The Genesis of the Mainframe

Bob O. Evans

Introduction
by Wilhelm G. Spruth

In 1927 the German author Stefan Zweig published a short book labeled “Sternstunden der Menschheit” (world hours of mankind). In it he descripes 14 historical events that changed history forever.

I have always believed the appearance of System /360 in 1964 has been such an event, changing the direction of the computer industry in more ways than anything else in the last 50 years.

The design of the S/360 architecture is rightfully credited to Gene Amdahl, Gerry Blaauw and Fred Brooks. Bob O. Evans has been the IBM executive who made it happen.

It is impossible to have met Bob O. Evans and not been impressed. He was forceful, decisive, intelligent, an excellent engineer with a broad vision, an uncanny capability to understand the most complex issues, and an outstanding leader for the people working for him. He also had an overpowering personality. I will be forever grateful for his guidance.

In 2002 I learned Bob O. Evans had written his memoirs, had considered to publish them as a book, and had been adviced by his friends not to do so, because the text contained too much sensitive material. I contacted Evans and asked him if I could read the manuscript. He gracefully consented and e-mailed me a copy. Bob O. Evans died in 2004.

I believe Bob O. Evans personal account on such a monumental event as the S/360 announcement should be made available. Thus in 2007 I asked his son Douglas B. Evans for permission to publish an extract of the memoirs on my website, to which he agreed.

I want to accept Bob O. Evans decision not to publish sensitive information. The original document is 259 pages long. The following extract has only 35 pages, but contains all the material which I believe is relevant looking at Bob O. Evans contribution to the S/360 development. The text has been extracted with no modifications except for reformatting. I have added a small number of comments in cursive typeface.

Chapter 1
Starting a Career

How did I come to be at the center of so much of IBM's greatest moments? It was sheer luck—being in the right place at the right time. I was born in Grand Island, Nebraska on August 19, 1927, and raised in Shelton, a mid-Nebraska farming community, population 900. I worked on the farms in the summers and was a rough-and-tumble lad from an area known for hard-working, basically honest and direct people. This place was not the epitome of culture and sophistication but had wonderful attributes: honest, hard working, caring people.

In January 1947 I enrolled in Iowa State's electrical engineering curriculum. I lived in a men's dormitory, Friley Hall, where I first saw Maria Bowman. She attended a dance in my dormitory, where I—a wallflower—watched the couples have fun. Maria stood out, dancing, beautiful, laughing, and popular with the entire group. I wished I could meet a girl like her.

We were engaged in the Spring of 1949, graduated from Iowa State University together in June and were married in Kansas City, Missouri in November 1949.

In July 1949 I found a position with Northern Indiana Public Service Company, a public utility serving the industrial north of Indiana . Later I heard from a cousin, Dan Evans, who had just joined a small accounting machine company—International Business Machines. He told me his company was going to build an electronic computer. I was intrigued.

Dan gave me the name of an IBM executive in Chicago, Les Turney. I wrote to Mr. Turney, expressing interest in IBM's new computer project. Wonder of wonders, I was invited to Poughkeepsie, New York, for interviews. Luckily, IBM accepted me as a junior engineer on the Defense Calculator project. On September 10, 1951, I joined the company and began my love affair with IBM, a most interesting career that lasted thirty-three years.

Chaper 2
Turning a Radical Corner

In 1914, Thomas J. Watson Sr. took over a small firm—Computing, Tabulating and Recording Company—that had a hodgepodge of products, including a correcting electric clock system, electric scoreboards, meat slicers, and call systems for nurses. A few years later, Mr. Watson became interested in the record keeping technology invented by Herman Hollerith. Watson discarded most of the product lines and directed the company to develop punched card accounting equipment. He renamed the company International Business Machines and over the years assembled a management team for his IBM. The renaming was typical of Watson Sr. and his brash, “big time” thinking.

IBM was a success by the time I joined the company in 1951. Sales were in the $100 million range. Profits were higher than average. Investors, employees, and customers all loved IBM.

Although relatively small, IBM was growing at a brisk pace, and was the world leader in the expanding business of electric accounting machines. Despite several competitors, including the Remington Rand Company, IBM had innovated, established strong marketing, sales, and service operations, and—by most users’ accounts—was “the best.” The patriarch, Tom Watson, Sr., had done quite a job.

Being in the right place at the right time had been pivotal to IBM’s success. Literally overnight, the advent of the U.S. social security system had created a demand for producing, updating, and maintaining massive alphanumeric records. The Social Security Agency’s technology of choice was essentially the Hollerith system of punched cards—small paper cards with holes punched in varying locations on the cards to connote numbers and alphabet with products to punch and interpret the holes, sort, tabulate and print the computational results, the product line in which Thomas Watson’s International Business Machines excelled.

Thomas Watson Sr.’s vision must receive the basic credit for this technology choice. The canny leader had “turned a radical corner” when he envisioned commercializing Herman Hollerith’s punched card invention and, relatively quickly, abandoned the mainstay products of his company to become the punched card czar. One marvels at how this leader could overcome the resistance of the members of the management team whose knowledge was of clocks, scoreboards, and assorted products, who had expansion plans in those areas, and who would naturally resist a diversion of company resources to strange new fields. Yet, this Svengali did just that—and, in the process, gave birth to IBM, which became the world leader in electric accounting machines.

Svengali is the name of a fictional character in George du Maurier's 1894 novel Trilby. The word "svengali" has entered the language meaning a person who manipulates another into doing what is desired.

In my early days at IBM, Tom Watson, Sr., was the demanding, revered and feared “maximum leader.” His president, John Phillips, and the executive team had cut their teeth and made their fortunes on electric accounting machines. That business had meant market leadership, high profits, endless new products to be designed, steadily increasing demand, and a trained company that excelled at what they did. One of the new products was the defense calculator—which ended up having a lot to do with turning the next “radical corner.”

The Defense Calculator

In 1948, the Korean War began. The patriotic and marketing-wise Thomas J. Watson, Sr., sent a telegram to President Truman offering the capabilities of IBM—a still small company—in support of the war effort. Government officials were perplexed as to what response they should give, but discussions ensued. These resulted in an IBM executive tour of government laboratories, aircraft companies, and others involved with the Department of Defense. Dr. Cuthbert C. Hurd, who had been hired by IBM in 1949 to build an Applied Science department, organized the tour. He took along Ralph L. Palmer, IBM's director of engineering, and James Birkenstock, IBM's "czar" of IBM business and contract relations. Together, they went out to try and gain an understanding of how IBM could best contribute.

Wherever Hurd, Birkenstock, and Palmer visited they found many complex problems waiting to be solved. In particular, they found completely inadequate computational capabilities.

At the time, Dr. John von Neumann—today known as the "father" of stored program computers—was building a new tool to solve complex mathematical problems. Working at the Institute for Advanced Study in Princeton, New Jersey, von Neumann's small team was making progress with their creative new electronic computer and, in the process, opening the new era of electronic computing. Dr. Hurd was a close friend of the eminent Dr. von Neumann, and knew of the ongoing project at the Institute.

Returning from their tour, Hurd and Palmer made a recommendation to Watson Sr.: IBM's contribution to the Korean War effort should be a new electronic computer modeled after the concepts being developed by von Neumann's team. IBM, small as it was, had already built one-of-a-kind computational showcase machines, including the Harvard Mark II and the Selective Sequence Electronic Calculator—all massive units that used magnetic relay component technologies that consumed high power, were slow and very limited in their capabilities.

Watson Sr. was a wise and experienced marketeer. He knew the value of producing notable special products, and he readily agreed with the Hurd-Palmer recommendation. With no government involvement, using the company's own funds IBM proceeded. And while other IBM executives probably cursed Watson's extravagance under their breaths, nobody dared issue a challenge.

So, the Defense Calculator project was organized. Cuthbert Hurd, Jim Birkenstock, and Ralph Palmer turned Watson, Sr.'s offer to President Truman into an action plan that was to remap IBM completely. I joined the project in September 1951 as a junior engineer working on design and testing.* The "25-hour days" were thrilling.

The Defense Calculator itself was massive with thirteen large cabinets of electronics and power supplies, most six feet tall and varying in width from three to twelve feet. The system filled a large room and below a raised floor were tens of complex cables interconnecting the various boxes of electronics.

We worked feverishly and, by the summer of 1952, Dr. Hurd had arranged some applications testing with independent research professionals.

In mid-1952 we decided to gamble that engineering could satisfactorily complete the testing while production began. Production of such a large and complex electronic system was a new adventure for IBM. Engineering's interface with manufacturing was Lars O. Ulfsparre, a long-suffering lad who had to work with incomplete drawings, a continuum of engineering changes, and quite inadequate parts lists and specifications. But he and his manufacturing compatriots survived.

By the fall of 1952 the Defense Calculator engineering model was beginning to operate in brief spurts. The mean free error path was probably three to five minutes as we continued to make engineering changes and debug the system.

In December 1952, the first production unit—now named the 701 Electronic Data Processing Machine, was delivered to the first floor of IBM's World Headquarters at 590 Madison Avenue in New York City—installed in a setting that was truly a showcase. In truth, design was incomplete. A team of engineers, myself included, worked to complete testing through March 1953. After that, the first production unit was used for demonstrations and high-profile problem solving. Dr. Hurd and his Applied Science group made certain the 701 was seen and used by the world's foremost scientists and mathematicians.

In the beginning of the computer industry, all of IBM's computer systems were leased rather than sold. When production commenced the 701 had a base rental of $17,000 per month. For that amount, the user got the computational unit plus 1024 36-bit words of cathode ray tube memory, a drum, four magnetic tape units, a card reader, card punch, printer, and power supplies. The demand forecast suggested that IBM could sell seventeen systems; in fact, nineteen were built and delivered to the crème de la crème of the national laboratories, government agencies, aircraft companies, and one or two leading-edge companies that would use the new computers for business applications.

What started as IBM's noble gesture in the national interest opened the electronic computer era for the company. The Defense Calculator began the second "radical corner" turned by the company, under the leadership of a Watson. Paced by this national interest project, Thomas J. Watson Jr. later moved IBM from electric accounting machines to electronic computers—a move that made IBM the envy of the business world.

Working on the Defense Calculator was probably the most exciting time a young engineer could experience. In addition to the camaraderie, the sense of purpose was high and the technical challenges were boundless. Indeed, the Defense Calculator project provides many fond memories as IBM's electronic computer age began. And with it, the stage had been set.

Enter Thomas J. Watson Jr. After a training stint in sales, he quickly rose to become the power in the company.

Chapter 3
Prelude to the S/360

Thomas J. Watson Jr.

Of all the executives I have met across the world, Tom Watson Jr. was the most electrifying and most powerful leader. He was the pivotal driving force that built IBM into the world's leading electronic computer company. A lot of the "old guard" fell by the wayside as Tom Jr.'s team took form.

Shortly after taking the CEO helm, Tom Jr. led a major restructuring of IBM

In 1959, Tom Watson Jr.'s major reorganization was finally completed when three new divisions were created to forge the company's electronic computer future, all under T. Vincent ("Vin") Learson, senior vice president for data processing. These were the General Products Division (GPD), led by Orland M. Scott. GPD, with sites at Endicott, N.Y. and San Jose, California was responsible for the development and manufacture of small systems; the Data Systems Division (DSD), led by William B. McWhirter. DSD, with the Poughkeepsie, N.Y. site was responsible for development and manufacture of larger systems; and the Data Processing Division (DPD, led by Gilbert E. Jones. DPD with regional, district and branch offices all over the United States, was responsible for U.S. sales and service. Forecasting, cost estimating, and pricing were placed within the two product divisions, GPD and DSD. Revenues accrued to the books of the two product divisions, whereas DPD—the sales and service operation—was financed by an allocation from revenue.

The loose line of demarcation separating the product divisions' missions centered on the rental costs associated with different systems. GPD was responsible for products with very low prices up to systems renting for $10,000 per month. DSD was responsible for products renting for more than $10,000 per month. A few products were shifted from one division to another to conform to the new organization.

This was the organization and management team that was to take IBM into leadership in the electronic computer industry and they did so quite swiftly!

In 1971 Tom Jr. had a heart attack. I was devastated to hear this news for many reasons, not the least of which was that I knew IBM was about to change—significantly. Tom Watson Jr. was highly intelligent, fair, and sensitive. He was a remarkable leader, an exemplary communicator, a man of great vision, and a man revered by IBM people across the world. I wonder whether there will ever be another corporate executive of his electrifying style and penetrating ability?

Ralph L. Palmer

A few people in the core team of executives lasted two entire decades under Tom Jr.'s leadership. Several dropped by the wayside along the way. In both groups were people who made meaningful contributions to IBM's success as they managed their individual responsibilities.

I have long been convinced that the person, next to Tom Jr., most responsible for IBM's brilliant success in the electronic computer age was Ralph L. Palmer, director of engineering at the most crucial of times. Of the hundreds of executives of many nationalities I have come to know since the early 1950s, none match Ralph's vision, incisiveness, and ability to move mountains. I credit this soft-spoken, focused man of few words with numerous concepts which, when he brought them to action, accelerated IBM into its leadership position.

