Biographies

J.A.N. LEE, EDITOR

In addition to biographies, this department publishes autobiographies, memoirs, and brief and extended obituaries. Each is intended to celebrate a life. They should be discursive, but concise, learned, and anecdotal, touching and informative, and occasionally imbued with compassionate humor. Naturally, they should include a thoughtful summation of the subject's life and virtue, as well as—where applicable—the cause of death.

We want these works to provide a mosaic of the social history of computing that will prove valuable to future scholars. Thus, subjects are not limited to the famous, but may commemorate the activities of lesser known participants in computing history. Indeed, the obituaries in particular are not limited to people, but may record the demise of organizations, computers, software, and significant concepts.

Readers are invited to submit additional comments concerning already published material, to write signed items for publication, and to suggest subjects that should be memorialized.

Unsigned items are the work of the editor.

IBM Field Engineering Experiences: A Personal Memoir by Glenn E. Meyers

In the winter of 1944/1945, I was a senior in high school contemplating my future in one of the armed forces. Several classmates had volunteered and were already in uniform. One of the things that influenced my decision was the fact that my brother had recently become a prisoner of war in Germany—he had been serving in the infantry. I tried joining the Navy, but my eyes were just as bad then as now, and I was told that only if I knew radar would the Navy take me. Not knowing what radar was, I was disappointed and resigned to going into the Army. One day shortly after my Navy turndown, I was listening to a broadcast advertising for radio operators in the merchant marine. I decided to explore that possibility and purchased two books: the textbook *Radio Engineering* by Terman and the *FCC Question and Answer* publication. The FCC publication had every possible question and the correct answers on all six elements involved in either radio broadcast or radiotelegraph broadcast licensing. Applicants also needed to know Morse code at 13 words per minute in order to become a shipboard operator.

There was a limited version of the radiotelegraph license available that tested on only two of the four elements needed for licensing to serve as the second-class or third-class radio operator on a ship. Two operators were normally carried on a cargo ship during World War II. I passed the test for the limited license in June 1945 and signed on the ship *Negley D. Cochran*, a liberty type with call sign KSSX. We sailed from Brooklyn, New York, on 16 August 1945, one day after the end of the war. (I was one of the crowd pictured in Times Square on 15 August 1945. It was a joyous celebration.)

I later upgraded to second-class and then first-class radiotelegraph licenses and continued to sail as a "sparky" on the SS *John A. Donald*, call sign KXDV, until November 1948, when I came ashore and accepted a position with RCA Service Co. as a TV technician in Pittsburgh, Pennsylvania.

In the spring of 1956, I attended an IEEE meeting in Memphis, Tennessee, having been invited by the lead technician at the RCA Service Co. branch office. I was the branch manager, having been through the technical ranks at RCA Service Co., and service manager at the Houston and New Orleans branches. At this IEEE meeting, an IBM engineer was diagramming the IBM 705 computer. My thoughts were that this certainly seemed like an upcoming field. During a trip to Pittsburgh later that year to visit relatives, I decided to inquire about employment at the IBM branch office. At that time, I was just turning 29 years of age and was frankly told that "It's a good thing you applied now, as next year we would not have looked at you." True story.

I started with IBM as a field engineer in August 1956 with the prospect of servicing the 704 computer that was to be delivered to Westinghouse East Pittsburgh and another to Westinghouse Bettis. One 704 was already installed in Monroeville, Pennsylvania, and there were prospects for another at Gulf Research, nearby at Harmarville, Pennsylvania.

After several weeks of toting another field engineer's tool bag, I was sent to Poughkeepsie, New York, for six months of intensive training on everything from the 026 keypunch and the 407 accounting machine to the 704 computer system. Most customers for the 704 had those card machines, and several of the card input and line printer devices for the 704 were derived from the normal

1058-6180/99/$10.00 © 1999 IEEE

IBM accounting machine lines. For instance, the 711 card readers and 716 line printers had most of their elements taken from the earlier accounting machines, such as the 402 and 407.

I was taught the 704 hardware and assembler language programming, since the diagnostics for the 704 were written in assembly language. I can clearly remember wondering what was going on for several days of programming class until the lightbulbs began to turn on and a TIX loop (transfer and index) became understandable.

