IEEE LEOS

Challenges for Photonics in Future Systems

A.F.J Levi

Abstract
As the physical scaling of transistors comes to an end, new approaches are needed to ensure continued enhancement in system performance. Photonics is uniquely positioned to enable innovative distributed architectures of otherwise all electronic systems.

INTRODUCTION

By the year 2020 Moore's law will come to an end and the minimum feature size of CMOS transistors will approach 10 nm. At this feature size individual transistors will have switching speed characterized by a unity gain cut-off frequency fT in excess of 400 GHz. It will be possible to create circuits with 100 million transistors per mm2 that have IO data rates of 80 Gb/s. IO at these data rates will use equalization and PAM-4 coding. At transistor densities of 100 million transistors per mm2 one may, in principle, accommodate two Intel P4 processors per mm2 of silicon. Electrical interconnects within such chips will be limited by wiring density and to mitigate the impact of this ICs will have 12 or more layers of metal interconnect. Electrical interconnects between chips will be severely limited by IO density, power, and cross-talk [1]. To avoid the expense of liquid cooling the 1 mm2 of circuitry for a dual processor will be distributed on a 10 mm2 die. The additional IC area will be available for passive, self-test, calibration, power regulation circuits, and high-speed IO pads. With such an approach it will be possible to expose several Tb/s bisection bandwidth IO per chip enabling new scalable architectures that are much more efficient than the P4 [2]. While IC packaging and backplane media is capable of supporting 80 Gb/s data rates over distances of 0.5 m, today the barrier to implementation is development of a high-density, near perfect, backplane connector.

On length scales greater than 0.5 m, photonics is the only technology likely to be capable of removing chip IO data bandwidth bottlenecks and thereby extending system performance beyond the end of Moore's law. The system impact of high-bandwidth density optical interconnects includes reduced message latency and improved system scalability [3]. The challenge is to develop a photonics technology that simultaneously provides significant system performance advantages, is transparent to the system designer, and costs less than an all-electrical approach.

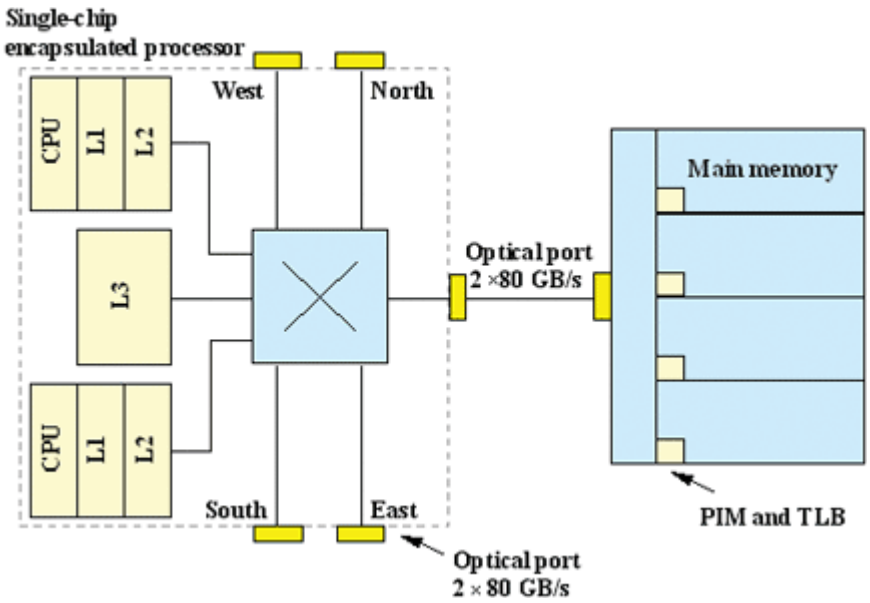

Fig. 1. The encapsulated processor is a single CMOS chip with optical ports as the only means of external high-speed data communication. The processor consists of two CPUs with L1 and L2 cache connected by a crossbar switch. The crossbar connects to shared on-chip L3 cache and multiple high-speed fiber-optic ports. Each optical port is capable of sustaining 80 GB/s (640 Gb/s) data throughput in each direction and one such port is dedicated to local main memory. Main memory could be configured to have its own processors and TLB. The remaining optical ports are available for IO and scalable SAN interconnect.

