BEYOND VAX: A CONVERSATION WITH GORDON BELL

Gordon Bell, 53, is a legendary figure around the halls of DEC, even though he last worked there four years ago. As head of DEC's engineering effort in the 1970s, Bell formulated the company's VAX strategy and shepherded its introduction and implementation. That strategy remains the foundation of DEC's product line and marketing efforts. The company insists this strategy is flexible and durable enough to take DEC into the next decade and beyond.

Computerworld Extra asked Bell, now assistant director for the Computer and Information Science and Engineering (CISE) Directorate at the National Science Foundation in Washington, D.C., to consider the future of VAX.

What is the origin and essence of VAX?
VAX came from a tiny task force I led in April 1975. The idea was to create a new computer family to be "culturally compatible" with the successful PDP-11. Its principal design goals were to be compatible with key operating systems and languages; to have a much larger address space than any existing computer; to be efficient at implementing high-level languages, including Fortran, C [for Unix] and Cobol; to be implementable over a wide range of sizes; and simply to be the highest performance computer in its class when first implemented.

In December 1978, after the VAX-11/780 had achieved immediate success, the company adopted the VAX strategy to provide a VAX homogeneous computing environment for a range of interconnected computers.

A user could compute in any of three styles from a cluster of large machines behaving as a single system, distributed traditional minicomputers and distributed clusters of workstations. The strategy also specified compatibility with other DEC computers and intercommunication with other standards and products.

Why has VAX been so successful?
The concept was incredibly simple, and hence everyone [customers and the company] could understand and support it. Also, the three-level computing hierarchy was right . . . even IBM discovered and endorsed it by the early 1980s. VAX provided the best, and only, totally compatible, single-interconnection environment. This required a range of computers, from VAX on a chip to the highest performance computers that could be built.

VAX gave DEC a product monopoly, since no other manufacturer has anything like this capability. It specifically exploited the fact that most manufacturers had a menagerie of product lines designed to segment the user base, fill product size and application gaps or help the manufacturer's organization.

Recently, IBM started to provide similar capabilities by having [IBM] 370-compatible minis and a plug-in card for a PC. But this is not enough because they have several operating systems, a worse problem than having multiple hardware architectures.

Also, given the complexity of the IBM architecture, including the I/O and operating systems, it's probably hard to make the architecture serve the wide range of users at this point in its life.

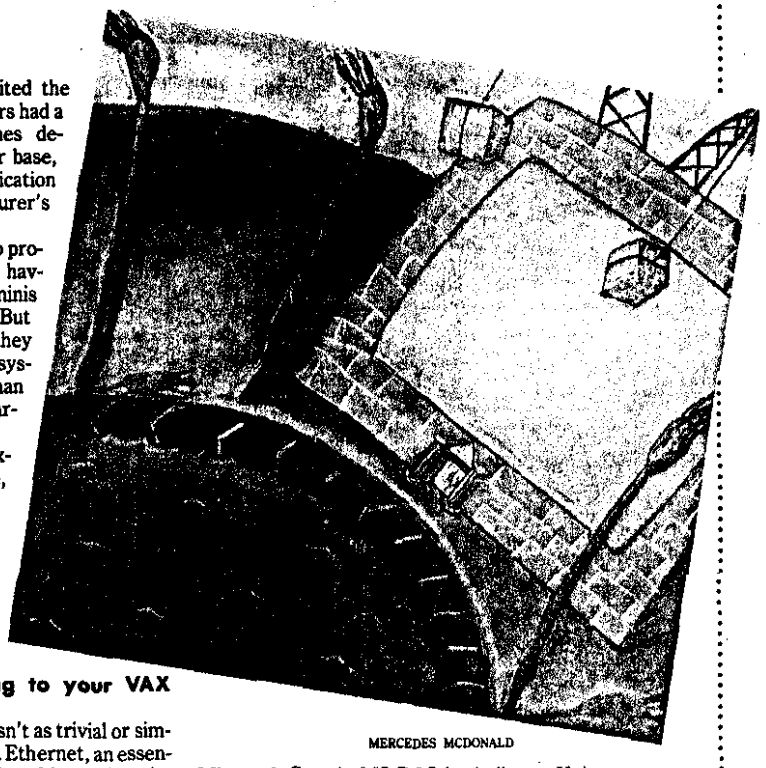

MERCEDES MCDONALD

Did things happen pretty much according to your VAX strategy?
Largely yes, although it wasn't as trivial or simple to do as one would think. Ethernet, an essential component, was questioned by various internal DEC committees, even after the whole system was working. Having adopted a VAX strategy in 1978, the company in 1980 decided it had to enter the PC market with a trilogy of non-VAX PCs, which only loosely fit the strategy. VAX was too large to build as a workstation until 1982 to '84.

