1983

Burton J. Smith: Architect of the supercomputer; Forecast: High touch

M. Daniel Rosen

John Naisbitt

Recommended Citation
Tempo, Vol. 28, no. 2 (1983), p. 34-36

Burton J. Smith

Architect of the Supercomputer

At eight-thirty in the morning, as is his custom on every working day, Burton J. Smith pulled up to his parking slot in his 1982 Jeep stationwagon. He cut the ignition, climbed down from the driver's seat, and headed for the low building ahead of him.

Around him were miles of brown-stubbed flatlands that led up to the Rockies in the distance. Once they were dotted with horse ranches; now they are the suburbs of Denver.

Oblivious to the scenery, Smith entered his office, gestured his visitor to a seat, and responded to a question about his role in today's history of the supercomputer. "I'm not an inventor; I'm an architect," he began. "I take a little bit of engineering, a little bit of art, some intuition and guesswork, and the tiniest bit of science—and I put things together in new ways."

He settled his bearish frame into his chair. "I work with existing materials, but I live in a world of abstractions and symbols. Everything I do goes directly from here," he said, pointing to his forehead, "to there," he said, pointing to the notepad on his desk.

At 42, Burton J. Smith, Sc.D., is vice-president, research and development, of Denelcor, Inc. And arguably, as principal "architect" of Denelcor's HEP 1 supercomputer, Burton Smith is among the most important people in the history of electronic computing since J. Presper Eckert and John W. Mauchly joined forces to develop the first practical electronic computer, the ENIAC, at the University of Pennsylvania in 1946.

What makes a computer a supercomputer is speed, not size. Supercomputers perform in millionths of a second (referred to as microseconds) what a typical home computer would take thousandths of a second to perform. To put a perspective on this sort of speed, consider the fact that the Denelcor HEP 1, operating at full capacity, can perform 160 operations per microsecond—and that there are more microseconds in a minute than there are minutes in a century.

These blinding speeds are faster by factors of unimaginable magnitude than any business computer on the market. The world's first supercomputer, a custom-built machine financed by the Defense Advanced Research Projects Agency, computed the dynamics of airflow around the engine of a rocket. The job took 18 hours. Machines are now in the planning stages that could do it in around 10 seconds.

The stakes are extremely high. As science writer William J. Broad recently pointed out in the *New York Times*, "No computer is now powerful enough, for example, to simulate the airflow around an entire aircraft, so aerodynamic designs are often put together in piecemeal fashion or by the repeated processing of two-dimensional slices. The first country with computers that can design a plane as a whole, according to a report by the National Science Foundation, 'will undoubtedly produce planes with superior performance.'"

Last year, a team of researchers from the Los Alamos National Laboratory and the Livermore National Laboratory came back from a fact-finding trip to Japan with a troubling report: backed by nearly a billion dollars in funding over a ten-year period, the Japanese are embarked on programs to produce supercomputers at least a thousand times more powerful than anything produced today. According to the U.S. investigating team, the Japanese see a need for large-scale computing capability in such areas as nuclear fusion, electrical power-system analysis, structural and thermal analyses of satellites, and the design and simulation of leading-edge semiconductor devices.

Breakthroughs in supercomputing are increasingly based on advances in "architecture," that is, on the way a computer goes about processing its problems; and this is the area in which Dr. Smith is making his main contribution. Cray Research and Control Data—the only other American companies making supercomputers—use conventional, sequential processing: one part of a problem must be completed before the next part can be started. Smith's supercomputer is based on a parallel architecture, in which the processing is done all at once.

To appreciate the difference, one must go back to the origins of electronic computing. "ENIAC, the first electronic computer, was a parallel design," Dr. Smith explained. "All the operations were carried out simultaneously. But the computer was programmed by plugging and replugging some 6,000 switches. The computer itself covered three walls, and any change in the program required a change in the switches.

"The change to sequential architecture came about through the work of John Von Neumann, the brilliant Hungarian mathematician. He found a way to store the programming inside the computer. At the same time, he changed the basic architecture of computers to accommodate a single processor and a single, huge memory—sequential architecture, in other words."

