Assessing the Energy Break-Even Point between an Optical NoC Architecture and an Aggressive Electronic Baseline

Luca Ramini
University of Ferrara, Italy
luca.ramini@unife.it

Paolo Grani
University of Siena, Italy
grani@dii.unisi.it

Hervé Tatenguem Fankem
University of Ferrara, Italy
herve.tatenguemfankem@unife.it

Alberto Ghiribaldi
University of Ferrara, Italy
alberto.ghiribaldi@unife.it

Sandro Bartolini
University of Siena, Italy
bartolini@dii.unisi.it

Davide Bertozzi
University of Ferrara, Italy
davide.bertozzi@unife.it

***Abstract*—Many crossbenchmarking results reported in the open literature raise optimistic expectations on the use of optical networks-on-chip (ONoCs) for high-performance and low-power on-chip communication. However, most of those previous works ultimately fail to make a compelling case for chip-level nanophotonic NoCs, especially for the lack of aggressive electronic baselines (ENoC), and the poor accuracy in physical- and architecture-layer analysis of the ONoC. This paper aims at providing the guidelines and minimum requirements so that nanophotonic emerging technology may become of practical relevance. The key differentiating factor of this work consists of contrasting ONoC solutions with an aggressive ENoC architecture with realistic complexity, performance, and power figures, synthesized on an industrial 40nm low-power technology. At the same time, key physical design issues and network interface architecture requirements for the ONoC under test are carefully assessed, thus paving the way for a well-grounded definition of the requirements for the emerging ONoC technology to achieve the energy break-even point with respect to pure electronic interconnect solutions in future multi- and many-core systems.**

I. Introduction

Optics could solve many physical problems of on-chip interconnect fabrics, including precise clock distribution, system synchronization, bandwidth and density of long interconnections, and reduction of power dissipation [1]. It may allow continued scaling of existing architectures and enable novel highly interconnected or high-bandwidth architectures.

However, despite the arguments in favor of optics for interconnects, and the promising monolithic integration routes with silicon, there is essentially no practical use today. The high cost targets for introducing this emerging technology, and the low-maturity of basic optical components, do not fully justify this scenario, since they only urge for more compelling cases where the benefits of chip-level nanophotonic interconnection networks can justify the cost barrier removal and a larger investment in technology development. It follows from this that such compelling cases cannot be directly derived from the benchmarking frameworks between optical NoCs (ONoCs) and their electronic counterparts (ENoCs) reported so far in the open literature.

While making an excellent point for the new interconnect technology, they lack of enough practical relevance to push it beyond the boundaries of an elegant research concept. In practice, they tend to deliver overly optimistic results for ONoCs for one or more of the listed reasons, which we recall from the conclusions in [2]. First, logical topologies are not well specified, hence preventing in-depth architecture review. Second, the baseline electronic NoCs exhibit naive or unoptimized architectures, hence overlooking that performance or power optimizations of ENoCs are many times far more practical than adopting an emerging technology. Third, specific instances of device parameters are currently meaningless for a fast-developing technology such as ONoCs. Fourth, the fixed power overhead is typically underestimated, directly or indirectly by assessing ONoC designs under high utilization regimes. Fifth, complex designs increase risk in terms of reliability, fabrication cost, and packaging issues. In addition to that, we would like to stress that place&route constraints are typically overlooked in ONoC topology design (hence underestimating layout-induced waveguide crossings and static power), and that electronic network interfaces for ONoC injection/ejection are often not considered in planning resource budgets.

This paper aims at a higher level of practical relevance in assessing the potentials of ONoCs for future multi- and many-core systems. This is fundamentally achieved in two ways. On one hand, *we make use of an aggressive electrical baseline*. We consider a realistic design point for the ENoC architecture in terms of complexity, area and power. Moreover, real synthesis runs of the target ENoC on a 40nm industrial low-power technology will provide the reference quality metrics the competing optical NoC solutions are contrasted with. On the other hand, *the ONoC is designed and accurately characterized based on both accurate physical-layer and architecture-layer analysis*. We select a wavelength-routed ring topology for the ONoC, whose simplicity can reduce the adoption risk of an emerging technology. At the physical layer, the increased accuracy in ONoC modeling is achieved by drawing the ring layout, especially its injection and ejection interfaces. At the architecture layer, the design of the network interface architectures needed to inject/eject electronic packets into/from the ONoC is made, thus capturing typically overlooked sources of performance and power overhead, such as flow control, clock resynchronization, or suitable FIFO sizing.

