

Area and Laser Power Scalability Analysis in Photonic Networks-on-Chip

Sergi Abadal^{*1}, Albert Cabellos-Aparicio^{*}, José A. Lázaro[†], Mario Nemirovsky[‡],
Eduard Alarcón^{*} and Josep Solé-Pareta^{*}

^{*}NaNoNetworking Center in Catalonia (N3Cat)

Universitat Politècnica de Catalunya, Barcelona, Spain

[†]Optical Communications Group (GCO)

Universitat Politècnica de Catalunya, Barcelona, Spain

[‡]ICREA Senior Research Professor at

Barcelona Supercomputing Center (BSC), Barcelona, Spain

Abstract—In the last decade, the field of microprocessor architecture has seen the rise of multicore processors, which consist of the interconnection of a set of independent processing units or cores in the same chip. As the number of cores per multiprocessor increases, the bandwidth and energy requirements for their interconnection networks grow exponentially and it is expected that conventional on-chip wires will not be able to meet such demands. Alternatively, nanophotonics has been regarded as a strong candidate for chip communication since it could provide high bandwidth with low area and energy footprints. However, issues such as the unavailability of efficient on-chip light sources or the difficulty of implementing all-optical buffering or header processing hinder the development of scalable photonic on-chip networks. In this paper, the area and laser power of several photonic on-chip network proposals is analytically modeled and its scalability is evaluated. Also, a graphene-based hybrid wireless/optical-wired approach is presented, aiming at enabling end-to-end photonic on-chip networks to scale beyond thousands of cores.

Index Terms—Network-on-Chip; Scalability; Nanophotonics; Area; Laser Power; Silicon-on-Insulator; Graphene; Hybrid

I. INTRODUCTION

In the ever-changing world of microprocessor design, multicore architectures are currently the dominant trend for both conventional and high-performance computing. Unlike in single-core designs, the performance of such processors is mainly determined by the capabilities of the on-chip network that interconnects its cores, usually referred to as Network-on-Chip (NoC). The communication requirements greatly increase as more cores are integrated in the same architecture, to the point of expectedly becoming the main performance bottleneck of multiprocessors.

Indeed, scaling a NoC to hundreds or thousands of cores presents important challenges in terms of bandwidth, area and energy: the NoC must provide enough bandwidth in order to support the non-linear increase in traffic, while maintaining affordable power and area overheads. Such challenges need to be addressed at the interconnect level first, as the performance of a given interconnect technology has a major impact on the overall network performance regardless of its architecture.

Conventional NoCs consist of a network of electrical on-chip wires and routers that convey the information from the transmitting core to the receiving core. However, it remains unclear whether such copper-based architectures will scale beyond several tens of cores, simply due to the energy efficiency of the underlying on-chip wires [1]. In this context, the breakthroughs accomplished in nanophotonics have not gone unnoticed, and the employment of optical on-chip communication has been proposed mainly due to its outstanding bandwidth and energy consumption capabilities with CMOS compatibility and reduced area footprint [2]. A summary of the state of the art in nanophotonics is presented in Section II.

Despite the huge potential of nanophotonic interconnects, the design of a scalable Photonic Network-on-Chip (P-NoC) remains as an important challenge. This is mainly due to two particular issues of this scenario, namely, the difficulty of implementing all-optical functions such as buffering or header processing, and the unavailability of efficient on-chip light sources. As a result, P-NoCs generally scale poorly due to overcomplexity.

In this work, we quantify such trend by reviewing several existing P-NoC proposals (references [3]–[6], see Section III). By means of simple analytical models, which are explained in Section IV, we first evaluate scalability of such architectures in terms of area and laser power, and then we briefly discuss their potential tradeoffs (Section V). Finally, in Section VI we introduce a hybrid wireless/nanophotonic architecture driven by graphene, aimed at enabling the design of scalable P-NoCs beyond thousands of cores for future generation multiprocessors. Section VII concludes the paper.

