The Promise of Information Technology – VERSION 9, 2/7/2006

The crucial role played by advances in information technology in our nation's world leadership is well documented. Advances in information technology have led to significant improvements in product design, development and distribution for American industry, provided instant communications for people worldwide, and enabled advances a whole range of health, security, and communications technologies. Recent analysis suggests that the remarkable economic growth the U.S. experienced between 1995 and 2000 was spurred by an increase in productivity enabled almost completely by factors related to IT – "IT drove the U.S. productivity revival," according to Harvard economist Dale Jorgenson. Advances in information technology also have changed the conduct of research, enabling our nation's continued leadership in a broad range of fields – from mapping the human brain to understanding climate change. Faced with problems that are ever more complex and interdisciplinary in nature, researchers are using IT to collaborate across the globe, simulate experiments, visualize large and complex datasets, and collect and manage massive amounts of data.

The crucial role played by federal support for fundamental research in stimulating these advances also is well documented. A 1995 report by the National Research Council referred to the "extraordinarily productive interplay of federally funded university research, federally and privately funded industrial research, and entrepreneurial companies founded and staffed by people who moved back and forth between universities and industry." That report, and a subsequent 1999 report by the President's Information Technology Advisory Committee (PITAC), emphasized the "spectacular" return on the federal investment in long-term IT research and development.

While the payoffs of past computing research have been dramatic, the field remains in relative infancy. Tremendous opportunities remain, animating a broadening core of researchers and promising even greater payoff for future research investments. What follows is a list of areas that have been noted in a broad collection of national and international studies as particularly rich with promise, should researchers make progress on the underlying research challenges. The list is by no means exhaustive, but should provide a deeper insight into the robust vitality of the field.

Driving advances in all fields of science and engineering

The role of high performance computing – simulation – in driving advances in all fields of science and engineering is well documented. Today, though, we are seeing the emergence of a new form of computational science: a computational science that is focused on the collection of massive amounts of data from sensors in the world around us, and by advances in techniques for storing, retrieving, mining, visualizing, and discovering knowledge in that data. Sensors are everywhere – in the oceans, in scientific instruments ranging from telescopes to medical

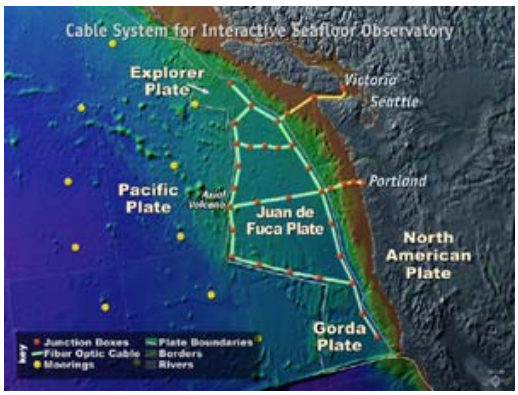

imaging systems, in our civilian infrastructure (buildings, roads, bridges). These sensors generate relentlessly increasing amounts of data. Discovery involves data analysis on a massive scale. Rapid advances in information technology are essential.

Designing a new Internet

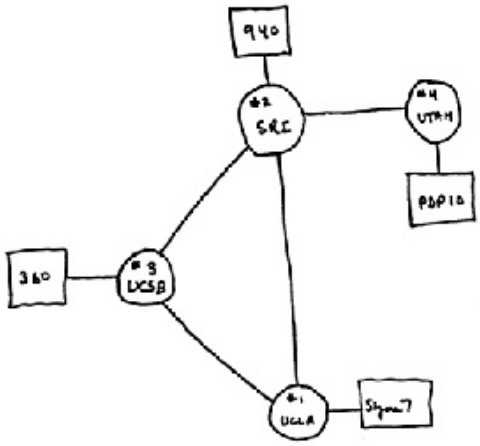

In 2005, Vint Cerf and Bob Kahn received computing's highest prize, the A.M. Turing Award, as well as the National Medal of Technology and the Presidential Medal of Freedom, for their creation in 1973 of TCP (the Transmission Control Protocol), the language of the Internet. It is remarkable that today's Internet utilizes the protocols that Cerf and Kahn devised more than 30 years ago (with many significant engineering improvements, of course). In 1980, there were roughly 200 hosts on the Internet – all of them operated by computer scientists and friends. In 1990, there were roughly 150,000 Internet hosts. Today, there are more than 300 million hosts, and an estimated 1 billion Internet users worldwide. The design of the Internet must be re-envisioned – it has reached its limits in terms of scalability, security, and robustness. Creating a "new Internet" that meets the demands of the 21st century is both a national priority replete with research challenges and opportunities.

