





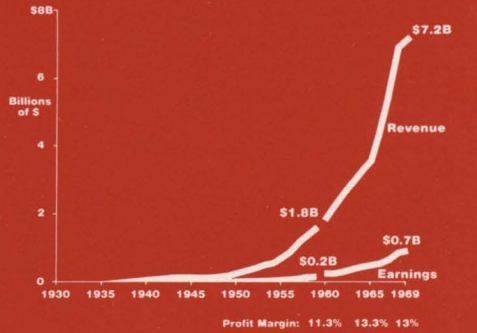

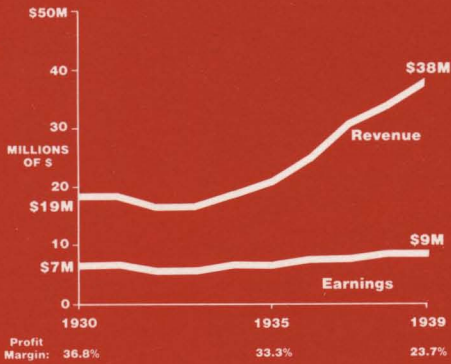
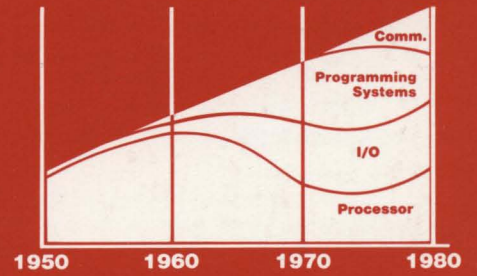
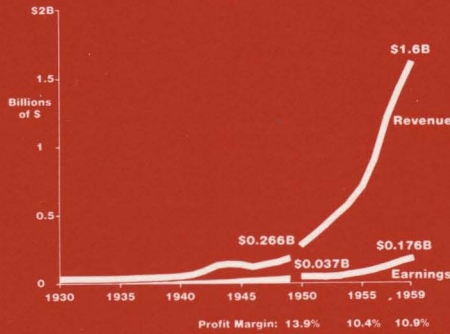



THE COMPUTER MUSEUM REPORT

NUMBER 9

SUMMER 1984

1911



Inside:

IBM System/360
Bob O. Evans

Computer Engineering Attitudes
Bernard Gordon

Contents

- 1 The Director's Letter
- 2 Computer Engineering Attitudes
From Eckert-Mauchly to Analogic
Bernard Gordon
- 9 IBM System/360
Bob O. Evans
- 19 Museum News
- 20 Membership Coupon
- 21 Founders
The End Bit
A newsbrief of the collection

Photo Credits:

The Hagley Museum and Library pp. 2, 4, 6
IBM Archives pp. 8, 10, 12, 14, 17
Phil Rosen pp. 7, 18
Meredith Stelling pp. 19, 20
Linda M. Watson (back cover)
Chris Kilby pp. 1, 19

THE COMPUTER MUSEUM

The Computer Museum is the only institution of its kind in the world, chronicling the evolution of information processing from the abacus to the silicon chip through exhibitions, publications, historical research and programs.

It recently moved from Marlboro, Massachusetts to Museum Wharf in downtown Boston, where it will open this fall. The Museum's new location is minutes from Logan International Airport and just a short walk from Boston's financial district and such historic landmarks as Faneuil Hall and the Freedom Trail.

The Museum offers individual and corporate memberships for \$25 and \$125 per year respectively. Non-profit organizations may join at \$25 per year. Other membership categories are available for those seeking a higher level of participation. All members receive a free subscription to The Computer Museum Report, a 10% discount on merchandise from The Computer Museum Store, free admission and invitations to Museum events.

For more information, contact Jana Buchholz, Membership Coordinator at The Computer Museum, 300 Congress Street, Boston, MA 02210, (617) 426-2800.

THE COMPUTER MUSEUM REPORT (ISSN 0736-5438)

The Computer Museum Report is published quarterly by The Computer Museum, 300 Congress Street, Boston, MA 02210. A yearly subscription to The Computer Museum Report is free with membership. Individual issues can be purchased through The Computer Museum Store for \$3 apiece.

The Museum staff is responsible for the contents of the Report. The views expressed do not necessarily represent those of The Computer Museum or its Board of Directors.

Design and production of the Report is done by Benson and Clemons.

STAFF

Director

Dr. Gwen Bell

Administration

Geri Rogers, Coordinator
Margherita Ciampa

Business

Eva Radding, Coordinator

Communications

Stephanie Haack, Coordinator
Brenda Erie

Exhibits and Archives

Meredith Stelling, Coordinator
Dr. Oliver B.R. Strimpel, Visiting Curator
Beth Parkhurst, Research Fellow
Paul Ceruzzi, Research Associate
John McKenzie
Gregory Welch
Bill Wisheart

Membership

Jana Buchholz, Coordinator

Store

Lea Cohen, Coordinator

The Computer Museum
300 Congress Street
Boston, Massachusetts 02210

BOARD OF DIRECTORS

John William Poduska, Sr., Chairman
Apollo Computer, Inc.

C. Gordon Bell
Encore Computer Corporation

Dr. Gwen Bell
The Computer Museum

Erich Bloch
International Business Machines

Harvey D. Cragon
Texas Instruments

David Donaldson
Ropes and Gray

Robert Everett
The MITRE Corporation

Sydney Fernbach
Computer Consultant

C. Lester Hogan
Fairchild Camera and Instrument
Corporation

Theodore G. Johnson

Mitchell Kapor
Lotus Development Corporation

Dr. Koji Kobayashi
NEC Corporation

John Lacey
Control Data Corporation

Patrick J. McGovern
CW Communications, Inc.

James L. McKenney
Harvard Business School

George Michael
Lawrence Livermore Laboratories

Dr. Arthur P. Molella
The National Museum of American History

Kenneth H. Olsen
Digital Equipment Corporation

Brian Randell
University of Newcastle-upon-Tyne

Jean E. Sammet
International Business Machines

Edward A. Schwartz
Digital Equipment Corporation

Kitty Selfridge
Henco Software, Inc.

Erwin Tomash
Dataproducts

The Director's Letter

As I sit and write this letter, carpenters are working in a gaping hole that will connect the two levels of the Museum with a stairwell. The Museum is between acts. Looking back six months, the memories of the fall programs in Marlboro are still fresh, while the staff's attention is focussed on the new exhibitions. This issue of the Report reflects the transition.

The major articles are versions of two outstanding programs from last year. Bob O. Evans, vice president of engineering, programming and technology at IBM presents an inside view of product development from meatgrinders to the IBM System/360. Then Bernard Gordon, president of Analogic reminisces about his second job, working for J. Presper Eckert at the Eckert-Mauchly Computer Corporation. Both of them breathe life into giant myths.


With the reopening months away, our concentration is focussed on planning new exhibitions. A half-inch scale model of the galleries filled with moveable machines, cases and people sits in the middle of the office space. It's a central gathering place where we test out ideas and give the rare visitor a "small scale" tour. When something doesn't fit or can't be explained, then it's modified. The model is rapidly filling with a variety of environmental, historical, operational and interactive exhibits. The major historical exhibit is a year by year timeline from 1950 to 1973 under the direction of Professor Paul Ceruzzi of Clemson University. The largest environmental exhibits are the AN/FSQ-7 and the SAGE installations being developed by Meredith Stelling of the exhibit staff. Working exhibits include printers and plotters and other devices that will operate on their own. Although the personal computer exhibit will have a concentration of interactive exhibits, every gallery will have visitor-generated programs that create understanding through hands-on experience. A number of college interns will be developing these interactive programs this summer.


The May 11th Pre-Preview Party provided an early deadline to get our stairwell and auditorium in place and to let us take advantage of the free space for both dining and dancing . . . activities that will be precluded by the full development of the exhibitions someday. It also kicked off the capital campaign described elsewhere in this issue.

This spring and summer, I'll be on the road visiting Minnesota, Washington state, California, New York, Europe and perhaps even Japan, spreading the message about our purpose as an international museum for the entire history of information processing. When the Museum reopens this fall, its exhibitions and programs should attract both national and international visitors when they're in Boston.

Gwen Bell

Bay 4 on floors five and six readied for a central staircase.





Computer Engineering Attitudes From Eckert-Mauchly to Analogic

Bernard Gordon

In 1948 Bernard Gordon graduated with a bachelor's and master's degree in electrical engineering from Massachusetts Institute of Technology (MIT). After starting with Philco Corporation, he joined the Eckert-Mauchly Computer Corporation. Today he is the president and technical director of Analogic Corporation which is engaged in the development and manufacture of high-precision, high-speed signal translation and information processing equipment. The following abbreviated and edited excerpts have been derived from a lecture presented by him at The Computer Museum on October 20, 1983. For historical purposes, the original presentation has been archived at the Museum on videotape.

About a year after I left MIT to start work at Philco Corporation, I received a call from Presper Eckert who told me I had been recommended by a professor at MIT and asked if I would come over for a job interview. Eckert, then about 28 years old, gave me such an intense technical and personal interview that even before he made me a job offer I told him I'd take the job just because he had motivated me to show him what I could do. He was so taken aback by this that, I guess, he felt he had to hire me and so he did.

