Implications of Historical Trends in the Electrical Efficiency of Computing

Jonathan G. Koomey
Stanford University

Stephen Berard
Microsoft

Marla Sanchez
Carnegie Mellon University

Henry Wong
Intel

The electrical efficiency of computation has doubled roughly every year and a half for more than six decades, a pace of change comparable to that for computer performance and electrical efficiency in the microprocessor era. These efficiency improvements enabled the creation of laptops, smart phones, wireless sensors, and other mobile computing devices, with many more such innovations yet to come.

Valentine's day 1946 was a pivotal date in human history. It was on that day that the US War Department formally announced the existence of the Electronic Numerical Integrator and Computer (ENIAC).[1] ENIAC's computational engine had no moving parts and used electrical pulses for its logical operations. Earlier computing devices relied on mechanical relays and possessed computational speeds three orders of magnitude slower than ENIAC.

Moving electrons is inherently faster than moving atoms, and shifting to electronic digital computing began a march toward ever-greater and cheaper computational power that continues even to this day. These trends proceed at easily measurable, remarkably predictable, and unusually rapid rates. For example, the number of transistors on a chip has doubled more or less every two years for decades, a trend that is popularly (but often imprecisely) encapsulated as Moore's law (see Figure 1).

Moore's law has seen several incarnations, some more accurate than others. In its original form, it was not a physical law, but an "empirical observation" that described economic trends in chip production.[2] (It has also sometimes served as a benchmark for progress in chip design, which we discuss more later on in this article.) As Gordon Moore put it in his original article, "The complexity [of integrated circuits] for minimum component costs has increased at a rate of roughly a factor of two per year,"[3] where *complexity* is defined as the number of components (not just transistors) per chip.

The original statement of Moore's law has been modified over the years in several different ways, as previous research has established.[4,5] The trend, as Moore initially defined it, relates to the minimum component costs at current levels of technology. All other things being equal, the cost per component decreases as more components are added to a chip, but because of defects, the yield of chips goes down with increasing complexity.[6] As semiconductor technology improves, the cost curve shifts down, making increased component densities cheaper (see Figure 1 in Moore's 1965 paper[3]).

In 1975, Moore modified his observation to a doubling of complexity every two years,[7]

 Published by the IEEE Computer Society 1058-6180/11/$26.00 © 2011 IEEE

which reflected a change in the economics and technology of chip production at that time as well as a change in his conceptualization of the law. That rate of increase in chip complexity has held for more than three decades, which is a reflection mainly of the underlying characteristics of semiconductor manufacturing during that period. As Ethan Mollick explained,[5] Moore's law is in some sense a self-fulfilling prophecy—the industry's engineers have used Moore's law as a benchmark to which they calibrated their rate of innovation. This result is partly driven by business dynamics, as David E. Liddle points out: "Moore's law expresses that rate of semiconductor process improvement which transfers the maximum profit from the computer industry to the semiconductor industry."[2]

The striking predictive power of Moore's law has prompted many to draw links between chip complexity and other aspects of computer systems. One example is the popular summary of Moore's law ("computing performance doubles every 18 months"), which correlates well to trends in personal computer systems to date but is a statement that Moore never made. Another is "Moore's law for power," coined by Wu-chun Feng to describe changes in the electricity used by computing nodes in supercomputer installations during a period of rapid growth in power use for servers ("power consumption of compute nodes doubles every 18 months").[8]

This article describes the implications of the relationship between the processing power of computers (which in the microprocessor era has been driven by Moore's law) and the electricity required to deliver that performance. Over the past 65 years, the steps taken to improve computing performance also invariably increased the electrical efficiency of computing, whether the logical gates consisted of vacuum tubes and diodes, discrete transistors, or microprocessors. The most important historical effect of this relationship has been to enable the creation of mobile computing devices such as laptop computers. If these trends continue, they will have important implications for more widespread future use of mobile computing, sensors, and controls.

Methods for Deriving Trends

Analyzing long-term trends is a tricky business. Ideally we'd have performance and energy use data for all types of computers in all applications since 1946. In practice, such data do not exist, so we compiled available data to piece together the long-term trends on the electrical efficiency of computation.

