

# The Transition to Chip-level Optical Interconnects

## Invited Paper

Michael W. Haney<sup>1</sup>, Daniel K. Sparacin<sup>2</sup>, Justin Hodiak<sup>2</sup>

<sup>1</sup>Defense Advanced Research Projects Agency, 3701 N. Fairfax Dr., Arlington, VA 22203

<sup>2</sup>Booz Allen Hamilton, 3811 N. Fairfax Dr.; Suite 600, Arlington, VA 22203

Mike.Haney@darpa.mil

**Abstract:** The increasing bandwidth demands of computing chips have spawned significant research into chip-scale photonic interconnects. We review the DARPA-sponsored efforts at the inter- and intra-chip interconnect levels and highlight the key challenges to be addressed.

©2010 Optical Society of America

**OCIS codes:** (200.4650) Optical Interconnects; (250.5300) Photonic Integrated Circuits

## 1. Introduction

Optical Interconnects (OI) are deployed in computing and signal processing systems ranging from Peta-Flop-class supercomputers (e.g., the IBM “Roadrunner”) with rack-to-rack optical cabling, to smaller military signal processing systems using optical backplanes. As CMOS IC performance continues to scale at the pace of Moore’s Law, on- and off-chip aggregate bandwidth requirements for high performance multi-core architectures are projected to hit ~10 TB/s by 2015 [1], with corresponding interconnect energy per bit (E/bit) budgets on the order of 100 fJ/bit [2]. Such bandwidth densities and E/bit requirements are problematic for the electrical chip-scale interconnect fabrics due to geometric and link performance limitations of metal interconnects for both on- and off-chip links; however the extrapolated performance of chip-scale photonic interconnects is thought to be within the required performance bounds [2-4].

As the level of OI integration continues to increase through the advancement of Photonic Integrated Circuit (PIC) technologies, we may now consider “flat” optical interconnect fabrics in which the chain of Electrical-to-Optical-to-Electrical transductions happens only once at any length-scale across the computer. Such a computing architecture would benefit from maintaining computation and communication bandwidth balancing, improved performance, and ease the programmability challenges of large systems – a critical issue impacting scaling of supercomputers to the Exascale [5]. The challenges in realizing such a flat architecture are three-fold: 1) Integration – realizing the necessary photonic technologies to enable dense, single-platform integration of energy-efficient chip-level OI, likely integrated directly within or on the microprocessor; 2) Packaging – development of alignment-tolerant optical interfaces that are compatible with microprocessor design and capable of high data throughput; and 3) Cooperative hardware design of the micro-processor and OI architectures – giving system designers the tools to optimize system performance.

## 2. DARPA Programs related to Optical Interconnects

The integration challenge has been a focus of R&D for more than a decade, witnessing a push to ever-shorter length use of optical interconnects. DARPA has played a significant role in the development of chip-scale photonic device and photonic integrated circuit technologies. The ability to leverage high-bandwidth optical communication to satisfy the needs of military applications and platforms has always been a large driver of this research as the size, weight, power, cost and environmental limitations of using available discrete components and technologies have traditionally been prohibitive factors in realizing this vision. As such, DARPA has sought to, and continues to, sponsor research in chip-scale, photonic devices with increasing levels of integration that will not only enable lower size, weight, power, and cost of photonic devices and functions – similar to that realized by the integration of discrete electronic components and functions in today’s microprocessors – but concurrently enable increased performance and new functionalities that would not be possible otherwise.

