

News from Nan

COMPUTER HISTORY

MY HIGH-TECH ADVENTURE: CHAPTER 06, SLAC

OCTOBER 6, 2018 | DICK

Dick wrote this post, which was part of a book, My High-Tech Adventure, to record his memories of working in the computer industry, starting in about 1960. It was first published in 2012.



My family and I at Endstation A at a SLAC Family Day celebration around 1969

The IBM 360/91

I left Bell Labs in 1969 and took a job at the Stanford Linear Accelerator Center (SLAC), a high-energy physics research center in Menlo Park, California. SLAC was operated by Stanford University under contract to the Atomic Energy Commission (AEC). SLAC did research on the fundamental particles of the universe. The research was done by sending electrons two miles down the accelerator and crashing them into targets so they could observe the shower of collision products.

When I got to SLAC they were running physics research computing batch jobs on an IBM 360/91 computer. In 1969, this was one of the most powerful computing machines on Earth. It could run at 3-16 mil-

million instructions per second (MIPS), depending on the workload, and had 2MB of core memory. The machine had cost SLAC roughly \$8 million. The console of the 360/91 had a flashing light display that was about 3 feet high and about 10 feet long. When the machine ran, the lights changed in all kinds of patterns. Only the IBM customer engineer (CE) knew what the lights meant.

The 360/91 was a very touchy machine. If it got powered off for any reason, it sometimes took days to get it working again. Main memory was divided into eight 256K core memory banks, called BSMs, that were each about the size of a modern refrigerator. IBM kept several spare BSM units around because the memory failed all the time. Failure was so common that IBM supplied SLAC with two CEs that were on site all the time just to keep the machine running. You can afford to do this if the customer pays \$8 million for a computer!

Also attached to the 91 were IBM 2305 drums. These held very little data, but had a head per track, so they were very fast.



IBM System/360 Model 91 at SLAC

In 1969, SLAC was running the HASP OS/MVT operating system on the model 91. Most physics jobs were large (300K) Fortran programs that processed physics experimental data stored on magnetic tape. Many of the SLAC computer users were still storing programs on punched cards, but there were a few IBM 2260 video terminals connected to a software system called CRBE (Conversational Remote Batch Entry) that ran on the model 91. CRBE was a primitive text editor and job submission system, running and storing programs, data and output on disk. Users accessed CRBE terminals in a large room we called a “bull pen.”

The operating system was unstable and we of the Systems Group were forever fixing bugs and installing IBM fixes and our own local enhancements for it. For this purpose we reserved the machine every weekday morning from approximately 7:00 to 8:30 a.m. Not exactly a non-stop operation! (There was notable exception to this regular morning outage. One day in 1974, when I came to work the 91 had been left in production. That was because the Richter physics collaboration had just discovered the charmed quark and were busy verifying their results! (This discovery was good enough to win Dr. Burton Richter the Nobel prize.)

I remember that we had a lot of problems with the Fortran IV compiler. What is funny is that when I left SLAC in 1986, they still had a full-time person installing IBM fixes for the Fortran compiler!

User and system data was stored on banks of IBM 2314 disk drives. Each disk drive held 29MB and a bank held 233MB. At this writing (2003) compact flash cards that fit in a shirt pocket hold 512MB.



IBM 2314 disk drives

There were a few IBM 2250 graphics displays attached to the 91. The 2250s had a screen about as big as a 19-inch TV and were very expensive. Use of the 2250 was severely limited because you needed to tie up 300K of the 2MB of memory to run the batch job driving the display. Sometimes on weekends, the physicists ran a Space War game on the 91 that used the 2250 displays.

The 91 was a real, not virtual, memory machine. SLAC users had to use many tricks to fit large programs into the limited memory using overlay structures. Some of the physics users were very clever at creating those overlays. Using overlays meant the user had to decide when code should be read in to overlay in memory other code that was no longer needed. This was what people did before computers supported virtual memory and paging.

When I first came to SLAC I became manager of User Services. We ran a consulting office where we answered questions and diagnosed system bugs affecting the users. This was my only management job in computing. My manager and mentor was Mel Ray. (Mel went on to work at the World Bank after he left SLAC.) My department included myself, Ted Syrett and Bernie Tice. We later added John Ehrman as a consultant.



Mel Ray

In addition to managing User Services, I also worked on a program execution profiler called PROGLOOK. After a while I moved into the Systems Group to work on computer measurement and dealt with such things as System Management Facility (SMF) data gathered by OS/360, computer accounting and measurement. Since the hardware was so expensive, we could afford to have people work a lot trying to optimize its use. My manager in the Systems Group was Ted Johnston. I worked for Ted for almost 15 years.