II. STATE OF THE ART IN NANOPHOTONICS

Nanophotonics has been proposed as the most promising technology for, on the one hand, providing small footprint devices to be integrated at the chip level, and on the other hand, potentially addressing the challenging targets of transmitter energies of 10-100 fJ/bit [1]. Several technologies are under intense research for low-energy, low drive voltages and compatibility with CMOS technology, which dominates consumer

¹Email: abadal@ac.upc.edu

TABLE I
OPTICAL PARAMETERS

Parameter	Value	Units	Ref.
Active/Passive Rings			
Pitch	8	μm	[5], [14]
Drop Loss	1/0.5	dB	[11], [12]
Pass Loss	0.01	dB	[15]
Waveguides			
Pitch	2	μm	[5], [8]
Propagation Loss	0.5	dB/cm	[8]
Bending Loss (radius: 10 μm)	0.15	dB/bend	[8]
Others			
Modulation Loss ($P_0 = P_1$)	3	dB	-
Number of Wavelengths	64	-	[16]
Datarate per Wavelength	10	Gbps	[16]
Splitter Excess Loss	0.04	dB	[9]
Switch Insertion Loss	1	dB	[6], [15]
Photodetector Area	20	μm^2	[3]
Photodetector Sensitivity	-30	dBm	[13]

electronics.

A. Nanoscale Silicon Photonics

The advancement of silicon as optical transmission medium, generally in conjunction with an insulator (e.g. silicon dioxide), has given rise to an outstanding progress in the development of the necessary building blocks for the creation of on-chip optical interconnects [2]. Several types of waveguides using silicon-on-insulator technology have been demonstrated [7], [8], with propagation losses ranging from 0.2 to several dB/cm and cross section dimensions below a few μm . Bends, crossings and power splitters have been also demonstrated [9]. For the application this paper is aimed for, components based on ring resonators are of special importance since they can be tuned to a unique target wavelength. This way, selective modulation [10], filtering [11] or switching [12] of a wavelength division multiplexed (WDM) signal can be achieved. Finally, photodetectors with responsivities of up to several A/W have been demonstrated [13]. With such responsivity, receivers with extremely low sensitivity can be achieved. In this work, we consider -30 dBm as a conservative figure for the usually required values of BER at tens of gigabits per second. This and other projected but conservative values, which will be used throughout the paper, are summarized in Table I.

B. Graphene Nanophotonics

Graphene is another material providing outstanding properties for nanophotonics, such as the capability to propagate strong confined light in the form of plasmons [17] or a strong gate-activated photoresistance [18]. Being graphene technology compatible with CMOS circuitry, devices with enormous potential for on-chip optical communication can be envisaged [19]. A first graphene-based optical modulator has been demonstrated showing an electro-optical bandwidth above 1 GHz at a broad spectrum from 1.35 to 1.6 μm [20]. It represents a significant step towards reduced footprint devices as the modulator area is reduced to only 25 μm^2 , thanks to the high absorption of graphene leading to a modulation capability

of 0.1 dB/ μm . At the receiver side, a graphene transistor-based photodetector at 40 GHz has been demonstrated, while the analysis of the results suggests potential intrinsic bandwidths in the order of 500 GHz [21].

III. PHOTONIC NETWORKS-ON-CHIP

Ever since the advent of nanophotonic interconnects, considerable efforts have been put into leveraging their outstanding properties to design on-chip networks for next generation multiprocessors. Conventional routed mesh designs are not feasible, since buffering and header processing must be done in the electrical domain, incurring high losses due to optoelectric conversions. Instead, several teams have proposed a wide variety of end-to-end optical schemes (see [3]–[6], [12], [22], among others). In the following, we introduce the architectures whose scalability will be analyzed in Section V:

ATAC [3]: strongly influenced by the work in [23], this architecture implements a full crossbar by replicating a single-writer multi-reader (SWMR) structure. Each transmitter is tuned to a unique wavelength and broadcasts its messages through a serpentine waveguide that connects all N cores, each of which must account for $N - 1$ detectors in the receiving end. Although the transmission medium is shared, contention is avoided since WDM is employed. Multibit transmission can be achieved by means of waveguide replication.

