Project Acorn: The Birth of the IBM Personal Computer

By David N. Saul - Last edited January 11, 2021

Introduction

Personal Computer. The generic term has become a common part of our vocabulary. Toddlers to seniors understand what it means. But if you go back forty years, the concept of single-user computers was reserved for electronics hobbyists, academics, and popular science speculation.

What I attempt to do in this set of historical remembrances is to take you back to 1981 and the introduction of the IBM Personal Computer. Personal computing was inevitable given the progress being made in microprocessors, solid-state memory, and software. But IBM gave legitimacy to the PC that hastened its acceptance. I played a small part in IBM's development of the first PC. I want to thank our classmate, Bob Gray, for providing his personal perspective on personal computing in academia in the late 1970s and his specific insights on CP/M, the most serious operating system for early microprocessors.

A bit of background is to understand how I was in the right place at the right time. Specifically, why was I a Project Manager in IBM's Cambridge Scientific Center as 1981 began and how did the Center become part of Project Acorn, the initial code name for the IBM Personal Computer?

I began work at IBM in 1965 as a Systems Engineer in its office just outside Harvard Square that supported federal government customers. As described in my earlier essay in our first Class compendium, I worked with the NASA Center in Technology Square on the Apollo ground guidance computer, MIT's Instrumentation Lab (now Draper), Air Force Cambridge Research Labs at Hansom Field and MITRE. While supporting MITRE I had my first experience programming IBM's recently announced System/360 computer. Prior to that announcement, IBM had a separate 36-bit word-based scientific computer series starting with the 701 and evolving to the 7094. For example, MIT had an IBM 7094 in Building 26 at the time of our graduation. MITRE had one of a very few Stretch computers which was a supercomputer version of the 7000 series. Since computers of the time did not have separate input/output processors they were slowed to a crawl when they had to read input, typically punched cards, or print to paper. Stretch speeded up I/O by spooling its input to magnetic tape and sending its output to magnetic tape. An IBM 1401 character-based business computer provided offline card input to tape as well as tape to paper printing. MITRE replaced their 1401 with one of the first System/360 Model 30 computers in the Boston area. Rather than writing a completely new spooler program I took advantage of the 1400 series emulation that came with System/360 and I only had to provide a shell in 360 Assembler which called the emulator with the correct parameters. Fifteen years later the IBM Asynchronous Communications program for the IBM Personal Computer used a similar architecture.

I moved from that position to a District systems-support role at IBM on Boylston Street in Boston. The District office provided second-level software assistance to customers in New England, upstate New York, and most of Connecticut. Our job was aptly described by a fellow MIT alumni co-worker, George McQuilken, as "birdmen". We flew in to the customer site, dropped a load of s**t, and flew out. I became a regular on Mohawk airlines for Buffalo and Rochester customers. I also put many miles on my car driving to southern Connecticut, except when the Arab oil embargo forced me to use the train.

In 1977 I got my first management job at IBM's Cambridge Scientific Center in Technology Square. CSC had two non-contiguous floors at 545 Technology Square. Most of the remainder of the building was occupied by MIT's Laboratory for Computer Science (LCS). Cambridge was one of three domestic Scientific Centers and others globally that engaged in advanced technology projects. We were the interface between research and development and were deliberately located adjacent to leading computer science universities. Cambridge is best known for its invention of the Virtual Machine, the extension of virtualization to an entire environment. In conjunction with the Palo Alto Scientific Center, Cambridge had added virtual memory hardware to an IBM System/360 Model 40, creating CP40. Later that prototype was productized into CP/67 and its operating system, CMS (Cambridge Monitor System). CP/67 was fortunate for IBM as the core virtual memory software, TSS (Time Sharing System) was mired in delays and performance issues. When the System/370 was announced VM/370 became a standard product and moved to a separate site in Burlington, MA. A historical footnote is that several of the CMS developers were recruited by DEC and you can see their hand in the development of VMS.