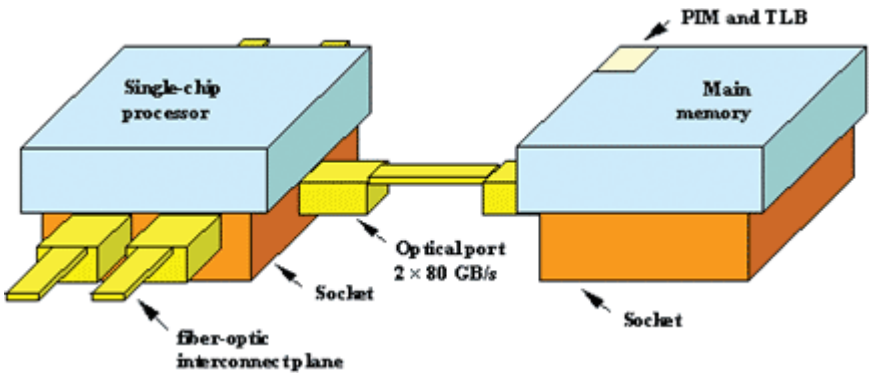

Fig. 2. The encapsulated processor includes a socket that supplies DC current and ground. Incorporated into the socket are the physical optical ports, each of which provide 2 X 80 GB/s = 1.28 Tb/s data bandwidth external to the encapsulated processor. One optical port is dedicated to local main memory

THE ENCAPSULATED PROCESSOR

The advantages of optical-solutions are by now well established [1]. There is reduced power dissipation from high-speed chip IO and improved edge-connection density bandwidth. The optical transmission medium has low cross-talk and zero EMI. Optical ports provide high-bandwidth to key resources such as main memory and the System Area Network (SAN). The same high-bandwidth allows scaling to larger multi-processor systems. Another important advantage of fiber-optic enabled distributed multi-processor systems is the reduction in overall system power density.

Incorporating the advantages of photonics with the integration capability of scaled CMOS electronics leads to a new microprocessor design-point illustrated in Fig. 1. Here an encapsulated processor is defined as a single CMOS chip with optical ports as the only means of external high-speed data communication. The processor IC and optical port have separate thermal management. There is a short electrical link from the processor IC to the optical port IC embedded in the socket shown in Fig. 2. The electrical link is low-power because there is no need for controlled 50 W impedance. The optical port IC decodes and multiplexes signals for the optical sub-assembly that contains low-power VCSEL

transmitters, PIN receivers and the optical interface. The bandwidth density of the optical communication channel is significantly greater than an electrical alternative. Each optical port is capable of sustaining 80 GB/s (640 Gb/s) data throughput in each direction and one such port is dedicated to local main memory. Main memory may have its own processors (PIM) and pipelined translation look-aside buffer (TLB) whose purpose is to efficiently feed the encapsulated processor. The remaining optical ports are available for IO and scalable SAN interconnect.
Scalability of the SAN is dependent on there being enough high-speed ports available for the network. In Fig. 1, the bisection bandwidth of an 8-port cross-bar switch integrated into the encapsulated processor is 128 GB/s (10.24 Tb/s). Innovative circuit design predicts that the switch-core consumes only a few watts of power.