Palmer believed computers would proliferate and, in the process, need massive data instantly available. He was certain magnetic tapes would be inadequate with their 1/2-inch wide, 2500 feet long, serial access requiring tens of seconds—even minutes—to retrieve a single character. Magnetic drums, used on the Defense Calculator and later on the IBM 650 product, had surface space limitations as well as mechanical impedances to utilizing multiple read-write heads. Palmer wanted fast, random access to voluminous data. While he believed magnetics was the likely technology, he wanted something better than tapes or drums. I was there in 1955 when he described his instinctive solution: the concept of stacks of rotating magnetic disks with multiple read-write heads. Just as bold as his idea was what he did next. Without exhaustive studies by ostensible experts, he simply established a laboratory in San Jose, California, just to work on random-access magnetic disk storage. Palmer selected Reynold Johnson as the San Jose lab's first director. Rey Johnson was a proven engineer in IBM's Endicott Laboratory who had shown unusual creative ability, and he was a good choice to reduce to practice the storage products Palmer wanted for IBM systems. Rey Johnson built the new San Jose Laboratory with bright young engineers such as Louis Stevens, fresh from Poughkeepsie's Defense Calculator, John Haanstra, a rising star in development and, later, general management, Roy Haug, Jack Harker, later to become an IBM Fellow, the most distinguished rank achievable by an IBM engineer, Alan Shugart, Bob Lawhead, Ralph Walker and others.

Today, IBM is credited with inventing the magnetic disk. Rey Johnson and his team achieved great success with Rey receiving the prestigious U.S. National Medal of Technology for his work on magnetic disks. From a business standpoint, San Jose's disk operations went on to produce billions of dollars of profit for IBM. It became a gargantuan enterprise, leading the world in large magnetic disk subsystems. It continues as such to this day, even though there are many new and successful competitors (not to mention the significant business opportunities IBM has missed in recent years, such as small disks for personal computers and workstations).

Ralph Palmer steered IBM to new technologies, putting in place IBM's supercomputer program—named "STRETCH" by Palmer to indicate "stretching the technologies." Unhappily, on a stand alone financial basis, STRETCH lost approximately $75 million. However, it was eventually recognized that STRETCH developed the technologies for the 7070 family, 7080 family and 7090 family of systems which, together, far more than repaid the investment in STRETCH. Six or seven of the STRETCH supercomputers were produced.

Tom Jr.'s "Ace", Vin Learson

I called him "the man with big feet!" And he certainly had them. T. Vincent Learson was over 6'3" and had an overpowering personality to go with his size. This "big" man hated staffs, hated bureaucracy, hated garrulous people and was as decisive as they come. Next to Ralph Palmer I believe Vin Learson was the most significant executive under Tom Watson, Jr. who led and drove IBM to the heights of success achieved in the 1955-1975 era. He was gruff, tough, intolerant of bad thinking, instinctive, demanding and wise. He assembled top-notch people for his team, grew them when they produced, and executed them when they had ample opportunity and failed.

In 1958 Vin named as Group Vice President heading all electronic computer operations except for the World Trade Corporation's sales and service functions. Vin was unique in that he used a very small staff to assist him. At that time there were only four people on Vin's Staff including Byron Havens, inventor of some early digital storage and shifting circuits that were important as well as leader of development of the Naval Ordnance Research Calculator, an early one-of-a-kind vacuum tube computer developed for the Naval Ordnance Laboratory for calculating firing tables for naval guns. Also on the staff was Walter Johnson who had a marketing background. The key person on Vin Learson's staff was Don Spaulding. Brought from a successful stint in the General Products Division, Spaulding was a Harvard-educated professional who had become a successful salesman in IBM and rose on a fast path. He was brilliant, and repeatedly played a key role in the maturation of IBM's electronic computer business. Under Learson and his small staff, in relatively short time IBM produced System 360, which was to have profound effect on IBM, indeed, the whole computer industry. There were many people who made that success however, certainly two pivotal executives were Vin Learson and Don Spaulding.

The 7070

The transistor, invented in Bell Telephone Laboratories in the late 1940s, brought tremendous potential to the growing computer industry. It offered substantial improvements in size, power consumption, speed, and reliability. This was not lost on IBM.

Business applications of transistorized computers, just about everyone agreed, had the highest potential, and that was where most of the engineering resources were focused. Competition grew rapidly; Sperry-Rand, Bendix, General Electric, General Precision, National Cash Register, and many other companies all offered products.

Ralph Palmer decided that a new business computer architecture was required and in 1957 established a design competition between IBM's Poughkeepsie and Endicott laboratories to develop IBM's first transistorized business computer. It was during this period that RCA, which had entered the arena with the Bismac system, in early 1958 announced its newest product, the RCA 501,—which was a transistorized technology entrant for the business applications area. This put pressure on IBM's laboratories to expedite development of a competitive product. Each lab had its particular strengths. Poughkeepsie had more extensive experience with large electronic computer design, with the 701, 702, 705, and 709 products—all reasonably successful. Endicott, though still heavily involved with electronic accounting machines, had produced the 650 Magnetic Drum Calculator. The 650 sold hundreds of systems when only a few tens had been forecasted thus, by 1958, had achieved the largest volume sales of IBM's vacuum tube calculators.

The inter-laboratory competition to win responsibility for the new business computer was vigorous.

Poughkeepsie advocated a new architecture. Endicott, starting from the 650 structure, had a more aggressive design group that added new features and functions to the original 650 design. Endicott won, to Poughkeepsie's dismay—one of Ralph Palmer's few mistakes as, in retrospect, the architecture was a jungle and caused the product to fail.

Shortly thereafter, Tom Watson Jr.'s reorganization of the company went into effect. As the Data Processing Group was set up under T.V. Learson, the 7070 was moved to Poughkeepsie from Endicott. This was a monumental change— given Poughkeepsie's disdain for the design that was to have immediate unsettling effect on the developing strategies of the two product divisions.

I was promoted to area manager for Processing Systems, GPD's stored program computer products operation, reporting to Donald T. Spaulding.

My Processing Systems' key product was the recently announced and wildly successful 1401, discussed in more detail later. I settled into my GPD Processing Systems assignment planning a strategy built on the 1401. One issue was whether compatibility was important enough to sacrifice some raw processor performance? I decided to see what the front line sales management believed. Thus, one of the senior planners in the Division, Gus Rathe, a former sales manager in New Orleans, arranged a tour of the east, central and west sales regions where we conferred with the regional Data Processing Division Vice Presidents confronting them with the question. It turns out none had strong conviction one way or another. The west coast had many of the aerospace companies who wanted performance but that VP knew enough about the frailties of software to at least be on the fence. The east coast VP, whose region had much of the banking and insurance industries was ambivalent however, the central region VP, Warren Hume, was more of a visionary and thought compatibility would become important and was willing to pay some performance price to achieve it. I left those discussions, certainly without a mandate, however, more convinced that compatibility was extremely important.

My first compatibility test came with the disk-centric ARS project started in the San Jose Laboratory and transferred to Endicott by Ralph Palmer who wanted the ARS system to have architectural synergism with the 1401 program. However, engineers seeking raw performance superiority usually ignored compatibility and such was the case with the ARS project in Endicott. The engineering manager, Jim McDonald, led the transferred project and it was neither fish nor fowl as McDonald and his team worked on a unique architecture. I concluded I wanted the new system to have upward compatibility with the 1401. This meant the ARS might have a richer instruction set but would maintain the 1401 instruction set and data flows as a nucleus so all 1401 software programs could run on the ARS with extra architecture added for performance reasons meaning that ARS programs would not run on the 1401. As I worked on the redirection of the project I encountered some opposition. Some that argued the change would cause several months' schedule loss and compatibility was not worth it; others argued that performance had to rule and they argued against compatibility. I took a relevant group of planners and engineers to a place in upper New York State called Rocky Point where we would concentrate on all the pros and cons and I could more carefully investigate the schedule and performance consequences. The more I learned the more I was convinced the penalties were less important and compatibility was the correct strategy. My boss, Don Spaulding, came to the resort to hear the conclusions Spaulding supported me and I stated the decision to build a compatible system, later named the 1410.

Meanwhile, Poughkeepsie did not want the 7070 and its problems. S.W. "Red" Dunwell in the Poughkeepsie engineering organization set out to terminate the program by initiating a project, internally dubbed the "70AB," which had very simple goals: build a system with twice the performance of the 7070 at half the cost and package it in one rollagon, the cabinets used then for housing the electronics and cabling. It was a noble endeavor, but one not practically achievable in light of time and the existing 7070 customer commitments. Guided behind the scenes by Ralph Palmer, Vin Learson—to avoid the 7070 program collapsing due to DSD's lack of interest—ordered DSD to accept 7070 and get it into production—which was done, but which added fuel to the fires of inter-laboratory competition that was to promulgate future divergent plans. I was temporarily taken out of my job as systems manager of Processing Systems and assigned as the Endicott principal who was to ensure the successful transfer of the 7070 to Poughkeepsie.

The Model T Ford of the Computer Industry

In the mid-1950s with the transistor age dawning, the engineers responsible for the evolution of electric accounting machines were trying to bring the new semiconductor technology to their product lines. Most of these products were controlled by a plug board, a mechanism that allowed customers to change certain functions in the machines by the way wires were plugged into the board, in effect, a limited method to achieve flexibility in control. The Endicott Laboratory had a project underway to technologically modernize EAM (Electronic Accounting Machine) equipment. Called the TAM project, meaning "The Accounting Machine," the project could not foresee much reduction in cost and little gain in performance because of the limitation of the plug board.

At that time, IBM's European business was principally electric accounting machines. Arthur K. (Dick) Watson, Tom Jr.'s brother, headed the World Trade Corporation and he longed for a more important role for his WTC development operations, small laboratories in Germany, France, Netherlands, Sweden and England, each generally working on modifications to U.S.-designed products to tailor them to European use.

In the 1950s punched card equipment was IBM's bread and butter. In Europe the punched card was enjoying success albeit somewhat lagging applications that had developed in the U.S. Arthur K. Watson wanted to build a World Trade Company that was vertically integrated with products optimized for European customers rather than modifying U.S. product designs for European users. It was a Utopian view that was to be proven quite wrong. Nonetheless, in the 1956-'57 era, Dick Watson was searching for product independence. The relatively young Boeblingen (Germany) lab was chartered to develop a new family of punched card machines, the 3000 series, based on a small paper card approximately one-third of the size of the standard IBM punch card. The plan was for a full spectrum of sorters, punches, key punches, tabulating machines and printers to go with the new card media. The arguments were miniaturization would bring user benefits in reduced space, higher speeds and, they promised, less expensive mechanisms. However, technical problems sank the project. The products had been announced and some produced however, they did not work reliably. A task force brought in from the U.S. concluded the likelihood was low they could be made to work. Moreover, the audit challenged a number of the assumptions thus a conclusion was reached to kill the project and give the 3000 series customers standard electric accounting equipment at 3000 series prices, another losing proposition.

Striving to recover from the embarrassment of the 3000 series demise and noting the lack of progress in the TAM project, Dick Watson proposed the Endicott Laboratory and the European Laboratories join together in a common project to bring the new semiconductor technology to electronic accounting machines which were still IBM's principle products for small businesses. Dubbed the WWAM project, World Wide Accounting Machine, the international group set out to find the answer.

Still, the plug board presented an obstacle that could not be overcome and the project languished. It came to a surprising end when, in the Endicott lab, a young engineer, Francis Underwood, working independently, found the solution and, in the process revolutionized the industry. Underwood eliminated the plug board and turned instead to a very simple stored program architecture. Performance was outstanding, costs were very attractive and that processor, combined with Jonie Dayger's revolutionary new 600 line per minute chain printer, plus a new card reader for card input and a new card punch for card output became a new generation computer system. A young Endicott engineer, Charles E. (Chuck) Branscomb, was assigned to lead the development program dubbed the SPACE project (Stored Program Accounting and Calculating Equipment). Branscomb was a quiet, reflective, Georgian mechanical engineer, skilled in the mechanical technologies of punch card accounting. Branscomb, Underwood and their small team developed the product, which was announced as the IBM 1401 EDPM in December 1958. An entry level 1401 system rented for $2475, far under the $30,000 and higher lease prices of the larger systems. The 1401 was forecast to sell 5000 systems. At last count it sold well over 20,000 units and, over time there were a succession of follow-on products such as the 1401G, 1410, 1440, 7010 and others. The 1401 brought stored program computing to a new strata of smaller companies and propelled IBM further into the big time of computing. The 1401 was truly the Model T Ford of the computer industry.

Chuck Branscomb rose in IBM management ranks eventually becoming President of the System Development Division, responsible for the world-wide development of IBM's commercial computer systems. He retired from IBM with distinction but his most notable contribution was the important 1401 EDPM, the "Model T Ford of the computer industry."

Other products

. In Poughkeepsie, the 7090 scientific system—a transformation of an established design from vacuum tubes to transistors—was very well received, and, as stated earlier, the 7070 was weak. Customers with the older vacuum tube 702, 705, 650, and 305 business computers were not converting to the new 7070 architecture, as IBM had expected. Poughkeepsie found it necessary to produce the 7080, a transistorized version of the 705, to satisfy immediate growth needs of the large-business customer set.

Finally, in the System/360 prelude, there is the STRETCH project mentioned earlier in relation to the NSA. The STRETCH, or 7030, computer was begun in 1955. Intended to stretch both technologies and performance with improvement goals of two orders of magnitude over the 704, its history is a fascinating story in itself. Suffice it to say here that when it was first delivered in 1961, without substantial reprogramming efforts, it failed to meet a performance claim of being eight times the 7090. It was concluded that the 7030 would produce an average of 5.5, not eight times, the 7090. As an example of Tom Jr.'s keen sense of fairness, IBM reduced by price by 5.5 divided by 8, from $13.5 million to $8.8 million.

Development of the aforementioned systems were a critical prelude to the important System/360 breakthrough, but one key element was absent: programming. Not yet an established science, programming remained rudimentary. While some progress had been made in developing compilers—the programming method of translating problem statements from a high-level language to a specific machine architecture—the portability of code was restricted considerably by lack of standards, differences in compiler structures and performance problems.

Programming support from systems manufacturers was growing, but was still generally limited to programs for tape to printer, card to tape, and so on—utilities for loading, sorting, unloading, and peripheral units. There were some advances, though. An Input Output Control System produced for the 7090 was an important step toward modern operating systems. The 1401 had a powerful Report Generator, which—among other functions—more automatically organized and formatted the computer's printouts.

