EETIMES

www.cmpnet.com
The Technology Network

Fiber weighed for chip interconnect

By Anthony Cataldo, EE Times
Oct 12, 2001 (12:27 PM)
URL: http://www.eetimes.com/story/OEG20011012S0061

SAN MATEO, Calif. — Semiconductor researchers and at least one startup are going out on a limb, eschewing copper interconnect to bring optical fiber directly to the microprocessor. Measuring distances in centimeters instead of the kilometers associated with fiber optics in telecommunications, this select group points to a promise of ultrafast data rates over a clean, low-power and low-noise pipe. It's still unclear when — or if — major MPU vendors will sign on, but Intel Corp., for one, is among the chip houses mounting research programs into fiber.

Proponents of fiber-to-the-processor don't discount the barriers the microprocessor industry will face to bring down the costs of fiber optics. And their claims that copper interconnect is about to run out of gas are debatable.

Yet the nascent interest in fiber reflects a sense of urgency within the chip industry to bolster processor I/O, which has emerged as a major bottleneck to system performance. This has become apparent as the microprocessor players rally behind a bevy of high-speed interconnect schemes, such as Infiniband, Rapid I/O, HyperTransport and Arapahoe.

Some argue that in a few years' time even those electrical interconnects will be pushed aside in favor of pumping light waves — which offer low crosstalk, zero electromagnetic interference and high bandwidth density — directly into a processor package.

Startup Primarion (Tempe, Ariz.), which gets some of its funding from Intel's venture capital arm, has set out to develop optical ports that would sit between a processor and an optical fiber, servicing MPUs approaching 10 GHz.

"A 10-GHz core is eating data at 30 to 60 Gbytes per second," said Bill Pohlman, chairman and chief technology officer of Primarion. "But if you can't feed it, you can't get the value out of the core. Even with simultaneous multithreading, the system is out of balance."

To Pohlman, fast electrical I/O schemes like low-voltage differential signaling will hit the wall once processors surpass 5 GHz. "It's a question of reducing the cost of optical. We see the crossover at 5 GHz," said Pohlman, a 19-year Intel veteran who retired as vice president of that company's microprocessor group.

Pohlman will describe Primarion's vision of fiber-to-the-processor next week in a keynote address at the Microprocessor Forum in San Jose, Calif.

Using fiber for short-distance interconnect has also caught the interest of researchers at Intel. The company has formed a team that works with several universities to investigate, for example, how to run a laser through a circuit board. "Optical is something we need to be investigating," said Wilfred Pinfold, technology director at Intel's microprocessor research lab. "There comes a point at which you are looking for a crossover."

Others see a more immediate need to kick-start the move to fiber. Tony Levi, a professor of electrical engineering at the University of Southern California and technical adviser to Primarion, has written a paper describing how microprocessors will need a fiber pipe to counter excessive power dissipation and main-memory bottlenecks. Levi contends that system architects will need to emphasize high-speed I/O and rein in the tendency to add more transistors to a chip.

"One problem is [that] we're burning too much power. The other is, there's a disparity between processing data and the inability of main memory to deliver it over the front-side bus," Levi said.

To Levi, electrical interconnect is running out of steam. Faster electrical I/Os will require a lot of controlled-impedance lines that need to be terminated, increasing power consumption. And the propensity to use microstripped lines in FR4 circuit boards will make it harder to overcome crosstalk and noise, he said.

One of the most compelling reasons to move to short-distance optical interconnect, observers said, is to boost a processor's bandwidth to memory, which has been cited as a weak link in system design. While high-speed main-memory devices are expected to deliver 6.4 Gbytes/s in a few years, the microprocessor industry needs to start thinking about a more radical solution, some say.

"You need to be in the range of 40 Gbytes/s to satisfy what MPUs are going to be doing in the next little while," Levi said. "The ultimate solution is to get away from a bus-based architecture and go to a switch-type solution."

Fiber optics should have no trouble meeting these bandwidth needs, say its proponents. Parallel fiber-optic modules in use today can deliver 3.7 Gbytes/s every centimeter, and are on track to provide 150 Gbytes/s/cm using wavelength-division multiplexing. Moreover, fiber has low crosstalk and takes less power at high frequencies than electrical interconnect, according to Levi.

Levi has proposed an "encapsulated processor" concept whereby a CMOS device uses fiber-optic ports as the only connection to external chip sets and DRAM. The processor, which itself could contain two CPUs and cache memory in the core, would integrate a crossbar switch that connects the ports to the processors and cache memory.

The ports, each of which could sustain 40 Gbytes/s of data throughput in each direction, decode and multiplex signals for an optical subassembly containing vertical-cavity surface-emitting lasers (VCSELs), PIN receivers and the fiber interface. There would also be a short, low-power electrical link from the port to the processor, according to Levi's proposal.

