Optical Chip-to-Chip Interconnects for Memory Systems

Oliver Boudreaux, Sang-Yeon Cho, Jaemin Shin, Ananthasayanam Chellappa, David Schimmel, Martin Brooke, Nan Marie Jokerst
School of Electrical and Computer Engineering, Georgia Institute of Technology
Atlanta, GA 30332-0250; oliver@ece.gatech.edu

Computer processor to memory interconnections currently supply data at bandwidths well into the gigabit range, and aggregate bandwidths into the terabit range are projected [1]. Larger data widths, greater issue widths (higher parallelism), and speculative fetching all contribute to the processor's demand for memory traffic. Other devices, such as display or communication subsystems, place additional demands on the interconnect to memory as well. To accommodate these demands, the interconnect to memory is either widened (i.e., more parallel lines), creating area and packaging problems, or, designed for higher data rates, necessitating the use of electrical transmission lines. Current board level electrical transmission lines are point to point interconnections. Processor to memory interface architectures can exhibit significant performance improvements when multi-drop interconnections are utilized. However, high data rate point-to-point electrical transmission lines make multi-drop interconnections very difficult to implement at high speeds. For high speed multi-drop processor to memory interconnections, this paper proposes that optical interconnections offer a simple to design and implement alternative to point to point electrical interconnections. In fact, the design complexity of these optical interconnections is similar to lower speed multi-drop electrical interconnections that are currently the dominant architectural implementation in current computing systems [2- 6].

First, we will explore the advantages of multi-drop interconnections over point to point interconnections. As processors increase in speed and complexity, they access larger regions of memory at higher rates. Unfortunately, the performance of DRAM does not increase as rapidly as the demands placed upon it because of delays intrinsic to the DRAM device. The three dominant delays are the time to decode the row and sense the data (t_{RCD}), the time to decode the column and output the data (CAS), and the time to write the data back to memory when moving to another row (t_{RP}). Assuming consecutive accesses are to distinct banks, the transactions can be overlapped in time to hide these delays, resulting in a lower average transaction time. Above some number of drops, the latency may be completely hidden and the performance will saturate. If we look at a generalized model of a memory system (Fig. 1a), we can see how the number of drops *m* and the number of channels *n* affects the average memory access time.

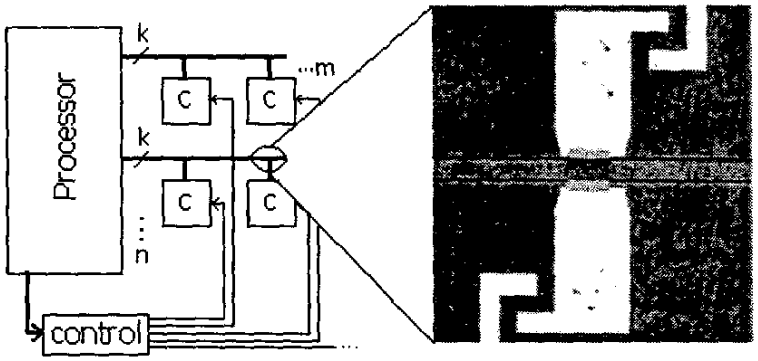

Fig. 1a General memory system **Fig. 1b Embedded photodetector**

Keeping the aggregate bandwidth to memory *kn* and the total size of the memory *mnc* constant, we can estimate the relative performance of a variety of configurations for different access sizes, delays, channel widths, and channel speeds. Systems with 1 Tb/s aggregate bandwidth, 100Gb/s channel bandwidth, and modules with 16 ns access times are compared in Figure 2. Configurations in the flat area show performance saturation while configurations in the curved area show sub-optimal performance. Of the optimal configurations, systems with more drops (7 in our example) and fewer channels (1 in our example) have fewer modules, and therefore, less overhead. This result holds regardless of the delays, size of the access or the data rate as the entire saturation curve is pushed out along both axes. The saturation occurs for fewer drops when there are more parallel channels because increasing the number of channels has the effect of reducing the total number of accesses to each channel and reducing the total latency, making it easier to hide.