Nonetheless, programming was still sparse and elemental. As evidence, consider IBM's programming development budget in 1961: less than 5 percent of the research and development budget. This was at the very beginning of volume shipment of the new transistorized systems.

Thus, as we knock on the door of System/360's development, the IBM transistorized computer products consisted of two classes—business and scientific—with more than one incompatible family within each class, and with limited programming provided with each family. Addressing the problems of incompatibility was to spur the next major advances for IBM—advances that shaped the computer industry as a whole.

Chapter 4
An Integrated Family of Products

With great foresight, following the completion of Tom Watson Jr.'s reorganization in 1959, Tom Jr. mused that, while the 7000 and 1400 series systems were serving IBM well, we should start planning the next generation of products immediately. Thus, Tom Jr. assigned the Poughkeepsie laboratory that task. Fred Brooks was assigned the leadership role at DSD. Fred is the most complete man I have ever known. Brilliant, with a Harvard Ph.D., he combines high intelligence, natural leadership abilities, and articulateness with warmth. He is a loving husband, attentive father, and deeply religious. He is a contributor to his community as well as professional organizations and worthy causes. Very few people seem to have everything, including the time to work for their fellow man.

In 1966 Fred resigned from IBM to join the faculty of the University of North Carolina and teach, something he loved. Along the way in 1975 Fred authored a most important book on software management, "The Mythical Man Month" that is widely read and used as a text in many countries of the world. And, in addition to other books and publications he has updated the excellent "The Mythical Man Month" book with a 1995 edition.

I had been privileged to share in the life of the most "complete man I have known."

It was not that way earlier as, in 1961 Fred's and my relationship was adversarial. Fred came to Endicott seeking the cooperation of the General Products Division with his emerging plan for a family of systems, internally called the 8000 series. Building on the 70AB, DSD first attempted to develop to replace the 7070, Fred planned that the Data Systems Division would develop a mid-range 8106 for business computing with a scientific attachment, the 8108 and a large scientific computer, the 8112. Fred wanted my group to develop a small business system, the 8103, and a small scientific system, the 8104. The family's first system, the 8106 mid-range system, was scheduled for announcement in second quarter 1961.

Flush with the success of the 1401 and the 1410 in process— I was not willing to abandon those winners to join the 8000 series plan, which did not sit right with me in the first place because the 8103, 8104, 8108 and the 8112 were architecturally incompatible and I was certain compatibility was fundamentally important.

Fred's subsequent decision that DSD would produce its own entry-level products for the 8000 series was a huge problem that threatened the structure of Learson's Data Processing Group. The dividing line between the two divisions made clear that such products fit into GPD's mission. The two computer product divisions were on diverging courses! Additionally, there was a third confrontational partner in the mix.

At its United Kingdom laboratory, IBM's World Trade Corporation was active. For more than two years, the marketing organization—especially in Europe—had been calling for a small binary computer to replace the decimal-structured 1620 scientific computer, which was unsatisfactory for certain types of scientific applications. At the small U.K. Hursley Laboratory in the southwest of England, John W. Fairclough responded.

I met John Fairclough in 1958, on my first trip to Europe. A bright young engineer with broad vision, John had a short stint in IBM-U.S. and was well aware of this call for a small binary computer to compete with the bright new company, Digital Equipment Corporation and its popular scientific computing products, the PDP series. I was in Hursley to gain an understanding of the lab's capabilities and plans and to determine how we could better work with Hursley as my General Products Division's Processing Systems group developed advanced products.

A former Ferranti Brothers' engineer, Dr. William Elliott, led the Hursley Laboratory. He was typically British—procedural and status conscious. To me, it was clear that young John Fairclough was the lab's intellectual leader. John had also come to IBM from Ferranti. Seeking a mission for the Hursley laboratory, John had proposed a small 48-bit scientific computer with the code name SCAMP, aimed for 1961 announcement, which he hoped would be accepted across the IBM world. In 1960, Dr. Elliott departed and went to Imperial College. John Fairclough rightfully took over the Hursley lab, which by then had three or four hundred people.

Fairclough's SCAMP design incorporated a novel read-only memory for controls. Previous IBM systems had controls consisting of standard logic circuitry. Since control was literally intertwined with all the arithmetic and logic elements, control electronics were complex. It was not unusual for needed engineering changes to proliferate throughout many areas of the central electronics—much like unraveling a knitted woolen sweater. Hursley's approach, adapted from concepts invented by a group at Cambridge University, used small wired transformers that, when combined with timing pulses, emitted streams of pulses in sequence that served as the control electronics. By changing the "1" and "0" patterns of the control memory rather than altering the control electronics, the engineers had found a way to reduce control costs while substantially simplifying engineering changes and functionality extensions. The new "microcode" became a way of life in the computer industry.

In mid-1961, SCAMP was facing difficulties; sales volume forecasts were insufficient to prove the business case for continuing the project. John Fairclough promptly proposed SCAMP II, a larger system with R&D costs expected to be lower. John planned to capitalize on the original design and expand the sales volume to improve the project's prospects. The business case for the new model, though, was still insufficient. Fairclough, having learned his lesson well and reasoning that the market was 10 percent scientific and 90 percent business, proposed SCAMP III— a version of the original design that the lab hoped would acquire the requisite sales volume with its added capabilities in business applications.

Confrontation

The stage was set for confrontation. The General Products Division had its very successful product line centered on the 1401, with plans for expanding that base upward—even into the mission area of the Data Systems Division. DSD had some successful products—and others that were less successful—armed with the corporate mission to plan the next generation of computers--- was busy with its proposed 8000 series, which was being extended into GDP's domain. The World Trade Corporation, in its quest for a small scientific product, was touting SCAMP as it expanded the project in both business and scientific performance to obtain sales volume.

As the three organizations—GPD, DSD, and WTC—pursued their independent programs, Don Spaulding quickly determined that multiple, overlapping, and incompatible product lines would be a disaster. He convinced Learson of the potential difficulties for the company and recommended as a first step that I be brought from GPD to head Systems Planning and Development in the Data Systems Division, with all advanced projects reporting to me. Learson agreed.

In January 1961 I was in Milwaukee, Wisconsin calling on happy GPD customers to learn more of what they needed in the future. I received a call from GPD's Vice President for Engineering, John Haanstra, to abort my schedule and go to New York City to meet yet that evening with Vin Learson. I caught a plane to LaGuardia and at approximately 8:00 p.m. met with Vin in World Headquarters. Vin's exact words were simply: "Bo, Poughkeepsie has this 8000 series plan. I want you to go to Poughkeepsie. If it is right, do it. If it is not right, do what is right!" It turns out the decisive Vin Learson had done this without first talking to DSD President Bill McWhirter or McWhirter's VP Engineering, Dr. Charles DeCarlo. I know this because Learson called McWhirter while I sat in his office that evening and it was clear this was new news to McWhirter. While I was sort of persona non grata in Poughkeepsie stemming from the recent 7070 transfer, McWhirter was a team player and he dutifully accepted me in

DSD. We became fast friends which remained so until Bill's death in 1994.

A day after my evening meeting with Learson, it was announced that I was replacing the Data Systems Divisions Director of Systems Development and Planning, Max Femmer, reporting to Dr. DeCarlo. This was the biggest job in DSD development operations with more than 1000 people in the organization. This announcement had to be discouraging to the able Fred Brooks because my position was over Fred, the 8000 series and the other DSD systems responsibilities. It became an undeclared war as Fred and his advocates worked around me to rush the 8106 to announcement to seal in the 8000 series plan. And my boss, Director of Engineering Dr. Charlie DeCarlo was the ring leader of my opposition. I was working hard on assessing the 8000 series and basically operated alone. The hostility was constantly evident. People who reported to me such as Frank Cumminsky who headed Planning were working with Dr. DeCarlo and Fred Brooks hoping to somehow break through the growing impedance I presented. Vin Learson's Group Staff were well aware of what was going on. Thus, to insure that I had a fair opportunity to do Vin Learson's bidding of "If it [8000 series] is right, do it. It is not right, do what is right" in April 1961 Vin Learson caused the replacement of Dr. DeCarlo with my former boss, Jerrier Haddad. By May 1961 I concluded the 8000 series would be a serious blunder, in part because of the lack of compatibility within the systems family. I did not buy Dr. Brooks' arguments that recompilation would be acceptable to make it possible for the programming from all the dissimilar architectures of existing products to work effectively on the dissimilar architectures of the 8000 series. There were other important reasons to scrap the 8000 series plan including technology choice. Jerrier Haddad backed my decision; the 8000 Series plan was killed.

Why Change System Structures?

Between 1952 and 1962, seven incompatible families of systems emerged at IBM, each with a number of serious problems both from the standpoint of users and the company. An enumeration of the problems reveals how serious a business problem this situation was becoming.

First, with so many types of architectures, IBM was spending most of its development resources propagating the wide variety of central processors. Few development resources went into either peripherals or programming. This had serious ramifications for users. While a user could move from one processor to another in a family perhaps as twice as fast, the improvement in throughput—the problem solving ability of the computer—might be only 10 percent, since the existing disk or tape peripheral devices and programming simply were not keeping pace with the central processors. My view at the time was that technology was moving fast enough that internal processing performance should increase each generation by 50 to 100 percent, prices should come down at least 25 percent, and throughput—taking peripherals into account—should improve by at least 25 percent.

Second, sales of any single system or family were in most cases too small to justify a disk, magnetic tape unit, or other peripheral optimized to that particular architecture. New peripheral devices, therefore, were suboptimized across differing architectures: business systems, with their serial-by-character, variable field length characteristics, and the scientific units, with their binary, highly parallel characteristics. This meant expensive electronic adapters had to be designed to attach available peripheral devices to the different systems. Some systems had to live with older peripherals; no justification could be found for new developments. This, too, constrained throughput.

Third, the demands of the varying architectures and the rapid invention and innovation taking place in transistors thwarted some of IBM's plans. When transistors came into service, IBM strove to take fullest advantage of the new capability through standardization of circuits and centralization of circuit design groups. In 1955, a new circuit packaging technology was developed for IBM's future transistorized products, called the Standard Modular System. Ralph Palmer and Jerry Haddad, leaders of a circuit standardization plan, hoped that approximately one hundred types of printed circuit cards would suffice for the 7000, 1400, and 1600 families to be designed. By 1963, though, there were more than 2,500 types of circuit cards, along with new problems in manufacturing logistics, testing, inventory control, spare parts stocking in the field, and customer engineer training. This proliferation of circuit cards also limited engineering's ability to improve the designs.

Fourth, IBM had increasing programming problems and demands. In the early 1960s, the company learned that functional capability, performance, and programming reliability were key to volume sales. Unfortunately, though, programming was not yet an established science and there was a relatively small group of programmers assigned to the individual systems in the seven different families. This only exacerbated the problems. Each system type needed its array of compilers for a variety of languages (COBOL, Fortran, etc.), usually designed separately for magnetic tape input-output and disk input-output. Each system type also needed utility programming such as loaders, memory dump, sort and peripheral control and so on. Users were generally unable to move across system boundaries and sometimes had difficulty moving from one processor of a family to another. Migration was difficult; conversion was very expensive, often prohibitive. Available resources had to be spread across a wide range of processor and peripheral equipment products.

Fifth, the split between scientific and business problems imposed real limits on users. Increasingly, the scientific areas needed the alpha-numeric and peripheral powers of the business systems, and the business systems needed the logical and computational powers of the scientific machines.

Sixth, main memory was chronically too small in each IBM system. As discussed in the section on Dr. Cuthbert Hurd, in 1955 Dr. John von Neumann estimated that "10,000 binary words of memory should be sufficient for any problem" he could foresee, and "allowing for inefficient programming, 20,000 would be ample." In the addressing structure of subsequent scientific computers, we generously allowed for 32,678 words. On successive designs, the company repeatedly increased the memory addressing structure and the amount of memory available. Nonetheless, in each case the memory capacity proved to be far less than users eventually required. It became clear than orders of magnitude improvements in main memory addressing and capacity would be required—which in and of itself demanded a substantial change in system architecture. Even today's low-cost personal computers, for example, have more 200 times the von Neumann estimate.

Finally, there was the problem of how many combinations of numbers, characters, and symbols were available for use. As of 1961, the internal coding of most systems used four binary bits—the 1s and 0s—to represent a single decimal digit and six binary bits to represent an alphabetic character. This allowed only 64 combinations, which was clearly restrictive.

As to the World Trade Corporation's divergent SCAMP plan, in June 1961 Arthur Watson came to my office in Poughkeepsie to discuss Hursley's plan and to understand my plans. He brought his executive assistant, Billy Christiansen, who argued for WTC not joining my plan but to proceed with SCAMP. However, Watson overruled him and WTC did join the New Product Plan, with a demand that the European Laboratories play a meaningful role—which I pledged.

I set about launching an international effort to build the new product line, the family of systems that came to market as Systems/360.

Once the decision was made to kill the 8000 series, I had to move swiftly as the Poughkeepsie engineering team was in chaos with their basic plan suddenly killed and no alternate plan in place. We had two problems: the first was to get going on the development of a unified systems product line that would serve all of IBM and the second was to set up programs to buy time by quickly extending the present products lines, the 7090, 7080, 7070 and to try and fill an enormous gap in IBM's mid-range scientific product line where Digital Equipment was eating IBM alive. I convened the key DSD laboratory management group to the serene old Inn at Saratoga Springs, N.Y. where we could concentrate for a few days and put plans in place. While I had to fly out in my plane for a day at one point, we accomplished a lot. One of my most important decisions was that I surprised Fred Brooks and offered him the key New Product Line position. The stoic Fred had expected to be moved out of DSD given the fight between us. However, I had such respect for Fred's intellect and leadership ability that he was the clear choice to lead the New Product Line development. I did have one condition in that I worried Fred's

Harvard training in architecture led him to more exotic and, I feared, expensive and complex designs. Thus, I wanted Dr. Gene Amdahl as the principal architect. I was insistent on Amdahl because of my view of Gene's brilliant architectural ability.

