contributed articles

DOI:10.1145/2001269.2001298

A new paradigm is needed to cope with the application, technology, and discipline challenges to our computing profession in the coming decades.

BY ZHIWEI XU AND GUOJIE LI

Computing for the Masses

THE FIELDS OF COMPUTER science and engineering have witnessed amazing progress over the last 60 years. As we journey through the second decade of the 21st century, however, it becomes increasingly clear that our profession faces some serious challenges. We can no longer solely rely on incremental and inertial advances. Fundamental opportunities and alternative paths must be examined.

In 2007, the Chinese Academy of Sciences (CAS) sponsored a two-year study on the challenges, requirements, and potential roadmaps for information technology advances into year 2050.¹² We present a perspective on a key finding of this study: A new paradigm, named *computing for the masses*, is needed to cope with the challenges facing IT in the coming decades.

Computing for the masses is much more than offering a cheap personal computer and Internet connection to the low-income population. It means providing essential computing value for all people, tailored to their individual needs. It demands paradigm-shifting research and discipline rejuvenation in computer science, to create augmented Value (V), Affordability (A) and Sustainability (S) through Ternary computing (T). In other words, computing for the masses is VAST computing.

The CAS study (including the five recommendations illustrated in the accompanying sidebar) focuses on China's needs. However, the issues investigated are of interest to the worldwide computing community. For instance, when considering the drivers of future computing profession, it is critical not to underestimate the requirements and demands from the new generations of digital native population. As of July 2010, 59% of China's 420 million Internet users are between the ages of 6-29 years old. The time frame of 2010-2050 is not too distant a future for them. These digital natives could drive a ten-fold expansion of IT use.

Challenges in the Coming Decades

The first challenge to address is the sobering fact that IT market growth appears to have reached a point of stagnation. The IT market size is measured by the total expenditure on computer and network hardware, software, and services. According to the Organization for Economic Cooperation and Development (OECD),¹⁶ the worldwide IT spending in 2008 grew only 4.49% to \$1,540 billion. The picture is even

» key insights

- China and worldwide data reveal the CS profession is facing many challenges. Future CS should include a focus on e-people, addressing the needs of the masses.
- Computing for the masses demands paradigm-shifting research and discipline rejuvenation in CS, to create ten-fold value growth, to break affordability barriers, and to achieve sustainability through human-cyber-physical ternary computing innovations.
- Computing for the masses has five pillars: A new CS for the human-cyberphysical ternary universe, a universal compute account for everyone, lean system platforms, a science-based Net ecosystem, and national information accounts.

bleaker if we look at the market's longterm history, as illustrated in Table 1. The compound annual growth rate (CAGR) of IT spending declined from double digits before 1980 to low single digits today. These numbers are based on *nominal* U.S. dollar values. Taking inflation into account, the IT market has seen barely any growth.

China is a relative new market for computing. Its IT market has shown

higher growth rates than the worldwide market in the past three decades. Still, we do see signs of slowing down. By estimation from International Data Corporation (IDC), China's IT market growth in 2008–2013 would drop to an annual rate of only 10.8%, in a context of 3.3% worldwide.

Realistic questions must be asked. Will the IT market both in China and worldwide, shrink (in real terms) between 2010–2030? Will it shrink further in the 2030–2050 timeframe? If IT becomes a shrinking sector, what will be the impact to the IT workforce, to computer science education, and to computing research?

The second challenge is that inertial and incremental technology progress faces limitations. The International Technology Roadmap for Semiconductors (ITRS) is the authoritative report on the 15-year assessment of future technology requirements for the semiconductor industry. Its 2009 edition states: "ITRS is entering a new era as the industry begins to address the theoretical limits of CMOS scaling" (www.itrs.net/Links/2009ITRS/ Home2009.htm). The current Internet architecture and protocols have inherent limitations in scalability, manageability, security, and quality of service. A number of "clean slate" future Internet research programs have started up, such as GENI and OneLab, to explore radically new approaches. Supercomputers face issues such as power and system complexity. Zettaflops supercomputers cannot be built without revolutionary devices and architecture.