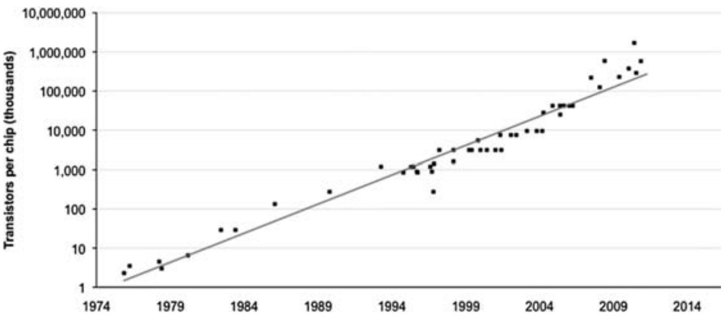

Figure 1. Transistor counts for microprocessors over time (thousands). The doubling time from 1971 to 2006 is approximately 1.8 years. (Data courtesy of James Larus, Microsoft Corporation)

To estimate computations per kilowatt-hour, we focused on the full-load computational capacity and the direct active electrical power for each machine, dividing the number of computations possible per hour at full computational load by the number of kilowatt-hours consumed over that same hour. This metric says nothing about the power used by computers when they are idle or running at less than full load, but it is a well-defined measure of the efficiency of this technology, and it helps show how the technology has changed over time.

Measuring computing performance has always been controversial, and this article will not settle those issues. In 2007, William D. Nordhaus published a sophisticated and comprehensive historical analysis of computing performance over time,[9] and that's the source for performance data on which we relied most heavily. We combined those data with measured data on the power use of each computer when operating at full load to calculate computations per kilowatt-hour. (More details on the data and methods are available in our Web Extra appendix, see http://doi.ieeecomputersociety.org/10.1109/MAHC.2010.28).

Historical Trend Results

Figure 2 shows performance per computer for all the computers included in Nordhaus' analysis from 1946 onward. We also added the 40 additional machines from this analysis for which measured power and performance were available. The figure does not include performance estimates for recent large-scale supercomputers (for example, those at www.top500.org), but it does include measurements for server models that

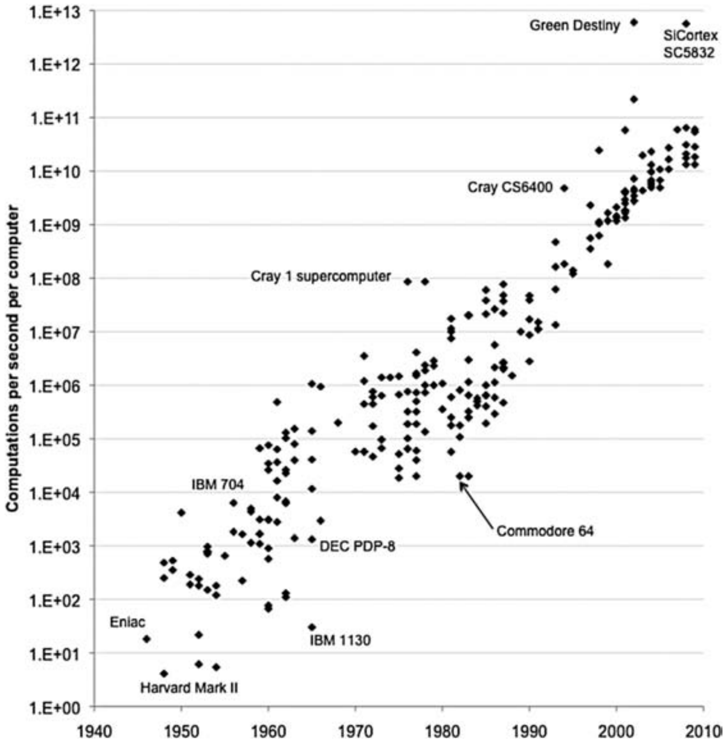

Figure 2. Computational capacity over time (computations/second per computer). These data are based on William D. Nordhaus' 2007 work,[9] with additional data added post-1985 for computers not considered in his study. Doubling time for personal computers only (1975 to 2009) is 1.5 years.

are often used as computing nodes for those machines.

The trends for microprocessor-based computers are clear. The performance per unit for PCs, regressed over time, shows a doubling time of 1.50 years from 1975 (the introduction date of the Altair 8800 kit PC) to 2009.[10] This rate corresponds to the popular interpretation of Moore's law, but not its exact 1975 formulation.

Figure 3 shows the results in terms of the number of calculations per kilowatt-hour of electricity consumed for the computers for which both performance and measured power data are available. These data include a range of computers, from PCs to mainframe computers.[11]

The transition from vacuum tube to transistorized computing is clearly evident in the data. During 1959, 1960, and 1961, as transistorized computers came to market in large numbers, there is a difference of about two orders of magnitude between the most and least electricity-intensive computers. Logical gates constructed with discrete transistors circa 1960 used significantly less power than those made with vacuum tubes and diodes, and the transition to transistors also led to a period of great technological innovation as engineers experimented with different ways to build these machines to maximize performance and improve reliability.

