from DATAMATION Jan 1982 STRETCH

IN FOCUS

FASTEST IN ITS TIME

by Edward K. Yasaki

The 1961 Stretch computer had a phenomenal list of "firsts." But the benefits to IBM extended far beyond that Ferrari of a machine.

Things just aren't the same anymore. It used to be that at this time of year there was a certain excitement as one anticipated seeing what the new cars from Detroit looked like. They would undoubtedly be longer and lower and certainly much sleeker. But the curiosity that once drew people to showrooms has passed.

And so it is with computers. For some reason, the computers of old stirred one's emotions more than the new ones do today. Maybe it's because they didn't make as many of any one model as they do today, and one could conceivably read about an exciting new mainframe and have it become obsolete before ever laying eyes on the real machine. How many people ever saw an Atlas computer, or a Univac Larc? Have you ever seen an Illiac?

One such classic is the IBM Stretch computer, in its time the world's fastest, the Ferrari of the computing scene. Some five years in gestation, Stretch's first customer shipment was in May 1961. That same month, IBM took no more orders for this sleek new computer. But more than sleekness, Stretch had a list of firsts that extended beyond one's reach. They pulled out all the stops on this one, entertaining any and every idea that young Turks just out of college could dream up. It was a sort of test bed for new features, the best of which were to appear in succeeding machines like the System/360s.

But perhaps most significantly it attracted bright young people to IBM, many of whom went on to design the 360s and 370s and their software; several were to attain management positions in the company. Some, like Frederick P. Brooks Jr. and Gerrit A. Blaauw, went into teaching, Brooks heading the computer science department at the University of North Carolina and Blaauw doing the same at Twente University in the Netherlands. An impressive number of them became IBM Fellows, among them Stephen W. Dunwell, Harwood G. Kolsky, John Cocke, James H. Pomerene, Robert A. Henle, and Edgar F. Codd.

One little-known first on the Stretch was its use of a Selectric typewriter as a console printer. The golf ball typewriter was not officially announced as an IBM product until 1961. As a result, whenever a visitor came to the machine room at Poughkeepsie, N.Y., where a Stretch had been set up, measures had to be taken to conceal or disguise this unannounced typewriter with its odd-looking typing element. IBMers did this by devising a piece of cardboard that fit over the slot on top of the typewriter. Harwood G. Kolsky recounts an amusing incident.

"I remember one time hosting some visitors at the Poughkeepsie lab and standing around giving the standard talk. One of the visitors walked over and picked up this piece of cardboard and looked inside. I thought everybody would die!" Fortunately the visitor didn't appear to have noticed anything new or different under the makeshift lid. "I looked around at the other IBMers and they were all turning pale."

People associated with the machine, either as designers, implementers, or users, had their own list of favorite firsts in the Stretch. When recalling some of those, they start with such features as the instruction look-ahead (four levels deep) or interleaved memories, the 8-bit byte and variable byte size. But Fred Brooks retrieves a list he has filed away and begins to read from it. There

Perhaps most significantly, Stretch attracted bright young people to IBM, many of whom went on to design 360s and 370s and their software.

are the supervisory facilities, memory protection, the maskable interruption system, the concept of having the console program interpreted, rather than hardwire defined. There was a separate input/output computer, the Exchange. The Stretch had bit addressing, boundary-free alignment, in the fashion of the 370s, and relative branching.

A significant feature, of course, was the use of error correction codes. The machine used 64-bit words plus eight bits of Hamming code to form a 72-bit word in memory. It had single-bit error correction, double-bit error detection. It had provisions for upper and lower case character sets, and could perform decimal, binary, and floating point arithmetic. The concept of a standard interface for I/O equipment was a first for IBM; before 1956 its computers had a different interface for each kind of I/O gear.

Of course, Stretch came originally with an oil-cooled core memory system.

Stephen W. (Red) Dunwell (at right) became IBM's scapegoat for Stretch, but was later apologized to and named an IBM Fellow. After 41 years with the company, he left to save from demolition Poughkeepsie's 110-year-old Bardavon Opera House.