Another feature of this paper is that it *carefully considers fixed-power overheads*, which are a significant percentage of total ONoC power [3]. Static power is especially important in those application domains where the network does not undergo high utilization, but it has to serve sporadic traffic peaks. This is the case of shared memory multiprocessors with distributed last-level cache, implementing hardware support for cache coherence. The use of an ONoC makes sense in this domain only if it can significantly cut down on the total application execution time, thus burning less static power. This paper considers the case study of a directory-based implementation of the MOESI protocol, and derives the requirements for both ENoC and ONoC design. Therefore, the compared interconnect solutions are fine tuned for the kind of messages they are supposed to route.

The enhanced level of accuracy pursued by this paper in crossbenchmarking optical vs. electronic interconnect technology primarily aims at *deriving guidelines of practical relevance to materialize the nanophotonic concept into an affordable technological solution for the next generation multi- and many-core systems*. For this purpose, we extend the crossbenchmarking

978-3-9815370-2-4/DATE14/©2014 EDAA

effort to the system level by backannotating the relevant physical and architectural metrics/effects into an abstract system-level functional simulation framework, capable of projecting the ultimate impact that optical interconnect technology may have on realistic execution scenarios. For this reason, we leverage on realistic traffic patterns such as Parsec 2.1 for the estimation of performance and energy metrics for both Electronic and Optical NoCs.

II. Target System

Our experimental setting consists of a multi-core processor composed by 16 Tiles. Each of them operates as both initiator and target for communications over the system interconnect. Tiles are disposed over the die area in a common 4x4 2D-Mesh structure. Each Tile is associated with its respective Network Interface (a master or a slave one), performing protocol conversion and (de-)packetization. In our study we assume that these may be omitted during power analysis since they are equally required for both NoC implementations, and will not be cause of differentiation.Our analysis begins from the point where the two architectures start to diverge, including for instance the buffering architecture and frequency converters.

III. Baseline Electronic NoC

In this section we introduce the baseline Electronic Network on Chip (ENoC), which we implement as an aggressive low-power 2D mesh for our reference technology and chip multiprocessor architecture. We chose a 2D mesh topology because it is a regular structure that maps well to the regularity of chip multiprocessor architectures and is well suited for general-purpose multi-core systems.
The switch architecture is inspired by the ×pipesLite architecture [5], which represents an ultra-low complexity design point in the space of electronic NoCs. ×pipesLite architecture is an input-buffered switch, implementing logic-based distributed routing and wormhole switching. In this design, one clock cycle is taken to traverse the switch and one clock cycle to traverse the link connecting two switches. Buffer capacity has been set to two slots for input buffers and six slots for output buffers. The flit width is 32 bits, which represents a good trade-off between area occupancy and provided bandwidth. This paper considers a directory-based implementation of the MOESI protocol, which requires at least 3 virtual channels to avoid message-dependent deadlock [6]. We implement Virtual Channel flow control by replicating the basic switch architecture without VC capabilities three times, based on the approach presented in [7]. This technique results in higher maximum operating speed, better performance and smaller area. In order to preserve the generality of the design and support cores with different operating frequencies that access a fixed-frequency NoC, dual-clock FIFOs have been included at the network interfaces.
Post-layout analysis has been performed to assess performance and power metrics of the Electronic NoC. For this, the switch architecture has been synthesized, placed and routed using *Synopsys Design Compiler* and *IC Compiler* tools. For the physical implementation we leveraged on an Ultra-Low-Power Standard VTh 40nm industrial technology library. After layout generation, the maximum operative frequency of our design is 1.2 GHz.

IV. Wavelength-Routed Optical Ring Design

In this section we describe our Wavelength-Routed Optical Ring topology, relying on a rigorous design methodology, addressing waveguide crossing concerns, assessing laser power, and considering the optical network interface architecture.