Re Amdahl, in a lifetime few have the opportunity to meet a genius. I have been fortunate to know several: Dr. Wernher von Braun of NASA, IBM's (and the Department of Defense's) Dr. Eugene Fubini, and IBM's Dr. John Cocke. Perhaps the most intriguing of all was Gene Amdahl, who came to IBM in 1952 after building a computer, the WISC, for his graduate thesis at the University of Wisconsin.

Nathaniel Rochester, who had been one of the principal planners on the Defense Calculator project, brought Gene to IBM. Gene's job was to improve the architecture of the 701 and help develop follow-on products. Working with other fine professionals, including Elaine Boehm, he produced indexing, floating point computation and a variety of more powerful computational abilities that became the 704, later the popular 709, and eventually the transistorized 7090.

Amdahl really hit his stride on the design of the proposed supercomputer for Livermore Laboratories, the LARC project, and was bitterly disappointed that it did not proceed. He became dissatisfied, as was I. Gene's disappointment, though, was enough to make him resign from IBM, and he joined Ramo Woolridge, a new company in California formed to be systems manager of the United States' new intercontinental ballistic missile program. However, Ramo Woolridge had a bigger appetite than just the ICBM program, though, and a new division was formed to go after broader business. Two brothers, Chris and Dean Wanless, were responsible for bringing a diversified R.W. business into being, and as they felt their way they did not realize they had a certifiable genius on their team, in the person of Gene. Instead, they used Dr. Amdahl to write proposals seeking new business, such as automated equipment for the Post Office.

Gene really wanted to live in sunny California, but eventually he tired of Ramo Woolridge's ineffective use of his abilities. He returned to the IBM Research Division in 1960 where he worked on high-speed computer architectures. That is where he was when I tapped him for the NPL team.

I was delighted when Fred accepted the New Product Line position I offered. I received Research Director Mannie Piore's approval to move Gene Amdahl to the New Product Line architecture position on the condition that I would put a supercomputer into the product line family. I would not commit that such a system could be available with the first of the planned family of products however did agree to put a high performance system in the plan.

The NPL project replaced the 8000 series and was to be a family of systems which I intended would serve the U.S. and World Trade Corporation alike with small systems that fit the General Products Division mission and larger systems that fit the Data Systems Division mission as well as fulfilling the needs of the World Trade Corporation.

One of several reasons I objected to the 8000 series plan was that it planned to use the aging SMS technology and superior alternatives were on the horizon. A leading candidate for the new work-horse technology was being developed in an internal research program dubbed Project Compac. That advanced effort was being led by one of the senior engineering specialists, Robert Domenico. This project was based upon solid logic technology using microminiature transistors and diodes with special resistors made with a paste and trimmed to the necessary resistive values. The solid logic technology components were fabricated in a new process that promised substantial size reductions, lower costs and higher speed although initially only one circuit on a small ceramic substrate. Later, multiple circuit versions were developed. S.L.T. was innovative with major gains expected. I wanted to proceed with the Project Compac technology however, Research Director, Dr. Mannie Piore, wanted to first consider whether integrated circuits could be brought in soon enough as I.C.s had promise of gains beyond the solid logic technology. Thus, a high level task force was established consisting of Erich Bloch, Dr. John Gibson who headed IBM's components operations and me. It was clear that I.C.s were probably three or more years distant thus we unanimously concluded the only choice for the New Product Line was S.L.T. With that recommendation, IBM commenced the extraordinary investment in this new technology and eventually it was produced as the foundation for the New Product Line. It became a success and demonstrated IBM component leadership. And, in hindsight, integrated circuits were much later than we had projected in the 1961 study; thus it is clear we made the correct decision at that time.

To Bob Domenico's anger, the head of components operations, Dr. John Gibson elected to put one of his "aces," Erich Bloch, in charge of S.L.T. development and production. There was no question that Erich Bloch was very competent, however, I had empathy for Bob Domenico as he had taken the innovative technology from concept to the brink of reality. Bob Domenico left IBM a few years later, joined Max Paley and Dr. Michael Flynn in a consulting venture, Palyn Associates, a company formed in approximately 1973 that continued to operate until Max Paley passed away in 1998. I was pleased that Bob Domenico found a happy life after the disappointment of the Project Compac leadership decision in 1961.

Erich Bloch led S.L.T. to production and then moved to various senior management positions in the mainstream development area. He was a Vice President in the Systems Development Division and later became a Vice President in the Components Division managing the giant E. Fishkill operations, the center of IBM's bipolar semiconductor technology. In 1982 Erich was moved to the Corporate Staff, replacing Jerrier Haddad, responsible for outside technical relationships. He did an excellent job and, among other things, was a founder and first Chairman of the Semiconductor Research Cooperative, later renamed the Semiconductor Research Corporation which continues today to do important U.S. semiconductor research.

Erich Bloch retired from IBM shortly after I departed and then entered his most important professional assignment, as he became Administrator for the National Science Foundation under the Bush administration. I hear scientists who pursued abstract research, were angry with the programs Bloch introduced however, Erich set in place a number of innovative nation-wide engineering initiatives that, to my way of thinking, made a lot of sense in better maximizing the value of U.S. funded research at NSF.

With the change to President Clinton's administration, Bloch departed NSF and today serves as a Distinguished Fellow on the Council of Competitiveness, a privately supported, non-partisan, non-profit forum of highly experienced senior leaders who focus on U.S. leadership in technology innovation and global markets. It is a fitting and distinguished role for a professional of Erich Bloch's experience and ability.

At Saratoga Springs, with the NPL leader in place, I turned to the task of extending the current product line and filling the scientific gap. Larry Kanter was assigned the project to speed up the popular 7090 scientific computer; the 7094 and, later, 7094 II products emerged rather quickly under Larry Kanter's leadership and were highly successful. Dick Case was assigned the project to build mid-range scientific computers based on the 7090 architecture, and in rather quick order the 7040 and 7044 emerged which were very successful; Richard Tracy extended IBM's high performance business computer, the 7080 and John Haanstra agreed to transfer the 1410 to DSD to better weld the divisions' product plans. To this DSD added a 7010 compatible extension to the 1410 family and these "temporizers" as I called them, helped keep the current the processor lines competitive in the market place as the core of DSD's engineering operations worked on the bold New Product Line.

During the second half of 1961 I concentrated on the New Product Line plan while keeping my eyes on the current product line extensions. The General Products Division's success continued and while their VP-Engineering, John Haanstra, pledged to develop the low end, high volume system in the NPL family, John had shown signs of wanting to command his own destiny and there were growing suspicions that John might cause the General Products Division to diverge to another plan. Divergence by GPD could bring the whole plan down as the New Product Line was being developed using the new hybrid microminiature Solid Logic Technology. That, in addition to the high software development costs, necessitated as much volume as possible for the New Product Line to be viable financially. The low end of the New Product Line was expected to require half the volume of the Solid Logic Technology capacity and, if John Haanstra, deferred GPD's NPL system for another round of contemporary systems using the mature Standard Modular System technology, the amortization of SLT would be left to the Data Systems Division's products and Hursley Lab's product thus the estimated volume was insufficient to make the business case. Once again, Don Spaulding of Vin Learson's staff was on top of the situation. To thwart the possibility that John Haanstra would abort the New Product Line plan by bolting off to do another round of contemporary products, Don Spaulding conceived a plan to secure John Haanstra's commitment. Spaulding's plan is worth a chapter in any textbook of creative management. He persuaded Vin Learson to set up an interdivisional study called the SPREAD study, meaning Systems Products Research and Development which we were to affectionately call "Spaulding's Plan to Reorganize Each and All Divisions." The SPREAD task force charge was to detail the international new product line plan and Spaulding cannily suggested that John Haanstra lead the study with a report scheduled to Tom Watson, Jr. and President Al Williams in January 1962. A group of a thirteen specialists were assembled from across IBM: Dr. Fred Brooks and Bill Heising, from my division, DSD, Martin Kelly and Jerry Svigals from the General Products Division, John Fairclough from the World Trade Corporation, Cy Rosen from DPD, the U.S. sales division, Dr. Herbert Hellerman from Research, Doug Newton from the Advanced Systems Development Division, Bruce Oldfield and Joel Aron from the Federal Systems Division and Walter Johnson from the Group Staff. The SPREAD study convened approximately Nov. 1, 1961 and we worked off-site at a motel in Stamford, CT. John Haanstra was Chairman and I was Vice Chairman. In retrospect I believe the stars of the study were the Data Systems Division's Dr. Fred Brooks and WTC's John Fairclough. Overall, this group put a lot of substance into the NPL plan that I had started and, more important, per Don Spaulding's strategy, produced a company-wide systems plan.

Don Spaulding's purpose almost came unhooked at the end of November 1961,mid-way through the SPREAD task force when Orland Scott was promoted to be President of the prestigious Data Processing Division and John Haanstra was named President of the General Products Division. We were all thrilled that John however he had to leave the SPREAD study to tend to GPD and I took over coordination for the month of December. Spaulding and I worried that John might not be committed to the plan. However, early in January, by Spaulding's plan, John Haanstra made the task force results presentation to Tom Watson, Jr., Al Williams, Vin Learson and other senior management who accepted the plan and John was "committed," or so we thought.

In January 1962, upon completion of the SPREAD study, I returned to the Data Systems Division assignment working on both the short term products and the New Product Line plan. There was a follow-on study in 1962 that was not as well known as the SPREAD study however, had its long term, positive effect. The study was chartered to examine the world of peripherals such as magnetic disks and magnetic tapes to insure homologation of the input-output equipment commensurate with the processor master plan set by SPREAD. This study, called STORE, was led by Jerrier Haddad and set in place a direction that had important influence, especially in the pivotal magnetic disk developments.

Moving Forward

I spent 1962 and 1963 concentrating on the New Product Line plan while keeping my eyes on extensions to the existing product line. The General Products Division's success continued, led by President John Haanstra, who had his division developing the low-end, high-volume system in the New Product Line family. However, trouble was to brew later as will be discussed.

Haanstra was one of the most electrifying—and frustrating—engineering leaders in IBM during the great years. I called him "Big John," not only because he was physically a big man, but because of my respect for him. He was an excellent engineer, strong-willed and a natural leader, and company heads noticed. John rose quickly through the ranks at the San Jose, California, disk engineering laboratory.

I first met John in 1955 when I was working for Jerrier Haddad in World Headquarters, seeking projects that were candidates for outside funding during the 1955 budgetary problems. At the fledgling San Jose lab, young John Haanstra stepped forward and offered to prepare the proposal and terms for the research contract I proposed to pursue—a task I usually undertook myself. I admired him from then on as a "take-charge" professional.

In 1958, John Haanstra was brought east to a high management position, director of engineering, in the newly formed General Products Division. Paced by the hugely popular 1401 system, GPD grew rapidly, was very profitable, and generally was more successful than the large systems operation, the Data Systems Division. A year after being formed, the division's success was recognized when Haanstra became the first vice president to be named in a division. We were all proud of John and the engineering force was electrified.

His promotion meant new heights were possible for engineers. His later elevation to President of the General Products Division made us even more proud of John. However, Haanstra's penchant for controlling his own destiny caused Don Spaulding and I to constantly worry that John would, sooner or later, try to take GPD down another path which would bring the NPL plan crashing down. These concerns proved to be well founded as will be discussed later.

It was an exciting time. With our plan to integrate scientific and business capabilities into single processors, businesses would no longer need to split computing facilities. They would accrue considerable efficiency and productivity gains. Scientific applications would have the alphanumeric and variable field length capabilities, and business applications would have scientific capabilities. IBM undertook to specify and develop a single high-level language, PL-1, to serve both types of applications, something that could replace FORTRAN and COBOL.

Our New Product Line was a bold plan to develop systems that were both upward- and downward-compatible, meaning that with the requisite memory and input-output configuration the software written for small systems could run on the large systems and vice versa. We were resolving the major problems that had plagued contemporary systems.

We planned a family of central processing units, each of which was to be adept at scientific and business applications as well as have facilities for the foreseen communications processing (teleprocessing) and event-driven (real time) applications areas. The plan called for the data formats, instruction repertoires, and principal interfaces within each processor to be identical. Each system would have the capability of processing both decimal and binary formatted information, have variable field length capability, and floating point arithmetic capabilities.

Undertaking both upward and downward capability was extremely ambitious. There were major technical questions: would the smallest processor's instruction repertoire and data flow be so rich as to thwart our cost objectives? Could the largest processors, constrained by the instruction repertoire limited by the low-cost family members, be powerful enough to be competitive? We knew that if we were successful, it would only be a few years before it would take fewer and fewer of IBM's development resources to administer processor evolution—freeing more resources to work on programming and peripheral device development.

Unfortunately, major problems began to crop up with the design plans. The most significant was the speed-cost problem: how to design a common architecture so the cost of the entry-level systems was low enough to be priced reasonably, while the architecture of the high-end systems would be powerful enough to achieve the necessary computational performance? Fred Brooks, Gene Amdahl, Gerrit Blaauw, Dr. William Wright, William Hanf and others struggled for months without success. Fred Brooks conceived a plan that, once again, demonstrated his superb leadership—he organized a design competition to solve the problem. Four teams worked independently to find the right balance between architecture, cost, and performance, the theory being the honor of succeeding would motivate the architects to do their best. It worked! From the design competition emerged a unicameral architecture, which the engineering teams immediately went to work to reduce to practice.

