Network on-Chip Optical Interconnects – Issues and Challenges

Khurram Aziz, Nasir-ud-Din Gohar
NUST School of Electrical Engineering & Computer Sciences, Islamabad, Pakistan
khurram@ieee.org, ndgohar@niit.edu.pk

Abstract

Copper wires, having served their purpose in communication networks for a long time, are now rapidly being replaced by photonics. Almost all current deployments employ photonics instead of copper. On-chip networks (NoC) seem to be headed the same way. Increasing gate speeds mean that interconnect delays have become significant and can no longer be ignored. In the 70 nm technology, the ratios between interconnect delays and gate delays can be very high. One of the effects of this increased interconnect delay is that one clock cycle is insufficient to propagate the signal across the entire chip, posing serious synchronization problems. Similarly, delay in copper interconnects and limited bandwidth forms a serious bottleneck between multi processor communications in Networks on-Chip, suggesting that the fate of copper on-chip will be similar to the fate of copper in communication networks. Recent studies suggest that photonics is a promising solution for on-chip interconnects. This study focuses on optical interconnects for Networks on-chip and looks at various techniques that have been proposed to address issues in optical interconnects.

Keywords-component; Network on-Chip; interconnects; optical interconnects.

Introduction

Networks on-Chip (NoC) is an exciting new paradigm [1] that takes advantage of the fact that various modules, such as processor cores, memory blocks and semi-conductor Intellectual Property (IP) blocks can be placed on a single chip and exchange data through the interconnects and routers that connect these modules as shown in Figures 1 and 2.

The need for Networks on-Chip (NoC) arises due to many factors. From a physical design perspective, the advances in nanometer CMOS technology [2] has meant that the interconnect delays, that used to be ignored in the past, have become significant and even surpassed the gate delays so much that they practically drive the performance

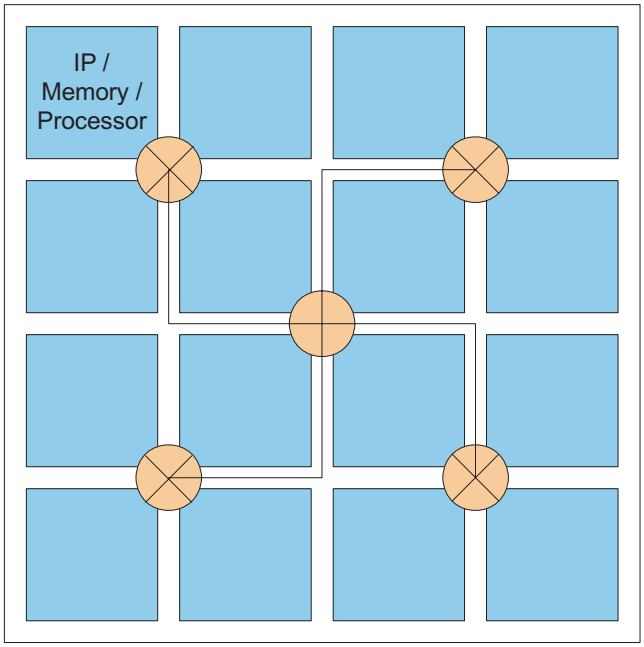

Fig. 1. Tile based, regular NoC architecture showing various modules (processors, memory, IP)

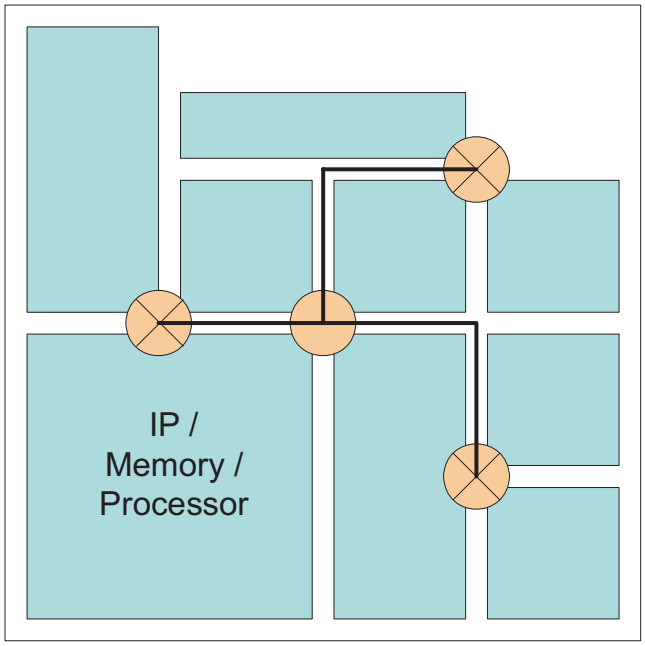

Fig. 2. Tile based, heterogeneous NoC architecture showing various modules (processors, memory, IP)