The IBM 704 was built of pluggable units that would normally hold eight dual-triode nine-pin tubes, most of them were 5965 or 5687 type. The system units were connected by cabling (normally) run under a raised floor. This raised floor was an air-conditioning duct with openings under each unit. Large blowers were utilized in the processor and auxiliary units to dissipate the heat from all the heat-generating tube filaments. Large power tubes were needed to supply plate voltages to those hundreds of tubes. Walking into a 704 system room was an introduction to loud blower noises with air flowing furiously into the units and out the tops. It was difficult to hear one another talk when shooting a bug near the mainframe. Our training at Poughkeepsie included work during the early morning, shooting applied bugs on the assembly line. It was always quite cold there, which helped keep me alert at six in the morning.

After six months, I was based at Westinghouse East Pittsburgh, home of LRA (Large Rotating Apparatus, or huge generators that shipped in sections on flatbed railroad cars) and then later at Westinghouse Bettis, where development of nuclear submarine engines was under way. Also, I worked at Gulf Research, where a 704 system was installed to do seismic exploration analysis and refinery run optimization.

During my first 10 years with IBM, I spent fully three solid years at various IBM classes in Poughkeepsie, Kingston, and Endicott (all in New York) and in Rochester, Minnesota, training on the 704, 709, 7040/44, 7090, 7094, 7094II, 1401, 360/30, 360/65, 360/75; the OS/360; the 360/7090 emulator program; and various peripheral devices such as the 714, 716, 717, 720, 1403, and 1442 printers; the 727, 729, and 2400 series tape drives; and the 1301, 2311, and 2314 disk systems. Coupled with assistance trips locally and out of town, I estimate I spent half my time away from home. There was also a good deal of overtime.

My very first trip in a jet was on a trip from Pittsburgh to Atlanta aboard the French Caravelle. It was a fantastic journey. Previously, it took quite a bit longer on a Viscount prop plane. The next leg of the trip, from Atlanta to Huntsville, Alabama, was on a DC3. The DC3 was flying against a strong head wind, prompting the IBM staffer meeting me to doubt it would ever land. I frequently got calls from the Chicago regional office to hop the next plane to somewhere. I was even rousted out of bed at 2:30 am to catch the first plane to somewhere. I sometimes silently wished there would be no empty seats on the plane. There were always seats available.

In 1961, I was promoted to field engineering technical specialist and in 1965 to advisory technical specialist (see Fig. 1). I was a district specialist on the 704/7094 and regional specialist on the 360/75.

I can still recall details of the various systems, but with no reference manuals nearby, some of my memory is more or less erased.

The 704 system consisted entirely of discrete components. One could see and touch each resistor, capacitor, tube, AND/OR circuit diode, wiring, etc. The cycle time was 12 microseconds, and an add to the accumulator took two machine cycles—one to decode the instruction and retrieve the data and one to do the add cycle. It typically had 4,096 storage locations (7,777 octal) of 36 bits each.

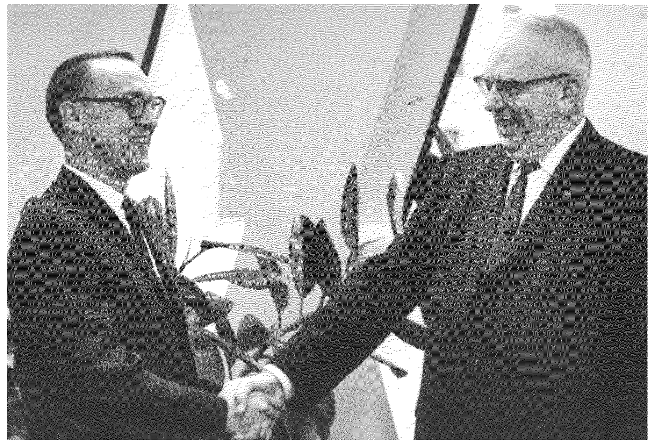

Fig. 1. The author (left) being promoted to advisory technical specialist in 1965.