resentative computer systems built by the Institute of Computing Technology, Chinese Academy of Sciences (ICT-CAS) over the past 50 years. They range from the vacuum tube computer (Model 103) to the latest parallel computer Dawning Nebulae. The figure reveals two critical concerns: power and system software complexity. For 45 years (before 2000), power needs per system never exceeded 100KW. But in the last 18 years, power requirements continuously grew, reaching 3MW in 2010. The system software complexity also grew rapidly in this time. If these trends continue, we may have an exaflops system by 2020, but the power requirement will exceed 100MW, and the system software will grow to over 100 million lines of source code. Can we overcome industrial inertia and fundamentally

Figure 1 depicts parameters of rep-

Table 1.	. IT mar	ket growt	h decline:	s to single	e digit (in compound	d annual growth rate)	
----------	----------	-----------	------------	-------------	-----------	-------------	-----------------------	--

World	Before 1980	1980-2000	2000-2008	2010-2030
	Double Digits	Near 10%	5.5%	1~3%?
China	1981–1990	1990-2000	2000-2008	2010-2030
	30%	44%	17%	5~10%

Compiled by the authors based on data from IDC, OECD and CCID

The China numbers are calculated based on data from China Center of Computer and micro-electronics Industry Development (CCID). The worldwide numbers are calculated based on data from IDC (before 1980 and 1980–2000) and OECD (2000–2008). See also calculations by Shane Greenstein of Northwestern University ("Are the Glory Days Long Gone for I.T.?" *New York Times*, Aug. 8, 2009).





address power, systems complexity, concurrency, cost, and reliability issues to reach realistic zettascale $(10^{21}$ flops) systems?

Computer science as an academic discipline also faces challenges. For the past five decades, computing has always been a highly regarded profession in China. However, we now see a disturbing phenomenon: Parents of university freshmen are advising their children to study business or law to make a good living. Moreover, parents urge students interested in pursuing a Ph.D. to choose a major in mathematics or physics. In China's Ph.D.- granting universities, computer science as an elected major has dropped from the top spot to out of the top five.²⁵

Our study reveals a primary reason for this trend is that young people today see computer science as a *tool*. That is, computing is a universal tool, thus offers good employment. However, computing is regarded as just a tool in many professions, not a primary value driver; computing requires long working hours doing boring tasks such as documentation, programming, and testing; and computing is a rapidly changing field, quickly turning skills learned in college obsolete.

In contrast, a medical school education provides life-long benefits. Medicine is not a boring tool, but a rewarding profession providing essential value to society. The coming generations of Chinese youth are manifestly more self aware and socially aware, caring about environment and community values. How do we convince perspective students that computing is an academically beautiful, intellectually stimulating, and socially rewarding profession?

These challenges do not mean that computing is a declining profession. In fact, its future potential is great because of the need to continually increase the direct and added value of IT, especially for those products and services targeting mass users. As a latecomer, China's IT use lags by several decades behind the developed countries. Its per-capita IT spending in 2008, as a measure of direct value of IT, was only \$85. Still, it has great growth promise, particularly in terms of adding value to other societal sectors throughout China, such as Internet services, energy, transportation, education, health care and entertainment.

China's urbanization in the next 30 years will move 1% of its population to the cities annually. In the next decade alone, China needs to build 80 million new homes to house the expanding urban population. Over 200 million smart meters will be installed in urban homes to support a smart grid and to save energy. More than 10 million cars were sold in China in 2009. All these modern automobiles, homes, smart meters need increasing IT use to add value.

China's Internet services sector (\$15 billion revenue in 2009) serves hundreds of millions of users but is not counted in the IT market. It continues to grow on a double-digit scale. Taobao, China's top online shopping Web site, predicts its number of page views per day will grow 80% annually for the next five years, to reach 33 billion page views per day. The Internet services sector is being augmented by cloud computing. China Education Television is using cloud computing to create a China Education Open Mall, to provide lifelong education to 800 million people via telecommunication, digital television, two-way satellite networks, and the Internet. The IDC 2009 Economic Impact Study report predicts that cloud computing and associated clients will add \$80 billion to China's IT market in 2014.

The same IDC report also predicts that China's IT job market will grow 7.2% annually for the next five years, from 4.5 million in 2008 to 6.4 million in 2013, thus maintaining the healthy IT job market growth seen during the period of 2000–2008 when the number of software workers in China grew from 30,000 in 2000 to over 1.5 million in 2008. The expected growth may be a bit slower in some markets, but IT areas that target the masses—like the Internet services sector—will see high double-digit growth in job openings.