and power dissipation of the chip. Thus communication across the chip can take several clock cycles to propagate, making it impossible to have a fully synchronous system [3]. From a system design perspective, performance gain is achieved by increasing the number of processor cores per chip, resulting in a key bottleneck being formed due to the global intra-chip communications infrastructure [4]. Networks on-Chip become a natural choice as they reduce the complexity of interconnection design, separate computation from communication, support IP reuse and handle synchronization issues.

Corresponding Author: Khurram Aziz
NUST Publishing, © (2009), ISSN: 2070-9900

Interconnection Issues

Network on-Chip systems require high-capacity short-reach interconnections capable of operating at 10 Gbit/s and above. The ratio between interconnection delay and gate delay was about 2:1 at the 180 nm technology. This will increase to 9:1 at the 65 nm technology [5] as shown in Figure 3. Traditional interconnects that employ copper are no longer feasible because at this scale, the inductance effects cannot be ignored and lateral coupling becomes dominant, degrading the signal as crosstalk noise increases. Providing a reliable and low-skew clock throughout the system becomes a very challenging task. Moreover, the interconnect delay means that communication across the chip cannot take place in a single clock cycle. This makes fully synchronous systems impossible.

Copper lines are also bandwidth-limited. Large bandwidths can be attained through parallelism but this requires wide busses and large buffers. This in turn, increases the power requirements of the system. All these issues indicate that communication and not computation is the key issue in future NoCs and the power consumption, cost, area and performance will be much more influenced by on-chip interconnects.

While several solutions have been proposed for alleviation of these issues, for example [6-8], on-chip photonic interconnects are seen as the most promising.

Optical on-Chip Interconnects

Photonic interconnects are emerging as the main candidates to replace copper for on-Chip communication. They offer several distinct advantages over copper for high-speed future systems. Some of the key advantages are:

- Bit-rate transparency – The energy dissipation in optical interconnects does not depend on the bit-rate, as the photonic switches switch on and off once per message [4], unlike the dynamic power dissipation that scales with bit-rate in routers based on CMOS technology.
- Low loss – Energy dissipation in optical waveguides is independent of the on-chip distances. No signal loss due to crosstalk between coupled interconnects.
- High bandwidth – Photonics offers huge bandwidth benefits as wavelength division multiplexing can be employed.
- Less power consumption – In optical interconnects, data can be transmitted from end to end without repeating, regenerating or buffering, hence resulting in less power consumption.

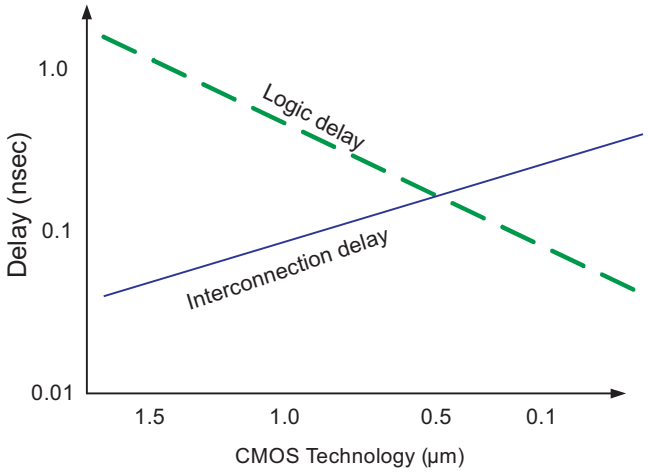

Fig. 3. Logic delay vs interconnection delay for different CMOS technologies different CMOS technologies

Optical interconnects do have their drawbacks, lack of buffer space being a major one. Signal processing abilities are also limited in photonics. Several approaches have been proposed in the literature to achieve the best of electronics and optical worlds by using hybrid approaches.

