



HAL
open science

Energy and Performance Trade-off in Nanophotonic Interconnects using Coding Techniques

Cedric Killian, Daniel Chillet, Sébastien Le Beux, Olivier Sentieys, Van-Dung Pham, Ian O'Connor

► **To cite this version:**

Cedric Killian, Daniel Chillet, Sébastien Le Beux, Olivier Sentieys, Van-Dung Pham, et al.. Energy and Performance Trade-off in Nanophotonic Interconnects using Coding Techniques. DAC 2017 - IEEE/ACM Design Automation Conference DAC, Jun 2017, Austin, United States. pp.6. hal-01495468

HAL Id: hal-01495468

<https://inria.hal.science/hal-01495468>

Submitted on 29 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Energy and Performance Trade-off in Nanophotonic Interconnects using Coding Techniques

Abstract—Nanophotonic is an emerging technology considered as one of the key solutions for future generation on-chip interconnects. Indeed, this technology provides high bandwidth for data transfers and can be a very interesting alternative to bypass the bottleneck induced by classical NoC. However, their implementation in fully integrated 3D circuits remains uncertain due to the high power consumption of on-chip lasers. However, if a specific bit error rate is targeted, digital processing can be added in the electrical domain to reduce the laser power and keep the same communication reliability. This paper addresses this problem and proposes to transmit encoded data on the optical interconnect, which allows for a reduction of the laser power consumption, thus increasing nanophotonics interconnects energy efficiency. The results presented in this paper show that using simple Hamming coder and decoder permits to reduce the laser power by nearly 50% without loss in communication data rate and with a negligible hardware overhead.

I. INTRODUCTION

Technology scaling down to the ultra-deep submicron domain enables the integration of hundreds of cores in a single chip. To enhance connecting many cores, chip-level communication needs disruptive interconnects to bring down the power budget. On-chip nanophotonics interconnects are an emerging technology considered as one of the key solutions for future generations of many cores [1], providing high bandwidth and low latency. Interconnect architectures, e.g. Multiple Writer Single Reader (MWSR) [2], are composed of laser sources, Micro-ring Resonators (MRs), waveguides and photodetectors. Active optical devices are power hungry compared to the traditional electrical functions. Considering this power consumption gap, new approaches to share the power consumption between the electrical and optical domains are needed in order to reduce the global power consumption.

This paper addresses this challenge and proposes to add data protection in the electrical domain, which allows reducing the power consumption of optical transmission without any reduction in the communication quality. For this purpose, Error Correction Code (ECC) are used on the transmitter side and faulty bits detected by the receiver are corrected. Since additional faulty bits are acceptable using ECC, the optical power emitted by the laser source can be significantly reduced. The laser being a major contribution of the power consumption, the global power of the optical interconnect can therefore be significantly reduced under specific scenarios without compromising energy per bit figures. The paper is organized as follows. Section II gives the start of the art. III details the approach for with a simple example and IV introduce the generic architecture and the models we use.

Power and performance trade-off are explored in Section V and Section VI concludes the paper.

II. RELATED WORK

ECC is a commonly used technique in wired and wireless communications [3], [4]. The association of power scaling of an emitter and the use of ECC allows to obtain an efficient energy per bit transmission given a targeted BER [5], [6]. Indeed, ECCs allow to correct errors in data communications, hence it is possible to maintain a given BER while decreasing the Signal-to-Noise Ratio (SNR). Moreover, these approaches are scalable to the Network-on-Chip (NoC) communication medium. In [6] the authors proposed to use a closed-loop transmission to dynamically tune the transceiver power and the ECC correctness capacity. This allows to guarantee a target BER while tracking the communication channel variability.

Regarding Optical NoC (ONoC), the laser output power is reduced under BER constraints, which allows for the reduction of the laser power [7]. This approach has been extended in [8] by considering the laser thermal sensitivity and the crosstalk to evaluate the BER. However, the authors did not investigate the use of ECC for data transmission. In this paper, for the first time, we jointly use ECC and laser output power tuning technique in order to further improve the ONoC energy efficiency. The proposed approach is complementary to existing techniques allowing to dynamically turn off lasers when no activity is detected [9].

III. ECC FOR POWER/PERFORMANCE TRADE-OFF

Figure 1 illustrates our approach in the context of a simple optical link with one laser source, one writer (source core) and one reader (destination core). We first detail the basic communication schemes without ECC. Then, we revisit the transmission schemes when ECC is used.

