only a small fraction can handle English. When we asked what services that smaller group could offer we hit a brick wall. Many are not skilled, and those who are, are busily running their nation's businesses. Maybe we were too ambitious when we envisioned a future workforce in Nepal selling office services to New York and London via the Web. The potential of the Information Age seemed overshadowed at every turn by the ancient forces that separate the rich from the poor.

Like others who have tried to do something in this area, we, too, came to the realization that the lack of communications, computers, and training is not the primary problem. The bigger obstacles are the same that have kept the poor from rising above poverty throughout history. Lack of education is at the helm. It is followed by lack of transportation, power, and telecommunications; absence of capital; misuse of whatever resources may be available; government inertia; and cultural taboos. Moreover, basic concerns over food, shelter, and health dominate poor people's plans and actions, as they should, ahead of the less tangible promises of information technology.

These observations and concerns were amplified by an MIT Laboratory for Computer Science survey about the uses of information technology in the developing world in 1999. The results showed that the biggest recent successes in developing countries, disguised under all sorts of information technology experiments, actually involved the introduction and use of POTS--plain old telephone service. And in cases where new information technologies beyond telephony seemed to be statistically active, we found that they were used mainly by the few relatively rich people among the poor--a faithful microcosm of what is happening globally, and hardly a model for addressing the larger problem. We have not yet latched on to an approach that can productively engage the poor in the global Information Marketplace.

If the world has to hold out until developing nations, and the poor in the industrial world's inner cities, fix in serial fashion the social, political, and economic problems that plague them, we will be in for a very long wait. What we must do instead is help through

donations, government aid, personal and corporate contributions, tax credits, loans, and all the mechanisms we can muster to improve education and infrastructure. Most important, we must explore creative "shortcuts" that have a chance of working. One possibility is to strengthen entrepreneurial initiatives among the poor through incubator programs that provide capital and other resources. Successes from within a community, as in the case of Bangalore, will stimulate duplication far more effectively than solutions from outside. Another shortcut may be the launching of short-term training and education programs aimed at preparing people directly for selling information work. Yet another shortcut involves the use of speech understanding technology to bypass illiteracy for people who, despite their inability to read and write, can contribute and benefit from the Information Marketplace.

A new world of human-centric computing must work for all humans. If the bulk of our planet's people are not interconnected, then humankind will not be able to do more by doing less. Only a few will have that privilege.

Monoculture and Overload

As much as we hope that human-centered computers may help level economic disparity across the world, the process will take time. In the shorter term, it is natural for us to wonder whether the technology might level cultural differences among the people who are interconnected. Collaboration, in the form of commerce, information work, entertainment, and education, plus individualized information access, open to the entire world the personal attitudes, customs, history, art, good and bad habits, and traits of peoples that are normally confined to citizens of single nations. Speech understanding will lead to translated exchanges that cross linguistic barriers. And automated, semantic exchanges among machines will spread shared concepts. Might these leveling forces push us toward one homogenized world culture?

When non-Americans ask this question, their dominant fear is that the answer will be "yes" and the resultant monoculture will be American. Nonsense! Tribalism is a far more powerful human force than any computing trend. Consider, for example, that although the member nations of the European Union have all been using English for a long time, it has barely affected the differences among their tribes. The Italians still differ from the British, who differ from the French, who differ from the Greeks, more or less as they have differed for centuries. What has happened among the people who participate in this sharing is the adoption of a shallow cultural layer that involves common sound bites and a few shared habits. That's exactly what I expect will happen as human-centric computing crosses national boundaries--a thin veneer of shared norms, not a monoculture.

A related fear is that the cross-border interactions will cause nations to vanish. Either their citizens will be globally distributed and won't care about national boundaries, or the ease with which the new technologies cross these boundaries will make national distinctions unnecessary. More nonsense! The police forces and armies of different nations are physically local and will remain so. They, along with their political leaders and their population, are dedicated to national survival with the same fervor that human beings are committed to personal survival. The likelihood of a military force, driven by a national political leadership, yielding its swords and bombs to some shared bits of information is pretty close to nil.

Still, the new technologies, by increasing communication, will foster a better understanding among tribes. A Greek and a Turk who love early music will join that musical "tribe" on the Net, and will get to know each other across the divide of their ancient national tribes. This could bode well for peace, since the more that people talk to one another, especially in casual settings, the less likely they are to kill their discussion

partners. At the same time, these technologies will also strengthen ethnic tribes by uniting local with distributed members. For example, the 7 million Greeks living in the United States, Australia, and elsewhere outside Greece could become better tied culturally, economically, and socially with the 10 million Greeks living in the country of Hellas. Human-centric computing has the schizophrenic ability to simultaneously strengthen diversity and tribalism.