Clock Synchronization Issues

Several techniques have been discussed in the literature for overcoming the strict synchronization issues in fully synchronous systems, such as relaxation of the need for timing margins and multi-cycle on-chip communication [3, 9].

Another approach is to replace the electrical clock distribution tree by an optical one [10], thus eliminating the need for repeaters or clock multiplier circuits. This is not entirely possible yet as optical clock signals cannot be routed down to the single gate level. A hybrid approach is hence proposed in [10], in which an optical layer is integrated on top of a CMOS IC, as shown in Figures 4 and 5.

This approach has its own issues in terms of power dissipation, as much less power is dissipated in transmission for global and intermediate clock circuits as compared to local, hence the optical layer only eliminates a minority of the clock distribution power while introducing more stages of electrical buffers [11].

Another approach is the Globally Asynchronous Locally Synchronous (GALS) system, first introduced in [12]. In this approach, the system is divided into several independent synchronous local blocks that are then connected globally using asynchronous interconnection techniques to make the communication possible among the various locally synchronous blocks. Several strategies can be used for implementing the GALS technique. These include pausible-clock generators that ensure no clock pulses are generated during data transfer, asynchronous FIFO buffers between locally synchronous blocks, and performing boundary synchronization [13].

A Coarse-grained HEterogeneous Regionalized

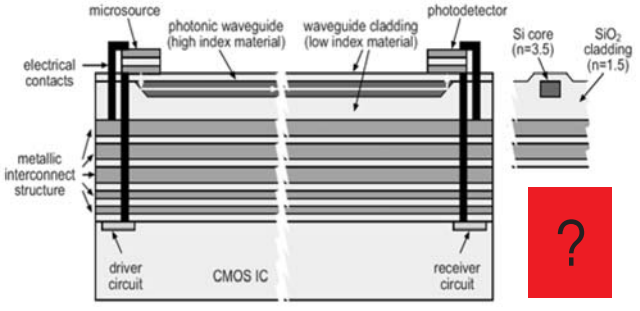

Fig. 4. [12] Hybrid interconnection structure with optical layer on top of CMOS IC

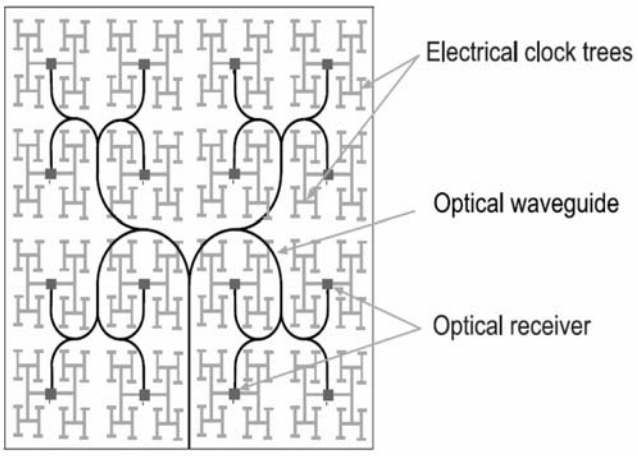

Fig. 5. [12] 1-16 point optical clock distribution tree

Adaptive GALS (CHERAG) is proposed in [14], which provides a comprehensive methodology for the design of GALS NoC-based Systems-on-Chip. A WLAN transmitter mapped to a CHERAG NoC is shown in Figure 6. In [15] an optical ring waveguide is proposed to replace global electrical interconnects.

Summary

Optical interconnects in Networks-on-Chip are still in their infancy and there are several open issues that need to be discussed. This paper highlighted the need for optical interconnects by discussing issues in copper interconnects. The clock synchronization issues were also discussed along with a survey of some existing techniques that address these issues. Better efficiency coupled with low energy consumption makes optical interconnects, the choice for future Network-on-Chip systems.

