

Challenges and Potentials for multi-Terabit-per-second Optical Transceivers

Ashok V. Krishnamoorthy, Dawei Huang, Theresa Sze, Robert Drost*, Ron Ho*, Howard Davidson and Rick Lytel

Physical Sciences Center, Sun Microsystems, Inc., 9515 Towne Center Dr., San Diego, CA 92121

*Sun Research Laboratories, Sun Microsystems Inc., 2600 Casey Ave., Mountain View, CA 94043

Abstract: We discuss the technical challenges and potential for implementing optical transceivers with multi-terabit-per-second capacities. We will examine typical approaches, discuss bottlenecks that have developed in the interconnect hierarchy, and identify potential solutions.

According to the ITRS 2003 roadmap, off-chip data rates per differential pair will have to increase from 2.5Gbps in 2003 to 4Gbps in 2006 and to 15Gbps in 2012 to achieve the required aggregate bandwidths for high-performance computing systems. Additionally, we expect the total number of high-speed off-chip signals to follow a similar trend and increase from 300 to 600. A potential roadmap for the total off-chip bandwidth requirements follows below.

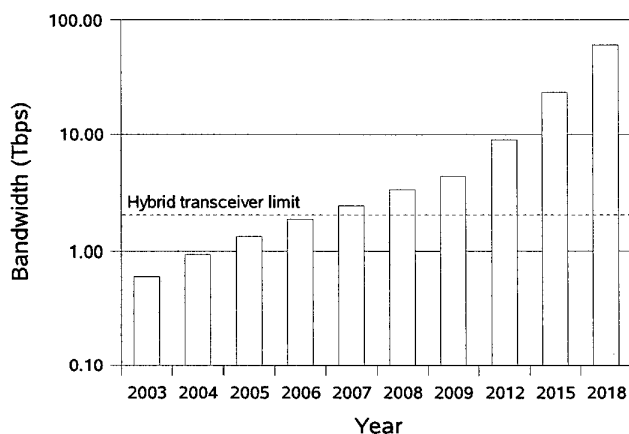


Figure 1: Scaling Off-chip Bandwidth Interconnect Requirements. Solid bars show projected off-chip bandwidth needs. The capacity limit for hybrid solutions resulting from electrical off-chip signaling constraints is indicated. We expect that a fully-integrated optoelectronic approach will be needed to meet the requirements.

Parallel optical interconnects have had significant penetration in box-to-box interconnect applications. To enter the box, there still exist four major issues for parallel optical interconnects: reliability, physical transport medium, integration and cost.

In a typical server application, reliability and availability are critical performance parameters. A typical metric is “five 9s availability”, which means the server should be available 99.999% of the time. This translates to less than 4 minutes of down time per year. If we assume 20 FITs per single VCSEL and 100,000 VCSELs per system, then the total FIT of the optically interconnected system would be 2,000,000, which translates to 1 failure per 500 hours. This falls far short of the five 9s requirement. Therefore, reliability improvements, redundancy and innovative device-level solutions are needed. Fortunately, strong progress in VCSEL reliability and recent innovations in high-density optical modules are improving system availability.

Physical transport medium is another important issue. The proposed approach should be able to scale to 10s of Tbps per chip, as shown in Figure 1. Fiber and polymer wave-guides are obvious candidates. However, none of these solutions can today support the sheer number of connections at the interconnect distance and density required by the modern processor chip. Free-space interconnects have been proposed, but a platform consistent with main-stream manufacturing flow or thermal constraints has not yet emerged. There are, however, promising efforts underway to provide scalable optical communication capability between chips that may address this weakness.

For the foreseeable future, electrical VLSI circuits will be responsible for processing information. Therefore, any optically interconnected system will involve optical-to-electrical conversion and vice-versa. Delivering data to the optical components and breaking *electrical* bottlenecks becomes one of the most critical issues for optical transceivers, particularly if the photonics and electronics are not tightly integrated. The first electrical bottleneck appears between the optoelectronic driver and receiver circuits and the photonic devices themselves. As discussed in [1], there are no accepted electrical interconnect solutions for off-chip bandwidth beyond 2Tbps. Therefore, for any non-integrated optical transceiver technology, i.e., a hybrid solution, the capacity limit will be identical to this electrical limit, as shown in Figure 1. One way to resolve this bottleneck would be to tightly integrate the optics and electronics. A second electrical bottleneck appears in the data transport from the processor to the input of the optical transceiver. This bottleneck is more difficult to solve, due to a lack of very high-speed bus standards among processor vendors, and additionally because there is no credible thermal or packaging solution for an integrated processor-plus-photonics offering.

The final issue for penetration on a massive scale is eventually one of cost. However this is highly dependent on target volumes and technology investment, hence not a core issue. Today’s cost curve stems from the fact that the integration level of optical transceivers is low and the production volume is low. As evidenced by the semiconductor industry over the past several decades, a higher integration level enables low-cost production. Hence both integration and cost issues may be simultaneously solved with a tightly integrated optical transceiver that removes both electrical and optical data transport bottlenecks.

In this presentation, we will survey and contrast various existing and novel optical integration and packaging approaches for optical transceivers and their potential to solve these four issues.

This material is based upon work supported by DARPA under Contract No. NBCH3039002.

[1] D. Huang, T. Sze, A. V. Krishnamoorthy, A. Del Alamo, D. Beckman, S. Fazelpour, H. Davidson, J. Cooley, and R. Lytel, “The Chip Multithreading Architecture and Parallel Optical Interconnects,” to appear in *IEEE LEOS Summer Topical Meeting on Optical Interconnects*, San Diego, 2004.