Pulses of Light Make Faster Computers

SO inexorable is the demand for ever-greater computing power that attempts are being made to exceed Moore's Law—that computer performance doubles roughly every 18 months. This effort is particularly evident in ultrascale computing, also known as parallel multiprocessing. Parallel supercomputers integrate as many as a thousand processors into a single system to achieve calculating speeds up to several trillion operations per second (teraops). However, further advances have been thwarted by bottlenecks in shunting data from processor to processor via traditional electronic interconnects.

A team of Lawrence Livermore researchers believes it has found a way to overcome communications limitations by replacing the flow of electrons interconnecting the processors of a supercomputer (or conceivably, the computers of a network) with pulses of light of different wavelengths. By combining Livermore advances in optoelectronics with off-the-shelf hardware, they are pointing the way to communication speed improvements up to 32-fold. And because optical interconnects can be packaged very tightly, additional microprocessors can be added to a supercomputer (or more workstations can be added to a network) for much greater overall performance with no decrease in communication speeds.

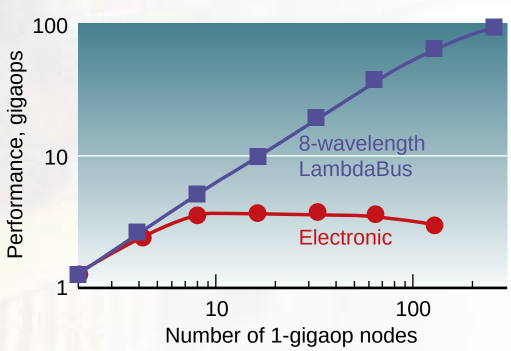

Lambda-connect transmitter–receiver modules handling only eight wavelengths easily outperform electronic interconnects, particularly as the number of nodes (processors or machines networked together) increases.

The technology development project is called lambda-connect (the Greek letter lambda represents wavelength in scientific notation). Funded by the Laboratory Directed Research and Development Program, the project has made rapid progress and resulted in the filing of four patents based on different aspects of the project. Lambda-connect appears so promising that several supercomputer and computer component companies have begun discussions with the Livermore researchers on ways to incorporate the new technology into their products. It has also been well received by government agencies that need new technologies capable of processing unprecedented volumes of data in as short a time as possible.

According to electronics engineer and principal investigator Robert Deri, ultrascale computers are essential for the Department of Energy's Accelerated Strategic Computing Initiative, which is developing capabilities to simulate nuclear weapon performance in lieu of nuclear testing. Ultrascale computers are also envisioned for climate and biomedical simulations as well as for specialized intelligence and Department of Defense missions.

The full potential of ultrascale computers has not been realized because standard approaches for sharing data among their many processors have limited their speed. Because traditional wire cable connections can carry only one "message" at a time, data become backed up while waiting to be processed or routed to another processor. These bottlenecks substantially degrade computational performance, complicate programming, and cause inefficient use of memory. Simply adding additional processors can compound the congestion without significantly improving performance.

Performance Gap Is Widening

Unfortunately, the performance gap between communications and other system components is widening. More demands have been placed on communications capabilities by more powerful computer boards (with more powerful processors and more processors per board), faster memory, increasing use of memory caches and shared memory, and new sensor systems that generate enormous amounts of data for processing.

Deri says that users need new tools to break the bottlenecks and meet the increasing communications demands to fully use computing power, memory, and sensor data. Several solutions have been offered, but they fail to relieve two key problems: inadequate throughput (the rate at which data flows) and high latency (initial time delay in transferring data).

The novel Livermore approach attacks both problems by building on the growing commercial practice of replacing electrical connections with optical signals of a particular wavelength. Optical signals transmitted along glass fibers are an attractive communication medium because they do not suffer from electromagnetic interference and other drawbacks associated with electrical signals.

