

HOW-TO

Chip Multiprocessing

By Alan Joch

Computerworld |

NOV 27, 2000 12:00 AM PST

Early next year, top-echelon power users will find out if two heads are better than one. That's when Sun Microsystems Inc., IBM, Compaq Computer Corp., Hewlett-Packard Co. and others will start to roll out high-end servers that take advantage of chip multiprocessing (CMP), a step forward from current systems that load up boxes with multiple discrete chip modules. (Conspicuously absent from this group is Intel Corp., which is betting on speed-enhancing instruction-level parallelism, a less-costly performance booster.)

"The driving force here is, rather than create more complicated processors, why not just put two in the same module?" says Linley Gwennap, principal analyst at microprocessor consulting firm The Linley Group in Mountain View, Calif. "The real problem is, unless operating systems really understand you have pairs of processors, it's not clear you'll see a big benefit." He adds that over time, operating systems will become multiprocessor-chip-savvy, but the programming hurdles will be difficult to overcome.

Early tests are showing that two processors in a single module outperform multiple discrete processors by 50% or more. By putting two CPUs on a single piece of silicon, engineers can take advantage of shorter distances and faster bus speeds when shuttling data between the two CPU cores. The net performance result for IBM's version - called the Power4 processor - is the ability to process 100GB of data per second, or the equivalent of 20 full-length DVDs, says Joel Tandler, director of technology strategy at IBM's server group in Austin, Texas.

[Buyer's guide: How to choose the right business laptops]

Data-crunching like that will likely come along with astronomical system prices - hundreds of thousands to millions of dollars - that will send CMP systems straight to the high-end technical and commercial market. This includes machines that process seismic data for oil exploration companies, e-business servers able to handle unpredictable traffic loads and spikes, data-intensive graphics imaging hardware and computers that crunch genomic data.



Not every high-end application is right for CMP, however. Financial batch-processing programs that sequentially march through a ledger one task at a time will still rely on single-processor systems.

The age for commercially viable CMP systems has arrived, thanks to continuing refinements in chip manufacturing techniques that let engineers pack circuits more densely. The extra die space opens up room for multiple chips - two in the initial systems but perhaps up to eight in later generations of CMP modules.

But simply squeezing two chips into one housing doesn't necessarily create an efficient multiprocessor. The biggest challenge to engineers is keeping these two-headed power plants stoked with data, and this is where some of the biggest design differences will surface among chip vendors.

Design Contrasts

Due out next quarter, Sun's MAJC-5200 module will include two 500-MHz CPUs, a graphics preprocessor and a data-transfer engine. Data rates for peak I/O will be 4.8GB/sec. The processors will share a 16KB four-way, set-associative data cache, and each CPU will also have its own 16KB, two-way, set-associative instruction cache.

An additional wrinkle in the MAJC-5200 will be multithreading: The hardware will be able to divide processing tasks into bite-size chunks that flow in an orderly way to each core to avoid any missed processing cycles. But Marc Tremblay, chief designer in Sun's processor product group, acknowledges that many software applications aren't optimized for multithreading. To compensate, the MAJC-5200 will use the Java Virtual Machine to speculatively generate threads in Java programs.

In contrast, IBM has chosen not to implement multithreading in its higher-speed, 1-GHz Power4 chip. To keep data flowing efficiently, the Power4 will cram 32MB of memory per chip into a second- and third-level cache to keep chip-to-chip communications flowing and to buffer information retrieved from system memory.

"The challenge you have with all systems when you increase frequency is that memory appears to be farther away," Tandler explains. "[The processing] cycle at 500 MHz is 1 nanosecond; at 1 GHz, it's 2 nanoseconds. So we're adding additional caching to the storage hierarchy."

Bus Speeds Beyond 1 GHz


Initial Power4 systems will use bus speeds of 500 MHz - or half the processor frequency - although the systems are designed for bus speeds greater than 1 GHz to anticipate rising processor speeds over time. IBM expects to release its CMP processor in the second half of next year.

Sun and IBM agree on one thing: Chip multiprocessing is the next big thing in CPU design. In the past, engineers squeezed out faster performance with better process technology, better microarchitectures and better compilers, Tremblay says. Chip multiprocessing adds another performance tool that directly addresses the Holy Grail of processor evolution.

"This could take us beyond Moore's Law," Tremblay says. That's great news for power users but one more challenge for software programmers.

Joch is a freelance writer in Frankestown, N.H. Contact him at ajoch@monad.net.

We're sorry.
This image is no
longer available



[Click here to view a larger chart.](#)

Copyright © 2000 IDG Communications, Inc.

Bing's AI chatbot came to work for me. I had to fire it.

SHOP TECH PRODUCTS AT AMAZON

SPONSORED LINKS

Cloud and Enterprise SSDs: Performance Tuned for Real-World Workloads

dtSearch® - INSTANTLY SEARCH TERABYTES of files, emails, databases, web data. 25+ search types; Win/Lin/Mac SDK; hundreds of reviews; full evaluations