I believe that this is a great thing for our world, where these opposing forces are basic to human nature and are becoming increasingly widespread in the cities and countries where people live. The simultaneous strengthening of tribalism and diversity is yet another interpretation of how doing more by doing less might affect our world.

Another common fear is that the new technologies will overwhelm us with information, rendering us ineffective. As fashionable as this fear is, don't worry about it. Since ancient times people have valued their own survival over all else. In a serious conflict between a debilitating amount of information and survival, there is no question as to what people will do: They'll trash the information without a moment of bad conscience. . . as they should!

The Technology Fountain

As we ask the basic questions of how far we might go with human-centered computers, and how much better off we may be, we should keep in mind that technology will not stand still, and will most likely create new avenues through new discoveries. Our future vision is necessarily limited, but from what we can see, two categories of potential developments stand out--machine learning and the merger of biology with computer science. Here's why.

If computer systems become capable of learning from practice and observation of their environment, rather than by being programmed by people, we are in for a very big change. Technically, this is not part of the human-centric tool kit we have been discussing. It will require new discoveries, and as I have repeatedly said, there is no basis to predict that it will or will not happen. But if it were to succeed, we would finally have achieved great progress toward the construction of intelligent systems. Each of us would have intelligent programs and knowledgeable advisers at our side. That would bring computers even closer to serving human needs, and would result in the ultimate human-centered systems, with dramatic consequences for all of us.

Some people believe that machine learning is a dated idea and computer intelligence will evolve just as human intelligence did. They argue that since computer processing power is accelerating so much more rapidly than the human brain's, it will only take a few decades before a computer's intelligence surpasses a human's. At this point, they conclude, a machine will no longer need a human to create its offspring, and we will become irrelevant. It's fun to raise such ideas for the mental stimulation they provide. But pretending that something like this is likely to happen is quackery. What does accelerating computer power have to do with intelligence? If you move your arms faster, do you get smarter? Of course not. The growing processing power of computers says nothing about how intelligent our machines may or may not become. As for machine learning being a dated idea, discoveries are not subject to fashion like clothes! A breakthrough in machine learning, if it were to happen, would instantly become a "modem" achievement.

The evolution of machine intelligence, to where machines can beget other machines, is a metaphor that shocks and seduces, because it ascribes to future machines capabilities that people believe are uniquely human. That's even further away from our understanding than machine intelligence! People should feel free to delight in such

musings. But they should not seriously worry about them any more or any less than they worry about our planet being struck by a gigantic asteroid.

The second big development that may lie ahead--a merger of biology and computer science--has nothing to do with the human-centric technologies we have been discussing. But if it were to happen, it would affect dramatically the way machines would serve us, especially for our health needs. This marriage seems plausible because biological organisms, including humans, can be characterized by their DNA structure--in other words, by information. Even though the amount of data needed to describe the molecular makeup of a single person is huge, it is still information. With the massive research effort known as the Human Genome Project as a base, scientists are increasingly able to describe in a digital blueprint the biological aspects of a person. In the imagined scenarios, this information could be used by our doctors and by us to forecast illnesses, presage hereditary strengths and weaknesses, fix or alter our human traits, and, in the extreme, to design a young fetus to have the characteristics we want it to have. In the other direction, too, biological techniques and materials could be used to fashion "computing machines" of a very different kind.

Developments like these could change the role of information in our lives, and would no doubt bring surprises. My own belief based on no facts whatsoever is that machine learning has a chance of succeeding at a partial level sometime this century. The more exotic possibility of a bio-computational merger toward the "boutique" design of living beings is too far in the future to be visible.

When we think of such possibilities, it is natural that we become frightened, to the point of asking for a moratorium on discovery, as some people have suggested, fearful of harming ourselves irreversibly with the unintended consequences of genetic engineering and machine intelligence. "Shut down the technology fountain," they say. I do not subscribe to this view, because the consequences of our discoveries are unpredictable and we are unable to chart a careful course through a universe we barely comprehend.