Milten, Wilbur, Orvyl

In the early 70s, SLAC moved to a computing environment developed at the main Stanford campus computer center that included:

- MILTEN – software that controlled hundreds of 2741 terminals connected to the 91
- WYLBUR – a line-mode text editor that could be used to create files, submit batch jobs, and look at job output from a terminal
- ORVYL – an interactive monitor that allowed people to do a simple form of interactive computing by running programs in various languages

John Halperin worked on MILTEN and Joe Wells took care of enhancing WYLBUR/ORVYL. WYLBUR was a big step forward from CRBE, the users liked it a lot, and for many years was in use at universities and other institutions all over the world. It provided more function and used a lot less computing resource than the IBM equivalent product, which was Time Sharing Option (TSO).

ADM3

About this same time, I bought my first home terminal. This was an ASCII terminal kit that we got through the Stanford University computer center at a discount. The kit had over 2000 solder joints and took many hours to complete. I logged on to the SLAC computer from home with the ADM3 using a 300-

baud acoustic coupler modem. Now people think 50KB is slow! Primitive as it was, this terminal saved me many trips to SLAC to fix problems.



My daughter Ellen at the ADM3

Co-workers

Here are some of the people I worked with at SLAC:

Ted Johnston. Ted was my boss and the manager of Systems Programming almost the entire time I worked at SLAC.

Chuck Dickens. Chuck was director of the computer center.

Joe Wells. Joe maintained and enhanced SLAC WYLBUR. He later went on to work at the IBM Yorktown Research Center.

Paul Dantzig. Paul was the son of George Dantzig, who helped invent linear programming, I still have a copy of George's *Linear Programming and Extensions* on my bookshelf. At this writing (2003), Paul is at the IBM Yorktown Research Center in Yorktown Heights, New York. Paul taught me quite a lot about automotive mechanics.

Sam Fuller. I knew Sam as a graduate student in the Stanford EE department. He later became vice president of R&D at DEC.

Bill Weeks. I hired Bill as a student from Stanford, where he had been a psychology major.

Joan Winters. Joan was a User Services consultant and our usability expert. She also worked on online help systems and was a leader and participant in many SHARE activities.

John Ehrman. John came to the computer center from the Computation Group and was very knowledgeable about IBM 360 Assembler language. He was also very active in SHARE. Like me, John later went to work at IBM.

Others. I once sat in a meeting at SLAC with Linus Pauling, and in other meetings with Donald Knuth and George Forsythe. These people never knew me; I was just sitting in the room. From the physics world I got to see Pief Panofsky, Burton Richter, and Marty Perl, among others, quite a few times. Since SLAC had so many PhD researchers on its staff, nobody was referred to as “doctor.” At SLAC we talked about Pief, not Dr. Panofsky. Marty and Burt both won Nobel prizes.

Note: Even though the Nobel prize was awarded to individual physicists, the work of each physics experiment was done by collaborations with hundreds of individuals in each. Only the top person of the collaboration got the award and the money.

I worked in Richter’s physics group once for two weeks helping them to convert to the new IBM VM/370 operating system. The group members treated Richter like a god and obviously treasured every word he said to them. During the whole time I was there, Burt never came out of his office to talk to anybody, so I didn’t get to meet him.

The Triplex

Later in the 70s, SLAC did a major computer acquisition (these happened about every 5 years) and bought two IBM 370/168s to augment the 360/91. We called the three computers the Triplex. The 168s each had 3MB of memory and were the first virtual memory machines I ever worked with. The whole Triplex system was controlled by IBM’s ASP MVT and SVS operating system, still with MILTEN/WYLBUR/ORVYL. The 2314 disk drives were upgraded to 3330s and 7-track tape drives were supplanted by 9-track drives.



Installation of the SLAC Triplex

When all the Triplex hardware arrived, it filled up the computer building we all had been working in, and most of the staff was forced to move to new offices. Those offices ended up being trailers that the AEC

had surplused from the Nevada Test Site where atom bomb testing had been done. Many SLAC people hated working in the trailers so much that they quit their job. I hung on and we did eventually get a new computer center built. Interestingly enough, this was because we failed a security audit for the computers, not because SLAC felt sorry for the people working in trailers.



SLAC trailers from the Nevada Test Site

IBM SHARE

The users of large IBM computers all belonged to a user group called SHARE. (Commercial customers had a similar organization called GUIDE.) SHARE was founded in August 1955. I went to my first SHARE meeting in the late 60s, but SLAC was very active at SHARE, and I went to over 50 SHARE meetings while at SLAC. My manager, Ted Johnston, was a big backer of SHARE attendance for the staff. The SHARE meetings were held in large hotels in big cities like Los Angeles, Anaheim, Houston, Chicago, and New York. For a while I knew my way around many of the big hotels in the U.S. At its peak, SHARE would host over 6000 people at a meeting. The attendees were from all the largest corporations, government organizations, and universities. When IBM was so dominant in the computer industry, prior to about 1985, SHARE was the computer meeting to go to.