CORONA [4]: this architecture implements a full crossbar by replicating a multi-writer single-reader (MWSR) scheme, as opposed to ATAC. Light power is split and guided into N different data waveguides, in each of which a unique core is able to receive data modulated by any of the other cores. In this case, two cores will contend if they want to communicate with the same destination. To avoid possible collisions, CORONA implements an all-optical token-ring-based arbitration scheme. Finally, multibit transmission is achieved by means of WDM.

DCOF [5]: rather than an architecture, DCOF can be regarded as a family of contention-free architectures implementable upon a full waveguide mesh. To radically reduce the number of waveguide crossings, multiple vertically coupled layers are employed. A parameter k determines to how many destinations each core is able to concurrently communicate, affecting both the total bandwidth and the power delivery to the transmitter. If $k = 1$, a crossbar is instantiated upon DCOF, while if $k = N - 1$ a full mesh is obtained. As in CORONA, multibit transmission is achieved by means of WDM.

Photonic Mesh [6]: in this proposal, cores are interconnected through a network of photonic switches linked by torus-like waveguides. The switches are driven by a parallel control plane that is implemented by means of a conventional electrical NoC. Contention is avoided by means of circuit-switching, wherein the path between transmitter and receiver is setup by the control plane and maintained during the whole optical transmission. Unfortunately, this means that the number of simultaneous transmissions, and therefore the available bandwidth, is strongly limited. In order to alleviate this issue, path multiplicity by means of component replication

[6], time-multiplexed interleaving [22] or wavelength-based routing [12] has been proposed.

IV. ANALYTICAL MODELS FOR P-NOC

In this section, we detail the well-known area and laser power models that will be employed in the scalability analysis.

A. Area Overhead

The area occupancy of a given optical architecture can be evaluated as the sum of the area of each individual component. This essentially includes modulators, waveguides, switches, filters and photodetectors. In the transmitting side, we will assume that modulators are made of one active ring resonator, whereas receivers will consist of a passive ring resonator-based filter and a photodetector. We also consider that all ring resonators are of the same size. Given these assumptions, the area of a given architecture can be approximated as:

$$A \approx N_{ring}A_{ring} + N_{det}A_{det} + \sum_i A_{wg,i}$$

where N_{ring} and N_{det} are the number of ring resonators and photodetectors. $A_{ring} = W_{ring}^2$ is the area of each ring, or the square of its pitch; while A_{det} is the photodetector area. The last term accounts for the area of each and every waveguide in the network, calculated as the product of its length L by its pitch W . As shown in Table II, the specific architecture of a given P-NoC will determine the number of components needed to implement it, generally as a function of parameters such as the number of cores, the link bandwidth or the total network bandwidth. For instance, the number of wavelengths that can be generated and transmitted in ring resonator-based interconnects is limited by the bandwidth of each wavelength and the free spectral range (FSR) of the ring resonators, given a maximum admissible crosstalk [16]. As a given architecture scales, links might require the employment of W_D parallel waveguides in order to accommodate a given number of wavelengths:

$$W_D = \left\lceil \frac{N_\lambda}{N_{\lambda,MAX}} \right\rceil$$

where N_λ is the number of wavelengths of the network and $N_{\lambda,MAX}$ is the maximum number of wavelengths. The $\lceil \cdot \rceil$ operator rounds the result upwards to the nearest integer. It is important to note that N_λ may depend on the number of cores of the network [3] or the targeted network capacity [4].

B. Laser Power

The power consumed in a P-NoC has diverse components, such as the energy required to drive the modulators or to perform optoelectrical conversion in a photodetector. However, the laser power will expectedly become the dominant power component in dense P-NoCs due to its poor scalability, which is shown below. A laser may consume several tens of Watts in this case, while transmitters and receivers with a combined efficiency of around 100 fJ/bit [12] would need to steadily process hundreds of Tbps of data to reach such figures. Therefore, in this work we will focus on the laser power.