Cars that can't crash

Advances in robotics enabled four cars to autonomously negotiate a challenging 130-mile obstacle course in the Mojave Desert in October 2005 – the "DARPA Grand Challenge." This milestone indicates that it is time to launch a program to create "cars that can't crash." In the U.S. alone, automobile accidents cost 40,000 lives and $55 billion each year. Worldwide, nearly 1.2 million die each year, 20 to 50 million more are injured or disabled, and more than $500 billion is lost. Given the tremendous recent progress illustrated by the DARPA Grand Challenge competition, it is reasonable to believe that investing 5% of the potential savings in computer science research for "cars that can't crash" could cut traffic accidents at least in half within a decade. By 2015, in the U.S. alone, we could annually be saving save tens of thousands of lives, hundreds of thousands of injuries, and tens of billions of dollars. Moreover, such an advance would give U.S. products a sizeable competitive advantage in the $1 trillion automotive market.

A teacher for every learner

Making the American dream a reality for all Americans, not just a privileged few, requires that educational excellence become the norm rather than the exception. Although information

technology is not a panacea for all of the shortfalls associated with our educational system, it offers the potential not only for significantly enhancing learning for all learners, but also for transforming the way we learn. We must couple educational practice and educational technology with recent advances in the learning sciences – knowledge of "how people learn." We must develop educational tools including adaptive tutors, massive multiplayer online games, collaborative authoring, learning in context / just-in-time learning, and flexible simulation. It is now within our reach to create a future in which all Americans can achieve the dream.

Augmented cognition

The amount of information with which we are bombarded on a daily basis is increasing relentlessly. Our ability to absorb, evaluate, and act upon that information is not. Information technology is largely responsible for our information overload. Information technology must also provide effective tools for coping – for absorbing and evaluating information, and for bringing to our attention at the appropriate time information that requires action on our part. This "augmented cognition" may be even more important for those who are cognitively impaired – for example, Alzheimer's patients, who with cognitive assistance and monitoring could live fuller, more independent lives.

Computational biology and medicine

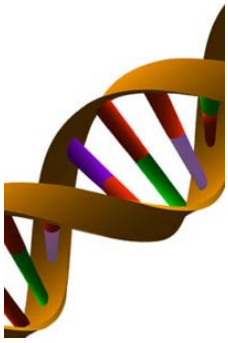

Watson and Crick's discovery of the double helix began the transformation of biology into an information science. The sequence of the human genome has 6 billion features. Such a 6 billion-dimensional vector, generated for each individual, contains the digital information that could be correlated with the health and/or disease states of that individual, leading to a revolution in health care – predictive, preventive, personalized medicine. Because the features of this vector are not independent, a full combinatorial search for the informative and non-redundant feature sets would be an impossible computational problem. Mathematical and computational tools will need to be developed to enable the enormously rich series of correlations to be revealed in a manner that is feasible. The big idea is to use a deep understanding of biology to reduce the informational complexity of these analyses to a computable size. It is clear that similar reasoning could also be applied to the multiparameter blood protein analyses that will also be a part of predictive medicine.

Computers that learn

The impact of computers on productivity has been enormous, but it has reached only a small fraction of its full potential. A key reason is that developing and maintaining software systems is so labor-intensive. What if computers could learn by themselves what to do, instead of having to be programmed and debugged in painstaking detail? Traditionally, computers take programs and inputs and produce outputs. The field of machine learning turns this around: it takes inputs and

outputs and figures out what programs are needed to turn one into the other. From modest beginnings, machine learning systems have become increasingly successful, to the point where today they power important parts of the operations of companies such as Google, Yahoo, Amazon and Microsoft. While highly promising, this is only the beginning. Additional research will dramatically broaden the classes of programs that can be automatically learned. Visiting MIT in late 2005, Bill Gates said "If you invent a breakthrough in artificial intelligence so machines can learn, that is worth 10 Microsofts."

