

Scalable Nanophotonic Interconnect for Cache Coherent Multicores

Randy W. Morris, Jr. and Avinash K. Kodi

Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701

E-mail: rmorris@cs.ohiou.edu, kodi@ohio.edu;

Category 1 Submission

Abstract: While snoop-based cache-coherence protocols provide fast cache-to-cache transfer, scaling the cache coherence protocols has been difficult due to the overhead of broadcasting requests using electrical networks. Therefore, as multicores increase exponentially with each technology generation, it is necessary to re-think the design of power and area-efficient interconnects to enable fast cache-to-cache transfers. One potential disruptive technology solution to maintain cache coherence is nanophotonics, primarily due to its high bandwidth, low power and simpler broadcast capabilities. Using nanophotonics, we propose a high speed and low power interconnection for future cache coherent multicores called CC-NPA. Our initial results indicate that CC-NPA increases performance by 25% and reduces power by 55% when compared to an electrical bus network.

1. Introduction

Present trends in computer architecture show that future multicores will be comprised of 10 to 100s of cores [1,2]. However, the continuous increase in multicores performance will be realized only if a scalable cache coherence protocol can be implemented on a power-efficient and high-performance communication network. Traditionally, the ease of programming found in snoopy cache coherence protocols is combined with the broadcast capability and natural ordering of the shared-bus [3]. However, bus-based networks are power-inefficient and have limited bandwidth that can sustain a small number of cores. A potential solution is to replace metallic interconnects with complementary-metal oxide semiconductor (CMOS) compatible power-efficient, area-efficient, high-bandwidth nanophotonic interconnects [4,5]. Nanophotonic technology allows for easier broadcasting of requests which can be used to design scalable cache coherent multicores.

In this paper, we propose a 16-core chip multi-processors (CMPs) call CC-NPA (Cache Coherent-Nanophotonic Architecture) that combines the benefits of snoopy cache coherent protocol with nanophotonics. We propose a tree-based nanophotonic network that combines cache requests using couplers and splitters, and broadcasts the requests to all cores simultaneously. CC-NPA use of a tree network allows for optical light with the same intensity to arrive at all destinations simultaneously without the need for either different intensity splitters (as in optical bus networks) or multi-cycle multicasts (as in electrical mesh-based networks). Moreover, CC-NPA allows memory controllers to be placed at the root of the tree network to maintain the total order of transactions. We also propose an optical power guiding system which allows for a reduction in optical power by only supplying power to the current active column of cores that are permitted to send out address requests. Prior nanophotonic network such as Shared-Bus architecture [1] combines electronics and optics for broadcasts whereas CC-NPA uses a one-hop optical network without any electrical broadcasts. Another optical network for symmetric multiprocessors (SMPs) called SYMNET has been proposed where an optical tree network is used to broadcasts requests [2]. While SYMNET has been proposed for board-level interconnects, CC-NPA is proposed for on-chip implementation with an emphasis on power reduction by utilizing power guiding techniques. The significant contributions of this work are as follows:

- We propose a 16-core nanophotonic snoopy cache coherent network call CC-NPA that is constructed similar to a tree network by combining and splitting signals, thereby ensuring the same intensity signals at all cores.
- We propose an optical power guiding system that routes optical power to only those cores that will transmit an address request. We combine the optical power guiding with optical token distribution and allocate power to those cores that can consume the token and transmit the request.
- Our results indicate that CC-NPA increases performance by about 25% and reduces power by 55% for select SPLASH2 benchmarks when compared to an electrical bus network using SIMICS running GEMS.

2. CC-NPA Architecture

In this section, we will give a brief description of CC-NPA. Figure 1(a) shows the layout of CC-NPA, which consists of 16-cores in a grid fashion. It should be noted that we can extend this configuration to 64 cores and beyond via concentration and/or adding additional waveguides for more bandwidth. Optical interconnects are routed in a tree configuration so that any data placed on the waveguide will arrive at the 16-cores simultaneously. In CC-NPA there is a total of 15 waveguides, 5 for address network, 2 for snoopy network and 8 for data network (not shown in the figure for clarity purposes). This gives 40 bits of address information, 16 bits of snoop information

and 80 bits of data information that can be transmitted in a single cycle. In addition, the use of separate waveguides for the address, snoop and data networks allow for split-transactions to take place.