TABLE II
OPTICAL COMPONENT COUNTS FOR DIFFERENT P-NOC ARCHITECTURES

	ATAC	CORONA	DCOF(k)	P-Mesh
Data Wguides / W_D	W	N+1	$N^2 + \frac{W\sqrt{N}}{W_D}$	$2\sqrt{N}$
Active Rings / N	W	WN+N	W N	$W+32W_D$
Passive Rings / N	W N	W+N	$W(N+f(k))$	W
Photodetectors / N	W N	W+N	W N	W

[N: number of cores, W: datapath width, W_D : waveguide replication]

Since integrating individual laser sources on a chip is currently unviable, it is commonly considered that light will be generated by an external multiwavelength source, coupled into the chip and then guided within the P-NoC. Also, the laser power is statically allocated as there is no feedback from the chip to the source. In this context, the procedure to evaluate such allocation is to perform a power budget, this is, to calculate the minimum laser power for which the power at any receiver is higher than its sensitivity. Since different wavelengths might account for different critical paths, the laser power at the output of the coupler must be expressed for each wavelength j :

$$P_{lstr,j}(dBW) = \left[S_{RX}(dBW) + \sum_i L_i(dB) \right]_j \quad (1)$$

where $P_{lstr,j}$ is the on-chip laser power needed for wavelength j , while S_{RX} is the receiver sensitivity and L_i is the loss of component i in the critical path of wavelength j . Apart from a loss of 3 dB in modulation due to symbol equiprobability, we will consider waveguide propagation, bending and crossing losses, ring resonator losses and the photodetector efficiency, as well as splitting losses if wavelength j must reach multiple destinations. The specific architecture of a given P-NoC will determine the number and type of components present in the critical path of each wavelength. If we express equation 1 in linear units, the total on-chip laser power requirements will be:

$$P_{lstr}(W) = \sum_j \left[S_{RX}(W) \prod_i L_i \right]_j \quad (2)$$

where $\prod(\cdot)$ denotes product and perfect laser power allocation is considered. Finally, in order to calculate the power consumed by the laser, the the laser wall-plug efficiency and the on-chip coupling efficiency have to be taken into consideration. Prospective values for these concepts are around 30% and 90%, respectively [12].

V. SCALABILITY ANALYSIS OF SELECTED P-NOC ARCHITECTURES

The analytical models introduced in the previous section are instrumental the main purpose of this paper: to perform a partial design space exploration of P-NoC based on the data available in the literature. Our aim is to show and compare area and laser power scalability trends among a set of architectures, regardless of their potential performance for a given application. To achieve this goal, parameters related




Fig. 1. Area scaling as a function of the number of cores considering end-to-end optical communication, for different architectures. The datapath width is fixed to $W = 32$.

to the size and losses of individual components are as summarized in Table I and will remain fixed for all architectures. The models have been also adjusted to yield the same link bandwidth in all architectures by fixing the datapath width. We acknowledge that results shown below might therefore slightly differ from that of the original works, and that it results in an unfair comparison from a throughput perspective. A throughput scalability study is expected in future work.

A. Area Overhead

Table II shows the number of components in the selected architectures according to our analysis. The number of waveguides is normalized to the replication factor, whereas the number of active rings, passive rings and photodetectors are normalized to the number of cores. It is worth noting that:

- The number of components is multiplied by the target datapath width W , except for the waveguides in designs where multibit transmission is performed through WDM.
- The need for arbitration in CORONA results in the addition of a waveguide and N modulators and detectors.
- DCOF needs $\frac{W\sqrt{N}}{W_D}$ extra waveguides for wavelength dropping in transmission. Also, the number of passive rings used to deliver individual wavelengths to each group of modulators depends on k as: $f(k) = k$ if $k < N/2$, or $f(k) = (N - 1) - k$ otherwise.
- In Photonic Mesh, four 4-port switches are needed per each core. Such switches are implemented using eight active ring resonators [6].