Post-Moore's-Law computing

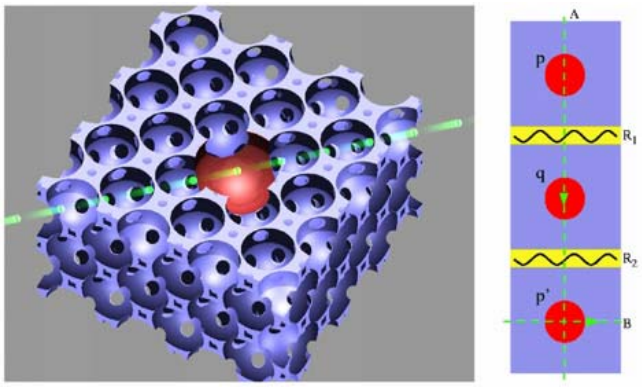

Moore's Law describes the exponential increase in integrated circuit density that we have enjoyed for more than 30 years. Of course, advances in computing have required more than just a decrease in feature size: we have had to invent design methodologies that handle hundreds of millions of transistors rather than tens of thousands … computer architectures that utilize these additional transistors to achieve a proportional increase in performance … system architectures that are synergistic with processor capabilities … system and application software to exploit these new capabilities. We have successfully done all of these things. But the game is changing. Although transistor density continues to increase for now, clock speeds are no longer increasing – "multi-core" is one response, architectures with far more dramatic increases in parallelism are another. Power consumption and heat dissipation are of increased importance, both for mobile and for data center systems. And post-silicon computing substrates – such as quantum computing – must be explored more aggressively.

Global virtual communities and their implications

Information and communications technology has made it increasingly possible for people around the world to share information and to interact with each other in global virtual communities. Even in the early stages of this revolution, we have already seen the enormous social, cultural, and political implications. As geographic boundaries crumble, will culture similarly globalize, or will there be a growing tendency for people to hide within their own virtual communities, leading to social fragmentation? What is the impact of a vast number of global virtual communities involving politics, commerce, and information dissemination? Many tools have already been built to support virtual communities, but these are primitive compared to what will be possible when the social and economic implications of the global reach of such communities are better understood. Research and development that combines computer and social science will be necessary.

Custom-made microbes, at your service

That was the title of a 2006 New York Times article introducing the emerging field of *synthetic biology* – the engineering of biological devices. Imagine a biological cell that could count: 1, 2,

3, 4, 5, 6, ... 256. Now imagine if every cell in the human body could be programmed to count. These engineered cells would find use in the study, and perhaps control, of aging and disease. For example, cells that divided too many times could be programmed to die before they grew into a tumor. Such simple, living computers made from proteins and DNA would not replace silicon-based computers. Rather, these living machines would enable information processing within a substrate for which we currently have no ability to compute – the living world. Making living computers a reality will require that researchers solve many compelling problems at the interface of biology, computer science, and engineering.

The Personal Memex

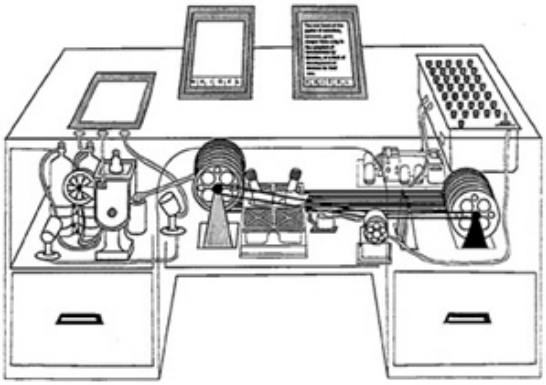

In his seminal 1945 paper "As We May Think," Vannevar Bush described the *Memex*, a device that stored all the information relevant to an individual and could be searched using spoken commands. Dramatic advances in storage are on the verge of making the Memex feasible in terms of cost and size. Equally dramatic advances in search and retrieval technology, though, are needed. Today's search engines (e.g., Google) are remarkable, but they only "pattern match" – they do not "understand." They cannot synthesize an answer to a question – they can only point you to a web page where an answer might be, if such a web page exists. They cannot retrieve relevant images (unless these images are text-labeled) or video – only text. A personal Memex should be able to store and effectively retrieve any digital information ever encountered by its owner, and bring this information to bear on relevant tasks. An enterprise Memex might be equally achievable – and at least equally valuable.