PHOTOGRAPH BY JAMES JOERN

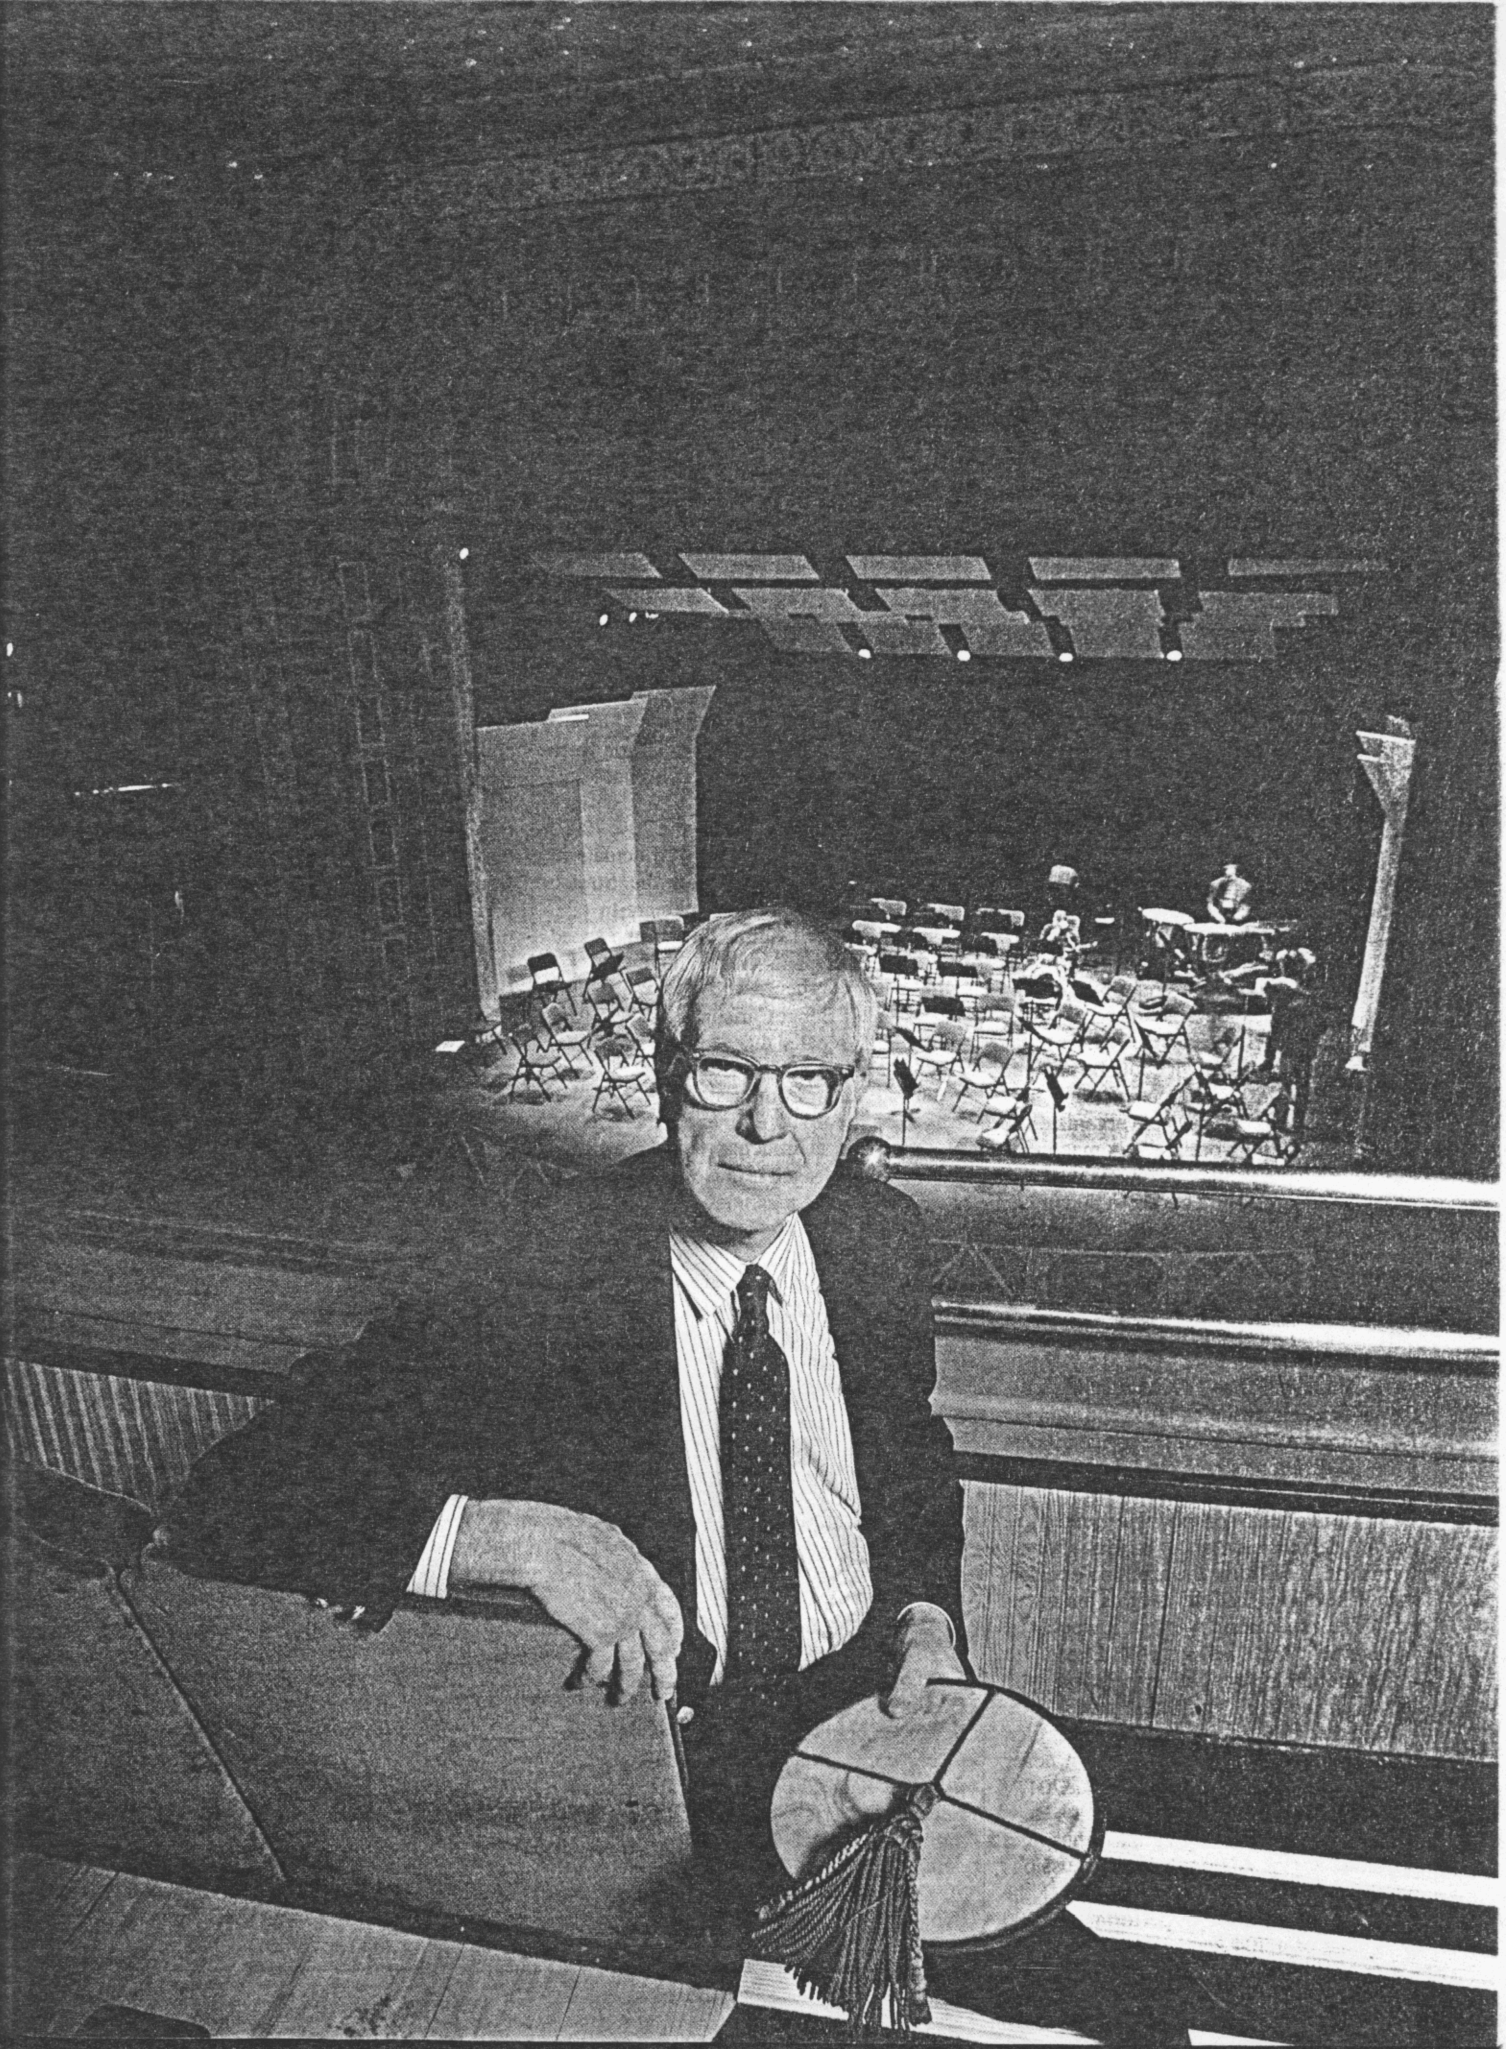

Jack Worlton of the Los Alamos Scientific Laboratory (LASL) recalls a problem experienced during the acceptance test phase, when there was a transient memory error in the cores. The engineers worked for days to remedy the error but were unable to figure out what was going wrong. They finally hit upon the problem. As Worlton explains it, there was a piece of solder loose in the oil bath, and because the oil was constantly in circulation the solder would move and attach itself onto a core and cause an error. Then it would move and lodge onto a different core and cause an error there.

"It was the only error I know of that was corrected in a machine by giving it an oil change," Worlton quips.

In a paper presented at the 1959 Eastern Joint Computer Conference in Boston, Erich Bloch said the objective with Stretch was to achieve an improvement in performance over the 704 by a factor of 100. Bloch, up who heads the corporate technical personnel development staff, said they could see a possible sixfold improvement in memory performance over the 704 and a tenfold improvement in basic circuit speed. In his paper, he even notes that "Simulation of Stretch programs on the 704 proved a performance of 100 times 704 speed in mesh-type calculations. Higher performance figures are achieved where double- or triple-precision calculations are required."

In the question-and-answer period at that '59 confab, Bloch was asked how much of the speed improvement could be credited to the use of faster components and how much to changes in the system organization. He replied: "I think one order of magnitude of improvement is due to faster devices and faster circuits. The other order of magnitude of improvement is due to system organization, multiplexing, and so forth. . . ."

