

# **ASSESSING TRENDS IN THE ELECTRICAL EFFICIENCY OF COMPUTATION OVER TIME**

Jonathan G. Koomey\*, Stephen Berard†, Marla Sanchez††, Henry Wong\*\*

\* Lawrence Berkeley National Laboratory and Stanford University

†Microsoft Corporation

††Lawrence Berkeley National Laboratory

\*\*Intel Corporation

Contact: JGKoomey@stanford.edu, <http://www.koomey.com>

Final report to Microsoft Corporation and Intel Corporation

Submitted to IEEE Annals of the History of Computing: August 5, 2009

Released on the web: August 17, 2009




## EXECUTIVE SUMMARY

Information technology (IT) has captured the popular imagination, in part because of the tangible benefits IT brings, but also because the underlying technological trends proceed at easily measurable, remarkably predictable, and unusually rapid rates. 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”.

This article explores the relationship between the performance of computers and the electricity needed to deliver that performance. As shown in **Figure ES-1**, computations per kWh grew about as fast as performance for desktop computers starting in 1981, doubling every 1.5 years, a pace of change in computational efficiency comparable to that from 1946 to the present. Computations per kWh grew even more rapidly during the vacuum tube computing era and during the transition from tubes to transistors but more slowly during the era of discrete transistors. As expected, the transition from tubes to transistors shows a large jump in computations per kWh.

In 1985, the physicist Richard Feynman identified a factor of one hundred billion ( $10^{11}$ ) possible theoretical improvement in the electricity used per computation. Since that time computations per kWh have increased by less than five orders of magnitude, leaving significant headroom for continued improvements. The main trend driving towards increased performance and reduced costs, namely smaller transistor size, also tends to reduce power use, which explains why the industry has been able to improve computational performance and electrical efficiency at similar rates. If these trends continue (and we have every reason to believe they will for at least the next five to ten years), this research points towards continuing rapid reductions in the size and power use of mobile computing devices.



**Figure ES-1: Computations per kilowatt hour over time**

# ASSESSING TRENDS IN THE ELECTRICAL EFFICIENCY OF COMPUTATION OVER TIME

Jonathan G. Koomey, Stephen Berard, Marla Sanchez, Henry Wong

## *INTRODUCTION*

February 14, 1946 was a pivotal day in human history. It was on that day that the U.S. War Department announced the existence of world's first general purpose electronic computer (Kennedy 1946). The computational engine of the Electronic Numerical Integrator and Computer (ENIAC) 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 towards ever-greater and cheaper computational power that even to this day proceeds at easily measurable, remarkably predictable, and unusually rapid rates. 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 S1). No other technology to our knowledge has improved as rapidly and over so long a period as IT.

Moore’s Law has seen several incarnations, some more accurate than others. It is not a physical law, but an “empirical observation” (Liddle 2006) that describes economic trends in chip production. As Moore put it in his original article (Moore 1965), “The complexity [of integrated circuits] for minimum component costs has increased at a rate of roughly a factor of two per year”, where complexity is defined as the number of components (not just transistors) per chip. The trend 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 (Kumar 2007). As semiconductor technology improves the cost curve shifts down, making increased component densities cheaper (Figure S2).

In 1975, Moore modified his observation to a doubling of complexity every two years (Moore 1975), which reflected a change in the economics and technology of chip production at that time. That rate of increase in chip complexity has held for about three decades since, which is a reflection mainly of the underlying characteristics of semiconductor manufacturing during that period. There is also a self-fulfilling aspect of Moore’s law, as summarized by Mollick (2006)—the industry’s engineers have used Moore’s law as a benchmark to which they calibrated their rate of innovation.

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 is a correct statement for the microprocessor era, but is one that Moore never made.

Another is “Moore’s law for power”, coined by Feng (2003) 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”).

This article explores 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. More specifically, it estimates how many calculations historical and current computers were (or are) able to complete per kilowatt-hour of electricity consumed, which is one way to measure the electrical efficiency of computation over time. We show data on these trends going back all the way to ENIAC. Of course, ENIAC was a very different device from a modern personal computer, and we must therefore use care in the inferences we draw. To avoid inconsistent comparisons, we rely on long-term performance trends developed in a consistent fashion, normalized per kWh of measured electricity use for each computer.

### ***CALCULATING COMPUTATIONS PER KWH***

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 simply do not exist, so we compiled available data in a consistent way to piece together the long-term trends.

To estimate computations per kWh we focused on the full load computational capacity and the active power for each machine, dividing the number of computations possible per hour at full load by the number of kWh 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 electrical efficiency of this technology, and it is one that can show how the technology has changed over time.

Measuring computing performance has always been controversial, and this article will not settle those issues. The most sophisticated and comprehensive historical analysis of computing performance over time is the work by Nordhaus (2007), which builds on the work of McCallum (2002), Moravec (1998), Knight (1963, 1966, 1968), SPEC <<http://www.spec.org>>, and others. We relied on Nordhaus’s benchmark of millions of computations per second (MCPS), to be consistent with his long-term trends. His analysis combined synthetic benchmarks in an attempt to mimic the increase over time in complexity of computing tasks.

Nordhaus estimated performance data for more than 200 different computers in the modern era (since 1946) ranging from the first vacuum tube machines, to the first identifiable personal computer (the Altair 8800), to the Cray 1 supercomputer, to modern day PCs and servers. Whenever possible we attached measured active power data to computers on Nordhaus’s list, but where such data did not exist we did not attempt to estimate power use. Instead, we located or performed power measurements for computers not on Nordhaus’s list and then estimated performance of those machines in MCPS by scaling using other published performance benchmarks, such as theoretical

FLOPS, Composite Theoretical Performance (CTP), or the SPEC benchmarks associated with these machines. Following Nordhaus, we used SPEC benchmarks for scaling when data were available, but only half a dozen machines for which we had power data also had SPEC benchmarks associated with them. As discussed in the supporting appendix, using CTP instead would have resulted in estimates for MCPS that were modestly different (between 8% smaller and 20% larger).

We chose to only include measured power data because the uncertainty associated with estimating power use is much greater than the uncertainties introduced by scaling the performance estimates.<sup>1</sup> The main sources for measured power data were Weik (1955, 1961, 1964) for computers from 1946 through the early 1960s, Russell (1978) for the Cray 1 supercomputer, Roberson et al. (2002), Harris et al. (1988), and Lovins (1993) for PCs, and Koomey et al. (2009) for servers. We also conducted new measurements for recent desktops and two laptops and compiled additional measured data from researchers in the technology industry (these data are presented and described in the supporting appendix, Tables S2 and S3, and in the final complete analysis spreadsheets downloadable at <<http://homepage.mac.com/jgkoomey>>).

## ***RESULTS***

**Figure 1** shows performance per computer for all the computers included in Nordhaus's analysis from 1946 onwards, as well as for the 36 additional machines for which measured power was available that we added for this analysis. It does not include performance estimates for recent large-scale supercomputers (e.g., those at <<http://www.top500.org/>>), but it does include measurements for server models that are often used as computing nodes for those machines. The trends for microprocessor-based computers are clear. The performance per unit for PCs shows a doubling time of 1.45 years from 1981 (the introduction date of the IBM PC) to 2009,<sup>2</sup> which corresponds to the popular interpretation of Moore's law but not its 1975 formulation.

**Figure 2** shows the results in terms of the number of calculations per kWh of electricity consumed for the computers for which both performance and measured power data are available. These data include a wide range of computers, from PCs to mainframe computers.<sup>3</sup> The transition from vacuum tube to transistorized computing is clearly evident in the data. During the years 1959, 1960, and 1961, as transistorized computers came to market in large numbers, there are about two orders of magnitude difference between the most and least electricity intensive computers. Logical gates constructed with discrete transistors use about a factor of ten less power than vacuum tubes, but the

---

<sup>1</sup> This statement is true as long as the computers on Nordhaus's list that we used to scale performance are of a similar type and vintage to the ones that we are adding to the list.

<sup>2</sup> All doubling times in the text are derived from the regression analyses described and documented in the supporting appendix and Table S1.

<sup>3</sup> For a broad discussion of the evolution of computer classes over time, see Bell (2007)

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 kWh doubled every 1.57 years over the entire analysis period, a rate of improvement only slightly slower than that for PCs, which doubled every 1.49 years from 1981 to 2009 (see **Figure 3**). 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 kWh doubled every 1.35 years. Computations per kWh increased even more rapidly during the shift from tubes to transistors, but the pace of change slowed during the era of discrete transistors.

### ***EXPLAINING THESE TRENDS***

Even current computing technology is very far from the minimum theoretically possible energy used per computation (Lloyd 2000). 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 (Feynman 2001). Since then, performance per kWh for computer systems has improved by a factor of  $4 \times 10^4$ , 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 like optical or quantum computing).

For vacuum tube computers, both computational speed and reliability issues encouraged 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 (like reduced capacitive loading, lower currents, and smaller tubes) that also reduced power use. And the simple economics of operating a tube-based computer led to pressure to reduce power use, although this issue was probably a secondary one 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. In order to accomplish this goal, power used per transistor also must be reduced; otherwise the power densities on the silicon rapidly become unmanageable. The power use of a CPU is directly proportional to the length of the transistor between source and drain, the ratio of transistor length to mean free path of the electrons, and the total number of electrons in the operating transistor, as Feynman (2001) pointed out. Shrinking transistor size therefore resulted in improved speed, reduced cost, and reduced power use per transistor (see also Bohr (2007)).

