

An Economic History of Computing¹

William D. Nordhaus

Department of Economics, Yale University

28 Hillhouse Avenue

New Haven, Connecticut 06520

Email: william.nordhaus@yale.edu

Yale University and the NBER

June 7, 2006

¹ This is a revised version of an earlier paper entitled, "The Progress of Computing." The author is grateful for comments from Ernst Berndt, Tim Bresnahan, Paul Chwelos, Iain Cockburn, Robert Gordon, Steve Landefeld, Phil Long, and Chuck Powell. Eric Weese provided extraordinary help with research on computer history.

An Economic History of Computing

Abstract

The present study analyzes computer performance over the last century and a half. Three results stand out. First, there has been a phenomenal increase in computer power over the twentieth century. Depending upon the standard used, computer performance has improved since manual computing by a factor between of 2 trillion to 73 trillion. Second, there was a major break in the trend around World War II. Third, this study develops estimates of the growth in computer power relying on performance rather than components; the price declines using performance-based measures are markedly larger than those reported in the official statistics.

The history of technological change in computing has been the subject of intensive research over the last five decades. However, little attention has been paid to comparing the performance of modern computers to pre-World-War-II technologies or even pencil-and-pad calculations. The present study investigates the progress of computing over the last century and a half, including estimates of the progress relative to manual calculations.

The usual way to examine technological progress in computing is either through estimating the rate of total or partial factor productivity or through examining trends in prices. For such measures, it is critical to use constant-quality prices so that improvements in the capabilities of computers are adequately captured. The earliest studies examined the price declines of mainframe computers and used computers that date from around 1953. Early studies found annual price declines of 15 to 30 percent per year, while recent estimates find annual price declines of 25 to 45 percent.²

While many analysts are today examining the impact of the “new economy” and especially the impact of computers on real output, inflation, and productivity, we might naturally wonder how new the new economy really is. Mainframe computers were crunching numbers long before the new economy appeared on the radar screen, and mechanical calculators produced improvements in computational capabilities even before that. How does the progress of computing in recent years compare with that of earlier epochs of the computer and calculator age? This is the question addressed in the current study.

² Table 8 below provides some documentation. See J. Steven Landefeld and Bruce T. Grimm, “A Note on the Impact of Hedonics and Computers on Real GDP,” *Survey of Current Business*, December 2000, pp. 17-22 for a discussion and a compilation of studies.

I. A Short History of Computing

Computers are such a pervasive feature of modern life that we can easily forget how much of human history existed with only the most rudimentary aids to addition, data storage, printing, copying, rapid communications, or graphics.

The earliest recorded computational device was the abacus, but its origins are not known. The Darius Vase in Naples (around 450 BC) shows a Greek treasurer using a table abacus, or counting board, on which counters were moved to record numbers and perform addition and subtraction. The earliest extant “calculator” is the Babylonian Salamis tablet (300 BC), a huge piece of marble, which used the Greek number system and probably deployed stone counters.

The design for the modern abacus appears to have its roots in the Roman hand-abacus, introducing grooves to move the counters, of which there are a few surviving examples. Counting boards looking much like the modern abacus were widely used as mechanical aids in Europe from Roman times until the Napoleonic era, after which most reckoning was done manually using the Hindu-Arabic number system. The earliest records of the modern rod abacus date from the 13th century in China (the suan-pan), and the Japanese variant (the modern soroban) came into widespread use in Japan in the 19th century.

Improving the technology for calculations naturally appealed to mathematically inclined inventors. Around 1502, Leonardo sketched a mechanical adding machine; it was never built and probably would not have worked. The first surviving machine was built by Pascal in 1642, using interlocking wheels. I estimate that less than 100 operable calculating machines were built before 1800.

Early calculators were “dumb” machines that essentially relied on incrementation of digits. An important step in development of modern computers was mechanical representation of logical steps. The first commercially practical information-processing machine was the Jacquard loom, developed in 1804. This machine used interchangeable punched cards that controlled the weaving and

allowed a large variety of patterns to be produced automatically. This invention was part of the inspiration of Charles Babbage, who developed one of the great precursor inventions in computation. He designed but never built two major conceptual breakthroughs, the “Difference Engine” and the “Analytical Engine.” The latter sketched the first design for a programmable digital computer. Neither of the Babbage machines was constructed during his lifetime. An attempt in the 1990s by the British Museum to build the simpler Difference Engine using early 19th century technologies failed to perform its designed tasks.³

The first calculator to enjoy large sales was the “arithmometer,” designed and built by Thomas de Colmar, patented in 1820. This device used levers rather than keys to enter numbers, slowing data entry. It could perform all four arithmetic operations, although the techniques are today somewhat mysterious.⁴ The device was as big as an upright piano, unwieldy,⁵ and used largely for number crunching by insurance companies and scientists. Contemporaneous records indicate that 500 were produced by 1865, so while it is often called a “commercial success,” it was probably unprofitable.

It seems unlikely that more than 500 mechanical calculators were extant at the time of rise of the calculator industry in the 1870s, so most calculations at that time were clearly done manually.⁶ By the 1880s, industrial practice plus the increasing need for accurate and rapid bookkeeping combined to give the necessary impetus for

³ See Doron Swade, *The Difference Engine*, Viking Press, New York, 2000.

⁴ An excellent short biography of this device is available in Stephen Johnston, “Making the arithmometer count,” *Bulletin of the Scientific Instrument Society*, 52 (1997), 12-21, available online at <http://www.mhs.ox.ac.uk/staff/saj/arithmometer/>.

⁵ The present author attempted to use a variant of the arithmometer but gave up attempting to perform addition after an hour.

⁶ A comprehensive economic history of calculation before the electronic age is presented in James W. Cortada, *Before the Computer*, Princeton University Press, Princeton, N.J., 1993.

the development of workable commercial adding machines and calculating machines. Corrado emphasizes the development of the technology underlying the typewriter as a key engineering element in calculator design.

Two different sets of designs were the circular machine and the keyboard design. The circular calculator was designed by Frank Baldwin in the United States and T. Odhner in Russia, both first built in the 1872-1874 period. The second and ultimately most successful early calculator was invented by Dorr E. Felt (1884) and William S. Burroughs (1885). These machines used the now-familiar matrix array of keys, and were produced by firms such as Felt Comptometer, American Arithmometer, Monroe, and Burroughs.

Production and sales of calculators began to ramp up sharply in the 1890s. The following table provides a rough estimate of the cumulative production of computational devices (excluding abacuses and counting boards) through 1920:

Decade ending	Cumulative Production
1800	50
1850	100
1860	150
1870	300
1880	700
1890	1,500
1900	8,000
1910	130,000
1920	900,000

Source: Combined from estimated sales of different devices from various sources. Worksheet available at <http://www.econ.yale.edu/~nordhaus/Computers/Appendix.xls>.

It is difficult to imagine the tedium of office work in the late 19th century. According to John Coleman, president of Burroughs, “Bookkeeping, before the

advent of the adding machine, was not an occupation for the flagging spirit or the wandering mind.... It required in extraordinary degree a capacity for sustained concentration, attention to detail, and a passion for accuracy.”⁷

Calculator manufacturers recognized that sales would depend upon the new machines being both quicker and more accurate than early devices or humans, but comparative studies of different devices are rare. A 1909 report from Burroughs compared the speed of trained clerks adding up long column of numbers by hand with that of a Burroughs calculator, as shown in Plate 1. These showed that the calculator had an advantage of about a factor of six, as reported:

Ex-President Eliot of Harvard hit the nail squarely on the head when he said, “A man ought not to be employed at a task which a machine can perform.”