Computations per kilowatt-hour doubled every 1.57 years over the entire analysis period, a rate of improvement only slightly slower than that for PCs, which saw efficiency double every 1.52 years from 1975 to 2009 (see Figure 4). The data show significant increases in computational efficiency even during the vacuum tube and discrete-transistor eras. From 1946 (ENIAC) to 1958 (when the last of the primarily tube-based computers in our sample came on line), computations per kilowatt-hour doubled every 1.35 years. Computations per kilowatt-hour increased even more rapidly during the shift from tubes to transistors, but the pace of change slowed during the era of discrete transistors.

In the recent years for which we have more than a few data points (2001, 2004, 2008, and 2009), there is a factor of two or three separating the lowest and highest estimates of computations per kilowatt-hour, which indicates substantial variation in the data in any given year. This variation is partly the result of including different types of computers in the sample (desktops, servers, laptops, and supercomputers), but the differences tend to be swamped by the rapid increase in performance per computer over time, which drives the results.

Explaining These Trends

Even current computing technology is far from the minimum theoretically possible energy used per computation.[12] In 1985, the physicist Richard Feynman analyzed the electricity needed for computers that use electrons for switching and estimated that there was a factor of 10^{11} improvement that was theoretically possible compared to computer technology at that time.[13] Since then, performance per kilowatt-hour for computer systems has improved by a factor of 4×10^4 based on our regressions, but there is still a long way to go with current technology before reaching the theoretical limits—and that doesn't even consider the possibility of new methods of computation such as optical or quantum computing.

For vacuum-tube computers, both computational speed and reliability issues encouraged