When we built time-shared computers and the Arpanet, we did it so we could avoid buying expensive machines, by sharing them. The efforts succeeded, not for these reasons, but because they helped people share information. The Internet was launched to interconnect networks of computers; no one expected that its biggest application would be the Web. Radar was designed for war, but ended up as a cornerstone of air transportation. Nuclear weapons research put nuclear medicine on the map. Thousands of innovations all share the same pattern--the early assessment is unrelated to the outcome. So limited is our ability to assess consequences that it's not even helped by hindsight. We can't judge whether cars, synthetic drugs, and nuclear power, all invented more than 50 years ago, are on balance good or bad for us today. Our track record of rationally assessing the future uses of science and technology is pretty lousy. How then are we going to tell what kind of research we should stop and when?

Maybe we should stop research altogether. This reminds me of a wise old airline employee. I had announced to him that I stopped flying with his company because of its poor safety record. "Listen sir," he said. "If your exit visa from this life is stamped 'death by aircraft,' even if you stay in your bed, the airplane will find you and crash upon you." At this, the dawn of the Technology Century, it is not fashionable to pay attention to forces and beliefs, like destiny, that lie outside current reason. We should reconsider. All the more so if we are arrogant enough to believe we understand our universe enough to successfully regulate its future course.

We should also remember that what we do as human beings is part of nature. I am not advocating that we do as we please, on the grounds that everything we do is natural, but

rather that we respect the natural human urge to probe and understand all that surrounds us.

I suggest that as we encourage the technology fountain to feed tomorrow's discoveries and their human uses, we stay vigilant, ready to stop when danger is imminent, not when our fears or premature rational assessments, which have failed us so often, scare us into doing so. And let's ponder what other help we might seek in reaching our decisions, especially since we are not the only determinants of change out there.

As we contemplate potentially earthshaking discoveries in the context of human-centered systems, let us remember that the primary role of information in our lives is to help us achieve our human goals. Information is, therefore, a means to getting there, rather than an end in itself. That is so powerful and fundamental a property of information that together with the unchanging nature of human purpose and human beings, it is likely to survive even the wildest of tomorrow's discoveries.

No Machines Beyond This Point

To fully understand the ultimate potential of human-centered computers, we should explore the limits of their uses. Is it possible that applying our new tools to certain tasks would result in our actually achieving less?

Yes. The tasks are the ones in which we convey to one another the primitive human emotions--primal forces that have been with us for thousands of years. These "forces of the cave," as I call them, range from fearing predators, seeking food and shelter, and nurturing our children to protecting our mate and trusting fellow tribe members.

By now, people who work as a team over the Internet have discovered that as long as they know and trust each other, the team functions well in its virtual forays. But when new team members join, the group loses its effectiveness. The team returns to progress only after the new members have bonded with the old ones in old-fashioned ways--by squeezing each other's hand, drinking beer together, exchanging personal stories, or giving one another a slap on the back. Building trust seems to be outside the limit of what we can do "at a distance," regardless of how faithfully the technology bridges space and time. The troubleshooter teams at British Petroleum, who use collaboration technologies to solve problems at remote oil well sites, have found this phenomenon to be true. So has MIT; as we began planning our own collaborative, distance education programs, we quickly agreed that our remote students would need to spend nearly as much time on the MIT campus as they did away from it, to partake of these deep forces that do not travel over the links of the Information Marketplace.

Why don't they? Well, imagine that your 14-year-old son has done something reprehensible. You grab him by the collar, squeeze his neck a bit, look him in the eye, and say, "Johnny, don't ever do that again." You then release your grip and explain why you were so menacing in your admonition. You could not have the same effect if Johnny were 3,000 miles away, even with the best collaboration technology that perfectly re-created your appearance, voice, and squeeze. Why? Because in the physical encounter, your son experiences a primitive fear As you grab him, his instincts tell him that the situation could progress toward greater physical danger Never mind if you have never struck him before. The primal forces of the cave, rather than reason, are at work. As far as these forces are concerned, there is no telling what you might do. But in the virtual scolding, your son knows, even if he is "scared" by your demeanor, that he can flip a switch and turn off the whole thing! The encounter is just a simulation. You are not transmitting primal fear to Johnny, only an image of that fear, which is no longer a primitive force.

The forces of the cave are with us all the time, regardless of the rational powers and sophisticated behaviors we invent to disguise them. And they cannot be easily tricked. Doctors healing patients, parents raising children, business associates building trust, lovers exchanging intimacy, friends accepting each other, enemies trading threats--all use the forces of the cave. Even though the information component of these activities could be communicated well with human-centered machines, the exchanges would be nowhere as effective.