There were many challenges, with innovative solutions that set new "standards" for computer compatibility. To solve the throughput problems brought about by the suboptimal and limited peripheral devices then in existence, the New Product Line design called upon a 7030 invention to have standard peripheral device interfaces across each of the processors in the family. Since all the interfaces were identical, peripheral devices such as disks could be designed for the entire family and thus enjoy the fullest volume potential. Electronics on the peripheral side of the interface would adapt the differing devices to the standard processor interface. This meant that when new technologies permitted a new peripheral device, it could be attached with relative ease and reprogramming for the new devices would be simplified.

But another show-stopping problem loomed. How would we migrate existing customers and their growing investment in applications programming to the new architecture? We believed that with compatible families, programming could be developed independently of individual processors. A control program, for example, could be planned to operate three different processors and could be used on all members in the family given each processor had the required memory and peripheral devices. Language compilers and utility programs such as loaders could be developed to serve the entire family. The new freedom from individual hardware promised very substantial gains in programmer productivity. More important, hardware independence signified that a customer would have to change little or no programming to move to higher performance processors or to move from a tape-oriented to a disk-oriented system. Similarly, if economics were adverse or a business decision to decentralize was made, an installation could move down in computing capacity without reprogramming. Programming independence was perhaps the single most important advantage of the New Product Line structure.

Would customers be convinced? Fred Brooks believed program translation was the answer, and put in place an extensive effort that proceeded for two years until it was clear the software translation problems were not readily surmountable. It was a time of crisis: if we lacked an expedient vehicle to allow current customers to easily migrate to the new family, the forecasts would have been dramatically reduced and the business case for investing in our development project would fall through the floor. But necessity breeds invention and, once again, that high-powered team came through. An engineer, Stewart Tucker, noted that the New Product Line architecture had sufficient registers and data paths that all the disparate existing systems' data flows could be contained within its structure. The issue became the different instruction sets. However, the Models 30, 40, 50, 60 and 62 instruction repertoires were implemented in the read-only memory concept first used by the Hursley Laboratory in Britain, making the incremental cost to add each prior system's instruction set to the NPL both doable and affordable. Emulation was born, and all but the one of the NPL's models could run the programs of existing systems as desired by each system's market niche.

It was only the most powerful of the family members, the Model 70, which did not have the emulation capability. The Model 70 did not use the microcode approach because the project manager, Daniel Doody, did not believe he could meet the performance specs that way. Instead, the Model 70's instruction repertoire was designed in the conventional method of wires and circuits. Part of the problem was the high performance solid logic circuits were not meeting the original specification, so Doody believed he could squeeze more performance into the design via hard-wired controls. That decision limited the Model 70's sales—it was to be placed in accounts that needed the performance and that were willing to develop new applications programming. In actuality, though, most customers ran their applications in emulation mode for several years.

One of the more profound aspects of the New Product Lines' architecture had to do with the addressing. As I mentioned earlier, experience had shown that every system IBM had produced had too little main memory. We knew we had to reach out and provide vast new memory capacities. The contemporary means of addressing was to have a set of bits (binary "1s" and "0s") sufficient to address each discrete location directly. However, the larger the memory to be addressed, the longer the stream of address bits, 32 or even 36 bits to plan for the future. If the past were prologue, even those address lengths would likely prove eventually to be insufficient. Address lengths like that would have been costly baggage, so Fred and his architects conceived the "base register-offset" addressing structure. In this structure, one set of address bits defined a general area and another set of address bits designated a specific address within the general area. This permitted vastly larger memory capacities with a structure that could be rather simply expanded in the future, if more memory capacity proved necessary.

The scientific types believed there would be a performance penalty without direct addressing, and complained fiercely to senior management about our approach. The opponents were formidable: for instance, Dr. John Bertram from Research and Dr. John Backus, the developer of FORTRAN, were critical and sowed seeds of doubt. We were able to defend the addressing structure successfully, though, and the project continued. In truth, management had no choice—the critics offered no alternative, and changing the design would have wrecked the New Product Line plan. I have been convinced of the merit of Fred's addressing architecture over the more than thirty years of its durability, which continues to this day.

The NPL project was stressful, and that stress began to wear on people. Eventually, a crisis developed—Fred Brooks and Gene Amdahl were at "sword's point." Gene was not proving to be the omnipotent architect that had been expected, and a lot of the invention was coming from others, including Gerrit Blaauw. The intellectual dispute between Fred and Gene increased. Appealing to each, I worked with them to resolve the problem. But the schism widened, finally coming to a head in the summer of 1963 when the two men somberly came to my office and announced they could not get along. I was told I would have to choose between them. I was inwardly irritated that Fred, as the project manager, had not solved the problem and was putting me in this spot. Because most of the critical architecture issues were resolved, it was tough to make the case for Gene—but I did not know what problems might lie ahead and I wanted to keep Gene on the project. Choosing Fred would mean Gene would be gone permanently, but I did not believe Fred would withdraw at this stage of the project were I to choose Gene. So, I folded my arms and chose Gene. Fred surprised me, and departed the project. I decided to let Fred "think about it" for a few days. Word quickly reached the Bill McWhirter, the DSD president, who was attending a senior management conference in Colorado. He panicked, rejected my "wait it out" strategy, and returned immediately to meet with Fred—who relented and returned. Perhaps my strategy would have worked; I cannot say for certain. Bill McWhirter's actions at least ensured Fred's return. Bill also gave Fred a raise, although he was not due for one. It became my duty to inform Fred. He thanked me, and in a moment of stupidity I said, "Don't thank me. Thank Bill McWhirter!" Fred now knew I had disagreed with the salary increase. I had foolishly insulted him and I know Fred remembers that incident to this day.

Giving Fred that undue raise was not the only thing Bill McWhirter did. There are many other things for which he deserves thanks. Under his reign, the Data Systems Division prospered and IBM computing grew abundantly.

William B. McWhirter was managing the obscure Supplies Division in 1958 when Tom Watson Jr. tapped him as president of DSD, the large systems division chartered in the Williamsburg reorganization. Bill had been a crackerjack salesman who had risen to senior management ranks. And while it is tough to get excited about punch card stock, typewriter ribbons, and tab machine forms, Bill ran that division well. Then, all of a sudden, there he was in the very vortex of IBM's computer future.

McWhirter was a consummate gentleman, well aware when his opponents held him in disdain—which, if anything, spurred him to try harder. Over time, however, Bill wore thin with the able Vin Learson, who came to distrust Bill's judgment. George Kennard replaced Bill as president of the Data Systems Division.

It was a Keystone Kops tragicomedy with tough consequences. Yet, it is also true, that in IBM's glory days, Bill McWhirter was one who contributed much to building the company. *(The Keystone Kops featured in a series of 1912 – 1913 silent film comedies about a totally incompetent group of policemen.)*

As for Gene Amdahl, he eventually left IBM again—in January 1970. Ostensibly, it was over a senior management decision not to put the supercomputer on which he had worked for years into volume production, opting instead for IBM to build only a couple—so the company could have the fastest computer on the planet. I think, though, that Gene was quite unhappy over a number of issues, and would have quit the next day had the hamburgers in the cafeteria been cold.

Gene went on to found Amdahl Corp. which became quite successful. In November 1972, Maria and I met Gene and his wife, Marian, at a National Academy of Engineering convocation in Washington. Over dinner the discussion turned to his efforts to find financing for his new company, and he asked my advice. "I never give free advice," I said. "My fee is $1.00." Gene hauled out a dollar bill and I asked both Marian and Gene to sign it as the "contract" for my advice. They did and I said, half in jest: "Go west"—meaning Japan. To this day I carry that autographed dollar bill as a souvenir of the past. I have no idea whether Gene had already started discussions with Fujitsu, but soon the Japanese owned a controlling interest in Amdahl. Later, Fujitsu moved Gene aside, and eventually he elected to resign.

The Gene Amdahl saga is bittersweet. This genius could envision the inner operations of a computer during the computational process in a manner that was astounding, and unmatched by anyone in my experience. He made great contributions to IBM in his designs of the 700 and 7000 series and he contributed in important ways to the design of System/360. And I am convinced that it was his brilliant mind and complete understanding of the complex 360 architecture that enabled the Japanese to copy IBM's designs three years sooner than if they had to earn that knowledge on their own. In effect, Gene Amdahl created the plug compatible central processor industry. However, since Amdahl Corp. Gene founded Trilogy which failed after spending hundreds of millions of dollars and, later, he founded Andor which also failed. On balance, his career has not matched his great intellect. Rather it has turned to a sad story for a man that really is a "certifiable genius!"

The Haanstra Problem Re-emerges

IBM had accepted the SPREAD plan, and Haanstra—true to his word—had the General Products Division working on the smallest, highest-volume system in the family as well as disk products to support the new product line. We were making great progress on the New Product Line. Then we were faced again with an old problem—John Haanstra's yen for independence. The canny Haanstra had a secret project developing a faster 1401, his division's workhorse product. It had gone on for a couple of years.

Unfortunately for him, Haanstra "blinked." In the fourth quarter of 1963, Honeywell announced a 1401-compatible system with better performance for the price than IBM's 1401 family offerings. As the Honeywell system began to sell well, Haanstra changed his mind, reneged on his commitment to the New Product Line, and came forward in January 1964 arguing for his "secret" project, the 1401S—a system five times faster than the existing product. He proposed to delay the Model 30, his division's member of the New Product Line family, in favor of the secret project.

I was crestfallen when I heard that Tom Watson Jr., in a mid-January 1964 meeting with Haanstra, reportedly described John's plan as "the finest fiftieth birthday present a man could wish for." I thought our New Product Line was finished. Late that month I departed for a scheduled trip to Europe in a complete state of depression, believing it was all over but the funeral march.

There were other problems compounding these feelings. As the new line moved closer to announcement, a number of contrary viewpoints arose. Some were financial and some were technical. Ironically, of all the opposition, the Strategic Planning people, instead of exhibiting bold vision, argued the compatible line meant putting "all the eggs in one basket"—and that IBM would thus face a double jeopardy. Failure of customer acceptance, the reasoning went, would of course be sweeping; this would be compounded by the fact that even if users accepted the architecture and the full line went into production, a superior competitive approach would adversely affect IBM's entire systems products. They argued that even a successful competitor attack on any of the existing seven families would not be disastrous in the same proportions.

However, the significance of the unified family prevailed. Behind the scenes, and in response to the specific Haanstra threat, Vin Learson and others concluded the lack of a low-end system in the NPL family would ruin the plan and in February the Corporate office reversed itself, killing the 1401S. In fact, the user appeal of compatibility, standard interfaces, programming independence and the perceived application power of the new architecture was too persuasive. IBM senior management never seriously considered the philosophy of "continuing with several architectures."

Fred Brooks phoned me in Europe to tell me the news. The full NPL plan was on, we were speed to announcement as early as March, and Haanstra was dropped as GPD president for this last-minute deviation from a plan to which he had earlier committed. In stepped Clarence Frizzell to replace Haanstra as GPD president.

"Friz" Frizzell was senior management's "all-purpose utility infielder/outfielder." Over his career, he handled special assignments and crises with aplomb. In charge of the input and output equipment on the Defense Calculator, he went on to distinguish himself leading special tasks such as the audit of the ill-fated European 3000 series, shepherding the 727 tape drives through a technical crisis, and being a division president. Quiet, warm, affable, and competent, Friz was a popular leader in IBM throughout the 1950s, 1960s and 1970s.

In the early 1980s, Clarence Frizzell retired as general manager of the GPD manufacturing plant in San Jose, California. Friz left me with so many fond memories. With his passing at the end of 1994 went one of IBM's major contributors during the great years.

Friz served well as John Haanstra's replacement. What happened to Haanstra is a good illustration of the fact that in Tom Watson Jr.'s IBM, a single mistake was not fatal. People with ability were given a second chance—as we will see.

The New Product Line Debuts … With Problems for Me

On April 7, 1964, IBM announced the New Product Line: System/360 Models 30, 40, 50, 60, 62, and 70. Why 360? The marketing people had come up with this imaginative name to infer 360 degrees, meaning the products covered the full circle of applications. However, sales were sluggish at first, as customers evaluated the new line. Eventually, though, System/360 sales turned on then turned to a landslide and the systems went on to success far beyond our wildest dreams and forecasts.

Shortly after the announcement, Tom Watson Jr. hosted a celebratory dinner in New York City. At the dinner, the chairman surprised me, announcing that he had junked his prepared script and instead would simply read a letter Fred Brooks had written to him citing my efforts in organizing the New Product Line project. The dinner turned into a moment of high honor for me, with my wife Maria and my System/360 compatriots there, all thanks to Fred's sensitivity and alertness.

However, as Dan Jenkins once wrote: "Life is tough and then you die!" Storm clouds were brewing: at the very moment System/360 sales were becoming a windfall for IBM, I received the first dire signal that change was underway. That change had a significant effect on my career.

The signal came in May 1964, a month after Systems/360 were announced. Dr. John Cocke was an IBM researcher for whom I had a lot of respect. He told me that MIT did not like System/360 and would probably not purchase our new products. His report spoke to Tom Watson Jr.'s prescience.

In 1961 Watson Jr. had commented to me that he thought the research at MIT was unique and probably relevant to our efforts at IBM. "Keep an eye on it," he said. I did, but not as effectively as I should have. Rather than travel to Cambridge myself, I delegated. At one point, I sent Gene Amdahl to a meeting at MIT. On another occasion, I sent Phil Stoughton, an IBM engineer, to the campus. They may well have reported the MIT thinking to me, but if they did I certainly did not absorb those reports. The problem: we did not incorporate a key element of MIT's research directly into the System/360 design. My failure to keep current with MIT's efforts was to become a major problem.