Figure 1 shows the evolution of the area footprint of each analyzed architecture as a function of the number of cores and normalized to the area of a 20x20 mm² die. In other words, considering perfect component distribution over multiple layers, Figure 1 evaluates the minimum number of layers needed for each specific P-NoC as a function of the number of nodes. The datapath width is fixed to $W = 32$, leading to a link bandwidth of 320 Gbps.

Despite needing an electrical control NoC, the area of which is calculated using ORION [24], our analysis shows that P-Mesh can be considered the only scalable option in terms of area. The other architectures have a quadratic behavior, as all of them count on at least one component scaling

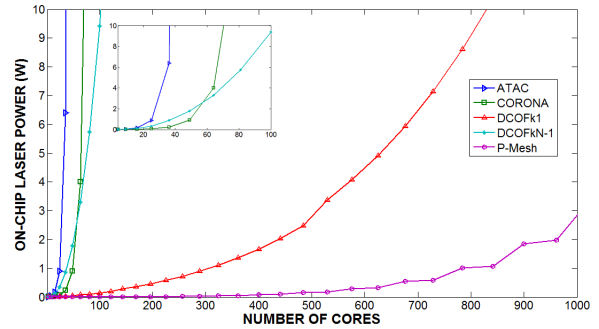


Fig. 2. Laser power scaling as a function of the number of cores considering end-to-end optical communication, for different architectures. The inset shows the same metric. The datapath width is fixed to $W = 32$.

as $O(N^2)$. DCOF performs specially poorly in this sense: since waveguides are dedicated, all cores must have both N transmitters and N receivers for all values of k .

B. Laser Power

Let us consider that, in all architectures, light is coupled into the chip and then split into two components that are routed to half of the nodes each. Upon reaching the transmitting core, a portion of this light is routed into a data waveguide and modulated. The modulated wavelengths propagate through the data waveguide towards its destination, probably passing through several non-intended receivers. At that point, light might be filtered to extract a given wavelength, photodetected and passed to the electrical receiver.

The optical power that must be delivered by the laser is calculated using equation (2). One can see that the result will be highly dependent on:

- *The link losses.* The kind reader will observe that losses are multiplicative. Therefore, the required laser power will scale exponentially if the number of components in the critical path grows with the number of cores. In fact, the laser power can be expressed as the product of exponentials ($P \sim \prod_i \alpha_i^{\beta_i}$, with β_i potentially being a function of N).
- *The number of possible simultaneous transfers.* We can consider that the power requirements grow linearly with the maximum number of simultaneous destinations, due to the increase of splitting losses.

These remarks will be useful to explain Figure 2, which shows the evolution of the required on-chip laser power as a function of the number of cores in the selected architectures.

One can see that ATAC and CORONA scale poorly mainly due to the N ring resonators per core, leading to an exponential penalty despite the relatively low passing losses of such components. Moreover and while both architectures allow N simultaneous transmissions, ATAC is designed to reach $N - 1$ destinations in each transmission, as opposed to CORONA's transmissions which are point-to-point. Therefore, ATAC needs additional laser power.

Despite counting on the largest component count, the power figures in DCOF are moderate since, unlike in other




Fig. 3. Laser power scaling as a function of the number of cores in the CORONA architecture, for different ring resonator pass loss values. The datapath width is fixed to $W = 32$.

architectures, the cores between transmitter and receiver do not contribute to the link losses. Still, results are highly dependent on the parameter k , which determines the number of simultaneous destinations per node.

Finally, P-Mesh shows good laser power scalability since, in the absence of path multiplicity, only one simultaneous transmission is possible. In the implemented torus topology, the maximum hop count between two nodes grows as $O(\sqrt{N})$. As path losses mainly depend on this factor, the power requirements scale better in P-Mesh than in other architectures.

C. Trade-off Discussion

The design of a P-NoC implies having to carefully evaluate certain trade-offs at the interconnect and network levels.