The projects I managed at the Cambridge Scientific Center when I arrived were mainframe oriented, which was the core of IBM's business. The minicomputer was just starting to eat away at the lower end of IBM's world with DEC's PDP being the most prominent inroad. Others had also cloned the IBM mainframe with former IBM architect, Gene Amdahl, building a high-end machine. I also worked on early networking projects which predated the Internet. Networking projects needed partners to communicate with and universities were the ideal collaborators. I managed the BITNET networking project with Grey Freeman at the Yale Computer Center and Ira Fuchs at CUNY.

One of my mentors in Cambridge was resident IBM Fellow, Nat Rochester. Nat had been one of the designers of IBM's first computers, the 701. In addition to telling us anecdotes about those early days, Nat could always be relied on for practical advice on

a project. He was also a bit eccentric. Nat liked to drive with no shoes on. He said it gave him better feel for the pedals. I remember the day Nat forgot to bring his shoes with him and spent the entire day walking barefoot around the Center. Nat died in 2001 and his memorial service in Duxbury attracted a crowd that overwhelmed traffic in that small seaside town with so many friends and colleagues wanting to pay tribute to him.

Another project I worked on turned out to be one that paved my path to the PC. When IBM introduced System 360 it also introduced the 8-bit byte and its encoding, EBCDIC. The rest of world continued, for the most part, with ASCII coding for characters. IBM did not compete at the low-end of the terminal market where Teletypes set the standard. Others were building CRT screen terminals that used the ASCII standard. IBM Japan needed a product and had developed the 3101 CRT terminal. To validate the 3101 in the United States an IBM location was needed with ASCII expertise and Cambridge was picked based on our university networking experience. We were able to validate the 3101 and it went on to some success, primarily outside of the US.

Personal Computers Begin to Emerge

Personal computers were starting to emerge at the time. Radio Shack's TRS-80, Commodore 64 and the Apple II computer were most popular with hobbyists. They were crude by today's standards with cassette tape for storage and upper/lower case keyboards optional. But they had also started to gain traction with academics who began to use them for scientific calculations, avoiding the delays from mainframe batch computers. Classmate Bob Gray was one of those pioneers. Here is what Bob has to say about those early days:

"In 1978 I purchased a Cromemco System 3, a Z80-based machine which ran CDOS (Cromemco DOS, one of the many CP/M imitators of the time). My home office was in a tiny rural town (La Honda, perhaps 500 people then, in unincorporated San Mateo County). I needed real computing power, including a decent word processor (Emacs and the WYSIWYG WordStar ran on CP/M and its imitators) as well as Basic and Fortran 4 (also great on CP/M). Modems then were too slow for me to work remotely from home (and the networks of the time too unreliable), so I needed a self-contained system, often exchanged data with the Stanford computers using 8-inch floppy disks and used the modem only for email. That machine cost me nearly $4K and weighed over 70 lbs. But it allowed me to do serious Fortran programming and then pass on the successes to graduate students who redid things in C or C++."

"By May 1978 I submitted a paper (published in January 1980) with two students to an IEEE Transactions with examples and preliminary results developed at home in Basic on the Cromemco System 3 (acknowledged in the paper) with careful expanded simulations obtained on a PDP 11/34 at Stanford. Personal/home computers running CP/M or an imitation were far from being for hobbyists alone. I and several colleagues were actively using them for statistical signal processing research by 1978. Another virtue of CP/M was that the early versions of LaTeX typesetting applications (combining Knuth's Tex with Lamport's macros) were quickly available for CP/M and hence for its imitators. I used the System 3 extensively until the early 1990s when I replaced it with a SUN Unix system, which Stanford paid for. I never owned a PC running MSDOS."