For more than three decades, despite staggering advances in technology, all commercially available computers but Denelcor's HEP 1 have followed the Von Neumann model. In the final analysis, conventional, Von-Neumann-style computers follow a plodding, step-by-step path.

One science writer puts it this way: "Suppose you wanted to add up corresponding numbers in two similar tables. The processor would be instructed first to add the first number in the top row of each table, go to the memory to get the data, add the two numbers, and store the result in the memory. Then it would start all over again, with an instruction to add the second number in the top row, and so on."

"The old Von-Neumann-style computing is something like a totalitarian dictatorship," Dr. Smith explains, "where one leader issues all the orders and the underlings carry them out in lockstep as quickly as they can. Parallel computing is more of a democracy: everybody has a task to perform, and they perform it by talking directly to each other, without consulting back with any boss. Why doesn't this end up in anarchy? Because we've programmed the workforce to be very conscientious."

To realize how this approach can speed things up, go back to the figures mentioned earlier. Suppose you had two tables of numbers, each with four rows across and six columns down, and you wanted to add the corresponding numbers in each table. Using a Von-Neumann-style computer, the operation would take 24 separate steps; using Burton Smith's approach —parallel processing—it would take only one.

Why was Smith the one to make the breakthrough? "Why not?" he responded with a grin.

Then he turned serious. "Remember, the idea was nothing new. And as a student, I had written my doctoral dissertation at MIT on certain aspects of parallel computing. There still were many questions to be solved in the abstract, and as a student and later as an electrical engineering teacher I continued to be fascinated by the concept.

"Why was I the one to develop a commercial parallel computer? I can't take sole credit. Here at Denelcor, at any given time, there were half a dozen people working with me on the project. The real question is, Why didn't someone do it before us? And the answer is, As long as computer technology was advancing the way it was, as long as quantum leaps were being made in the speeds of computer circuitry, as long as the costs of semiconductor chips were falling at the rates that they were, no one saw the need to design computers any differently."

Smith's opportunity came in 1973, when Denelcor signed him on as an outside consultant. He was an assistant professor of electrical engineering at the University of Colorado at the time, known for his background in advanced computer design.

Founded as the Denver Electronics Corporation in 1968, Denelcor produced a line of analog computers. Analog computers had their vogue in the 1960s, when their high speed and relatively low cost made them much in demand for scientific applications. But as time went on, their shortcomings outweighed their advantages: they still used old-fashioned, patch-board programming; they weren't nearly as accurate as their digital cousins; and they were extremely hard to keep working properly.

Smith's main assignment, then, was to come up with a successor to Denelcor's obsolete product line—a digital computer as fast as analog, competitively priced for scientific uses. A supercomputer, that is.

Development money for the project came from the U.S. Army's Ballistic Research Laboratory in Aberdeen, Maryland—the same organization that underwrote development of ENIAC.

HEP 1 (the acronym stands for heterogeneous element processor) was the outcome of this effort. The first multiprocessor system was shipped to the Ballistics Research Lab last December for use in the development of new weapons systems. By the end of 1982, Denelcor had installed a total of six processors—including one at the Messerschmidt-Bolkow-Blohm aerospace firm near Munich, West Germany, where it is being used in the development of new aircraft.

Smith paused to glance at his watch. It was ten o'clock, time for him and Anthony Casorso to head off for a meeting on the floor above. Casorso, a slender man with piercing green eyes and blond hair cut slightly over his ears, is Smith's only full-time assistant. "I hate being a manager," explained Smith. "I'm an architect, a designer. I don't even like managing things around my house."

The meeting involved some cost discussions concerning HEP 2, Denelcor's second-generation supercomputer. Om Gupta, formerly a research engineer at IBM's Poughkeepsie, New York, development lab, reported on some production prices he'd gathered.

Smith turns hard-nosed when matters are hard-nosed. "What assumptions are you making about production runs?" he asked. "Are you adding in the development costs?" And: "We need facts, not opinions, from those people you spoke to."