DARPA’s investments into chip-level OI have evolved along two paths. One path has focused on utilizing arrays of modulated VCSELs and photodetectors with out-of-plane coupling to enable parallel optical links with very high bandwidths. Starting in the mid-1990’s, DARPA funded the Parallel Optical Link Organization (POLO) program to develop low-cost, high performance, multi-Gb/s, parallel optical interconnect modules, based on VCSEL

## OMV2.pdf

technology, for server applications. This led to the Parallel Optical Network Interconnect (PONI) program, from which Agilent transitioned the technology to a commercially available 30 Gb/s optical link product comprised of 12 parallel, optical fiber links running at 2.5 Gb/s each. More recently, IBM, funded under the DARPA chip-to-chip optical interconnect (C2OI) program, has extended this approach to an inter-chip fabrics in which arrays of VCSELs and photodetectors, controlled with CMOS drivers, communicate through board-level, polymer waveguides and off-board polymer waveguide ribbons. In 2009, IBM demonstrated a 160 Gb/s optical link comprised of 16 parallel 10 Gb/s board-level, waveguide links (up to 30 cm) with CMOS control logic consuming only 4.4 pJ/bit [6] and is continuing to develop and scale this technology towards higher bandwidths with higher efficiency [3].

DARPA's other development path arose from the desire to miniaturize and ruggedize telecommunication technology for military communication systems. This research led to some of the first photonic integrated circuits (PICs) by integrating semiconductor-based, chip-scale, active and passive photonic components onto a single substrate. In 2001, the Chip-Scale Wavelength Division Multiplexing (CS-WDM) program commenced with the objective of significantly increase the number of WDM functions that could be integrated onto a single chip in order to radically reduce the footprint of components for fiber optic networks. Significant progress in the CS-WDM program showcased the potential for PIC technologies, including high power, chip-scale, tunable lasers with integrated 40 Gb/s modulators; multiplexors; and wavelength converters. These advancements set the foundation for the Data in the Optical Domain Network (DOD-N) program in 2004, which aims to develop energy efficient all-optical data routers, and the Electronic & Photonic Integrated Circuits (EPIC) program in 2006, which focused on developing CMOS-compatible photonic circuits from transceiver applications. Under DOD-N, UCSB researchers have arguably demonstrated the world's most photonically-complex PIC: a monolithic InP-based tunable router chip comprised of 8 SOA-based wavelength converters integrated with an 8x8 AWG with a 640 Gb/s routing capacity [7]. Similarly notable, Luxtera has commercialized their Si-photonics technology with monolithically integrated CMOS drivers, currently offering the lowest power active optical cable at 40Gb/s (4 x10 Gb/s links) with an overall energy efficiency of 20 pJ/bit for fiber lengths of up to 4 km [8]. This development was accelerated under the EPIC program. The Photonic Analog Signal Processing Engines with Reconfigurability (PhASER) program, initiated in 2007, continues the push toward increased density and functionality in PICs. In 2008, DARPA began funding the Ultraperformance Nanophotonic Intrachip Communications (UNIC) program to develop CMOS-compatible photonic technology for high-throughput, power efficient, intra-chip photonic communication networks. Here the vision is to employ many ultra-efficient (160 fJ/bit) optical links embedded within a microprocessor [4].

### 3. Discussion

At present there is no common "flat" OI architecture that spans all of the traditional packaging/interconnect levels of a high-performance computing system. In fact there are distinct contrasts between the C2OI multimode waveguide/VCSEL array concept being developed for the inter-chip domain, and the UNIC single-mode waveguide/WDM multiplexed concept for the intra-chip domain. Much of this discontinuity in architectures has to do with the way packaging trade-offs associated with interfacing micron-accurate OI elements on the chip address the 10's of micron accuracy required for elements on the chip package. For the C2OI concept, the use of lower bandwidth-density multimode waveguides (i.e., no WDM) is traded-off for more packaging-friendly alignment accuracy. In contrast, the UNIC approach aims to exploit the potential of high bandwidth, WDM signals on single-mode waveguides, but the complexity of the network is sensitive to the insertion loss of the intra- and inter-chip couplers. In both cases, packaging is a critical challenge driving the optical loss, thereby bounding the E/bit performance of the interconnect fabric, and therefore system architecture.