Recognizing the business potential, IBM started a project in its Raleigh, North Carolina plant to build an Apple II clone, called the 3101P. It used the 3101 keyboard and screen with a modified chassis to hold adapter cards. No software would have to be written for the new product as it could reuse anything written for Apple. Given our work with the 3101, several of us flew down to Raleigh and met with the 3101P developers. We asked them what distinguished their product from Apple, and they were hard pressed to come up an answer other than that their chassis was easier to insert adapter cards into and they had a better keyboard. Assuming the price would have been higher; we did not see the 3101P as a winner. What was the alternative?

TRS -80

Apple II

Every year our Cambridge Scientific Center head, Dick MacKinnon, would hold a kickoff meeting in January to celebrate the accomplishments of the previous year and brainstorm about the year ahead. In January 1981 we were offsite at a local hotel. One of the stalwarts of our Center was Fritz Giesin. Fritz was a former IBM Customer Engineer with experience maintaining hardware. More than that, Fritz was the person you went to if you needed something built. He was the ultimate tinkerer. For his 1980 Christmas holiday, Fritz had taken his family to Boca Raton, Florida. Never one to relax, Fritz had stopped by the IBM plant to visit with old friends. While there, he heard about Project Acorn, a "skunk works" project to build a personal computer. Fritz described to us what he knew about Acorn and asked whether the Center wanted to get involved. Of course, we did. Our experience with the 3101 was our entry. It turned out that being in Cambridge gave us another entry that we did not realize at the time.

Project Acorn

The IBM Personal Computer would never have existed without the creativity, salesmanship, drive, and persistence of one man – Don Estridge. He was the force behind Project Acorn and convinced IBM's then chairman, Frank Cary, that IBM needed to innovate rather than copy. I only met Don once and it was after the work I describe below. The three US Scientific Centers in Cambridge, Palo Alto and Los Angeles would gather several times a year at an IBM laboratory or development site for a symposium. In late 1981 we went to Toronto where, upon introduction, Don Estridge commented that he had already seen me on the IBM PC announcement video. At that meeting I sensed the force of his personality and why he earned the title of "Father of the IBM Personal Computer". IBM's future PC struggles might have been lessened if Don Estridge had not been tragically killed in a private airplane crash in 1985. Anecdotally, going through Canadian customs and immigration, Nat Rochester had his floppy disks confiscated as undeclared business goods. The disks contained Nat's presentation slides, so he had to improvise during the meeting.

Acorn was based on an Intel 8088 chip rather than the Motorola 6800 that Apple used. We were told by the Acorn architects that they chose to use the 8088 chips, a variant of Intel's 8080 with a smaller data path, primarily due to cost. That choice meant that new software would have to be written from scratch – operating system, programming languages, applications, etc. Developing the PC hardware could borrow from other parts of IBM like keyboards and monitors, but internally developed software had never been a historical IBM strength. The solution was to go outside for everything. One of the first software contracts was with a Washington state startup, Microsoft, for a BASIC interpreter which would be the primary programming language. Other contracts were let with outside companies for accounting software from Peachtree and a word processor, WordStar. The Center saw an opportunity for us to contribute communications software, which was not part of the initial plan.

Software

We assumed that the first choice for an operating system would be CP/M, the first successful operating system for microprocessors, invented by Gary Kildall and sold by his company, Digital Research of Pacific Grove, California. We were told that the IBM negotiators were not able to reach a satisfactory deal for CP/M with DRI during their initial visit, so IBM turned again to Microsoft with a request for a suitable operating system. Microsoft then purchased a CP/M-like operating system from Seattle Computer Products and transformed it into MSDOS, which they provided to IBM. Our software development was not dependent on the choice of operating system. Following an agreement with Digital Research, early IBM PCs offered customers a choice of which of the two operating systems they wished to purchase, with MSDOS being one-sixth the price of CP/M.