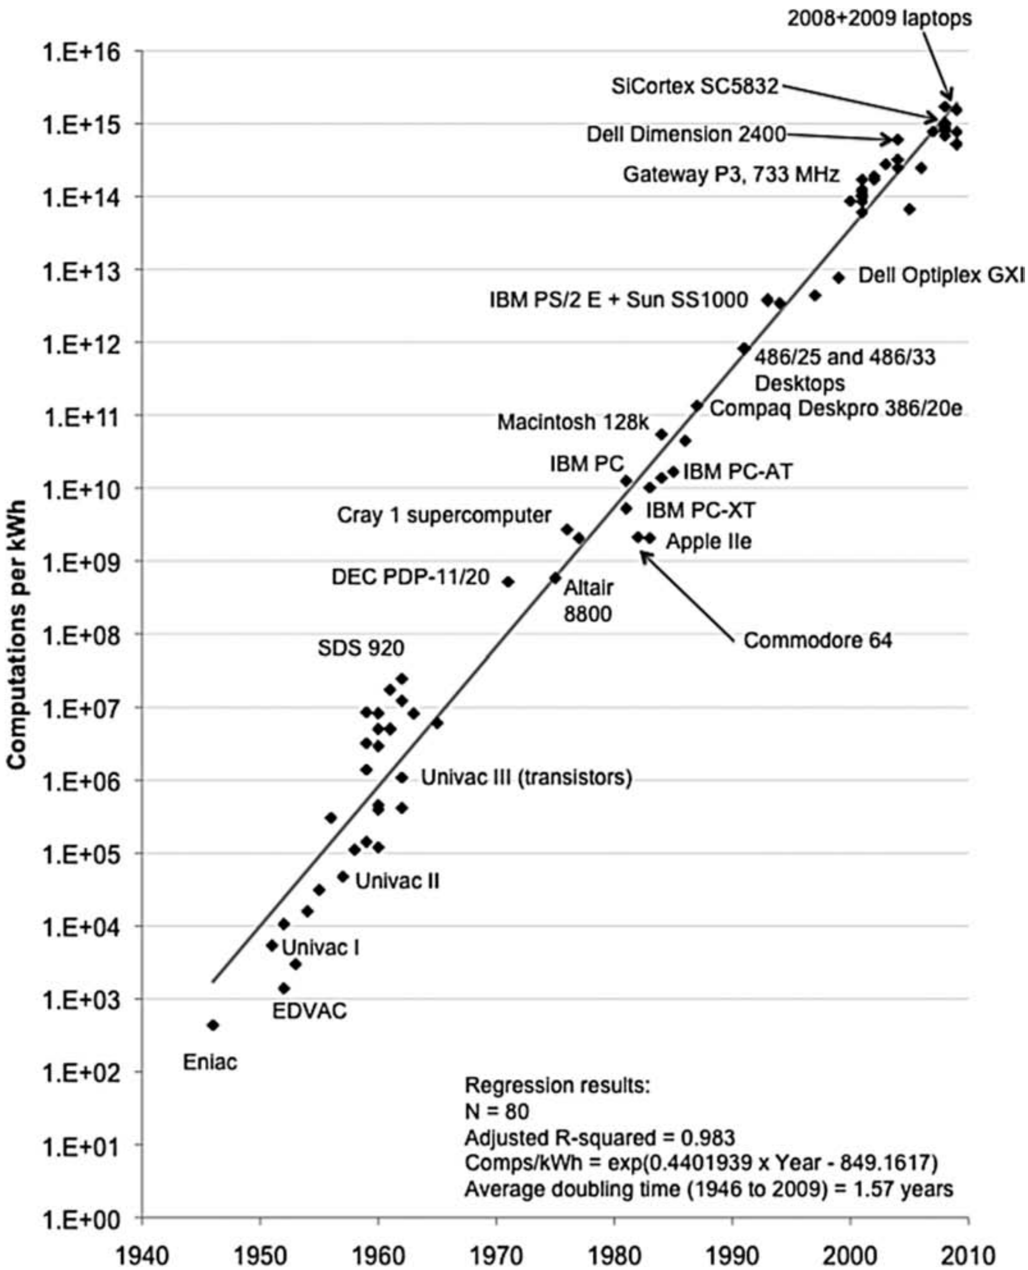

Figure 3. Computations per kilowatt-hour over time. These data include a range of computers, from PCs to mainframe computers and measure computing efficiency at peak performance. Efficiency doubled every 1.57 years from 1946 to 2009.

computer designers to reduce power use. Heat reduces reliability, which was a major issue for tube-based computers. In addition, increasing computation speeds went hand in hand with technological changes (such as reduced capacitive loading, lower currents, and smaller tubes) that also reduced power use. And the economics of operating a tube-based computer led to pressure to reduce power use, although this was probably a secondary issue in the early days of electronic computing.

For transistorized and microprocessor based computers, the driving factor for power reductions was (and is) the push to reduce the physical dimensions of transistors, which reduces the cost per transistor. To accomplish this goal, power used per transistor also must be reduced; otherwise the power densities on the silicon rapidly become unmanageable.

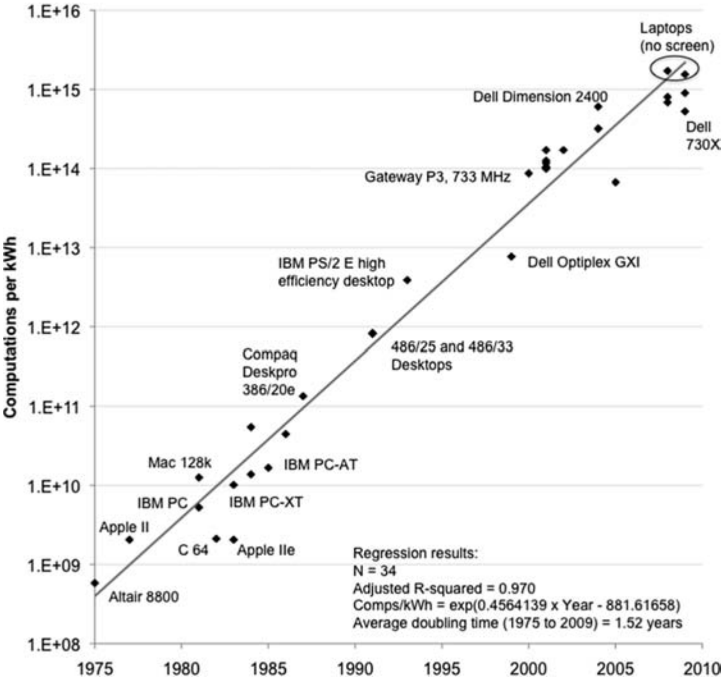

Figure 4. Computations per kilowatt-hour over time for personal computers alone. Efficiency doubled every 1.52 years from 1975 to 2009.

Per-transistor power use is directly proportional to the length of the transistor between the source and drain, the ratio of the transistor length to the electrons' mean free path between collisions, and the total number of electrons in the operating transistor, as Feynman pointed out.[13] Shrinking transistor size therefore resulted in improved speed, reduced cost, and reduced power use per transistor.[14]

Power use is driven by more than just transistor size, however. Computer systems include losses in power supplies and electricity used by disk drives, network cards, and other components. And the energy efficiency associated with these components does not necessarily improve at rates comparable to the trends identified in this article. More research is needed to understand the relative contributions of these different components to progress in the electrical efficiency of computer systems as a whole.

Historical and Future Implications

These trends have been critical for the historical development of mobile computing. They also have implications for the total power used by computers over time and for the availability and ubiquity of battery powered mobile computing devices in the future.