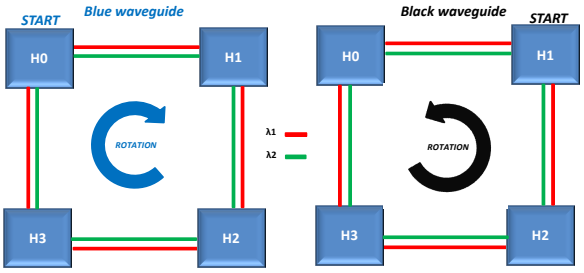

TRUTH TABLE

TX/RX	H0	H1	H2	H3
H0	X	λ1	λ2	λ1
H1	λ1	X	λ1	λ2
H2	λ2	λ1	X	λ1
H3	λ1	λ2	λ1	X

Fig. 1. Principle of the designed Optical Ring Architecture

A. Design Methodology

Simplicity and low implementation cost make the optical ring topology one of the most appealing interconnection networks proposed in the open literature to interconnect initiators and targets of a given multi-processor system-on-chip (MP-SoC). In addition, and especially in a 3D-stacked scenario, the ring topology efficiently meets the place & route constraints unlike other solutions such as multi-stage networks [4], and filter-based ones [8].
For these reasons, we engineered our Wavelength-Routed Optical Ring Architecture inspired by Le Beux's algorithm [9]. The key property is that the same wavelengths can be reused on a single waveguide to establish multiple communications, as depicted in Figure 1.
Here, we have an optical ring network structured into 4 **Hubs** *H0, H1, H2, H3)* which are both initiators and targets. In the proposed example, two different wavelengths ($\lambda 1$, $\lambda 2$) and two distinct waveguides (the blue and black one) are sufficient to realize 12 contention-free optical paths. This is demonstrated through the wavelengths assignment reported in the truth-table of Figure 1. We denote that, no wavelength is listed along its diagonal as self-communications are not allowed. By pursuing the presented principle we scaled up our ring architecture up to **16 Hubs**. For the design of such a network 13 wavelengths were reused and multiplexed on 16 different waveguides to enable 240 contention free optical paths.
Ultimately, as it was described here, by leveraging on both SDM (Spatial-Division-Multiplexing) technique, based on replicating physical waveguides, and the WDM (Wavelength-Division-Multiplexing) one which enables several optical communications on a single waveguide at the same time, and alternating also different rotations (clockwise and counterclockwise), it was possible to deliver multiple contention-free optical paths, while minimizing the amount of physical resources and of optical power losses.

B. Waveguide Crossings Concern

Le Beux et al. developed an Optical Ring architecture in [9], called ORNoC. The key property of a such a structure is that in principle it has neither waveguide crossings nor photonic switching elements, thus leading to an appealing interconnection network. However, there are physical effects that come into play when its actual implementation is pursued. **First**, reachability of all SDM waveguides from injection and ejection interfaces of optical packets has to pay the price of undesired crossings. **Second**, MRRs (Micro-Ring-Resonators) are needed not only at the destination stage (Ring Filters) to

PHOTONIC COMPONENTS AND DEVICE PARAMETERS	VALUE
COUPLER LOSS	0.46dB
MODULATOR INSERTION LOSS	4.0dB
PHOTODETECTOR LOSS	1.0dB
FILTER DROP LOSS	1.0dB
THROUGH RING LOSS	10^{-3}dB/ring
PROPAGATION LOSS	1.5dB/cm
BENDING LOSS	0.005dB
WAVEGUIDE CROSSING LOSS (@REALISTIC)	0.52dB
WAVEGUIDE CROSSING LOSS (@AGGRESSIVE)	0.18dB
RECEIVER SENSITIVITY	-20dBm
LASER EFFICIENCY (@REALISTIC)	8%
LASER EFFICIENCY (@AGGRESSIVE)	20%

TABLE I. PHOTONIC COMPONENTS AND DEVICE PARAMETERS.

selectively eject the optical signal, but also to couple it into the ring after modulation (Ring Couplers). In an injection interface, the optical power is at first modulated and then coupled into one or more physical waveguides the ring is broken into. This gives rise to additional waveguide crossings to reach the ring waveguides that are further away from the injection interface. Similar considerations hold for ejection interfaces. Overall, this effect becomes so relevant that when the ring is scaled up to 16 Hubs, the worst case number of crossings in an optical network interface is 15. Another important issue is given by the total number of MRRs, penalized by its ring modulators, ring filters and ring couplers. The 16x16 Ring is not able to stay below 2.16K MRRs, hence resulting in a large thermal tuning overhead.

C. Laser Power Assessment

The preliminary step to evaluate the efficiency of an optical network relies on the estimation of the insertion loss across all wavelengths involved in the given design. Such a metric is extremely important to quantify the total amount of laser power needed to reliably detect the optical message at the destination node. For this reason, we calculated the insertion loss (IL) as the sum of physical components which affect the optical signal along the path, starting from Modulator, coupling filter, propagation distance, bending waveguide, ring filter, photodetector without overlooking crossing waveguides. In the end, we quantified the worst case ILmax on each wavelength and laser power was derived accordingly. In our study, die size is assumed to be *8mmx8mm* while loss parameters are listed in table I.

For the sake of a more comprehensive analysis of nanophotonic devices, and of their evolution over time, we distinguished two relevant cases: realistic and aggressive ones. For the former one we considered a wall-plug laser efficiency of 8% while crossing waveguides were supposed to be optimized with the standard elliptical taper [10]. In the aggressive case, we instead projected metrics assuming a laser efficiency of 20% whereas crossings were improved through MMI (Multi-Mode Interference) tapers [11]. In both cases, the detector sensitivity was considered to be the same at -20 dBm. Figure 2 illustrates the laser power trend across wavelengths. This figure highlights that laser sources (assumed to be Continuous Wave) must be treated in a different way. In other words, as the insertion loss is not the same for each path, CW lasers can be sized accordingly. Therefore, there will be some laser sources that turn out to be more power hungry than others.