Therefore, in 1948 on a hot summer day I reported to work at the Eckert-Mauchly Computer Corporation in Philadelphia in an old building near Wissahickon Park. One of my first memories is that of seeing Al Auerbach, now a long-time friend, standing literally in his underwear working in the middle of the heat of the circuitry which was supposed to become the BINAC, forerunner of the UNIVAC. As I recall most of the small group of engineers were nearly all in their twenties. The chief engineer was Jim Weiner who had come down from Raytheon. Jim ruled over us like a master sergeant and engendered in us reactionary passions . . . but he made us do our jobs. In later years I learned to bless him because he and Eckert inculcated in me, and I believe in the others who worked at the Eckert-Mauchly Computer Corporation, engineering disciplines which have served me well during the past 35 years.

It is interesting to note that Eckert-Mauchly had figured out that they would need about \$100,000 to engineer the UNIVAC and ready it for production. They had raised about this amount of money from the American Totalizator and had figured out this amount of money based literally on the number of solder joints in the machine and multiplying that by so many pennies. They therefore had predetermined the rate at which all of the work must be accomplished from logical design, the software, the electronic design, the construction, and the debugging.

Eight to ten engineers were to build, not knowing any better, all of the original circuitry for the UNIVAC and as well the first high-speed start-stop digital tape mechanisms, the

tape plating and manufacturing facilities for those tapes, the first-known card-to-tape converters, and the many other major sub-units of the UNIVAC system.

The machine was to have approximately 5,700 vacuum tubes, used primarily for amplification and pulse forming and 18,000 semiconductor diodes used primarily for high-speed gating. (It may be interesting to recall the semiconductor diodes utilized were purchased as war surplus materials from Western Electric.) When I arrived for work one of the engineers, Bob Shaw, had already essentially single-handedly drawn all of the detailed logic diagrams. I recall Eckert saying to me: "You are going to design the circuits, standard flip flops, standard gates, and so forth." He had allowed only a few working days to do this. I didn't know I couldn't do it, so I set out to do it. In a relatively short time, no more than a few weeks, we had designed and proven the capabilities of the standard gates; I then designed the I/O circuitry, supervisory control circuitry and tape control circuitry, standard flip flops and what we'd call pulse formers.

Eckert then set me to work to design the crystal transducer system for the acoustic memories of the UNIVAC and then all of the electronics for the memory system. The time allowed for each major design was always measured in days, not weeks or months. At that time I thought I was working on the world's first acoustic memories and it wasn't until a considerable time later that I found out that Maurice Wilkes, who is present at this lecture, had actually built a unit earlier in England. While I was carrying out this work together with the other engineers at the Eckert-Mauchly Computer Corporation, Pres Eckert and Jim Weiner taught me via their direction a number of factors about engineering and engineering supervision. I do recall that at the time we were receiving this type of direction we felt that they were very tough. But in the process of being apprentices to these master engineers, most of us went through a maturing and learning process which, in retrospect, I wouldn't have traded for anything. If in my later years I have myself developed a reputation for being a

J. Presper Eckert Jr.

tough engineering task master, I am pleased to say—and I hope that he will be pleased by my saying it—that Eckert was responsible.

For example, after Eckert more or less gave me a "gold star" for doing the acoustic memory, he put me in charge of a few other even younger engineers who were then being hired into the company. He gave me the following directive: "If you ever see an engineer studying during work hours, I want you to give him his first warning. If he does it a second time, terminate him." His view was, and it still remains mine today, that people owe it to themselves to further their career, to study at home, and that they should come to work prepared to get the physical work done.

The philosophy of "worst case design" probably originated, or at least was formalized, at the Eckert-Mauchly Computer Corporation. Eckert and Weiner insisted that when we design something, we must design it thoroughly, into the ground so to speak, and release our circuitry to production without ever breadboarding. In the first UNIVAC they established rules for derating such that every 25L6 vacuum tube must properly function in its circuitry with its emission dropping to approximately 50 percent with the screen voltage varying, with the heater voltages varying, with carbon resistors changing 20 percent, etc.

Although I didn't really prepare for this lecture in any formal way, as I stand here, I can remember the derating numbers of 35 years ago like a catechism. For example, every germanium diode which had a nominal back resistance of about a megohm with a back voltage of 30 volts had to continue to work satisfactorily if that back resistance went to 18,000 ohms. Every carbon resistor had to be able to change 20 percent and each power supply voltage had to change in the worst possible combination, about five percent. As a result, we were able to design with parts that really weren't very good and design equipment that could be predicted to work right essentially the first shot.


Eckert taught me to pay great attention to every detail. He taught me that the design engineer was responsible for every aspect of the design. The engineer should know how the components were made. What were their strengths and what were their weaknesses. There should be extreme tolerances on everything. He knew that only by doing this was it possible to make a machine with 5,700 vacuum

tubes each with a nominal emission life of about 5,000 hours work at all. However, by applying the rules of derating everything, it was possible to make a machine at that time which worked for acceptable periods of time.


At the end of every week, Eckert and Weiner would come around and we'd show them our big schematics with 40 to 100 vacuum tubes on them. He would look at a drawing, almost closing his eyes, and point to a resistor at random and say: "Why is that resistor that value? Why isn't it five percent higher? Why isn't it five percent lower? Show me in your notebook where you proved absolutely that that resistor is exactly the right value." I think I almost got fired one day because I had a grid resistor returned to ground, and he asked me why. I said that it was half way between plus and minus infinity, which was an unsatisfactory answer.

Every once in a while something humorous related to the disciplines that were put in effect would take place. For example, whenever the power came on the UNIVAC, a red light went on at the top of the machine's frame. Jim Weiner established the rule that whenever anybody made a mistake such as putting a screw driver or a scope probe in the wrong place and blew up a diode, he would have to buy a Coca-Cola for all the employees of the company, approximately 30. However, one day Jim Weiner himself put his screw driver into the wrong place and blew up all 18,000 diodes! It made us all feel much better. No one ever found out how he was able to blow them all up simultaneously, but he sure did.

I have always felt that Eckert conveyed a particularly important engineering philosophy to us. He felt, I believe, that any engineer worth his salt should be able to design anything at any time, either electrical or mechanical. If he didn't know how to do it, then it was his responsibility to go out and learn how to do it. I remember his saying to me: "When you go home tonight, your wife is going to want you to cut the grass. Don't do it. Hire somebody else to cut the grass who is a grass cutter, and you study and design for the company." He said: "This effort will come back to you many times in the future." I never did cut the grass and always felt as a result of his direction that it was my mother-in-law's job to take out the garbage and not mine! In any event, I have always spent continuously over the last 35 years two hours a day studying at home or at the MIT library or elsewhere . . . every day.



Al Auerbach and Jim Weiner (right), who according to Bernard Gordon, "established the rule that whenever anybody made a mistake such as putting a screwdriver or a scope probe in the wrong place and blew up a diode, he would have to buy a Coca-Cola for all of the company, approximately 30."



J. Presper Eckert Jr. is shown with a BINAC Mercury Memory Tank.

To engender his attitude, every once in a while Eckert would notify all the engineers that they would be given a written test. The test material generally had nothing to do with our then current work. The test material would touch upon a variety of subjects, such as the workings of an alternator or a power station or how to design a filter. If an engineer could not pass such a test, he was likely to be terminated. This, I believe, was Eckert's way of making sure that his engineers had a very broad interest and would be prepared intellectually to tackle anything that they had to. It was not unusual that one engineer such as myself would design wide band IF amplifiers one week and stainless steel tanks with crystal transducers for sonic mercury systems another week.

Eckert, through his personality and the fact that we were building the first commercial computer, got us very excited and interested in our work. Not only the theoretical and technical aspects but also the economic aspects. He used to get us to think in terms of how much everything cost, how much did the solder joint cost, how much did it cost to make a drawing, how much did it cost to have a secretary prepare a technical document, how many lines should a draftsman put on a piece of paper each day, etc.

To try to keep within his original \$100,000 budget, it was required at the Eckert-Mauchly Computer Corporation that every day one vacuum tube's worth of circuitry be released into production about every half hour by every engineer. There was no getting around it. Those were the standards set and that is what was expected. As I recall it was less than a year after the design started that the UNIVAC fully stood on the floor at the Eckert-Mauchly Computer Corporation complete with the first high-speed start-stop tape mechanisms, first acoustic memories, tape-to-card converters, ready to be system tested.

We probably didn't know it at that time, but nearly all of the engineers at the Eckert-Mauchly Computer Corporation were highly motivated by the atmosphere which I have briefly described. About the time of the completion of the UNIVAC 1, the then Sperry Rand Corporation bought out the company and the culture began to change as "big time management" attitudes began to permeate the company. Many of the original engineers, including myself, then began to leave the company. Eckert, who had been my mentor, said to me when I left: "You may

never build another computer again, but it is probably true that everything you build in the future will in one way or another resemble a computer." He was right.

After Eckert-Mauchly

Eckert's prediction was proved on my next job. I moved back to Boston because the weather was too hot and muggy in Philadelphia. In Boston I worked for a company called Laboratory for Electronics, founded and populated with very famous names from MIT's radiation laboratory and who indeed had made major contributions to the series of well-known books entitled RADIATION LABORATORY SERIES. The company wanted to build a computer. But since the guys from MIT wanted the job, they were given the opportunity. I was assigned to work on the development of a doppler navigating radar. One day we realized that every half cycle of the doppler return signal represented the distance that the plane had traveled. So we thought that if we could count these half cycles, in turn we could build a digital doppler radar. Thus, consistent with Eckert's prediction, the doppler radar which would normally have been an all analog system ended up resembling a digital computer.