Historical Development of Mobile Computing

The trends identified in this research have strongly affected the development of mobile computing technologies because these devices are constrained by battery storage. As computations per kilowatt-hour increase (holding battery capacity constant), more mobile devices became feasible. Performance and efficiency improvements are inextricably linked, and in some sense, mobile computing is the inevitable result of long-term improvements in computations per kilowatt-hour.

The most visible beneficiaries of these trends have been laptop computers, cellular phones, and personal digital assistants. For example, sales of laptop computers (which use significantly less power than desktop machines) exceeded sales of desktops for the first time in 2009, according to IDC data,[15] demonstrating that portable computers are displacing desktop machines in many applications. This development would not have been possible without long-term improvements in computational efficiency because battery technologies have not improved in the past nearly as rapidly as semiconductor technologies.

Total Electricity Used by Computing Equipment

The total electricity used by computers is not just a function of computational efficiency; the total number of computers and the way they are operated also matter. Table 1 shows the total number of PCs in 1980 and 1985, estimated from historical shipments (see http://arstechnica.com/old/content/2005/12/total-share.ars), and for 1996, 2000, and 2008 as estimated by IDC.[16]

The table shows that the installed base of personal computers doubled on average about every three years between 1980 and 2008. Performance growth per computer has just about cancelled out improvements in performance per kilowatt-hour in the PC era (the doubling times are both about 1.5 years), so we would expect total PC electricity use to scale with the number of PCs. However, that simple assessment does not reflect how the technology has evolved in recent years.

First, the metric we analyzed here focuses only on the peak power use and performance of computers—it says nothing about the energy use of computers in other modes (which are the dominant modes of operation for most servers, desktops, and laptops). Servers in typical business applications approach 100% computational load on average for

Table 1. Installed base estimates for desktop and laptop computers (millions of units).*

Form factor	Region	1980	1985	1996	2000	2008
Desktop PC	US			80.4	151.3	194.4
	Western Europe			58.4	92.3	130.9
	Japan			12.4	21.4	30.8
	Asia Pacific (excluding Japan)			34.0	71.1	249.4
	Latin America			10.5	26.6	79.7
	Canada			8.5	16.0	20.8
	Central and Eastern Europe			7.1	13.3	47.6
	Middle East and Africa			4.2	9.6	30.5
	Total			*215.5*	*401.7*	*784.1*
Portable PC	US			14.3	30.9	121.8
	Western Europe			6.4	14.9	103.4
	Japan			5.9	17.2	38.3
	Asia Pacific (excluding Japan)			2.8	7.3	78.1
	Latin America			0.6	1.5	18.7
	Canada			0.8	2.8	12.6
	Central and Eastern Europe			0.4	0.8	25.7
	Middle East and Africa			0.4	1.1	15.5
	Total			*31.6*	*76.5*	*414.2*
Grand total		**2.1**	**23.1**	**247.1**	**478.2**	**1198.3**
Index (1980 = 1)		*1.00*	*11*	*117*	*227*	*568*
Average annual growth since 1980			61%	35%	31%	25%
Doubling time since 1980 (years)			1.45	2.33	2.56	3.06
Index (1985 = 1)			*1.00*	*10.7*	*20.7*	*52.0*
Average annual growth since 1985				24%	22%	19%
Doubling time since 1985 (years)				3.21	3.43	4.03

* The data for 1996 to 2008 originated from David Daoud at IDC.[16] The installed base in 1980 and 1985 is based on historical shipments data from http://arstechnica.com/old/content/2005/12/total-share.ars and an assumed CPU lifetime of five years, which is comparable to IDC's assumptions.

only 5 to 15% of the time,[17] and desktop and laptop machines usually have even lower utilization.

Second, laptop computers (which typically use one-third to one-fifth of the power of a comparable desktop, as shown in Table S2 in the Web Extra appendix) have started to displace desktops in many applications. That trend is confirmed by the data in Table 1. Liquid crystal display (LCD) screens, which use about one-third of the power of comparable cathode-ray tube (CRT) monitors, have largely displaced CRTs for desktop computers since 2000. More recently, LCD screens have seen significant efficiency improvements with the advent of LED backlighting.

Finally, the EPA's Energy Star program for office equipment has had a substantial impact on the electricity used by this equipment since its inception in the early 1990s,[18,19] particularly when computers are idle (which is most of the time). The program has promoted the use of low-power innovations in desktop machines that were originally developed for laptops.

A complete analysis of electricity used by computing over time would tally installed base estimates for all types of computers and correlate those numbers with measured power use and operating characteristics for each computer type over all their operating modes, including the low-power modes promoted by Energy Star.