From this paper, too, we learn that 169,100 transistors were to be used in the machine, mounted on two types of circuit boards or cards. There were 18,747 so-called single cards and 4,025 double cards, the latter being twice as large and packing four times the capacity of the former. There were 24 different single card types and 18 different double cards.

But Jack Worlton, who joined the Stretch design team about a year after its formation, recalls that the machine's performance never reached the heights anticipated. "The expectation was that it would run faster than it did," he says now. Even up to a few months before delivery, people who had been modeling the Stretch were forecasting that it would run 75 times faster than a 704. "In fact, to my knowledge it never ran more than 25 times faster than a 704," he says.

Harwood Kolsky of IBM, who was a physicist at Los Alamos when discussions on the Stretch began with IBM, says the initial projections of a performance 100 times that of a 704 was merely a ball park figure. But as time went on, it became an unrealistic target in the minds of the people involved. "This was one of the reasons the machine was later considered not to be successful."

He continues: "I should quickly add that it would be very easy to pick a problem" that ran on Stretch and transfer it to the 704, where it could take a thousand times longer—because the problem would overflow the memory. The Stretch had more than 100K words of memory, versus something like 32K for the 704.

Worlton explains that Stretch was one of the first machines with a broad performance spectrum. Anyone who took advantage of some of the machine's features could get it to run fast, but if the job were programmed ineptly, the performance improvement might be only five or six times greater than the 704. He says the 704, 7090, and 7094, for example, didn't fluctuate in performance that much. But the Stretch made it possible for programmers to take advantage of its features to gain the speed that was inherent in the machine.

One little-known first on the Stretch was its use of a Selectric typewriter as a console printer.

Ed Voorhees, who was assistant to group leader Bengt Carlson of the Los Alamos team, recalls that a performance up time of at least 90% was demanded by the Atomic Energy Commission (AEC), which was to be the first customer for the Stretch. He guesses that in actual fact the up time percentage averaged in the low 90s. "Not as good as today's machines, certainly," he says. But there's no doubting where Voorhees' heart is. Like so many of the LASL user community who ran jobs on the Stretch, he thought highly of the machine.

Worlton recalls the sum paid by the AEC as \$4.2 million, considerably below the price tag later set on the Stretch. But he said no one paid the list price. Ed Lafferty of the Mitre Corp. in Bedford, Mass., recalls that his organization acquired a new Stretch on a lease-purchase plan. He doesn't remember how long they had it on a lease, but says when the decision was made to buy the machine the final payment was for \$6 million, the first and only time he held such a sum in his hands.

The need to lower the price of a Stretch and the financial drubbing being taken by the vendor with each order received were disclosed by IBM's chairman, Thomas J. Watson Jr., at the Western Joint Computer Conference in 1961. The chairman, using the occasion to convene a press conference at the Ambassador Hotel in Los Angeles, set a cutoff date of May 15, 1961, after which no more orders were going to be taken. At that time, DATAMATION reported the price reduction was to some \$8 million from an original \$13.5 million, saying this was proportional to the shortfall in performance of the machine. As reported by DATAMATION, Watson said:

"We undertook the Stretch contract for the Atomic Energy Commission some years ago. They asked us for certain specifications that they wanted met. We said we could meet them within a certain time and then we went about doing it. The cost of building a computer was completely underestimated so that the government funds we have in Stretch are minor compared to IBM

STRETCH'S INHERITANCE FROM THE 709

In the midst of the Stretch development project, after many of the component subassemblies had been built but before they could be assembled into a workable whole, the federal government issued an RFP for a solid-state computer system to be used in the ballistic missile early warning system (BMEWS), part of the nation's defense setup. IBM proposed a transistorized version of the 709 computer, offering to supply a 709 for program development until the solid-state version could be built.

The new computer, which came to be called the 7090, inherited its system design from the 709 and its hardware from the Stretch, including the emitter-coupled logic circuits, the packaging, frames, covers, power supplies, and memories. Those pieces came directly out of the supply of parts collected to produce the first Stretch. Recalls Frederick P. Brooks Jr., "Even though Philco announced a transistorized computer first, IBM delivered one first—the 7090. And George Monroe's team did that thing in six months from start to first customer delivery." Later IBM was to similarly develop the 7080 from the vacuum tube 705 III, using Stretch componentry.

Harwood G. Kolsky, now at the IBM Palo Alto Scientific Center, recalls sitting in meetings with Monroe and suggesting that some of the weaknesses of the 709 be corrected in the new version. But Monroe, a stern engineer from the old school, was adamant in retaining the same features. Kolsky says now that Monroe was probably wise in not budging because if he had started making changes, everyone would have descended on him with their pet design features for incorporation in the new machine.

Kolsky observes, too, that the 7090 had "one tremendous advantage, which we now understand very well but didn't at that time. The new computer already had its software written." It came from the 704 and 709 and would run on the new machine unchanged. By contrast, developers of the Stretch were still struggling with a compiler, for example. So the new 7090 was able to leapfrog the software issues, says Kolsky, "which turned out to be much larger on Stretch than anyone had anticipated."

funds. At the end of the period [of the original contract], we were late on delivery date. And when we finally began to assemble the computer we found that though we had the world's fastest and most capable computer, the specifications were not met. . . .