On Saturday, June 5, 1964, Dr. Cocke and I flew to Boston in the Cessna 210 I shared with Joe Logue. That afternoon we met with two MIT professors, Robert Fano and Fernando Corbetto. An articulate Corbetto confirmed Dr. Cocke's earlier input that MIT did not like our systems architecture and most likely would not purchase 360 systems. He cited four problems: three were easily fixable, and one was not—the absence of an architecture MIT had developed in their Multics software system research to allow the system to do most of the complex work in managing the allocation of memory and storage addresses, essential to on-line, terminal-centric environments. It was called "dynamic address translation." Somewhere along the way, that message had not gotten through to me. Now the System/360 architecture was set in concrete.

Soon, the word was out and Headquarters was furious with me. Meanwhile, an opportunistic General Electric—anxious to strengthen its fortunes in the computer industry—promised new systems with dynamic address translation. MIT elected to purchase those products.

Back at the Poughkeepsie Laboratory, we planned a solution. The fine 360 architect, Dr. Gerrit Blaauw, began to design modifications to a System/360 Model 62 to provide dynamic address translation. An angry IBM Headquarters, though, was not content to let us solve the problem and established a special project managed by Watts Humphries, who had not even been on the original System/360 team. Control of the new project, Model 67, was moved out of my domain.

IBM allocated huge sums to develop the operating system, which later came to be known as the Time Sharing System. It played a role for a few years until, ironically, I scrapped it in 1970. IBM was fortunate: GE had severe difficulties with its new systems and the complex software and, as I recall, never delivered to meet the orders from MIT and other leading-edge users. Eventually, IBM entered the time sharing environment with the Model 67 with TSS, the first set of IBM time sharing products.

Comments

by Wilhelm G. Spruth

The Time Sharing System (TSS) had major problems when it became available in 1967. It took several years to make it into a solid product.

It is interesting to look in more detail at the relationship between IBM and the universities, especially the MIT.

During and after the second world war, most computer science research occurred at universities like Iowa State University, MIT, University of Pennsylvania, University of Manchester in UK, and others. Several start ups like UNIVAC and Ferranti worked in close cooperation with universities. IBM had financed the building of the Harvard University Mark-1 computer and had contributed to the MIT Project MAC. The people at academia had made impressive contributions to the development of the nascent computer industry and considered themselves the high priests of computer science.

However, times were changing. Around 1964, more than 50% of all then active computer scientists were employed by IBM. If you wanted to contact a fellow computer scientist on some problem, chances were, he worked for IBM as well. There was less and less communications with people outside IBM.

People at academia resented bitterly not being consulted on an important development project like System /360. Moreover, they had developed a number of concepts, they felt should be incorporated in the System /360 design but were not. There were two reasons for this.

Some concepts were too advanced and ahead of their time to be introduced into mass production by 1965. The decision to use SLT instead of integrated circuits is a case in point. Not to use dynamic address translation is another one. In 1964 dynamic address translation (virtual memory) was an inmature technolog; it took IBM and GE several more years to solve the problems.

In addition, there was a broad consensus within the academic community as to the importance of a number of concepts, which Amdahl, Blaauw and Brooks decided to ignore. Direct addressing is an example. Adding injustice to injury, as time evolved, Amdahl, Blaauw and Brooks were proven right in almost all their decisions.

As a consequence, academia believed, System /360 was an obsolete design because they had not been consulted, and because it did not implement many concepts they felt were important. Thus the mainframe dinosaur myth originated, which persists in academia until today.

What should B.O.Evans done differently ? Maintaining close contacts with MIT might have helped. But maybe not. It is difficult to see what major contributions academia could have made to improve the IBM design. In the following nearly 50 years, nobody has been able to come up with an alternative superior to System /360.

End of comments

My world took another step downward in mid-1964, even as System/360 sales were skyrocketing. In the overall System/360 plan we were working on supercomputers—Models 91 and 92—to be the family's high-performance end. We were not far enough along with these systems to announce them with the Models 30-40-50-60-62-70, but we did give it some consideration. Then the situation exploded.

Dr. Harwood Kolsky was a former Los Alamos mathematician who had joined IBM to work on supercomputing. He wrote to Tom Watson Jr. to complain that Control Data Corp. had its 6600 supercomputer in production, with the 7600 model reportedly on the way. IBM, claimed Kolsky, was not competitive.

Characteristically, Tom Jr. swept aside the agenda of an executive conference taking place at Jackson Hole, Wyoming, to discuss supercomputers. The result was another new project, and a slap in the face to me. Dr. Charles DeCarlo, who had been removed in 1961 as VP-engineering in the Data Systems Division, returned to the scene as the high-performance computing "czar."

Dr. DeCarlo began his new assignment with wild thoughts. Some of the scientific computing types did not like System/360's base register-offset addressing. They demanded direct addressing—even though this came with high performance penalties—and so DeCarlo planned to develop new products in a different architecture. Soon, though, he came back to earth and concluded the best route was to get the Models 91 and 92 announced and in the hands of the sales force—and turned his energies in that direction.

I was gone from DSD by the time the high-performance systems were announced. In later lawsuits, the opponents claimed premature announcement to thwart competition. After an arduous course of litigation over several years, the courts sustained the decisions made in 1964 and 1965. It mattered little, though, since CDC had already won the day. That company's 6600 and 7600 systems soundly whipped the IBM Models 91 and 92—a defeat that plagued Tom Watson Jr. for many years.

I had other troubles even as System/360 sales were booming. Tom Jr.'s ears were still ringing from the complaints of the sales force about the inroads being made by the competition. Tom Jr. became convinced I had paid insufficient attention to evolving the current product lines as I concentrated on developing the System/360 family. He charged my fellow engineer, Charles Bashe, with the assignment to evaluate current products against the competition. Bashe had been a development manager on IBM's first vacuum tube computer designed for business applications, the 702 Tape Processing Machine, and later led IBM's work in designing check processing equipment for banking applications. He was well qualified to do the analysis Tom Jr. wanted. The results, presented in December 1964, were devastating for me.

We had indeed fallen behind the competition in some areas. Bashe concluded and reported this from his examination of speeds and price-performance of current business and scientific processors, printers, card readers and card punches, and magnetic tape products. Each of these had progressed only minimally during the years 1962-1964. The charge was true but also inevitable because IBM simply did not have the resources both to develop the System/360 family and keep the current product lines moving to the maximum. I had made the correct tradeoff of securing System/360 while only moderately extending the current product lines. However, the demanding Tom Watson Jr. wanted both and believed I should have "kicked down his door and demanded more resources to develop both the new product line and keep all the current products moving competitively." That actually was an impossibility from both fiscal and human resource standpoints, however, that is how Tom Jr.'s mind worked. Thus, when Watson reviewed Bashe's analysis, I could see my tombstone. And soon it came.

Tom Watson Jr. restructured IBM again, this time to better insure the corporation's ability to deliver the bold new Systems. He decided it was time for me to move on and he gave John Haanstra a second chance. In January 1965, a year after Haanstra had gone to IBM purgatory, Tom Jr. reorganized, unifying U.S. manufacturing within the Systems Manufacturing Division, creating a Components Division to proceed with Solid Logic Technology and semiconductors, and establishing a unified division for worldwide engineering, the Systems Development Division (SDD). John Haanstra was reborn as president of SDD, the new integrated development division. Tom Jr. surprised me by naming me President of the then struggling Federal Systems Division. Given a choice, I would rather have stayed with my Data Systems Division team to see System/360 into production and through the many software problems. However, that was not to be.

Poughkeepsie from 1961 to 1964 had been the experience of a lifetime, one I am exceedingly lucky to have lived.

Chapters 5
Federal Systems Division

Becoming president of the Federal Systems Division was a demotion for B.O.Evans.The division was loosing money and it is reported Evans got the assignment to shut it down. Instead of doing so, Evans decided to turn it around.

I was happy in the Federal Systems Division, which had made progress in my four-and-a-half years there. Division sales had increased almost three-fold and we had been profitable every year.

The US Space Program

The Federal Systems Division (FSD) had a single customer: The US Government. Among many other projects, FSD was involved in the US Space Program.

It was a privilege for any engineer to be associated with the U.S. space program. I was most fortunate to participate from the vantage point of president of IBM's Federal Systems Division. In the process I worked with terrific NASA professionals, including the famous Dr. Wernher von Braun and many of his German team.

Dr. Werhner von Braun was a most unusual professional in that he had pursued rocketry all of his life. As chief scientist for the German V-1 and V-2 programs at Peenemunde during World War II, von Braun's work nearly tipped the balance of the war. "Rescued" by the Americans just hours before the Russians, von Braun and his team of colleagues moved to Huntsville, Alabama where they led the embryonic U.S. programs in rocket development at the U.S. Army's Redstone Arsenal.

Dr. von Braun wrote extensively about the promise of space and, for a period was a feature writer in the old Saturday Evening Post magazine. Chesley Bonesteel put von Braun's thinking into art; his illustrations were a memorable feature of the magazine's series. Von Braun also made extra money lecturing on space, crisscrossing the United States in his Beech Bonanza, city by city, expounding his dream of space travel. In 1957, Dr. von Braun came to Endicott, N.Y., and delivered his lecture to perhaps 300 people in the cafeteria of the Endicott-Johnson Shoe Company. I listened, enthralled, to his description of interplanetary exploration and the distant promise of space. Little did we know that, within a few months, Sputnik would change the world forever and thrust von Braun and his theories and dreams into a new limelight.

My first meeting with the eminent Dr. von Braun was in 1966 when he and many of his Huntsville team came to FSD's Owego, N.Y., laboratory and plant to review the work we were doing on the Apollo man-to-the-moon program. IBM Owego developed and manufactured a special triple redundant computer and a data adapter. In that meeting, von Braun showed the depth of his excellence, taking the lead as the NASA group's unleashed a barrage of very detailed questions on the design and test of the equipment we were developing. The project engineer, Munroe Dickenson, was one of FSD's finest; he answered every question precisely. I left that meeting very proud of that fine engineer.

FSD played an important role in both the Gemini and Apollo projects. We produced the on-board spacecraft computer for Gemini and we managed the Real Time Computer Complex in Houston. With Apollo, FSD's role expanded considerably. We designed and manufactured the triple modular redundant computers and data adapters that were mounted in the rocket system's instrument unit, the instrument unit integration and test facility at Huntsville, the much-expanded Real Time Control Center in Houston, as well as an integration and test group at Cape Kennedy.

John Haanstra

Tom Watson Jr. restructured IBM again, establishing a unified division for worldwide engineering, the Systems Development Division (SDD). John Haanstra was named president of SDD, the new integrated development division.

When I first joined FSD, Bruce G. Oldfield was vice president of FSD's Federal Systems Center. Bruce, along with one of his key systems engineers, Joel Aron, had been on the SPREAD task force that welded the interdivisional plan for the New Product Line. Oldfield and Aron, with Henry White, Robert Crago, Allison Todd, George Gerrish, Gen. (ret) Thetus Odom and others, built this diversified business, which became approximately 4000 people all working on advanced systems projects across the FSD complex of responsibilities. Over time, however, I concluded a change was necessary at the top of the Federal Systems Center.

I never knew exactly why, however, after two years at the helm of the System Development Division, John Haanstra was removed and replaced by Charles Branscomb. Perhaps it was the old mistrust stemming from the NPL/1401S argument or perhaps there were new problems. However, I respected and understood "Big John," thus I lobbied to bring him to FSD. In mid-1966 I appointed John W. Haanstra vice president of FSD's Federal Systems Center.

When John Haanstra came to work for FSD in 1966, we both knew that at IBM "two strikes and you are out" and that his career at IBM was dead as far as advancing further in the company's mainstream. Still, John worked hard in FSD and contributed. It was only a matter of time, though, until he resigned to take a position as head of General Electric's computing operations—where he made a similar impact.

John was moving in wide circles to establish a new plan for GE's computer operations, and the "old guard" folks were having trouble with his assertive ways. My money would have been on John to be a renaissance man at GE. Unhappily, we never found out. John had purchased a Beech Baron twin-engine airplane and learned to fly. After a year at GE, John, along with his wife, June, and one of his three children, Glen, crashed in the New Mexico desert near Climes Corners, N.M. after losing an engine. All aboard perished, ending a fine career well before the real potential was realized.

The Ill Fated ACS Project

Early in 1965, as I went off to the Federal Systems Division, Tom Watson Jr. wanted IBM to have the world's most powerful supercomputer. Tom Jr. had erroneously become convinced the IBM development "bureaucracy" was too stuck on things such as compatibility and standards, and smarting from a business journal article that lauded Control Data Corporation's Seymour Cray for developing the CDC 6600 with "29 people and a janitor," Tom Watson Jr. wanted an unfettered environment for IBM to develop the world's leading supercomputer. The Advanced Computer Systems group was established in Menlo Park, CA. The first management chosen was Dr. Gene Amdahl and Max Paley who were soon joined by other well qualified professionals drawn to the excitement of building the fastest supercomputer and, ostensibly, with an open check book.

