Designing High-Performance, Power-Efficient NoCs With Embedded Silicon-in-Silica Nanophotonics

Elena Kakoulli, Vassos Soteriou, Charalambos Koutsides, Kyriacos Kalli
Department of Electrical Engineering, Computer Engineering and Informatics
Cyprus University of Technology, Limassol 3603, Cyprus
{elena.kakoulli, vassos.soteriou, c.koutsides, kyriacos.kalli}@cut.ac.cy

ABSTRACT

On-chip electrical links exhibit large energy-to-bandwidth costs, whereas on-chip nanophotonics, which attain high throughput, yet energy-efficient communication, have emerged as an alternative interconnect in multicore chips. Here we consider silicon nanophotonic components that are embedded completely within the silica (SiO_2) substrate as opposed to existing die on-surface silicon nanophotonics. As nanophotonic components now reside *subsurface*, within the silica substrate, non-obstructive interconnect geometries offering higher network throughput can be implemented. First, we show using detailed simulations based on commercial tools that such Silicon-in-Silica (SiS) structures are feasible, and then demonstrate our proof of concept by utilizing a SiS-based mesh-interconnected topology with augmented diagonal optical channels that provides both higher effective throughput and throughput-to-power ratio versus prior-art.

Categories and Subject Descriptors

B.4.3 [**Input/Output and Data Communications**]: Interconnections (Subsystems)

General Terms

Performance, Design

Keywords

On-chip Nanophotonics, Silicon-in-Silica, Topology

1 Introduction

On-chip nanophotonics have recently emerged to address the scalability challenges faced by all-electrical Networks-on-Chips (NoC) of chip multicores, offering higher bandwidth and shorter message delivery latency at manageable power budgets [1, 2]. Prior-art has demonstrated various hybrid silicon-based electrical/nanophotonic designs where all components, i.e., photonic, electrical, and optoelectronic, all reside on the silicon chip's surface [1, 2]. As such, two chip fabrication technologies exist: (1) where all components inhabit a single die, or (2) where optical and electrical components each occupy a dedicated silicon plane(s) that are connected vertically, forming an architecturally dense 3D die structure.

Such 3D architectures, however, entail both high design complexity and fabrication intricacy as die have to be perfectly aligned, along with exacerbated heat dissipation requirements. In this paper we consider an alternative architecture where silicon (Si) nanophotonic components are *embedded completely* in silica (SiO_2), i.e., reside *subsurface*, as opposed to existing *on-surface* silicon photonics. As nanophotonic components now reside *subsurface*, within the silica substrate, a greater portion of a chip's area can be utilized by cores and routers, while non-obstructive interconnect geometries offering higher network throughput can be fabricated with standard lithography methods, including elaborate topologies that contain and utilize diagonal waveguides (nanophotonic channels or links). The *contributions* of this paper are that: (1) we show using detailed simulations based on the RSoft commercial software tool [3] that such subsurface Silicon-in-Silica (SiS) structures are feasible, with optical signal integrity highly maintained, achieving up to 99% efficiency, (2) provide their design features, and (3) then demonstrate our proof of concept by using a SiS-based mesh-interconnected nanophotonic topology with augmented diagonal waveguides that attains higher effective throughput and throughput-to-power ratio vs. prior-art.

NOCS '15, September 28-30, 2015, Vancouver, BC, Canada
ACM 978-1-4503-3396-2/15/09.
http://dx.doi.org/10.1145/2786572.2786588

2 SiS-Embedded Nanophotonics: Background, Components Characterization and Design

Photonic NoCs comprise: (1) waveguides, i.e., photonic signal-carrying channels, that interconnect routers, (2) Micro-Ring Resonators (MRRs) that are connected to waveguides serving as either optical modulators or as filters for sending/extracting data to/from a waveguide, respectively, and (3) transceivers that mark the electro-optical boundary, converting signals from the optical domain to the electrical domain (to be utilized by routers) and vice-versa.

MRRs are key to the design of optical NoCs. They act as modulators, where optical data can be inserted into a bus waveguide that is coupled to the MRR's closed-loop waveguide at a specific wavelength through electro-optical modulation. They also act as versatile resonant optical filters, extracting light of a specific wavelength from an input bus waveguide carrying optical data to an output bus waveguide that is coupled to the same closed-loop waveguide at its opposite side. A photo-detector then measures the light intensity at the device output for a given filtered wavelength. Since MRRs resonate at distinct wavelengths, and remain quiescent otherwise, this behavior serves as the primary mode in realizing ultra-high bandwidth wavelength-division-multiplexing, where waveguides simultaneously carry multiple optical signals without interference.