"We will make delivery of the machines because we do not want to break our promise to our customers. We are going to take a good, fat loss on Stretch, but we hope that it will be the fastest and most capable computer on the market. . . . If we get enough orders at this price, we could go out of business. . . ."

Not to worry. According to one report, IBM's total loss on Stretch was a mere $20 million. While that may have seemed like a significant amount of money for IBM in the early 1960s, in retrospect the benefits that accrued to the company far outweighed any damage that may have been inflicted.

Indeed, the one damaged most may have been the man responsible for the Stretch development task, Stephen W. (Red) Dunwell, who became the scapegoat and, after the first customer shipment in May '61, was banished to a research position at Yorktown Heights. It wasn't until five years later that his contributions to the company came to be recognized. Thomas Watson Jr. made a public apology to Dunwell and awarded him the prestigious IBM Fellowship. But at the age of 62, after 41 years with IBM, Dunwell took early retirement.

Along with his wife, he took over Poughkeepsie's 110-year-old Bardavon Opera House, which was about to be demolished to make room for a parking lot. They breathed new life into the theater, made it a nonprofit, year-round operation, and turned it over to a full-time manager and staff. Now, Dunwell says, he and his wife have begun something they know a little more about—a timesharing business.

The machine's performance never reached the heights anticipated.

Dunwell recalls that IBM made eight Stretch computers, all in Poughkeepsie, but he could account for the whereabouts of only three or four. Most sources confirm that in addition to shipments to Los Alamos and its sister lab in Livermore, one went to the National Security Agency, one to the Atomic Energy Authority in the U.K., one to the Weather Bureau, and one to Mitre. Other sources believe the Dahlgren Naval Base got one and that the eighth went to the Argonne National Laboratory in Illinois.

When told that Fred Brooks credits him for the Stretch and calls him "the hero of the piece," Dunwell replies, "There were lots of heroes." Then he adds, "It was a heroic effort, I might say."

For the record, it must also be noted that some people at Control Data Corp. had different ideas about Stretch. In years past, they have said that the mainframe was really designed to keep CDC out of the superscale, scientific computer market. Perhaps one could call it a "knockout" machine, not so much Stretch as Smash.

GENESIS OF STRETCH

IBM alone could not afford to develop Stretch, so it asked the NSA to share the costs.

The genesis of the Stretch project, as can best be determined, seems to trace back to the National Security Agency and its need for more computing power than was available. It was easily determinable that such power could not be developed at an affordable price by using vacuum tubes, and yet it was equally obvious that an enormous investment would be required to develop the infant transistor technology. Unfortunately, IBM's policy was that the cost of such technology development had to be borne by the product for which it was incurred.

"In 1954 I believed that the only solution to that dilemma was to obtain support for early development work from an organization which could afford the new technology," recalls Stephen W. Dunwell. "Two of those organizations were the National Security Agency and the Atomic Energy Commission." In testimony presented at the IBM-Justice Dept. antitrust trial in New York City, Dunwell related how a group of engineers brainstormed the problem of overcoming the inadequacies of transistors and of manufacturing the types of solid-state devices required to build new and better computers. IBM management, including Dunwell, was then able to inform the NSA of what the company could do.

"That delegation," he recalls, "made it clear to NSA that IBM alone could not afford to do what was required and asked NSA to share in the cost of developing the necessary components. Dr. Solomon Kullbach, on behalf of NSA, agreed to do so."

It was late in 1954 or early in '55, he continues, that the folks at the Lawrence Livermore Laboratory in California asked for a proposal for the fastest computer IBM could build. A similar request also went to the makers of Univac computers, with Remington Rand winning that development contract. Disheartened but not deterred, IBM turned to Livermore's sister lab in Los Alamos, New Mexico, which expressed an interest in sharing the cost of developing the necessary technology.

"In January 1956 that computer became known as Stretch," Dunwell said in his testimony at the trial "and sometime thereafter was called the IBM 7030." The computer was designed jointly by engineers at IBM and senior scientists at Los Alamos.

The timing on this development project was very fortuitous. Had it been considered two or three years later, circumstances would have ruled against it, for it was a time when government procurement procedures were getting stricter. Harwood G. Kolsky, now at the IBM Palo Alto Scientific Center, says, "At the time the Stretch project was getting started, it was still possible for a major laboratory like Los Alamos to just enter into a contract," saying this is what we want and if you'll build it we'll buy it. "Two or three years later, they would never have been able to do something like that."

Ed Voorhees of Los Alamos, who was on the Stretch design team, would agree with that. "I always felt [the Stretch] was one of the best bargains the government ever got," says Voorhees. "But for some reason, efforts at Livermore and Los Alamos to undertake later development-type activities like this just got the cold shoulder from the AEC."