Dr. Amdahl believed that System 360's basic instruction repertoire, configuration of registers and data paths were quite satisfactory for the high performance sought, thus commenced his work on the Advance Computer System using the architecture of the recently announced System 360. However, Dr. Emanual Piore managed Research and other divisions and he listened to voices of dissent that did not like the 360 architecture, particularly the base register-offset addressing structure. Piore wanted some of his Research professionals on the supercomputer project, particularly a research professional with a mixed history of success in his research endeavors, Dr. John E. Bertram. Piore succeeded in convincing

Tom Jr. and soon a Research contingent led by Dr. Bertram and including the brilliant Dr. John Cocke, joined the Menlo Park project. Dr. Cocke came to work on his theory of optimizing compilers while Dr. Bertram had stronger conviction about the architecture and rejected the 360 approach. In short order Amdahl and Paley were swept aside and the Research contingent took over the Advanced Computer Systems project. Rebuffed, Gene Amdahl and a very bright engineer, John Earle, retreated into obscurity to pursue the 360 architecture approach. In the meantime, under Bertram's direction, the ACS project proceeded enthusiastically. At first they talked of a computer capable of executing one billion instructions per second, perhaps a hundred-fold improvement over the top of the System 360 product line. Soon, however, reality set in and the goal came down to 500 million instructions per second, then 250 Mips. Finally, in 1968 the goal became 70 Mips with some hope for doubling that performance with Cocke's new compiler. Clearly the projects' goals had tumbled precipitously and corporate interest was dimming. Their troubles were not over. In December 1968, still President of the Federal Systems Division, I was at a neighborhood Christmas party in Bethesda, Md. when I received a call from Gene Amdahl. The conversation started with: "I have the son-of-a-bitch," language not characteristic of the gentle Gene and showing how deeply resentful he was of Dr. Bertram's rejection of his ideas and who had ousted Gene from the ACS architectural leadership. In this conversation Gene related that John Earle and he, using the same technologies as that planned for ACS, and using 360 architecture, had designed a supercomputer that was plus or minus 5% of the speed of the Bertram ACS design. Moreover, the Amdahl-Earle design required only 90,000 semiconductor circuits vs the 270,000 semiconductor circuits required for the Bertram ACS design. This was profound! The news broke quickly and the System Development Division headquarters was in turmoil. Erich Bloch, the VP responsible for engineering in SDD assigned a trusted colleague to investigate. In the meantime, the Research contingent driving the ACS project, having clearly overdesigned the system, were now under the pressure of Amdahl's claim and were frantically striving to reduce the semiconductor circuit count in their design. They brought the count down to 200,000 semiconductor circuits without losing much performance, blunt testimony to overdesign. In March 1969 Bloch's man reported: his count was that Amdahl's design was 108,000, not the 90,000 Amdahl claimed, however still a remarkable difference. In short order, the Research contingent departed ACS, and at my suggestion to Vin Learson, Gene Amdahl became "CEO" of ACS and Max Paley, "COO" as I believed Gene Amdahl was a weak manager and needed experienced management under him to translate his thinking into action. Gene never agreed with my observation of him and this has been a bone of contention between us for many years. Worse, by June 1969, Max Paley and Gene Amdahl were fighting thus, in the summer of 1969, placed Max Paley in a Federal Systems Division position in the Los Angeles area. Gene continued to lead the ACS project although Frank Cary did not want to put the ACS system into production; rather he wanted to build one or a few machines so that IBM would have the fastest system in existence. This made Gene unhappy as he wanted the product to proceed to volume production and eventually led to Gene resigning again as discussed later. In net, the ACS project was a failure and once again, Watson Jr. was rebuffed in his strong desire that IBM have the most powerful supercomputer in the world.

Chapter 6
Back to the Mainstream

In September 1969, Tom Watson Jr. flattered me, asking me to return to lead the Systems Development Division (SDD), IBM's worldwide mainstream engineering operations. It was absolution from my ignoble move from the Data Systems Division four years earlier. However, I wished to decline.

I was happy in the Federal Systems Division, which had made progress in my four-and-a-half years there. Division sales had increased almost three-fold and we had been profitable every year. Moreover, I thought I could be important to "mainstream" IBM's future by steering FSD to become a trailblazer for advanced applications. We had a number of software development and complex systems programs underway, which I believed to be in both IBM's and the national interest. Also, the assemblage of executives in SDD was not my view of the leadership required—and I had no interest in executing a "bloodbath." That is what I told Tom Jr.

He rejected my polite effort to decline the post. IBM paid me a lot of money and Tom Jr. was insistent, so I had no choice. Two days later the change was announced.

Making a Critical Correction

Now at SDD, I had the opportunity to make an important correction. I had long believed we had failed by not designing into System/360 the storage management architecture MIT had wanted. Called dynamic address translation, or DAT, this capability made it much easier for the computer to manage allocation of storage without the programmer having to keep track of all the physical locations for storage of data and instructions. As multiple distant terminals and remote computers communicated with central computers in multiprogramming and multiprocessing environments, storage management would be exceedingly complex. The System/360 mistake had seemed so clear, and I was pleased to hear while at FSD that the Systems Development Division team had designed dynamic address translation capability into the 360 successor, System 370—which was to be announced imminently.

I had dinner with Chuck Branscomb, my SDD predecessor, the night before Tom Jr.'s announcement that I would become the new SDD president. While disappointed in the change, Branscomb—always a gentleman and a professional—was totally cooperative. I was astounded to learn from him that dynamic address translation had been dropped from the main System 370 plan. Instead, SDD planned to develop a separate group of processors on a later schedule to provide dynamic address translation for those leading-edge customers that demanded such capability.

The SDD plan was for 370 systems without dynamic address translation to run under an improved version of System/360's OS/MVT operating system. These were dubbed "vanilla" 370s. A more limited product line with dynamic address translation would use the Time Sharing System. TSS had been specially developed for MIT, Bell Labs, and others to operate on a version of the System/360, Model 65, which had been modified to add dynamic address translation capability. This product was named the Model 67.

The SDD plan was fatally flawed, and I knew it immediately. By leaving dynamic address translation out of the mainstream 370 product line, IBM would be denying that capability to the broad customer set on the very eve of the computer-communications age. The strategy was equally bad given that it meant pursuing a separate product line for the time sharing environment. So, my first act as SDD president was to establish in October 1969 a small task force to plan the incorporation of DAT into the base 370 product line. That task force included some of SDD's brightest engineers and programmers: Richard Case, Don Gavis, Dr. Edward Sussenguth, and Robert Ruthruaff. Their charge was to address the changes that would be made to both hardware and software, based on the assumption that software would be built on OS/MVT—the main System/360 operating system.

It was a tumultuous time. The SDD "establishment" rallied to oppose my insistence on making the change. I remember the Poughkeepsie product planners coming in with rolls of flip charts that could choke a boa constrictor, all to argue that dynamic address translation had no meaningful value. I was astounded; after all, IBM had been taught an important lesson on System/360 when MIT, Bell Labs and other leading-edge customers rejected 360 because of DAT's absence. Moreover, in the five years since that lesson, the computing world had marched inexorably forward in attaching both distant and local workstations and terminals to the computers. This had increased exponentially the complexity of memory and storage management as users would come on line, retrieve past work or create new files, and store both temporary and permanent files.

The desire was to have the computer appear to each user as if the system was dedicated to that one user, with instant response times. Absent an architecture to help solve the complex problems of storage management, it was clear there would be a major hardware and software problem in environments with tens or hundreds of users on line. Dynamic address translation was the key component of moving IBM systems into the future and fundamental for IBM to grow in the terminal/communications age.

The Poughkeepsie product planner's opposition to dynamic address translation and convoluted arguments was unbelievable. Their baloney was just so many lemmings walking en masse off the cliff. I plugged on with the DAT plan.

Even more disconcerting, a respected forecaster, Jim McDermit, briefed my boss, Spike Beitzel, using reams of subjective data and concluded that DAT would possibly mean only 3 to 5 percent more sales, hardly justification for the massive change in strategy I had launched. Indeed, DAT had to be designed into the hardware of the new 370 family, which also meant major changes in the operating system software. This would mean significant schedule slippages over the existing System 370 "vanilla" plan, with serious expense, revenue and profit consequences for the year 1970.

Meanwhile, the programming group developing the evolved operating system for the new 370 "vanilla" family simply did not believe my desired change to add dynamic address translation would ever come about. They, like the planners, dragged their collective feet, with a hundred different reasons for why adding the capability was unnecessary.

O. Scott Locken was a senior manager who had grown up in System/360 operating system programming. He was aware the current OS programming team was opposing the change. Locken had been working on a plan to develop a next-generation operating system and came to me in November 1969 with an attractive proposal to develop his proposed new operating system. Rather than modifying the older designs, Locken's operating system was tailored to dynamic address translation. It was an appealing story and—pressed by the reticence of the current operating system group to get into gear with DAT—I bought Locken's plan, launching the Advanced Operating System project. That proved to be a false start.

A month or so into that project, Don Gavis came to me to argue that I was making a serious mistake. Gavis, a former Product Test manager and one of the professionals who had been responsible for getting the 360 software stabilized and delivered, had high credibility with me. He came to IBM with no college degree; however, in the mid-1960s I had a role in selecting him for the MIT Sloan program where, if successful, one is granted an MBA. Gavis had participated in that program successfully, which gave him a new confidence. That, coupled with his particular background, made me listen to him.

Gavis made some pivotal points. While Locken's new system might be optimized to the dynamic address translation architecture, he argued, a new system could not be built close to the schedules promised. In addition, migration from the existing 360 software would undoubtedly be more complex, if doable at all. Moreover, the current operating system group were now working on how to incorporate dynamic address translation. I decided to reverse myself, held Locken's hand through the explanation, and switched from the revolutionary to the evolutionary approach. With the Poughkeepsie planner's negative views about DAT, the minuscule forecast, and now a switch in operating system plans, corporate management was in a snit. However, senior executives Vin Learson and Gilbert Jones accepted my reasoning and we set out to produce MVS.

I wish I could claim to be smart enough to have used Locken as a simple ploy to motivate the Poughkeepsie programming group. That was not the case. The result, though, was to get DAT going.

The decision was in place to integrate dynamic address translation into the System 370 family, add dynamic address translation to the mainstream operating system, and abandon the half-baked plan to have a secondary product line with a few processors evolved from the System/360 Model 67 using TSS. I then met with a large group of our most distinguished customers including Bell Labs, American Express, leading banks, and others and told them of the plan to build a unified product line. While there were suspicions and grumbling, I do not believe we lost any of those key customers by making the change in plans. In fact, I am certain we secured hundreds of new customers because the design was more attuned to the coming age of computer communications.

Initially, I had wanted to delay announcing System 370 until we could get the dynamic address translation hardware into the whole family and complete the new MVS operating system. However, for reasons related to 1970 and 1971 operating plan revenue and profits, Spike Beitzel insisted on proceeding in April 1970 to announce the two largest systems of the former plan, the 155 and the 165, together with a new disk family. These were the “vanilla” System/360 architecture, sans dynamic address translation, running the old operating systems, MFT and MVT. It was left to SDD to plan how to complete and phase in the new systems and software adding dynamic address translation. Fortunately, engineering went well and the replacement high-end systems, the 158 and 168, were available in 1972 together with the new MVS operating system. As it turned out, we probably could have waited: there was a significant recession in 1970 and IBM’s computer sales were way off plan despite the “vanilla” 370 systems.

Saving DOS

A similar thing happened at the mid-range: we first shipped “vanilla” architecture Systems 135 and 145, followed by the dynamic address translation versions, the 138 and 148. However, in this case, I made another change in the operating system.

To understand this change, we need to go back to 1965. The Disk Operating System—DOS—was a workhorse operating system for the smaller systems in the highly successful Systems 360 family. Thousands were installed and were serving the customers well. Inside IBM, however, DOS was widely viewed as a technical mess, a system that had not been developed in a process environment where it could have been made durable and evolvable. So, the Endicott Laboratory set out to develop a new low-end operating system, LEOS, intended for release in 1970 with the new Systems 370 family.

When I arrived in SDD in October 1969, plans were almost complete to announce the 370 family with LEOS. My attention was focused on changing the whole systems plan to include DAT, so I did not engage much in the LEOS project other than what it would need to accommodate the new storage management capability that I was determined to include in the 370 systems.

As the IBM organization slowly swung to the dynamic address translation plan, I began to focus more on specifics. How would we migrate thousands of DOS users to LEOS? This became a haunting question—one with no good answers. I feared a technically elegant product that would not sell because existing customers could not move easily from DOS to LEOS.

I turned to the question of just how bad the DOS code really was. Could it be made durable and evolvable? An excellent software manager, Jim Frame, came out of the woodwork to argue that DOS could indeed be saved. I had a high regard for Jim and good past experiences with him from his work on the 7070 program. This experience became pivotal in my decision to kill the newly designed LEOS and build on the existing DOS base.

If I ever made correct decisions and contributed to IBM, it was to demand and make dynamic address translation happen and to have the sense to reverse a decision I made and revert to develop what became MVS and the modernized DOS. All were great successes. I was allowed to do this because Vin Learson was still in power and he had confidence in me. With the IBM bureaucracy that had grown by the late 1970s, I am certain that no one manager could make such a dramatic change essentially single-handedly in the face of multiple adverse inputs.

Some Other Key Decisions

IBM understood that most countries had different standards, processes, and languages and that products designed for the United States would likely require changes to be acceptable. It was also well understood that other countries had fine professionals who would not leave their homelands. So, where future business justified the effort and expenditure, IBM added laboratories outside the United States—in the United Kingdom, France, Netherlands, Germany, Austria, and Sweden.

Even in 1970, the Japanese were looming as tough future competitors. In my view, it was very important for IBM to have a Japanese development laboratory that could take advantage of Japanese professionals who might otherwise not be available to the company, as well as help us understand and design for the special requirements of the Japanese market.

I went to the Corporate Management Committee seeking approval to establish a lab in Fujisawa, near Yokohama, adjacent to our manufacturing plant. The committee agreed, and I moved swiftly.

At the time, Nobu Mii, the bright young engineer who had come to us from NHK, was in rotational assignments learning more of IBM. I did not think he was ready to head the Fujisawa Laboratory. Seeing no one else from Japan in the company qualified to take the job, I assigned a U.S. engineer, Edward V. Hofler, to be the first director. Ed did a good job of organizing that laboratory, bringing in qualified talent and building the first mission. Later, Nobu Mii returned to Japan and took over the leadership. That laboratory became one of IBM's best.