Power use is driven by more than just the microprocessor, however. Computer systems include losses in power supplies and electricity used by disk drives, network cards, and other components, and the power efficiency associated with these components does not necessarily improve at rates driven by Moore's law. 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.

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 the highest estimates of computations per kWh, which indicates substantial variation in the data in any given year. These differences are partly the result of including different types of computers in the sample (desktops, servers, laptops, supercomputers), but they tend to be swamped by the rapid increase in performance per computer over time, which drives the results.

### ***IMPLICATIONS OF THIS RESEARCH***

The computer industry has been able to sustain rapid improvements in computations per kWh over the past sixty years, and we expect those improvements to continue in coming years. This research suggests that doubling of computations per kWh every 1.6 years is the long-term industry trend, but we believe (because of the large remaining potential for efficiency) that achieving faster rates of improvement is within our grasp, if we make efficiency a priority and focus our efforts on 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 (doubling every 1.5 years or so) is an ongoing subject of debate in the computer industry (Bohr 2007), but near-term improvements are already “in the pipeline”. Continuing the historical trends in performance is at this juncture 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 (Asanović et al. 2006), which is a relatively new development for the IT industry and it is another reason why whole system redesign is so critical to success.

The trends identified in this research have important implications for mobile computing technologies because these devices are constrained by battery storage. The power needed to perform a task requiring a fixed number of computations will fall by half every 1.5 years (using the trend from the PC era), 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 batteries don’t improve). These two scenarios define the range of possibilities. Some applications (like laptop computers) will likely tend towards the latter scenario, while others (like mobile sensors) will take advantage of increased efficiency to become less power hungry and more ubiquitous.

Of course, the total electricity used by computers is not just a function of computational efficiency as defined here—the total number of computers and the way they are operated also matter. **Table 1** shows the total number of PCs in 1985, estimated from historical shipments (from <http://arstechnica.com/old/content/2005/12/total-share.ars>), and for 1996, 2000, and 2008 as estimated by IDC (Daoud 2009). That table shows a doubling time for installed base of personal computers of about 4 years from 1985 through 2008.

Performance growth per computer has just about cancelled out improvements in performance per kWh in the PC era (the doubling times are approximately the same), 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 analyzed here focuses only on the peak power use and performance of computers—it says nothing about the power use of computers in other modes (which for most servers, desktops, and laptops are the dominant modes of operation for these machines). Servers in typical business applications approach 100% computational load on average for only 5-15% of the time, and desktop and laptop machines have similarly low utilization numbers.

Second, laptop computers (which typically use one third to one fifth of the power of a comparable desktop, as shown in Table S2) have started to displace desktops in many applications. That trend is confirmed by the data in Table 1. And liquid crystal display (LCD) screens, which use about a third of the power of comparable cathode ray tube (CRT) monitors, have largely displaced CRTs for desktop computers since 2000.

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 (Johnson and Zoi 1992, Sanchez et al. 2008), 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.

## ***CONCLUSIONS***

The performance of electronic computers has shown remarkable and steady growth over the past 60 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.

Performance growth per computer has just about cancelled out improvements in computations per kWh in the PC era, so all other things being equal, PC electricity use should scale with the installed base of PCs (which increased by a factor of more than fifty from 1985 to 2008). All other things are not equal, however. Sales of laptop computers (which use significantly less power than desktop machines) are due (for the first time) to exceed sales of desktops in 2009, according to IDC data. LCD screens, which are two to three times less electricity intensive than the old CRTs, have almost completely displaced CRTs in the marketplace. And the EPA's Energy Star Computers program has had

substantial success in promoting power saving technologies for computers, monitors, and other office equipment.

Remarkably, the average rate of improvement in the electrical efficiency of computing from ENIAC through 2008 (doubling about 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.

In 1985, the physicist Richard Feynman identified a factor of one hundred billion ( $10^{11}$ ) possible theoretical improvement in the electricity used per computation. Since that time computations per kWh have increased by less than five orders of magnitude, leaving significant headroom for continued improvements. The main trend driving towards increased performance and reduced costs, 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. If these trends continue, they presage continuing rapid reductions in the power consumed by mobile computing devices, accompanied by new and varied applications for mobile computing.

## REFERENCES

Asanović, Krste, Rastislav Bodik, Bryan Catanzaro, Joseph Gebis, Parry Husbands, Kurt Keutzer, David Patterson, William Plishker, John Shalf, Samuel Williams, and Katherine Yelick. 2006. *The Landscape of Parallel Computing Research: A View from Berkeley*. December 18. <<http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>>

Bell, Gordon. 2007. *Bell's Law for the birth and death of computer classes: A theory of the computer's evolution*. MSR-TR-2007-146. November 13. <<http://research.microsoft.com/apps/pubs/default.aspx?id=64155>>

Bohr, Mark. 2007. "A 30 Year Retrospective on Dennard's MOSFET Scaling Paper." *IEEE SSCS Newsletter*. vol. 12, no. 1. Winter. pp. 11-13.

Daoud, David. March 31, 2009. *Personal Communication*: "IDC spreadsheet titled "IDC WW PC Tracker\_InstalledBase\_2008Q4\_IDC.xls" sent in email." Contact: P. C. w. J. Koomey.

Feng, Wu-Chun. 2003. "Making the case for efficient supercomputing." In *ACM Queue*. October. pp. 54-64.

Feynman, Richard P. 2001. *The Pleasure of Finding Things Out: The Best Short Works of Richard P. Feynman*. London, UK: Penguin Books.

Harris, Jeff, J. Roturier, L.K. Norford, and A. Rabl. 1988. *Technology Assessment: Electronic Office Equipment*. Lawrence Berkeley Laboratory. LBL-25558. November.

Johnson, Brian J., and Catherine R. Zoi. 1992. "EPA Energy Star Computers: The Next Generation of Office Equipment." In *Proceedings of the 1992 ACEEE Summer Study on Energy Efficiency in Buildings*. Edited by Asilomar, CA: American Council for an Energy Efficient Economy. pp. 6.107-6.114.

Kennedy, T. R. 1946. "Electronic computer flashes answers, may speed engineering." *The New York Times*. New York, NY. February 15. p. 1+.

Knight, Kenneth E. 1963. *A Study of Technological Innovation—The Evolution of Digital Computers*. Thesis, Carnegie Institute of Technology.

Knight, Kenneth E. 1966. "Changes in Computer Performance." *Datamation*. September. pp. 40-54.

Knight, Kenneth E. 1968. "Evolving Computer Performance 1963-67." *Datamation*. January. pp. 31-35.

Koomey, Jonathan G., Christian Belady, Michael Patterson, Anthony Santos, and Klaus-Dieter Lange. 2009. *Assessing trends over time in performance, costs, and energy use for servers*. Oakland, CA: Analytics Press. August 17.

Kumar, Rakesh. 2007. "The Business of Scaling." *IEEE SSCS Newsletter*. vol. 12, no. 1. Winter. pp. 22-27.

Liddle, David E. 2006. "The Wider Impact of Moore's Law." In *IEEE SSCS Newsletter*. September. pp. 28-30.

Lloyd, Seth. 2000. "Ultimate physical limits to computation." *Nature*. vol. 406, no. 6799. pp. 1047-1054.

Lovins, Amory B. 1993. *What an Energy-Efficient Computer Can Do*. Old Snowmass, CO: Rocky Mountain Institute. Aug. 10, 1993 (revised Oct. 10, 1993). <[http://www.rmi.org/images/PDFs/Energy/E93-20\\_EnergyEffComputer.pdf](http://www.rmi.org/images/PDFs/Energy/E93-20_EnergyEffComputer.pdf)>

McCallum, John C. 2002. "Price-Performance of Computer Technology." In *The Computer Engineering Handbook*. Edited by V. G. Oklobdzija. Boca Rotan, FL: CRC Press. pp. 4-1 to 4-18.

Mollick, Ethan. 2006. "Establishing Moore's Law." *IEEE Annals of the History of Computing (Published by the IEEE Computer Society)*. July-September. pp. 62-75.

Moore, Gordon E. 1965. "Cramming more components onto integrated circuits." In *Electronics*. April 19.

Moore, Gordon E. 1975. "Progress in Digital Integrated Electronics." *IEEE, IEDM Tech Digest*. pp. 11-13. <<http://www.ieee.org/>>

Moravec, Hans. 1998. "When will computer hardware match the human brain?" *Journal of Evolution and Technology*. vol. 1, <<http://www.transhumanist.com/volume1/moravec.htm>>

Nordhaus, William D. 2007. "Two Centuries of Productivity Growth in Computing." *The Journal of Economic History*. vol. 67, no. 1. March. pp. 128-159. <[http://nordhaus.econ.yale.edu/recent\\_stuff.html](http://nordhaus.econ.yale.edu/recent_stuff.html)>

Roberson, Judy A., Richard E. Brown, Bruce Nordman, Carrie A. Webber, Gregory K. Homan, Akshay Mahajan, Marla McWhinney, and Jonathan G. Koomey. 2002. *Power levels in office equipment: Measurements of new monitors and personal computers*. Proceedings of the 2002 ACEEE Summer Study on Energy Efficiency in Buildings. Asilomar, CA; Washington, DC: American Council for an Energy Efficient Economy (also LBNL-50508). August.