Kolsky, who was also on the Los Alamos design team before joining IBM, recalls that day when a group from IBM went to Los Alamos to make a presentation on the state of the computer art and the type of computer they thought they could build. It was Sept. 20, 1955, and the delegation was headed by Cuthbert Hurd. "They talked in terms of a 10-megapulse machine," he says, referring to the speed of the underlying transistors. Lloyd Hunter gave a presentation on magnetic cores. Dunwell spoke on machine organization, of the idea of having interlaced memories to compensate for the fact that the logic was much faster than the memory. They were talking about a two-microsecond memory and the final product ran at something like 2.25 usec, so the IBMers were very accurate on that technology forecast.

"A lot of the things that were discussed would best be classified as harebrained schemes."

"Their estimate on the transistors turned out to be optimistic," Kolsky recalls, "not because the transistors didn't switch in the times they thought," but because the long lines that ran from one frame to the next tended to slow the clock time. But he says one must understand that they were talking about something (the transistors) just out of the research stage and destined for a giant machine. "It takes a real act of faith to do something like that," he chuckles.

In January 1956 Dunwell was appointed manager of the Stretch development program, and the following November the contract was signed by IBM and the AEC/Los Alamos. It called for delivery in 42 months, which made it May 1960. Kolsky recalls the planning meetings. "A lot of the things that were discussed would best be classified as harebrained schemes," he says. "Somebody would come in and say, 'Why can't we do the following,' and they would spell out something or other, but it would have completely undermined the whole structure of the machine if you did something like that."

He adds: "The Stretch project attracted large numbers of fresh young graduates coming out of schools who had heard about the project and wanted to work on the biggest computer in the world. This is the sort of benefit to IBM that is hard to measure. I keep running into those people over the years. They slowly drift up into high positions in the company. They probably would not have joined IBM if it hadn't been for the Stretch project."

Among them, of course, was Frederick P. Brooks Jr., who says, "I went straight from Aiken's lab to work as an architect on Stretch." That was his first job at IBM; he went on to become principal architect of the System/360 and now heads the computer science department at the University of North Carolina. Of his experiences on the Stretch project, he says, "It was an exciting project. You had a chance to try everything you could dream up." He pauses momentarily and then adds, "And we did."

This apparently was true, for Kolsky says if he had to criticize the project at all it would be for the ease with which people were able to put features into the machine. Some of the features were later removed for the 360s, which Kolsky calls "a cleaned-up version of the Stretch." He says a look at the instruction sets on the Stretch and the 360 will show the strong family resemblance.

Recalls Jack Worlton of Los Alamos, who joined the design team about a year after its formation, "One of the beautiful instructions we put in there that we finally had to give up because it was so slow was branch-on-bit. You could pick out any bit in memory, examine it, and branch if it was a one or if it was a zero. But the trouble was that it took about five multiply times to accomplish this.

"Conceptually it was just a beautiful instruction," he continues, "but absolutely worthless. There are still computer designers who haven't learned that: you don't put too much complexity into the order set because it's difficult to build and maintain. If it's too complex, it'll never be used."

But Kolsky, like Brooks, is quick to heap praise on Dunwell for his management of the project. "His real genius was the fact that he saw where [IBM] should be five years hence and put together a project over the endless objections of everybody," Kolsky says. And when people came to him with technical problems, which they did daily, Dunwell would "turn them around and send them back out with the idea that 'yes, it can be solved.' " Until this project, IBM had moved cautiously, making evolutionary advances. But it was Dunwell who sought to make a factor-of-100 improvement in mainframe performance in one giant step.

Dunwell, of course, wasn't concerned only with the design features of the new computer, for engineers at IBM were also tackling the basic hardware technology on which the entire design would rest. The substitution of transistors for vacuum tubes was to reshape the system design of computers. It was like a new ball game.

It was necessary to design transistors suitable for use in computers. The solid-state devices of that time, Dunwell said in his testimony, "were neither fast enough nor had they the current-carrying capabilities necessary to control the ferrite core memories which would be needed." As in any pioneering role, it was also necessary for them to figure out how to manufacture such devices.

It was found, too, that some engineers just could not think in terms of the new solid-state technology, having been brought up on vacuum-tube devices. In an attempt to get them to redirect their thinking, Dunwell recalls, "for a time the laboratory expressly forbade anyone to have a piece of vacuum-tube equipment visible within his work area."

The design team under Dunwell tackled other problems. They had to come up with a new design for a power supply system, abandoning a 60-cycle system with transformers to go instead to a 400-cycle system with a motor generator. The back panel wiring, it was determined, was too complex to expect anyone to do the job correctly, so they got the Gardner-Denver Co. to make an automated wire-wrap machine. This device was driven by a punched card reader. And they went to Texas Instruments for the initial lot of transistors, IBM in 1956 having no such manufacturing capabilities. Burndy Corp. was the supplier for the tens of thousands of specially designed connectors needed for each computer.

"Up to that time," Dunwell said in his testimony, "all logical design had been recorded by draftsmen, but it was clearly out of the question to record the design of a machine of such complexity by manual means. A computer-generated design was necessary and a process for that purpose was developed."

When reached at his home in Poughkeepsie, where he retired in 1965, Dunwell said, "One of the fundamental things we were up against, having to do with manufacture and design, was that this machine was big enough and complex enough so we knew we would never get it together if we didn't automate the design and the manufacture." He explained that there would be so much wiring in the machine and it would involve so many drawings that they knew

For Frederick P. Brooks Jr., the Stretch project was what enticed him to join IBM.