Call for a New Paradigm

The challenges and opportunities we have noted in this article suggest that computing must transform itself, yet again, from the social perception of a high-tech tools discipline to the manifestation of its authentic core of providing essential informational value to Young people today see computer science as a tool. That is, computing is a universal tool, thus offers good employment. However, computing is regarded as just a tool in many professions, not a primary value driver. all people; that is, computing for the masses.

Computing for the masses should include the following four specific features:

► Value-augmenting mass adoption: By 2050, China's IT users should cover 80% of the population, with the percapita IT expenditure increasing 13 times over.

► Affordable computing: Each mass user's total cost of ownership of IT use should go down significantly, when experiencing the same or increasing IT value.

► Sustainable computing: The order of magnitude growth in IT use should not see the order of magnitude growths in energy consumption and emission.

► *Ternary computing*: These features can only be achieved with transformative innovations. A major source of opportunities for such innovations is the coupling and interaction of the human society, the cyberspace, and the physical world.

Value-augmenting mass adoption is the overarching goal of computing for the masses. Affordability and sustainability are two important constraints for the 21st century. Ternary computing represents the transformative innovations needed to achieve the goal under the two constraints, utilizing trends in computer science research. Let's look at these four features in more detail.

Value Augmenting Mass Adoption. An obvious goal of computing for the masses is to expand the user base to 80% of the population. For China, this translates to 1.2 billion users by 2050. To achieve this goal, computer use should go far beyond institutions such as companies, governments, or science labs. It should target the masses as personal or community users. This implies that future computing workloads should change significantly. According to IDC's data, the worldwide servers market reached its peak of \$65.5 billion in 1997, where over 95% ran institutional workloads, but only 4.3% ran personal workloads. Computing for the masses advocates that by 2050, personal workloads may expand many times to use over 50% of servers worldwide.

Furthermore, computing should offer value services personalized to the

contributed articles

masses' individual requirements, thus enlarge the per-capita IT value. We are far from really understanding what IT value is for the masses. IT expenditure only partially reflects IT value. IT provides added value to other sectors which is not counted as IT expenditure. IT also provides personal and societal value that is not measured in economic numbers. The French economist Yann Moulier Boutang recently likened IT to bees, where the added value (pollination) is much larger than the direct value (honey).

Studies in other fields could give us some hints on defining IT value. Lazo, et al.¹⁰ surveyed U.S. households to measure how much dollar value households place on weather forecasts, to show that the total value of weather services (\$31.5 billion) is much larger than their total cost (\$5.1 billion). Nordhaus¹⁵ used Lumen-Hour to measure the value of lighting, to show that advances in lighting technology improved value/price by five orders of magnitude in 1800–2000. The CAS study uses the annual IT expenditure per user as an approximation of IT value. Five classes of IT value are defined:

The *IT poverty line* measures the minimal value that should be provided to any citizen. Today such value manifests in an inexpensive personal computer enabling basic applications and Internet access. The *commodity value*, in addition, offers commodity IT products and services targeted to a wide range of users. The *ubiquity value* adds mobile or "anywhere" IT services to the commodity value. The *expertise value* provides extra IT value for a specific field of expertise, such as an animation system for a cartoon artist. In

Table 2. Two scenarios of IT growth in China 2008–2050.

Historic Data and Projections	IT Market (\$ Trillion)	IT Market (CAGR)	IT Users (Million)	IT Spending per Capita
2000 (Actual Data)	0.026	25.0%	22.5	\$21
2008 (Actual Data)	0.11	12.7%	270	\$85
2050 (Poverty Line Growth)	0.25	2.0%	1,200	\$190
2050 (Value-Augmenting Growth)	2.0	7.1%	1,200	\$1,321

CAGR: Compound Annual Growth Rate.

Figure 2. Predictions of China's fossil fuel consumption and CO_2 emission (left axis), compared to IT expenditure (right axis). All are per-capita data normalized to year 2010 values.



2008, we estimate the IT poverty line value in China to be \$150, the commodity value to be \$300-500, the ubiquity value to be \$500-1,000, and the expertise value to be \$1,000-10,000, respectively. The personalized value refers to the user-centric situation, where IT hardware, software, and services are customized to an individual user's needs. For example, the Loongson (Godson) CPU research team at the authors' institute uses a highly tuned computing environment in designing microprocessor chips. The EDA software providers and the downstream manufacturer provide onsite, customerspecific services. The team's annual IT spending is over \$20,000 per user.