Power-Area: While ATAC and CORONA show large albeit affordable area overheads, their power requirements will render them unfeasible for high network sizes until technology advancements allow a dramatic reduction of the dominant loss component (see Fig. 3). DCOF with $k = 1$ reduces such losses and achieves scalable laser power by employing dedicated waveguides, at the expense of being unfeasible in terms of area. P-Mesh shows excellent scalability results in both metrics.

Power/Area-Throughput: Even though P-Mesh shows the outstanding area and power scalability, its limited bandwidth will probably become a roadblock for future multiprocessors. Allowing a larger number of simultaneous transmissions would increase the available bandwidth at the expense of higher laser power requirements, as commented in Section V-B. Moreover, additional area is needed to accommodate the infrastructure for such transmissions, e.g. complex routers or additional switches. The question here is, at which point and under which conditions, having more bandwidth does not necessarily imply a higher throughput. In future work, we expect to further investigate this trade-off.

Power/Area/Throughput-Cost: Another important aspect to consider is the cost of a given architecture. A large number of optical passive elements can be included without a significant increment of the cost. In contrast, including active elements represents a higher cost due to more complex fabrication processes and wire-bonding post-processing. Consequently,

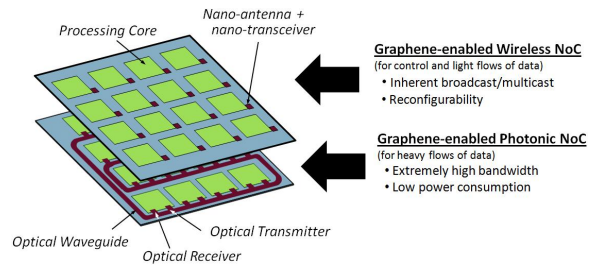


Fig. 4. Example of a graphene-enabled hybrid wireless-optical NoC.

architectures accounting for a large number of active elements may have unaffordable costs and may be discarded.

VI. A GRAPHENE-ENABLED HYBRID ON-CHIP NETWORK ARCHITECTURE

Drawing an analogy from optical access networks, some authors propose to scale P-NoCs by employing optical communication as the *backbone* of the network, whereas conventional on-chip wires are used in the *edges* [3], [25]. In other words, P-NoCs would provide communication among clusters of cores instead of doing it at the core level. However, the performance of the electrical side might become a bottleneck in such clustered approach. Alternatively, our ultimate goal is to enable the creation of P-NoC architectures which are both scalable beyond thousands of cores and capable of delivering optical communication at the core level. In this regard, we propose a graphene-enabled wireless/optical approach, conceptually represented in Figure 4):

Nanophotonic Plane: In our approach, a P-NoC will be implemented by means of state-of-the-art nanophotonic components for the efficient transmission of heavy flows of data. Graphene-based components could be employed due to their expectedly reduced footprint [20]. Broadband modulators could replace entire banks of ring-based modulators in some cases, leading to a substantial reduction in area. Also, if graphene nanophotonics would demonstrate ultrafast optical modulation and switching at lower losses, acceptable laser power values for high node counts could be obtained in a wide variety of architectures.

Wireless Plane: Due to the stringent requirements of the on-chip scenario in terms of area and bandwidth, wireless communication at the core level cannot be implemented by means of conventional metallic antennas. Instead, the wireless plane can be enabled by graphene-based nano-antennas due to their size and potential bandwidth [26]. Preliminary results show that, due to the plasmonic effects present at the surface of a graphene patch, a graphene antenna of a few microns would be able to radiate in the Terahertz band, this is, two orders of magnitude lower than a metallic antenna of the same size. We refer the interested reader to [26], [27] for more details.

On the one hand, the wireless plane could be used for the transmission of selected flows of data. As discussed in [3], [28], a multicore processor would greatly benefit from efficient broadcast and all-to-all communication capabilities. Wireless communication offers such possibility without the

need of wiring infrastructure and provided that a contention mechanism is implemented. Fortunately, wireless medium access control is a well-researched area and efficient solutions can be expected in this respect.

