2011

Extrapolating the Price to Performance Frontier for Computer System Components: Processing, Storage, Memory, and Network Interface

G. Kent Webb, *San Jose State University*

Extrapolating the Price to Performance Frontier for Computer System Components: Processing, Storage, Memory, and Network Interface

Dr. G. Kent Webb, San Jose State University, San Jose, CA

ABSTRACT

As new computer hardware becomes available offering better performance at a lower price, computer accessibility rapidly improves resulting in dramatic changes to society. Planners in business and other organizations need an estimate of future prices and performance to help design their systems or to anticipate the effect of these changes. This paper presents a new set of historical annual data from 1987 to 2010 defining basic price to performance measurements for computer components including processors, hard drives, random access memory, and network interface cards. Two approaches to extrapolating price to performance are evaluated, the industry learning curve and a constant rate of increase implied by Moore's Law. Regression analysis of this new dataset shows long-term, stable improvements in price to performance consistent with Moore's Law provide a very good fit of historical data and a better approach to extrapolating futre price to performance than a learning cuve approach. Practioners can apply basic percentage changes to make reasonable forecasts that may be modified by short-term market fluctuations

INTRODUCTION

Over the past 25 years, computers that were once large and expensive with limited availability have become widely available as smaller, cheaper devices. How long will this trend continue and at what speed? Improvements in the manufacturing of computer hardware combined with an expanding market resulting from lower prices are key factors determining the trend. In the semiconductor and disk drive industries, the most important feature of each new generation of hardware results from manufacturers learning to work at smaller and smaller scales. The result is a long-term average total cost curve for the industries that is steeply downward sloping, a reflection of the manufacturer's learning curve.

In order to provide a backdrop to the historical data collected for analysis, price and performance measures, a short history of computer industry events that have influenced the shape of this learning curve appears in the next section. The heat barrier that recently disrupted the long-term trend of increasing processor speed provides an example of the circumstances that may bring an end to the benefits of the learning curved for this industry This information provides for major components of a computer system: processor, disk drive storage, random access memory, and network interface.

A review of the literature on technology learning curves reveals a standard formulation applied to a variety of environments, although much of the literature is devoted to improving specification of the learning curve model. Previous studies on the learning curves or average total costs curves for computer equipment have been remarkably stable, although these studies have been limited to fewer component categories or a shorter period. Using regression analysis to calibrate the learning curve equations for each computer component group confirms the strong pattern of improvement in price and performance evident in the data and a factor analsis suggests a stong commonality among the component learning curves. The remarkable stability of the advancement of the learning curve over time provides a basis for extrapolating these trends into the future, but special circumstances related to the manufacture of the network interface will likely dampen the trend of improvement for these components.

Market History

The first electronic digital computer that could be reprogrammed was contracted in June 1943 by the U.S. Army to calculate artillery-firing tables for use in World War II (Herman, 1972). The ENIAC (Electronic Numerical Integrator And Computer) was completed at a cost of $500,000 the year after the war ended, constructed in the basement of the engineering school at the University of Pennsylvania where it sits today -- over budget and too late for the war. A descendent of the ENIAC became the first commercially available computer, manufactured by Univac in 1951. The company later merged with Sperry Rand, and then with Burroughs, and is now known as UNISYS.

Moore's Law

The development of personal computers in the 1970s significantly expanded the market for technology. Among early personal computers, the Apple I introduced in 1976 was just an assembled circuit board; the keyboarded and display were added by the user. For $500 to $600 (about $2,500 in 2007 inflation adjusted dollars), a programmable computer was available to a mass market. The availability of lost cost components used to build the Apple computer was made possible by a trend that Gordon Moore of Intel had summarized as "Moore's Law" in 1965. Moore noted that the cost minimizing number of components per circuit in the manufacturing process was continually rising and that this should result in a doubling of price to performance every year (Moore, 1965). More recently, Moore's Law as applied to microprocessors has been revised to a doubling about every two years as improvements have slowed somewhat to a fairly constant rate since the 1970s (Webb, 2004).

Intel had also experienced a similar pattern of improvement in the cost per bit of memory, now the dynamic random access memory (DRAM) used for fast, short-term storage in computer systems. Moore makes a number of prescient predictions about these technologies in his 1965 article. Much of the article is devoted to an explanation of why the long run average total cost curve for the integrated circuit industry is steeply downward sloping and why this situation can be expected to persist for many years to come. Just a few years ago, Moore acknowledged that while an exponential pattern of improvement cannot last forever, the pattern still appears to have many years left to run (Moore, 2003).