Table 2 contrasts two growth paths for China, both extending IT users to 1.2 billion people from 270 million in 2008. The poverty-line growth scenario offers the 930 million new users only poverty line IT value. This is a common conception for reaching the masses, but an unlikely scenario. It ignores the fact that by 2050, the majority of China's population will enter the middle class and most of the Chinese population will be digital natives. They will demand more than the IT poverty line value. The value-augmenting scenario makes two assumptions: by 2050, China's per-capita annual IT spending should approach that of the U.S. in 2000, which was about \$1,400; and the digital divide will not worsen; that is, the user distribution among the value classes stays the same.

The value-augmenting projection is not overly optimistic. Even when the projection is achieved, China's information welfare in 2050 will still be 50 years behind that of developed countries. Another fundamental support lies in China's education effort for the masses. In 2008, seven million students graduated from colleges. They all took one-semester to one-year's worth of computer literacy courses. This trend is likely to continue for decades. The implication is that by 2050, China will have educated over 300 million new computing literate college graduates. They will require at least ubiquity value in IT use.

Affordable Computing. Computing for the masses is affordable computing for everyone's value creation and consumption. For instance, future high school students in China may learn protein folding in a biology course based on computational thinking, by accessing petaflops computer simulation service for a week, with a lab cost of \$10 per student. However, four barriers to such value affordability exist today. They are:

▶ Cognition barrier. Computing must offer readily cognizable value, for example, a learning experience in protein folding, not just a petaflops tool. Showing value instead of ware to the society is often difficult. Computing may also be associated with negative value. A 2009 study by Beijing Women Association showed that 17% of parents strongly oppose children's access to the Internet, and 66% allow supervised access. The top concerns are Internet games (cited by 45% of parents) and online porn (40% of parents). How to make the cognizable value of computing outweigh its negative effect is a challenge.

► *Cost barrier.* Petaflops capability can be rented via cloud computing services today, but the price must be reduced by six orders of magnitude.

► *Control barrier.* A user's computing activities today are often tied to a few platforms or vendors. User creativity is hindered by platform or vendor control. Switching a computing provider or platform is much more difficult than switching a TV channel.

► Usability barrier. Computing needs to provide value-level user interface, and hide low-level details such as lowlevel coding, deployment, configuration, maintenance, monitoring, debugging, optimization, and adjustments to technical changes.

Affordable computing implies that we should aim for reducing the total cost of ownership for the masses, including purchasing cost, learning cost, use cost, and platform switching cost. An affordable computing product or service not only comes with a low price, but also provides readily identifiable value, freedom from platform control, and ease of use.

Sustainable Computing. An important constraint for computing in the 21st century is sustainability requirements. From 1980 to 2008, China's energy consumption per dollar GDP decreased 208%. A more desirable scenario is to maintain a reasonable GDP

Five Sample Tests for Transformative Innovations

TEST 1. RANK LOW-CARBON HOUSEHOLDS IN BIG CITIES AUTOMATICALLY. Automatically rank the households in Beijing and London by the best practices in terms of energy efficiency, and contrast the two cities' top 10 best practices. Learning from communities can enable behavior changes of mass users, which could significantly reduce energy waste and carbon emission.

A Net IT system connects the physical world (electric grids and home appliances) and human behaviors (home electricity usage patterns) of the millions of households in Beijing and London. This ternary computing system ranks the households in each city and reveals the top 10 best practices. The computation is fully automatic, from data sensing to behavior discovery and ranking. The computation is fast in that the knowledge is instantaneously retrievable and sufficiently up to date. The household behaviors may be strategic and evolutional, thus the results are dynamic, different every week, day, hour, or even minute. Data objectivity is maintained and household privacy is protected.

TEST 2. USE INFORMATION ASSETS VIA UCA, EVEN WHEN TIME AND SPACE HAVE SHIFTED.

Any mass user, for example, a college graduate, can seamlessly use all his current information assets 10 years later via the same Universal Compute Account. The graduate will have accumulated substantial personalized informational value and assets, all accessible in the future, despite the fact his device, location, provider, network technology may have changed many times. Similarly, the graduate should be able to seamlessly use all his information assets when he moves to another country.