The forces of the cave set a clear limit as to how far human-centric computing can go toward helping us do more by doing less. Even when we finish the Unfinished Revolution, they will still hold sway.

Greater Humanity?

We want to go beyond the efficiency, ease of use, fun, and productivity implications and explore whether the human-centric technologies can "do more" to enhance our humanity, to truly make us "better off." To ponder this lofty question, we must declare what we consider being "human" signifies. Each of us assigns a highly individual interpretation to this term, since it defines the meaning and purpose of our unique lives.

We can't get a universal definition, but we can characterize several of the dimensions that constitute what humanity might mean to different people. Then we can assess how the new technologies may or may not help us along each dimension. By selecting which of the dimensions you deem important, you can get an idea as to how human-centric computing might affect your own sense of being human.

During the Enlightenment, people decided to separate reason from faith and from the literature of the ancients. This freed science and technology from the shackles of religion and humanism. It fueled the Industrial Revolution and later the Information Revolution. The success of industrialization confirmed the wisdom of separating these dimensions of humanity, and reinforced the three-way separation among technologists, who put their faith in reason; humanists, with their focus on the arts, literature, and human feelings; and believers centered on spirituality.

Here, then, are three historically vetted dimensions of what it might mean to be human: the *reason* part that stands behind science, technology and rational thinking; the *feeling* part that lifts the arts and the humanities; and the *faith* part that helps us cope with what cannot be explained or felt. Add our physical *action* and we cover a good deal of what it means to be human. Where do you fall along these dimensions? Which do you consider more important in your own life? As you formulate your answer, let's take a look at whether, or how, the human-centric technologies might affect each dimension.

The rational part of being human will benefit greatly, because it is the stuff of which the technologies are made. We have seen many ways to enhance our reason through greater access to information, better communication, customization, and much more.

Automation can amplify the action part of our humanity by bringing the physical world under our greater control, and by harnessing our machines to act in our stead. Planning, crucial to future action, is also dramatically enhanced by having access to good information and being able to process and share it effectively.

How about the feeling dimension? We have just established that the primitive forces are outside the reach of the new technologies. However, that doesn't mean that emotions can't be conveyed by the virtual world. We all laugh and cry at a good story or movie that reaches us over the Net, so certainly lighter-than-primal emotions can be sustained.

We can intensify sensory perception, too, by brightening colors and sounds, and perceiving sensations across greater distance. We can read more, access a great deal of the world's art, and use aids that help us when we create a poem or a picture. But we can't emote more deeply through the new technologies. The audience of a large-screen, 3-D, multimedia video packed with visual and sound effects cannot be made to feel more sincere empathy with the victims of a plane crash than you do when you read a good, plain-text newspaper article about the tragedy. The new technologies can amplify the feeling part of our humanity in a quantitative and somewhat perfunctory sense, but they cannot make us feel more deeply.

That brings us to faith. It is hard to imagine how a person's spirituality could be enhanced by technologies that deal with information. After all, faith, in those that have it, is essentially defined as something internal to our being and outside the realm of human reason, feeling, or action. The new human-centric technologies cannot amplify the spiritual dimension of our humanity.

If you are a hard-core technologist who believes that rationality is the essence of being human, or if you are a driven person who believes in action, then human-centric computing will greatly enhance your humanity. If you are an artist who thrives on feelings and new ways of expressing the world, you will find partial enhancement from the new technologies. If you are a monk whose life revolves around spirituality, you will look elsewhere for help. But if you possess varying amounts of these human dimensions-which describes most of us--then you can determine how much "better off' your humanity will be by analyzing how each of the dimensions you care about will be affected.

Beyond the Information Revolution

To my thinking, the ultimate way in which we can do more by doing less goes beyond the Information Revolution, but is made all the more urgent by its growing dominance. It involves the way we reconcile the human dimensions within us.

The millennium that just ended was dominated by God and faith, reflected in religious wars from the Crusades to the ongoing Middle East crisis, the split from Orthodoxy, the Reformation, and centuries of music and art that stemmed overwhelmingly from religion. Now, as the new millennium begins, this dominance is shifting toward a new "god"--technology--which began its powerful ascent toward the end of the 20th century. People stand awestruck by the miracles of information technology, biotechnology, medicine, and materials science, which promise to transform our behavior, our being, and our surroundings. They increasingly place their faith in this new god to address their human needs for better health, protection from danger, explanations of our surrounding world, and greater happiness. Since technology, and especially information technology, left unchecked, will further enhance reason at the expense of feeling and faith, aggravating the separation among these three pieces of humanity.