As only one core is allowed to transmit data at any given time, an arbitration technique is required so two or more cores will not try to access the bus at the same time and cause a conflict. In CC-NPA, we prevent two or more cores from accessing the bus at the same time through the use of optical tokens. Optical tokens circulate around a select column of cores and if a core in the column requires one of the three buses (address, snoop or data) it will activate its corresponding micro-ring resonator for the correct token. Optical tokens are first injected onto the inject token waveguide for the select column of cores by the control center (explained later) shown in Figure 1(b). As the token circulates it will be captured if a core needs to place data on one of the three buses. After placing data on the correct waveguide, the token is injected back onto the inject token waveguide for other cores in the column to capture it. Once the token arrives back to the control center, the control center will inject the token into a different column and start the above mentioned process for a different column of cores. It should be mentioned, that there are two different and independent optical tokens circulating around the waveguides. Each of the optical tokens represents the right to communicate on either the data or address networks and different tokens can be circulating around different column of cores at the same time. They are not required to circulate around the same column together. This allows for increased utilization of the network as all cores in the selected column will not require both tokens. In order to allow fair sharing of network resources, the control center will shut down optical power to the current active row if it takes significant time for the token to return back to the control center. After enough time for the captured token to be used has past, the control center will power up and inject the token into the next column waveguide. For addition increased utilization of network resources, CC-NPA can incorporate more advance optical token techniques such as Flexishare [5] and Fair Token Slot [6] to overcome resource starvation that is possible in the current implementation of CC-NPA.

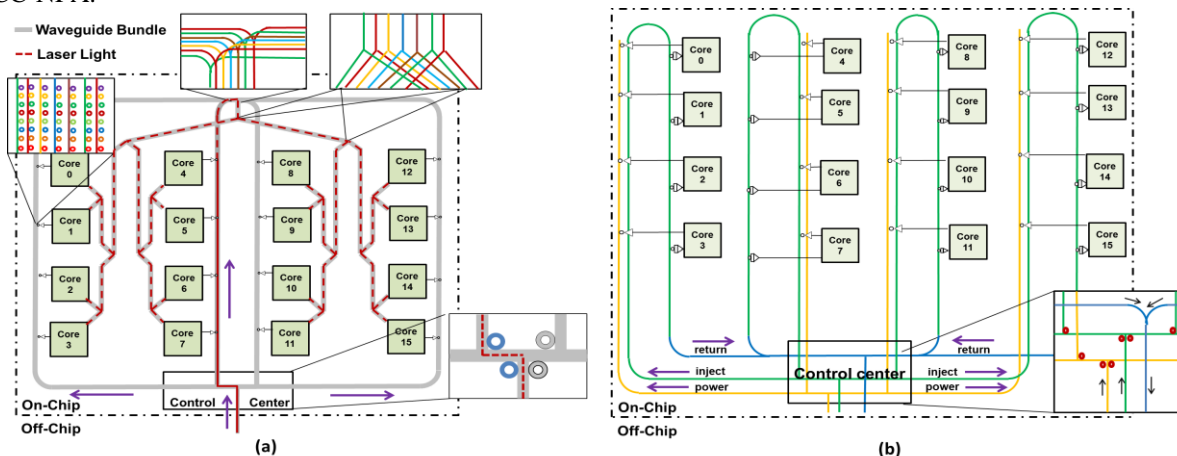


Figure 1:(a) Proposed layout of CC-NPA and (b) Optical token network use arbitration in CC-NPA.

In both Figure 1(a), and Figure 1(b), there is a control center at the bottom. The control center is used to guide optical power to the correct column of cores. This allows for a reduction in optical power as only the corresponding active column is supplied with optical power. For a better explanation of how power guiding operates lets use an example that core 0 needs to communicate and the address token is presently circulating the first column. Since the address token is circulating around the first column, the control center guides optical power only to the first column of cores as these cores only have the potential to capture the token. Here, the control center allows for reduction in power of about 75% as the three other column of cores are not powered up.