The 7090 was my introduction to Standard Modular Systems (SMS), which were transistorized circuitry with individual transistors replacing the tubes—all the resistors, diodes, and capacitors were still discrete components. Some wiring was now deposited on the cards. Individual cards typically held one or two circuits such as AND/OR gates. Storage still consisted of individual cores for each bit position. Typically, the 7090 had 32,768 words of 36 bits (12 octal positions). This availability of limited storage certainly contributed to the year 2000 problem inasmuch as it was important to conserve storage locations.

The 709x series also brought on independent I/O channel processors, freeing CPU from that task. With the system 360 came IBM's first integrated circuits mounted in metal containers about one-inch square that contained deposited resistors and transistors and diodes. It was designated SLT for Solid Logic Technology. With the 360, IBM also used microcoded decoding of instructions in either punched 80-column card capacitor formats (called CCROS, card capacitor read-only storage) or transformer circuits, but the high-performance 360/75 still retained all hardware decoding of instructions.

As a technical specialist, and later advisory technical specialist, I spent many weeks out of town on assist calls and local calls on overtime. Certain fixes stand out in my memory even today, some 40-plus years later. I presume the intensity of concentration had burned the situation into my memory.

My very first programming fix was on Fortran at Westinghouse Bettis, when one night I was called in because the 704 had halted with a red light halt. Op code "00" would do that, and as I scanned storage, I discovered it was cleared out. Not a single "1" bit was in storage. Herbert Bright was the lead programmer at Westinghouse Bettis, and in reading the highlights of 1996 (IEEE April 1996 Highlights), I see that Bright had gotten the very first Fortran compiler out of development on 19 April 1957. It was shortly after that, sometime in 1957, when I was called out to determine the problem. Now I did not really know what Fortran was, other than a compiler, but Bright had an assembler listing, and he had it opened to the page where the "STOP" card was processed. The STOP card was the last card read into the card reader prior to the halt on the 704. I picked up the listing and reviewed it as Bright was on the telephone with someone. It did not take long to spot a problem with the loop that was processing the card.

The loop was looking for a blank position in the 6 position MQ register to stop it. On a 704, the MQ register also served as the I/O register. My thought was, "But what if it never finds a blank?" Looking at the card deck and the particular card, I could see immediately what the problem was. The STOP card had extra parentheses around it, thus "(STOP)"—which would fill the MQ register with other than blanks—so this loop would never end until it wiped out all the storage positions. The extra parentheses would normally be ignored, but in this loop they were disastrous. Easy solution, just remove the parentheses. Bright also told the Fortran developers, I am sure.

As we gained experience with the 704 and as the systems aged, we found that for quite a large percentage of problems, one simply had to remove the covers and look at the particular bit position for a burned-out tube. One entire gate of the 704 consisted of the 36 bit positions of the Storage Register, Accumulator, Adders, and MQ register. Burned-out tube filaments were the main trouble with the systems as they aged. Another problem that showed up on older systems was called silver migration. The wiring into each gate was terminated in edge connectors. The pins were silver-coated. In time, the silver migrated between pins with potential differences. This caused frequent outages.

The customer at Westinghouse Bettis reported a failure to us that was a bug in the 704. As we investigated this situation, we found it also was repeatable on the other 704 systems in the area. It concerned the least-possible floating-point number where there was only a "1" bit in the last position on a floating-point add operation—the very smallest possible fraction. We eventually determined that the ripple time, through all 27 floating-point fraction positions, of the adders to accumulator for a floating-point add was the problem. The four-microsecond pulse gating the adders was not long enough due to the powering of a particular line gating the carry pulse across 27 positions. Carry look-ahead circuitry was not in use yet.

We installed an on-the-spot engineering change of adding another power tube (5687) to the circuit to better shape the pulse and reported it to Poughkeepsie engineering. It fixed the problem. I remember there were at least three 5687s already powering the line. Since the customer had reported the problem, we had to fix it. There was never a formal engineering change as far as I know. This situation was rather remote. I have recently seen that the 704 could execute about 5,000 flops with its 12-microsecond cycle times. The fact that it would fail on one particular numerical value was unique, but this customer had found the failure. This was probably the only 704 with that change on it.