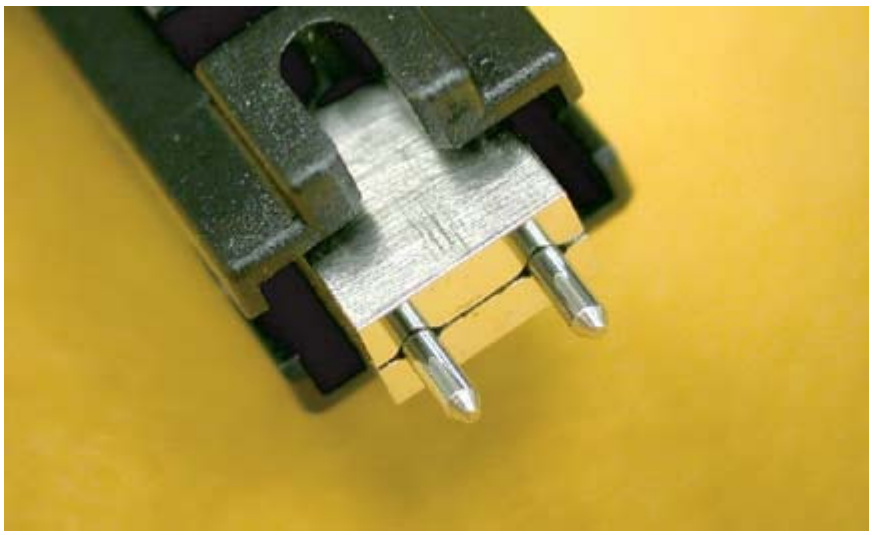

Fig. 3. Precision stamping can be used to manufacture very low cost optical connectors. The photograph shows a MT-RJ connector for SMF with a precision stamped metal ferrule developed at USC.

LOW COST VOLUME MANUFACTURING

While the DARPA sponsored MAUI program is developing pre-competitive optical port technology for the encapsulated processor [4], there still remains the important issue of reducing both the price and cost of photonic components. Conventional assembly methods that use machine vision and robotics are not well suited to volume production of several million units per month. New approaches need to be explored.
One of the few volume manufacturing methods capable of reducing manufacturing expense to levels close to cost of materials is stamping. Until recently, it was not possible to reliably stamp metals or other materials to the 100 nm accuracy required for snap-together assembly of single-mode fiber-optic components. While today milling and grinding machines are available that can machine punch and die sets to the required accuracy, understanding the plastic deformation of materials at the nm scale, developing design methodologies, and ensuring reproducibility is still a challenge.
At USC we have made some initial progress in applying precision stamping to photonics. Fig. 3 shows a MT-RJ ferrule for single mode fiber (SMF) that has been precision stamped using a closed punch and die. Future directions for this technology include snap-together assembly of optical transmitters and receivers suitable for use as the optical ports integrated in the processor sockets illustrated in Fig. 2. Availability of low price, high-performance, optical ports embedded in sockets is only one piece of the needed infrastructure. For example, printed circuit board technology to route optical and electrical signals between CMOS chips is also required and is being developed by a number of groups [5]. Of course, only if these solutions are brought to the marketplace will the future for optics in systems be bright!

REFERENCES

1. A. F. J. Levi, "Optical interconnects in systems," Proc. IEEE vol 88, 750 (2000).

2. M. Horowitz and W. Dally, "How scaling will change processor architecture," ISSCC, vol 47, 134, 2004.
3. P. Wijetunga and A. F. J. Levi, "The Case For Fiber-to-the Processor," Integrated Optoelectronics: Proceedings of the First International Symposium, eds. M. J. Deen, D. Misra, and J. Ruzyllo. 201st Meeting of The Electrochemical Society in Philadelphia, PA from May 12-17. Electrochemical Society Proceedings 2002-4, pp. 381-397 (Electrochemical Society, Pennington, NJ, 2002).
4. B. E. Lemoff, M. E. Ali, G. Panotopoulos, G. M. Flower, B. Madhavan, A. F. J. Levi, and D. W. Dolfi, "The MAUI program - Enabling fiber-to-the-processor with parallel multiwavelength optical interconnects," J. Lightwave Technol. (2004).
5. See, for example, R. T. Chen, L. Lin, C. Choi, Y. J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M. K. Hibbs-Brenner, J. Bristow, and Y. S. Yu, "Fully embedded board-level guided-wave optoelectronic interconnects", Proc. IEEE, vol 88, 780 (2000).

** CuDepartment of Electrical Engineering, University of Southern California, Los Angeles, CA 90089-2533*
Phone: (213) 740 - 7318, email: alevi@usc.edu, web: http://www.usc.edu/alevi

---

return to contents