THE STRETCH CLUB: IT NEVER CAME TO PASS

A nonprofit, cooperative venture by which a number of aerospace companies in Southern California could own and share in the use of an IBM Stretch system was given serious consideration, but never came to pass. The job of studying the economic feasibility of such a plan was assigned to DATAMATION editorial advisor Robert L. Patrick, who was just starting out as an independent consultant. It was given to him by Paul Armer, who was head of computer sciences at Rand Corp., the prestigious Santa Monica, Calif., think tank. Armer, of course, was until recently the executive secretary of the Charles Babbage Institute.

"Southern California was the hotbed of cooperative ventures in the early days of computing," explains Patrick. Many space companies could not justify getting their own Stretch, but it was thought that they could all use a piece of one. These were companies such as Lockheed, Northrup, Douglas, and North American, as well as Rand. About that same time, Herb Grosch was working for a large consulting company called CEIR Inc., which also announced its intention to install a Stretch in a Los Angeles service center. Both would have accomplished the same thing, allowing those same users to subscribe for some time on the machine, except that one was a commercial venture, the other a nonprofit arrangement. Neither plan was fulfilled.

In his study, with pages dated from January 1960 to June 1960, Patrick reported that the Stretch did not look attractive as a straight out-and-out computing engine, that it might cost the users more to run their workloads on the Stretch than on the machines they were then using.

"That wasn't necessarily negative as far as the Stretch was concerned," Patrick now explains. "It was just a phenomenon of the workload"—mostly short jobs. If the jobs had been weather analyses or forecasts involving massive partial differential equations, then the evaluation would likely have swung the opposite way, in favor of the Stretch. That, says Patrick, was the problem with sales of Stretch, as well as the superscale Control Data 6600 and its ilk. Everyone had shorts jobs. Only the atomic energy community had jobs massive enough to require a big number cruncher.

they would be forever changing drawings, correcting mistakes, and would get into a mad loop from which there would be no escape—unless things were automated.

When reminded of the book on Project Stretch, *Planning a Computer System*, Dunwell said, "One of the rules we had on that project was that nothing was done without first documenting carefully why it was done . . . so there was a great deal of documentation done, justifying the particular choices made, as we went along. And then abstracts were made from that for the Stretch book."

Fred Brooks remembers the book and especially a review of it by Lytton Strachey in the *Computer Journal*. In that review, Strachey said in part, "I get the impression that Stretch is in some way the end of one line of development. Like some early computer programs, it is immensely ingenious, immensely complicated, and extremely effective. But somehow at the same time crude, wasteful, and inelegant. And one feels there must be a better way of doing things."

Brooks, who used that quote in his own book, *The Mythical Man-Month*, says of Strachey's words, "I think that's an accurate assessment."

STRETCH OUT: Bill Ivie tried to get others to pull the switch on Stretch for the final time at BYU's September 1980 shutdown of the system. To the left with coat and tie is computer services manager Willard Gardner, and the bearded gentleman in the center is Joe Wise, who first found the Stretch listed in a government surplus inventory and campaigned for its acquisition.

STRETCH MARKS AT BYU

"A lot of people were betting money that we would never get [Stretch] operational."

"The prophets of doom were legion," says Gary Carlson, former director of computer services at Brigham Young University. Here was a small university nestled in the western foothills of the American Rockies, and it's about to take title to a 10-year-old Stretch computer installed near Boston. The intention is to dismantle it, move it to Provo, Utah, put it back together again, which would be no small feat, and get usable work out of it. No way.

"A lot of people were betting money that we would never get it operational," Carlson recalls. There was so much negative comment from his friends in the industry that he began to question his own decision. He figured it would cost the university about $50,000 to get into the game, just to see if it could be done. "So there was at least that much of a clear gamble on my part." But there appeared to be no alternative. It was 1970 and there clearly was a growing need for scientific computing capabilities on campus. "And we, like all universities, were always broke."

BYU had installed an IBM 7040 in 1963, and in 1968 installed a 360/50 that opened up computing on campus. It not only made it possible to provide computing services all over the campus but also got people interested in its applications. So, by 1970, two years after acquiring the mod 50, there was a growing demand for computing capabilities, and the Stretch would satisfy that need "at a price we could afford."

Carlson, of course, looked around to see what was commercially available, but found "the numbers were just mind-boggling." They looked into a 360/65, a Univac 1108, and a Burroughs 5700. He recalls all the prices were in the $3 million to $4 million range, and they couldn't afford that.

So how much did he figure it would cost to acquire the Stretch? "Well, Joe kept telling me that for a $5 registration fee we could get it." Carlson says Joseph L. Wise, manager of the scientific computing facility on campus at that time, was the primary instigator. Carlson had the final say and was supported by his assistant, Willard Gardner.