That separation grew as the Industrial Revolution became increasingly successful, and let to problems. Technologists began questioning their purpose. Humanists became disaffected with gadgets and materialistic ideas. The spiritually inclined resented the loss of beliefs. Youth, sensing that something was missing inside them, turned to apathy and drugs. People focused increasingly on themselves, celebrating possessions, lamenting depressions, and fragmenting families. Governments separated faith from reason in the school curricula. A politically correct population become increasingly reluctant to say "God." Universities isolated technologists from humanists in watertight compartments across campus from each other. Today the separation has become so ingrained we don't even see it or the problems in has engendered. We simply accept it as "natural."

If we allow this trend to continue, our problems will increase and we will miss the prospect of being better off in the biggest possible sense of being human. We simply can't go far if we stay fragmented. Take humanism; until recently, the essence of being well educated was, in the words of the English poet Matthew Arnold, "to know the best that has been thought and said in the world." If you needed technology, you bought it, like potatoes, to serve your loftier humanistic goals. That's how technologists became known as practitioners of "the servile arts." This humanist-dominant view made sense when technology was a small part of our lives--a notion that is no longer valid! Today, higher purpose may *originate* with technology, as in the invention of the Web by a full-fledged technologist. Many sites with a purely social purpose, developed by technologists, are already in operation. No pure humanist could ever have come up with these ideas, without also understanding technology. It's time for Matthew Arnold's words to be qualified. Technology will be as important a contributor to noble endeavors and understanding our world as humanistic ideals were and will continue to be. Keeping the technologists separated from the humanists will keep us from discovering these new territories.

People also have an inherent need for spirituality, which offsets the powerlessness we feel before the many mysteries that surround us. In an increasingly rational world, how might our children fulfill this human need, which has led billions to religion throughout the centuries? Never mind grandstanding on the industrial world's easy answer that church and state should stay separated, and the latter shouldn't glorify any particular sect in the schools. Good. Let's keep doing that. But then what? Will learning in the next millennium stay chained to reading, arithmetic, and reason? What of birth, friendship, love, marriage, illness, divorce, conflict, death, purpose?

If we remain fragmented, we'll be unable to fulfill our full human potential, because we will be running on only some of our cylinders. People lived for thousands of years without this internal separation. And we were not always as impressed with reason, morality, and all that we have built on the shaky foundation of human thought as we have been in the last few centuries. It is ironic, yet inescapable, that so many "thinkers," especially Western philosophers, stayed chained to reason and built their theories upon it, as if it were the only solid ground. Granted, we can't help but be impressed by this unique capability of our brain, which in its exquisite architecture and processes holds our awesome power to think. Yet, viewed from afar, it is just another property of a few ounces of meat tucked inside the skulls of antlike creatures that roam a huge earth in an infinite universe. What does reason have to do with the love of a child, the beauty of a flower, the eternity of stone, our origin, our destination? The new century of technology is amplifying our tendency to overrate reason at the expense of spirituality, and technological reason at the expense of humanistic ideas.

Do you find such philosophical considerations too abstract? Do you prefer to stay practical? Then here's something for you: How do we cope with children who use guns to kill their classmates? What do we do when genetic engineering can alter the personality of a fetus? How do we deal with trans-border crimes over the Internet? And how about all the other "ordinary" problems we will face that won't be as famous as these, but just as hard? Every decision we make, whether it's choosing a school for our children, managing people, cementing or breaking relationships, facing illness, running a household or a company or a country, will increasingly involve issues and considerations that are intertwined across these artificial divisions. Pure technology can't solve these problems. Nor can pure humanism or pure faith. We need to bring these back together if we want to find our way through the maze of an increasingly complex world.

This is especially true as we begin our journey to finish the Unfinished Revolution. The human-centric technologies will bring computers closer to us and give us power to do more by doing less. But the highest meaning of "human-centric," and its biggest benefit to us will be determined by what we do to achieve the human goals we set. We will be

better off and we will be finishing the ultimate Unfinished Revolution if we reach for these goals using all our human dimensions in concert, standing once again in awe before the sunset, the wheel and what may lie behind them.

Kurzweil vs. Dertouzos[3]

by
Ray Kurzweil
Michael L. Dertouzos

In this Technology Review article, Raymond Kurzweil and Michael Dertouzos debate Bill Joy's Wired article urging "relinquishment" of research in certain risky areas of nanotechnology, genetics, and robotics.