Russell, Richard M. 1978. "The CRAY- 1 Computer System." *Communications of the ACM*. vol. 21, no. 1. January. pp. 63-72.

Sanchez, Marla, Gregory Homan, and Richard Brown. 2008. *Calendar Year 2007 Program Benefits for ENERGY STAR Labeled Products*. Berkeley, CA: Lawrence Berkeley National Laboratory. LBNL-1217E. October.

Weik, Martin H. 1955. *A Survey of Domestic Electronic Digital Computing Systems*. Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 971. December. <<http://ed-thelen.org/comp-hist/BRL.html>>

Weik, Martin H. 1961. *A Third Survey of Domestic Electronic Digital Computing Systems*. Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 1115. March. <<http://ed-thelen.org/comp-hist/BRL61.html>>

Weik, Martin H. 1964. *A Fourth Survey of Domestic Electronic Digital Computing Systems (Supplement to the Third Survey)* Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 1227. January. <<http://ed-thelen.org/comp-hist/BRL64.html>>

## ***ACKNOWLEDGMENTS***

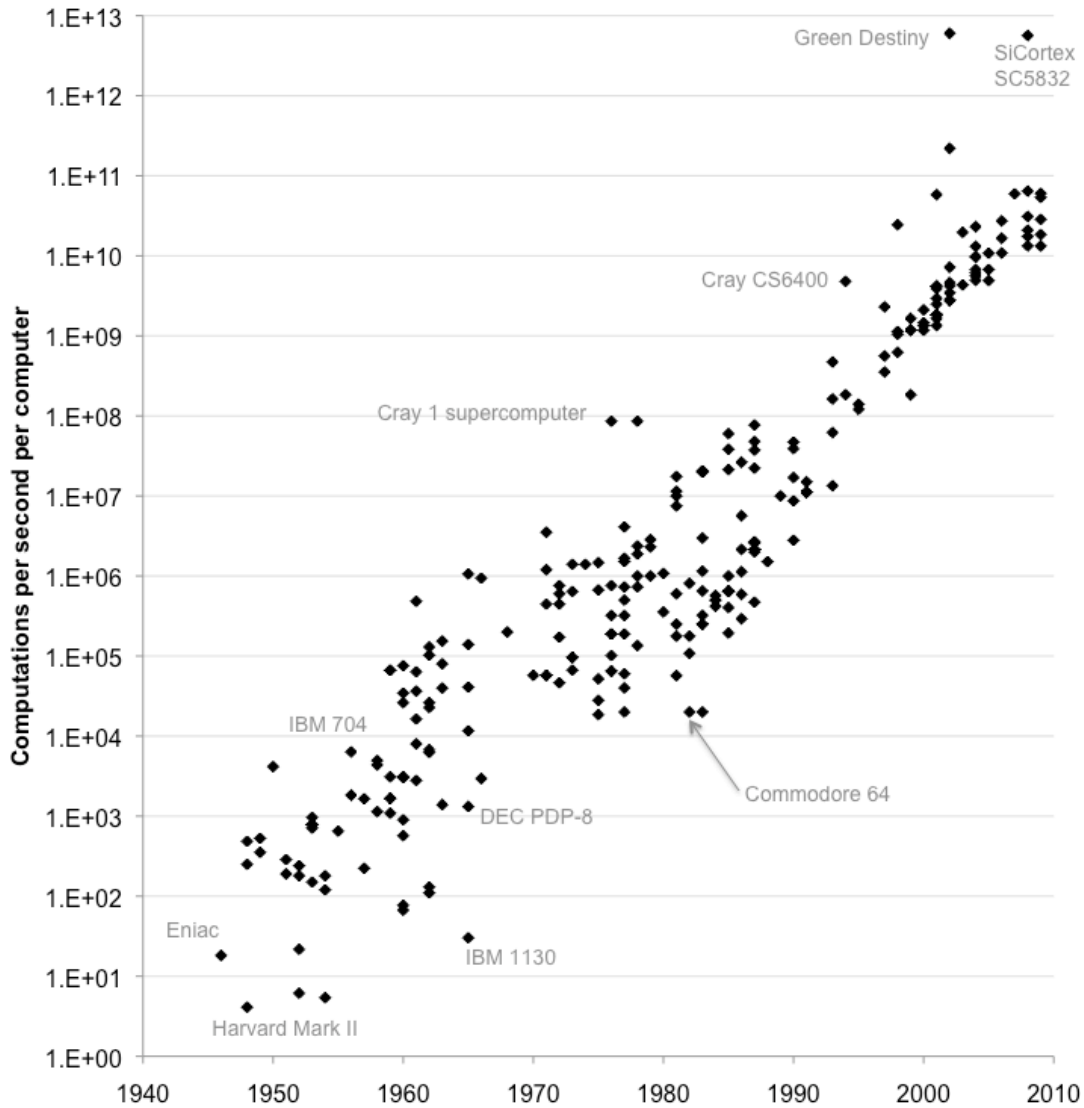
This report was produced with grants from Microsoft Corporation and Intel Corporation and independent review comments from experts throughout the industry. All errors and omissions are the responsibility of the authors alone. The authors would like to thank Rob Bernard of Microsoft Corporation and Lorie Wigle of Intel Corporation for their financial support of this project.

We give our 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.

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, Wuchun 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, and colleagues at LBNL for digging through their archives for relevant materials. 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.


And 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 to Vernon Turner for facilitating that data sharing.




**Figure 1: Computational capacity over time (computations/second per computer).**

Data source: Nordhaus (2007), with additional data added post-1985 for computers not considered in his study. Doubling time for PCs (1981 to 2009) is 1.45 years.





**Figure 2: Computations per kilowatt-hour over time**



**Figure 3: Computations per kilowatt-hour over time for personal computers alone**

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

<i>Form factor</i>	<i>Region</i>	<i>1985</i>	<i>1996</i>	<i>2000</i>	<i>2008</i>	
Desktop PC	USA		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	
	<i>Total</i>			<i>215.5</i>	<i>401.7</i>	<i>784.1</i>
Portable PC	USA		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	
	<i>Total</i>			<i>31.6</i>	<i>76.5</i>	<i>414.2</i>
<i>Grand total</i>			<i>23.1</i>	<i>247.1</i>	<i>478.2</i>	<i>1198.3</i>
Index 1985 = 1		1.00	10.72	20.75	51.99	
Avg annual % growth since 1985			24%	22%	19%	
Doubling time since 1985 (years)			3.21	3.43	4.03	

(1) Source for 1996-2008: IDC, from file "IDC WW PC Tracker\_InstalledBase\_2008Q4\_IDC.xls", supplied by David Daoud of IDC to JK in email, March 31, 2009.

(2) Installed base in 1985 based on historical shipments data from <<http://arstechnica.com/old/content/2005/12/total-share.ars>> and an assumed CPU lifetime of 5 years, which is comparable to IDC's assumptions.

**SUPPORTING APPENDIX:**  
**ASSESSING TRENDS IN THE ELECTRICAL EFFICIENCY OF  
COMPUTATION OVER TIME**


Jonathan G. Koomey, Stephen Berard, Marla Sanchez, Henry Wong

***TWO CONTEXTUAL GRAPHS SUPPORTING STATEMENTS IN THE TEXT***



**Figure S1: Transistor counts for microprocessors over time (thousands)**

The doubling time from 1971 to 2006 is about 1.8 years. Data source: James Larus, Microsoft Corporation.



**Figure S2: Integrated circuit manufacturing cost as a function of components per chip (from Moore (1965))**

### ***DATA AND METHODS***

To avoid inconsistent comparisons, we rely on long-term performance trends developed in a consistent fashion as well as measured electricity use data for historical computers. The data and results from the analysis are summarized in **Table S1**.

Our goal for this analysis is an accurate general overview of trends in compute capabilities and power use over time, and for this purpose, the metric of computations per kWh is a reasonable one. We calculate this metric for dozens of different computers, ranging from laptop PCs to mainframes and integrated supercomputers.

### ***Computations per kWh***

To estimate computations per kWh we focus on the full load computational capacity and the active power for each machine, dividing the number of computations per hour by the number of kWh consumed over that same hour. That requires estimates of peak computational performance and power use while the computer is delivering that performance.

## *Performance*

Measuring computing performance is not an easy task. The most sophisticated and comprehensive historical analysis of computing performance is the work summarized in Nordhaus (2007). For most of the computers analyzed here, we rely on Nordhaus's benchmark of millions of computations per second (MCPS), to be consistent with his long-term trends. Nordhaus estimated performance data for more than 200 different computers in the modern era (since 1946) ranging from the first vacuum tube machines to modern day PCs and servers.

Where possible we attached measured active power data to computers on Nordhaus's list, but when such data did not exist we did not attempt to estimate it. Instead, we located measured power data for computers not on Nordhaus's list and then estimated performance of those machines in MCPS by scaling using other performance benchmarks, such as LINPACK FLOPS, theoretical FLOPS, Composite Theoretical Performance (CTP), or the SPECint\_rate and SPECfp\_rate benchmarks.<sup>1</sup>

We scaled performance for new machines not on Nordhaus's list using Equation A1:

$$\text{Performance}_{\text{new}} (\text{MCPS}) = P_{\text{new}} (\text{FLOPS/CTP/SPEC}) \times P_{\text{ref}} (\text{MCPS}) / P_{\text{ref}} (\text{FLOPS/CTP/SPEC}) \quad (\text{A1})$$

Where  $P_{\text{ref}}$  is the performance of the reference system, expressed in MCPS, FLOPS, CTP, or SPEC, and  $P_{\text{new}}$  is the performance of the new system, expressed in FLOPS, CTP, or SPEC. **Table A-2** shows the data we used to scale performance.