I gave quite a few talks at the SHARE meetings. This is how I got to be comfortable speaking in front of a large audience. One talk I gave on hardware measurement of an IBM mainframe had several hundred people in attendance.



IBM SHARE 71 Registration Area

At SHARE, I first worked in the Computer Measurement and Evaluation (CME) Committee, where we discussed how to monitor and instrument the large mainframe computers of the day. Tom Bell from RAND managed the CME group. CME worked with IBM designing an SMF (System Management Facility), which was part of OS/360. Later I moved to being a member of the CMS Project when I became more involved in interactive computing in the 70s and 80s.

The thing some people at SHARE meetings aspired to was to become a SHARE officer. I was never a SHARE officer. Officers got to hear IBM “secrets” that normal SHARE members did not get to hear. Officers wore colored ribbons on their suits and were called ribbon-wearers. I later learned when I went to work at IBM that some of the “secrets” were carefully planted by IBM development groups that wanted to drum up customer excitement over their latest development project. Some of the IBMers were very good at getting customers to say they wanted to buy things that IBM developers wanted to work on.

At the peak era of its computer industry dominance, IBM was very tight-lipped about what it was going to do and it rationed the introduction of technology so as to maximize revenue. There were actually companies that bought first-day order positions for new IBM hardware that they would sell to other companies like shares of stock. The machines always cost about \$5 million and came out every 5 years.

After I left SLAC for IBM, I served as the IBM representative to the VM System Management Project. I got a very different perspective on SHARE attending as an IBMer!

The following figure shows the description of a talk I gave at SHARE in Atlanta in 1992.

Overview of AIX/ESA DASD Storage Subsystem Support O922

This session will provide an overview of functions provided by the DASD support in AIX/ESA. Items described will include the hardware supported, the DASD device driver main functions, ICKDSF under AIX/ESA and DASD EREP support under AIX/ESA. The external interfaces to the device driver, including all supported ioctl calls, will be discussed. In addition we will describe the system files that control configuration of DASD devices.

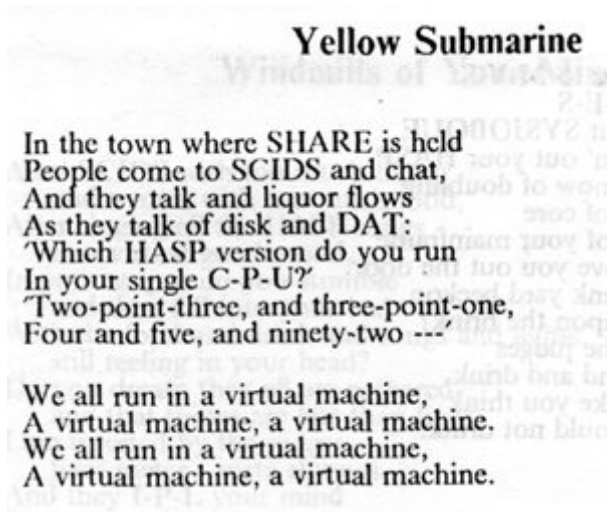
Speaker: Richard H. Johnson (IBM)
 Chaired By: Harold Pritchett (UGA)
 Sponsored By: AIX Operating Systems Project

Marriott, Conv. Lvl., Summit

Description of one of my SHARE talks

SHARE meetings were intense. I would typically attend 5-6 sessions during the day on various IBM technical topics or project working sessions. At night there were also impromptu meetings called “birds-of-a-feather” that lasted as late as 9:00 p.m. At 6:30 p.m. every night of the meeting there was an open bar in a large ballroom where everyone tended to collect to talk and to make plans for going out to dinner. This was called SCIDS, and I don’t remember what it stands for. On Thursday night at SCIDS around 9:00 p.m.

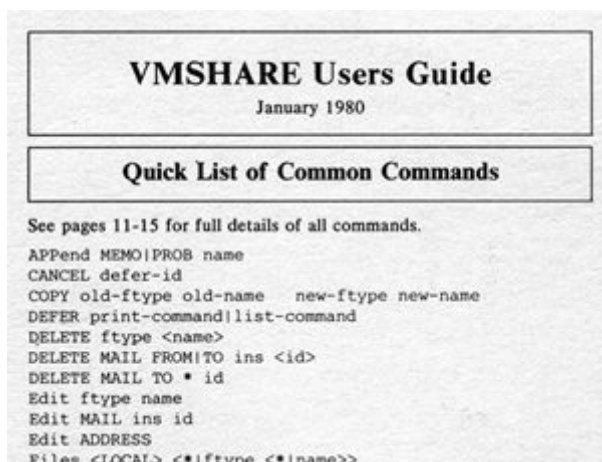
everybody in the room sang silly songs from the HASP (Houston Automatic Spooling Program) song-book. Here is one I sang many times.