Another major redirection I believed essential was to pay more attention to communications and terminals. I set about to take appropriate action. I cut the SDD headquarters staff by 40 percent to reduce the interminable bureaucracy. I brought John Fairclough from the Hursley Laboratory in England to lead the laboratory in Raleigh, North Carolina, which would be the point laboratory in terminals and communications. I sent Dr. Ed Sussenguth to Raleigh to lead the design of Systems Network Architecture, which would be our unified communications subsystem architecture and moved Jack Kuehler from the Raleigh lab—where he did not have the requisite systems experience—to the critically important and demoralized San Jose lab, where he flourished. Kuehler was charged with revitalizing the disk operations—which he did, putting in place appropriate advanced programs and turning around morale.

I had returned to SDD full of conviction that communications and terminals were IBM's future. And I knew that banking terminals had to be first on my agenda given the importance of banking to IBM's future. I thus revised the mission for the lab in Kingston, New York, to focus on banking, standard terminals, and terminal controllers. John Fairclough allied with Earl Wheeler, director of the Kingston lab, and together they forged a new world of IBM terminals and communications subsystems. John Fairclough's technical ability, vision and leadership was a principle factor in IBM's communications subsytems success.

One of the important automation areas we wanted to enter was point-of-sale terminals. Supermarkets had the most potential for this new technology. We had an extensive product line in development, including scanners as well as the hardware and software tools for inventory and resupply optimization. At the time, large supermarkets had hairline-thin profits: a percent in the black one quarter, a percent or two in the red the next. Our automation products conservatively raised profits to the range of 3 to 5 percent.

I put Jerrier Haddad—my old boss—in charge of the giant Poughkeepsie laboratory to work on the new 370 Advanced Function architecture, inserted new development programs in the Endicott Laboratory to develop a family of printers for general purpose and application-specific terminals, and evolved the other laboratories' missions to coincide with this redirection.

The laboratories in Europe came under SDD's responsibility. Following Dr. Gardiner Tucker the leadership of the Swedish, U.K., Dutch, French, Austrian and German Laboratories was assigned to Byron L. Havens who reported to me when I became President of SDD. Havens had been a luminary in early electronic computer activities. A member of the small Watson Laboratory adjacent to Columbia University, Havens had led the team that, in the early 1950s, developed the Naval Ordnance Research Calculator as mentioned earlier. Havens invented the "Havens' delay circuit" a very clever design of a storage register that could both load and transmit a binary "one" or "zero" and could also shift the binary bit left or right. That circuit was the centerpiece of the arithmetic unit in the Defense Calculator. By Havens, highly respected, became one of Vin Learson's small, powerful Group Staff in the early 1960s then, late in the 1960s, moved to Nice, France where he managed the European Laboratories.

One Special Result of 360 Standardization

Consolidation of system product lines into the unified Systems/360 had allowed IBM to apply the systems in new directions some of which we could not clearly foresee when the series was designed in the early 1960s. One application area we did anticipate was the communications environment. In 1962 and 1963 we were not tremendously clear on the specifics of that application thus the 360's initial design tried to implant the fundamental "hooks" for the communications hardware and software to come.

The forecast for communications attachments was strong; IBM estimated that as many as one-third of the Models 40 and larger would have remote terminal/communications applications by 1970. This turned out to be far off the mark: that percentage was reached in 1968! By 1970, almost all Model 40 and larger systems were used in communication applications, far beyond the 1964 forecasts. The "teleprocessing" era was well underway.

IBM had been saturated with the task of creating the basic 360 processor family, and did not have a consistent communications subsystem architecture nor sufficient development of communications control programming. Users, impatient to expand their operations, moved on their own into communications applications with communications control programming in bits and pieces done by their own programmers, by local IBM systems engineers, or by IBM laboratories. The result was a period in the late 1960s when there were many dissimilar communications control programs, each serving some particular purpose but none welded into a master communication subsystem plan.

This rampant growth of communications applications continued throughout the 1970s. Users connected computers into networks and attached distant terminals by the hundreds. Moreover, rapid growth of minicomputers gave rise to attached remote terminals used in many applications across the spectrum of what is called batch processing (payroll, accounts receivable, inventory control, sales statistics, etc.); interactive processing (airline reservations, on-line banking terminals, retail point of sales applications, etc.); and real time processing (such as refinery control). A typical application most of the public encounters is the powerful American Airlines SABRE system, which operates across much of the world and allows a ticket agent anywhere to make flight reservations, secure seat assignments, reserve special meals, etc—all in real time, so that, for example, if a seat is sold in Tokyo it is instantly taken out of the inventory of seats available worldwide.

These demanding new applications of computers inexorably required that large computer applications have immediate access to databases accumulating in remote minicomputer sites, or mini-computers needed access to central databases as well as remote databases in distant minicomputers. It is no wonder communications applications quickly became a large and fast-growing segment of the computer industry.

IBM responded to the need for standardization of the communication subsystem in 1972 with its Systems Network Architecture. In 1970 I had sent Dr. Edward Sussenguth to the Raleigh laboratory to lead the development of a unified communications subsystem architecture, and with Systems Network Architecture he succeeded. SNA eventually had tens of thousands of installations across the world.

Steadily extended over time, SNA by 1980 had substantial capabilities in terms of networking, and a wide variety of corporate networks using central processors, remote minicomputers, and distant terminals employed the SNA architecture. In effect, the standardization of processor architecture freed development resources and allowed IBM to develop new communications hardware and software as well as substantially improve magnetic disk and magnetic tape storage subsystems. Indeed, I believe this "standardization" within IBM gave impetus to new subsystem unification that now proliferates across other manufacturer's equipment as well as other IBM product lines.

A Strategic Blunder

It was long believed that very low-priced computers would produce enormous volumes. The concept of stored program computing for a million small businesses stirred visions of hundreds of thousands of systems sold, a dream the personal computer has today made a fact. Over the years IBM pursued "small business computers" in dozens of approaches. Most never saw the light of day, however, and the seeming inability to find the solution increasingly disturbed IBM's senior management.

When we were developing System/360 we aimed at an entry-level system that would rent for perhaps $1,500 per month and sell for $50,000. When costs were finalized and prices set, however, we missed that entry price goal by a factor of two. Later there were designs—specifically Models 25, 22, and 20—developed to take the 360 systems entry-level to lower price levels, but the price reductions ended up being small. We failed to achieve the utopian breakthrough. By 1969, President Frank Cary had run out of patience and decided to create a new division charged with going after that mass market promised by very low-priced systems. In effect, he was searching for the personal computer.

Had I been in the Data Processing Group at the time, I would have opposed this approach: it threatened IBM's systems cohesiveness. However, by the time I returned to head the Systems Development Division in October, the new division was a done deal. The General Systems Division was established under the leadership of Jack Rogers, a fine executive with sales experience. The division's mission was small business computers, as well as an IBM response to the very successful PDP family developed by Digital Equipment Corporation (DEC).

Frank Cary was a bright, confident leader with a record of accomplishment in each position he held. But his GSD decision turned out to be a bad one. The mistake was in his reasoning the compatible System/360 was right for its time, but that the need for standardized products was passing. That was wildly divergent from the compatibility direction set by System/360 and 370. Cary believed new design tools, more powerful software, common interfaces, standard components and mass production capabilities should now let IBM customize our products to user's needs independent of the system's architecture.

While he was reaching for the "open systems" concept, a correct intuition, his charter for GSD failed to establish the entry price goal to have GSD move down into the untapped mass market. Rather, GSD bolted off to develop product after product, each one basically competing with the low-end and mid-range of the mainstream products. They missed the very low-priced market completely.

Over the decade of the 1970s GSD repeated the errors of earlier times by producing an assortment of products that were largely incompatible with each other, unable to communicate with each other, incompatible with the mainstream product line, and each requiring unique software. These products included the System 32, 34, 36, and 38 as well as the Series 1, the long-sought binary computer to compete with DEC's highly successful PDP series and its successors. It is true that some of the GSD systems became a significant business—such as the System 36, which grew into the popular and profitable AS 400. However, from 1970 to 1980 IBM's share of the mid-range market in the United States dropped from approximately 65 percent to less than 20 percent as Digital, Data General, Hewlett Packard, and other companies each produced families of generally compatible systems. IBM's prior position as mid-range leader was devastated by the GSD strategy: not only had the mid-range systems been IBM's bread and butter, these systems built the huge customer base from which grew the large IBM mainframe market.

In my view, GSD completely missed the worthy mission Frank Cary intended. In the process, the division needlessly tore asunder IBM's cohesive systems products structure. While some people who were in GSD at the time will shout in protest of my evaluation, the loss of market share is unassailable—and stands as stark testimony to the error of the strategy.

Career Apogee

I reached the apogee of my IBM career in 1972 when, on top of being a division president, I was named a corporate vice president. I was very proud, not realizing then that my time on the peak of Olympus was to be short-lived.

Out of Line Management

In March 1972, the new IBM President John Opel changed the mainstream organization radically. In May 1973, there was a mini-reorganization wherein systems architecture and much of systems software was taken from SDD and split between the three divisions, GPD, DSD and my division, renamed the Systems Communications Division (SCD).

In March 1977, I was removed me from the presidency of the Systems Communications Division. I was put out to pasture in a sort of face-saving job in charge of the corporate staff for Engineering, Programming and Technology (EP&T)—an embarrassing and traumatic event that affected me deeply. I wanted the authority and accountability of line management, and I disliked the frustrations of staff positions.

Management specialists usually argue that a well-structured organization is impervious to the loss of a single individual. I believe, however, that IBM's road downhill in the late 1980s and early 1990s actually began way back in 1970, when Tom Watson Jr. suffered a heart attack and quickly transferred the reins of IBM's leadership to others.

The first was the able T. Vincent Learson. He was an excellent "Mr. Inside." He knew the business, had all the correct instincts, and was a driving leader who knew how to delegate. He was both feared and respected, and had a basic integrity. His chairmanship was short lived under the policy set by Tom Jr. of mandatory retirement at age sixty for corporate officers. When Vin Learson had to retire there was no question that his replacement would be: Frank T. Cary, president under Chairman Learson.

Frank had grown in sales ranks, most always the prerequisite for IBM's top posts. He managed with distinction IBM's vaunted sales arm, the Data Processing Division, and became Group Executive over all data processing activities, where again he did well and was elevated to the presidency. When Vin stepped aside in 1973 at age 60, Frank became chairman.

Frank Cary is an unusually intelligent, sophisticated, reserved, and a quiet leader. He left indelible marks on IBM during his watch. For instance, he was personally responsible for IBM's entry into the personal computer business. I thought it uncanny that, with legions of planners across the corporation, no one—myself included, even with a prod from MIT's Michael Dertozous—came forward demanding the PC product. Rather, Frank sensed the importance and personally drove the PC project. He set up and protected a small task group under Bill Lowe charged to get IBM a PC product as quickly as possible. And, under Lowe's team's work and Frank's support—and only thirteen months from commissioning the study—IBM was shipping its first personal computers. They turned into astonishing volumes, beyond anyone's expectations.

On June 1, 1984, I retired from IBM.

The following part of Bob O Evans memoirs deal with events following his reassignment to lead the corporate staff for Engineering, Programming and Technology and later his retirement from IBM.

Bob O. Evans dated his memoirs May 1999. He died at the age of 77 in Hillsborough, Calif., on September 2, 2004.

Epilogue
by Wilhelm G. Spruth

The S/360 Secret

The development of the S/360 is an unbelievable event. It is something that just should not have happened.

Information technology has been developing at an extremely fast pace. There are very few developments which were not obsoleted by newer developments after a few years. Thus it is understandable that anything older than a few years is questioned as to its obsolescence. Looking at todays information technology status, the landscape has changed beyond recognition since 1964.

With one exception: the S/360 architecture. The S/360 architecture was introduced in 1964. Nobody has been able to improve the basic design during the following nearly 50 years! Many people have tried multiple times to develop an improved hardware architecture, claiming significant performance advantages. No development has ever succeeded to show a noticable improvement over the S/360 architecture, and many (e.g. Burroughs B5000, VAX) have been a disappointment.

We did an exercise with a group of experts a few years ago. We decided to travel back in time to 1963 and submit to Amdahl, Blaauw, and Brooks a proposal with a list of items, which – based on the knowledge available in 2008 – would improve the S/360 design. We compiled the list: it was surprisingly short and included no major items. Some examples were: Do implement 32 bit addressing immediately, eliminate the EDIT instruction, do not implement imprecise interrupts (on the model 91)….

I had a chance to question Gerry Blaauw on the 24 Bit addressing decision. He answered: " I had an uneasy feeling on the 24 bit addressing issue, but I was too tired to start another fight, and I knew, we could correct it later if needed. And nobody in 1963 could imagine a computer with more than 16 MByte of main memory".

I sometimes wonder, what would habe happened, if Bob O. Evans had been given the chance to continue with IBM in an exetuvive position until his natural retirement around the late 80s. Would Evans have announced the IBM PC in 1980 with a 32bit S/360 subset architecture, avoiding industry problems like little endian/big endian, ASCII/EBCDIC, code pages, buffer overflows, security patches, virtualization problems, I/O drivers…….

We will never know. Still, todays mainframes implement the most modern, leading edge architecture. For details, see http://tobias-lib.uni-tuebingen.de/frontdoor.php?source_opus=4710 .

An independet view on the S/360 development is availabe in 2 articles published in 1966:

T. A. Wise:"I.B.M.'s $ 5,000,000,000 Gamble". FORTUNE Sept. 1966, p.118
T. A. Wise: "The rocky Road to the Marketplace". FORTUNE Oct. 1966, p 138

Boeblingen, June 2010, Wilhelm G. Spruth