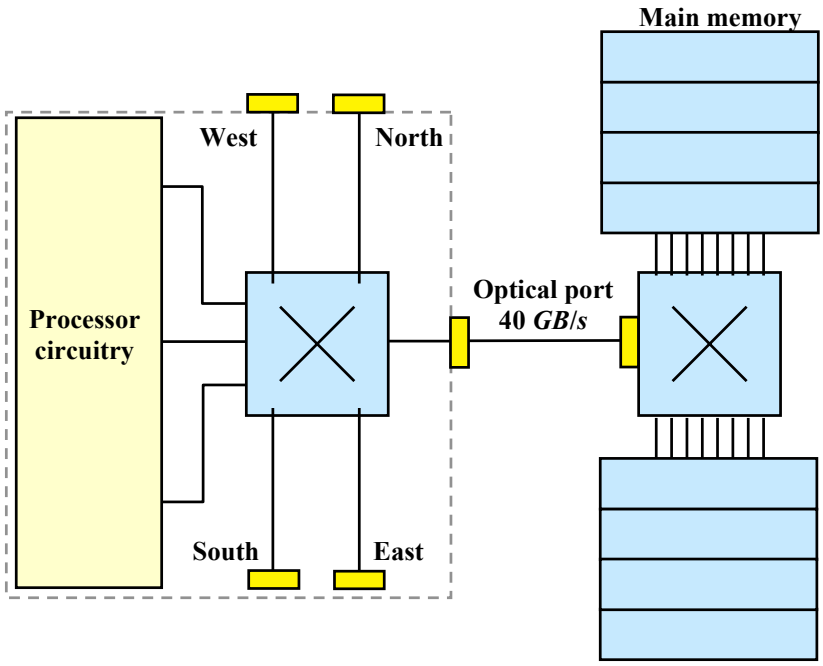

The ultimate chip-to-chip interconnect?

Optical fiber link provides 40 GB/s to main memory.

At DRAM speed

Main-memory access stands to improve dramatically under this scheme. The memory subsystem could have its own processor and pipelined translation-lookaside buffer. A pair of memory banks could feed into a crossbar switch over point-to-point, half-duplex electrical links running at 2 Gbytes/s each. The aggregate bandwidth would then be sent over an optical link to the processor at 40 Gbytes/s in both directions, more closely matching the internal speed of the DRAM core.

"If you look inside a DDR [double-data-rate] DRAM itself, the internal burst rate is 32 Gbytes/s. The problem is getting the data from the memory to the processor," Levi said. "The way this works is, you can access many banks of memory and aggregate them to the processor."

Today, larger cache memories are being exploited to help solve the memory bandwidth problem. Intel, for one, plans to absorb a level-three cache into its future 64-bit McKinley processor. The problem here, some say, is power consumption, which has emerged as a major concern for MPU makers in recent years.

Cache can also be an expensive proposition. "To make the caches bigger and operate out of cache helps to a certain extent, but there's only so much silicon you can put on a chip before it becomes cost-prohibitive," said Cary Snyder, a senior analyst with MicroDesign Resources.

Not only does cache memory take up vast amounts of space on the die, it doesn't necessarily solve the memory bandwidth problem. Latency that occurs when a cache miss forces a processor to fetch data from main memory only gets worse the faster you run the processor.

"If you take 10 GHz and extrapolate the bus, it will take hundreds of clock cycles to refill a line of cache," said Pohlman of Primarion.

And while electrical I/O is turning to complex schemes like low-voltage differential signaling to eke out more bandwidth, it's finding it difficult to keep up with microprocessor speeds.

"One thing really attractive about optical is that it will scale with Moore's Law, whereas electrical is out of gas," Pohlman said. At 10-GHz processor speeds, "if electrical can go a couple of centimeters it will be a miracle," he said.

But it may not take divine intervention to get more mileage out of copper interconnect. Intel claims it can reach speeds of 10 GHz and beyond in five to eight years using copper. "We're confident we can get to 10 GHz. And there's reason to believe we can double that," Pinfold said.

Yet he admitted that bringing fiber to the processor would be desirable if the cost came down. "If there's a real breakthrough in optical, we may be able to bring it in sooner," Pinfold said.

The biggest barrier in terms of cost is making the conversion from optical to electrical and back again, as would be needed on a processor package or module. But proponents contend that the advent of VCSELs, lower-cost interface electronics and lower-cost packaging is rapidly driving down the price tag for the conversion. "These are going to be 25-cent items in the near future once the volume is up there," Polhman said.

VCSELs represent a big step toward lowing conversion costs, because they can be tested at the wafer level and take advantage of the manufacturing efficiencies of batch processing. Beyond that, Primarion says it has come up with a way to use plastic semiconductor packaging technology for laser-cum-electrical components, which now require a technician to peer through a microscope and align the laser fiber by hand. Primarion said it expects to introduce the technology by next year. "This is not something that's out of the mainstream. Any assembly house in Asia would say 'Yes, I can do this,' " Pohlman said.

The move to layered high-speed I/O standards, in which the physical layer and transport protocols are separated, should make it easier for standards groups overseeing Infiniband and perhaps Arapahoe to wrap in optical. "It's just another media type. You won't have to change any software," Pohlman said.

Fiber-to-the-processor proponents claim their proselytizing is starting to pay off. Pohlman said he knows of two major processor vendors that have started optical I/O programs.

USC's Levi, meanwhile, said he is trying to build a coalition of companies, universities and government agencies to build up the fiber infrastructure. So far he said he's had discussions with officials at the Defense Advanced Research Projects Agency, which mentioned fiber-to-the-processor as one of the possible applications for a chip-scale wave-division multiplex project it recently kicked off.

But the real task ahead will be to convince market movers like Intel that backing optical interconnect for chip-to-chip communications is the right move. For now, the jury's still out.

"I don't think the technology has reached the point where you can draw a trend line to predict where the crossover will occur," said Intel's Pinfold.