It was also on this job that I met An Wang who had just started his own company. He and I built a sequenced number generator which resulted in patents for wire core memories. We pulsed stacks of magnetic cores in sequence and read them out to generate arbitrary codes for controlling our rate multiplying navigational computer.

In 1953 I decided that it would be useful to tie computers together with analog signals and built a device called a DATRAC, the first known shift programmed successive approximation A/D converter. At that time together with another gentleman named Joe Davis, I started a company called EPSCO, Inc. and began hooking up analog to digital converters to computers . . . an activity I have been heavily involved in ever since. Today at Analogic we build a variety of measurement devices that compute, varying from very sophisticated phased-array ultrasound medical imaging machines to high-speed signal processing computers.

I have consciously and unconsciously tried to follow some of the principles that I originally learned in my younger days when I worked at the Eckert-Mauchly Computer Corporation. At Analogic we expect that proj-

ect engineers should personally be able to do the variety of tasks required on their projects. We rarely put more than three or four engineers on even the most complicated equipment that we design, such as the very first instant imaging CAT scanner or signal processing communications computers that make hundreds of millions of computations a second. Our project engineers who can be assigned from project type to project type are the keystones of our company.

Very often people from all around the world ask us: what do we at Analogic do that is different to get the engineering productivity and stability of our engineering staff. I always answer by saying: "We don't do anything different. You are doing things different. We are doing the same old things that we learned 30 or more years ago."

Let's briefly look at how things used to be and how they are today. In 1948 there were about 2,000 electronic engineers being graduated in the United States. Today with 250,000 electronic engineers nominally at work in our society and about 17,000 graduating each year, a hue and cry is heard across the nation that there is a shortage of engineers. What is wrong? In the 40's, engineers were taught, in addition to the type of disciplines that I have referred to, a breadth of mathematics and physics. They could be prepared to do anything because they'd been taught fundamental principles.

Recently I attended a seminar where a speaker stated that "the complexity of current projects is such that the mind of a single project engineer cannot encompass the breadth of the work." That fellow was talking about an engineering work station. Another fellow made a similar point about personal computers. Those of you who are in the audience who are about my age know that this is nonsense, because we were all called upon, when we were younger, to build and be totally responsible for much more complicated things. Certainly there is not a heck of a lot of real physical hardware engineering in any personal computer. Any good engineer could design a personal computer hardware-wise in a few weeks. The software would clearly take longer. But for the hardware, he needs to have an organizational concept utilizing available chips or have them laid out in gate arrays, buy a display and storage elements, and essentially "glue" it together. It would probably take longer to get the tooling for the plastic case than to actually

design the personal computer. Bear in mind that with Eckert's \$100,000 engineering goal (even if that translates to \$500,000 today) he intended to design from scratch the world's first commercial computer, the world's first card-to-tape converter, the world's first commercial acoustic memories, etc. Keep in mind that there was typically a half to one engineer working on each subsystem.


Now, let's look forward a couple or so years when you will probably be able to hold in the palm of your hand a 10 megaflop 32-bit high-speed computer with about a million bits of memory whose factory cost will be \$200 or \$300. When such building blocks are available, much of the "beauty" of this fellow's computer architecture or that fellow's computer architecture will fall by the wayside. The tasks for computer-related systems will be more and more related to being able to harness that computing power and design and build useful real-world machines encompassing a breadth of technology.

Now, what has happened to engineers? I would like to state my opinions and I am aware that not everyone will agree. In most companies the attitudes of Pres Eckert or Jim Weiner are no longer taught nor is the mentor relationship available to most young engineers. It is very rare for a youngster out of school to go to work for a 28 year old truly experienced engineer. He is liable to go to work for another youngster who has only been out of school for two years, who in turn has worked for a youngster with a similar limited level of experience.

I believe that with about five percent supervision by a broadly experienced motivating engineer, less experienced engineers can increase their productivity somewhere between two and three times. At Analogic we jokingly call this "Gordon's Rule" and are certain that the theoretical parameter of improvement is "e" or 2.7183.


Now, some people such as the people developing work stations claim that by the appropriate use of engineering work stations it should be possible to increase engineering productivity by 4 to 1. Possibly they are right and possibly Gordon's Rule is right. Of course, if they are both right, then it must be possible to achieve a ten-fold increase in engineering productivity. If this is so, you would think that this combination would easily solve the engineering shortage!

However, in my opinion the reality is that the true problem is that there is a grave shortage of engineers whose



Eckert-Mauchly Computer Corporation, Engineering Group, 1949. Row 1: Fran Morello, Bob Shaw, Pres Eckert, Brad Shepard, Frazer Welsh, John Mauchly, Jim Weiner, Al Auerbach, Betty Snyder; Row 2: John Sims, Marv Jacoby, Paul Winsor, Gerry Smoliar,

Art Gehring, Betty Jay, Ed Blumenthal, Bob Mock, Jean Bartik, Herman Lukoff, Bernie Gordon, Ned Schreine; Row 3: George Gingrich, Marv Gottlieb, Lou Wilson, Doug Wendell, Charlie Michaels, Ben Stad, Si Levitt, Larry Jones.



BINAC Laboratory tables, May 12, 1949.

education and orientation gives them a very broad view. Recently I became concerned about the breadth of capability of many software engineers. The following may be instructive. When Eckert interviewed me in 1948, I had learned a great deal about "pole and zero-based transfer functions" by working at Philco. When I had to design an IF amplifier for the UNIVAC, I was able to achieve an overall transfer function by matching the effective poles of the transducers to an optimum complementary transfer function of the amplifier. Recently I started playing with my little home computer and just to exercise myself decided to write a program for an arbitrary number of poles and zeros to calculate the phase and amplitude transfer functions and the transient response. Having once known how to do this very well mathematically and particularly knowing the graphical interpretation of pole zero relationships, it took me only a few hours to achieve the result I desired. The next morning upon arrival at Analogic I asked one of our relatively new but previously experienced software engineers how long it would take him. His answer was six months! At first I was startled, but as I proceeded to talk to him, I discovered that he could probably write the program in two hours also . . . if he knew something about poles and zero mathematics . . . but that he felt it would take him five months 29 days and six hours to learn about poles and zeros! Then he could write the program.

We can all recount examples of projects where hardware engineers, software engineers, marketing people or customers could not interact effectively because they did not understand each other's needs. It is my belief that in the 36 years since I went to work for Eckert, I have witnessed a continual decline in the average productivity of engineers. Let's take a measure of it. Only 25 years ago it was common to say that there should be a development engineer for every million dollars worth of electronic production in the United States. Now, 25 years later, with an inflation factor of at least four and with the availability of CAD/CAM techniques, LSI and VLSI and all the other modern wonders, a computation of the total electronic output in the United States divided by the number of electronic engineers at work yields a value of only about a half a million dollars. This combined factor of eight is of great economic significance. It should cause most business managers and technical leaders to pause and

give consideration to whether they have allowed the standards of engineering excellence and productivity to decline in their own organizations.

Question and Answer Period

Q: Did Eckert ever really fire someone for failing one of his tests?

A: Yes.

Q: How could you keep working if you thought your job was on the line?

A: I'm not suggesting that somebody should be fired because they don't know something, but if they won't learn something, that is different. Not too long ago I fired a mechanical engineer who would not draw. He said that he thought up designs "in his head" and he would then translate his thoughts to a draftsman . . . and that it was beneath his dignity to draw. We found that he really *couldn't* draw and didn't want to learn. He had a degree in mechanical engineering . . . but had never taken a drafting course!! He's not atypical.

Q: What was the role of John Mauchly?

A: I believe that Mauchly was the original driving force behind the ENIAC. He was a professor at the University of Pennsylvania, and Eckert was a graduate student. They founded the Eckert-Mauchly Computer Corporation. At the time I was employed, Mauchly was somewhat less active for reasons, as I recall, that were very personal.

Q: Was there a strict hierarchy and structure?

A: Although Jim Weiner was the chief engineer, Eckert would often jump up to the top of a filing cabinet and sit on it and squat. He would take on the characteristics of a guru to anyone that was around at the time. As I recall there really wasn't a pecking order at all. He used to have what I thought was a wonderful idea of saying to people, "Say anything that comes to your mind. Idea. Idea. Idea. You have 99 inadequate ideas and maybe the 100th will be invaluable." Eckert would always engender an atmosphere where people would not be afraid to be wrong about anything. We all had a lot to learn and to conceive.

Note: Recently the Massachusetts Board of Regents has authorized the formation of a new institute to be called The Gordon Institute, a school of engineering leadership to be located in Wakefield, Massachusetts. Its aim will be, consistent with Eckert's philosophy, to teach engineers a

breadth of knowledge involving technology, ethics, and philosophy, considered to be musts for true leaders and to develop an orientation toward the successful economic accomplishment of projects undertaken.

October 20, 1983



Gordon and Wilkes. Bernard Gordon (left) president of Analogic Corporation is greeted by EDSAC designer Maurice Wilkes after Gordon's lecture on "Computer Engineering Attitudes From Eckert-Mauchly To Analogic" at The Computer Museum on October 20, 1983.