For resonance to occur in MRRs the condition $2\pi r n_{eff} = m\lambda_m$ must be met, where n_{eff} is the effective refractive index of the ring (implemented in silicon) that must be greater than the surrounding material (silica in our case vs. air in conventional silicon nanophotonics implemented onto the silicon die surface), r is its radius, m is the mode number of the MRR (a positive integer number), and λ_m is the resonant wavelength. Unlike previous works in which silicon nanophotonic components reside on the silicon surface along with electrical components [1, 2], here we take a substantially different fabrication approach where MRRs and waveguides made of silicon are embedded within the silica surface, forming a Silicon-in-Silica (SiS) structure, with electrical components placed onto the silica surface on top of them. Our SiS structure consists of two waveguides and a ring resonator made of pure silicon (Si) embedded within a bulk volume of silica (SiO_2) using lithographic techniques described by White et al. [4]. To implement our SiS structure we calculated the time-dependent electromagnetic field, using the RSoft commercial software tool [3]. For optimal results we carefully defined the dimensions and refractive index profile of the proposed subsurface structure. With regard to the behavior of the electromagnetic field at the differing material boundaries, we applied the perfectly matched layer bound-

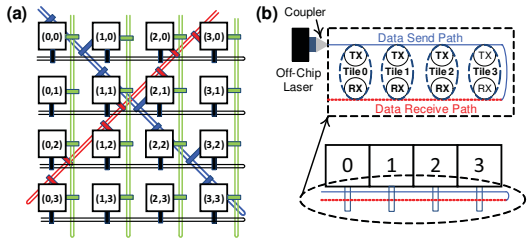

Figure 1: (a) a 4×4 topology with horizontal (black) and vertical (green) subnets augmented with diagonal (red) and anti-diagonal (blue) subnets, (b) a subnet of 4 nodes with transmit (TX) and receive (RX) groups of rings (circles) related to each tile (ovals).

ary conditions. The following structural data were found or set for our SiS design: working wavelength of 1.593635 μm, MRR ring diameter of 3.6 μm, waveguide length of 5 μm and width of 0.2 μm, component width of 4 μm, where the separation distance between the waveguide and the ring resonator is 0.05 μm; this gap leads to optimum coupling between the waveguide and the ring. The SiS MRR works effectively, extracting 99% of the waveguide-propagated light power due to the high refractive index contrast between silicon and silica, with just 0.98 dB/cm waveguide loss. Lastly, our SiS structure can work bi-directionally, i.e., either as a light receiver or as a light transmitter.

3 Mesh Topology With Augmented Diagonal Waveguides and Adaptive Routing Algorithm

Our interconnect architecture utilizes SiS-embedded waveguides and MRRs exclusively. To avoid the large power overheads of chip-spanning optical channels utilized in prior-art [2], we instead break the network into shorter and hence power consumption-limited bus waveguides, a.k.a. "subnets," similar to those used in the LumiNOC architecture [1]. As waveguides can now be fabricated at any angle within the silica substrate, i.e., *subsurface*, we extend the mesh-like subnet topology of LumiNOC to utilize *diagonal* waveguides that are augmented to the base mesh interconnect. Figure 1-(a) shows such a symmetric 16-node realization, as a proof of our concept, where a Diagonal (D) subnet and an Anti-Diagonal (AD) subnet are augmented to the base array-interconnected topology that comprises Horizontal (H) and Vertical (V) subnets. Each optical-electrical router assumes a conventional 4-stage pipelined organization, with (1) Optical to Electrical (O/E) conversion and routing computation, (2) Virtual Channel (VC) arbitration and photonic subnet allocation, (3) switch allocation, and (4) crossbar traversal and Electrical to Optical (E/O) conversion; finally, waveguide propagation follows. Each such router either uses two (H+V), or three (H+V+D, or H+V+AD), or all four (H+V+D+AD) input/output ports, according to its waveguide bus-connecting count, plus the processor injection/ejection e-ports. Each such input port extracts data using its SiS waveguide-MRR structure, and converts it to electrical signals (using a local O/E block) to be absorbed by the router (or to be further relayed); a router relays or sends optical data, pre-converted from their electrical form by a local E/O block, through its output ports, and are then modulated and inserted into the subnet via the associated MRR.