At a more abstract level, the optimal OI packaging solutions for a flat architecture will likely lead to unique optimizations, configurations, and topologies of the OI fabric which differ from the traditional electrical interconnection paradigm that has prevailed over the past couple of decades. Such alterations will likely require modification of the microprocessor for optimized system performance. As such, it is critical that the computing community develop tools for cooperative design of the microprocessor and OI fabric as a whole. PIC technology, which allows the dense integration of the required sources, modulators, multiplexors, detectors, and fabric to interface to CMOS processors, will enable this. Cooperative hardware design of micro-processor and OI architectures will intimately involve the tradeoffs that this technology presents. DARPA has invested heavily to prove that CMOS-compatible PICs can be successfully fabricated alongside CMOS electronics, however most CMOS foundries are averse to altering their process flow and in order to incorporate PICs. For cooperative hardware design to succeed, the CMOS fabrication requirements and trade-offs of incorporating photonics and

## OMV2.pdf

electronics must be considered and optimized to enable necessary widespread CMOS foundry access for PIC fabrication.

### 4. Conclusion

Future high-performance commercial and military computing architectures are projected to require 10's, and conceivably even 100's, of TB/s of bandwidth to the chip. The critical chip-scale interconnect challenges will therefore be common to high-performance computers ranging in size from a few chips to large supercomputing platforms with thousands of chips. For example, at the "high-end," we can imagine a supercomputer built within a decade from 10 TFLOP chips with an OI fabric that delivers ~10 TB/s of bandwidth to each chip. A Peta-flop class computer could be made from approximately 100 such chips, assuming the OI fabric capacity is sufficient. Assuming approximately 100 W power consumed per chip, such a computer may eventually be housed entirely within the volume of a single rack. However, such a "*Peta-flop-in-a-rack*" will be possible only if high-density, ultra-low E/bit, OI developed to handle the massive interconnect requirements, are embedded at the chip-scale. To achieve this, the ability to optimize computational performance in a flat architecture will require coordinated investment and development of PIC technology, packaging solutions, and cooperative hardware design of microprocessor and the OI fabric as a single architecture.

### 4. Acknowledgement

The authors are grateful for many useful discussions with Dr. Jagdeep Shah of DARPA/MTO, and for his insightful leadership of several of the programs mentioned in this summary.

### 5. References

- [1] ITRS Roadmap, see <http://www.itrs.net/Links/2007ITRS/Home2007.htm>
- [2] D. Miller, "Device Requirements for Optical Interconnects to Silicon Chips" Proceedings of the IEEE 97 (7), 1166-1185 (2009)
- [3] J.A. Kash, P. Pepeljugoski, F.E. Doany, C.L. Schow, D.M. Kuchta, L. Schares, R. Budd, F. Libsch, R. Dangel, F. Horst, B.J. Offrein, Y. Vlasov, W. Green, F. Xia, C.W. Baks, Y.J. Kwark, D. G. Kam, M.B. Ritter, "Communication Technologies for Exascale Systems" Proceedings of SPIE, Vol 7221, 72210F-12 (2009)
- [4] A.V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, J. Cunningham, "Computer Systems Based on Silicon Photonic Interconnects" Proceedings of the IEEE 97 (7), 1337-1361 (2009)
- [5] "ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems", P. Kogge, Ed., see [http://users.ece.gatech.edu/mrichard/ExascaleComputingStudyReports/exascale\\_final\\_report\\_100208.pdf](http://users.ece.gatech.edu/mrichard/ExascaleComputingStudyReports/exascale_final_report_100208.pdf)
- [6] C.L. Schow, F.E. Doany, C. Chen, A.V. Rylyakov, C.W. Baks, D.M. Kuchta, R.A. John, J.A. Kash, "Low-Power 16 x 10 Gb/s Bi-Directional Single Chip CMOS Optical Transceivers Operating at < 5 mW/Gb/s/link" IEEE J. Solid-State Circuitry 44(1), 301-313 (2009),
- [7] S.C. Nicholes, M.L. Masanovic, B. Jevremovic, E. Lively, L.A. Coldren, D.J. Blumenthal, "The world's first InP 8x8 monolithic tunable optical router (MOTOR) operating at 40Gbps line rate per port", 2009 Conference on Optical Fiber Communication – OFC 2009, pp. 1-3 (2009)
- [8] See <http://www.luxtera.com/>