On-Chip Photonic Communication for High-Performance Multi-Core Processors

Keren Bergman[1] and Luca P. Carloni[2]

Department of Electrical Engineering[1] and Department of Computer Science[2], Columbia University

500 West 120th Street, New York, NY 10027

bergman@ee.columbia.edu; luca@cs.columbia.edu

Abstract

The quest for high-performance and low-power has brought computer architects to design multi-core architectures where an increasing number of parallel processing cores are integrated on a single die to operate in a tightly coupled fashion. With nanometer technologies, a chip multi-processor (CMP) based on a multi-core architecture delivers better performance-per-watt than a traditional deeply-pipelined superscalar microprocessor running at the highest possible clock frequency. However, in order to fully exploit the processing capabilities offered by the integration of an increasing number of cores, three major challenges must be addressed: the increasing on-chip power dissipation, the limited I/O bandwidth, and the complexity of parallel programming. We argue that the design of the on-chip communication infrastructure plays a critical role across these three challenges and that the insertion of an optical network-on-chip based on nano-scale integrated silicon photonic technology offers a unique opportunity to address them. We conclude by sharing our vision for the application of on-chip photonic communication to high-performance embedded computing.

1. Introduction

New high-performance chip multi-processors (CMP) based on multi-core architectures, such as IBM Cell BE, Intel Duo and Sun Niagara, are fundamentally different from traditional microprocessors and require addressing new design challenges. As the number of processing cores integrated on the die continues to grow, CMPs increasingly resemble parallel computers where global on-chip communication plays a dominant role in the ultimate system performance. In this new communications-bound design, global wire scaling becomes a major driver of high computing performance. While local interconnects scale approximately in accordance with transistors, global wires do not because they need to span across the entire chip to connect distant logic gates. Meanwhile the bandwidth requirements for global on-chip communication scale up with the growing number of cores. As packaging-technology constraints limit the total on-chip power-dissipation budget, the increasing fraction of power devoted to global communications becomes a key bottleneck to realizing the high bandwidths capacities and stringent latency requirements that are demanded in CMPs. (*power challenge*). The power dissipation problem is greatly exacerbated for off-chip electronic interconnects as they typically consume at least one order of magnitude more power even for rather short distances and do not scale significantly with technology node. Thus, the off-chip communication bottleneck will become even worse in future high-performance computing systems that will require interconnecting large numbers of CMP nodes and memory banks (*I/O bandwidth challenge*). Convergent failings of the communication infrastructures both at the intra-node and inter-node domains force designs of highly imbalanced parallel computing systems in terms of the communications bandwidths (Bytes/sec) to computation (Flops) ratio. As larger fractions of computation capacity is taken up in managing the limited communication resources, the system imbalance directly translates to inefficient computing task parallelization, poor programmability, and ultimately detrimental application performance and machine productivity (*programmability challenge*).

2. The Case for Photonic Networks-on-Chip

Networks-on-chip (NoC) made of carefully-engineered communication links and providing packet-switched transmission services have been proposed as a shared communication medium that is highly scalable and can offer enough bandwidth to replace many traditional bus-based and point-to-point links [1]. However, NOCs that are purely based on electronic communication face major challenges in satisfying the increasing bandwidth demands of future CMPs while staying within the stringent power constraints.

Thanks to the properties of low loss in optical waveguides and bit-rate transparency [3], a photonic interconnect network can deliver considerably higher bandwidth and lower latencies with significantly lower power dissipation than an interconnect network based on electronic signaling. Photonic channels can support large amounts of data traffic across longer distances in a bandwidth-oriented NoC that interconnects many processing cores and memories. Furthermore, seamless delivery of off-chip communication bandwidth for optical message switching across multiple chips and DRAM memories can be obtained with minimal additional power consumption.

The *photonics opportunity* is becoming possible now by recent advances in nano-scale silicon photonics. High speed optical modulators at data rates exceeding 12.5Gb/s have been reported in the literature [6] and the integration of modulators, waveguides and photo-detectors with CMOS

circuits for off-chip communication has recently become commercially available [2].