Originally published January 1, 2001 in *High Technology*. Published on KurzweilAI.net March 7, 2001.

Raymond Kurzweil

Although I agree with Michael Dertouzos' conclusion in rejecting Bill Joy's prescription to relinquish "our pursuit of certain kinds of knowledge," I come to this view through a very different route. Although I am often paired with Bill Joy as the technology optimist versus Bill's pessimism, I do share his concerns about the dangers of self-replicating technologies. Michael is being shortsighted in his skepticism.

Michael writes that "just because chips...are getting faster doesn't mean they'll get smarter, let alone lead to self-replication." First of all, machines are already "getting smarter." As just one of many contemporary examples, I've recently held conversations with a person who speaks only German by translating my English speech in real time into human-sounding German speech (by combining speech recognition, language translation and speech synthesis) and similarly converting their spoken German replies into English speech. Although not perfect, this capability was not feasible at all just a few years ago. The intelligence of our technology does not need to be at human levels to be dangerous. Second, the implication that self-replication is harder than intelligence is not accurate. Software viruses, although not very intelligent, are self-replicating as well as being potentially destructive. Bioengineered biological viruses are not far behind. As for nanotechnology-based self-replication, that's further out, but the consensus in that community is this will be feasible in the 2020s, if not sooner.

Many 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. When people think of a future period, they intuitively assume that the current rate of progress will continue for the period being considered. However, careful consideration

[3] http://www.kurzweilai.net/meme/frame.html?main=/articles/art0136.html; 09th May, 2006

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. 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 any serious consideration of the history of technology shows that technological change is at least exponential, not linear. There are a great many examples of this, including exponential trends in computation, communication, brain scanning, miniaturization and multiple aspects of biotechnology. One can examine this data in many different ways, on many different time scales and for a wide variety of different phenomena, and we find (at least) double exponential growth, a phenomenon I call the "law of accelerating returns." The law of accelerating returns does not rely on an assumption of the continuation of Moore's law, 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. 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 exponential growth is ignored).

This observation also applies to paradigm shift rates, which are currently doubling (approximately) every decade. So the technological progress in the 21st century will be equivalent to what would require (in the linear view) on the order of 20,000 years.

Michael's argument that we cannot always anticipate the effects of a particular technology is irrelevant here. These exponential trends in computation and communication technologies are greatly empowering the individual. Of course, that's good news in many ways. These trends are behind the pervasive trend we see toward democratization, and are reshaping power relations at all levels of society. But these technologies are also empowering and amplifying our destructive impulses. It's not necessary to anticipate all of the ultimate uses of a technology to see that there is danger in, for example, every college biotechnology lab having the ability to create self-replicating biological pathogens.

However, I do reject Joy's call for relinquishment of broad areas of technology (such as nanotechnology) despite my not sharing Michael's skepticism on the feasibility of these technologies. Technology has always been a double-edged sword. We don't need to look any further than today's technology to see this. If we imagine describing the dangers that exist today (enough nuclear explosive power to destroy all mammalian life, just for starters) 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 2001 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?

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.

The continued opportunity to alleviate human distress is one important motivation for continuing technological advancement. Also compelling are the already apparent economic gains, which will continue to hasten in the decades ahead. There is an insistent

economic imperative to continue technological progress: 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. Forgoing fields such as nanotechnology is untenable. Nanotechnology is simply the inevitable end result of a persistent trend toward miniaturization that 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.

Furthermore, 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 (for example, the terrorists) who would have all the expertise.

The constructive response to these dangers is not a simple one: It combines professional ethical guidelines (which already exist in biotechnology and are currently being drafted by nanotechnologists), oversight by regulatory bodies and the development of technology-specific "immune" responses, as well as computer-assisted surveillance by law enforcement organizations. As we go forward, balancing our cherished rights of privacy with our need to be protected from the malicious use of powerful 21st-century technologies will be one of many profound challenges.

Technology will remain a double-edged sword, and the story of the 21st 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.

Michael Dertouzos

In my column, I observed that we have been incapable of judging where technologies are headed, hence we should not relinquish a new technology, based strictly on reason. Ray agrees with my conclusion, but for a different reason: He sees technology growing exponentially, thereby offering us the opportunity to alleviate human distress and hasten future economic gains. From his perspective, my point is "irrelevant," and my views on the future of technology are "skeptical." Let's punch through to the underlying issues, which are vital, for they point at a fundamental and all-too-often ignored relationship between technology and humanity.