Implications for the Future

The computer industry has been able to sustain rapid improvements in computations per kilowatt-hour over the past 65 years, and we fully expect those improvements to continue in coming years. This research suggests that doubling of computations per kilowatt-hour about every year and a half is the long-term industry trend. Because of the large remaining potential for efficiency, we believe that achieving faster rates of improvement is within our grasp if we make efficiency a priority and focus our efforts on a holistic compute system approach, constantly revisiting the notion of what Amory Lovins of Rocky Mountain Institute calls "clean slate, whole system redesign."

Whether performance per CPU can grow for many years more at the historical pace is an ongoing debate in the computer industry,[20] but near-term improvements, such as 3D transistors, are already "in the pipeline." At this juncture, continuing the historical trends in performance (or surpassing them) is dependent on significant new innovation comparable in scale to the shift from single core to multi-core computing. Such innovation will also require substantial changes in software design,[21] which is a relatively new development for the IT industry and is another reason why whole-system redesign is so critical to success.

The most important future effect of these trends is that the power needed to perform a task requiring a fixed number of computations will fall by half every 1.5 years, enabling mobile devices performing such tasks to become smaller and less power consuming and making many more mobile computing applications feasible. Alternatively, the performance of mobile devices could continue to double every 1.5 years while maintaining the same battery life (assuming battery capacity doesn't improve).

These two scenarios define the range of possibilities. Some applications (such as laptop computers) will likely tend toward the latter scenario, while others (such as mobile sensors and controls) will take advantage of increased efficiency to become less power hungry and more ubiquitous.

Conclusions

The performance of electronic computers has shown remarkably steady growth over the past 65 years, a finding that is not surprising to anyone with even a passing familiarity with computing technology. In the personal computer era, performance per computer has doubled approximately every 1.5 years, a rate that corresponds with the popular interpretation of Moore's law. What most observers do not know, however, is that the electrical efficiency of computing (the number of computations that can be completed per kilowatt-hour of electricity) also doubled about every 1.5 years over that period.

Remarkably, the average rate of improvement in the electrical efficiency of computing from ENIAC through 2009 (doubling approximately every 1.6 years) is comparable to improvements in the PC era alone. This counterintuitive finding results from significant increases in power efficiency during the tube computing era and the transition period from tubes to transistors, with somewhat slower growth during the discrete-transistor era.

The main trend driving increased performance and reduced costs in recent decades, namely smaller transistor size, also tends to reduce electricity use, which explains why the industry has been able to improve computational performance and electrical efficiency at similar rates. Similarly, reduced capacitive loading, lower currents, and smaller tubes helped vacuum-tube computers significantly improve their energy efficiency over time. The existence of laptop computers, cellular phones, and personal digital assistants was enabled by these trends, which if they continue, presage continuing rapid reductions in the power consumed by mobile computing devices, accompanied by new and varied applications for mobile computing, sensors, and controls.

Acknowledgments

This report was produced with grants from Microsoft and Intel and independent review comments from experts throughout the industry. All errors and omissions are the responsibility of the authors alone. We thank Rob Bernard of Microsoft and Lorie Wigle of Intel for their financial support of this project. We also give special thanks to Ed Thelen of the Computer History Society, who endured innumerable questions about early computers and was remarkably patient in sharing his knowledge and insights (he also put us in touch with Bill Newman, Robert B. Garner, and Rick Dill, who educated us about the power use of early transistors compared to tubes). We would like to thank Kaoru Kawamoto for supplying recent measured data for PC active power; John Goodhue at SiCortex for supplying and explaining data on his company's largest supercomputer;

Wu-chun Feng at Virginia Tech for supplying data on the Green Destiny supercomputer; Mark Monroe of Sun Microsystems for supplying power data on three older Sun servers; Gordon Bell of Microsoft for supplying power data on some early DEC machines; Saul Griffith and Jim McBride for useful discussions about information theory and power use; Jay Young of Cray for measured power use of the Cray CS6400; Erik Klein (collector of historical computers) for allowing Koomey to measure the power use of his Altair 8800, Apple II, and Osborne 1 machines; Mary Ann Piette, Richard E. Brown, and Bruce Nordman at Lawrence Berkeley National Laboratory (LBNL) for digging through their archives for relevant materials; and Erik Brynjolfsson at MIT for providing helpful guidance about the regressions. Other colleagues at LBNL allowed us to meter their computers, and for that access we are also grateful. We would also like to thank Bill Nordhaus of Yale for his superbly documented historical analysis of computing performance and for his help in understanding its subtleties. Finally, we would like to thank David Daoud and Tom Mainelli of IDC for generously sharing their data on sales and installed base for PCs and sales of monitors, and Vernon Turner of IDC for facilitating that data sharing.