IBM System/360

Bob O. Evans

Bob O. Evans is IBM vice president, engineering, programming and technology. He joined IBM in 1951 as a junior engineer in Poughkeepsie, New York, where he took part in the development of IBM's first large scale computers. After various assignments in computer development, he was promoted in 1962 to vice president, development, for the Data Systems Division which included overall management responsibility for development of IBM System/360. The following article is based on a lecture presented by him at The Computer Museum on November 10, 1983. For historical purposes, the original presentation has been archived at the Museum on videotape.

The System/360 and its direct descendants have accounted for more than a hundred billion dollars worth of revenue and considerable profit for IBM and has been the foundation of our basic business for years longer than we anticipated. I wish to tell you something of the environment, actions and people who made System/360 happen.

IBM was formed in 1911. At that time it was called the Computing, Tabulating and Recording Company and was an amalgamation of three tiny companies that worked on products such as meat slicers, scales and nurse call systems. One part of the small firm was the Tabulating Machine Company that had been built upon the intellect of Herman Hollerith, inventor of the punched card. This little company, recording a few tens of thousands of dollars of revenue, grew slowly in those days. By the early 1930's, CTR had grown and, amazingly, had shed itself of most of the prior products, the nurse call system, the scoreboards and the meat slicers, concentrating upon the Hollerith concept to become an electric accounting machine company.

Several factors accounted for CTR's success: first, the strength of the Hollerith concept itself; second, the young leader who ran the company, Thomas J. Watson, who had come from the National Cash Register Company; and, third, the U. S. Social Security Act of 1931, which created an enormous demand for the types of machines CTR built.

The Computing, Tabulating and Recording Company's name was changed to International Business Machines in 1924. In 1930 IBM's revenue was \$19 million a year and then grew by 1939 to \$38 million a year. A more important measure of the effectiveness of the company is net earnings after tax—which were 36.8 percent in 1930. Of course the tax structure in those days was substantially different. Nonetheless in net: IBM was a healthy small company, producing electric accounting machines for a growing demand.

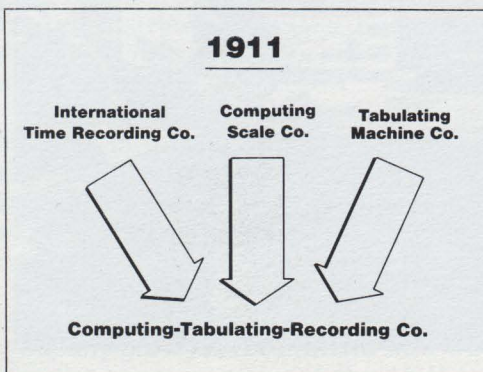
By 1949 IBM had grown to be a \$200 million a year business that primarily leased electric accounting machines. The view within was that IBM was the product leader in electric accounting machines; it was a profitable institu-

tion and investors loved IBM. If you had bought a few dollar's worth of stock then, you would not have to work now. IBM had very strong user loyalty, and most importantly, there were abundant opportunities for new electric accounting machines.

Let us examine IBM in the decade of the 40's in more detail. There was a revenue bulge that came during the war years as the company—like all U.S. industries—turned to the national effort. Then there was some downturn as the company recaught its breath after the war to return to its basic business direction. Profit was 20 percent of earnings in 1939 and by 1949 profit had grown to \$33 million or 18 percent net after tax.

In the national interest work during the war years IBM produced fire control systems and navigation and bombing systems among other products. From this IBM's Military Products Division grew, and was later renamed the Federal Systems Division, although the revenue of that division is today a small percentage of IBM's total.

New events led IBM to turn another radical corner. One often wonders how these things happen and, on reflection, the change was most unusual, for here was IBM doing well with electric accounting machines when the Korean war started. Shortly after the war broke out, Mr. Watson sent a telegram to President Truman offering IBM resources for the national effort. The consequence of this telegram was that two IBM executives surveyed the National Laboratories and other national interest work around the United States to determine what IBM might do. One was an engineer named Ralph Palmer, in my viewpoint one of the geniuses that IBM was fortunate to attract, who established the foundation of the IBM development community as it still exists today. The other was a master salesman, Dr. Cuthbert Hurd. Dr. Hurd and Mr. Palmer, under the aegis of the Watson telegram to the President, toured the U.S. They visited such places as Livermore, Los Alamos, National Security Agency and aerospace companies to determine how IBM might contribute. When they returned they told Mr. Watson the best thing IBM could do was build a high-speed computer much like the high-speed com-



Evolution of CTR, the Company that Became IBM.


puter that Dr. John von Neumann was building at the Institute for Advanced Study and Professor Maurice Wilkes was building at Cambridge University. They concluded there was great need for such computer power in national interest areas and that IBM should do it.

The government was not all that interested, so Mr. Watson, anxious to keep his pledge, decided that IBM would fund the effort, thus in 1950, the project began.

A principal viewpoint then in IBM was that such a project was an intrusion on the mainstream. The estimated demand for such electronic systems was ten or so and the prices were certain to be astronomical. Thus the view was the project was indeed a sacrifice, but IBM should get on with it and then get back to our basic EAM (electronic accounting machines) business as swiftly as possible.

The project was called the Defense Calculator and was formally named the IBM 701 Electronic Data Processing Machine. The first system was installed at IBM's World Headquarters in New York City in December 1952. I was lucky to be one of the engineers who went to New York City to get that system installed and operating. Nineteen 701's were built between 1952 and 1954. The rental for the system was a staggering \$20 thousand a month at a time when other IBM machines rented for \$300 a month or so. Thus, the 701 did not seem to have much promise. Fortunately, Mr. Watson's son, Tom Watson, Jr., saw the potential of electronics. He had become President of the company and pressed for more effort in electronic computers. You can imagine the reaction of some senior management. They knew the accounting machine business, they loved it and there were long lists of new EAM features and equipment needed to meet customer requirements. Thus many pressed to continue focusing on EAM. But Tom Watson, Jr. led the business into electronics.

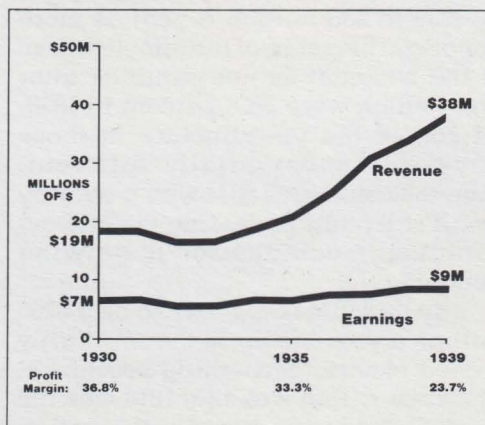
In the 1952 and 1956 era of vacuum tube technology, a number of computers came from IBM. The business computers were characterized by being alphanumeric, handling both alphabet and numbers, and operating serially by character on those voluminous strings of variable character length data. Business systems also had more extensive peripherals, usually tape drives, card machines and printers. In contrast to the business systems were scientific systems such as the IBM 701, which were parallel, binary and had more limited peripherals.



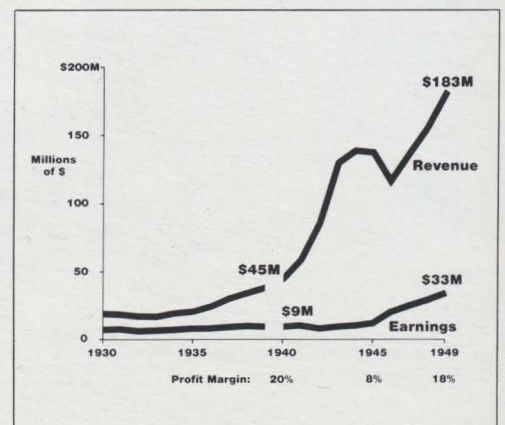
IBM 701 Electronic Data Processing Machine, 1952.



IBM 704, 1955.



IBM Growth: 1930-1939.



IBM Growth: 1940-1949.


In a short time an improved version of the 701 was produced called the IBM 704. Gene Amdahl had come to IBM from the University of Wisconsin, where, as his doctoral thesis he built a machine called the WISC. He, John Griffith, and a small group worked on the architecture of what became the 704 with new innovations such as floating point, indexing, and other bright new functions. Later core memory replaced the old Williams tube memories and, still later, the IBM 709 evolved from the 704 base. In that era another business computer was produced, the IBM 650, centering about a magnetic drum storage device. More than a thousand 650 computers were sold, far more than the forecast.

The 305 RAMAC was a new system conceived by Ralph Palmer, IBM's engineering genius. He wanted to see business data immediately accessible to the processors and he envisioned a disk device. Palmer set up a laboratory, IBM's third, in San Jose, California, to develop disk products. The 305 RAMAC became the first disk system that IBM produced. The sales forecast was for four or five thousand although fewer were sold.


Also on the business systems side, several hundred 702 and 705 systems were produced. They rented for more than \$30,000 a month, taking the place of a lot of sorters, collators, gang punches and calculators that were then the mainstream of IBM's business. Some 250 704's and 709's were sold to scientific users.

These big rental, big ticket items brought in a lot of revenue to IBM in that exciting period. So Tom Watson, Jr.'s hunch about electronics proved correct and IBM was on its way into a new era.

How did the business do? Through the decade of the 30's and the 40's the company grew to \$200 million. Now we see the consequences of the shift to



IBM Growth: 1950-1959.



IBM Vacuum Tube Computer Families: 1952-1956.

electronics for, in the decade of the 50's, IBM grew swiftly to approximately a billion dollars in 1957, and in the following two years to \$1.6 billion, fueled by our movement into electronics.