Put an eight dollar a week clerk at listing and adding figures, and the left hand column [see Plate 1 below] is a fair example of what he would produce in nine minutes if he was earning his money.

The column on the right shows what the same clerk could do in one-sixth the time, or one and a half minutes.⁸

The early calculators were not well designed for mass data input and output. This problem was solved with the introduction of punched-card technology, adapted circuitously from the Jacquard power loom. The Electrical Tabulating System, designed by Herman Hollerith in the late 1880s, saw limited use in hospitals and the War Department, but its first serious deployment was for the 1890 census. The Tabulator was unable to subtract, multiply, or divide, and its addition was limited to simple incrementation. Its only function was to count the number of individuals in specified categories, but for this sole function, it was far speedier than all other

⁷ Quoted in Cortada, *Before the Computer*, p. 26.

⁸ Burroughs Adding Machine Company, *A Better Day's Work at a Less Cost of Time, Work and Worry to the Man at the Desk: in Three Parts Illustrated*, Third Edition, Detroit, Michigan, 1909, pp. 153-154.

available methods. During a government test in 1889, the tabulator processed 10,491 cards in 5½ hours, averaging 0.53 cards per second. In a sense, the Hollerith tabulator was the computer precursor of IBM's "Deep Blue" chess-playing program, which is the reigning world champion but couldn't beat a 10-year-old in a game of tic-tat-toe.

Over the next half-century, several approaches were taken to improving the speed and accuracy of computation, and the tales of mechanical and electrical engineering have been retold many times. The major technologies underlying the computers examined here are shown in Table 1. Some of the major technological milestones were the development of the principles of computer architecture and software by John von Neumann (1945), the first electronic automatic computer (the ENIAC in 1946), the invention of the transistor (1947) and its introduction into computers (1953-56), the development of the first microprocessor (1971), personal computers (dated variously from the Simon in 1950 to the Apple II in 1977 or the IBM PC in 1981), the first edition of Microsoft Windows (1983), and the introduction of the world wide web (1989).

While the engineering of calculators and computers is a much-told tale, virtually nothing has been written on the economics of early calculating devices. The economics of the computer begins with a study by Gregory Chow.⁹ He estimated the change in computer prices using three variables (multiplication time, memory size, and access time) to measure the performance of different systems over the period 1955-65. Many studies have followed in this tradition, and Jack Triplett provides an excellent recent overview of different techniques.¹⁰ Overall, we have identified 253

⁹ Gregory C. Chow, "Technological Change and the Demand for Computers," *The American Economic Review*, vol. 57, No. 5, December 1967, pp. 1117-1130.

¹⁰ Jack E. Triplett, "Performance Measures for Computers," in Dale W. Jorgenson and Charles W. Wessner, Eds., *Deconstructing the Computer: Report of a Symposium*, National Academy Press, Washington, D.C., 2005, pp. 97-140.

computing devices in this study for which minimal price and performance characteristics could be identified. The full set of machines and their major parameters are provided in an appendix available online.¹¹

II. Measuring Computer Performance

The fundamental question addressed in this study is the price of a standardized unit of computation. In early computers, this task might be adding or multiplying a set of numbers. In modern computers, the work might be solving numerical programs or operating a word-processing program. In all cases, I measure “computer power” as the amount of computation that can be performed in a given time; and the cost of computation as the cost of performing the benchmark tasks.

A. Background on measuring performance

Measuring computer power has bedeviled analysts because computer characteristics are multidimensional and evolve rapidly over time. From an economic point of view, a good index of performance would include both measures of performance on all important tasks and a set of weights that indicates the relative economic importance of the different tasks. For the earliest calculators, the tasks involved primarily addition (say for accounting ledgers). To these early tasks were soon added scientific and military applications (such as calculating ballistic trajectories, design of atomic weapons, and weather forecasts). In the modern era, computers are virtually everywhere, making complex calculations in science and industry, helping consumers surf the web or email, operating drones on the battlefield, and combating electronic diseases.

An ideal measure of computer performance would follow the principles of

¹¹ The data sources for this study are contained in an Excel series of worksheets available on the web at <http://www.econ.yale.edu/~nordhaus/Computers/Appendix.xls>. The page labeled “Contents” on that spreadsheet describes the contents in detail. It contains the major variables as well as descriptions of the derivations of variables and performance of different machines.

standard index number theory. For example, it would take an evolving mix of tasks $\{X_1(t), \dots, X_n(t)\}$ along with the prices or costs of these tasks $\{P_1(t), \dots, P_n(t)\}$. The tasks might be {addition, subtraction, multiplication, regression analysis, ..., flight simulation, Internet access, playing chess, ..., DNA sequencing, solving problems in quantum chromodynamics, ...}. The prices would be the constant-quality prices of each of these activities (using the reservation price when the activity level is zero). In principle, we could use Törnqvist indexes to construct chained cost indexes.

In practice, construction of an ideal measure is far beyond what is feasible with existing data. There is virtually no information on either the mix or relative importance of applications over time or of the market or implicit prices of different applications. The absence of reliable data on performance has forced economic studies of computer prices (called “hedonic” pricing studies) to draw instead on the prices of the input components of computers. The hedonic approach is not taken in this study but will be discussed in a later section.

As a substitute for the ideal measure, the present study has linked together price measures using changing bundles of computational tasks. The tasks examined here have evolved over time as the capabilities of computers grew. Table 2 gives an overview of the different measures of performance that are applied to the different computers.

The earliest devices, such as counting boards, the abacus, and adding machines, were primarily designed for addition; these could sometimes parlay addition into other arithmetic functions (multiplication as repeated addition). The earliest metric of computer performance therefore is simply addition time. This is converted into a measure of performance that can be compared with later computers using alternative benchmark tests. For computers from around World War II until around 1975, we use a measure of performance developed by Knight that incorporates additional attributes. For the modern period, we use computer benchmarks that have been devised by computer scientists to measure performance

on today's demanding tasks.

B. Details on Measures of Computer Performance

This section describes the different measures in detail. The major purpose is to develop a time series of performance from the earliest days to the present. We designate *CPS* or "computations per second" as the index of computer power, and *MCPS* as "millions of computations per second." For ease of understanding, I have set this index so that the speed of manual computations equals 1. As a rough guide, if you can add two five-digit numbers in 7 second and multiply two five-digit numbers in 80 seconds, you have 1 unit of computer power.

Addition Time

The earliest machines, as well as manual calculations by humans, were primarily capable of addition. Plate 1 shows the results of a typical task as described in 1909. In fact, until World War II, virtually all commercial machines were devoted solely to addition. We can compare the addition time of different machines quite easily as long as we are careful to ensure that the length of the word is kept constant for different machines.