References and Notes

1. T.R. Kennedy, ''Electronic Computer Flashes Answers, May Speed Engineering,'' *New York Times,* 15 Feb. 1946, p. 1+.
2. D.E. Liddle, ''The Wider Impact of Moore's Law,'' *IEEE Solid-State Circuits Newsletter,* Sept. 2006, pp. 28–30.
3. G.E. Moore, ''Cramming More Components onto Integrated Circuits,'' *Electronics,* vol. 38, no. 8, 1965; http://download.intel.com/research/silicon/moorespaper.pdf.
4. T.R. Halfhill, ''The Mythology of Moore's Law: Why Such a Widely Misunderstood 'Law' Is So Captivating to So Many,'' *IEEE SSCS Newsletter,* Sept. 2006, pp. 21–25.
5. E. Mollick, ''Establishing Moore's Law,'' *IEEE Annals of the History of Computing,* vol. 28, no. 3, 2006, pp. 62–75.
6. R. Kumar, ''The Business of Scaling,'' *IEEE Solid-State Circuits Newsletter,* vol. 12, no. 1, 2007, pp. 22–27.
7. G.E. Moore, ''Progress in Digital Integrated Electronics,'' *IEDM Tech. Digest,* IEEE Press, 1975, pp. 11–13.
8. W. Feng, ''Making the Case for Efficient Supercomputing,'' *ACM Queue,* Oct. 2003, pp. 54–64.
9. W.D. Nordhaus, ''Two Centuries of Productivity Growth in Computing,'' *J. Economic History,* vol. 67, no. 1, 2007, pp. 128–159; http://nordhaus.econ.yale.edu/recent_stuff.html.
10. All doubling times in the text are derived from the regression analyses described and documented in the Web Extra appendix and Table S1. See http://doi.ieeecomputersociety.org/10.1109/MAHC.2010.28.
11. For a broad discussion of the evolution of computer classes over time, see G. Bell, *Bell's Law for the Birth and Death of Computer Classes: A Theory of the Computer's Evolution,* tech. report MSR-TR-2007-146, Microsoft Research, 13 Nov. 2007; http://research.microsoft.com/apps/pubs/default.aspx?id=64155.
12. S. Lloyd, ''Ultimate Physical Limits to Computation,'' *Nature,* vol. 406, no. 6799, 2000, pp. 1047–1054.
13. R.P. Feynman, *The Pleasure of Finding Things Out: The Best Short Works of Richard P. Feynman,* Penguin Books, 2001.
14. See also M. Bohr, ''A 30 Year Retrospective on Dennard's MOSFET Scaling Paper,'' *IEEE Solid-State Circuits Newsletter,* vol. 12, no. 1, 2007, pp. 11–13, and Carver Mead's thinking in the late 1960s, as summarized in D.C. Brock, ed., *Understanding Moore's Law: Four Decades of Innovation,* Chemical Heritage Press, 2006, pp. 98–100.
15. T. Mainelli, personal communication with J. Koomey, 14 Apr. 2009.
16. D. Daoud, personal communication with J. Koomey, 31 Mar. 2009.
17. J.M. Kaplan, W. Forrest, and N. Kindler, *Revolutionizing Data Center Efficiency,* McKinsey & Co. report, 2008; www.mckinsey.com/clientservice/bto/pointofview/Revolutionizing.asp.
18. B.J. Johnson and C.R. Zoi, ''EPA Energy Star Computers: The Next Generation of Office Equipment,'' *Proc. 1992 ACEEE Summer Study on Energy Efficiency in Buildings,* Am. Council for an Energy Efficient Economy, 1992, pp. 6.107–6.114.
19. M. Sanchez, G. Homan, and R. Brown, *Calendar Year 2007 Program Benefits for Energy Star Labeled Products,* tech report LBNL-1217E, Lawrence Berkeley National Lab., Oct. 2008.
20. K. Asanovíc et al., *The Landscape of Parallel Computing Research: A View from Berkeley,* tech. report UCB/EECS-2006-183, Univ. of California at Berkeley, 18 Dec. 2006; www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf.

21. M. Bohr, "A 30 Year Retrospective on Dennard's MOSFET Scaling Paper," *IEEE Solid-State Circuits Newsletter,* vol. 12, no. 1, 2007, pp. 11–13.

Jonathan Koomey is a consulting professor at Stanford University. He is one of the leading international experts on the economics of reducing greenhouse gas emissions and the effects of information technology on resource use. He is also the author of *Turning Numbers into Knowledge: Mastering the Art of Problem Solving*, 2nd edition (Analytics Press, 2008). Koomey has a PhD from the Energy and Resources Group at the University of California at Berkeley. Contact him at jgkoomey@stanford.edu.