Lyrics to SHARE SCIDS Yellow Submarine VM song

BITNET and VMSHARE

Long before the advent of the Internet, around 1981, SLAC was part of a large computer network called >BITNET. This network linked thousands of computers all over the world and allowed users to exchange files and e-mail. The network fostered collaboration from customers on improving the utility of VM. Many of us were part of a collaboration called VMSHARE that included a bulletin board system devoted to various topics on improvements to VM. We also used it as a way to exchange information about problems and fixes we had found. IBM distributed all the source code for VM to VM customers, so they could make fixes and improvements of their own very quickly. In the present era a similar thing is happening with the Linux operating system and other open source projects. As the saying goes: this has all happened before and it will happen again.



VMSHARE User's Guide

1984, the year of travel

In 1984, the work on VM at SLAC had attracted interest from other physics laboratories in Europe and Japan. SLAC started to get requests for exchanges between our programmers and those of the other labs. As a result I was able to go on some interesting trips that year.

IN2P3 and CERN

In July 1984, I traveled to physics labs in Geneva and Paris as a consultant to help get the SLAC version of VM running at both places. I started at CERN (Centre Europeenne pour la Recherche Nucleaire) in Geneva. CERN is the largest particle physics research lab in Europe. It is located near Lake Geneva on the border between France and Switzerland. I spent most of my working time at CERN talking to Sverre Yarp and people on his computer system group staff, mostly explaining how the SLAC system worked. After a while it became clear that they were not really going to just install the SLAC software, but were going to rewrite it all to their own needs. CERN had a big budget and could afford to do things like that. In my off hours I got to see lots of the sights in Geneva and did a train tour all around Lake Geneva ending up at Chillon.



A view of CERN

My next stop after CERN was a French physics lab, called IN2P3 (Institut National de Physique Nucleaire et de Physique des Particules). It was located at the University of Paris. I spent two weeks there helping the IN2P3 Computer Center (Centre de Calcul) install some of the SLAC modifications to the base VM operating system. My main contact there was John O'Neill, a U.S. ex-patriot. I enjoyed working with John and other people on the staff, and I got to wander around Paris on nights after work and one weekend for many long walks. I was truly smitten by the Paris bug and have remembered the place fondly ever since.



IN2P3 and the Notre Dame cathedral

IBM Japan and KEK

In October of 1984, I went to Tokyo, Japan, to help IBM Japan put together a bid for computing facilities at the KEK physics lab (north of Tokyo). I spent my time explaining to people on the IBM Japan staff how SLAC had used VM for physics work. I found out later that IBM Japan lost out in the bidding process against Fujitsu, who basically gave away their computers to get the bid. I was able to tour Tokyo evenings and on one weekend.



Office of IBM Japan in Tokyo

◀ 2018 ◀ BERNIE TICE ◀ BILL WEEKS ◀ BITNET COMPUTER NETWORK ◀ BURTON RICHTER ◀ CERN PHYSICS LAB
 ◀ CHARMED QUARK ◀ CHUCK DICKENS ◀ FORTRAN PROGRAMMING LANGUAGE ◀ HASP SONGBOOK
 ◀ IBM 2214 DISK DRIVE ◀ IBM 2250 GRAPHIC DISPLAY ◀ IBM 2260 VIDEO TERMINAL ◀ IBM 2305 DRUM
 ◀ IBM 3330 DISK DRIVE ◀ IBM ASP OPERATING SYSTEM ◀ IBM CRBE SOFTWARE SYSTEM ◀ IBM CUSTOMER ENGINEER (CE)
 ◀ IBM JAPAN ◀ IBM SHARE ◀ IBM SVS OPERATING SYSTEM ◀ IBM SYSTEM MANAGEMENT FACILITY (SMF)
 ◀ IBM TIME SHARING OPTION (TSO) ◀ IBM VM/370 OPERATING SYSTEM ◀ IN2P3 ◀ JOAN WINTERS ◀ JOE WELLS
 ◀ JOHN EHRMAN ◀ JOHN HALPERIN ◀ KEK PHYSICS LAB ◀ MEL RAY ◀ MILTEN CONTROL PROGRAM ◀ MYHITECH
 ◀ ORVYL INTERACTIVE MONITOR ◀ PAUL DANTZIG ◀ PROGLOOK COMPUTER PROGRAM ◀ SAM FULLER ◀ SLAC TRIPLEX

◀ STANFORD LINEAR ACCELERATOR CENTER (SLAC) ▶ TED JOHNSTON ▶ TED SYRETT ▶ VMSHARE COLLABORATION
◀ WYLBUR TEXT EDITOR