Moravec's Information-Theoretic Measure of Performance

A measure of performance that relies primarily on arithmetic operations but has a stronger conceptual basis is the information-theoretic measure devised by Hans Moravec. To compare different machines, Moravec defined computing power as the amount of information delivered per second by the machine.¹²

This can then be put on a standardized basis by considering words with a standard length of 32 bits (equivalent to a 9-digit integer), and instructions with a length of one word. Moravec assumed that there were 32 instructions, and included measures on addition and multiplication time, which were weighted seven to one in the operation mix. Using this definition, the information-theoretic definition of

¹² See Moravec, *Mind Children: The Future of Robot and Human Intelligence*, Harvard University Press, Cambridge, MA, 1988, especially Appendix A2 and p. 63f.

performance is:

$$\text{Computer power (Moravec)} = \frac{\{[6 + \log_2 (\text{memory}) + \text{word length}/2]/[(7 \times \text{add time} + \text{mult time})/8]\}}$$

The attractiveness of this approach is that each of these parameters is available for virtually all computers back to 1940, and can be estimated or inferred for manual calculations, abacuses, and many early calculators. The disadvantages are that it omits many of the important operations of modern computers, it considers only machine-level operations, and it cannot incorporate the advantages of modern software, higher-level languages, and operating systems.

Knight's measure

One of the earliest studies of computer performance was by Kenneth Knight of RAND in 1966.¹³ He wanted to go beyond the simplest measures of addition and multiplication time and did a number of experiments on the capacity of different machines to perform different applications. His formula for computer power was as follows:

Knight's Index of Computer Power

$$\approx 10^6 \frac{\{[(\text{word length} - 7)(\text{memory})]\}}{[\text{calculation time} + \text{input-output time}]}$$

Knight's formula is quite similar to Moravec's except that he includes a larger number of variables and particularly because he calibrates the parameters to the actual performance of different machines.

MIPS

One of the earliest benchmarks used was MIPS, or millions of instructions per

¹³ Kenneth Knight, "Changes in Computer Performance," *Datamation*, vol. 12, no. 9, Sept. 1966, pp. 40-54 and "Evolving Computer Performance 1963-1967," *Datamation*, vol. 14, no. 1, Jan. 1968, pp. 31-35.

second. In simple terms, instructions per second measures the number of machine instructions that a computer can execute in one second. This measure was developed to compare the performance of mainframe computers. The most careful studies used weighted instruction mixes, where the weights were drawn from the records of computer centers on the frequency of different instructions. These benchmarks were probably the only time something approaching the ideal measure described above was constructed.

A simplified description of MIPS is the following. For a single instruction:

$$\text{MIPS} = \text{clock rate} / (\text{cycles per instruction} \times 10^6)$$

To understand the logic of this measure, recall that computers that use the von Neumann architecture contain an internal clock that regulates the rate at which instructions are executed and synchronizes all the various computer components. The speed at which the microprocessor executes instructions is its “clock speed.” For most personal computers up to around 2000, operations were performed sequentially, once per clock tick.¹⁴ An instruction is an order given to a computer processor by a computer program. Computers with complex instruction sets might have between 200 and 400 machine-language instructions, while computers with reduced instruction sets would have only 30 to 50 unique instructions.

Instructions differ in terms of the size of the “word” that is addressed. In the earliest computers (such as the Whirlwind I), words were as short as 16 binary digits or 5 decimal digits. Most personal computers today use 32-bit words, while mainframes generally employ 64-bit words.

Modern Benchmark Tests

¹⁴ Many of the major topics in computer architecture can be found in books on computer science. For example, see G. Michael Schneider and Judith L. Gersting, *An Invitation to Computer Science*, Brooks/Cole, Pacific Grove, California, 2000.

Measures like additions or instructions per second or more complex indexes like those of Knight or Moravec clearly cannot capture today's complex computational environment. Computers today do much more than bookkeeping, and a performance benchmark must reflect today's mix of activities rather than that of a century ago. For this purpose, we turn to modern benchmark tests.

A *benchmark test* is an index that measures the performance of a system or subsystem on a well-defined set of tasks. Widely used benchmarks for personal computers today are those designed by SPEC, or the Standard Performance Evaluation Corporation. As of mid-2006, the version used for personal computers was SPEC CPU2000.¹⁵ SPEC CPU2000 is made up of two components that focus on different types of compute intensive performance: SPECint2000 for measuring and comparing computer-intensive integer computation and SPECfp2000 for measuring computer-intensive floating-point computation.

Table 3 shows the suite of activities that SPEC2000 tests. These are obviously not routine chores. The benchmark fails to follow the elementary rule of ideal indexes in that the performance on different benchmarks is clearly not weighted by the economic importance of different applications. We discuss below the relationship between the SPEC and other benchmarks. To make current tests comparable with early ones, ratings have been set by comparing the rating of a machine with the rating of a benchmark machine.

III. Measure of Computer Performance

This study is an attempt to link together computational performance of different machines from the nineteenth century to the present. A unit of computer performance is indexed so that manual computations are equal to 1. A standard modern convention is that the VAX 11-780 is designated as a one MIPS machine. In

¹⁵ See <http://www.spec.org/osg/cpu2000/> .

our units, the VAX 11-780 is approximately 150 million times as powerful as manual computations. Different modern benchmarks yield different numbers, but they are essentially scalar multiples of one another.

Constructing metrics of performance is difficult both because the tasks and machines differ enormously over this period and because measures of performance are very sketchy before 1945. The data since 1945 have been the subject of many studies since that period. Data for this study for computers from 1945 to 1961 were largely drawn from technical manuals of the Army Research Laboratory, which contain an exhaustive study of the performance characteristics of systems from ENIAC through IBM-702.¹⁶ Additionally, studies of Kenneth Knight provided estimates of computer power for the period 1945 through 1966. Data on the performance of computers through 2003 have been carefully compiled by John C. McCallum and are available on the web.¹⁷ Machines since 2003 were evaluated by the author.

Reliable data for the earliest calculators and computers (for the period before 1945) were not available in published studies. With the help of Eric Weese of Yale University, data from historical sources on the performance of 32 technologies from before 1940 were obtained, for which 12 have performance and price data which I consider reasonably reliable. I will discuss the data on the early technologies because these are the major original data for the present study.

¹⁶ See particularly Martin H. Weik, *A Survey of Domestic Electronic Digital Computing Systems*, Ballistic Research Laboratories, Report No. 971, December 1955, Department of the Army Project No. 5b0306002, Ordnance Research and Development Project No. Tb3-0007, Aberdeen Proving Ground, Maryland available at <http://ed-thelen.org/comphist/BRL.html>. This was updated in Martin H. Weik, *A Third Survey of Domestic Electronic Digital Computing Systems*, Report No. 1115, March 1961, Ballistic Research Laboratories, Aberdeen Proving Ground, Maryland, available at <http://ed-thelen.org/comphist/BRL61.html#table-of-contents>.

¹⁷ See <http://www.jcmit.com/cpu-performance.htm> .