Soon the conference table in the office was piled high with notes and charts. There were legal-pad sheets filled with engineering sketches. Smith thumbed quickly through them, nodding in agreement with some things, shaking his head at others.

He stepped out of the meeting and immediately into another—a hallway encounter with Denelcor's vice-president of marketing, J. Philip Carley. After a quick exchange about ways to develop market-size estimates for the new supercomputer, Smith bolted around the corner to the office of Russ Saunders, Denelcor's vice-president of sales.

"There aren't any 'Great Moments of Discovery' here," Smith explained on the way back to his own office. "My key to innovation is being able to talk to somebody else about what I've got on my mind. That's why I spend a lot of my time popping in and out of offices, asking and challenging, bouncing off ideas.

"Many of my ideas come while I'm driving back and forth to work. I try not to think about my work when I'm home. It's important, I think, to turn it off—important for your well-being and the well-being of your family.

"When an idea does come, I write down an equation or a sketch in a ring-binder notebook. I've already filled two notebooks working on HEP 2. Hundreds of ring binders went into the development of HEP 1.

"But most of my thinking time is spent filling in the blanks. The biggest problem is finding people who can work well together. HEP 1 took five years to develop. With a product-development cycle as long as that, you've got to have people who can deal with give and take.

"The next big priority is freedom to experiment—freedom to make mistakes. Everything has to be judged on its merits. The worst thing to do is to take the first idea that pops into your mind and run with it. You should always look for other ways."

Smith's beginnings were anything but auspicious. The son of the chemistry department head at the University of New Mexico, he was, by his own admission, a classic "underachiever." His high school grades were mediocre, and he barely squeaked into Pomona College.

By the end of his freshman year, he was on academic probation. Back he came to Albuquerque for a last try at the University of New Mexico, but by the middle of his sophomore year he gave it up, dropped out, and enlisted in the navy. Four years later he returned and turned things around.

At the University of New Mexico, majoring in electrical engineering, he graduated *summa cum laude* in 1967. His master's degree and his Sc.D. were also in electrical engineering—punctuated by a National Science Foundation fellowship in 1968-69 and a National Science Foundation summer traineeship in 1970.

Next came a turn at teaching. "I never really considered a business career," he admitted. "All I wanted to do was to get out of Cambridge, come back out West, and do research and work with students."

Despite his crucial consulting role in the development of HEP 1, Smith resisted all entreaties to join Denelcor as a full-time executive for nearly six years. "It was a question of principle," he contended. "I wanted to make tenure before I left the University of Colorado, and I held out until I did in 1979."

Later, over lunch, he spoke of personal things.

"What motivates you?" he was asked.

"If I picked one thing, I'd probably be misleading. It's certainly not money. I'm comfortable with my lifestyle. It's fame of a sort, I suppose, being known by my peers. But mainly it's making a contribution to science. I think if that aspect were missing, the rest couldn't make up for it."

Then he paused and the corners of his mouth crinkled upward. "And it's fun," he said. "I enjoy what I'm doing and if I ever got to the point where I didn't, I think I'd quit." He has another reason to try: a wife and two daughters. "My family's as much fun as my work," he said.

"How far are you from reaching the full potential of supercomputers?" he was asked, as the check arrived.

"I don't think we'll ever reach it," he answered. "There's a saying in this business that a computer is always ten times slower than the problem you're trying to solve. The appetite will always be bigger than the size of the meal on the table. There will always be people who need to solve big problems—and people like me who get a kick out of solving them."

—M. Daniel Rosen

FORECAST/ High Touch

The scenario for sitting at home with our computer and tapping out messages to the office is a very limited one. My line about all that is that it's okay for emergencies, like Mondays. But for the most part people want to be with people. We'll be going to the office, going shopping, going to movies, going to restaurants, getting together for any excuse as the counter-balance to an incredible amount of technology in society. □ What's resulting is that we're more and more examining our own humanity as we more and more technologize our communities.

JOHN NAISBITT
Author, Megatrends