The "killer application" for all PCs turned out to be the spreadsheet. Another startup in Cambridge, Software Arts, was developing VisiCalc. Originally developed for Apple by Dan Bricklin and Bob Frankston, VisiCalc was the prototype for all of today's spreadsheet products that have transformed data manipulation. Later at Lotus (acquired by IBM), Bob developed Lotus Express before moving to Microsoft. After many years I reconnected with Bob Frankston at the Boston CTO Club. He remains an active contributor to the IEEE and a seminal thinker on multiple topics.

The original IBM Personal Computer had hardware communications adapter cards for both binary synchronous communications and asynchronous communications but no plans for software to drive them. That provided an opening for a Scientific Center project. We volunteered to write asynchronous communications support that would allow the IBM PC to operate over telephone lines to an IBM mainframe, minicomputer, etc. And we committed to finishing in time to be included for the mid-year product announcement. The Acorn software manager, H. L. "Sparky" Sparks, was dubious but agreed. He did not think we would make it. We may not have but the announcement slipping until late summer gave us more time.

We needed a prototype IBM PC to test with and they were in truly short supply. One could not be procured for us. The solution was due to our proximity to Software Arts in Central Square. Shipments to their shared space were not secure. The solution was when the latest version of PC hardware was produced it was sent to our secure IBM location. Fritz would then hand deliver it to Software Arts and take the previous hardware back to our Technology Square facility to use for software development. The oldest machine would be sent back to Boca. The PC prototype was locked in a windowless room with two tables and a few chairs. Only I and the two developers had keys. My hardware contribution was to neatly print characters on stickers which were placed on the keyboard. The prototype keyboard had come from an IBM word processor and was EBCIC instead of ASCII. Never underestimate the importance of taking an engineering drawing course.

Networking

Modems were not yet available to connect the asynchronous adapter to a phone line, so we used an acoustic coupler operating at 300 bits per second. Later 2400 bps modems seemed lightning fast to us.

Asynchronous Communications Support was developed by Frank Bequaert and Charlie Salisbury (MIT Class of 1964). The software emulated an ASCII terminal. To speed development, Frank wrote most of the software using the native BASIC interpreter and Charlie wrote an Assembler driver for the adapter card, being called and interrupting when needed. They had to limit features as the Boca Raton management insisted that the program fit in the default 64K memory. We did produce a Version 2 for 128K with those additional features before the end of the year.

At some point the Acorn code name was dropped and Chess became the new label. There was also a Checkers project which became the ill-fated PC Jr. Plans were being made for selling the PC in IBM stores as well as Sears, distribution channels entirely new to IBM. One of the stores opened, briefly, across from the Burlington Mall. After much discussion it was decided that the PC would also be sold by IBM's Data

Processing Division which normally marketed mainframes et al. DPD VP Dave Hanna was an early supporter of integrating the PC with IBM's other products.

I became a regular traveler to Boca Raton, Florida to meet with the Acorn software management. The Acorn group was housed in rental space away from the main IBM development site which helped to maintain its confidentiality. The building abutted the railway tracks, and, due to the noise, we would have to pause meetings with Sparky Sparks, in his single-pane windowed conference room whenever a freight train would pass by. He later left IBM for newly founded Compaq as vice president of sales and service.

PC Announcement Day

As PC announcement day neared, we were finishing up our testing and documentation. The Data Processing Division realized they needed some announcement material to educate their sales force. A professional video would be produced with DPD President, George Conrades, taking the lead role. Earlier I had demonstrated our Asynchronous Communications Support to him on a visit to Cambridge. We had connected to the New York Stock Exchange over a dialup line and were receiving stock prices, looking like a Teletype. I had to convince Conrades that we were doing this in real time and had not pre-recorded the data feed. When the time came to record the announcement video I was selected to act as the PC user. We filmed in our locked room with me sitting at the PC and describing the uses to which it could be put. I had a brief talk to memorize. Mostly I remember having them apply stage makeup so I would appear human on the video. Thankfully, I bore no resemblance to Charlie Chaplin who the IBM advertising agency chose as its silent campaign spokesman. Years later I met George Conrades when he was CEO of BBN in Cambridge and reminded him of our shared video past. He did not seem to remember me.