The data on manual calculations were taken from a Burroughs monograph, from estimates of Moravec, and from tests by the author.¹⁸ The computational capabilities of the abacus are not easily measured because of the paucity of users in most countries today. One charming story reports a Tokyo competition in 1946 between the U.S. Army's most expert operator of the electric calculator in Japan and a champion operator of the abacus in the Japanese Ministry of Postal Administration. According to the report, the addition contest consisted of adding 50 numbers each containing 3 to 6 digits. In terms of total digits added, this is approximately the same as the tests shown in Plate 1. The abacus champion completed the addition tasks in an average of 75 seconds, while the calculator champion required 120 seconds. They battled to a standoff in multiplication and division. The abacus expert won 4 of the 5 contests and was declared the victor.¹⁹

This comparison suggests that, in the hands of a champion, the abacus had a computer power approximately $4\frac{1}{2}$ times that of manual calculation. Given the complexity of using an abacus, however, it is unlikely that this large an advantage would be found among average users. We have reviewed requirements for Japanese licensing examinations for different grades of abacus users from the 1950s. These estimates suggest that the lowest license level (third grade) has a speed approximately 10 percent faster than manual computations.²⁰

We have estimated the capabilities of early machines based on then-current procedures. For example, many of the early machines were unable to multiply. We therefore assume that multiplication was achieved by repeated addition. Additionally, the meaning of memory size in early machines is not obvious. For

¹⁸ The calculations are available at

<http://www.econ.yale.edu/~nordhaus/Computers/Appendix.xls>.

¹⁹ The contest and its results are described in Takashi Kojima, *The Japanese Abacus: Its Use and Theory*, Charles E. Tuttle, Rutland, Vermont, 1954.

²⁰ Takashi Kojima, *op. cit.*

machines that operate by incrementation, we assume that the memory is one word. There are major discrepancies between different estimates of the performance of early machines, with estimates varying by as much as a factor of three. Given the difficulties of collecting data on the earliest machines, along with the problems of making the measures compatible, we regard the estimates for the period before 1945 as subject to large errors.

The construction of the performance measures was described above. The only other assumptions involve constructing the cost per operation. These calculations include primarily the cost of capital. The data on prices and wage rates were prepared by the author and are from standard sources, particularly the U.S. Bureau of Labor Statistics and the U.S. Bureau of Economic Analysis. We have also included estimates of operating costs as these appear to have been a substantial fraction of costs for many computers and may be important for recent computers. For the capital cost, we use the standard user cost of capital formula with a constant real interest rate of 10 percent per year, an exponential depreciation rate of 10 percent per year, a utilization factor of 2000 hours per year, and no adjustment for taxes. These assumptions are likely to be oversimplified for some technologies, but given the pace of improvement in performance, even errors of 10 or 20 percent for particular technologies will have little effect on the overall results.

IV. Results

Overall trends

I now discuss the major results of the study. The following table shows a summary of the overall improvement in computing relative to manual calculations and the growth rates in performance. The quantitative measures are computer power, the cost per unit computer power in terms of the overall price level, and the cost of computation in terms of the price of labor. The overall improvements relative to manual computing range between 2 and 73 trillion depending upon the measure used. For the period 1850 (which I take as the birth of modern computing) to 2006,

the compound logarithmic growth rate is around 20 percent per year.

Improvement from Manual (1850) to 2006	Improvement	
	Ratio	Annual growth rate (percent per year)
Computer power (MCPS per second)	2,050,000,000,000	18.3
Price per calculation (MCPS per 2003\$)	7,100,000,000,000	19.1
Labor cost of computation (MCPS per hour)	73,000,000,000,000	20.6

We now discuss the results in detail. Start with Figure 1, which shows the results in terms of pure performance – the computing power in terms of computations per second. Recall that the index is normalized so that manual computation is 1. Before World War II, the computation speeds of the best machines were between 10 and 100 times the speed of manual calculations. There was improvement, but it was relatively slow. Figure 2 shows the trend in the cost of computing over the last century and a half. The prices of computation begin at around \$500 per MCPS for manual computations and decline to around $\$6 \times 10^{-11}$ per MCPS by 2006 (all in 2006 prices), which is a decline of a factor of 7 trillion.

Table 4 shows five different measures of computational performance, starting with manual computations through to the mid 2000s. The five measures are computer power, cost per unit calculation, labor cost per unit calculation, cycles per second, and rapid memory. The general trends are similar, but different measures can differ substantially. One important index is the relative cost of computation to labor cost. This is the inverse of total labor productivity in computation, and the units are therefore CPS per hour of work.²¹ Relative to the price of labor,

²¹ The advantages of using wage as a deflator are twofold. First, it provides a measure of the relative price of two important inputs (that is, the relative costs of labor

computation has become cheaper by a factor of 7.3×10^{13} compared to manual calculations. Given the enormous decrease in computational cost relative to labor cost, it can hardly be surprising that businesses are computerizing operations on a vast scale.

Trends for different periods

We next examine the progress of computing for different subperiods. The major surprise, clearly shown in Figures 1 and 2, is the discontinuity that took place around World War II. Table 4 shows data on performance of machines in different periods, while Table 5 shows the logarithmic annual growth rates between periods (defining manual calculations as the first period). Table 5 indicates modest growth in performance from manual computation until the 1940s. The average increase in computer productivity shown in the first three columns of the first row of Table 5 – showing gains of around 3 percent per year – was probably close to the average for the economy as a whole during this period.

Statistical estimates of the decadal improvements are constructed using a log-linear spline regression analysis. Table 6 shows a regression of the logarithm of the constant-dollar price of computer power with decadal trend variables. The coefficient is the logarithmic growth rate, so to get the growth rate for a period we can sum the coefficients up to that period. The last column of Table 6 shows the annual rates of improvement of computer performance. All measures of growth rates are logarithmic growth rates.²²

and computation). Additionally, the convention of using a price index as a deflator is defective because the numerator is also partially contained in the denominator.

²² The growth rates are instantaneous or logarithmic growth rates, which are equivalent to the derivatives of the logarithms of series with respect to time for smooth variables. This convention is used to avoid the numerical problems that arise for high growth rates.

The regression analysis shows that the explosion in computer power, performance, and productivity growth began around 1945. Tables 5 and 6 provide slightly different estimates of the sub-period growth rates, but it is clear that productivity growth was extremely rapid during virtually the entire period since 1945. Using decadal trend-break variables, as shown in Table 6, we find highly significant positive coefficients in 1945 and 1985 (both indicating acceleration of progress). The only period when progress was slow (only 22 percent per year!) was during the 1970s.

The rapid improvement in computer power is often linked with “Moore’s Law.” This derives from Gordon Moore, co-founder of Intel, who observed in 1965 that the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented. Moore predicted that this trend would continue for the foreseeable future. When he revisited this question a decade later, he thought that the growth rate had slowed somewhat and forecast that doubling every 18 months was a likely rate for the future (46 percent logarithmic growth). Two remarks arise here. First, it is clear that rapid improvements in computational speed and cost predated Moore’s forecast. For the period 1945-1980, cost per computation declined at 37 percent per year, as compared to 64 percent per year after 1980. Second, computational power actually grows more rapidly than

A word is in order for those not accustomed to logarithmic growth rates: These will be close to the conventional arithmetic growth rate for small numbers (2 or 3 percent per year) but will diverge significantly for high growth rates. For example, an arithmetic growth rate of 100 percent per year is equivalent to a logarithmic rate of 0.693. A further warning should be given on negative growth rates. There is no difficulty in converting negative to positive rates as long as logarithmic growth rates are used. However, in using arithmetic growth rates, decline rates may look significantly smaller than the corresponding growth rate. For example, a logarithmic growth rate of -0.693 represents a decline rate of 50 percent per year.