Some companies working on the early computers were ahead of IBM, and I would have to say that IBM was able to succeed so well because of our marvelous sales force and outstanding service which were the keys to IBM's ability to grow from a small company to the very significant company we became in the 1950's.

Profit margin for that period was somewhat reduced as heavy investments were going into electronics. After tax margin declined to 10.9 percent, still healthy by most measures.

Then we entered the transistor age. IBM announced its first semiconductor system in 1957 and delivered it in 1958. Through the period of 1959 and 1960, IBM brought out a number of systems, some with new architectures and some with evolved architectures based on their vacuum tube predecessors.


For example, the 7080 was a semiconductor version of the 702 and 705 business systems. It was brought out because the new architecture 7070 business system had not done as well as had been expected. Customers that had 702's and 705's were not converting their programs to the radically different architecture of the 7070, thus the compatible 7080 was produced. Less than 100 of the 7080's were produced, yet the system was a business success.

A special story, however, was the 7070 which was IBM's new business architecture entry for the semiconductor age. RCA had produced their vacuum tube BISMAC series and then moved to their transistorized 501 series. The 501 had good performance and price, and IBM was racing to com-

pete before we lost initiative in the business systems area as business applications were viewed as being 80 or 85 percent of the demand in those days while scientific applications provided the rest. Thus the 7070 was a new era system that we hoped would retain IBM's position and allow us to grow from that base.

Ralph Palmer had done something that was typical of him: he held a competition to determine which laboratory was going to design the 7070. Poughkeepsie, IBM's large systems laboratory, had a design that was attractive and they vied for the responsibility of building IBM's new transistorized business entry, essentially the plum of the development community.

The Endicott laboratory, which had earlier produced the 650 system, had its own version of what to do: they proposed to build upon the 650 architecture and Endicott worked hard to win the prize. When the dust settled, Endicott had won the mission with a lot of aggressiveness in proposing features and function in what was to become the 7070. It turned out, however, the 7070 was such a complex system that it did not sell as well as had been expected.



Dawning of Transistor Age for IBM Computers: 1957-1960.

In the meantime, in Endicott there was work on replacing electric accounting machines with stored program computers. IBM had been struggling for seven years to find a way to consolidate in an electronic system the capabilities that were found in the assorted unit record machines such as gang punches, collators, sorters and calculators. Several approaches had failed because the people working on the designs had tried to build systems with plug boards which were the control unit in the electric accounting machines.

A bright engineer in the Endicott laboratory, Fran Underwood, con-

ceived a von Neumann stored program system that became the IBM 1401. That system, announced in 1957 and shipped first in 1958, went on to become, from IBM's standpoint, the Model T of the computer industry. It rented at \$2495, an unprecedented bargain in contrast to the \$20, \$30, \$40 and \$50 thousand per month that customers were paying to rent the bigger systems of the times. We expected to sell 5000 1401's but eventually installed more than 12,000. The 1401 led IBM into the computing big time, bringing to the company a much broader set of systems customers.


On the scientific side, the 7090 was a transistorized version of the 704 and 709, just like the 7080 was a transistorized version of its vacuum tube predecessor. Something like 300 7090's machines were installed. They were profitable and were very popular in the scientific and aerospace communities and that had something to do with some of the arguments that arose during System/360's design.

There had been a gap in the middle of IBM's scientific product line and a lot of clamor came from the demanding sales force for small scientific computers. A group at Poughkeepsie developed a machine called the 1620. However, instead of a small binary design they produced a decimal design. Its rental was \$1600 making it the first IBM system with a rental price smaller than its serial number. We sold more than a thousand of those systems to the fledgling minicomputer area.


The 7030 was a special machine. Years earlier, Dr. Edward Teller had wanted a new scientific system for three-dimensional hydrodynamic calculations, and Dr. Teller talked about his need to IBM super salesman Cuthbert Hurd. Dr. Hurd had guessed that such a system might take a couple of years to build, might cost \$2.5 million and might run at one or two million instructions per second. Dr. Teller went to Congress and got the funds. And so a small group that included John Griffith and Gene Amdahl, worked on a design that we called LARC for Livermore Automatic Reaction Calculator. A Univac team also worked on their version of LARC. We thought we had a great design and were on the way out the door of the Poughkeepsie laboratory to present our design to Dr. Teller when Ralph Palmer stopped us and said, "It's a mistake." Transistor technology was changing rapidly, and we were going to build this system with point contact transistors or surface



IBM 305 RAMAC, 1950's.



IBM 7090, 1958.



IBM 1401, 1958.

barrier transistors, the semiconductors that produced the best speed in the early days. Palmer had noticed the newly invented diffusion process promised better control of the speed of semiconductors and thought it would be a mistake to build the LARC system with obsolete semiconductors and occupy the estimated 350 people required to build it. So Palmer, to our dismay, forced us to tell Livermore's Lou Nofrey and Dr. Edward Teller that we had decided not to build the design we had worked on. We showed Livermore our design approach to illustrate the kind of system we were capable of building but said, "We are not going to build that machine for you; we want to build something better! We do not know precisely what it will take but we think it will be another million dollars and another year, and we do not know how fast it will run but we would like to shoot for ten million instructions per second." So Dr. Teller bought the Univac machine, and we went back to lick our wounds.

Later, with the Univac LARC system commencing development for the AEC and the able Sperry salesman selling it, IBM concluded that we had better fund a new system ourselves. The thesis was to build the fastest system. It was internally called Project Stretch, for stretching the technology. We did design the Stretch system ultimately producing a total of seven. Its IBM type number was the 7030 and it was the fastest system in the world for a period. The 7030 was quite expensive to build, costing IBM tens of millions of dollars. However the technology and the architecture that flowed from Stretch later had important influences. All of us in the IBM development community have a soft spot in our hearts for taking on such "one-of-a-kind, break-the-sound-barrier" projects.

It would be relevant to describe the company organization in the 1950's when IBM was still very small. Although it was a \$200 million firm, there was one vice president for engineering and he handled all engineering business such as the national interest business, supplies, typewriters and electric accounting machines, the largest engineering activity and, the few engineering tasks in electronic computers. And so it was with manufacturing with one VP overseeing all aspects and so it was with marketing in that a VP oversaw both sales and service. That was an inappropriate structure for the growing IBM which crossed a billion dollars of sales in 1957, thus the organization was changed. A

major reorganization started in 1955 and in four years the change was completed. In essence, the company decentralized and formed new divisions.

The World Trade Corporation, that had started years earlier was beginning to grow. It had its own marketing for the countries in which IBM was present, its own service, its own manufacturing and its own development with its own laboratories and engineers. World Trade had rationalized their countries needed products that were different from what the Americans were producing, so it set out to build its own products for its customers.


In the mainstream was a senior vice president for data processing, T. Vincent Learson. His organization was set up in a new structure consisting of three divisions. The General Products Division in Endicott, New York and San Jose, California had the mission of developing and manufacturing products with rentals up to \$10,000 per month. In Poughkeepsie, New York the mission of the Data Systems Division was the development and manufacture of systems renting above \$10,000 per month. The Data Processing Division handled sales and service and was headed by a super professional, Gilbert E. Jones. In its heyday it was as fine a marketing force as ever existed.

One important point: In this structure the financial books were controlled by the product divisions; marketing and service were run on apportionments that were doled out by the product divisions. Thus the product divisions did the market forecast; set prices and had general responsibility for the financial health of the products they produced.


Now let us consider the IBM product offerings at the time System/360 development was commencing. There we were in 1960 with six families of new systems, most of them doing well. The 7070 was not selling as well as we had hoped but the rest were selling well and some, such as the 1401, far exceeding our forecasts.

The major reorganization had just been completed in 1959 when Tom Watson, Jr. called the new senior management together and, in what I thought was real vision, said that our new products should serve IBM well but we should start thinking about where we are going in the future and should have someone start working on that future. His conclusion was the Data Systems Division would be given that mission.

Now some irony: the General Products Division, which had won the inter-



IBM's Overlapping Product Line: 1961.



IBM's Immediate Products to Strengthen the Product Line: 1961-1963.

nal competition to build the 7070 was struggling with that system's design and release to manufacturing. It was late in schedule and its architectural complexity was affecting programming.

However in the major reorganization of the 1950's as luck would have it, the Data Systems Division took over responsibility for the 7070 and its problems. Some of DSD's leaders thought the best thing to do was to get rid of 7070 so they started a project in Poughkeepsie to build a better system. The development leader in Poughkeepsie, Steven Dunwell, gave a simple charge to the engineers under him: "I want a machine that is twice as fast as the 7070, at half the cost." He had another little codicil on his charge: he wanted it packaged in one rollagon, which was one of the packages we used then in larger systems.

So the people in Poughkeepsie began the new design. Bolstered by Tom Watson's assignment of a corporate mission to plan the next series, they expanded their 7070 replacement into a family called the 8000 series.

The proposed 8106 was the specific product Poughkeepsie conceived to replace the 7070, and it was furthest

along. As a matter of fact, it was being prepared for announcement in March 1961. To fulfill their worldwide mission, Poughkeepsie quickly planned other systems around the 8106. They added a scientific attachment called the 8108; it was not a standalone machine—you had to buy an 8106.