Shortly after the PC announcement, IBM had one of its quarterly mainframe users' meetings, SHARE, at the Palmer House hotel in Chicago. Headquarters DP marketing people came up with the idea of shipping our prototype PC to the hotel and have me give a PC introduction "birds of a feather" session to a small group. We were added to the schedule but the day before the session our prototype was seized by security guards sent to our hotel room by IBM. Our development PC was not Underwriters Laboratory certified and could not be shown in public. We were told a certified model would arrive that night. After midnight I drove to O'Hare Airport and I met an engineer from Boca Raton. We took the boxes off the luggage conveyor, piled them into the rental car and drove back to the hotel. We spent the rest of the night assembling the two PCs, under his guidance. It was the first time I had ever inserted memory chips into a circuit board. Fortunately, I did not bend any of the pins. I also had to install the 5.5-inch floppy disk drive in its bay which required cantilevering the PC over a table and inserting a screw from underneath. An early engineering change with the production model PC changed that procedure to a safer one, inserting screws from the side.

The next morning, I set the PC up on a table in front of a small hotel meeting room. When we opened the door the entire hallway outside was filled with people who wanted to get in. I had to repeat my presentation on overhead transparencies (no PowerPoint yet) to larger and larger meeting rooms for the entire day. The only software I had to demonstrate was a program that drew random colored lines on a TV monitor and a BASIC program that calculated payments, given amount, term and simple interest. IBM did not yet have its own color monitor available. The SHARE membership presentations were followed by a late-night trip to Lake Geneva, Wisconsin, to present to the SHARE Board of Directors after their dinner. After a heavy meal and much wine, the Board did not show as much interest as their members had.

On August 12, 1981 IBM announced the model 5150, IBM Personal Computer. The only piece of IBM internally developed software included in that announcement was Asynchronous Communications Support. I was and continue to be proud of the work that our team did in 1981 to make that happen. Like many of the accomplishments of the Cambridge Scientific Center during that period, our Personal Computer project would not have happened without the leadership of Dick MacKinnon and his ability to see how technology could be applied to relevant problems.

With the announcement IBM had made provisions for an employee Personal Computer purchase program and our family received its first home computer later that year. It had two 5.5-inch floppy drives, a monochrome monitor, and an Epson dot matrix printer. Disk hard drives did not come along until the XT model in early 1983. Our first role-playing game was Zork, entirely text based. We soon purchased enhancements from a start-up mail order company in New Hampshire, PC Connection. They are still in business and an early example of the huge industry that has prospered from the personal computer.

Building on this first terminal communications program from Cambridge, in 1982 staff member Jim Perchik developed a full featured terminal communications program, 3101 Emulation Program, with full-screen capabilities as well as the ability to emulate other manufacturers Teletype compatible devices. 3101 Emulation had the distinction of having zero bugs reported during its two-year lifespan.

As a postscript, in 1992 IBM closed all three remaining domestic Scientific Centers as a cost cutting measure. That closure ended their twenty-eight years of advanced technology contributions to computer science. Ironically, in 1989 the IBM Systems Journal (Volume 28; Issue 4) had published "History and Contributions of the IBM Scientific Centers", an entire issue dedicated to the innovative work done in their twenty-five-year history. I was one of the guest editors on that issue.

David Saul '64, his first job after graduating from MIT in 1965 was with IBM's Federal Systems office in Cambridge. His first management position was with IBM's Cambridge Scientific Center, where he oversaw advanced technology computer projects in virtualization, multiprocessing and networking. In 1981, his group developed IBM's first software product that allowed the IBM

Personal Computer to communicate over telephone lines to mainframes. In 1993, he joined State Street and applied his computer skills to financial services. He was State Street's Chief Scientist, where he focused on applying innovation to their business.

Volume 2 - Then, Now, and Beyond