When the aggressive case comes into play, even if the absolute power gap is strongly reduced (4x), there is still a relative gap across wavelengths, thus confirming that laser sources should be treated separately.

Finally, had we ideally designed our ring topology without accounting for crossings at each optical interface, the total laser power would have been clearly lower than our accurate approach. More in details, there would have been a reduction

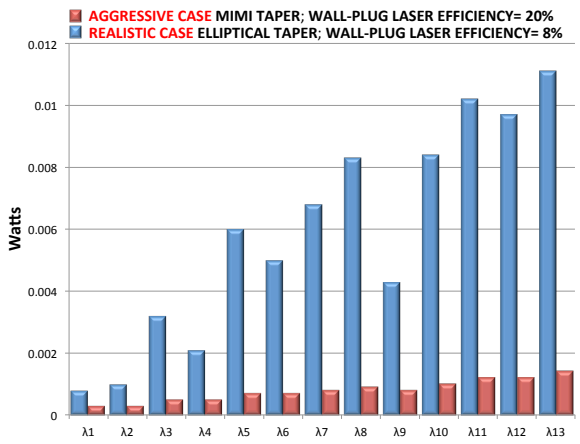

Fig. 2. Laser power results across wavelengths: aggressive vs. realistic

of 80%, and 41% in the realistic and aggressive scenarios respectively.

D. Optical Network Interface Architecture

This section describes, to the best of our knowledge, the first complete network interface architecture for optical networks as depicted in Figure 3. As a consequence, the objective is not to present the best possible design point, but rather to start considering the basic components, and indicating which one deserves the most intensive optimization effort for prime time of optical interconnect technology in industry.

To avoid message-dependent deadlock, every network interface needs separate buffering resources for each one of the three message classes of the MOESI protocol. This should be combined with the requirements of wavelength routing: each initiator needs an output for each possible target, and each target needs an input for each possible source. As a result, in the baseline version of the NI, each initiator comes with 3 FIFOs for each potential target, and each target with 3 FIFOs for each potential initiator. In a more optimized version of the NI (the one in Figure 3), all destinations share the same set of 3 FIFOs and the flits are sent to different paths afterwards (all logic components after 1x15 demuxes are replicated for each destination). All the FIFOs at both the transmission and the reception side must be dual-clock FIFOs to move data between the processor frequency domain and the one used inside the NI. The serializers are responsible for translating the flit into a 10 GHz bit stream. The reception side is specular: flits must follow the deserialization process and another set of dual-clock FIFOs. Clearly, wavelength-routed ONoCs move most of their complexity to the NIs, which should therefore not be overlooked by means of overly abstract models.

Another key issue to be considered in NIs concerns the resynchronization of received optical pulses with the clock signal of the electronic receiver. In this paper we assume source-synchronous communication, which implies that each point-to-point communication requires a strobe signal to be transmitted along with the data on a separate wavelength, and used to correctly sample received data. Optical transmission of clock signals is an active research field: see for instance [12]. This strobe signal is generated starting from the electrical clock of the transmitter, and removes the need for phase-locked loops (PLLs) or delay-locked loops (DLLs). In this work, we assume that a form of clock gating is implemented, therefore when no data is transmitted, the optical clock signal is gated.

Another typically overlooked issue consists of the backpressure mechanism. We opt for credit-based flow control because it

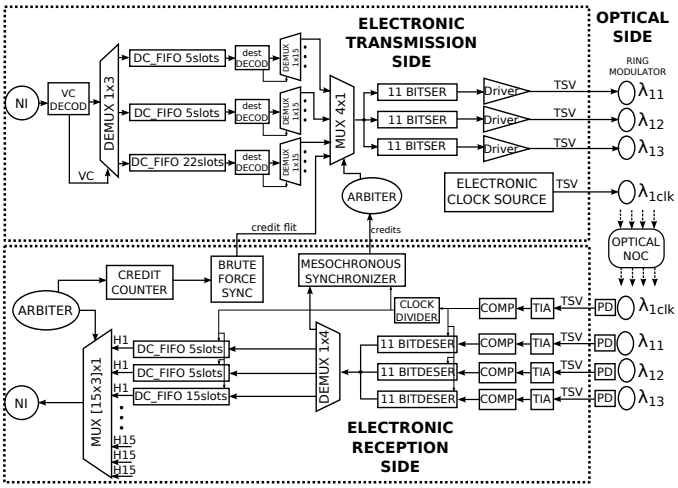

Fig. 3. Optical Network Interfaces Architecture

does not rely on timing assumptions, and credit tokens can reuse the existing communication paths, thus avoiding any additional waveguide, and resulting in a milder impact over static power.