We used three reference systems from Nordhaus's list. The Dell PC Limited 386-16, the Dell Precision Workstation (PW) 420-1000, and the Dell Precision Workstation 690. We used the first to scale performance for the Compaq Deskpro 386 and four other computers through the year 1993, the second to scale performance for 16 machines from 1999 through 2004, and the third to scale performance for 14 machines from 2005 to 2009.

As Nordhaus pointed out in his article, it is always better to use real benchmarks that measure the actual time for computers to perform certain tasks than to rely on theoretical benchmarks. Unfortunately, SPEC benchmarks were only available for half a dozen computers for which we also had measured power data. For those machines, we found that using CTP instead of the SPEC benchmark would have resulted in a MCPS that was 2 to 8% lower in three cases, and from 2% to 20% higher in the other three cases. The differences are not huge, and they go in both directions. Further analysis of the errors introduced by use of theoretical benchmarks would of course be useful, but for purposes of this analysis, we're satisfied that the use of CTP does not bias the results for MCPS significantly.

---

<sup>1</sup> For details on CTP for Intel processors, see <<http://www.intel.com/support/processors/sb/CS-017346.htm>>. For AMD processors, see <[http://www.amd.com/us-en/Processors/ProductInformation/0,,30\\_118\\_8796\\_8800~124990,00.html](http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8800~124990,00.html)>. For SPEC benchmarks, see <<http://www.spec.org>>.

## *Energy use*

Energy use of IT equipment has been a major focus of research for more than two decades (Baer et al. 2002, Blazek et al. 2004, Dandridge 1994, Harris et al. 1988, Kawamoto et al. 2002, Koomey 2008, Koomey et al. 2002, Koomey et al. 2004, Koomey et al. 1996, Lovins and Heede 1990, Mitchell-Jackson et al. 2002, Mitchell-Jackson et al. 2003, Nordman et al. 1996, Norford et al. 1990, Piette et al. 1991, Roth et al. 2002, Roth et al. 2006). The most common error in assessing energy use for computers is to rely on the nameplate power use printed on the computer's power supply, which is generally two to three times larger than typical power use for that device in operation. In this analysis, we rely only on measured power use.

For the early machines, we used data from Weik (1955, 1961, 1964).<sup>2</sup> Where there were multiple estimates for power use of those machines, we took a simple average across installations, as noted in the footnotes to Table A-1. We examined the detailed descriptions for each such machine to make sure we obtained the correct power estimates. There was little variation in power use for these machines as a function of computing load.

Bell (2009) supplied power data for PDP-1, PDP-4, and PDP-8 minicomputers. Russell (1978) gave power use for the Cray 1 supercomputer. Roberson et al. (2002), Harris et al. (1988), Kawamoto (2009), Lovins (1993), Sanchez (2009b), and Ecos Consulting via Sanchez (2009a) supplied measured power data for PCs, and Monroe (2009) and Koomey et al. (2009) supplied measured power data for servers. We also conducted new measurements for some old desktops, as well as very recent desktops and two laptop machines. These recent computers (circa 2008 and 2009) generally use the Intel Core 2 Duo CPU, which was the dominant processor for both desktops and laptops in 2008 (more than one-third of desktops and more than one half of laptops used some form of that CPU in 2008, according to IDC data (Daoud 2009)).

For this analysis, the measured power needed to reflect a maximum load when the CPU was fully utilized. Most of the time the computer is on it is in what we call "active idle", which is typically much lower than the full load power we seek. Almost all available studies measure active idle power. Some studies also measure the maximum power the computer used booting up, and when full load power was not available, we used maximum boot power. Some studies do measure active power, typically by measuring power use while the computer is opening or running software that pushes the computer to maximum CPU utilization.

For a handful of machines circa 2001, we increased active idle measurements by 20 W, which was the median power difference between active idle and full load power use in the 2001 data from Ecos Consulting via Sanchez (2009a). The Ecos study measured

---

<sup>2</sup> These three reports by Weik are all on-line, thanks to Ed Thelen of the Computer History Society <<http://ed-thelen.org/comp-hist/on-line-docs.html>>. We're attempting to locate the 2d edition of this survey so it can be scanned, with no luck as of June 2009.

power use when opening software (which we treat as the full load power measurement), and the new measurements conducted for this study measured power use when 15 browser windows were opened in rapid succession.

For the two laptops currently in the sample (MacBook 2.4 GHz, circa 2008, and a Dell latitude E6400, circa 2009), we used maximum measured power use and subtracted 5 watts for the 13.3" + 14" LCD screens, to make the measurement consistent with those for desktops.

One subtlety in these measurements is the treatment of computers with high-end graphics processing units (GPUs). Many business computers use the CPU itself for graphics processing, but some machines destined for use by gamers or designers have high-powered dedicated graphics boards. GPUs add both electricity use and computational power, but it is difficult to determine how the additional performance should factor in to MCPS. We only measured power used by a few machines with significant dedicated GPUs (the Dell 700XL, custom ASUS PC, and the Dell 730X, all of which are owned by a photography business). For these three machines the electricity used by the GPU is captured in the measured power use, but the GPU is assumed to make no contribution to MCPS. This assumption makes sense as long as the GPU is not used to do other processing beyond that what's needed for display purposes. In special cases GPUs are able to perform actual computing tasks, but this situation has generally been a rare one for typical computers (although that is changing as software becomes more sophisticated in tapping multiple processors of different types within each computing system). In principle, the trends in computations per kWh identified here for general computing technology should also apply to GPUs, but the complexities of how the computational power of GPUs affects *useful* computing power make further research necessary on this topic.

### *Comparing trends over time*

To allow straightforward comparisons, we use the metric of doubling time, defined as the number of years it takes for a parameter (performance per watt, for example) to double. We first calculate the instantaneous growth rate  $g$  as in Equation A2<sup>3</sup>:

---

<sup>3</sup> It is more common in most situations to use simple growth rates, calculated as  $g = \left( \frac{Y_t}{Y_o} \right)^{\left( \frac{1}{t} \right)} - 1$

but this method gives erroneous answers for growth rates higher than about 10% per year. For the high growth rates common to information technology equipment, instantaneous growth rates are more appropriate and accurate (Nordhaus 2007). The instantaneous growth formula is derived from the equation  $Y_t = Y_o e^{gt}$ . To convert a simple annual percentage growth rate (P) to a continuously compounded instantaneous rate, take the natural logarithm of (1+P). We are indebted to Philip Sternberg of IBM for helping to sort out the subtleties of these growth calculations.



$$g = \frac{\text{LN} \left( \frac{Y_t}{Y_o} \right)}{t} \quad (\text{A2})$$

where

$Y_t$  is some quantity at time  $t$ ,

$Y_o$  is that quantity at time 0

and  $t$  is the time over which growth occurs, measured in this case in years (from year 0 to year  $t$ ).

Instantaneous growth rates assume continuous compounding, which is necessary when dealing with the rapid growth rates common in computer technology. An instantaneous growth rate of 69.3% implies a doubling every year.

We can then calculate the doubling time using Equation A3:

$$\text{Doubling time} = \frac{\text{LN} ( 2 )}{g} \quad (\text{A3})$$

Using the doubling time allows us to compare the trends in servers to another important parameter popularly reported in this fashion (Moore's law), which in its most precise form states that the number of transistors on a chip doubles roughly every two years.

### ***Regression Analysis***

In order to derive trends in the data (from which we derived doubling times) we took the natural log of computations per kWh or computations per PC and then used Excel 2004's analysis tool pack to do a linear regression on the log of these parameters, with time as the independent variable. We performed these regressions for different time periods and sets of the data in an exploratory fashion. The doubling time results summarized in the main text were the data most relevant to understanding the underlying trends.

Once the regression parameters were determined we then plotted the lines on our scatter charts (Figures 2 and 3).

The regression parameter associated with the slope of the line in a semi-log regression turns out to be equivalent to the instantaneous growth rate defined above. This result follows from the regression procedure, which yields the parameters  $m$  and  $B$  in the following equation.

$$Y = m \times T + B$$

Where  $Y = \text{LN}(\text{computations per kWh})$ ,

T is time in years,

m is the slope, and

B is the y-intercept of the regression line.

To calculate computations per kWh, raise e to the power  $mT+B$ , as follows:

$$\text{Computations per kWh} = e^{(mT+B)}$$

Since  $e^B$  is a constant, the equation is the same as that for the standard one for instantaneous exponential growth, with m equal to the instantaneous growth rate.