TEST 3. INCREASE DATACENTERS EFFICIENCY BY ORDERS OF MAGNITUDE.

Typical million-server datacenters each consume much less than 100 million watts, and achieve Performance-Energy-Efficiency (PEE) greater than 10–20%. This should result from understanding and innovations in devices, hardware architecture, software, services, and sector ecosystem policies.

TEST 4. FACILITATE MASS CREATIONS OF VALUE NET SERVICES.

A science-based Net ecosystem is in place such that a class of high-school students in Sichuan province, utilizing community assets, can learn to create and deploy a value Net service, such as a trusted, individualized computational protein folding service, within a month. They should be able to do it independently, and other schools can use the service to teach proteins folding, without requiring a vendor or provider's approval.

TEST 5. ENABLE SCIENCE-BASED MACRO INFORMATION POLICIES.

A national information accounts system is in place to allow a country to have effective information fiscal policies and monetary policies. A scientifically designed accounts system supported by a meticulously instrumented ternary world could help answer the question: How to price the Net to maximize 1.2 billion people's welfare? A more concrete example is the invention of a science-based messaging tax policy following the computer science principle of being "strict on output and tolerant on input," which could drastically reduce spam while facilitating legitimate communication.

growth while achieving zero growth in both fossil fuel consumption and CO_2 emission. A projection based on this scenario is illustrated in Figure 2, utilizing data from a CAS energy roadmap study.²⁴ China's per-capita IT expenditure will reach \$1,300 by 2050, an increase 13 times over the 2010 value. But growths in fossil fuel consumption and CO_2 emission will significantly slow down, so the per-capita numbers in 2050 will return to the 2010 levels.

Sustainability is a difficult constraint that computing for the masses must address. Koomey⁹ calculated that electricity use for datacenters doubled worldwide from 2000 to 2005, and the Asia Pacific region grew faster than the worldwide annual rate of 16.7%. IDC predicted that from 1998 to 2012, the number of servers installed worldwide will grow at an average annual rate of 11.6% to reach 42 million servers, with power and cooling costs reaching \$40 billion.²¹ In 2008, China's telecommunication carriers consumed 24 billion KWH of electricity, growing 14% annually. Such growth trends of IT energy consumption are contrary to the objectives of Figure 2, and cannot be sustained. When computing users in China grow 3.4 times to 1.2 billion in 2050 with intensified IT use, current practices will require 150 billion KWH just to power China's datacenters.

In the past 40 years, energy and environment impact for computing research and IT use in China was not a main issue, because the IT user base was only a small part of the population, and the IT sector is more efficient and clean than the rest of the economy. In designing future IT systems for the masses, saving energy, reducing pollution, and using recyclable material will become hard design objectives, as important as functionality, cost, and speed. Innovations are needed to achieve these objectives while maintain healthy growth of the IT market. Utilizing the diversity factor in the workloads of the mass users offers an opportunity.

Computing can also help sustainability in other sectors that involve mass users. IDC predicts that technologies such as smart meters, energy management systems for buildings, intelligent building design, and teleworking, all involving IT and affecting mass users, can help reduce China's CO_2 emission by 532 million tons in 2020.²¹

Ternary Computing. Computing for the masses is not just a lofty ideal to conquer the digital divide. It is based on and required by an economical and social trend in China and in the world: people are increasingly living in an intermingling ternary universe of the physical world, the human society, and the cyberspace. Computing is becoming a common fabric of the ternary universe that relates to value creation through information transformation.8,19,23 People expect computing services and computational thinking to enrich their physical and societal lives, through IT devices, tangible interfaces, and intangible interfaces.

Thus, computing for the masses calls for *ternary computing*—transformative innovations in computing that utilize the ternary universe trend. We focus here on five requirements pertinent to ternary computing for the masses, ranging from computer science foundation, ease of personalized We are far from really understanding what IT value is for the masses. A French economist recently likened IT to bees, where the added value (pollination) is much larger than the direct value (honey). access, efficiency in use, effectiveness in value creation and consumption, and ways to measure and regulate. From the academic angle, these form the five pillars of ternary computing for the masses.