Just a few miles south of where Moore and his group at Intel were working on processors and memory, the IBM Research facility in San Jose had been experiencing similar improvements in the price and performance of hard disk drives. In 1956, IBM sold the first commercially available hard drive for $10,000 per megabyte (MB), a $50,000 five MB drive. By 1980, hard drive costs had fallen to as low as about $200 per MB with current prices around $0.0002 per MB. In 2008, Western Digital introduced a terabyte drive available for personal computers.

A number of technologies have improved the networking of computers. One landmark development occurred when Bob Metcalf, while working at Xerox in the 1970s, invented a simple technology to connect copying machines, the beginning of Ethernet. Getting little support to develop products, he left to form 3Com, a company that began shipping network interface cards in the early 1980s. The network interface card (NIC) allows computer connections for communication over local networks or other wider networks like the internet.

Although designed as a cheap, uncomplicated technology to connect copying machines, Ethernet connections are now commercially available at 10-gigabit (Gb) speeds with higher speeds being tested in development labs. Integrating the network interface into the computer motherboard is another trend that substantially reduces the cost and improves reliability.

The improvements in the manufacturing process that reduce price per transistor for processors and the price per megabyte for storage and memory related to Moore's Law require manufacturers to learn to use processes that operate at increasingly smaller scales. Theoretical limitations on the scale of manufacturing related to quantum theory allow for more than a hundred years of development at the current pace; however other technical considerations have sometimes disrupted the advancement of the manufacturing learning curve.

At the 1996 Comdex show, an annual computer industry meeting, Intel predicted that by 2011 their processors would integrate one billion transistors and operate at 10 GHz. In October, 2004, however, the company scrapped plans to introduce a 4 GHz microprocessor and announced that it would focus on other ways to improve performance. The company has since languished in the 3 GHz performance range so it seems that the 10GHz speed range has become unattainable in the foreseeable because of problems in handling the increasing heat created by running processors at higher speeds. In February, 2008, the company announced a new chip with 2 billion chips for its quad core Itanium processor, well ahead of the 1996 forecast (Intel, 2008). While the improvement in price per transistor that was the basis for Moore's Law did hold true, the advancement in speed that many have incorrectly attributed to Moore's Law came to a dramatic end (Webb, 2004).

The Learning Curve

Lieberman discusses the strategic business importance of understanding the nature of the industry learning curves that drive long run costswhich can provide competitive advantage (Lieberman, 1987). He also summarizes the results of much of the early work and model specifications. A common form used for the learning curve to model technology and other processes is the power function

$$Y=aX^b$$ Learning Curve, Equation 1

where Y is the cost of the technology, b measure the progress over X which will be specified as time in this analysis but which has also represented cumulative installation of the product over time in other studies. For the learning curve on individual product life cycles, the cumulative installation variable appears to some as more theoretically appealing than the time trend approach, but it is more difficult to apply in extrapolating the trend forward since it requires a forecast of future installations. By comparison, the constant rate of change implied by Moore's Law can be summarized by equation 2.

$$Y = ab^X$$ Constant Rate of Change (Moore's Law), Equation 2

A number of learning curve approaches have been used (Doms and Forman, 2005; Kim, Jeon-Dong, and Tai-Yoo, 2005) relying on various transformations of the dependent variable of price to performance over time. For example, a logistic regression model to measure the change in a price index created to measure the price per megabyte for hard disk drives (Webb, 2003), results similar to a simple extrapolation forward of the average percentage change in price. Some researchers have used a hedonic regression (Victor and Ausubel, 2002) to model the learning curve for computer equipment, but the limited timeframe and lack of technological variables in the data used by the study limited the results.

A number of studies have investigated the price per transistor over time. Aizcorbe (Aizcorbe, 2005) notes that during the mid 1990s competition with processor manufacturer AMD seemed to stimulate an increase in the learning curve as a result of more rapid product introductions and shortend product life cycles. Flamm provides a clear analytical framework explaining how "technological trends in the semiconductor industry are ultimately reflected in the dynamics of costs and, in the long run, semiconductor integrated-circuit prices" (Flamm, 2003). Barnett, Starbuck and Pant (Barnett, Starbuck, and Pant, 2003) verify that Moore's Law has provided good long-term forecasts for processors and argue that on factor that has helped keep the forecasts accurate for almost 40 years and find that forecasts for a variety of industries tend to be more accurate in highly concentrated and controlled industries.