Moore's Law would predict because computer performance is not identical to chip density. From 1982 to 2001, the rate of performance as measured by computer power grew 18 percent per year faster than Moore's Law would indicate.

One of the concerns with the approach taken in this study is that our measures might be poor indexes of performance. We have compared MCPS with both addition time and cycle time (the latter comparison is shown in Tables 4 and 5). Both simple proxies show a very high correlation with our synthetic measure of MCPS over the entire period. Computer power grows at very close to the speed of addition time (for observations from 1900 to 1978) but 10 percent per year more rapidly than cycle speed (1938-2006).

In this regard, it is natural to ask whether the changing character of computers is likely to bias the estimates of the price of computer power. The earliest calculators had very low capability relative to modern computers, being limited to addition and multiplication. Modern computers perform a vast array of activities that were unimaginable a century ago (see Table 3). In terms of the ideal measure described above, it is likely that standard measures of performance are biased *downward*. If we take an early output mix – addition only – then the price index changes very little, as discussed in the last paragraph. On the other hand, today's output bundle was infeasible a century ago, so a price index using today's bundle of output would have fallen even faster than the index reported here. Put differently, a particular benchmark only includes what is feasible, that is, tasks which can be performed in a straightforward way by that year's computers and operating systems. Quantum chromodynamics is included in SPEC 2000, but it would not have been dreamt of by Kenneth Knight in his 1966 study. This changing bundle of tasks suggests that, if anything, the price of computation has fallen even faster than the figures reported here.

V. Alternative Approaches

Comparison of Alternative Modern Benchmark Tests

Using direct measures of computer performance raises two major problems. First, a properly constructed benchmark should weight the different tasks according to their economic importance, but this property is satisfied by none of the current benchmark tests. For example, as shown in Table 3, the SPEC benchmark that is widely used to test PCs contains several exotic tasks, such as quantum chromodynamics, which are probably not part of the family computer hour. Most benchmarks simply apply equal geometric weights to the different tasks. Second, the rapid evolution of computer performance leads to rapid changes in the tasks that the benchmarks actually measure. For example, the SPEC performance benchmark has been revised every two or three years. In one sense, such changes represent a kind of chain index in tasks; however, because tasks are not appropriately weighted, it is impossible to know whether the chaining improves the accuracy of the indexes.

With the help of Eric Weese, I investigated the results of using different benchmark tests over the last decade. For this purpose, we examined (1) the SPEC benchmarks, (2) a series of tests known as WorldBench, which have been published by *PC World*, and (3) SYSmark98, a measure that evaluates performance for 14 applications-based tasks. To illustrate how PC benchmarks work, the SYSmark98 test of office productivity is the harmonic mean of the time to open and perform set tasks on the following programs: CorelDRAW 8, Microsoft Excel 97, Dragon Systems NaturallySpeaking 2.02, Netscape Communicator 4.05 Standard Edition, Caere OmniPage Pro 8.0, Corel Paradox 8, Microsoft PowerPoint 97, and Microsoft Word 97.

We first examined 30 machines for which both PC benchmarks had results over the period December 1998 to November 1999. The two benchmarks were reasonably consistent, with a correlation of 0.962 in the logarithm of the benchmark scores over the 30 machines. However, as shown in comparison one of Table 7, the rate of improvement of the two indexes differed markedly, with the SYSmark98 showing a 38 percent per year improvement over the 11-month period, while the

WorldBench tests showed a 24 percent per year improvement for the same machines. This difference was found even in the individual benchmarks (e.g., the result for Excel 97), and queries to the vendors produced no reasons for the discrepancies.²³

A second comparison was between the SPEC benchmark results and the WorldBench tests. For this comparison, we were able to find 7 machines that were tested for both benchmarks over a period of two years, using the 1995 SPEC test and 3 different WorldBench tests. For these machines, as shown in comparison 2 of Table 7, the rate of improvement of the SPEC and WorldBench tests were virtually identical at 67 and 66 percent per year, respectively.

The final test involved a comparison of the WorldBench score with the improvement in computations per second calculated for the present paper. For this purpose, we gathered different WorldBench tests for the period from 1992 to 2002 and spliced them together to obtain a single index for this period. We then calculated the growth of WorldBench performance per constant dollar and compared this to the growth of CPS per constant dollar from the current study. As shown in comparison 3, the WorldBench performance per real dollar over 1992-2002 showed a 52 percent per year increase. This compares with a 62 percent per year increase for the computers in our data set over the same period.

To summarize, we have investigated the results of alternative benchmarks tests. None of the benchmarks is well constructed because the weights on the different tasks do not reflect the relative economic importance of tasks. We found some discrepancies among the different benchmarks, even those that purport to measure the same tasks. The WorldBench test, which is oriented toward PCs, showed slower improvements in constant-dollar performance over the 1992-2002 period than the CPS measure constructed for this study. However, an alternative

²³ For example, we compared the raw scores for the two benchmarks for six identical machines and three identical programs. The harmonic means differed by as much as 17 percent between the two sets of tests.

test, SYSmark, show more rapid growth than the WorldBench and was more consistent with the CPS measure in this paper and with the SPEC benchmark. In any case, the improvement in constant dollar performance was extremely rapid, with the lower number being a 52 percent per year logarithmic increase for the WorldBench and the higher number being a 62 percent per year increase over the last decade for the CPS.

Comparison with Alternative Indexes of Computer Prices

Economists today tend to favor the use of hedonic or constant-quality price indexes to measure improvements. The hedonic approach attempts to measure the change in the “quantity” of goods by examining the change in characteristics along with measures of the importance of the different characteristics.²⁴

A pioneering study that investigated hedonic prices of performance was undertaken by Paul Chwelos. He investigated the characteristics of computers that were important for users and information scientists in 1999 and found the top six characteristics were (1) performance, (2) compatibility, (3) RAM, (4) network connectivity, (5) industrial standard components, and (6) operating system.²⁵ He then estimated the change in the cost of providing the bundle.

Clearly, such an approach is not feasible over the long span used here. In the

²⁴ There are many excellent surveys of hedonic methods. A recent National Academy of Sciences report has a clear explanation of different approaches. See Charles Schultze and Christopher Mackie, *At What Price? Conceptualizing and Measuring Cost-of-Living and Price Indexes*, National Research Council, Washington, D.C., 2002.

²⁵ See Paul Chwelos, *Hedonic Approaches to Measuring Price and Quality Change in Personal Computer Systems*, Ph. D. Thesis, the University of Victoria, 1999, p. 43. Performance was defined as a “characteristic of the a number of components: CPU (generation, Level 1 cache, and clock speed), motherboard architecture (PCI versus ISA) and bus speed, quantity and type of Level 2 cache and RAM, type of drive interface (EIDE versus SCSI).”

present study, we examined only the price of a single characteristic, performance. This decision reflects the fact that only two of the six performance characteristics discussed in the last paragraph [number 1 (performance) and number 3 (RAM)] can be tracked back for more than a few decades. Network connectivity is a brand-new feature, while operating systems have evolved from tangles of wires to Windows-type operating systems with tens of millions of lines of high-level code. This discussion indicates that computers have experienced not only rapid improvements in speed but also provide additional goods and services.