Pillar 1: Computer Science for the Ternary Universe. When our traditional computer science was established, we assumed a man-machine symbiosis system consisting of a human user interacting with a computer.¹³ In the 21st century, the masses will live in a ternary universe of the cyberspace, the human society, and the physical world, clustered into various value communities. This calls for an augmented computer science that provides a foundation to investigate the computational processes and phenomena in the ternary universe. The main target of computer science study will be the Net, that is, the common computational fabric of the ternary universe, which is a Net of people, bits, and things. Here, we suggest five research areas:

► Algorithm networks. We need to enhance traditional computability and algorithmic theories to account for time and space bounds of networks of interacting algorithms. New complexity metrics, such as energy complexity and effort complexity, are required to measure and study energy consumption and human labor needs.

► *Ternary systems modularity.* We need to discover modularity rules in ternary computing systems, similar to the Liskov substitution principle¹⁴ in object-oriented systems. Such rules involve human-cyber-physical resources and should enable seamless substitution of a Net service with a better one.

► Fundamental impossibility. We need more impossibility results like Brewer's CAP theorem (the impossibility of simultaneously satisfying Consistency, Availability, and Partition tolerance)⁵ that rigorously relate basic properties of Net services. We especially need results that bridge qualitative issues (for example, privacy and safety) to quantitative issues (for example, scalability and energy saving), that are important in the ternary universe.

► *Emergence and locality.* We need to positively embrace complex systems, by learning to utilize their emergent properties as an advantage, not a hurdle. An example is to map various net-

work phenomena to new locality principles^{2,11} to enable the design of better Net services and to facilitate the emergence of desired ternary communities.

• Computing in Nature. We need to understand Les Valiant's evolution problem:²² How did Nature "compute" the Homo sapiens so efficiently, generating a genome with 3×10^9 base pairs in only 3×10^9 years? The answer could help us learn how Nature computes. The principles learned can be used in building good ternary systems.

Pillar 2: A Universal Compute Account for Everyone. Each of the 1.2 billion users in China will have a tetherless, seamless, life-time universal compute account (UCA) to access the Net. The UCA is not tied to any device, location, network, resource, or vendor. A user will "log on" to his or her UCA. The UCA is not merely a user's identifier, but a personal information environment, a uniform grip to the local and Net resources that the user is entitled to use. In effect, it offers a user his personal Net, where resources are used on demand and the services are charged by usage. The UCA could also behave as an entry to a personal server, enabling the user to contribute and share. Over one billion such UCAs can generate great community value, leading to the 2^w network effect.¹⁸

For computing to reach the masses, IT must offer some fundamental stability and constancy. The UCA is such a fundamental invariant as it could help deliver ever occurring new values while making technical idiosyncrasies and upgrades invisible. It provides a stable focal point for continued innovation and value delivery, improvement in ease of use, and personality accumulation. We have seen similar precedents, but targeting resources instead of users, in URI and RESTful services.⁴

Pillar 3: Lean System Platforms for the Masses. Once a user logs onto the Net, he or she should see a lean system, where most resources are used for generating application value. We do not yet have good metrics for measuring *applications value* delivered per watt. For datacenters, we can define an approximation called the *Performance-Energy-Efficiency* (PEE) measure, where the percentages of energy spent on application speed is used as a value indicator: as shown in Figure 3. The Power Usage Effectiveness (PUE) is defined as the amount of total power entering a datacenter divided by the power consumed by the IT hardware (servers, storage, networks, among others). Assume a total life time of four years for datacenter IT hardware. Utilization is defined as the percentage of this total time when the IT hardware is busy running applications, instead of running system tasks or staying idle. Efficiency is defined as the application speed achieved divided by the peak speed of the IT hardware, when the hardware is running applications.

Data generated from the CAS field studies are summarized in Table 3. The most efficient scenario refers to an efficient supercomputing center running optimized Linpack-like applications. For every watt spent, 0.4 watts are used for applications. The efficiencies are much lower for typical usage scenarios. The potential for improvement could grow two to three orders of magnitude.

We are far from fully understanding existing platform inefficiencies. We need new micro and macro benchmarks and metrics that characterize workloads for the masses, to guide the design of parallel, distributed, and decentralized computer architectures, especially the cooperation and divide of labor among application software, compiler, execution model, and hardware. We lack a design formula for the servers and datacenters that can relate value, resource, and energy in a precise way, as Hennessey and Patterson's formula¹⁷ did for the CPU microarchitecture designs by precisely relating performance to the number of instructions, cycle-per-instruction and clock frequency.