On the other hand, the photonic opportunity can be realized only after overcoming the inherent restrictions of optical technologies, i.e. limited buffering and signal processing capabilities. Keeping these limitations in mind, we have made the case for the design of a photonic NoC that is based on a novel *dual-network hybrid NoC design* [4, 5]. By combining a photonic data transmission network with an electronic control network on a single CMOS integrated circuit, we can leverage the best properties of each technology. As a result fast optical message switching is accomplished without packet level alignment or synchronization. The photonic network is comprised of broadband 2×2 photonic switching elements which are capable of switching wavelength parallel messages (i.e. each message is simultaneously encoded on several wavelengths in a single waveguide) as a single unit, with a sub-ns switching time. The switches are arranged as a two dimensional matrix and organized in groups of four. Each group is controlled by an electronic circuit termed electronic router to construct a nonblocking 4×4 switch. This structure lends itself conveniently to the construction of planar 2-D topologies such as a mesh or a torus. In [4] we discussed the architecture detail of the proposed photonic NoC, including network topology, routing algorithms and flow control mechanisms. In [16] we estimated the power of a photonic NoC, and compared it to an electronic NoC designed to provide the same bandwidth to the same number of cores. The compelling conclusion of the study was that the power expended on intrachip communications can be reduced by two orders of magnitude when high-bandwitdh communication is required among the cores.

.

3. The Potential of Photonic Communication for High-Performance Embedded Computing

We argue that the realization of intra- and inter-chip photonic communication infrastructure will bring major benefits for several classes of high-performance embedded applications. In particular, for streaming applications that process massive amounts of data in real time, computation is bandwidth-intensive with limited data reuse. An important application that stresses the system's communications bandwidth is the one-dimensional Fast Fourier Transforms (1D-FFT), which belongs to the HPC benchmark suite and is a pervasive computational kernel in a variety of high-performance applications. The 1D-FFT offers a natural way to estimate the potential of photonic communication for high-performance computing. We completed a simple yet credible quantitative analysis assuming that in a future technology node (22nm) we will be able to integrate on a single die (possibly relying on 3D technology): (a) an instance of our photonic NoC delivering about 8.6TB/s of on-chip or off-chip communication bandwidth; (b) 30GB of memory and (c) 36 processing cores that combined are capable of offering about

9.2TFLOP/s of computational power. Then we considered the largest 1D-FFT in terms of the number N of complex data points that can be computed on this future CMP using the classic radix-2 Cooley-Tukey algorithm, which asymptotically requires $5 \cdot N \cdot \log_2 N$ steps. We assumed that the computation can be organized to exploit the photonic NoC capabilities such that the movement of data among the CMP processing cores occurs while the computation is performed. This places an inherent requirement on the photonic NoC to behave as a fully connected network.

Based on our traffic simulations the NoC design outlined in [4] achieves approximately 70% of this performance and would require additional 2X of spatial over-provisioning to approach full connectivity. We note that this is not an unreasonable design, as the spatial over-provisioning would require additional network paths with photonic switch elements leading to higher yet still within reasonable budget optical losses, but no additional gateway I/O hardware. In conclusion, we estimated that this future CMP can execute a 1.88-billion 1D-FFT in 38.8ms. This impressive result can be better appreciated if one considers that it represents a 1000x speed-up over what state-of-the-art supercomputers can to today and 30x better than a future equivalent electronic node. This result is made possible by the photonic NoC that, besides offering unique capabilities in terms of bandwidth–per-watt, in this scenario plays two key roles: (1) during the loading of the on-chip memory it provides the necessary off-chip I/O bandwidth (2) during the FFT computation it enables the necessary high-bandwidth on-chip data transfer among distant processors with the same throughput as if the data is "local to" the processor.

4. Conclusions

Efficient design of intra- and inter-chip communication is key to the success of future multi-core processors. While several challenging design problems remain to be addressed, photonic NoCs offer a unique solution for high-performance embedded computing in terms of combined ultra-high bandwidth and limited energy dissipation.

References

[1] W. J. Dally and B. Towles, *Route packets, not wires: On-chip interconnection networks,* In Proc. of Design Automation Conf., June 2001.

[2] C. Gunn, *CMOS photonics for high-speed interconnects*. IEEE Micro, 26(2):58-66, Mar./Apr. 2006.

[3] R. Ramaswami and K. N. Sivarajan. *Optical networks: a practical perspective*. Morgan Kaufmann, 2002.

[4] A. Shacham, K. Bergman, and L. P. Carloni, *On the design of a photonic network on chip*. In The 1st IEEE Intl Symp. on Networks-on-Chips (NOCS), May 2007.

[5] A. Shacham, K. Bergman, and L. P. Carloni, *The case for low-power photonic networks-on-chip,* In Proc of Design Automation Conference, June 2007.

[6] Q. Xu et al., *12.5 Gbit/s carrier-injection-based silicon microring silicon modulators*. Optics Express, 15(2):430–436, 22 Jan. 2007.