Wise says that in those days he regularly scanned government publications that listed surplus equipment. In one such listing, he saw an IBM 7094-II system, so he called a man in Washington with whom he frequently chatted about surplus gear. The man said, "Why do you want a 7094 when there's a Stretch system available?" Wise took the idea to Carlson, explaining that the system was available at no cost except for those related to shipping and reassembling. Whereupon Carlson is supposed to have said, "So what if the Navy wants to give me a battleship?" Wise says he still uses that rejoinder whenever anyone talks about getting something for nothing.

But Wise, sympathetic to the needs of researchers for computational power, was insistent. He talks of users who periodically needed four or five hours of 360/65 time and could get it only on Thanksgiving Day or New Year's Day. And there were some very large simulation runs on campus; one in particular he remembers ran on the campus Librascope L-3055 computer for some 150 hours. When asked if there were that many large jobs to be run, Wise explains that if the capacity is there, people come up with the jobs.

Gary Carlson recalls that in his presentation to the university's board of trustees he estimated the cost of getting and installing the computer at $100,000. It apparently sounded better to them than the several millions required for a new machine. He also inquired to see if IBM would maintain the Stretch and seems to think their quoted fee was almost $10,000 a month.

"You get a bargain now, but it'll eat you alive in operating costs," his detractors said. So one can imagine the reaction when Carlson told them he'd maintain it with his own people, plus a couple of students. As it turned out, BYU was able to get by with two full-time staff plus a few students.

The longest job BYU ran on the Stretch was a chemistry problem that lasted for 523 hours.

"Bill Ivie is the superstar of this whole show," says Carlson of his manager of operations at the Stretch center. Ivie assembled the machine and, for the final seven or eight years, kept it running. Willard Gardner, who succeeded Carlson as director of computer services, says the entire operating costs, including the salaries of Ivie and students and supplies, has been less than $100,000 a year. "So we've operated it for something less than people thought it would

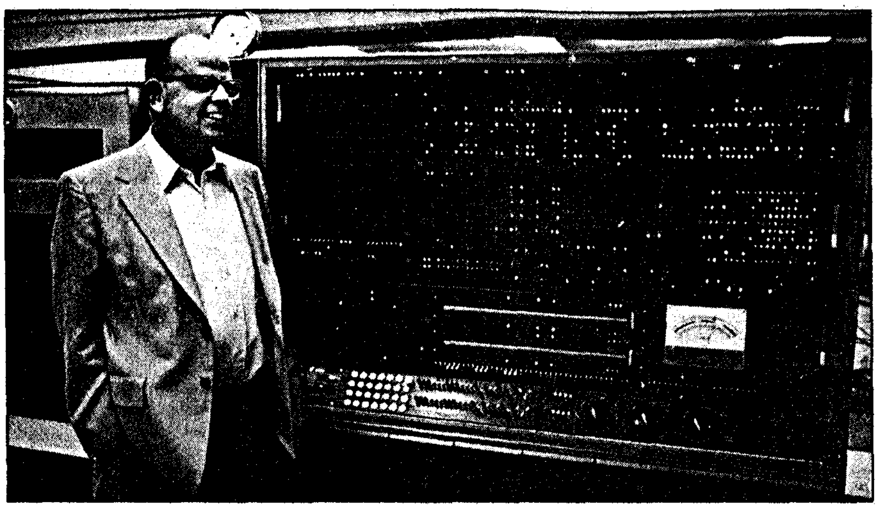

DR. GARY CARLSON, who approved BYU's Stretch acquisition in 1971, stands in front of the maintenance panel, which had more than 3,000 lights. About three or four of the bulbs burned out each day.

cost to maintain it,'' Carlson says proudly.

Approval to acquire the Stretch came in mid-March 1971. The dismantled hardware from Mitre Corp. in Lexington, Mass., had arrived by May, new false flooring at the site had been installed by July 1, and the main units had been reassembled and recabled by the end of July. It required some 14 months, however, before the first user job could be run as a test. (In mid-November 1971, BYU also acquired the Los Alamos Stretch, and it was soon cannibalized for spare parts.) A report by Wise in January 1973, about the time the installed system appeared ''capable of running programs in a general mode,'' shows the university's expenditures at some $165,000, including acquisition of the Los Alamos Stretch, site preparation, installation labor costs, and software development.

Bill Ivie recollects that there were some 8,000 pounds of cabling that came with the Stretch. The cables were so long that he just ran them up and down the length of the machine room before connecting them up to their destinations. According to Willard Gardner, Ivie and his crew got the hardware running long before BYU chalked up any operational time. What held back the initiation of service was the poor software and software documentation that came with the machine.

But Ivie says the longest job they ran on the Stretch was a chemistry problem that lasted for 523 hours. This was made possible by a facility developed at BYU that allowed everything in main memory to be read out onto tape, leaving the processor free to run just one job. When that job stopped, other jobs could be rolled in off tape to be run. Ivie says they ran a number of jobs that lasted for 30, even 40 hours, and some for more than 100.

This Stretch, believed to be the last one to be operational, was shut down by Ivie in September 1980. Replacing it for scientific computing on campus are a Digital Equipment Corp. PDP-10 and a VAX-11/780. Joe Wise says that although the VAX will run some jobs faster than the Stretch, ''for the large compute job I don't feel that a VAX is an appropriate replacement for the Stretch, even now.'' ✻