**TABLE S1: REGRESSION RESULTS**

**SUMMARY OUTPUT**      1946 to 2009

<i>Regression Statistics</i>	
Multiple R	0.991510108
R Square	0.983092294
Adjusted R Square	0.982863812
Standard Error	1.233562466
Observations	76

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	6547.320819	6547.320819	4302.702602	2.56209E-67
Residual	74	112.6040504	1.521676357		
Total	75	6659.92487			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	849.2593751	13.30876743	-63.81202314	1.90131E-66	875.7776726	822.7410775	875.7776726	822.7410775
X Variable 1	0.440243485	0.006711541	65.59498915	2.56209E-67	0.426870448	0.453616523	0.426870448	0.453616523

Doubling time (yrs)      1.57

**SUMMARY OUTPUT**

1946 to 1958

<i>Regression Statistics</i>	
Multiple R	0.882228676
R Square	0.778327436
Adjusted R Square	0.750618366
Standard Error	1.005603474
Observations	10

**ANOVA**

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	28.40494239	28.40494239	28.08926547	0.00072845
Residual	8	8.089906777	1.011238347		
Total	9	36.49484916			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	990.5149397	188.6704669	5.249973437	0.00077386	1425.589816	555.4400632	1425.589816	555.4400632
X Variable 1	0.511896683	0.096585543	5.299930704	0.00072845	0.289170022	0.734623343	0.289170022	0.734623343

Doubling time (yrs) 1.35

**SUMMARY OUTPUT**

1958 to 1963

<i>Regression Statistics</i>	
Multiple R	0.456449794
R Square	0.208346414
Adjusted R Square	0.161778556
Standard Error	1.590884274
Observations	19

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	11.32340251	11.32340251	4.474039027	0.049481231
Residual	17	43.02551712	2.530912772		
Total	18	54.34891963			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	-1140.29999	545.9702758	2.088575221	0.052098202	2292.196572	11.59659287	2292.196572	11.59659287
X Variable 1	0.589073422	0.278496374	2.115192433	0.049481231	0.00149744	1.176649404	0.00149744	1.176649404

Doubling time  
(yrs) 1.18

**SUMMARY OUTPUT**

1963 to 1981

<i>Regression Statistics</i>	
Multiple R	0.979757381
R Square	0.959924526
Adjusted R Square	0.946566035
Standard Error	0.792160162
Observations	5

**ANOVA**

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	45.09264135	45.09264135	71.85875347	0.003446758
Residual	3	1.882553169	0.627517723		
Total	4	46.97519452			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	863.5317902	104.1473077	8.291446117	0.003675302	1194.975005	532.0885756	1194.975005	532.0885756
X Variable 1	0.447872796	0.052834165	8.476954257	0.003446758	0.279730904	0.616014688	0.279730904	0.616014688

Doubling time  
(yrs) 1.55

**SUMMARY OUTPUT**

1981 to 2009

<i>Regression Statistics</i>	
Multiple R	0.980110249
R Square	0.960616099
Adjusted R Square	0.959700195
Standard Error	0.836507467
Observations	45

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	733.9039669	733.9039669	1048.816693	7.76898E-32
Residual	43	30.08902394	0.699744743		
Total	44	763.9929908			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	873.2425534	27.91367945	31.28367778	3.24341E-31	929.5358523	816.9492545	929.5358523	816.9492545
X Variable 1	0.452189827	0.013962751	32.38543951	7.76898E-32	0.424031257	0.480348398	0.424031257	0.480348398

Doubling time  
(yrs) 1.53

**SUMMARY OUTPUT**

1981 to 2009, PCs only

<i>Regression Statistics</i>	
Multiple R	0.981268116
R Square	0.962887115
Adjusted R Square	0.961607361
Standard Error	0.896456372
Observations	31

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	604.6540649	604.6540649	752.3997802	2.70882E-22
Residual	29	23.30538677	0.803634026		
Total	30	627.9594517			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	897.5524378	33.80797089	26.54854504	6.74305E-22	966.6975009	828.4073746	966.6975009	828.4073746
X Variable 1	0.464381036	0.016929734	27.42990668	2.70882E-22	0.429755842	0.499006231	0.429755842	0.499006231

Doubling time  
(yrs) 1.49



**SUMMARY OUTPUT**

1946 to 1981

<i>Regression Statistics</i>	
Multiple R	0.928542513
R Square	0.862191198
Adjusted R Square	0.857597571
Standard Error	1.488243436
Observations	32

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	415.7151324	415.7151324	187.6929163	1.90094E-14
Residual	30	66.44605573	2.214868524		
Total	31	482.1611882			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>	<i>Lower 95.0%</i>	<i>Upper 95.0%</i>
Intercept	1047.195297	77.43969353	13.52272006	2.66584E-14	1205.348249	889.0423443	1205.348249	889.0423443
X Variable 1	0.541314627	0.03951171	13.70010643	1.90094E-14	0.460620951	0.622008303	0.460620951	0.622008303

Doubling time  
(yrs) 1.28

## ***NEXT STEPS***

One way to improve this analysis would be to collect more measured power data for machines in the 1960s through the 2000s, as well as to include measured power data for large scale supercomputers. It would be particularly helpful to collect more data on older laptop computers, to see if the trends in those machines are comparable to those for desktops. Analysis of microprocessor power use (watts) compared to number of transistors on a CPU would give insight into how trends in CPU power use differed from trends in power used by other parts of computer systems.

Further statistical analysis would help in deciphering the underlying drivers for the electrical efficiency and computation trends as well as any potential errors introduced in using theoretical benchmarks like CTP instead of task-based benchmarks like SPEC. And correlating the computational efficiency data with computing costs would also be illuminating.

**Table S2: Performance and cost data for Figures**

Model	Year	Performance	Active	Computations	Notes
		M Computations/s MCPS	Power watts	per kWh	
ENIAC	1946	1.82E-05	150000	4.37E+02	1, 2
Univac I	1951	1.90E-04	126521	5.41E+03	3, 4
EDVAC	1952	2.17E-05	56000	1.40E+03	5, 6
ORDVAC	1952	1.80E-04	61000	1.06E+04	7, 8
Whirlwind II	1953	1.50E-04	180000	3.00E+03	9, 10
Burroughs 204	1954	1.20E-04	27100	1.59E+04	11, 12
IBM 702	1955	6.50E-04	74900	3.12E+04	13, 14
IBM 704	1956	6.35E-03	75000	3.05E+05	15, 16
Univac II	1957	1.65E-03	124700	4.76E+04	17, 18
UNIVAC 1105	1958	4.95E-03	160000	1.11E+05	19, 20
Burroughs D204	1959	1.67E-03	1870	3.21E+06	21, 22
NCR 304	1959	1.67E-03	42211	1.42E+05	23, 24
IBM 7090	1959	6.65E-02	28000	8.55E+06	25, 26
GE 210	1959	3.10E-03	8000	1.40E+06	27, 28
Honeywell 800	1960	2.62E-02	32000	2.94E+06	29, 30
IBM 1620	1960	6.69E-05	2000	1.20E+05	31, 32
CDC 160	1960	7.70E-05	700	3.96E+05	33, 34
CDC 1604	1960	3.45E-02	15000	8.28E+06	35, 36
Digital PDP-1	1960	3.03E-03	2160	5.05E+06	37, 38
IBM 1401 (card)	1960	5.70E-04	4500	4.56E+05	39, 40
IBM 7074	1961	3.65E-02	26226	5.00E+06	41, 42
IBM 7030 (Stretch)	1961	4.84E-01	100000	1.74E+07	43, 44
RCA 601	1961	6.36E-02	45000	5.09E+06	45, 46
UNIVAC III	1962	2.28E-02	75200	1.09E+06	47, 48
UNIVAC 1107	1962	1.03E-01	30000	1.23E+07	49, 50
SDS 920	1962	6.77E-03	990	2.46E+07	51, 52
DEC PDP-4	1962	1.30E-04	1125	4.16E+05	53, 54
Honeywell 1800	1963	7.99E-02	35000	8.22E+06	55, 56
DEC PDP-8	1965	1.32E-03	780	6.09E+06	57, 58
DEC PDP-11/20	1971	5.74E-02	400	5.17E+08	59, 60
Cray I	1976	8.60E+01	115000	2.69E+09	61, 62
IBM PC	1981	2.50E-01	72	1.26E+10	63, 64
Commodore 64	1982	2.00E-02	34	2.12E+09	65, 66
IBM PC/XT	1983	2.50E-01	89	1.01E+10	67, 68
Apple IIe	1983	2.00E-02	35	2.06E+09	69, 70
Apple Macintosh	1984	5.00E-01	33	5.45E+10	71, 72
Compaq Deskpro	1984	4.19E-01	110	1.38E+10	73, 74
IBM PC/AT	1985	6.40E-01	138	1.67E+10	75, 76