How do the performance-based indexes used here compare with price indexes for computers? A summary table of different price indexes for recent periods is provided in Table 8. There are six variants of computer price indexes prepared by the government, either for the national income and product accounts by the Bureau of Economic Analysis (BEA) or for the Producer Price Index by the Bureau of Labor Statistics.²⁶

During the 1969-2004 period, for which we have detailed price indexes from the BEA, the real price index for computers fell by 23 percent per year relative to the GDP price index (using the logarithmic growth rate), while the real BLS price index for personal workstations and computers fell by 31 percent. Academic studies, using hedonic approaches or performance measures, show larger decreases, between 35 and 40 percent. Our real price index of the price of computer power fell by between 50 and 58 percent depending upon the subperiod.

How might we reconcile the significant discrepancy between the hedonic measures and the performance-based prices reported here? A first possible discrepancy arises because government price indexes for computers are based on the

²⁶ See J. Steven Landefeld and Bruce T. Grimm, "A Note on the Impact of Hedonics and Computers on Real GDP," *Survey of Current Business*, December 2000, pp. 17-22 for a discussion and a compilation of studies. The data are available at <http://www.bea.doc.gov/>.

prices of *inputs into computers*, while the measures presented here are indexes of the *cost of specified tasks*. The hedonic measures will only be accurate to the extent that the prices of components accurately reflect the marginal contribution of different components to users' valuation of computer power. It is worth noting that current government hedonic indexes of computers contain *no* performance measure.²⁷

A second and more important difference is that computers increasingly are doing much more than computing, so that our indexes capture many non-computational features. To illustrate, in late 2005, a Intel® Pentium® 4 Processor 630 with HT (3GHz, 2M, 800MHz FSB) was priced at \$218 while the Dell OptiPlex™ GX620 Mini-Tower personal computer in which it was embedded cost \$809. The \$591 difference reflects ancillary features such as hard-drives, ports, CD/DVD readers, pre-loaded software, assembly, box, and so forth. A perfectly constructed hedonic price index will capture this changing mix of components. To illustrate this point, assume that a 2005 computer is 25 percent computation and 75 percent razzle-dazzle, while a 1965 computer such as the DEC PDP-8 was 100 percent computation. Using our estimates, this would change the real price decline from 45 percent per year to 48 percent over the four decades. It seems unlikely that the prices of the non-computational components are falling as rapidly as the computational parts. Hence, the discrepancy is partly because “computers” are now doing much more than computing.

²⁷ The variables in the earlier BLS hedonic regression for personal desktop computers (designed in 1999 but discontinued after 2003) contained one performance proxy (clock speed), two performance-related proxies (RAM and size of hard drive), an array of feature dummy variables (presence of Celeron CPU, ZIP drive, DVD, fax modem, speakers, and software), three company dummy variables, and a few other items. It contained no performance measures such as the SPEC benchmark. The new BLS pricing approach contains no performance measures at all and instead uses attribute values available on the Internet as a basis to determine appropriate quality adjustments amounts.

Supercomputing

While this study has emphasized conventional computers, it will be useful to devote a moment's attentions to the elephants of the computer kingdom. Scientists and policy makers often emphasize supercomputing as the "frontier" aspect of computation, the "grand challenges of computation," or the need for "high performance computing." These are the romantic moon shots of the computer age. What exactly are the grand challenges? Generally, supercomputers are necessary for the simulation or solution of extremely large non-linear dynamic systems. Among the important applications discussed by scientists are applied fluid dynamics, meso- to macro-scale environmental modeling, ecosystem simulations, biomedical imaging and biomechanics, molecular biology, molecular design and process optimization, and fundamental computational sciences.²⁸ To pick the second of these areas, environmental modeling, there are enormous demands for improvements in modeling of climate systems and interactions between oceans, the atmosphere, and the cryosphere; our understanding of many issues about the pace and impact of climate change will depend upon improving the models and the computers to solve the models.

The progress in supercomputing has paralleled that in smaller computers. As of November 2005, for example, the largest supercomputer (IBM's Blue Gene/L with 131,072 processors) operated at a maximum speed of 280,600 gigaflops (billions of floating-point operations per second or Gflops). Using a rough conversion ratio of 475 CPS per Flop, this machine is therefore approximately a 133,000,000,000 MCPS machine and therefore about 53,000 times more powerful than the top personal computer in our list as of 2006. The performance improvement for supercomputers has been tracked by an on-line consortium called "TOP500." It shows that the top

²⁸ See the discussion in National Research Council, *High Performance Computing and Communications: Foundation for America's Information Future*, 1996.

machine's performance grew from 59.7 Gflops in June 1993 to 280,600 Gflops in June 2005.²⁹ Over this period, the peak performance grew at a rate of 97 percent per year. This is higher than the rate in our sample of smaller computers. However, it is likely that the performance was tuned to the benchmark, and the large systems are clearly not as versatile as personal computers.

The price of supercomputing is generally unfavorable relative to personal computers. IBM's stock model supercomputer, called "Blue Horizon," is clocked at 1700 Gflops and had a list price in 2002 of \$50 million – about \$30,000 per Gflops – which makes it approximately 34 times as expensive on a pure performance basis as a Dell personal computer in 2004.

V. Conclusions

The key purpose of this study is to extend estimates of the price of computers and computation back in time to the earliest computers and calculators as well as to manual calculations. Along the way, we have developed performance-based measures of price and output that can be compared with input-based or component-based measures.

Before reviewing the major conclusions, we must note some of the major reservations about the results. While we have provided performance-based measures of different devices, we note that the measures are generally extremely limited in their purview. They capture primarily computational capacity and generally omit other important aspects of modern computers such as connectivity, reliability, size, and portability. In one sense, we are comparing the transportation skills of the computer analogs of mice and men without taking into account many of the "higher" functions that modern computers perform relative to mice like the IBM 1620 or ants like the Hollerith tabulator.

In addition, we emphasize that some of the data used in the analysis,

²⁹ See www.top500.org.

particularly those for devices before 1945, are relatively crude. Additionally, the measures of performance or computer power used for early computers have been superseded by more sophisticated benchmarks. While conventional equivalence scales exist and are used when possible, the calibrations are imperfect. Subject to these reservations, the following conclusions seem warranted.

First, there has been a phenomenal increase in computer power over the twentieth century. Performance in constant dollars has improved relative to manual calculations by a factor in the order of 2×10^{12} (that is, 2 trillion). Most of the increase has taken place since 1945, during which the average rate of improvement has been at a rate of 45 percent per year. The record shows virtually continuous extremely rapid productivity improvement over the last six decades. These increases in productivity are far larger than that for any other good or service in the historical record.³⁰

Second, the data show a sharp break in trend around 1945 – at the era where the technological transition occurred from mechanical calculators to what are recognizably the relatives of modern computers. There was only modest progress – perhaps a factor of 10 – in general computational capability from the skilled clerk to the mechanical calculators of the 1920s and 1930s. Around the beginning of World War II, all the major elements of the first part of the computer revolution were developed, including the concept of stored programs, the use of relays, vacuum tubes, and eventually the transistor, improved software, along with a host of other components. Dating from about 1945, computational speed increased and costs

³⁰ Scholars have sometimes compared productivity growth in computers with that in electricity. In fact, this is a snails-to-cheetah comparison. Over the half-century after the first introduction of electricity, its price fell 6 percent per year relative to wages, whereas for the six decades after World War II the price of computer power fell 47 percent per year relative to wages. An index of communications prices for 1200 – 2002 constructed by Eric Weese shows a decline of about 10^5 or 1.4 percent per year.

decreased rapidly up to the present. The most rapid pace of improvement was in the periods 1985-95 and 1945-55.