There exist much space for enhancing efficiency in the entire Net, including networks, client devices, and sensors. New internetwork architectures supporting network virtualization must be investigated. Client machines should support more intuitive interactions via three-dimensional, multimodal and semantic interfaces. When comparing today's PC systems to the Xerox PARC personal computer, Alan Kay estimated that "approximately a factor of 1000 in efficiency has been lost."³

Pillar 4: A Science of the Net Ecosystem. IT growth in the past 30 years benefited significantly from a rich IT ecosystem, manifesting as the whole of guiding principles, infrastructure interfaces, technical standards, and the interactions of users, academia, industry, and volunteer communities. However, we do not yet have a science of IT ecology. Computing for the masses will lead to mass creation and innovations utilizing network effects. This calls for a science of the Net ecosystem to study the effectiveness of mass creation and consumption in the ternary universe.

Two research topics are important to start establishing such a science. The first is to identify the basic objective attributes of the Net ecosystem. For creators, we must inherit the Internet *openness* and *neutrality* principle¹ to enable mass creation and to prevent monopoly. For users, we need a safe Net, with *trust*, *security*, and *privacy*. Both the creators and users should not feel isolated, but free to share and contribute to *community assets*, such as shared data, software, hardware, groups, and processes. A Net ecosystem with these six attributes is called a

Table 3. PEE values for different datacenter usage scenarios.

Usage Scenario	PEE	Efficiency	Utilization	PUE
Most Efficient	40%	0.75	0.70	1.3
Typical	0.12%-1.2%	0.05-0.10	0.07-0.25	2.1–3.0

Figure 3. Performance-Energy-Efficiency (PEE) measure.

 $\mathsf{PEE} = \mathsf{Efficiency\ Utilization/PUE} = \frac{\mathsf{App\ Speed}}{\mathsf{Peak\ Speed}} \quad \frac{\mathsf{Time\ for\ App}}{\mathsf{Total\ Time\ }} / \frac{\mathsf{Total\ Power}}{\mathsf{IT\ Power}}$

harmonious ecosystem.

The second topic is to develop a rigorous computing model that can precisely define and classify these six attributes, and relate them to new technology. We can start with simple ternary computing systems, where component elements can be superimposed. For instance, the studied system can be a centralized cloud service for exchanging household best practices of saving energy. Model checking techniques can be developed to verify the service has sound mechanisms, such that adding a community asset will not negatively affect household privacy.

Pillar 5: National Information Accounts. As IT use becomes an indispensible part of people's life and a primary value driver of national economy, it demands meticulous and scientific accounts, in a way similar to how economists created the national economic accounts system²⁰ now used worldwide. This national accounts system for the Net, supported by countrywide information meters, can give us objective pictures of the stocks and flows of information materials, goods, services, and users. It helps identify and understand IT's use, cost, value, bottlenecks and opportunities, as well as societal and environment impacts. It serves as a foundation for objectively assessing individual IT projects and national IT policies, by scientifically linking IT resources to IT values.

Such a national information accounts system is also good for rejuvenating the computing discipline. It could add to a basis for a pervasive yet unified computing discipline based on computational thinking, by scientifically organizing massive data from the ternary universe. Micro and macro informatics insights could be gained from such organized and linked data, as our discipline evolves from numerical computing, symbolic computing, process computing to data computing.^{6.7}

Conclusion

Over the decades our profession did invent many transformative technologies benefiting the masses, such as the personal computer and the Web. For instance, the PC significantly reduced computing cost, enlarged the user base, *and* augmented the IT market. It achieved all these benefits not by dumbing down a mainframe computer, but through transformative innovations such as interactive computing, desktop-based graphical user interface technology, object-oriented programming, and the Ethernet. In the 21st century, we must study and understand such successful endeavor and do it again, but in a more fundamental and systematic way. In the next 40 years, the rich demands, workloads, and experiences of the mass users will generate unprecedented economic scale and intellectual stimulus, expanding computing's width and depth in ways that we are now barely able to see.

Acknowledgments. This work is supported in part by China Basic Research Program (2011CB302500, 2011CB302800). The authors would like to thank the CAS IT Roadmap Study Group for stimulating discussions, and the anonymous referees for their constructive comments.