Compaq Deskpro 386	1986	2.15E+00	174	4.46E+10	77, 78
Compaq Deskpro 386/20e	1987	2.69E+00	72	1.34E+11	79, 80
AST Bravo 486/25	1991	1.13E+01	50	8.10E+11	81, 82
Gateway 2000 486/33C	1991	1.50E+01	65	8.31E+11	83, 84
IBM PS/2 E	1993	1.34E+01	13	3.87E+12	85, 86
SUN SS1000 x8	1993	4.72E+02	456	3.73E+12	87, 88
SUN Ultra450-300	1997	5.61E+02	459	4.40E+12	89, 90
Dell OptiPlex GXI	1999	1.83E+02	86	7.68E+12	91, 92
Gateway ATXSTFGP7733	2000	1.35E+03	56	8.66E+13	93, 94
SUN Blade 1000	2001	4.19E+03	248	6.08E+13	95, 96
DL360 G1	2001	2.94E+03	124	8.52E+13	97, 98
HP Pavilion 7920	2001	1.65E+03	51	1.17E+14	99, 100
Compaq 5000	2001	1.84E+03	53	1.25E+14	101, 102
Dell Optiplex GX400	2001	3.91E+03	83	1.70E+14	103, 104
Micron Client Pro	2001	1.84E+03	67	9.90E+13	105, 106
Compaq iPaq	2001	1.35E+03	48	1.01E+14	107, 108
Compaq DeskPro EN SFF	2001	1.84E+03	63	1.04E+14	109, 110
Green Destiny	2002	2.20E+05	4200	1.89E+14	111, 112
Gateway 700XL	2002	7.21E+03	152	1.71E+14	113, 114
Whitebox 1 U compute node	2003	1.97E+04	255	2.77E+14	115, 116
Intel Platform SE7520AF2 Board (3.6 GHz/1M L2 Intel Xeon processor)	2004	2.31E+04	336	2.48E+14	117, 118
DELL Dimension 2400	2004	1.31E+04	78	6.04E+14	119, 120
DELL Optiplex GX270	2004	9.63E+03	109	3.18E+14	121, 122
DELL Optiplex GX260	2004	9.83E+03	111	3.19E+14	123, 124
Custom ASUS P5 AD2-E motherboard	2005	4.91E+03	264	6.69E+13	125, 126
Dell PowerEdge 1950 Woodcrest	2006	2.73E+04	398	2.47E+14	127, 128
PowerEdge 2950 III (Intel Xeon E5440) Intel Xeon E5440	2007	5.95E+04	276	7.76E+14	129, 130
Dell Precision T3400	2008	2.09E+04	94	8.01E+14	131, 132
Dell Optiplex 765	2008	1.75E+04	92	6.85E+14	133, 134
SiCortex SC5832	2008	5.65E+06	20000	1.02E+15	135, 136
Proliant DL160 G5 (3.0 GHz, Intel Xeon processor E5450) Intel Xeon E5450	2008	6.43E+04	269	8.61E+14	137, 138
IBM System x3200 M2 Intel Xeon X3360	2008	3.11E+04	119	9.39E+14	139, 140
Macbook laptop, 13.3 inch screen	2008	1.33E+04	28	1.71E+15	141, 142
Dell Latitude E6400 laptop, 14 inch screen	2009	1.33E+04	31	1.54E+15	143, 144
Dell PowerEdge 1950 III Harpertown	2009	5.38E+04	371	5.22E+14	145, 146
DL360 G5	2009	6.01E+04	282	7.67E+14	147, 148
Dell 730X	2009	2.85E+04	195	5.26E+14	149, 150
Dell Optiplex 960	2009	1.84E+04	74	8.97E+14	151, 152

## **NOTES FOR TABLE S2**

1. Performance: Nordhaus (2007). 17468 tubes.
2. Power: Feb 14, 1946 NYT article says 150 kW but BRL study says 174 kW (apparently the ENIAC was modified from its original version: <http://ed-thelen.org/comp-hist/BRL-e-h.html>).
3. Performance: Nordhaus (2007). 5200 tubes, no transistors.
4. Power: [http://en.wikipedia.org/wiki/UNIVAC\\_I](http://en.wikipedia.org/wiki/UNIVAC_I). The power estimate is an average over 7 installations.
5. Performance: Nordhaus (2007). 5937 tubes, 328 transistors, 12,000 diodes.
6. Power: <http://en.wikipedia.org/wiki/Edvac>.
7. Performance: Nordhaus (2007). 3430 tubes, 2091 transistors, 915 diodes.
8. Power: <http://ed-thelen.org/comp-hist/BRL61-o.html#ORDVAC>. Power use is sum of computer (40), core memory (15), and magnetic drum (6).
9. Performance: Nordhaus (2007). 14,500 tubes, no transistors.
10. Power: <http://ed-thelen.org/comp-hist/BRL61-w.html#WHIRLWIND-II>. Power factor is what is used in Table 13 of BRL64 report--Computer use given as 200 kVA.
11. Performance: Nordhaus (2007). 1202 tubes, no transistors.
12. Power: <http://ed-thelen.org/comp-hist/BRL61-b.html#BURROUGHS-205> (<http://ed-thelen.org/comp-hist/BRL61-b.html#BURROUGHS-204> says to use B205 data for power)--this is an average over two installations.
13. Performance: Nordhaus (2007). 10,000 tubes, no transistors.
14. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm07.html#IBM-702>.
15. Performance: Nordhaus (2007). 5,000 tubes, no transistors. Arithmetic unit uses tubes, not clear about other parts from BRL-61. This was a very widely used machine.
16. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm0704.html#IBM-704>.
17. Performance: Nordhaus (2007). 5200 tubes, 1200 transistors.
18. Power: <http://ed-thelen.org/comp-hist/BRL61-u4.html#UNIVAC-II> (power factor derived from vA data in that same source).
19. Performance: Nordhaus (2007). 8200 tubes, 1100 transistors, 15,000 diodes.
20. Power: [http://en.wikipedia.org/wiki/UNIVAC\\_1105](http://en.wikipedia.org/wiki/UNIVAC_1105).
21. Performance: Nordhaus (2007). 8500 transistors, no tubes.
22. Power: <http://ed-thelen.org/comp-hist/BRL61-b.html#BURROUGHS-D-204>.
23. Performance: Nordhaus (2007). 4000 transistors, no tubes.
24. Power: <http://ed-thelen.org/comp-hist/BRL61-n.html#NATIONAL-304>. Power is an average over 7 machines.
25. Performance: Nordhaus (2007). 20000 transistors, no tubes.
26. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm7070.html#IBM-7090>. Use the LRL and Space tech labs installations, which cost between 2.3 and \$3M and use 35 KVA @ 80% PF.
27. Performance: Nordhaus (2007). 10000 transistors, no tubes.
28. Power: <http://ed-thelen.org/comp-hist/BRL61-g.html#GE-210>. Power factor is assumed. VA given by manufacturer as 10 kvA @ 208V. I didn't use the 40 kVA number for the sole installation listed because it looks like that's the capacity delivered to the room, not the amount used by the computer.
29. Performance: Nordhaus (2007). 2000 transistors, no tubes.
30. Power: <http://ed-thelen.org/comp-hist/BRL61-h.html#HONEYWELL-800>.

31. Performance: Nordhaus (2007). 3300 transistors, no tubes.
32. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm1401.html#IBM-1620>.
33. Performance: Nordhaus (2007). 1400 transistors, no tubes.
34. Power: <http://ed-thelen.org/comp-hist/BRL61-c.html#CDC-160>.
35. Performance: Nordhaus (2007). 25,000 transistors, no tubes.
36. Power: <http://ed-thelen.org/comp-hist/BRL61-c.html#CDC-1604>. Assumes National Bureau of Standards installation is typical (Boulder CO).
37. Performance: Nordhaus (2007).
38. Power: From Gordon Bell's email to Koomey on 090501.
39. Performance: Nordhaus (2007).
40. Power: From Ed Thelen's email to Koomey on 090501. He rebuilt an old 1401 machine and measured it as 5 kVA with estimated 0.9 PF.
41. Performance: Nordhaus (2007).
42. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm7070.html#IBM-7074>. Power factor is assumed--VA is given as 29.14 kVA. Power summary (Table XIII in BRL61 gives 26 kW).
43. Performance: Nordhaus (2007). 200,000 transistors, no tubes.
44. Power: <http://ed-thelen.org/comp-hist/BRL61-ibm7070.html#IBM-STRETCH>.
45. Performance: Nordhaus (2007).
46. Power: <http://ed-thelen.org/comp-hist/BRL61-r.html#RCA-601>.
47. Performance: Nordhaus (2007). Univac III was the first fully transistorized version of Univac.
48. Power: <http://ed-thelen.org/comp-hist/BRL61-u4.html#UNIVAC-III>.
49. Performance: Nordhaus (2007). 25522 transistors, no tubes.
50. Power: <http://ed-thelen.org/comp-hist/BRL61table13.html>.
51. Performance: Nordhaus (2007).
52. Power: <http://ed-thelen.org/comp-hist/BRL64-s.html#SDS-920>. Power factor assumed to be 0.9, use given as 1.1kVA.
53. Performance: Nordhaus (2007).
54. Power: From Gordon Bell's email to Koomey on 090501.
55. Performance: Nordhaus (2007).
56. Power: <http://ed-thelen.org/comp-hist/BRL64-h.html#HONEYWELL-1800>, 35kW given as typical system.
57. Performance: Nordhaus (2007).
58. Power: From Gordon Bell's email to Koomey on 090501.
59. Performance: Nordhaus (2007).
60. Power: p.82 of <http://www.research.microsoft.com/users/GBell/Digital/PDP%2011%20Handbook%201969.pdf>.
61. Performance: Nordhaus (2007).
62. Power: Russell (1978),L42 <http://portal.acm.org/citation.cfm?doid=359327.359336>.
63. Performance: Nordhaus (2007).
64. Power: Power data from Harris et al. 1988 (average of peak measured power, CPU only).
65. Performance: Nordhaus (2007).

66. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--includes CPU (23 W) and commodore 64 floppy drive (11 W).
67. Performance: Nordhaus (2007).
68. Power: Power data from Harris et al. 1988 (average of peak measured power, CPU only).
69. Performance: Nordhaus (2007).
70. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--includes CPU (24 W) and commodore 64 floppy drive (11 W).
71. Performance: Nordhaus (2007).
72. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--includes max power for CPU (48 W) less built-in monitor power (15 W).
73. Performance: Nordhaus IBM PC performance (8088/86), scaled by ratio of clock speeds (8 MHz/4.77 MHz). Original Deskpro used an 8 MHz 8086.  
[http://en.wikipedia.org/wiki/Compaq\\_Deskpro](http://en.wikipedia.org/wiki/Compaq_Deskpro).
74. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--used maximum boot power.
75. Performance: Nordhaus (2007).
76. Power: Power data from Harris et al. 1988 (average of peak measured power, CPU only).
77. Performance: Nordhaus, assuming same performance as Dell PC limited 386-16. The 386 came in 12, 16, 20, 25 and 33 MHz versions (though the 12 MHz versions had quality problems and weren't used much). I assume the Deskpro 386 used the 16 MHz version. GFLOPS not available from Intel CTP site.
78. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--used maximum boot power.
79. Performance: Nordhaus scaled by CTP relative to Dell PC limited 386-16. year 1987 is a guess. The 386 came in 12, 16, 20, 25 and 33 MHz versions (though the 12 MHz versions had quality problems and weren't used much). GFLOPS not available from Intel CTP site.
80. Power: Power data from Koomey measurement at the Microsoft computer archives on 081024--used maximum boot power.
81. Performance: Nordhaus scaled by CTP relative to Dell PC limited 386-16. Processor info from <http://www.intel.com/support/processors/sb/CS-020868.htm#9>.
82. Power: Power data from Roberson LBNL measurements 2001. Used maximum boot power.
83. Performance: Nordhaus scaled by CTP relative to Dell PC limited 386-16. Processor info from <http://www.intel.com/support/processors/sb/CS-020868.htm#9>.
84. Power: Power data from Roberson LBNL measurements 2001. Used maximum boot power.
85. Performance: Nordhaus scaled by CTP relative to Dell PC limited 386-16. Assumed processor is 80486SX2 at 50 MHz which Henry Wong says is equivalent to the IBM manufactured processor for this machine (486SLC2 (TM) 50/25MHz).
86. Power: Power data reported by Lovins (1993) from PC World Aug 1993, p.62. Estimated increase in power of 2.5W for full load.
87. Performance: Nordhaus (2007).

88. Power: Power data supplied in an email from Mark Monroe of Sun Microsystems June 21, 2009.
89. Performance: Nordhaus (2007).
90. Power: Power data supplied in an email from Mark Monroe of Sun Microsystems June 21, 2009.
91. Performance: Nordhaus scaled by CTP relative to Dell Precision Workstation (PW) 420-1000. Specs at <http://support.dell.com/support/edocs/systems/dzer/Specs.htm>. Choose middle Pentium processor with MMX (200 MHz/66).
92. Power: Power data from Roberson LBNL measurements 2001. Used maximum boot power.
93. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Pentium 3, 733 MHz.
94. Power: From LBNL measurements 2002 (active idle), added 20W to account for full load power (derived from Ecos measurements).
95. Performance: Nordhaus (2007).
96. Power: Power data supplied in an email from Mark Monroe of Sun Microsystems June 21, 2009.
97. Performance: Nordhaus scaled by CTP rel to Dell PW420-1000. Machine has 2 processors. Processor is Pentium 3, 800 MHz.
98. Power: From Koomey et al. 2009.
99. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Intel Celeron 900 MHz..
100. Power: From Ecos measurements, maximum power starting an application.
101. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Intel Pentium III 1000 MHz.
102. Power: From Ecos measurements, maximum power starting an application.
103. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. CTP and GFLOPS from Henry Wong, Intel. Processor is Intel Pentium 4, 1300 MHz.
104. Power: From LBNL measurements 2002 (active idle), added 20W to account for full load power (derived from Ecos measurements).
105. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Intel Pentium III 1000 MHz.
106. Power: From LBNL measurements 2002 (active idle), added 20W to account for full load power (derived from Ecos measurements).
107. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Pentium 3, 733 MHz. This desktop machine [<http://en.wikipedia.org/wiki/IPAQ\\_\(desktop\\_computer\)>](http://en.wikipedia.org/wiki/IPAQ_(desktop_computer)) should not be confused with the current HP IPAQ handheld computer [<http://en.wikipedia.org/wiki/IPAQ>](http://en.wikipedia.org/wiki/IPAQ).
108. Power: From LBNL measurements 2002 (active idle), added 20W to account for full load power (derived from Ecos measurements).
109. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Intel Pentium III 1000 MHz.
110. Power: From LBNL measurements 2002 (active idle), added 20W to account for full load power (derived from Ecos measurements).
111. Performance: Nordhaus scaled by MFLOPS rel to Dell PW420-1000. MFLOPS is based on peak theoretical FLOPS, to be consistent with the reference system.



112. Power: From email from Wuchun Feng to Koomey on 090502 (see tab server performance trends). Assumes full load power does not include disks, as recommended by Wu.
113. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. MTOPS not available for 2.4 GHz processor so we scaled linearly from MTOPS for 2.8 GHz Pentium 4. Processor is Pentium 4 @ 2.4 GHz.
114. Power: Measured power at full load (opening 15 browser windows in rapid succession) measured by Koomey at 14 Grove St., Winchester MA, June 6, 2009.
115. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Intel Pentium 4 Xeon at 3.06 GHz.
116. Power: From Koomey et al. 2009.
117. Performance: Nordhaus scaled by CTP rel to Dell PW420-1000. MFLOPS and CTP from <http://www.intel.com/support/processors/sb/CS-017346.htm>. Processor is Intel Xeon 3.6 GHz.
118. Power: Maximum Power from SPEC power run.
119. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Celeron at 2 GHz.
120. Power: From Kawamoto measurements for 2003/2004 computers, measured at full load.
121. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Pentium 4 at 3 GHz.
122. Power: From Kawamoto measurements for 2003/2004 computers, measured at full load.
123. Performance: Nordhaus scaled by CTP relative to Dell PW420-1000. Processor is Pentium 4 at 3.06 GHz.
124. Power: From Kawamoto measurements for 2003/2004 computers, measured at full load.
125. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Pentium 4 @ 3.4 GHz.
126. Power: Measured power at full load (opening 15 browser windows in rapid succession) measured by Koomey at 14 Grove St., Winchester MA, June 6, 2009.
127. Performance: Nordhaus scaled by simple average of SPECint\_rate and SPECfp\_rate 2006 relative to Dell Precision Workstation (PW) 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <http://www.spec.org>). Processor is Dual Core Xeon 5150 4MB Cache at 2.66 GHz.
128. Power: From Koomey et al. 2009.
129. Performance: Nordhaus scaled by simple average of SPECint\_rate and SPECfp\_rate 2006 relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <http://www.spec.org>). MFLOPS and CTP from <http://www.intel.com/support/processors/sb/CS-017346.htm>. Processor is Intel Xeon 5440 @ 2.83 GHz.
130. Power: Maximum Power from SPEC power run.
131. Performance: Nordhaus scaled by simple average of SPECint\_rate and SPECfp\_rate 2006 relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <http://www.spec.org>). Processor is Core 2 duo E8500 at 3.16 GHz.

132. Power: Measured power at full load (opening 15 browser windows in rapid succession) measured by Jonathan Koomey at LBNL, May 21, 2009.
133. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Core 2 duo E8500 at 3.16 GHz.
134. Power: Measured power at full load (opening 15 browser windows in rapid succession) measured by Jonathan Koomey at LBNL, May 21, 2009.
135. Performance: Nordhaus scaled by MFLOPS rel to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). MFLOPS is based on theoretical FLOPS, to be consistent with the reference system. I don't have LINPACK numbers for the Del PW690. Theoretical flops are from the SiCortex SC5832 data sheet, supplied by John Goodhue, April/May 2009.
136. Power: Power use measured while running LINPACK, from SiCortex SC5832 data sheet, verified by John Goodhue, April/May 2009.
137. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). MFLOPS and CTP from <http://www.intel.com/support/processors/sb/CS-017346.htm>. Processor is Intel Xeon 5450 @ 3.0 GHz.
138. Power: Maximum Power from SPEC power run.
139. Performance: Nordhaus scaled by simple average of SPECint\_rate and SPECfp\_rate 2006 relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <<http://www.spec.org>>). MFLOPS and CTP from <http://www.intel.com/support/processors/sb/CS-017346.htm>. Processor is Intel Xeon 3360 @ 2.83 GHz.
140. Power: Maximum Power from SPEC power run.
141. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Intel Core 2 Duo at 2.4 GHz.
142. Power: Measured power of J. Koomey's laptop at full load (opening 15 browser windows in rapid succession) minus 5 watts for screen power. Measured in February 2009.
143. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Intel Core 2 Duo at 2.4 GHz (P8600).
144. Power: Measured power of M. Koomey's laptop at full load (opening 15 browser windows in rapid succession) minus 5 watts for screen power. Measured June 26, 2009.
145. Performance: Nordhaus scaled by simple average of SPECint\_rate and SPECfp\_rate 2006 relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <<http://www.spec.org>>). Processor is Quad Core Xeon E5410, 2x6MB Cache at 2.33 GHz.
146. Power: From Koomey et al. 2009.
147. Performance: Nordhaus scaled by SPECint\_rate 2006 relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). SPEC benchmark downloaded from <<http://www.spec.org>>). Processor is Intel Xeon 5450 @ 3.0 GHz.
148. Power: From Koomey et al. 2009.
149. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Core i7 CPU, quad core, 2.66 GHz.
150. Power: Power measured by Koomey at 14 Grove St., Winchester MA, June 6, 2009. It was not possible to make the processor reach 100% utilization with the simple

approaches I tried, so instead I extrapolated linearly to 100% using active idle power (152W) and measured power at 60% load (178 W, opening about 10 pieces of software in rapid succession).