Third, these estimates of the growth in computer power, or the decline in calculation costs, are more rapid than price measures for computers used in the official government statistics. There are likely to be two reasons for the difference: first, the measures developed here are indexes of performance, while the approaches used by governments are based on the prices of components or inputs; and, second, “computers” today are doing much more than computation.

Fourth, the phenomenal increases of computer power and declines in the cost of computation over the last three decades have taken place through improvements of a given underlying technology: stored programs using the von Neumann architecture of 1946 and hardware based on Intel microprocessors descended from the Intel 4004 of 1971. The fact that this extraordinary growth in productivity took place in a relatively stable industry, in the world most stable country, relying on a largely unchanged core architecture, is provocative for students of industrial organization to consider.

157.38	15738
762.91	76291
435.75	43575
800.72	80072
43.02	4302
987.25	98725
500.00	50000
1002.50	100250
245.65	24565
82.47	8247
4250.86	425086
1014.75	101475
243.92	24392
914.75	91475
5475.80	547580
14850.07	1485007
410.25	41025
.74	.74
9.10	910
27.72	2772
896.35	89635
1238.63	123863
7800.00	780000
10000.00	1000000
127.34	12734
77.01	7701
303.24	30324
3808.89	380889
4480.92	448092
1456789.34	145678934
11025.22	1102522
600.10	60010
2250.85	225085
14823.45	1482345
8207.12	820712
2100.00	210000
<u>1552723118</u>	1552723118
2	3

Facsimile of Figures Added and Listed by a Clerk in 9 Minutes.

Facsimile of Figures Added and Listed by a "Burroughs" in 1 Min. 30 Sec.

Plate 1. Comparison of Manual Calculation with Manual Calculator

This plate shows a comparison of manual calculators and computations by a clerk in adding up a column of numbers such as might be found in a ledger. The calculator has an advantage of a factor of six. (Source: Burroughs Adding Machine Company, *A Better Day's Work at a Less Cost of Time, Work and Worry to the Man at the Desk: in Three Parts Illustrated*, Third Edition, Detroit, Michigan, 1909, pp. 153-154.)

1. Manual (see Plate 1) - up to around 1900
2. Mechanical - circa 1623 to 1945
3. Electromechanical -1902 to 1950
4. Relays - 1939 - 1944
5. Vacuum tubes - 1942 - 1961
6. Transistor - 1956 - 1979
7. Microprocessor - 1971 - present

Table 1. The Seven Stages of Computation

The dates in the table represent the dates for the technologies that are represented in this study.

1. Addition time (up to about 1944)
2. Millions of instructions or operations per second (1944 to 1980s)
3. Moravec's formula:
Performance a function of (add-time, multiplication-time, memory, word size)
4. Knight's formula (1944 to 1972)
Performance a function of (word size, memory, calculation-time, IO time, ...)
5. Synthetic benchmarks:
Dhrystone (1984 - 1990)
WorldBench
SYSmark
SPEC (latest being SPEC2000): 1993 to present

Table 2. Alternative measures of performance used in this study

Benchmarks used in measuring computer performance have evolved from the speed of addition or multiplication to the performance on complex tasks. The tasks used in the latest SPEC benchmark are shown in Figure 3.

SPECint2000

Compression
FPGA circuit placement and routing
C programming language compiler
Combinatorial optimization
Game playing: Chess
Word processing
Computer visualization
Perl programming language
Group theory, interpreter
Object-oriented database
Compression
Place and route simulator

SPECfp2000

Physics: Quantum chromodynamics
Shallow water modeling
Multigrid solver: 3D potential field
Partial differential equations
3D graphics library
Computational fluid dynamics
Image recognition/neural networks
Seismic wave propagation simulation
Image processing: Face recognition
Computational chemistry
Number theory/primalty testing
Finite-element crash simulation
Nuclear physics accelerator design
Meteorology: Pollutant distribution

Table 3. Suite of Programs Used for SPEC2000 Benchmark

This table shows the benchmarks used to evaluate different computers. The first set use largely integer applications while the second are largely floating-point scientific applications.

Source: John L. Henning, "SPEC CPU2000: Measuring CPU Performance in the New Millennium," *Computer*, July 2000, p. 29.

Period	Computer power (units per second)	Total cost per million unit computer power (2006 \$)	Labor cost of computation (hours per unit computer power)	Cycle speed (cycles per second)	Rapid access memory (bits)
Manual	1.00E+00	1.00E+00	1.00E+00	1.00E+00	1.00E+00
Late 19th C	6.48E+00	3.20E-01	1.92E-01	2.00E+00	6.02E+00
1900-1939	1.78E+01	2.00E-01	1.12E-01	7.50E+00	8.00E+01
1940s	1.67E+03	5.61E-02	1.14E-02	2.50E+05	1.12E+03
1950s	1.18E+05	1.67E-03	2.55E-04	1.80E+06	1.12E+04
1960s	2.92E+06	7.39E-05	9.76E-06	1.00E+07	3.08E+05
1970s	7.48E+07	1.29E-06	1.44E-07	5.56E+07	1.23E+06
1980s	1.50E+08	1.03E-07	1.18E-08	8.02E+07	6.17E+05
1990s	4.02E+10	4.54E-11	5.17E-12	1.82E+09	1.54E+08
2000s	8.39E+11	2.99E-13	2.97E-14	1.80E+10	4.93E+09

Table 4. Basic Performance Characteristics by Epochs of Computing

Source: Each year takes the median of computers or devices for that period. Each series is indexed so that the value for manual computing equals 1.

Period	Computer power (units per second)	Total cost per million unit computer power (2006 \$)	Labor cost of computation (hours per unit computer power)	Cycle speed (cycles per second)	Rapid access memory (bits)
----- Growth, percent per year, logarithmic -----					
19th C v manual	6.2	-3.8	-5.5	2.3	6.0
1900-1939 v 19th C	2.4	-1.1	-1.3	3.1	6.1
1940s	19.0	-5.3	-9.6	43.7	11.1
1950s	47.2	-38.9	-42.1	21.8	25.6
1960s	44.0	-42.7	-44.7	23.5	45.4
1970s	25.1	-31.3	-32.6	13.3	10.7
1980s	7.9	-28.9	-28.6	4.2	-7.9
1990s	53.2	-73.5	-73.5	29.7	52.5
2000s	36.7	-60.6	-62.3	27.7	41.8

Table 5. Growth Rates of Different Performance Characteristics of Performance in Different Epochs of Computing (average annual logarithmic growth rates, percent)

Source: Table 4.