Burroughs was working on a technique called push down stacks and Polish notation and that concept enamored some of our people. Thus Poughkeepsie decided to build an analogous high performance system called the 8112. The General Products group was so successful with the 1401 that they did not want anything to do with the 8000 series but Poughkeepsie required small systems to handle peripheral management and to provide growth for their bigger systems. Therefore, they wanted a small commercial machine and started a design called the 8103, a small business computer. To fill the gap in the scientific area, Poughkeepsie proposed a machine called the 8104. These systems had some architectural similarities but, by and large, were quite dissimilar and that was perhaps the fatal flaw.

Other groups in IBM were working away too. The General Products Division, with their 1401 success, had planned to take that machine in all directions, down and up. They proposed a 1401G, 1440, 1410 and 7010. They had a 1620 model II, and because of the success of the 1401 and 1620's it appeared that General Products was headed for success with a line of systems competing with Data Systems' proposed 8000 series.

The World Trade Corporation did not like the 1620, it was a decimal machine and World Trade wanted a small binary machine. Thus the Hursley, England Laboratory started a design of a 48-bit, small binary machine called SCAMP-I, a credible machine that might have succeeded had it proceeded. Unhappily, the computer demand in Europe in those days could not generate enough volume to pay SCAMP's way, so the machine was in financial trouble. The aggressive Hursley Laboratory then said, "We can build a faster version called SCAMP-II on the SCAMP-I base, get more volume and fix the business case." They tried just that but it was not enough to fix the business case. So, undaunted, they hypothesized a business version of SCAMP called SCAMP-III, and were evaluating that approach.

In net then, World Trade had its evolutionary plan, Data Systems had the corporate mission and its 8000 se-

ries plans and General Products had its plans based on the success of the 1620 and the 1401. All the camps were in competition. It appeared as if a time would come when a customer would call up and say, "I would like to hear about an IBM machine," and three salesmen would get stuck in the door waving their catalogues saying, "Don't listen to him, listen to me."


My role in this came in January 1961 when Vin Learson asked me to leave Endicott, where I was working on the 1401, 1620 and the 1410, and to go over to Endicott's rivals in Poughkeepsie. His instructions were simple: "Look at that 8000 series—if it is right, build it; if it is not right, build what is right." That is about the length of the discussion I had with Learson.

One of the problems we had with all those architecturally dissimilar systems, was that peripherals had to be customized by family. If you wanted to build a peripheral that was optimized for parallel binary machines, that was tough to justify business-wise. If you were going to build something that was serial by character for commercial machine, that was another design. None of these systems had enough volume to sustain new investment in a variety of types of peripherals, so the peripheral groups in San Jose, Endicott and Poughkeepsie worked at what they believed best to build, and the system adapted those devices to the processors.

Since most of the technology work was going into the processors the peripherals were not keeping pace with the processors. It was possible to go from one processor to another and get 100 percent gain in internal performance, but because of slow peripherals a user might realize only a 10 or 15 percent gain in thruput performance and that is before you take programming into account.

Circuit technology was also different by type of machine. Here I must say that Ralph Palmer and senior development management had strived to standardize our semiconductors from the beginning. Previously in IBM every project had its own designers who would design the circuits for their projects, optimizing their products for their intended applications. In 1955 Ralph Palmer established central circuit-design laboratories, with the centralized group providing circuits to the systems groups. It caused much disagreement in the laboratories but, in hindsight, it was the right decision.

To aid standardization we designed a printed circuit card called

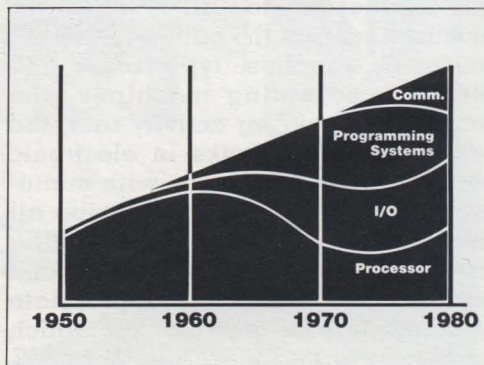


Two COMPAC Modules (bottom), Equivalent to the Standard Modular System Card (top): 1961.

SMS—the Standard Modular System. One card was approximately the size of your hand and held one circuit of discrete transistors, resistors, diodes and capacitors. We developed a lot of automated equipment to insert components, to solder them in place and to test the cards. In the early days of transistors and the Standard Modular System, the management theory was that if we did it right, about a hundred of these SMS card types would serve all the IBM systems which would be just fine for service, service training, engineering refinement and further evolution.

However by 1960, the requirements had exploded out of control and had grown to more than 2500 card types. The Standard Modular System plan had missed its target significantly. There were so many card types the circuit engineering force spent its time designing new circuits. And, of course, field inventory, field engineer training, and such things were expensive and complex.

Perhaps the worst problem that plagued our many types of systems was programming. In 1960, during the heyday of the 1600 and the 7000 series, our programming budget was \$5 million, less than five percent of the development budget. With so many types



IBM Development Resources: 1970-1980.

of architectures, not only did we have to produce FORTRAN for each type of architecture but there had to be a FORTRAN for the disk version of the 7090, and one for the tape version of the 7090, as well as special assemblers and utilities. We were in trouble with respect to programming in 1960 and we knew it.

Moreover, we had split our customers' computing with scientific and business machines. Boeing is a typical example. It had two very able yet separate computing shops—one had 7080's, one had 7090's—vying for funds, vying for applications and vying for people. What really was happening, we perceived, was that business systems needed more of the logical and computing abilities of the scientific systems, and the scientific operations needed more of the variable field length and alphanumeric capabilities of the business systems. We had unwittingly put our customers into two camps and the camps were competing.

The user programming investment was high and growing rapidly, and our customers had sent us a signal with the 7070: no matter how powerful the architecture, no matter how much better the price-performance ratio was in contrast to older systems, they were not going to make the move. Most users could not afford to convert and did not.

In 1960 most IBM development resources went into the evolution and propagation of processors. Only a small amount went into peripheral research and enhancements. Most peripheral R&D went into tapes, a bit into disks and printers, and a tiny amount—\$5 million in 1960—for programming.

Thus, with all these problems, in considering the 8000 series in 1960, we concluded it had frailties such as the incompatibilities between the architectures themselves, had other missing elements in the program and were planning implementation in existing technology. In May 1961, a decision was made to build a new family of systems in new technology. Each system in the family would be equally adaptable to business and scientific use. And while it was easy to produce machines that were upward compatible, we were going to try and design the new systems to be both upward and downward compatible. Thus if any systems had the required peripherals and the amount of memory specified by the programming, it could run the same programs, whether it was a big machine or a small one. More importantly, the approach unfettered

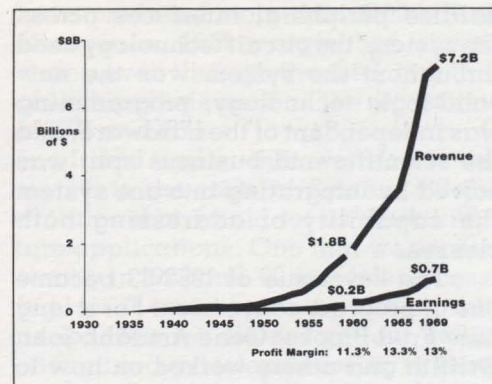
programming from the specific systems themselves. The entry-level programming could run on the whole family, and large systems programming—more complex programming with higher function—could also run on the whole family.

And to fix the I/O problem, the new systems' thesis was standard interfaces for peripherals. We decided to have the peripheral devices adapt to the standard interfaces so that control programming would not have to be changed extraordinarily by new peripherals, and we hoped the new peripherals could achieve high volumes.


Lastly, the plan was to build the new systems in a new technology that was under development in IBM. Internally it was called the Compact technology, later named SLT—Solid Logic Technology. Basically, it was a hybrid, micro-miniaturized technology which, instead of using the palm-sized SMS card to package the circuits, Compact used fingernail-sized chips, each containing a single circuit. Erich Bloch, John Gibson and I agonized a lot in 1961 about whether we should go to large-scale integration instead of pursuing the hybrid micro-miniaturized technology. Fortunately, we elected to build what we had in hand. Heavy investment went into automating the production of SLT technology and production was very sophisticated. Significant volumes were turned out at high quality and low cost.

The architecture of the systems had a decimal and variable field length base structure with optional binary and floating point. Each system could perform scientific as well as business calculations and we also tried to design in the basics needed to allow us to expand to new applications such as real time or event driven applications as they unfolded.

Another problem: IBM has an aggressive sales force and they were paid largely on commissions. Our salespersons did receive a base amount which would buy baloney sandwiches, but if they wanted to eat steak, they had to sell. Our sales force's long range viewpoint was that "tomorrow is too long." They certainly had a tough time waiting for a few months, let alone a few years. However, anything as significant as shifting gears to a new technology, new architecture and new programming was going to take a lot of time. We estimated that we would announce in 1964. It turns out we did announce in 1964 and shipped in 1965. But in 1961, such a delay seemed like an eternity to the sales force.



IBM Growth: 1960-1969.



IBM Growth: 1970-1982.

In the meantime Seymour Cray at CDC and lots of able companies were beginning to succeed, bringing out competitors for the IBM product lines. Our sales force felt their homes were burning down and they wanted some solutions quickly. So we put in place some programs I called "temporizers"; I hate the word, but that is what we called them then. The project consisted of extensions to the current product lines. There was to be a higher speed version of the 7090, called the 7094, which turned out to be so successful that we built a 7094-2 and we actually worked on a 7094-3. Two new extensions of the 7070 were built—the 7072 and the 7074—intended to aid the lagging 7070 sales.