4. CC-NPA Optical Power Dissipation

CC-NPA is constructed with nanophotonic components and as such the optical power dissipation of CC-NPA needs to be taken into account. Table 1 shows the optical losses and parameters for select optical devices used to construct CC-NPA. CC-NPA maximum optical power loss is given by $5 \times L_S + 7 \times L_W + L_C + L_N + 3 \times L_I + L_F + 8 \times L_B + 100 \times L_{WC}$. This gives a maximum optical loss of approximately -43.1 dB or 204mW per wavelength for a total electrical laser power 5.44 W. This is well within the power budget found in today CMPs. From the above calculated optical power loss, waveguide crossing constitute a significant portion of overall optical power. One technique to overcome this is to use multiple optical layers as this will significantly reduce waveguide crossings.

Table 1: Optical device parameters

Device	Loss (dB)	Device	Loss(dB)	Device	Loss(dB)	Device	Loss(dB)
Coupler(L _C)	1	Modulator Insertion (L _I)	1	Bending (L _B)	1	Splitter (L _S)	3
Non-linearity (L _N) (at 30mW)	1	Waveguide (per cm) (L _w)	1.3	Waveguide Cross (L _{wc})	0.05	Laser Efficiently	30%
Photo- (L _P) detector	1	Filter drop (L _F)	1	Receiver (L _{RS}) sensitivity	-20 dBm		

3. Results

We simulate CC-NPA using the full system simulator, called SIMICS with the GEMS memory system for select SPLASH2 benchmarks. We compared CC-NPA to a standard electrical bus with a similar cache coherent design except there is an addition of two clock cycles (four cycles total) added for network communication due to a slower electrical bus. For calculating the electrical bus delay, we assumed a bus length of 3cm constructed with global interconnects. Since we assume a clock speed of 5 GHz, the time it take for a cache message to traverse the whole bus is calculated to be 780 ns. This results in a 4 cycle delay or twice the delay relative to CC-NPA. Figure 2 shows the application speed-up relative to the electrical network for the select SPLASH2 benchmarks. As you can see, CC-NPA has about a 25% speed-up of SPLASH2 applications. This is due to the faster optical interconnects allowing for quick delivering of cache coherent messages. Using Orion 2.0 and ITRS data, we estimate CC-NPA has a 55% reduction in power over an electrical bus. For calculating the electrical power dissipation in the electrical bus, we assumed a data bus of 160 bits, a bus length of 3cm and global interconnects parameters.




Figure 2: SPLASH2 speed-up of select applications relative to an electrical bus.

4. Conclusion

In this paper, we propose a nanophotonic network called CC-NPA. CC-NPA overcomes the limited bandwidth, high latency and high power found in electrical bus networks through the use of power-efficient and high-bandwidth nanophotonics. We also propose a technique for guiding the optical power to the select column of cores, thus reducing the power dissipation. Our results indicate CC-NPA increases performance by 25% and reduces power consumption by 55% when compared to an electrical bus.

Acknowledgement: This work was supported in part by the National Science Foundation (NSF) Grants CCF-0915418, CCF-0538945, CCF-0915537 and ECCS-0725765.

5. References

- [1] N. Kirman and et. al, "Leveraging optical technology in future bus-based chip multiprocessors," in Proceedings of the 39th International Symposium on Microarchitecture, December 2006, pp. 492-503.
- [2] A. Louri and A. Kodi, "An Optical Interconnection Network and a Modified Snooping Protocol for the Design of Large-Scale Symmetric Multiprocessor (SMPs)," IEEE Transactions on Parallel and Distributed Systems, vol. 15, no. 12, pp. 1093-1104, December 2004.
- [3] G. Raymond and et. at., "Nanoelectronic and nanophotonic interconnect," *Proceedings of the IEEE*, vol. 96, no. 2, pp. 230-247, Feb. 2008.
- [4] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," in *IEEE Transactions* 1. 57, no. 9, pp. 1246-1260, September 2008.
- [5] Y. Pan, J. Kim and G. Memik, "FlexiShare: Energy-Efficient Nanophotonic Crossbar Architecture through Channel Sharing," IEEE International Symposium on High-Performance Computer Architecture (HPCA), Bangalore, India January 2010
- [6] D. Vantrease, N. Binkert, R. Schreiber, and M.H. Lipasti, "Light speed arbitration and flow control for nanophotonic interconnects," Microarchitecture, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on , vol., no., pp.304-315, 12-16 Dec. 2009