Figure 1: Product Pricing and the Learning Curve

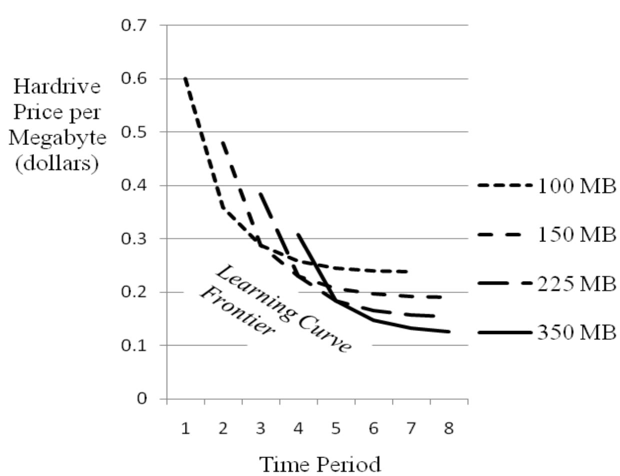

Figure 1 illustrates how the basic market mechanism for improving price and performance by illustrating recent representative prices for four hard disk drives each with different capacities. A hard drive with 100 Megabytes (MB) of capacity is introduced in period one for a price of $0.6 per MB, or about $60 for the hard drive.

Given the economics of scale in the production process, the price of this product declines over the product life. Prices tend to decrease rapidly in the first years, but the percentage rate of decline moderates over the life of the product. In the next year, a higher capacity 150 MB hard drive is introduced with a higher price both per hard drive and per MB, but the more rapid rate of price decline for the new model soon overtakes the price to performance of previous version of the product. Within a year or two, the new 150 MB drive represents the best price and performance combination, the frontier of the manufacturing learning curve in Figure 1. With each new successive product introduction, the frontier of the learning curve allows for a lower price per megabyte for the product.

Historical Data

Table 2 Historic "Best Prices" for Computer Components

Year	Processor price per Transistor (cents)	Hard Drive price per Megabyte (dollars)	DRAM price per Megabyte (dollars)	Network Price per Megabit (dollars)
1987	0.0595	65	15.01	39.9
1988	0.0467	45.2	14.16	34.5
1989	0.0334	33.6667	13.94	37
1990	0.0248	9.25	6.18	25
1991	0.0179	7.11	4.43	21.9
1992	0.0134	4.34	2.79	19.9
1993	0.0106	2.02	2.99	6.4
1994	0.0083	0.95	3.00	3.49
1995	0.0075	0.8707	3.02	1.76
1996	0.0058	0.2394	0.95	1.11
1997	0.0038	0.11252	0.41	0.78
1998	0.0016	0.06537	0.18	0.67
1999	0.0011	0.02565	0.185	0.53
2000	0.00097	0.01244	0.083	0.47
2001	0.00071	0.00605	0.0187	0.42
2002	0.00065	0.00263	0.0293	0.39
2003	0.00053	0.0014	0.0142	0.20
2004	0.00047	0.00126	0.01172	0.089
2005	0.00032	0.00113	0.0079	0.064
2006	0.00025	0.00085	0.0133	0.059
2007	0.00012	0.00057	0.0072	0.046
2008	0.000050	0.00019	.0020	0.035
2009	0.000038	0.000117	0.00074	0.0244
2010	0.000026	0.000069	0.00056	0.0189

Table 2 provides average annual prices for performance measures of major hardware components of networked systems from 1987 to 2010. This data was collected from an extensive search of archived industry publications and internet sites. Individual product announcements in industry publications account for most of the pricing information. The data represent full retail pricing and so may be somewhat higher than studies relying on wholesale prices. Prices for each year represent the best available price to performance available. For example, processor prices appear in the table as the price per transistor, consistent with the original concept from Moore's law, and represent the lowest price per transistor available during mid-year. Hard drive and DRAM prices are reported as price per megabyte, a standard industry measure. Price for network interface cards are measure per megabit (Mb) of transmission speed. Prices were generally lower than trend in 2008 with a weak economy and signficant over capacity in the DRAM market where prices fell about 75 percent after a relatively strong 2007.

TREND AND LEARNING CURVE ANALYSIS OF THE DATA

As summarized in Table 3, the average annual decline in prices for hard disk drives has been about 40 percent, with prices for other components falling in the twenty percent range per year. Variability of price changes as measured by the standard deviation of the percentage change has been highest in the DRAM market where volatility in demand and price fixing among manufacturers has resulted in a few years of unusual price behaviour.