A bigger version of the 1410 was built for 1401 growth, the 7010. A 1620 Model 2 was built, and for that gap in the small scientific area, two systems were built that were related to the 7090 architecture—the 7040 and 7044.

All these systems were undertaken starting in mid-1961. Some were announced in 1962 and the rest by May of 1963. IBM suffered competitive losses but we were able to keep the sales force alive during the time the gears were being shifted to System/360.

In net: System/360 solutions in terms of the problem was to stan-

standardize peripheral interfaces across the system; the circuit technology used throughout the system was the new solid-logic technology; programming was independent of the hardware, and the scientific and business split was solved by integrating into one system the capability of addressing both classes.

The key issue of 1962-63 became one of program conversion. For a long time Fred Brooks, Gene Amdahl, John Griffith and others worked on how to do this. The first thought was to have machine translation. Bright people worked on a conversion program that would allow one to dump a program in a hopper and have the conversion produced run effectively on the new architecture System/360. After a couple of years of hard work and several million dollars of investment, we concluded automatic conversion was not going to make it. The theory then was that we had better back off to machine-assisted translation where we would translate as much as we could and signal the items that had to be handled manually.

We knew our customers were not going to convert manually; we had to have a tool. Necessity breeds invention, and a couple of professionals found the solution. We found that if we examined the 1401's registers and data flow in the light of the 360, the 360 had all the registers and more, and all the data paths and more. Since we had decided to use some of Professor Wilkes' work in the controls of these machines, namely read-only memory instead of hard-wired logic, the controls were vastly simplified. Thus it was relatively easy to add to a 360 machine the instruction set for a 1401 and literally throw a switch so the System/360 would run credibly as a 1401. Emulation proved out for the 1401, 7070, 7090, and 7080—fortunately for IBM.

Systems running in emulation mode did not run at full 360 performance, of course. But, by and large, through the combination of read-only memories for control that let us add the instruction repertoires of the older machines, the 360 machines did take on the form of the older systems and customers could run the old machines' programs with reasonable price-performance and then convert at their leisure to the newer architecture when they wanted. But believe it or not, some users are still running in emulation mode after all these years.

1962 was a period in which we found ourselves asking can we make

it? Can we design the family? For a while it appeared that we could not design a processor that was inexpensive enough at the low end while containing the instructions of the big machines; similarly for a processor at the high end, would their performance be limited by staying compatible?

But senior management realized that if we produced a five-times 1401 and Honeywell produced a four-times 1401, the whole question would be quickly reduced to plant capacity. Honeywell might sell 5000, we might sell 5000, and conceivably there would be a price war. Worse, the 1401, invented years earlier, was inadequate for future applications.

In contrast to the existing product lines, there were so many attributes in the new product family that in February 1964 IBM decided to go ahead. We announced System/360 on April 7, 1964. We announced five machines; the Model 30 was developed in Endicott and the Model 40 was developed in Hursley, England. As a side point, World Trade wanted to play a role in the 360 development. Its labs were full of bright people, but young and inexperienced, thus I wanted to give them supporting roles. However, Vin Learson said "Absolutely not. They have to have a head-held-high role; we want to give them a whole system." So we did. We exported a number of U. S. people to help Hursley, and after that Hursley became one of the senior labs in IBM's development community.

Poughkeepsie developed the Models 50, 60 and 62, and Model 70.

Later, through the last part of the 1960's, there were successors and additions announced: the entry System 20 and the 22, the 25, and a scientific optimization Model 44. Some new memory came into the 65, which replaced the 60 and the 62, and the 75 with the new memories replaced the 70.

The 360 model 67 grew out of MIT's criticism of System/360. MIT scientists were important in computer research, and we wanted to be certain we stayed close to MIT's thinking. And during this period, as busy as we were, 360 design people would go occasionally to MIT. However, in retrospect, MIT did not hear us, we did not hear them, and I presume we did not speak clearly enough to them.

When the 360 system was announced on April 7, we all settled down to the happy task of making it happen. But on June 6, 1964, I traveled to MIT to see what they thought of 360, which by then had been announced for a couple of months. To my dismay, Professors

Corbato and Fano told me that they did not like System/360.

Three of MIT's four criticisms were trivial and could have been fixed quickly but, criticism one was deep in the concrete and that was MIT's view that time sharing was just around the corner, thus dynamic address translation would be a fundamental part of any system's architecture in the future. Without it, management of the storage by the programmers would be an impossibility.

There was some debate in IBM, but I decided that MIT was right, and we had missed it. It took us several years, but we did fix it and finally got dynamic address translation across the family. However, back in 1965-66, we produced a special version of the Model 65 called the Model 67 which was built for leading-edge customers like Bell Labs that wanted time sharing and demanded dynamic address translation.

Unhappily for us, MIT decided to buy a General Electric machine and not the 67 that we were designing to supplement the 360 family and answer their requirements. Through the 1960's, the only 360 machine that had dynamic address translation was the Model 67. A special version of that design, called the 9020, became the system used in the FAA's enroute traffic control system.

We thought in those days we would be lucky if the series would last one generation—3, 4 or 5 years—and if we were really lucky it would last 8 or 10 years. However System/360 has lasted 20 years, and we are working now to extend its life into the 90's. Possibly it will not make it, but the durability of the 360 architecture has far surpassed our expectations.

By the late 1960's, technology had marched on to the point that instead of one circuit per solid logic chip, we could do three or four circuits per chip: the early days of large-scale integration. So we produced a family of follow on 360 systems: the 115, 125, 138, 148, 158 and 168 and, in between, there is some detail of what were called "vanilla" machines that I will skip. The bottom line is that all these machines had dynamic address translation and our control programming was substantially evolved to accommodate virtual systems capability.

The mid-range and high-performance systems of the 1970's were all direct members of the 360 architecture family. And since 1979 the 43XX machines and the 308X's were added and they are all members of the 360 architecture. These systems, over the years,



IBM System/360.

have produced more than \$100 billion of revenue. IBM margin has stayed strong even through the thick and thin of such periods such as the recessions of 1971 and 1975.

I said earlier that we tried to design into the roots of System/360 the abilities that would let us work in future applications. One that we sensed clearly in 1962 and '63 was teleprocessing, for it was beginning during that period. But we did not get our hands enough around teleprocessing to know just what to do, so we put hooks into System/360 to add teleprocessing capabilities later.

Our estimate was that in the United States we would sell 2500 of the 40, 50, 60 and 70 systems, and by 1970 a third of those would have remote terminals and thus require communications, hardware and programming. What actually happened, fortunately for IBM, was that we sold twice as many of those systems as we had expected by 1970, and by 1968 we had already passed in teleprocessing what we had expected to reach by the end of 1970. And by 1970, we had sold two and a half times what we had expected to sell in terms of teleprocessing.

In hindsight, just building those 360 machines and the complexities of the technology, new peripherals and control programming so consumed our resources that we really did not tend swiftly enough to communications. And that explains the alphabet soup that existed in 1970 teleprocessing for one laboratory or another would develop a piece and a customer would produce something else and the assemblage was inadequate and inconsistent for teleprocessing in 1970.

Thus, starting in the early 1970's, we set out to do the same thing to the communications subsystem that we had done to the central processing subsystem. It was called Systems Network Architecture (SNA), and some of you may be familiar with SNA. We shipped SNA first in 1974 and, it has been generally accepted by the International Standards Organization as an architecture that straightens out the protocols, disciplines and structures of the communications subsystem.

In 1973, when we were finishing work on SNA, our hope was that we might install 3500 SNA systems worldwide early in the 80's. Last year an IBM team gathered to celebrate our 10,000th SNA customer.

At present in the hey-day of PC's and the exploding world of work stations, we are talking in terms of

hundreds of thousands of SNA installations.

SNA has had a succession of sophisticated additions to the structure, the features you would expect once a base is in place; alternate routing in the case of line outages, dynamic re-configuration non-IBM terminal attachment, and those types of abilities.

There is an explosion taking place in computers and communications. Today we find computers connected to computers by communication lines and control units connected by communication lines to hundreds and thousands of terminals. Then, of course, there are minicomputers pioneered by Digital. Whether it is real-time applications, batch applications or interactive applications, minicomputers also require communications from distant terminals, and more and more, these terminals need access to central data bases and vice versa. Thus there is great need for computer communications.

If I would characterize where we are today in allocating our resources, we spend a good deal more on communications and still spend a hand-

some amount on programming and peripherals.

The computer-communications explosion caused us to decide to do something more significant in communications.

We worried about AT&T for, if they controlled all communications and also provided computers, IBM might be at a disadvantage. Thus we invested in a communications satellite company that is providing new communication services. It will be a good business in its own right and is bringing new communication capabilities to teleprocessing users, keeping pressure on the telephone companies. We think that is good for the industry.

How did general management operate? First there was a strategically minded management in the 1960's. Tom Watson assigned a team to work on the next generation.

A broad direction was set but the senior management delegated detail; they did not strive to manage the architecture. They heard the debates and worked to resolve problems but never stepped in to dictate designs such as 36-bit words.

The 360 undertaking stressed IBM to the limits and senior management organized and reorganized IBM to meet the needs of the times.