Stephen Berard is a senior program manager in the Windows kernel group at Microsoft, where he was responsible for energy efficiency and power management features in Windows. His primary focus is on reducing energy needs and operating costs through innovations in Windows power management. Berard has a BS in computer science from the University of Rhode Island. Contact him at sberard@microsoft.com.

Marla Sanchez is a PhD candidate in engineering and public policy at Carnegie Mellon University. Her research areas include assessing potential for environmental life-cycle reductions in products, assessing environmental impacts and policy strategies for the information sector, and advancing standards and procedures to support use-phase efficiency and life-cycle carbon reporting efforts. Sanchez has a MS in energy and resources from the University of California at Berkeley. Contact her at mcsanche@cmu.edu.

Henry Wong is a senior staff platform technologist at Intel. He has more than 16 years of industry experience in digital and mixed signal processor development and more than five years of experience in technology development and enablement of high-efficiency and high-reliability power-conversion techniques for Intel's Itanium and Xeon processor platforms. Wong has a BS in semiconductor physics and econometric modeling from Yale University. Contact him at henry.l.wong@intel.com.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.
MEMBERSHIP: Members receive the monthly magazine *Computer,* discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.
COMPUTER SOCIETY WEBSITE: www.computer.org

Next Board Meeting: 13–14 Nov., New Brunswick, NJ, USA

EXECUTIVE COMMITTEE
President: Sorel Reisman*
President-Elect: John W. Walz;* **Past President:** James D. Isaak;* **VP, Standards Activities:** Roger U. Fujii;† **Secretary:** Jon Rokne (2nd VP);* **VP, Educational Activities:** Elizabeth L. Burd;* **VP, Member & Geographic Activities:** Rangachar Kasturi;† **VP, Publications:** David Alan Grier (1st VP);* **VP, Professional Activities:** Paul K. Joannou;* **VP, Technical & Conference Activities:** Paul R. Croll;† **Treasurer:** James W. Moore, CSDP;* **2011–2012 IEEE Division VIII Director:** Susan K. (Kathy) Land, CSDP;† **2010–2011 IEEE Division V Director:** Michael R. Williams;† **2011 IEEE Division Director V Director-Elect:** James W. Moore, CSDP*
**voting member of the Board of Governors* *†nonvoting member of the Board of Governors*

BOARD OF GOVERNORS
Term Expiring 2011: Elisa Bertino, Jose Castillo-Velázquez, George V. Cybenko, Ann DeMarle, David S. Ebert, Hironori Kasahara, Steven L. Tanimoto
Term Expiring 2012: Elizabeth L. Burd, Thomas M. Conte, Frank E. Ferrante, Jean-Luc Gaudiot, Paul K. Joannou, Luis Kun, James W. Moore
Term Expiring 2013: Pierre Bourque, Dennis J. Frailey, Atsuhiro Goto, André Ivanov, Dejan S. Milojicic, Jane Chu Prey, Charlene (Chuck) Walrad

EXECUTIVE STAFF
Executive Director: Angela R. Burgess; **Associate Executive Director, Director, Governance:** Anne Marie Kelly; **Director, Finance & Accounting:** John Miller; **Director, Information Technology & Services:** Ray Kahn; **Director, Membership Development:** Violet S. Doan; **Director, Products & Services:** Evan Butterfield; **Director, Sales & Marketing:** Dick Price

COMPUTER SOCIETY OFFICES
Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928
Phone: +1 202 371 0101 • **Fax:** +1 202 728 9614
Email: hq.ofc@computer.org
Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314
Phone: +1 714 821 8380 • **Email:** help@computer.org

MEMBERSHIP & PUBLICATION ORDERS
Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641 • **Email:** help@computer.org
Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan
Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553
Email: tokyo.ofc@computer.org

IEEE OFFICERS
President: Moshe Kam; **President-Elect:** Gordon W. Day; **Past President:** Pedro A. Ray; **Secretary:** Roger D. Pollard; **Treasurer:** Harold L. Flescher; **President, Standards Association Board of Governors:** Steven M. Mills; **VP, Educational Activities:** Tariq S. Durrani; **VP, Membership & Geographic Activities:** Howard E. Michel; **VP, Publication Services & Products:** David A. Hodges; **VP, Technical Activities:** Donna L. Hudson; **IEEE Division V Director:** Michael R. Williams; **IEEE Division VIII Director:** Susan K. (Kathy) Land, CSDP; **President, IEEE-USA:** Ronald G. Jensen

revised 18 July 2011

IEEE