We utilize LumiNOC's [1] flow-control protocol, where each waveguide double-backs (loops) from the external laser source, connects to each router input port (acting as the receiver), and then to each router's output port (acting as the transmitter). In this way a single subnet, which contains as many wavelengths as the number of routers it connects, is used first to arbitrate which router should gain access to the bus waveguide, using one-hot encoded control messages as in LumiNOC, and in the second phase the winner router with granted permission transmits its data (see Figure 1-(b)).

Our SiS-based topology presents opportunities for routing optimization, so as to maximize network throughput, since in most cases alternative progressive routing paths may exist, where at most, a packet has to traverse a two-hop route along a path. As such, we develop an adaptive photonic routing algorithm where dynamic metrics, in terms of the number of pending arbitrations per alternative waveguide, are considered. The waveguide, or waveguide pair, with the least such number is chosen as the projected fastest route.

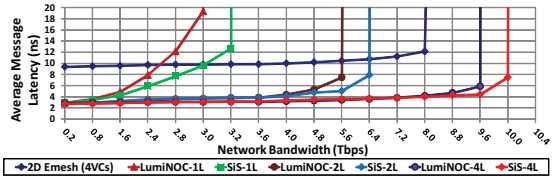

Figure 2: Latency-bandwidth curves under bit-complement traffic. The term "L" denotes the number of photonic layers.

4 Experimental Setup and Results

We implemented a cycle-accurate microarchitectural-level simulator that supports (1) 2D all-electrical meshes (EMesh), (2) the LumiNOC architecture, and (3) our proposed SiS-based topology; all utilize four-stage pipelined routers with input Virtual Channels (VCs) (see Section 3), and each comprises a 64-node CMP with its tiles arranged in an 8 × 8 array - the SiS-based topology was scaled linearly using the base 4 × 4 SiS topology of Figure 1-(a). We assume a 10 GHz waveguide modulation rate, with routers clocked at 5 GHz [1]. Synthetic bit-complement traffic traces, with packets comprising four 128-bit flits, are utilized. Figure 2 shows our SiS topology outperforming LumiNOC under all equivalent photonic layer counts. All photonic interconnects outperform EMesh in terms of latency before reaching their saturation, while SiS-4L and LumiNOC-4L (4 optical layers), also outperform EMesh in terms of sustainable throughput.

We next estimate both the electrical power consumption of NoC routers, using the Orion 2.0 power models library [5] at 45 nm CMOS process technology, and the photonic energy consumed using the Corona [2] architecture power models, except for where SiS parameters are available. We focus on the realistic throughput-per-power metric, in terms of Tbps/W. Under 1, 2, and 4 photonic Layers ("L"), our SiS-based topology achieves respective performances of 3.28 Tbps/W, 2.67 Tbps/W, and 2.58 Tbps/W, while LumiNOC equivalently attains 3.29 Tbps/W, 2.48 Tbps/W, and 2.41 Tbps/W; our SiS-based architecture, except with 1L, achieves both improved performance and efficiency vs. LumiNOC.

5 Conclusions

We presented a photonic on-chip interconnect where silicon nanophotonic components are embedded in silica. We first characterized such SiS structures and then constructed a photonic mesh NoC with augmented diagonal waveguides that provides higher throughput-to-power versus prior-art.

Acknowledgments

This work falls under the Cyprus Research Promotion Foundation's Framework Programme for Research, Technological Development and Innovation 2009-10 (ΔΕΣΜΗ 2009-10), co-funded by the Republic of Cyprus and the European Regional Development Fund, and specifically under Grant ΠΕ/ΠΕΝΕΚ/0311/06.

6 References

[1] C. Li et al. *LumiNOC: A Power-Efficient, High-Performance, Photonic Network-on-Chip*. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol. 33, No. 6, pp. 826-838, June 2014.

[2] D. Vantrease et al. *Corona: System Implications of Emerging Nanophotonic Technology*. Proc. 35th ACM/IEEE International Symposium on Computer Architecture, pp. 153-164, June 2008.

[3] RSoft Optical Software Solutions, Synopsys Inc. Available [online] http://optics.synopsys.com/rsoft/, 2014.

[4] T. P. White et al. *Silica-Embedded Silicon Photonic Crystal Waveguides*. Optics Express, Vol. 16, No. 21, pp. 17076-17081, Oct. 2008.

[5] A. B. Kahng et al. *ORION 2.0: A Power-Area Simulator For Interconnection Networks*. IEEE Transactions on Very Large Scale Integration Systems, Vol. 20, No. 1, pp. 191-196, 2011.