At the same time, the high-end implementation of VAX — Venus [the 8600] — was more than two years late as engineers hit the complexity wall and essentially forgot the recipe of how to design computers. These two events accounted for DEC's poor financial performance in the early '80s.

Do you see anything that could challenge the VAX strategy yet?
No. In 1978, I thought the only possible threat was Unix, because it provides compatibility at a higher level, somewhat like VAX. I imagined that innovative or small companies would develop Unix systems for interconnect computing environments by the mid-'80s. Now I'll push that back three to five years.

A critical hole is in the PC space where Microsoft Corp.'s MS-DOS is similar to Unix, but isn't compatible. Unix needs to evolve in range, human interface and applications. Having AT&T control it doesn't help — it has to truly be a public standard. The government support of Unix [Posix] still could have an impact.

Also, I don't see a single large computer company coming up with anything like VAX because of the cost and commitments of preserving their code museums for running old programs.

How far can DEC go with the VAX architecture?
I don't believe all the capabilities in the architecture, as constrained by its addressing, have been exploited yet. DEC still has uniqueness.

Critics point out that even if clustered, the architecture will sooner or later top out. What do you think?
Here, history is a good guide. Every architecture has sooner or later either run into a limit or been inappropriate to the technology. With the

Continued on page 78

BEYOND VAX

FROM PAGE 80

exception of the IBM 360, which had an inherent 32-bit address to get it to the 1990s, history also tells us that companies try to evolve their architectures for too long. They end up with 100% of their user base market but a declining share of the entire market. Eventually, even that market declines as users desert the obsolete machines.

A major computer technology generation lasts about a decade. I believe it is hard to design an optimum architecture that lasts much longer. While VAX may top out, it should be a fine base for evolution.

What will the topping out of the VAX mean to the thousands of sites committed to DEC's single-architecture prescription?

Again, let me rely on history. VAX was a major new hardware architectural evolution from the PDP-11, yet it preserved programming interfaces, languages and data bases. The same concept could be reapplied even if DEC changes the underlying hardware architecture. The program and data base interface must be preserved — in effect, it should be transparent if users adhere to certain VAX and VMS standards.

It's probably important to define VAX, or VMS, compatibility and whether a new, basic hardware architecture could be used to implement this environment — that is, without object compatibility. The problem is much easier than with the PDP-11 or with the 370 because VMS is a single interface which subsumes the network but includes the command language, DCL, and various languages. Fortunately, nearly all programs are written in a high-level language today and would be compatible.

"RADICALLY new applications should be sought that build on the environment and do things no other environment can support."

Can DEC engineers develop a totally new architecture for the '90s and beyond that will play on Decnet and run software from existing machines?

Companies with different underlying hardware architectures provide existing proof of VMS user-level compatibility. Certainly DEC should be able to do this, too.

Are the engineers at work on such a scheme now, and if so, what is the best guess as to what it is?

I hope so, because it is essential. The key is to identify the crucial limits of VAX and to eliminate them. Again, it might have goals similar to those we used to create VAX in the first place. The only goal I would add to the original VAX set would be the independence of the Instruction Set Processor hardware architecture. Just as VAX added new dimensions of comparison, a new architecture also must add those dimensions of comparison.

I would hope a new plan would address parallelism of all forms and performance for the scientific and engineering community, including the ability to collaborate effectively via the computer using high-speed interconnects. It would handle large scientific and engineering data bases.

A radical view of data integrity and data bases is also needed. Improvements in the cost of ownership and availability dimensions are quite possible. In addition, DEC could address the mass market for users who want a great computing environment but don't want to become system programmers or administrators. This would rule out any compatibility with MS-DOS and the IBM PC. The PC has allowed everyone to relive and retrace computing history and to become system administrators with all the accoutrements, including large manuals. I'm happy to avoid this trip back to the '70s; I use an Apple Macintosh.

You've made several comments about needing higher performance VAXs. What is the biggest VAX you can build?

There are two basic measures of performance: total processing power available to a single job stream — in other words, throughput; and power available to a single job. For the former, Vaxclusters partially provide this power, but multiprocessors extend the range even more and in a more cost-effective fashion. Furthermore, multiprocessors are starting to use parallel processing to provide speedup of a single job, which can be done either by the compiler or the user.

DEC should have already introduced a significant multiprocessor with dozens of microprocessors, a "multi" like those from Encore, Masscomp, Sequent, Stratus and others. VMS as a multiprocessor operating system shouldn't be the limit. By using the CMOS Microvax, more than 100 MIPS [million instructions per second] could be put in a small box. This approach would provide at least one or two non-"me-too" products. Moreover, it gets the price into the $10,000 per MIPS range vs. the $100,000 to $200,000 range typical of the large mainframe. These ridiculous prices aren't sustainable except for large mainframes, where users are locked into buying code museums — and someday the users may get smart.