Table 3: Average Annual Change in Price 1987 to 2010

	Processor price per Transistor (cents)	Hard Drive price per Megabyte (dollars)	DRAM price per Megabyte (dollars)	Network Price per Megabit (dollars)
Percent Change	-27.06%	-42.12%	-25.05%	-25.70%
Standard Deviation	0.141	0.199	0.397	0.191

A number of regression models were run against the data to identify patterns in the percentage change of the price to performance metric for each equipment group. Prior year percentage change of all variables proved to be insignificant in estimating percentage change for any of the variables indicating that data from previous years do not help predict one year forward. A number of time trend specifications were also tested to see if the percentage change in price is constant or changing over time, but none of these efforts provided any statistically significant patterns.

The following tables report the results of applying the basic learning curve, equation 1, and the constant rate of improvement for each of the computer components: processor, hard drive, memory, display, and network interface. For the learning curve: a linear regression was applied using the natural log of the price data from table 2 and the log of the time trend with the idea that prices change at a changing rate. For the constant improvement approach, the independent variable is simply the time trend, based on the idea of Moore's Law that prices change at a constant rate.

The results for processors in table 4 show the basic pattern. While both models fit the data well based on the P-values and R-Square, the constant rate of improvement model results in a significantly higher R-Square of 0.984 compared to the learning curve approach with an R-Square of 0.832. P-values are also significantly better for the constant improvement model.

Table 4: Processor, Regression Analysis of Historic Data

Processor Learning Curve Dependent Variables: Natural Logarithm of Price per Transistor Independent Variables: Natural Log of Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	-0.997	-1.870	0.0762
Learning Rate "b"	-2.261	-9.947	.000000003
R-Square = 0.832			
Processor Constant Improvement Dependent Variables: Natural Logarithm of Price per Transistor Independent Variables: Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	-2.400	-20.876	4.67 E-15
Improvement Rate "b"	-0.311	-35.623	1.43 E-19
R-Square = 0.984			

A similar result is reported in table 5 for the price per megabyte of hard disk drives. The R-Square for the constant improvement model is 0.988 while the learning curve R-Square is 0.864. The better P-Values for the constant improvement model are also significantly better, supporting the constant improvement approach as the best model.

Table 5: Hard Drive, Regression Analysis of Historic Data

Hard Drive Learning Curve Dependent Variables: Natural Logarithm of Price per Megabyte Independent Variables: Natural Log of Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	7.990	8.213	7.75 E-08
Learning Rate "b"	-4.677	-11.237	4.06 E-10
R-Square = 0.864			
Hard Drive Constant Improvement Dependent Variables: Natural Logarithm of Price per Megabyte Independent Variables: Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	4.968	24.614	1.99 E-16
Improvement Rate "b"	-0.633	-41.206	8.10 E-19
R-Square = 0.988			

In Table 6, the results again favor the constant improvement change to the price per megabyte for DRAM. The R-Square for the constant improvement model is 0.965 compared to the 0.786 for the learning curve approach.

Table 7, investigating the price per megabit for the speed of a network interface card, also shows that the constant improvement model works better with an R-Square of 0.987 compared to the learning curve R-Square of 0.846. Again, P-values are also better in the constant improvement model.

Table 6: DRAM, Regression Analysis of Historic Data			
DRAM Drive Learning Curve Dependent Variables: Natural Logarithm of Price per Megabyte Independent Variables: Natural Log of Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	5.707	6.598	1.99 E-06
Learning Rate "b"	-3.162	-8.573	3.94 E -08
R-Square = 0.786			
DRAM Drive Constant Improvement Dependent Variables: Natural Logarithm of Price per Megabyte Independent Variables: Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	3.842	15.549	1.24 E-12
Improvement Rate "b"	-0.444	-23.583	4.55 E-16
R-Square = 0.965			

Table 7: Network Interface Card (NIC), Regression Analysis of Historic Data			
NIC Drive Learning Curve Dependent Variables: Natural Logarithm of Price per Megabit of Speed Independent Variables: Natural Log of Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	6.0080	10.158	2.43 E-09
Learning Rate "b"	-2.676	-10.485	1.42 E-09
R-Square = 0.846			
NIC Constant Improvement Dependent Variables: Natural Logarithm of Price per Megabit of Speed Independent Variables: Time Period (starting with period 1)			
	Coefficient	**t-statistic**	**P-value**
Intercept "a"	4.363	24.863	1.63 E-16
Improvement Rate "b"	-0.363	-41.206	2.85 E-17
R-Square = 0.987			

CONCLUSIONS

A simple constant rate of improvement model summarizes with significant accuracy the change in the price of computer components over long time periods. The learning curve approach may be important to modelling prices over the product life cycle, but does not seem to apply as well to the long-term pattern of improvement.