I was frequently sent out of town on assistance calls. After so many hours spent on a down situation, we had a procedure to call the next level of assistance. Such a call occurred during a trip to NASA Ames near Cleveland that, I remember, was during a snowstorm. I drove the turnpikes from Pittsburgh toward Cleveland, wondering if I would ever get there. I arrived and immediately went to the system room. This was when NASA Ames had a 704. Three field engineers were intently drooped over the system console panel trying to decipher what was wrong.

As I entered the room, I noticed a blinking neon on the CPU wiring back panel. These neons were hooked between the heater and cathode of each row of pluggable units to detect heater-to-cathode shorts. This neon was steadily blinking. I called their attention to the light. We uncammed each pluggable unit in that position and found the one with the bad tube. We changed all eight tubes and ran the diagnostics and went home. Now I am not blaming anyone—I just had a different perspective.

One of the items that field engineers did (that one may have suspected but could not prove) was the time a 716 printer on the 704 system would not operate correctly at its rated speed. We found that if we reduced the speed from 150 lines per minute to 120, it would operate correctly. We had no idea at the time how to make it operate at the rated 150 lines per minute. We spent many hours of our preventive maintenance time trying to decipher that one. Since this was used simply as a line printer for operator information, we did not feel bad that we could make it work at only 120 lines per minute versus a rated 150. Failures stand out also along with the victories.

I once was at this installation when a programmer who was looking over the printout from the 716 said, "Oops, we built a bomb." This was where they built submarine nuclear engines. It was a time when we felt very close to what was happening in the world during the cold war.

Once at Huntsville's Redstone Arsenal, I was called in on a 7094 bit drop problem that was simply an adjustment of the timing pulse gating the output of the storage sense amplifiers. I noticed that a 1403 printer was being worked on at a nearby 1401 system. I left and planned to return the next morning to see if my diagnosis on the 7094 was correct. The next morning, I was asked to look at the 1403 problem, which by this time was on another system with identical results—"sync checks," which on a 1403 meant that the chain printer position was in doubt. As this chain flew by the hammer, a counter kept track of what character was possible to be printed. If it lost track, then gibberish was printed. Since this 1403 had been moved to a different 1401 system, we knew the problem was definitely in the 1403 printer. I figured it was possibly due to electrical noise. I therefore put the scope probe on a voltage bus and immediately saw very large spikes. We started disconnecting motor fuses. One fuse that controlled the hydraulic reservoir pump eliminated the noise. Taking the top from the hydraulic reservoir, I could see a line had been run over the top of a metal pulley and had worn through the insulation, causing sparking and large voltage spikes on the line that fed the logic circuits. Whoever installed the changes to the wiring did not realize the importance of properly dressing the wiring.

Another Redstone Arsenal trip was when IBM first brought out the 7094II with double-precision floating-point operation. Redstone's 7094 had been brought up to 7094II in a field conversion that had taken over a week. I was assisting and had been at the installation for 10 days. During final checkout before turning the system back to the customer, we had run diagnostics. A Double-Precision Floating-Point Add diagnostic failed for perhaps 30 seconds. We did everything we knew to make the system fail again, but to no avail. We were told that the system would have to be fixed as "We are sending a man to the moon." Those German scientists/programmers would accept nothing less than all problems resolved. Thanks to them, we got to the moon. Now, as I had been 10 days at Huntsville on this change and was eager to return home, I was determined to solve this problem. The printout had been tossed in the garbage, so I retrieved it. I examined the diagnostic printout, the circuitry, and the manual of operations on Double-Precision Floating-Point Add to see if I could determine what was possibly wrong for the 30 seconds. After about five hours of research, I could see that the only probability was a three-legged AND circuit with one leg likely having a slow fall, allowing a spike to be emitted from the AND circuit. I relayed this information to an IBM engineer who had been the main factory contact on the large change to the 7094. I suggested that this AND circuit (MQ1 to ACC 35) was emitting a spike due to the slow fall of this other leg to the AND. Sure enough, it was the correct call. There were some incredible engineers on that one.