What markets would such a machine address?

DEC seems enamored with the commercial and transaction processing markets. Multis are the best computers for these markets because the applications only demand total MIPS for a large collection of jobs. The system has advantages for a general, interactive job stream such as program development as demonstrated by the multi suppliers. The microprocessor inherently provides the best cost/performance by almost two orders of magnitude — we simply look at the MIPS per chip.

By ganging them and matching them to a memory, one can get the most power in a single system at a small fraction of the cost of an emitter-coupled logic-based computer with a few expensive processors. It also offers, inherently, much better availability characteristics.

Would such a system address all of your concerns about inadequate scientific and engineering performance?

Not entirely, but two multis could replace an entire product line and provide 100-MIPS-level performance and substantially better price/performance for the user than the current "model" approach.

In addition, consistently competitive compute servers are needed, which would run technical work in the [Cray Research, Inc.] Cray-1 speed range. In the long run, a multi might do the job, but for now, the vector multiprocessor is the main line . . . in effect, another Crayette.

What is the largest uniprocessor VAX that can be built?

The speed of a uniprocessor, such as the VAX or a 370, is correlated with the clock speed.

A high-end machine with a 40- to 60-MHz clock could probably be built and still be in the mini price range with a power of two to three times the current models.

Note that the current [IBM] 3090 uses about a 60-MHz clock, and the Cray XMP clock is almost twice as fast, although both have roughly the same scalar speed. Clock speed isn't always a good indicator.

Could you look retrospectively on what DEC might have done in the four years since you left?

Let me provide my own reference point first.

I've been involved with a bunch of new computers, three of which are on the market, plus several start-ups that are creating new markets. All of the computers provide more capabilities than VAXs, and the engineering has been done in a small fraction of the time and budget of DEC product designs. Aside from the evolutionary extensions and products, I would have probably urged for greater innovation and carried on enough experiments to have selected a VAX II architecture by 1986, with benchmarking now and delivery in '88 — a decade after the 780.

"I would have probably urged for greater innovation and carried on enough experiments to have selected a VAX II architecture by 1986."

Would VAX II be reduced instruction set computing, or RISC, based?

Probably.

What other big issues face DEC in the future?

1. Thinking VAX is the end, not simply the best thing around today, is an enormous hurdle. While nothing is yet in the marketplace to challenge it, several new systems do and will. This thinking leads to arrogance. VAX gave DEC a monopoly in much the same way that the 360 gave IBM a product monopoly by the '70s that only lasted a decade. DEC should compare its products with the best small companies, not old-line suppliers.

2. Thinking VAX can do it all by itself . . . or by relying on the old IBM applications being converted to VAX just because VAX is better than the 360. Radically new applications should be sought that build on the environment and do things no other environment can support. Also, understanding the limits that come from new uses is critical to VAX II.

3. Being enamored with the commercial interests and not attending to the scientific and engineering base, especially in the universities. The commercial market tolerates high prices for higher performance, but they are unique. The technical marketplace is far more demanding on products.

4. Poor presence on the desk, and even picking MS-DOS and [Intel Corp.'s] 80286 or 80386 to implement. I don't see what another clone brings to the marketplace — certainly not profitability. Service revenue can be obtained simply by going into that segment of the service business. Integration with the Apple Mac is also important.

5. Responsive, efficient and creative manufacturing still appear to be nonexistent. While DEC is probably no worse than the average, it's not adequate to compete in the '90s when the Japanese and others arrive.

That's a big set of worries. Are you optimistic?

Certainly. They are making lots of money, have lots of cash and exceptional people. All they need is a challenge. The plethora of new start-ups certainly provides that.

THE VAX FAMILY

In the past year, DEC has introduced seven VAXs (performance figures are based on the VAX-11/780 as the company's unit of measure; DEC does not use MIPS)*

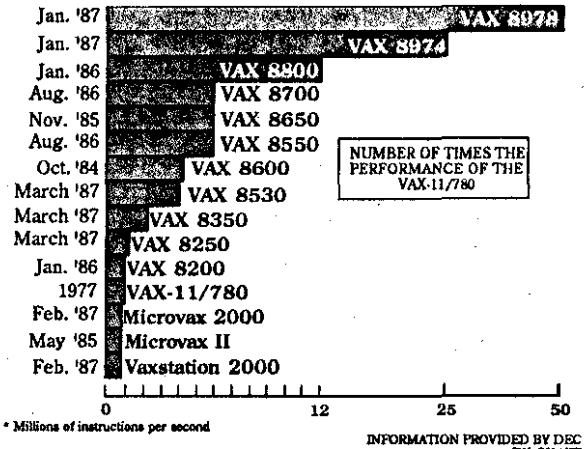

* Millions of instructions per second

INFORMATION PROVIDED BY DEC
CW CHART