References

- 1. Cerf, V. The open Internet. *Telecommunications Journal of Australia 59*, 2 (July 2009), 18.1–18.10.
- Easley, D. and Kleinberg, J. Networks, Crowds, and Markets: Reasoning About a Highly Connected World. Cambridge University Press, 2010.
- Feldman, S. A conversation with Alan Kay. ACM Queue 2, 9 (Dec. 2004), 20–30.
- Fielding, R. T. and Taylor, R. N. Principled design of the modern Web architecture. ACM Transactions on Internet Technology 2, 2 (May 2002), 115–150.
- Gilbert, S. and Lynch, N. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant Web services, ACM SIGACT News 33, 2 (June 2002), 51–59.
- Hendler, J., Shadbolt, N., Hall, W., Berners-Lee, T., Weitzner, D. Web science: An interdisciplinary approach to understanding the Web. *Commun. ACM* 51, 7 (July 2008), 60–69.
- Hey, T., Tansley, S., and Tolle, K., editors. *The Fourth* Paradigm: Data-Intensive Scientific Discovery. Microsoft Research, Redmond, WA, 2009.
- Karp, R. Understanding science through the computational lens. J. of Computer Sci. and Tech 26, 4 (2011), 569-577.
- Koomey, J. G. Worldwide electricity used in data centers. Environmental Research Letters 3, 034008 (Sept. 2008), 1–8.
- Lazo, J. K., Morss, R. E., and Demuth, J. L. 300 billion served: Sources, perceptions, uses, and values of weather forecasts. *Bulletin of American Meteorological Society* (June 2009), 785–798.
- Lee, R. and Xu, Z. Exploiting stream request locality to improve query throughput of a data integration system. *IEEE Transactions on Computers 58*, 10 (Oct. 2009), 1356–1368.
- Li, G., editor. Information Science and Technology in China: A Roadmap to 2050. Science Press Beijing and Springer-Verlag Berlin, 2010.
- Licklider, J. Man-computer symbiosis. IRE Transaction on Human Factors in Electronics HFE-1, 1 (1960), 4–11.
- 14. Liskov, B. and Wing, J. A behavioral notion of subtyping. ACM Transactions on Programming
- Languages and Systems 16, 6 (Nov. 1994), 1811–1841.
 15. Nordhaus, W. D. Do real output and real wage measures capture reality? The History of lighting suggests not. In T.F. Bresnahan and R.J. Gordon, editors. The Economics of New Goods. The University of Chicago Press, 1997.
- 16. OECD. The OECD Information Technology Outlook

2008. 2009.

- O'Hanlon, C. A conversation with John Hennessy and David Patterson. ACM Queue 4, 10 (Dec.-Jan. 2006–2007), 14–22.
- 18. Raman, T.V. Toward 2^w, beyond Web 2.0. *Commun. ACM 52*, 2 (Feb. 2009), 52–59.
- Snir, M. Computer & information science & engineering—What's all this? Keynote Speech at the 2nd NSF-NSFC Sino-USA Computer Science Summit. Washington, DC, July 2008.
- Stone, R. The accounts of society. Nobel Memorial Lecture (Dec. 8, 1984); http://nobelprize.org/nobel_ prizes/economics/laureates/1984/stone-lecture.pdf.
- Turner, V., Bigliani, R., and Ingle, C. Reducing Greenhouse Gases through Intense Use of Information and Communication Technology. International Data Corporation, 2009.
- Valiant, L. Examples of computational models for neuroscience and evolution. Speech at the Princeton Workshop on the Computational Worldview and the Sciences. Dec. 11, 2006.
- 23. Wing, J. Computational thinking. *Commun. ACM* 49, 3 (Mar. 2006), 33–35.
- Chen, Y. editor. Energy Science and Technology in China: A Roadmap to 2050. Science Press Beijing and Springer-Verlag Berlin, 2010.
- Zhang, M. and Lo, V. Undergraduate computer science education in China. In *Proceedings of ACM SIGCSE* 2010 (Mar. 2010).

Zhiwei Xu (zxu@ict.ac.cn) is a professor and CTO of the Institute of Computing Technology, Chinese Academy of Sciences.

Guojie Li (lig@ict.ac.cn) is a professor and director of the Institute of Computing Technology, Chinese Academy of Sciences.

© 2011 ACM 0001-0782/11/10 \$10.00