On the other hand, the wireless plane can be used to control the nanophotonic plane, following a similar approach than in [6]. Contention is resolved in a cost-effective way, avoiding the need for other arbitration schemes or contention-free architectures that consume considerable area and power. For instance, CORONA requires $2N^2$ ring resonators for arbitration, leading significant area figures in large networks. In P-Mesh, the wireless plane would replace the electrical control NoC, enhancing performance since the path could be setup in a single broadcast transmission. We could also improve the throughput of the system by setting up several non-intersecting paths using such broadcast scheme. In future work, we will quantitatively investigate the possible impact of introducing a wireless control plane in such architectures.

VII. CONCLUSIONS

We have shown, at least for a set of selected architectures, that meeting certain communication requirements in terms of bandwidth and connectivity implies a large number of on-chip components and, consequently, limited scalability in area or power. Reducing the area footprint or the losses of certain key components will have a significant impact in such metrics, but might not suffice. We have also analyzed some of the technological and architectural trade-offs that can be found in the inspected architectures. Finally, we have proposed to adopt a hybrid approach aiming at the creation of scalable P-NoCs, wherein a wireless network will both control the nanophotonic plane and transmit selected flows of data. In future work, we expect to analyze and quantify the impact of this approach to the scalability of existing and novel P-NoCs.

ACKNOWLEDGMENT

This work has been partially funded by Generalitat de Catalunya (SGR 2009-1140) and by the Spanish MINECO under the projects TEYDE (FEDER TEC2008-01887) and DOMINO (TEC2010-18522).