Lastly, I must say that through a lot of countering viewpoints, senior management such as Tom Watson, Vin Learson and Al Williams, had a lot of tenacity and did risk a lot.

As to whether it was worth it, I will just say that from the period 1964 to 1980, the profit after tax on 360 systems was far greater than the total sales we had anticipated back in 1964. This System/360 was an outstanding business success. More importantly, it gave us the foundation to move resources into new peripherals, to do the things like SNA and all that went with SNA in terminals and teleprocessing, to specialize in certain industry areas and to diversify into businesses such as satellites.


It has also given us a new complexity for in the 1960's came the compatible peripheral competitors. A small company in Oklahoma, Telex, started making copies of IBM's magnetic tape. A number of customers bought the copies. And soon, manufacturers produced copies of our disks, multiplexers and main memory and by the 1970's we saw copies of our terminals and finally, the pièce de résistance, compatible central processors from Fujitsu, Hitachi, Amdahl, Magnusson and others.

Those copies were expected. When we started to work on System/360 our rationalization was that, in the face of copies we had to insure that IBM was constantly the best, that we had the best technology and the best programming and the best price performance. Those ideas sold in IBM and we still believe it.

One negative consequence was the anti-trust litigation that was very costly and stressful.

In the last days of Lyndon Johnson's administration a law suit was filed by the Department of Justice. Also, Telex had filed suit saying that we had damaged them with our "predatory" practices. We filed against Telex for stealing and in a curious decision in 1972, the District Court in Tulsa found for both companies. It found Telex guilty of stealing and fined them \$20 million and found us guilty of damaging Telex and fined us \$120 million. After trebling under U.S. antitrust law that fine went to \$360 million.

At the end of 1972 IBM stock went from \$365 to \$140.



Evans viewing the Whirlwind. Museum board member Gordon Bell (left) points out the Whirlwind computer exhibit to IBM vice president Bob O. Evans at the Museum's former site in Marlboro, Massachusetts. Designed in 1951, the Whirlwind was the first real-time, parallel-processing computer with core memory.



Office Preparations. Before the staff arrives from Marlboro, workers ready the newly renovated office space at Museum Wharf.



Moving in. Almac movers slide a word processing system into the new offices at Museum Wharf on moving day, March 7th.



The last load. An Almac moving truck arrives at Museum Wharf in downtown Boston March 7th to deliver the second and final truckload of office furniture and equipment.

\$10 Million Capital Campaign Kicked Off with Pre-Preview Party

A gala pre-preview party at The Computer Museum's new home in downtown Boston kicked off its \$10 million capital campaign May 11th. The evening opened with a lecture by Intel founder Robert Noyce followed by dinner, dancing and films.

The capital campaign is designed to raise the \$10 million needed to purchase half interest in Museum Wharf, create new exhibits, and start an en-

dowment. It will establish the Museum in its new quarters and provide a foundation for its continued growth.

A new brochure outlining the campaign asks potential donors to "invest in the future of The Computer Museum." Suggested levels of giving start at 4K (\$4,096) and continue in the standard industry memory units of 8K, 16K, 32K, 64K . . . 256K.

Individuals and corporations from

throughout the industry and Museum members will be asked to participate in the campaign.

For more information on the campaign, contact: Gwen Bell, Director, The Computer Museum, 300 Congress Street, Boston, MA 02210.



Under wraps. This is not a trendy new exhibit, but Museum artifacts upon arrival at Museum Wharf.

Museum to Host Engelman Lectures on Artificial Intelligence

Judith A. Clapp

A provocative new lecture series on artificial intelligence (AI) in honor of AI pioneer Carl Engelman begins at The Computer Museum this fall.

For more than 20 years, Carl Engelman and other researchers in industry and academia have worked on creating intelligent computer systems.

Engelman's contributions to the field are among the major successes that have occurred since the inception of artificial intelligence. He developed MATHLAB, one of the first systems to perform symbolic mathematics in the mid-1960's. He headed a group of researchers at the MITRE Corporation in development and application of expert systems, techniques for computers to understand English, and research in the automatic generation of computer software.

Only a few "intelligent" computer systems are commercially available to

date, but the interest in AI is greater than ever. In the United States, the Defense Advanced Research Projects Agency spends some \$50 million a year on AI research. In Japan, the Ministry of International Trade and Industry is spending \$450 million in a 10-year project to develop a Fifth Generation computer particularly suited to AI applications.

The Computer Museum's lecture series will trace how the fundamental approaches to artificial intelligence have been developed, show how they have been successfully applied, and explain why there are so many challenges remaining.

If you share an interest in artificial intelligence or would like to make a contribution in Engelman's memory, please send your donation to Engelman Lectures, The Computer Museum, 300 Congress Street, Boston, MA 02210.

July

The Computer Museum and Graphics Pioneers will co-sponsor a panel session entitled "A Retrospective: Six Perennial Issues in Computer Graphics" at SIGGRAPH '84 on July 25th in Minneapolis, Minnesota. Museum director Gwen Bell, Atari chief scientist Alan Kay and others will trace the evolution of the technology and the major issues in the field.

September

Members will receive the fall issue of The Computer Museum Report featuring a complete fall calendar of events and a transcript of Robert Noyce's Museum lecture on his invention of the integrated circuit. The new 1984-85 Museum Store catalog will also be off press and on its way to members.

November

An opening party for Museum members is planned for Tuesday evening November 13th. The Museum will open to the public on Wednesday November 14th.

Join the Museum

Name	_____
Address	_____
City/Town	_____
State	Zip

Please enter a membership in the following category:

Member	\$25
Corporate Member	\$125
Founder	\$250
Corporate Founder	\$2500

All members receive the quarterly Museum Report, a 10% discount on Museum Store purchases, and announcements of Museum programs and events. All membership contributions are tax-deductible within the limits provided by law.

To become a Member or Founder fill out this coupon and return it with your check or money order to:

Membership Coordinator
The Computer Museum
300 Congress Street
Boston, MA 02210

New Founders


(January 24 through April 24, 1984)

Applied Magnetics Corp.
AVX Corporation
John C. Barstow
Alan G. Bell
Gregory C.F. Bettice
Marshall D. Butler
Thomas W. Crosby, Jr.
Joe Cychosz
Bob O. Evans
Tse-yun Feng
J. Thomas Franklin, Esq.
Regina H. Garb
Charles W. Johnson
M. Kapor/E. Poss
Bryan S. Kocher
Jim McIntosh
James N. O'Boyle, Jr.
James F. Shaughnessy, Jr.
The Small Business Foundation
Rex Rice

Corporate Founder List

(January 24 through April 24, 1984)

Almac Moving and Storage, Inc.
American Federation of Information
Processing Societies, Inc.
Apollo Computer, Inc.
Association for Computing Machinery
Automatic Data Processing, Inc.
Benton and Bowles
Bolt, Beranek and Newman
Boris Color Labs
British Computer Society
Robert Cipriani Associates
Clint Clemens
Codenoll Technology Corporation
Computer Science Press
ComputerWorld
Control Data Corporation
Convergent Technologies
Coopers and Lybrand, Boston
Data General
Dataproducts Corporation
Digital Equipment Corporation
Digital Equipment Computer Users Society
Expoconsul International, Inc.
Ford Motor Company
Fujitsu Limited
General Systems Group, Inc.
GTE Data Services, Inc.
Hewlett Packard Company Foundation
IEEE Computer Society
Intel Corporation
International Telephone & Telegraph Corporation
Los Alamos National Laboratory
MDB Systems, Inc.
Microsystems Engineering Corporation
MITRE Corporation
MOCO, Inc.
Motorola, Inc.
National Semiconductor Corporation
OMNI Publications International, Ltd.
PCK Technology Division
Richard Reno
Schlumberger Foundation
Seldin Publishing
Software Results Corporation
System Development Foundation
Tobin Vending Service
The Travelers
Wang Laboratories



What is a Founder?

Any individual or corporation who gives more than \$250 or \$2500 respectively during the Museum's first two years.

On March 14, 1982, The Computer Museum was given provisional status until June 24, 1984 by the IRS as a public, non-profit charitable foundation. During this time, the Museum has to prove that it is indeed a widely supported public institution.


With almost 500 Founders at present, the Museum's goal is to close the list on June 24th with 1,000.

This is a unique opportunity to help establish the only international museum devoted exclusively to computers and the history of information processing. If you are a Founder, we thank you, and if not, urge you to become one.

A newsbrief of the collection

The MITS Altair 8800 was introduced in 1975 as a low-cost mini-computer available either assembled (\$621) or in kit form (\$439). Designed around the Intel 8080 microprocessor, the Altair was marketed as the processor for an extensive system of products, ranging from line printers to TV-camera inputs. Because few of these peripherals were ever developed, the Altair became a hobbyist's computer.

It was the first computer to use a cassette tape as auxiliary memory; the information was recorded acoustically as modem tones. The Altair was also the first computer to offer BASIC on the Intel 8080 processor. "The Altair's success is really based upon Intel's quality," former Altair salesman Rick Jevon related in a gallery talk at the Museum in September 1982. Nonetheless, according to Geoff Feldman, another early Altair salesman, "It required a tremendous amount of very creative twiddling to make it work."



**The
Computer
Museum**

300 Congress Street
Boston, MA 02210

Non-Profit
Organization
U.S. Postage
PAID
Permit No. 110
Fitchburg, MA
01420