Ray's exponential-growth argument is half the story: No doubt, the number of transistors on a chip has grown and will continue to grow for a while. But transistors and the systems made with them are used by people. And that's where exponential change stops! Has word-processing software, running on millions of transistors, empowered humans to contribute better writings than Socrates, Descartes or Lao Tzu?

Technologies have undergone dramatic change in the last few centuries. But people's basic needs for food, shelter, nurturing, procreation and survival have not changed in thousands of years. Nor has the rapid growth of technology altered love, hate, spirituality or the building and destruction of human relationships. Granted, when we are in the frying pan, surrounded by the sizzling oil of rapidly changing technologies, we feel that everything around us is accelerating. But, from the longer range perspective of human history and evolution, change is far more gradual. The novelty of our modern tools is counterbalanced by the constancy of our ancient needs.

As a result, technological growth, regardless of its magnitude, does not automatically empower us. It does so only when it matches our ability to use it for human purposes. And that doesn't happen as often as we'd like. Just think of the growing millions of AIDS cases in Africa, beyond our control. Or, in the industrial world, ask yourself whether we are truly better off surrounded by hordes of complex digital devices that force us to serve them rather than the other way around.

Our humanity meets technology in other ways, too: In forecasting the future of technology, Ray laments that most people use "linear thinking" that builds on existing patterns, thereby missing the big "nonlinear" ideas that are the true drivers of change. Once again, this is only half the story: In the last three decades, as I witnessed the new ideas and the 50-some startups that arose from the MIT Laboratory for Computer Science, I observed a pattern: Every successful technological innovation is the result of two simultaneous forces-a controlled insanity needed to break away from the stranglehold of current reason and ideas, and a disciplined assessment of potential human utility, to filter out the truly absurd. Focusing only on the wild part is not enough: Without a check, it often leads to exhibitionistic thinking, calculated to shock. Wild ideas can be great. But I draw a hard line when such ideas are paraded in front of a lay population as inevitable, or even likely.

That is the case with much of the futurology in today's media, because of the high value we all place on entertainment. With all the talk about intelligent agents, most people think they can go buy them in the corner drugstore. Ray, too, brings up his experience with speech translation to demonstrate computer intelligence. The Lab for Computer Science is delightfully full of Victor Zue's celebrated systems that can understand spoken English, Spanish and Mandarin, as long as the context is restricted, for example to let you ask about the weather, or to book an airline flight. Does that make them intelligent? No. Conventionally, "intelligence" is centered on our ability to reason, even imperfectly, using common sense. If we dub as intelligent, often for marketing or wishful-thinking purposes, every technological advance that mimics a tiny corner of human behavior, we will be distorting our language and exaggerating the virtues of our technology. We have no basis today to assert that machine intelligence will or will not be achieved. Stating that it will go one way or the other is to assert a belief, which is fine, as long as we say so. Does this mean that machine intelligence will never be achieved? Certainly not. Does it mean that it will be achieved? Certainly not. All it means is that we don't know-an exciting proposition that motivates us to go find out.

Attention-seizing, outlandish ideas are easy and fun to concoct. Far more difficult is to pick future directions that are likely. My preferred way for doing this, which has served me well, though not flawlessly, for the last 30 years, is this: Put in a salad bowl the wildest, most forward-thinking technological ideas that you can imagine. (This is the craziness part.) Then add your best sense of what will be useful to people. (That's the rational part.) Start mixing the salad. If you are lucky, something will pop up that begins to qualify on both counts. Grab it and run with it, since the best way to forecast the future is to build it. This forecasting approach combines "nonlinear" ideas with the "linear" notion of human utility, and with a hopeful dab of serendipity.

Ray observes that technology is a double-edged sword. I agree, but I prefer to think of it as an axe that can be used to build a house or chop the head off an adversary, depending on intentions. The good news is that since the angels and the devils are inside us, rather than within the axe, the ratio of good to evil uses of a technology is the same as the ratio of good to evil people who use that technology...which stays pretty constant through the ages. Technological progress will not automatically cause us to be engulfed by evil, as some people fear.

But for the same reason, potentially harmful uses of technology will always be near us, and we will need to deal with them. I agree with Ray's suggestions that we do so via ethical guidelines, regulatory overviews, immune response and computer-assisted surveillance. These, however, are partial remedies, rooted in reason, which has repeatedly let us down in assessing future technological directions. We need to go further.

