Optical Interconnects for Chips Possible

May 10 2004

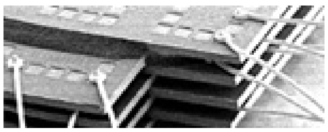

"Optical can potentially attain much higher system bandwidth with lower electromagnetic interference", says Ian Young, Intel Fellow and Director of Advanced Circuit and Technology Integration.
Researchers at Intel's Components Research Lab have taken a significant step in the direction of optical interconnect technology by developing a prototype device that is based on low-power CMOS transceivers and high-volume commodity packaging technology and attained an aggregate bandwidth of more than 8 giga-transfers per second. By building many of the device components on a standard CMOS process, the researchers have sought to remove some of the inherent difficulties of optical interconnect technology and provided a foundation for chip-to-chip interconnects over the next decade.

Optical interconnect technology, in the form of fiber optics, has been used for decades in long-distance applications, such as telephony and wide area networks, and in recent years has made inroads into shorter-distance applications, such as communications between servers in a data center. Now, optical interconnect is showing promise for ultra-short-reach applications in its ability to support far higher bandwidth than the metal wires that now carry data from board to board, chip to chip, and component to component within a single chip.

While still some years from viability in a production environment, the technology has been demonstrated by a team from the Intel Components Research Lab in a fully functioning prototype device that attained an aggregate bandwidth of more than 8 giga-transfers per second.

Peek inside the computer of the future. Here data travels on beams of light. Photons rather than electrons pass signals between one destination and another. Within a decade such optical interconnects may replace the copper metal interconnections in computers.

Although it's more expensive today to communicate with light than with electric current, Intel researchers are working to change that and say that optical data transfer holds future promise as a more efficient means to move large amounts of data inside a computer rapidly.

With the continued growth in the integration density of CMOS (complementary metal-oxide semiconductor) and clock frequency of microprocessors, the aggregate bandwidth required between future-generation microprocessors and chipsets or between multiprocessors on a motherboard will increase sharply.

For example, in about 5 years CMOS-based transistors will be fast enough for transceivers to operate at clock speeds of roughly 14 GHz, fast enough to support data-transfer rates of 20 Gbps. This will seriously challenge the current copper-trace-based technology, which is considered likely to top out at 15-20 Gbps due to the significant signal degradation, power dissipation, and electromagnetic interference that are unavoidable at higher clock speeds.

Manufacturers have sought to address the limitations of copper by increasing the thickness of the traces, using more exotic substrate materials with lower dielectric loss tangents, and employing differential signaling techniques or more sophisticated input/output (I/O) drivers at the transmitter and receiver. However, all these potential solutions are costly, thereby making optical interconnect technology an increasingly attractive alternative.

Optical interconnect technology could begin to play a significant role by the middle of this decade. Around 2006 or 2007, we expect to see commercial-level optical interconnect technology at the board-to-board level, and sometime in the 2010s at the chip-to-chip level.

Optical interconnect technology has been used for decades in long-distance applications such as telephony and the Internet, and in recent years it has been applied to shorter-distance applications such as computer-to-computer, rack-to-rack, and network-to-network communications. Such applications have proven that even at high frequencies, optical fibers and waveguides can support an unlimited distance-bandwidth product as compared with electrical cables, with minimal loss in signal and almost no cross-talk.

Still, optical interconnect technology in the ultra-short-reach setting poses its own cost problems. The technology depends largely on components made of gallium arsenide and germanium, which are more expensive than silicon. Alignment is also a more delicate matter in optical technology, and for this reason optical equipment has been difficult to design and build. Consequently, current research is focused on ways to make the technology more cost-effective, particularly from a manufacturing standpoint. Ultimately, an optical solution could replace electrical interconnect when it has higher performance at lower cost and strong manufacturability in high volume.

Toward that end, researchers at the Intel Components Research Lab are matching high-performance optical components, such as vertical-cavity surface-emitting lasers (VCSELs), with a cost-effective and industry-compatible packaging solution based on low-power CMOS transceivers and standard high-volume microprocessor packaging technology. Most recently this research culminated with the January 2004 demonstration of a fully functioning prototype device that reached aggregate speeds of more than 8 giga-transfers per second.

The device is a high-speed 12-channel link (8 data channels) with a parallel CMOS optical transceiver package. The optical I/O is based on an optoelectronic flip-chip pin grid array (FCPGA) package, and the key components of the hybrid package are gallium arsenide VCSELs, p-type intrinsic n-type doped silicon (PIN) photodiode arrays, acrylate polymer waveguide arrays, multiterminal fiber-optic connectors, and the CMOS transceiver chip.

These components are flip-chip mounted on top of the FCPGA organic substrate to enable parallel point-to-point optical transmission. During complete-link optical transmission, VCSEL arrays are directly modulated by CMOS drivers with non-return-to-zero (NRZ) data signals with source-synchronous clocking and coupled to multimode polymer waveguide arrays for transmission to a receiver array that consists of gallium arsenide photodiodes and on-chip transimpedance amplifiers. Test circuits residing on the CMOS chip enable bit-error testing of the optical transmission link.

Read more details on Intel web-site.