151. Performance: Nordhaus scaled by CTP relative to Dell PW 690 (intel Xeon 5160, 3.0 GHz, 1 processor/2 cores). Processor is Core 2 duo E8600 at 3.33 GHz.

152. Power: Measured power at full load (opening 15 browser windows in rapid succession) measured by Chuck Goldman with Koomey's guidance at LBNL May 21, 2009.

**Table S3: Performance scaling factors for computer systems not in Nordhaus's performance database**

Computer system	Year	Composite			Benchmark Performance	Benchmark Performance	Scaled MCPS	Reference System	Metric used
		Nordhaus Millions of comps/sec MCPS	Theoretical Performance MFLOPS	Theoretical Performance CTP MTOPS					
Compaq Deskpro 386	1986	2.2		1.8		2.2	PC limited 386-16	Equivalence	
Compaq Deskpro 386/20e	1987			2.2		2.7	PC limited 386-16	CTP	
AST Bravo 486/25	1991			9.3		11.3	PC limited 386-16	CTP	
Gateway 2000 486/33C	1991			12.4		15.0	PC limited 386-16	CTP	
IBM PS/2 E	1993			11.1		13.4	PC limited 386-16	CTP	
Dell OptiPlex GXI	1999		200	233		183	PW420-1000	CTP	
Gateway ATXSTFGP7733	2000		1,477	1,710		1,345	PW420-1000	CTP	
DL360 G1	2001		3,200	3,734		2,936	PW420-1000	CTP	
HP Pavilion 7920	2001		1,800	2,101		1,652	PW420-1000	CTP	
Compaq 5000	2001		2,000	2,333		1,835	PW420-1000	CTP	
Dell Optiplex GX400	2001		2,600	4,966		3,905	PW420-1000	CTP	
Micron Client Pro	2001		2,000	2,333		1,835	PW420-1000	CTP	
Compaq iPaq	2001		1,477	1,710		1,345	PW420-1000	CTP	
Compaq DeskPro EN SFF	2001		2,000	2,333		1,835	PW420-1000	CTP	
Green Destiny	2002		240,000			220,200	PW420-1000	MFLOPS	
Gateway 700XL	2002		4,800	9,168		7,210	PW420-1000	CTP	
Whitebox 1 U compute node	2003		12,240	24,988		19,650	PW420-1000	CTP	
Intel Platform SE7520AF2 Server Board	2004		14,400	29,398		23,118	PW420-1000	CTP	

DELL Dimension 2400	2004	8,000	16,667			13,107	PW420-1000	CTP
DELL Optiplex GX270	2004	6,000	12,249			9,632	PW420-1000	CTP
DELL Optiplex GX260	2004	6,120	12,494			9,825	PW420-1000	CTP
Custom ASUS P5 AD2-E motherboard	2005	6,800	13,882			4,906	PW690	CTP
Dell PowerEdge 1950 Woodcrest	2006	42,560	83,346	49.8	40.7	27,283	PW690	SPEC avg
PowerEdge 2950 III (Intel Xeon E5440)	2007	90,560	171,686	127	70	59,511	PW690	SPEC avg
Dell Precision T3400	2008	25,280	49,507	38.3	31.1	20,922	PW690	SPEC avg
Dell Optiplex 765	2008	25,280	49,507			17,497	PW690	CTP
SiCortex SC5832	2008	8,165,000				5,651,269	PW690	MFLOPS
Proliant DL160 G5 (3.0 GHz, Intel Xeon E5450)	2008	96,000	182,000			64,324	PW690	CTP
IBM System x3200 M2 Intel Xeon X3360	2008	45,280	85,843	59.8	43.2	31,052	PW690	SPEC avg
Macbook laptop 2.4 GHz Core 2 Duo	2008	19,200	37,600			13,289	PW690	CTP
Dell Latitude E6400 laptop	2009	19,200	37,600			13,289	PW690	CTP
Dell PowerEdge 1950 III Harpertown	2009	74,560	141,354	113	66	53,813	PW690	SPEC avg
DL360 G5	2009	96,000	182,000	110	NA	60,106	PW690	SPECint_rate
Dell 730X	2009	42,560	80,687			28,517	PW690	CTP
Dell Optiplex 960	2009	26,640	52,170			18,438	PW690	CTP

*Reference systems*

Dell PC limited 386-16	1987	2.2		1.8				
Dell PW420-1000	2000	1835	2000	2333				
Dell PW690	2006	16611	24000	47000	30.4	24.7		

PW420-1000 = Dell Precision workstation 420-1000 MHz

PW690 = Dell Precision workstation 690 (Intel Xeon 5160, 3.0 GHz, 1 processor/2 cores)

SPEC avg implies that scaling performance was done using the simple average of int\_rate and fp\_rate data, following Nordhaus, who used an average of SPECint and SPECfp.. SPECint\_rate and SPECfp\_rate are more appropriate for assessing performance of multiprocessor machines..

## **REFERENCES**

Bell, Gordon. May 1-3, 2009. *Personal Communication*: "Emails on measured power data of DEC PDP-1, PDP-4, and PDP-8 minicomputers." Contact: J. Koomey.

Daoud, David. April 10, 2009. *Personal Communication*: "IDC spreadsheet titled "IDC PC Shipments.xls" sent in email." Contact: J. Koomey.

Harris, Jeff, J. Roturier, L.K. Norford, and A. Rabl. 1988. *Technology Assessment: Electronic Office Equipment*. Lawrence Berkeley Laboratory. LBL-25558. November.

Kawamoto, Kaoru. May 17., 2009. *Personal Communication*: "Measured power use and performance data for personal computers in file named "PC data(kawamoto).xls"." Contact: J. Koomey.

Koomey, Jonathan G., Christian Belady, Michael Patterson, Anthony Santos, and Klaus-Dieter Lange. 2009. *Assessing trends over time in performance, costs, and energy use for servers*. Oakland, CA: Analytics Press. August 17.

Lovins, Amory B. 1993. *What an Energy-Efficient Computer Can Do*. Old Snowmass, CO: Rocky Mountain Institute. Aug. 10, 1993 (revised Oct. 10, 1993). <[http://www.rmi.org/images/PDFs/Energy/E93-20\\_EnergyEffComputer.pdf](http://www.rmi.org/images/PDFs/Energy/E93-20_EnergyEffComputer.pdf)>

Monroe, Mark. June 21, 2009. *Personal Communication*: "Email on measured power data of Sun E450, Blade 1000, and SS 1000 servers." Contact: J. Koomey.

Moore, Gordon E. 1965. "Cramming more components onto integrated circuits." In *Electronics*. April 19.

Nordhaus, William D. 2007. "Two Centuries of Productivity Growth in Computing." *The Journal of Economic History*. vol. 67, no. 1. March. pp. 128-159. <[http://nordhaus.econ.yale.edu/recent\\_stuff.html](http://nordhaus.econ.yale.edu/recent_stuff.html)>

Roberson, Judy A., Richard E. Brown, Bruce Nordman, Carrie A. Webber, Gregory K. Homan, Akshay Mahajan, Marla McWhinney, and Jonathan G. Koomey. 2002. *Power levels in office equipment: Measurements of new monitors and personal computers*. Proceedings of the 2002 ACEEE Summer Study on Energy Efficiency in Buildings. Asilomar, CA: Washington, DC: American Council for an Energy Efficient Economy (also LBNL-50508). August.

Russell, Richard M. 1978. "The CRAY- 1 Computer System." *Communications of the ACM*. vol. 21, no. 1. January. pp. 63-72.

Sanchez, Marla. April 27, 2009a. *Personal Communication*: "Email containing spreadsheet from Ecos Consulting supporting Energy Star impacts analysis with measured data on PCs (filename: Computer\_Power\_Test\_results.Ecos.xls)." Contact: J. Koomey.

Sanchez, Marla. April 29, 2009b. *Personal Communication*: "Email containing spreadsheets with measured PC power data created by Judy Roberson at LBNL (filenames: PC metering\_Judy\_01.xls and PC\_Metering\_Results\_from\_Judy\_020109..xls)." Contact: J. Koomey.

Weik, Martin H. 1955. *A Survey of Domestic Electronic Digital Computing Systems*. Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 971. December. <<http://ed-thelen.org/comp-hist/BRL.html>>

Weik, Martin H. 1961. *A Third Survey of Domestic Electronic Digital Computing Systems*. Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 1115. March. <<http://ed-thelen.org/comp-hist/BRL61.html>>

Weik, Martin H. 1964. *A Fourth Survey of Domestic Electronic Digital Computing Systems (Supplement to the Third Survey)* Aberdeen Proving Ground, Maryland: Ballistic Research Laboratories. Report No. 1227. January. <<http://ed-thelen.org/comp-hist/BRL64.html>>