The 7090/7094 systems had an oil-cooled storage unit. The array of cores was immersed in an oil bath to keep temperatures even throughout the array. To remove this array was not something a field engineer would look forward to—but of course it happened. Intermittent errors were traced to the probability that we had a shorted array pin or pins. This was at Westinghouse East Pittsburgh. Removing this storage array was no small matter, as it was about three feet tall and inside a sealed tank that was very messy to remove. It also took a small crane. Upon removal, we had a fairly good idea where the array would have been shorted. I looked carefully at the suspect area and saw a solder splash. The splash had the shape of a hand with four fingers. Carefully clutching the wiring, this little splash of about one-quarter-inch width had done its dirty deed. I always wished I had saved it as a souvenir.

On a call to assist at NASA Ames, we were looking for a grounding problem on a 7094. The system was grounded to frame at only one place to control ground currents. This system had developed a short in another spot, and we were looking for where it was. We had the gates out of several CPU frames on the 7094 lying about the machine room. Interconnecting cables were piled atop one another. It occurred to me that if the customer's personnel would appear, there would be extreme doubt that we would have the machine up and running on Monday morning, but we did. I can just imagine what the customer's upper management would have said.

I went to Detroit to work at Chrysler HQ on a 7094. I recall landing in a snowstorm. I had taken a minimum of clothing, that is, I had only one suit. The seat of the pants split open on my way to the hotel. What to do? All I could round up quickly were some straight pins. These I fashioned into a repair of my seat. The next day at Chrysler, I carefully always wore my coat and got jabbed a time or two but did not draw blood. Don't ask me about the trouble call results, I have no idea now. I remember only the trauma.

When the 1401 system first came out, the diagnostic programs were card input, each card with one instruction—which made for very large input decks. I reasoned that this was wasteful in several ways, so I submitted a suggestion to issue diagnostic decks using the condensed version that utilized much more of the card and would cut down the size of the card decks. After an initial turndown of my suggestion, it was later the practice to issue the decks in the condensed version. I resubmitted my suggestion and was paid $1,400 for the suggestion, quite a sum in the early 1960s.

When the 7094 and 1401 systems were phased out with the introduction of the 360, our customers ordered several models. So in 1964 and 1965, I went back to Poughkeepsie and Kingston to be trained, specifically on the 360/30, 65, and 75 as well as the operating system OS/360. An incident on the 360/30 stands out. We had put on an engineering change, and the operator reported he could no longer get his system to operate correctly. All diagnostics were OK. The "Diagnose" instruction had been changed. This instruction, previous to the change, read in the keys from the front panel, and the customer had programmed the "Diagnose" instruction to determine if the operator wanted to do a "Tape to Printer, or Tape to Punch, or Reader to Tape" as an offline system. The operator dialed a specific number to get the selection. After our engineering change, this no longer worked. I remember calling the Endicott office with the problem and asking for a fix. None was forthcoming, since "We might make a mistake over the phone." Meanwhile, the system was useless to my customer. My partner had been reviewing the circuitry while I was on the phone. He said, "I see what we can do to change the microcode to make it work." So I said "never mind" to the person on the other end of the telephone. "We see what to do." So we changed the microcode card and fixed the system so the customer could use it. What else but "customer service."

One of the more insidious bugs I fought was one that took 10 minutes into the program to halt a 360/65 with a red light. The 360/65 Emulator program was running a 7090 compatibility package that was running a 704 program. I scheduled myself in on a Sunday to take the system and try to track down the cause. I ran the system up as close to the halt as possible with a stopwatch, then *single cycled for nearly eight hours* as I looked through various listings to see what was happening. The red light could come on for several reasons, some not well-documented. It turned out to be a simulated/emulated 704 "Copy Check." I was really disappointed that the 7090 compatibility package running the 704 program under the 360/65 7090 emulator would even consider that it had to simulate a "Copy Check" that on a 704 was a disaster but on a 7090 or 360 system was an impossibility. So why not ignore it? As the 704 CPU served as the I/O processor, the timing between a Write or Read Select and the "Copy" command was critical, since records could easily be skipped by the "Copy" command coming too late. For instance, a tape drive would start tape movement upon detecting the "Read Select." CPU processing was allowed between the Select and Copy. The "COPY" command gated the bits from the I/O device to the MQ. However, on the 360 and 7094, the I/O was processed independently of the CPU, and nothing happened until the I/O commands were loaded into the I/O processors. Therefore, no timing problems ever existed in the manner they were possible in the 704. My fix for this one was to put in a patch in the 360/65 emulator software to ignore the condition of a late copy command, which was undoubtedly due to the three layers of programming plus a layer of microcode. Now one can see that field engineering sometimes has to take a practical approach.