As human beings, we have a rational, logical dimension, but also a physical, an emotional and a spiritual one. We are not fully human unless we exercise all of these capabilities in concert, as we have done throughout the millennia. To rely entirely on reason is to ascribe omniscience to a few ounces of meat, tucked inside the skull bones of antlike creatures roaming a small corner of an infinite universe--hardly a rational proposition! To live in this increasingly complex, awesome and marvelous world that surrounds us, which we barely understand, we need to marshal everything we've got that makes us human.

This brings us back to the point of my column, which is also the main theme of this discussion: When we marvel at the exponential growth of an emerging technology, we must keep in mind the constancy of the human beings who will use it. When we forecast a likely future direction, we need to balance the excitement of imaginative "nonlinear" ideas with their potential human utility. And when we are trying to cope with the potential harm of a new technology, we should use all our human capabilities to form our judgment.

To render technology useful, we must blend it with humanity. This process will serve us best if, alongside our most promising technologies, we bring our full humanity, augmenting our rational powers with our feelings, our actions and our faith. We cannot do this by reason alone!

****

Books by Michael L. Dertouzos

The Unfinished Revolution: How to Make Technology Work for Us--Instead of the Other Way Around
Michael L. Dertouzos, with a new forward by Tim Berners-Lee
ISBN: 0066620686
Format: Paperback
240 pages
Reprint edition
Pub. Date: October 8, 2002 by HarperCollins Publishers Inc., New York

Unfinished Revolution: Human-Centered Computers and What They Can Do for Us
Michael L. Dertouzos
ISBN: 0066620678
Format: Hardcover
225 pages
1st Edition
Pub. Date: January 9, 2001 by HarperCollins Publishers Inc., New York

Symbiotic Man: A New Understanding of the Organization of Life and A Vision of the Future
Joel de Rosnay, Phyllis Aronoff (Translator), Foreword by Michael L. Dertouzos
ISBN: 0071357440
Format: Hardcover

299 pages
Pub. Date: March 27, 2000 by McGraw-Hill Companies

Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web
Tim Berners-Lee, With Mark Fischetti, Foreword by Michael L. Dertouzos
ISBN: 006251587X
Format: Paperback
256 pages
First Paperback Edition
Pub. Date: November 7, 2000 by HarperCollins Publishers Inc., New York

What Will Be: How the World of Information Will Change Our Lives
Michael L. Dertouzos and Dert Dertouzos, Foreword by Bill Gates
ISBN: 0062514792
Format: Hardcover
320 pages
1st edition
Pub. Date: March 26, 1997 by HarperSanFrancisco
.....
ISBN: 0062515403
Paperback
384 pages
Pub. Date: March 11, 1998 by HarperSanFrancisco

What Will Be: How the New World of Information Will Change Our Lives
Michael L. Dertouzos, Forward by Bill Gates
ISBN: 0062514792
Format: Hardcover
320 pages
1st HarperCollins hardcover edition
Pub. Date: March 26, 1997 by HarperSanFrancisco 1st edition
.....
ISBN: 0062515403
Format: Paperback
384 pages
1st HarperCollins paperback edition
Pub. Date: March 11, 1998 by HarperSanFrancisco

Made in America: Regaining the Productive Edge
Michael L. Dertouzos, Richard K. Lester, Robert M. Solow, and the MIT Commission on Industrial Productivity
ISBN: 0262041006
Format: Textbook Hardcover
358 pages
Pub. Date: April 3, 1989 by the Massachusetts Institute of Technology
.....
ISBN: 0060973404
Format: Paperback
368 pages
1st Harper Perennial Library edition
Pub. Date: September 1990

The Computer Age: A Twenty-Year View
Michael L. Dertouzos (Editor) and Joel Moses (Editor)
ISBN: 0262540363
Format: Paperback

510 pages
Pub. Date: July 9, 1980 by The MIT Press

Systems, Networks and Computation: Basic Concepts
Michael L. Dertouzos, Michael Athans, Richard N. Spann, and Samuel J. Mason
ISBN: 0070165300
Format: Hardcover
608 pages
Pub. Date: August, 1972 by McGraw-Hill Education

Threshold Logic: A Synthesis Approach
Michael L. Dertouzos
ISBN: 0262040093
Format: Hardcover
256 pages
Pub. Date: December, 1965 by The MIT Press