The Livermore technology calls on wavelength division multiplexing, or WDM, to vastly increase the utility of optical connections. Instead of a single wavelength, many different wavelengths are carried by parallel, multimode glass fibers (MMFs) that are already in use in local area networks. In that respect, says Deri, “We don’t need to invent a whole new infrastructure.” An MMF connection is about six to ten times larger than the ubiquitous glass fiber that carries telecommunications signals. The larger cabling reduces cost and improves reliability because it requires significantly looser alignment tolerances.

With lambda-connect, every parallel optical fiber within a cable carries data of different wavelengths, with each wavelength assigned a destination. Thanks to filters developed by Livermore engineers, data are “source-routed,” with their wavelength determining the ultimate destination.

Running like an Express Train

In this way, says Deri, each processor can communicate simultaneously with a large number of others without significant increases in cabling or processor complexity. Lambda-connect makes possible optical-fiber “express channels,” which like express trains, go to their designated destinations directly, requiring no electronic routing. What’s more, the number of processors can be increased significantly for powerful performance boosts with no communications delay.

With standard electronics, Deri says, every communication is like a local train making numerous intermediate stops. Electronic express channels are difficult because they strain the processing capabilities of electronic interconnects and in some cases require such long cables that electrons cannot travel their lengths effectively.

The Livermore technology achieves unprecedented gains in bandwidth combined with significant decreases in latency. “The fact that all data travel simultaneously solves the bandwidth problem, and the fact that data all travel to different destinations solves the latency problem,” says Deri. He also notes that the data error rate has been measured at less than 10^{-11}, or 1 bit in 100 billion, meaning that the technology transmits data essentially error-free.

The team is presently developing key components for transmitter–receiver modules that can handle optical data of up to 32 wavelengths. To date, their modules can route four different wavelengths on optical-fiber cabling and are integrated with standard processor boards.

Achieving the project goal requires significant innovation in filter and microoptics technologies as well as advances in vertical-cavity, surface-emitting laser (VCSEL) diode technology. Laser diodes and associated electronics inside the transmitter–receiver modules turn electrical codes into optical pulses at distinct wavelengths for data transmission to