ELECTRONIC DEVICES (// bit parallelism)	STATIC POWER (mWatts)	DYNAMIC ENERGY (fJ/bit)
DC_FIFO_TX_5 //3	0.12	10.65
DC_FIFO_RX_5 //3	0.12	8.54
DC_FIFO_TX_22 //3	0.12	39.00
DC_FIFO_RX_15 //3	0.12	26.50
MUX4x1_ARB //3	0.08	0.36
MUX45x1_ARB //3	0.9	5.09
SERIALIZER //3	0.0475	9.41
DESERIALIZER //3	0.0289	7.74
MESO_SYNCH //3	0.082	8.00
BRUTE_FORCE //3	0.004234	1.4
DC_FIFO_TX_5 //4	0.12	12.72
DC_FIFO_RX_5 //4	0.12	10.2
DC_FIFO_TX_22 //4	0.12	46.41
DC_FIFO_RX_15 //4	0.12	31.65
MUX4x1_ARB //4	0.11	0.49
MUX45x1_ARB //4	0.9	5.09
SERIALIZER //4	0.0417	2.63
DESERIALIZER //4	0.0281	6.12
MESO_SYNCH //4	0.113	11.1
BRUTE_FORCE //4	0.00503	1.66
DEMUX1x3	0.000725	0.92
DEMUX1x15	0.0021	25.21
DEMUX1x4	0.00056	6.72
COUNTER@4bits	0.02964	1.014
TSV	/	2.5
TRANSMITTER (aggressive)	0.025	20
TRANSMITTER (realistic)	0.100	50
RECEIVER (aggressive)	0.050	10
RECEIVER (realistic)	0.150	25
THERMAL TUNING /MRR @20K	0.020	/
E-SWITCH (3VCs)	17.9	193

TABLE II. STATIC AND DYNAMIC POWER OF ELECTRONIC DEVICES.

V. POWER MODELING

This section describes the power modeling assumptions on which our crossbenchmarking framework is based. Every electronic component has been synthesized, placed and routed using a Low-Power 40nm industrial technology library, in order to provide realistic power measurements (not derived from optimistic or ideal estimations). Power metrics have been calculated by backannotating the switching activity of block internal nets, and then importing waveforms in the *PrimeTime* Tool. It is worth observing that we have applied clock gating for the sake of realistic measurement of static power.
Energy-per-Bit has been computed by removing the Static Power by the Total power on a component-basis, under 50% switching activity assumption.
Overall, the power consumption of the Electronic NoC is built upon replicating the power contribution of its basic switch components. As mentioned in section III, we also considered the power consumption of both electronic NI buffering and frequency converters (dual-clock FIFOs) which contribute around 11.5 mW. The static power dissipated (Idle power) by the entire network (16 switches), is around 286 mW (only the top-level clock tree is omitted). In addition, the energy required for transmitting data over each hop of the ENoC is 193 fJ/bit.
Similarly, the power dissipation of Optical Network Interfaces is computed by composing the power consumption of each of its sub-blocks (DC_FIFOs at the transmission sides, Demultiplexers, SERs, Synchronizers, DESERs, DC_FIFOs at the reception sides, Multiplexers, and Credit counters).
The static power contribution of all optical components is given by: Laser sources, Thermal tuning, Transmitter (i.e., the driver-ring modulator couple), Receiver (i.e., Photodetector, Trans-Impedance Amplifier, and Comparator) and the source-synchronous clock. The latter addendum is internally composed by further laser sources, Transmitters, Receivers, and MRRs as well. For static and dynamic power parameters, as well as for their relative ratio, we consistently assume values from the same literature source [16], [2]. In order to transmit each bit there is the need for: 13 CW laser sources, 240 TXs and RXs and 720 MRRs. These resources must be replicated as many times as the target bit parallelism, and also for the optical clock support. Power metrics of all basic blocks of our architectures are summarized in Table II. The derived static and dynamic power values for electronic and optical components are combined with system-level simulation results (see section VI) to obtain comprehensive metrics under the effect of functional traffic.

VI. EXPERIMENTAL RESULTS

This section proposes our results in terms of **performance** and **energy consumption**. We followed two fundamental analysis strategies, hereafter referred as: CMF (Common Modeling Framework) vs. AMF (Accurate Modeling Framework). Their comparison is very instructive.
The first one reflects common modeling assumptions in the open literature, which lead to an overly optimistic assessment of optical interconnect technology. In particular, network interfaces are typically oversimplified, and end up being abstracted by simple input/output FIFOs of infinite length. Similarly, the blocking effect of the backpressure mechanism is overlooked. As a consequence, the ONoC easily proves much more performance-efficient than the electronic counterpart. Moreover, the lack of a layout analysis in addition to a physical-layer analysis in ONoC design is another important source of optimism in previous evaluations.
In contrast, the key strength of this work (AMF methodology) consists of a careful exploration of E/O and O/E interfaces, accounting for the contributions and effects of every building block: routing, buffering, serialization and deserialization processes, as well as optical transmitters and receivers, clock domain synchronizer, backpressure cost. Last but not least: the propagation of source-synchronous clocks in optical networks and their impact on the overall system power consumption. It will be very interesting to compare the design trade-offs with the two experimental methodologies. They will be hereafter presented with both the conservative and the more aggressive parameters of the optical technology.