Another fix that points out this approach was one in which the 360/65 emulator would go into a loop while printing out error messages on the 1443 printer. Possibly this error would never happen if the printout device had been a 1403 printer that was much faster than the 1443. I do not know. In examining the loop, I saw that the error buffer would hold about 40 error messages, and the 360/65 emulator program would then loop in a rather long sequence that I did not care to try to figure out. It was late in the day. I reasoned that if we already had 40 errors, why would it be important to print out more? Surely it could be figured out from all those messages. Therefore, I patched the routine to ignore more errors if the buffer was full. Why beat it to death? The 360/65 7090 emulator package made it very easy to insert patches. Was someone looking ahead?

One of the more difficult 360 errors took several weeks to track down. This was early in the installation of the 360/65 at Westinghouse East Pittsburgh. The system would come up with a very infrequent missing record on tape. We were stumped for quite a while, until examining the error recovery routine, we discovered the I/O command string was depending on only a count to stop the processing of a record. It was possible on the 2400 series of tapes to drop a character or two and not detect it until the end of the record. This meant that any unreadable characters the tape might skip over would leave the count unfulfilled, and the tape would merrily go on trying to find another character to fulfill the count command. Unfortunately, this would skip over an interrecord gap, read one or more characters from the following record, then as the count command was fulfilled, the disconnect would come. However, the tape control unit would continue reading on until it came to the next interrecord gap, which meant the tape was now mispositioned and had skipped a record. The error light would come on, and the error recovery routine would try to recover, but it would be over the wrong record. The I/O command string was changed to look for tape error (e.g., a dropped character) before proceeding blindly, just counting four-character words.

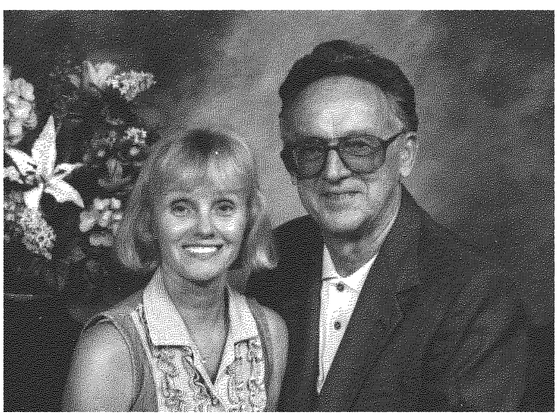

Fig. 2. A recent photo of the author and his wife, Sonja.

My very last fix as a pseudo–field engineer was in 1975 at the IBM Development Laboratory at Rochester, Minnesota. This was after I had been out of the field for six years. I had supported this 360/65 7090 emulator system for the six years I had been there, working on several software and engineering changes/problems. I was asked to assist when the local field engineers got stumped on the system. It was a Convert-to-Binary instruction. It appeared that the results were random, but I reasoned that the input was probably a ripple pattern, and if I backed the results back through the adders, I would come up with the real problem. As I did that, I found it was a bit 7 in a 0 to 7 bit number character being picked up under certain conditions, then scattered throughout the adders as carries moved the bit pick to other apparent positions. At that time, I managed a department in Rochester that did custom programming for various projects. I never really lost my attitude or aptitude on the 360/65. So regardless of the circumstances, we got the system running and did everything to get it there. Regardless of who said what, the uppermost thing was to get the customer's system fixed.

I retired in 1990 (see Fig. 2).

This concludes my "True Confessions" but the Field Engineering Division's motto "IBM Means Service" was taken to heart.

Glenn E. Meyers can be contacted at
19246 Pine Bluff Ct.
Port Charlotte, FL 33948, U.S.A.
email: gemgramp@sunline.net