THE EDITOR'S VIEW

The Death of the Superprocessor

Cray, Chen End Quest to Build World's Fastest Processors

The news seemed ironic, even twisted. Intel plans to build the world's fastest supercomputer from the same microprocessors it expects to put into next-generation PCs. Not too long ago, Intel's x86 microprocessors were regarded by scientists as little more than toys, while the giants of computer architecture built exotic designs from extravagant technology to produce the fastest processors in the world.

One such giant is Seymour Cray. His name became synonymous with supercomputing as his eponymous company developed widely used systems from the Cray-1 to the Cray-YMP. These machines were built around "superprocessors" that were far faster than the microprocessors of their day. These superprocessors used bipolar logic to achieve high clock speeds and combined masses of discrete chips to create fast, wide computing units and fast, wide buses. Liquid circulation systems kept these processors cool. In a world of single-processor systems, the fastest CPU won.

As its name implies, the Cray-YMP linked a handful of superprocessors together to generate even greater computational horsepower. This move kept Cray ahead of others who were also linking processors together; the performance advantage of the YMP still came from the power of its individual processors.

What Cray didn't anticipate was that this trend would continue, as others linked dozens, then hundreds of processors together in a single system. Cray stubbornly pursued his lifelong task of improving the performance of his superprocessors, eventually leaving his own company in 1989 and starting Cray Computer, as his former compatriots carried on at Cray Research. At the new company, he built the Cray-3 and began work on the Cray-4, using expensive gallium-arsenide chips built in the company's own fab. But in its entire corporate lifetime, Cray Computer never sold a single machine. The company finally closed its doors last April, and Cray himself is in retirement, at least for now.

Cray's most prominent protegé was Stephen Chen, who eventually broke away to form his own company, SSI (Supercomputer Systems, Inc.). With backing from IBM and others, Chen and his team continued to develop superprocessors. But after a few years, Chen also ran out of funding before selling any machines.

The rage in the supercomputer world is parallel processing. Superprocessors are not fit for large-scale parallel processing, due to their cost and size. (A four-processor Cray-YMP would take up most of the average living room and consumes about $10,000 of electricity each month.) Instead, most vendors today rely on cost-effective microprocessors. Although a single microprocessor may provide only half, or less, of the performance of a superprocessor, the aggregate performance of a conglomeration of small chips is well beyond what a handful of superprocessors can produce.

For example, the new Intel machine, which was commissioned by the U.S. Department of Energy for nuclear weapons simulations, will execute up to 1.8 trillion operations per second using 9,000 Pentium Pro (P6) processors. Intel beat out rivals such as Cray Research and IBM for the $46 million deal.

These competitors are taking microprocessor-based approaches as well. Cray Research's T3D systems each contain up to 64 Alpha processors *(see **0715MSB.PDF**)*. IBM's supercomputer is based on 32 Power2 processors. Silicon Graphics has made a big splash in the scientific market with its 32-processor Challenge system, powered by R8000 CPUs. Convex, soon to be a division of HP, builds parallel PA-RISC supercomputers.

Chen has resurfaced with a new company called SCI (SuperComputers International). He, too, has adopted the new model, building servers based on Pentium Pro. All these microprocessor-based systems can provide high performance at a relatively low price, due to the efficiencies of using mass-market chips. They can also reach the market faster, as superprocessors usually take longer to design and verify than microprocessors.

A few companies are continuing to design superprocessors, including NEC's supercomputer division and startup Tera. Some of the so-called Grand Challenge problems remain suited to these types of machines, although the market for them is small.

Advances in parallelizing compiler technology and other new techniques have allowed many traditional supercomputer problems to be distributed among lots of inexpensive processors. What Eugene Miya dubbed "the attack of the killer micros" has overcome most of the superprocessors. Their spirit lives on in muscular microprocessors like the PA-8000 and Power2. ♦

Will microprocessors completely take over the supercomputer market, or is there room for superprocessors? Send your opinion to editor@mdr.zd.com.