Mastering IT system complexity

The ever-increasing capabilities of computing systems have managed to keep pace with the ever-increasing aspirations that we have for these systems. However, this remarkable progress has been accompanied by ever-increasing complexity. As a result, today's computer systems are tremendously difficult to install, configure, operate, and maintain. The situation is inconvenient, risky, and expensive – typically, annual outlays for maintenance and operations far exceed total hardware and software costs. Research has finally begun to focus on these issues, and there have been some notable successes: companies such as Akamai and Google, for example, efficiently operate collections of tens of thousands of systems that span the globe. These are special situations, though; for the typical home or business desktop system or server facility, the costs of ownership – and the risks – continue to be far too great. A grand challenge in computer systems for the next decade is to reduce these costs and risks – to make as much progress on security, privacy, dependability, and ease of use as we have made on computational cost-performance.

Flattening the world

One of the greatest opportunities we have to foster economic growth and security for our own nation is to improve the status of the billion people on the planet currently living in abject poverty. On the surface, it may seem that IT has little role to play; most of the trappings of IT are far more expensive than can be affordably replicated at this scale. Digging deeper, however, many of the afflictions of these people – inadequate health care, lack of clean water, deficient education, and lack of economic opportunity, for example – can be systematically addressed over time by access to well-designed solutions that incorporate IT as a key element. The design of contextually appropriate information and communication technology to address third world development issues has recently become a major research focus for a number of institutions.

Personalized health care

The combined trends of Moore's Law, MEMs sensors, and low-power radios are enabling an explosion of opportunities to create "sensors for everyone." Embedding sensors in everyday devices such as cell phones, wrist watches, and household appliances can provide a wealth of important information on individuals' personal activity patterns. As an example, obesity researchers want to see day-long and week-long activity patterns so that they can better advise patients on how to alter their behavior. The specifics of where and when people walk, run, use stairs, etc., are important because "lifestyle advice" must be customized to each individual in order to be most effective. In elder care, long-term patterns in the frequency, duration, and mix of an elder's activities can lead to early warning signs of various conditions, both physical and cognitive. Techniques for processing this kind of sensor data range from basic signal processing to sophisticated statistical machine learning. Creating visualization tools and user interfaces that consumers and health care providers can use and perform what-if analyses is another important direction, coupled with psychology research on appropriate motivation strategies.

Information technology, innovation, and creativity

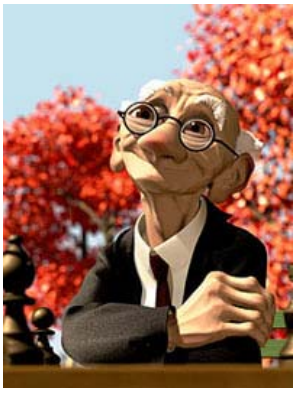

Already information and communications technology has had a major impact on many areas of entertainment and the arts. Computer animation and virtual special effects have transformed movie making. Composing software has allowed even amateurs to create sophisticated new music. New kinds of multi-media art forms are being devised. Computer games are so realistic and engaging that they now often generate more revenue than blockbuster movies, and show tremendous promise for education and training. Interactive drama can adapt to the interests of the audience. Not only is information technology having a major impact on artistic expression itself, but communications technology is changing the way artistic works are distributed, with well known implications for intellectual property rights.

Information technology thus raises a plethora of opportunities and challenges for entertainment and the arts, from building new tools that allow an even greater expansion of artistic expression, to ensuring that artists can maintain intellectual control of their work.

Summary

Advances in information technology have transformed our lives, powered our economy, and changed the conduct of science. But the field remains in relative infancy, and even greater opportunities lie ahead. The preceding examples are by no means exhaustive – they merely serve to suggest the benefits to society of further investments in fundamental research in information technology. The best, truly, is yet to come.

For additional information on these topics, see http://lazowska.cs.washington.edu/promise/