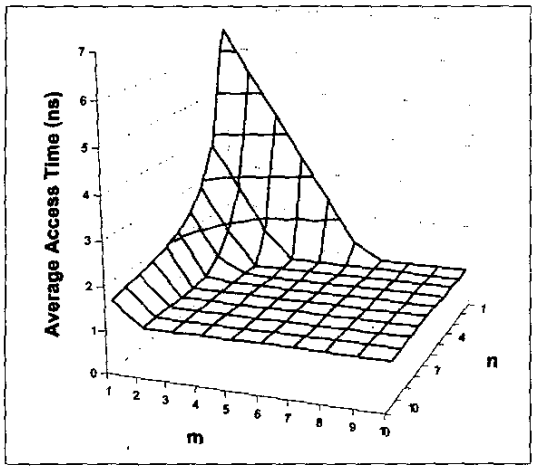

The next critical question relates to the physical implementation of multi-drop optical interconnections. A number of approaches have been studied for chip-to-chip optical interconnection. Most of the studies have been focused on the integration of the optical signals into an electrical interconnection system including free space interconnects with diffractive optical elements [7], fiber optic waveguides [8], and integrated waveguides [9]. The drawbacks to free space and silicon bench optical interconnection include misalignment and reliability problems. Guided wave approaches, including substrate guided mode interconnects, fiber optic waveguides, and integrated optical waveguides offer highly reliable signal routing since transmitting optical signal are confined and guided in the waveguide. However, the previous guided wave optical interconnects also have misalignment issues between transmitters/receivers and waveguides. To solve this alignment problem, fully embedded thin film optoelectronic interconnections have been proposed [10]. For the proposed interconnection system, the emitter and photodetector are heterogeneously integrated onto the interconnection substrate and embedded in the optical waveguides to create planar lightwave circuits (PLCs).

The simplest (nonoptimized) physical implementation of an optical multi-drop architecture is a linear array of thin film photodetectors embedded in the direction of propagation in an optical waveguide. If each active embedded thin film photodetector (as shown in the photomicrograph in Figure 1b) couples a fraction of the optical signal in the waveguide, then a simple multi-drop scheme can be implemented. By changing system variables such as the waveguide/photodetector interaction length, it is also possible to balance the photodetector outputs. Alternative schemes include fanout implementations such as multimode interference (MMI) couplers and H-tree structures. Through proper design of the embedded integration structure, the number of drops in the proposed structure can be determined by the following simple equation:

$$(1-\eta)^{N-1} \times 10^{-0.1\alpha L_2(N-1)} = \frac{I_{PD}}{P_0\, 10^{-0.1\alpha L_1} \times \eta \times R}$$

where η is the coupling efficiency at each drop, R is the PD responsivity, P_0 is the emitter output power, α is the optical propagation loss, $L_1$is the distance between microprocessor and the first drop, $L_2$is the separation between each drop, I_{PD} is the minimum required photocurrent for receiver at the last drop, and N is the number of drops. For example, a fairly typical OE system would consist of: emitter output power: 5mW, η: 10%, R: 0.6 A/W] α: 0.1 dB/cm, L_1: 5cm, L_2: 2cm, and I_{PD}: 10 μA. This system could support up to 38 drops, which is more than sufficient for a 1Tb/s system.

1. D. Burger, J.R. Goodman, A. Kagi ,"Limited bandwidth to affect processor design" *Micro, IEEE*, 17, pp. 55-62, 1997.
2. G. I. Yayla, P. J. Marchand, and S. C. Esener, "Speed and energy analysis of digital interconnections: comparison of on-chip, off-chip, and free-space technologies," *Applied Optics*, vol. 37, no. 2, pp.205-227, Jan. 1998.
3. D. Skirmont, "Optical Interconnection Backplane Technology for Terabit Routers", *Pluris*, Jan, 2000.
4. B. R. Rothermel, D. W. Helster, A. M. Sharf, "practical Guidelines for implementing 5Gbps in Copper Today, and the Roadmap to 10Gbps", *AMP DesignCon*, 2000.
5. Christopher Tocci, H.John Caulfield, Optical Interconnection – Foundation and Application, *KARTECH HOUSE,INC*, ISBN 0-890006-632-9,1994.
6. J. Shin, C. Seo, A. Chellappa, M. Brooke, A. Chatterjee and N. Jokerst, "Comparison of Electrical and Optical Interconnect", Electronic Component and Technology Conference, pp.1067-1072, 2003
7. S. Tewksbury and L. Hornak, "Optical clock distribution in electronic systems," *J. VLSI Sig. Pro.*, V.16, 1997, pp. 225-246.
8. P. J. Delfyett, D. H. Hartman, and S. Z. Ahmad, "Optical clock distribution using a mode-locked semiconductor laser-diode system," *Journ. of Light. Tech.*, vol. 9, 1991, pp. 1646-1649.
9. B. Bihari, J. Gan, L, Wu, Y. Liu, S. Tang, and R. T. Chen, "Optical clock distribution in supercomputers using polyimide-based waveguides," *Proc., Opto. Intercon. VI.*, San Jose, CA, Jan. 1999, pp. 123-133.
10. Sang-Yeon Cho, Sang-Woo Seo, Martin A. Brooke, Nan Marie Jokerst, "Integrated Detectors for Embedded Optical Interconnections On Electrical Boards, Modules, and Integrated Circuits," IEEE Journal of Special Topics in Quantum Electronics, Vol. 8, No. 6, pp. 1427-1434, 2000.