A. Methodology

Experimental results were obtained through GEM5 full-system simulator [13] where we modeled in details both the

electronic baseline and the optical architecture described in the previous sections. Modeling included functional behavior, timing accuracy, and energy consumption. Performance and energy were evaluated for the PARSEC 2.1 benchmark suite, a collection of heterogeneous multithreaded applications spanning different emerging application domains [14]. These benchmarks are representative not only of chip multiprocessor workloads, but also of the current and near future usage of high-end embedded devices such as smartphones and tablets [15]. We adopted the medium input-set to have a significant workload size and maintain a reasonable simulation time (within a few days per benchmark). All benchmarks have been run on a Linux 2.6 operating system, which was booted on the simulated architecture, and were instantiated with a degree of parallelism 16. Benchmarks were modified to enforce core affinity as to avoid non-determinism due to operating system scheduling.
As for performance metrics, we considered execution time of the entire parallel region of each benchmark as representative of the end-user perceived performance, and the overall energy consumption of the considered networks. Finally, the shared cache resources are coordinated by a state-of-the-art MOESI coherence protocol.

B. Result discussion

As shown in Figure 4, the optical solution, for both 3 and 4 bit parallelism, is able to deliver performance speedups over the electronic baseline. It achieves up to 23% improvement in case of the larger parallelism (4-bit), with peaks of more than 30% for *canneal* and *swaptions* applications. This speedup is indirectly useful to reduce the overall static energy consumption of the optical network. The 3-bit parallelism scores slightly worse obtaining a 18% performance improvement. We did not consider ONoCs with less than 3-bit parallelism, since the bandwidth of optical paths ends up being less than that of electronic counterparts. More than 4 bits were equally not considered since the corresponding static power of the ONoC became unacceptable. From the energy consumption point of view, Figure 5a and Figure 5b show the achieved results for the common optimization analysis. In this setup the aggressive and the realistic case show a very different behavior. In the former case, the ONoC saves energy with respect to the electronic baseline by almost 70% on average for the 3-bit case and 60% for the 4-bit one, by exploiting its reduced static consumption. In the latter case, ONoC also obtains a good energy improvement (avg. of 28% for 3-bit setup and 13% for the 4-bit one) over the electronic baseline demonstrating that, even with the realistic setup, it is still able to get some benefits compared with the ultra-low-power electronic baseline. Although the CMF raises expectations that are not justified in practice, it is already able to point out a realistic effect: in the presence of a coherence traffic, the great dynamic energy savings of the ONoC do not count a lot in the final energy balance, since static power is the dominating factor, and it is mainly associated with the amount of instantiated resources, as well as with technology maturity. Figure 6a and Figure 6b show the achieved results for the accurate model. In this case the aggressive setup causes the ONoC to track closer (avg. overhead of 11.6%) the break-even with the very low-power electronic baseline, by exploiting the benefits derived by the execution time speedup of the optical network. The results achieved in the realistic case are instead far from those of the electronic baseline, and this is due to the more relaxed technology used in this setup, and to the more accurate modeling of the optical structures needed. The energy overhead worsens by more than 2x with respect to the electronic energy figures for the 3-bit configuration. With

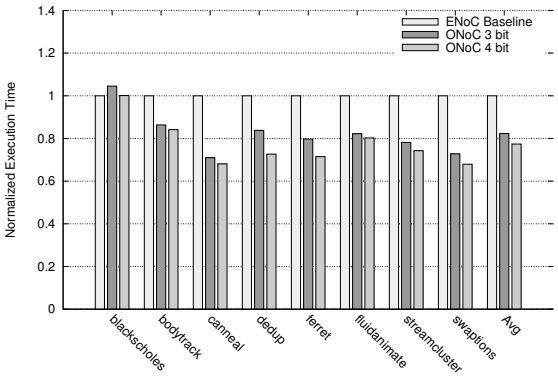

Fig. 4. Performance comparison of the ONoC with the electronic baseline

respect to the AMF, the CMF points out a clear underestimation of static power, which further emphasizes the dynamic power benefits, and especially that the actual complexity and overhead of network interfaces are typically overlooked. Indeed, these results question the common conclusion that ONoC prime time will depend only on the progress of technology maturity. In contrast, such prime time will only come through an in-depth optimization of network interfaces, where the real complexity is hidden under subtle issues such as buffer numbers and sizing, protocol-dependent deadlock, synchronization and flow control. The obtained performance benefits are not able to reduce the static consumption to get results comparable with the very low-power electronic baseline with the nowadays available technology. The good performance results could be however a good starting point in case of future enhancement, giving a hint to better explore and refine the technology applied to these new optical solutions, trying to obtain the maximum advantages from them. Especially, it will be mandatory to explore static power gating solutions, similarly to what happened in the past with electronic circuits.