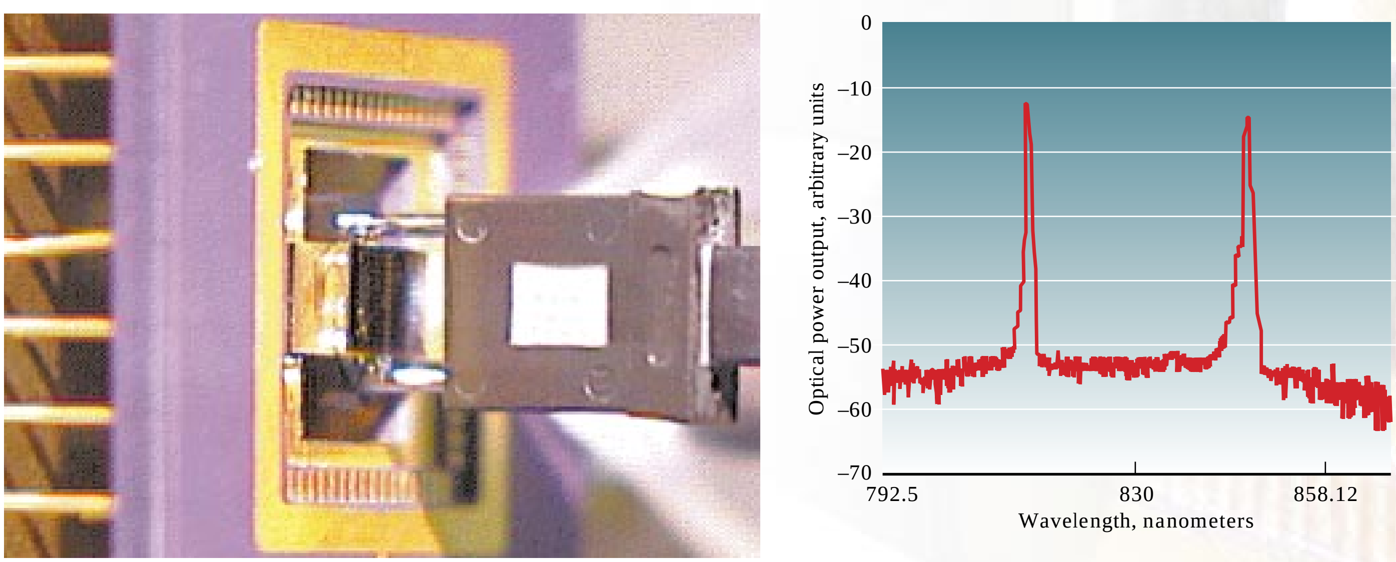

Transmitters with four wavelengths of optical data have been developed; the project’s goal is a transmitter capable of handling 32 wavelengths. The figure shows a dual-wavelength transmitter and its optical output. Wavelengths are separated by about 30 nanometers (billionth of a meter).

wavelength-encoded destinations. The optical output of the VCSEL lasers is emitted perpendicular to the semiconductor wafer surface (which is the dominant plane shown in the figure below). This surface-normal emission requires package mounting innovations over the more conventional edge-emitting laser diodes.

Deri says that the team is well suited to develop the required components because microoptics, microassembly, and photonics are all Livermore strengths. Team members have four R&D 100 awards in the area of photonics and over 75 years of accumulated photonics experience. The team includes co-principal investigator Mark Lowry and investigators Mike Larson, Steven Bond, Mike Pocha, Raj Patel, Rick Ratowsky, Mark Emanuel, Henry Garrett, Holly Peterson, Bill Goward, Claire Gu (from the University of California at Santa Cruz), and Rhonda Drayton (from the University of Minnesota).

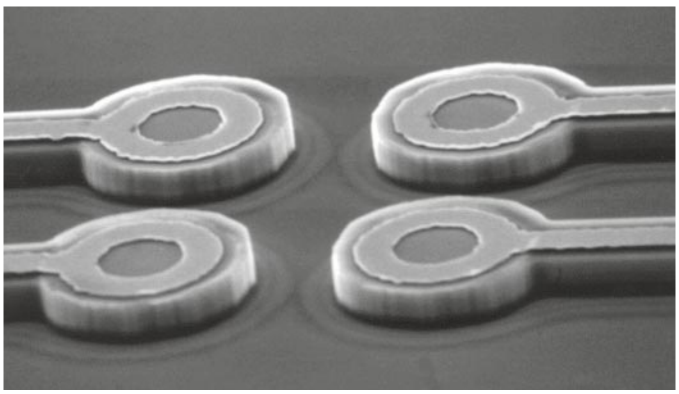

A cluster of vertical-cavity, surface-emitting lasers (VCSELs) that turn electrical codes into optical pulses for data transmission. The cluster is a closely spaced, two-by-two configuration that allows coupling into a single multimode fiber. Twelve of these clusters are coupled into fiber-ribbon arrays that are 12 fibers wide.

Deri says that "leaders in the supercomputing field have told us this approach is the most innovative and highly leveraged use of optical interconnects they have seen." The first commercial products incorporating lambda-connect technology may appear as early as 2003. Deri notes, however, that adoption of the Livermore approach depends on continued demonstration of its effectiveness to industry leaders. The team has also developed relationships with organizations that are traditionally aggressive, early adopters of advanced computing technology.

Deri expects lambda-connect advances to be used by other Livermore programs. The surface-mounted laser diodes, for example, will be used in advanced diagnostics and sensors for physics experiments. Furthermore, the project is generating interest in optical interconnects at the semiconductor chip level. However, the greatest impact on Livermore research programs will be the arrival of commercial machines using lambda-connect technologies to boost computer performance to record heights.

—Arnie Heller

Key Words: embedded systems, lambda-connect, Moore's Law, multimode glass fiber (MMF), optoelectronics, photonics, vertical-cavity, surface-emitting laser (VCSEL) diode, ultrascale computing, wavelength division multiplexing (WDM).

For further information contact
Mark Lowry (925) 423-2924 (lowry3@llnl.gov).