Variable	Coefficient	Std. Error	t-Statistic	Decline rate for decade (percent per year)
Year	-0.006	0.010	-0.632	-0.6%
DUM1945	-0.445	0.065	-6.798	-45.1%
DUM1955	0.026	0.108	0.240	-42.5%
DUM1965	0.062	0.098	0.629	-36.3%
DUM1975	0.147	0.078	1.875	-21.6%
DUM1985	-0.527	0.083	-6.333	-74.3%
DUM1995	0.135	0.114	1.190	-60.8%

Note: Year is calendar year. Dum"year" is a variable that takes on value of zero up to "year", and value of year minus "year" after "year".

The growth rate is logarithmic.

Dependent variables is ln (cost per CPS in 2006 dollars).

N = 235

Table 6. Regression Analysis for Trends in Computing Power

Regression shows the trend in the logarithm of the deflated price of computer power as a function of year and time dummies. The last column shows the cumulative sum, which can be interpreted as the rate of decline in cost for the decade or period shown in the first column.

	Comparison 1		Comparison 2		Comparison 3	
	WorldBench per constant dollar	SYSMark98 per constant dollar	WorldBench per constant dollar	SPEC per constant dollar	WorldBench per constant dollar	Computations per second constant dollar
Sample period	Dec 98 - Nov 99		1995-1996		1992 - 2002	
Number of computers	30	30	7	7	47	37
Logarithmic growth rate						
Coefficient	0.24	0.38	0.66	0.67	0.52	0.62
Standard error	0.19	0.22	0.05	0.20	0.02	0.03
Arithmetic growth rate in computer performance per year	27.1%	46.2%	93.5%	95.4%	68.2%	85.9%

Table 7. Comparison of Different Benchmarks

This table shows the results of three sets of comparison between alternative benchmark tests of computer performance. In each case, the variable examined is computer performance per constant dollar, deflated by the consumer price index. The first two comparisons use exactly the same machines, while comparison 3 uses different machines over the same period. The different benchmarks are described in the text.

Study	Period	Method	Rate of real price decline (percent per year) [g]	Source
Government price data				
Price index for computers and peripherals (BEA)	1990-2004	Hedonic	-17.8%	[a]
	1969-2004	Hedonic	-18.7%	[a]
Price index for personal computers (BEA)	1990-2004	Hedonic	-28.6%	[b]
Producer price index (BLS)				
Semiconductors and related devices	1990-2004	Hedonic	-50.5%	[c]
Personal workstations and computers	1993-2004	Hedonic	-31.2%	[c]
Academic studies				
Berndt and Rappaport, personal computers	1989-1999	Hedonic	-38.3%	[d]
Chwelos, desktop computers	1990-1998	Performance	-37.2%	[e]
This study				
Price of computer power (\$ per MCPS)	1969-2005	Performance	-50.7%	[f]
Price of computer power (\$ per MCPS)	1990-2002	Performance	-57.5%	[f]

[a] BEA web page at www.bea.gov, Table 5.3.4

[b] BEA web page at http://www.bea.gov/bea/dn/info_comm_tech.htm

[c] BLS web page at www.bls.gov.

[d] Landefeld and Grimm., op. cit.

[e] Chwelos, op. cit.

[f] From regression of logarithm of price on year for period.

[g] All prices use price index for GDP to deflate nominal prices.

Table 8. Comparison of Price Indexes for Different Studies

This table shows estimates of the decline in prices of computers from different studies and methodologies.




Figure 1. The progress of computing power measured in computations per second (CPS)

The measure shown here is the index of computing power. For a discussion of the definition, see text. The unit is manual computing power equivalents. The larger circles are estimates that have been judged relatively reliable, while the small circles are estimates in the literature that have not been independently verified.

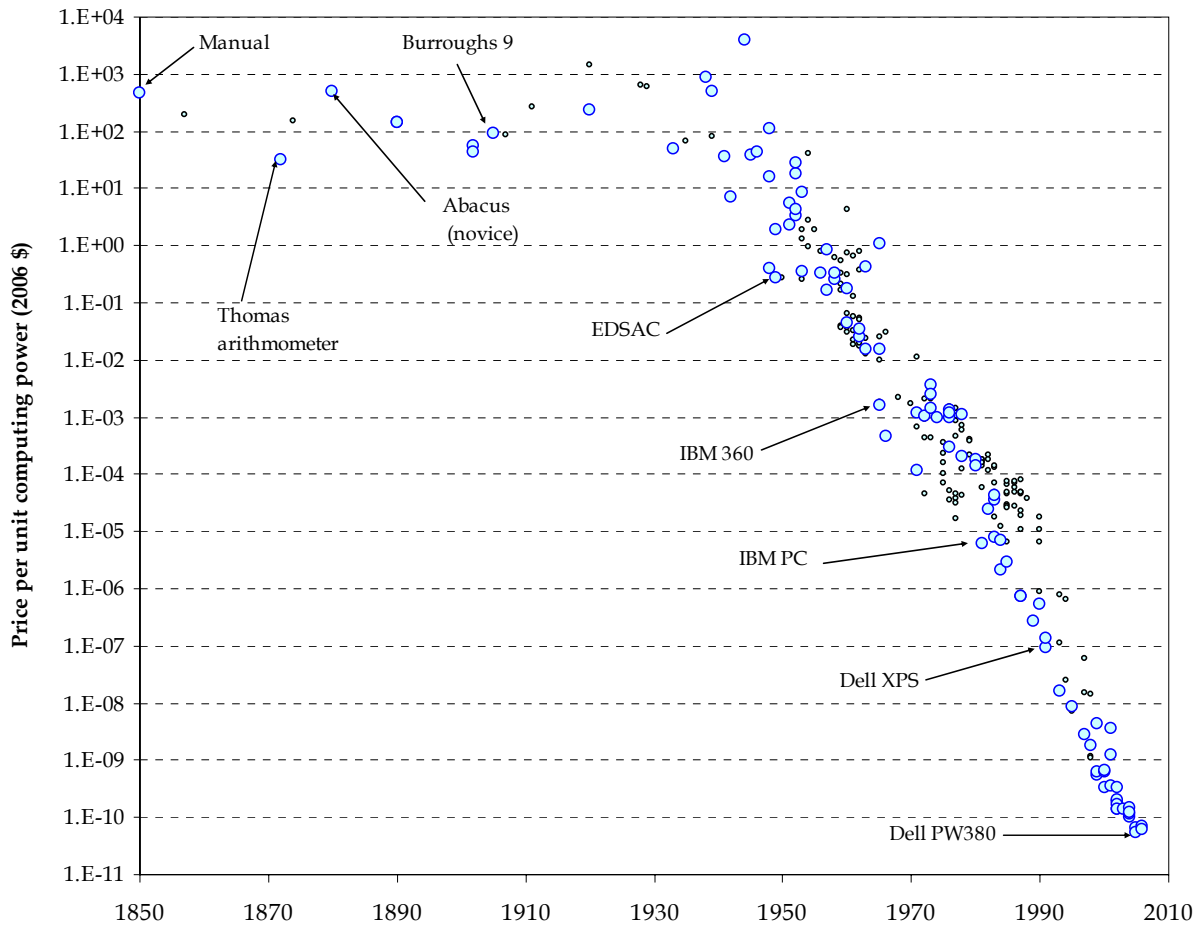


Figure 2. The progress of computing measured in cost per computation per second deflated by the price index for GDP in 2006 prices

Source: The larger circles are estimates that have been judged relatively reliable, while the small circles are estimates in the literature that have not been independently verified. Data table as described in text.