VII. Conclusion

This paper aims at a high level of practical relevance in the crossbenchmarking of an optical NoC vs. its electronic counterpart. The key novelty consists of the use of an electronic baseline aggressively optimized for low-power. As a result, the optimistic assessments of many previous works are put in discussion. Nonetheless, the experimental results do not paint a dismal picture on optical interconnect technology. In fact, it is proven to achieve relevant performance speedups even with bursty communication workloads (as opposed to the high utilization rates typically assumed), which are common in shared memory multiprocessors. With conservative projections for optical component parameters, the major role played by static power is apparent. This calls for new power gating techniques. With more aggressive projections, the network interface turns out to be the clear bottleneck to achieve the break-even point with low-power ENoCs, hence it should be thoroughly analyzed for optimization. In future work, we will investigate more communication-dominated scenarios, in an attempt to capitalize on the far lower dynamic power consumption of ONoCs.

Acknowledgments

This work was supported by the PHOTONICA project (RBFR08LE6V) under the FIRB 2008 program, funded by the italian government. All authors would like also to thank Marta Ortín Obón (from the University of Zaragoza) for her valuable feedbacks.

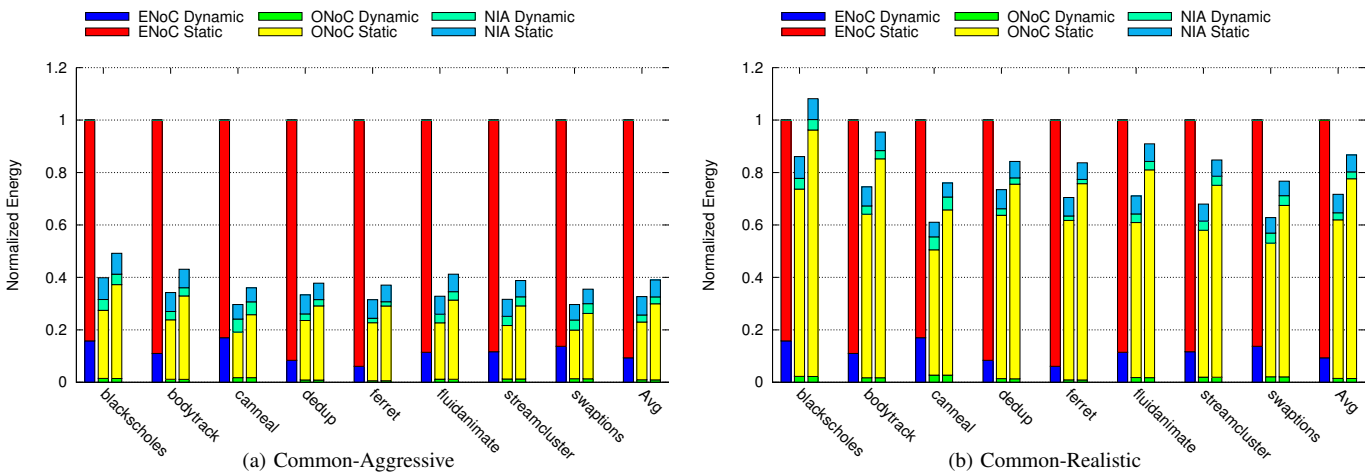

(a) Common-Aggressive (b) Common-Realistic

Fig. 5. Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wrt. ENoC baseline both for the common-aggressive (a) and the realistic (b) case