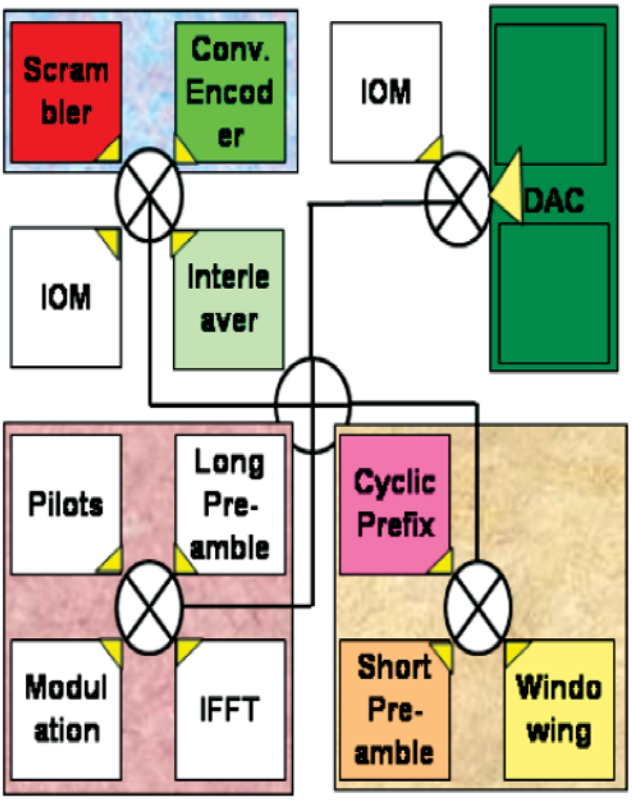

Fig. 6. [10] WLAN transmitter mapped to a CHERAG NoC with VFI regions

REFERENCES

1. L. Benini, G. De Micheli, "Networks on chips: A new SoC paradigm.", *DATE, Design, Automation and Test in Europe*, Paris, March 3 - 7, 2002
2. B. Ackland, A. Anesko, D. Brinthaupt, S. J. Daubert, A. Kalavade, J. Knobloch, E. Micca, M. Moturi, C. J. Nicol, J. H. O'Neill, J. Othmer, E. Säckinger, K. J. Singh, J. Sweet, C. J. Terman and J. Williams, "A single chip, 1.6-billion, 16-b MAC/s multiprocessor DSP," *IEEE J. Solid-State Circuits*, Mar. 2000, pp. 412-424.
3. T. Bjerregaard and S. Mahadevan, *A survey of research and practice of NoC*, ACM, 2006.
4. A. Shacham, K. Bergman and L. P. Carloni, "On the design of a photonic Network-on-Chip", *First International Symposium on Networks-on-Chip (NOCS)*, May 7 – 9, 2007.
5. Semiconductor Industry Association: *International technology roadmap for semiconductors* (2005)
6. E. G. Friedman, "On-Chip interconnect: The past, present, and future" *1st NoC Workshop*, March 2006
7. M. A. Horowitz, E. Alon, D. Patil, S. Naffziger, R. Kumar and K. Bernstein. "Scaling, power, and the future of CMOS." *IEEE Intl. Electron Devices Meeting*, Dec. 2005, pp. 9-15.
8. A. Naeemi, R. Sarvari and J. D. Meindl, "On-Chip interconnect networks at the end of the roadmap: Limits and nanotechnology opportunities", *International Interconnect Technology Conference (IITC)*, pp. 221-223, 2006
9. J. Cong, Y. Fan, G. Han, X. Yang, and Z. Zhang, "Architecture and synthesis for on-chip multi-cycle communication," IEEE Trans. On Computer-Aided Design of Integrated Circuits and Systems, pp. 550 – 564, 2004.
10. I. O'Connor and F. Gaffiot, "Advanced research in on-chip optical interconnects", *Lower Power Electronics and Design*, C. Piguet, Ed., , CRC Press,2004, pp 51-56
11. N. Ranganathan and N. P. Jouppi, "Evaluating the potential of future onchip clock distribution using optical interconnects", *Hewlett Packard White Papers*, January 2008.
12. D. M. Chapiro, "Globally-asynchronous locally-synchronous systems", *Ph.D. Thesis, Stanford University, CA. Dept. of Computer Science*, 1984.
13. M. Krstic, E. Grass, F. K. Gurkaynak and P. Vivet, "Globally asynchronous, locally synchronous circuits: overview and outlook", *IEEE Design & Test of Computers*, Vol. 24, Issue 5, Sept. – Oct. 2007
14. N. D. Gohar, J. M. Chabloz and A. Hemani, "Coarse-grained heterogeneous regionalized adaptive GALS NoC-based SoC: A flexible design methodology", accepted for publication in *First International Workshop on Network on Chip Architectures (NoCArc)*, Nov. 8, 2008
15. S. Pasricha and N. Dutt, "ORB: An on-chip optical ring bus communication architecture for multi-