Fundamental advancements in manufacturing economies as summarized by Moore's Law will likely continue to drive significant improvements in price and performance. If these trends persist, in five years the price per transistor for a processor would be at .0028 cents each, about 23 percent of the current price. The price of a transistor, a relatively sophisticated device, is now about $1/50^{th}$ the price of a simple metal paper clip and the price continues to fall.

The manufacturing advancements related to network interface cards are somewhat different from those of processors, DRAM, and hard drives which bascially involve manufacturing at smaller scales while the network card requires an advancement of speed. As with processors where a speed barrier changed the strategy of advancement a few years ago, network interface cards may also encounter a similar speed barrier. One symptom is that the product life cycle for high-speed Ethernet cards has been lengthening in the past few years. Although this pattern was not strong enough to produce a statistically significant effect in the model estimated for this study, it does suggest a potential slowing of the future improvements for this technology. No technology barriers appear evident that would disrupt the progress of improvement for the price to performance measures for processors, storage, or memory over the next 5 to 10 years.

One important conclusion of ths research for practicioners is that the constant rate of change approach suggests that simple average percentage rates of change (see table 3) can be applied to forecast future price to performance. This approach will give the same result as the more complicated log model discussed in this paper, but can be easier to work with and much easier to explain to a non-technical audience.

REFERENCES

Aizcorbe, A. (2005). "Moore's Law, Competition, and Intel's Productivity in the Mid- 1990s" *The American Economic Review* 95(2) 305 – 309.

Barnett, M. L., Starbuck W.H., and Pant, N.R. (2003). "Which dreams come true: Endogeneity, industry structure and forecasting accurary," *Industrial and Corporate Change*, 12(4), 653 – 672.

Doms, M. and Forman, C. (2005). "Prices for Local Area Network Equipment" *Information Economics and Policy*. 17, 365 – 38

Flamm, K. (2003). "Moore's Law and the Economics of Semiconductor Price Trends," *International Journal of Technology, Policy, and Management*. 3(2), 127-141.

Goldstine, Herman H. (1972). *The Computer: from Pascal to von Neumann*.

Princeton, New Jersey: Princeton University Press. ISBN 0-691-02367-0.

Intel Inc. (2008). http://www.intel.com/pressroom/kits/quickrefyr.htm#2008 (also the source of the transistors per processor data).

Jerz, D. G. (2007). "Somewhere Nearby is Colossal Cave: Examining Will Crowther's Original "Adventure" in Code and in Kentucky" *Digital Humanities Quarterly*,: v1, n2, Summer http://www.digitalhumanities.org/dhq/vol/001/2/000009.html

Kim, W., Jeong-Dong L., and Tai-Yoo K. (2005). "Demand forecasting for multigenerational products combining discrete choice and dynamics of diffusion under technological trajectories" *Technological Forecasting and Social Change*. 72(7), 825-849

Lieberman, M. B. (1987). "The Learning Curve, Diffusion, and Competitivve Strategy," *Strategic Management Journal*. 8(5), 1987, pp. 441 – 453.

Moore, G. (1965) "Cramming More Components Onto Integrated Circuits," *Electronics*, (38) 8. (Available at: http://www.intel.com/research/silicon/moorespaper.pdf)

Moore, G. No Exponential Is Forever…But We Can Delay 'Forever'" (2003). *International Solid State Circuits Conference*, (The slideshows from this presentation are available on Intel's Website at:

ftp://download.intel.com/research/silicon/Gordon_Moore_ISSCC_021003.pdf

U.S. Display Constortium (2003). http://www.usdc.org/technical/USDCroadmap_ExcutiveSummary.htm

Victor, N. and Ausubel, J.H. (2002) "DRAMs as Model Organisms for Study of Technological Evolution" *Technological Forecasting and Social Change* 69(3):, 243-262.

Webb, G. K. (2003) "A Rule Based Forecast of Hard Disk Drive Costs", *Issues in Information Systems*, 4(1), 337 – 343.

Webb, G. K. (2004). "Predicting Processor Performance", *Issues in Information Systems*, 5(1), 340 – 346.