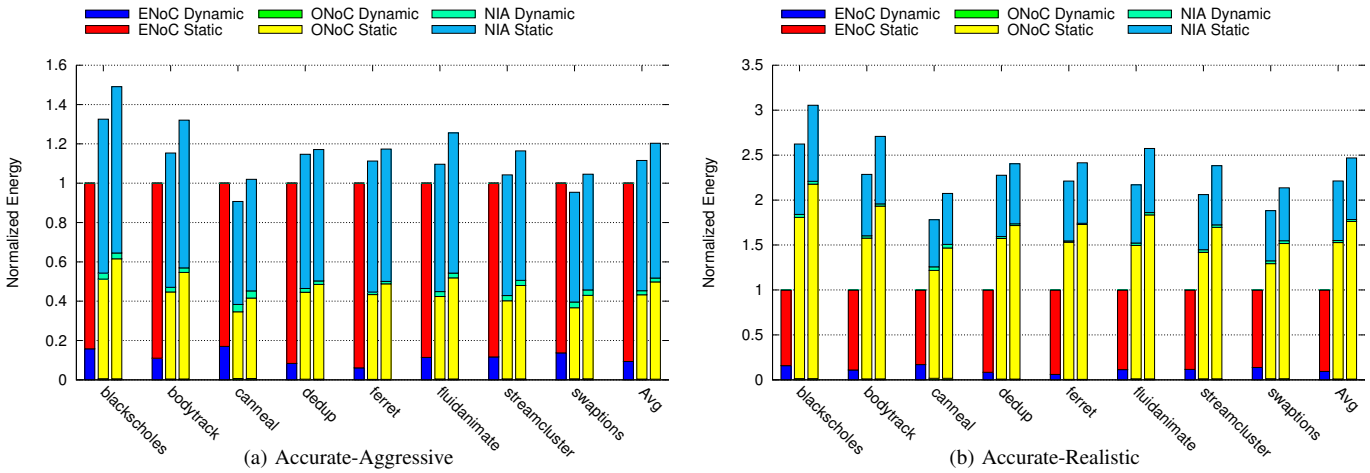

(a) Accurate-Aggressive (b) Accurate-Realistic

Fig. 6. Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wrt. ENoC baseline both for the accurate-aggressive (a) and the realistic (b) case

REFERENCES

[1] D. Miller. *Rationale and challenges for optical interconnects to electronic chips.* Proceedings of the IEEE. vol. 88, no. 6, pp. 728-749, 2000.

[2] C. Batten, A. Joshi, V. Stojanovic, K. Asanovic. *Designing chiplevel nanophotonic interconnection networks.* Emerging and Selected Topics in Circuits and Systems, IEEE Journal. vol. 2, no. 2, pp. 137-153, 2012.

[3] A. Udipi, N. Muralimanohar, R. Balasubramonian, A. Davis, N. Jouppi. *Combining memory and a controller with photonics through 3d-stacking to enable scalable and energy efficient systems.* ISCA 11.

[4] I. O' Connor, M. Brire, E. Drouard, A. Kazmierczak, F. Tissafi-Driss , D. Navarro, F. Mieyeville, J. Dambre, D. Stroobandt, J.-M. Fedeli, Z. Lisik, F. Gaffiot. *Towards reconfigurable optical networks on chip.* ReCoSoC 2005. pp. 121-128.

[5] S. Stergiou, F. Angiolini, S. Carta, L. Raffo, D. Bertozzi, G. De Micheli. *timespipes lite: a synthesis oriented design library for networks on chips.* Design, Automation and Test in Europe, 2005. pp. 1188-1193 Vol. 2.

[6] A. Hansson et al. *Avoiding messagedependent deadlock in network-based systems on chip.* VLSI Design, vol. 2007. 2007

[7] F. Gilabert, M.E. Gomez, S. Medardoni, D. Bertozzi. *Improved utilization of noc channel bandwidth by switch replication for cost-effective multi-processor systems-on-chip.* (NOCS), 2010. ACM/IEEE International Symposium on, 2010, pp. 165-172.

[8] X. Tan, M. Yang, L. Zhang, Y. Jiang, J. Yang. *On a scalable, non-blocking optical router for photonic networks-on-chip designs.* (SOPO), 2011. pp. 1-4.

[9] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, P. Paulin. *Optical ring network-on-chip (ornoc): Architecture and design methodology.* (DATE), 2011, 2011, pp. 1-6.

[10] N. Sherwood-Droz, H. Wang, L. Chen, B.G. Lee, A. Biberman, K. Bergman, M. Lipson. *Optical 4x4 hitless slicon router for optical networks-on-chip (NoC).* Opt. Express, vol. 16, no. 20, pp. 15 915-15 922, Sep 2008.

[11] G.R. Hadley. *Effective index model for vertical-cavity surface-emitting lasers.* Opt. Lett., vol. 20, no. 13, pp. 1483-1485, Jul 1995.

[12] J. Leu, V. Stojanovic. *Injection-locked clock receiver for monolithic optical link in 45nm SOI.* (A-SSCC), 2011 IEEE Asian, 2011, pp. 149-152.

[13] N. Binkert, R.G. Dreslinski, L.R. Hsu, K.T. Lim, A.G. Saidi, S.K. Reinhardt *The m5 simulator: Modeling networked systems.* Micro, IEEE, vol. 26, no. 4, pp. 52-60, 2006.

[14] C. Bienia, S. Kumar, J.P. Singh, K. Li. *The parsec benchmark suite: characterization and architectural implications.* PACT 08.

[15] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, D. Estrin. *Diversity in smartphone usage.* Proceedings of the 8th international conference on Mobile systems, applications, and services.

[16] S. Beamer, C. Sun, Y. Kwon, A. Joshi, C. Batten, V. Stojanovic, K. Asanovic. *Re-architecting DRAM memory systems with monolithically integrated silicon photonics.* ISCA '10.