REFERENCES

- [1] D. A. B. Miller, "Device Requirements for Optical Interconnects to Silicon Chips," *Proceedings of the IEEE*, vol. 97, no. 7, pp. 1166–1185, Jul. 2009.
- [2] R. G. Beausoleil, P. J. Kuekes, G. S. Snider, S.-y. Wang, and R. S. Williams, "Nanoelectronic and Nanophotonic Interconnect," *Proceedings of the IEEE*, vol. 96, no. 2, pp. 230–247, Feb. 2008.
- [3] G. Kurian, J. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. Kimerling, and A. Agarwal, "ATAC: A 1000-Core Cache-Coherent Processor with On-Chip Optical Network," in *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*. ACM, 2010, pp. 477–488.
- [4] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. Beausoleil, and J. Ahn, "Corona: System implications of emerging nanophotonic technology," *ACM SIGARCH Computer Architecture News*, vol. 36, no. 3, pp. 153–164, 2008.
- [5] C. Nitta, M. Farrens, and V. Akella, "DCOFAn Arbitration Free Directly Connected Optical Fabric," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 2, no. 2, pp. 169–182, Jun. 2012.
- [6] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.
- [7] Y. Vlasov and S. McNab, "Losses in single-mode silicon-on-insulator strip waveguides and bends," *Optics express*, vol. 12, no. 8, pp. 1622–31, Apr. 2004.
- [8] J. Cardenas, C. Poitras, and J. Robinson, "Low loss etchless silicon photonic waveguides," *Optics Express*, vol. 17, no. 6, pp. 4752–7, 2009.
- [9] C. Manolatos, S. Johnson, S. Fan, P. Villeneuve, H. Haus, and J. Joannopoulos, "High-density integrated optics," *Journal of Lightwave Technology*, vol. 17, no. 9, pp. 1682–1692, 1999.
- [10] Q. Xu, S. Manipatruni, and B. Schmidt, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Optics Express*, vol. 15, no. 2, pp. 430–436, 2007.
- [11] S. Xiao, M. Khan, H. Shen, and M. Qi, "Multiple-channel silicon micro-resonator based filters for WDM applications," *Optics Express*, vol. 15, no. 12, pp. 7489–7498, 2007.
- [12] N. Kirman and J. F. Martínez, "A Power-efficient All-optical On-chip Interconnect Using Wavelength-based Oblivious Routing," *ACM Sigplan Notices*, vol. 45, no. 3, pp. 15–28, 2010.
- [13] S. Sahni, X. Luo, J. Liu, Y.-h. Xie, and E. Yablonovitch, "Junction field-effect-transistor-based germanium photodetector on silicon-on-insulator," *Optics letters*, vol. 33, no. 10, pp. 1138–40, May 2008.
- [14] B. Little and J. Foresi, "Ultra-compact Si-SiO₂ microring resonator optical channel dropping filters," *IEEE Photonics Technology Letters*, vol. 10, no. 4, pp. 549–551, 1998.
- [15] D. Ding and D. Pan, "OIL: a nano-photonics optical interconnect library for a new photonic networks-on-chip architecture," in *Proceedings of the 11th international workshop on System level interconnect prediction*, 2009, pp. 11–18.
- [16] K. Preston, N. Sherwood-droz, J. S. Levy, and M. Lipson, "Performance Guidelines for WDM Interconnects Based on Silicon Microring Resonators," in *Proceedings of Conference in Lasers and Electro-Optics (CLEO)*, 2011, pp. 4–5.
- [17] F. H. L. Koppens, D. E. Chang, and F. J. G. D. Abajo, "Graphene Plasmonics : A Platform for Strong Light-Matter Interactions," *Nano Letters*, pp. 3370–3377, 2011.
- [18] M. C. Lemme, F. H. L. Koppens, A. L. Falk, M. S. Rudner, H. Park, L. S. Levitov, and C. M. Marcus, "Gate-activated photoresponse in a graphene p-n junction," *Nano letters*, vol. 11, no. 10, pp. 4134–7, 2011.
- [19] Q. Bao and K. P. Loh, "Graphene photonics, plasmonics, and broadband optoelectronic devices," *ACS Nano*, vol. 6, no. 5, pp. 3677–94, 2012.
- [20] M. Liu, X. Yin, E. Ulin-Avila, B. Geng, T. Zentgraf, L. Ju, F. Wang, and X. Zhang, "A graphene-based broadband optical modulator," *Nature*, vol. 474, no. 7349, pp. 64–7, Jun. 2011.
- [21] F. Xia, T. Mueller, Y.-M. Lin, A. Valdes-Garcia, and P. Avouris, "Ultrafast graphene photodetector," *Nature nanotechnology*, vol. 4, no. 12, pp. 839–43, Dec. 2009.
- [22] G. Hendry, J. Chan, S. Kamil, L. Oliker, J. Shalf, L. P. Carloni, and K. Bergman, "Silicon nanophotonic network-on-chip using TDM arbitration," in *2010 18th IEEE Symposium on High Performance Interconnects*. IEEE, 2010, pp. 88–95.
- [23] N. Kirman, M. Kirman, and R. Dokania, "Leveraging optical technology in future bus-based chip multiprocessors," in *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, 2006, pp. 492–503.
- [24] A. Kahng, B. Li, L. Peh, and K. Samadi, "Orion 2.0: A fast and accurate noc power and area model for early-stage design space exploration," in *Proceedings of Design, Automation & Test in Europe*, 2009, pp. 423–8.
- [25] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating Future Network-on-Chip with Nanophotonics," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 429–440, 2009.
- [26] I. Llatser, C. Kremers, A. Cabellos-Aparicio, J. M. Jornet, E. Alarcón, and D. N. Chigrin, "Graphene-based nano-patch antenna for terahertz radiation," *Photonics and Nanostructures - Fundamentals and Applications*, vol. 10, no. 4, pp. 353–358, 2012.
- [27] I. F. Akyildiz and J. M. Jornet, "Electromagnetic Wireless Nanosensor Networks," *Nano Communication Networks (Elsevier) Journal*, vol. 1, no. 1, pp. 3–19, 2010.
- [28] S. Abadal, E. Alarcón, M. C. Lemme, M. Nemirovsky, and A. Cabellos-Aparicio, "Graphene-enabled Wireless Communication for Massive Multicore Architectures," *Accepted for Publication in IEEE Communications Magazine*.