A. Communication Without ECC

To communicate, the source core i) sets the corresponding MR to the modulation state and ii) sends the data (d_0 , d_1 and d_2 in the figure). Data are serialized and the corresponding stream of information bits leads to an On-Off-Keying (OOK) modulation of the optical signal at the wavelength λ_x , which is emitted by the laser. The modulation is realized by electro-optic effect on the MR. Forward biased is applied to perform voltage tuning, which leads to a blue shift of the MR resonant wavelength. In our model, data '1' and '0' lead to OFF state and ON state, respectively:

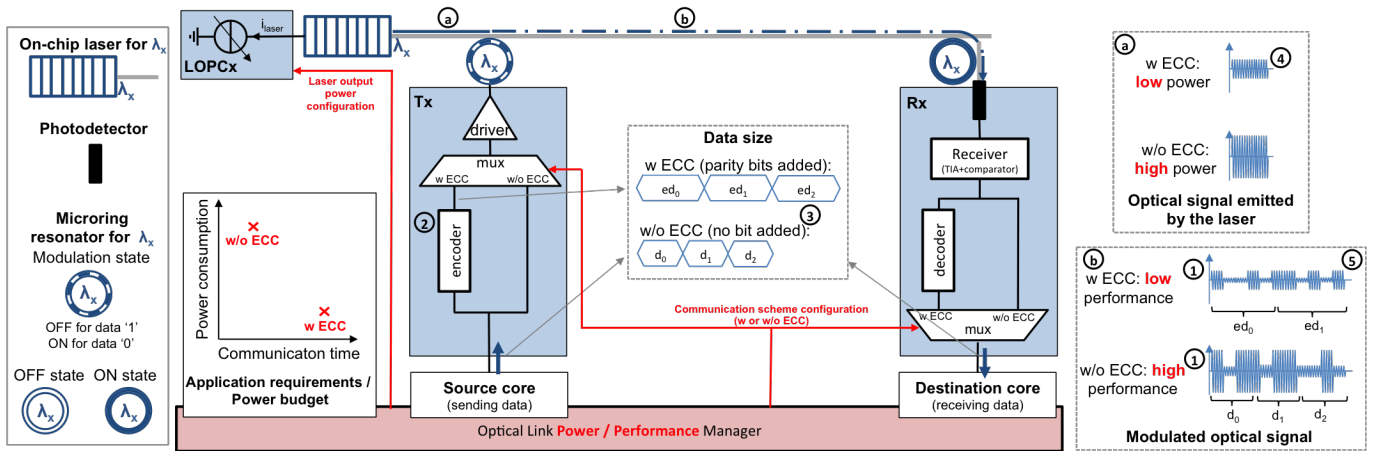


Fig. 1. General principle of the use of ECC in an ONoC

- OFF state: the MR resonant wavelength is detuned from the optical signal and the optical continues propagating with low losses.
- ON state: the MR resonant wavelength and the optical signal are aligned. Most of the optical signal power is absorbed by the MR, which leads to a strong attenuation on the output.

The difference between OFF and ON attenuation defines the extinction ratio. As an example, a 9.2dB extinction ratio at 9Gb/s with 4.32mW and for a 2Vpp CMOS driver has been reported in [10]. In Figure 1, this is illustrated by the gap between the two optical signal power levels (mark ①). The modulated signal then propagates along the waveguide until the reader, where it is dropped to the photodetector (the optical signal is dropped by MR, by putting it in ON state). The photocurrent is then converted back to a digital data transmitted to the destination core. The transmission can be characterized as follows:

- Power: during the transmission, the modulated optical signals experience losses (not shown in the figure but discussed later on). The laser thus must emit a high enough optical power in order to reach the targeted BER (e.g. 10^{-9}), knowing the losses, the MR extinction ratio, and the photodetector dark current. The laser efficiency being relatively low (e.g. around 5%), its contribution in the optical link power budget is significant.
- Performance: thanks to the direct modulation of the data (i.e. without ECC), the transmission time is reduced to the minimum if the target BER is reached.

The following section discusses how the use of ECC impacts the optical communications.

B. Communication With ECC

When an ECC is used, the data to be emitted are first encoded (ed_0 , ed_1 and ed_2 in Figure 1, mark ②) before serialization and modulation. Due to the added bits to be transmitted, the communication time is longer (mark ③). In the receiver, the optical signals are converted back into electrical signals and the data are decoded. This allows detecting and

correcting errors, which in turns improves the BER. Hence, assuming a same target BER as for the communication without ECC, the constraints on the SNR are leveraged when ECC is used. This leads to reducing the power of the optical signals, as suggested in the figure with the decreased amplitudes of the emitted signal (mark ④) and the modulated signal (mark ⑤); it is worth mentioning that the extinction ratio is the same). Hence, a global reduction of the optical link power is achievable if the energy saved on the laser side is larger than the energy spent in the data coding/decoding and in the transmission of the added bits.

C. Performance/Power Trade-offs

In short, not using ECC leads to fast but power hungry communications, while ECC increases the communication time but is potentially more power efficient. This complementary trend well suits execution constraints of nowadays applications running on complex 3D-Integrated Circuits. While execution deadlines have to be met for real-time applications, energy saving strategies can be applied for power hungry multimedia-like applications (e.g. by degrading the BER).

This paper proposes to configure nanophotonics interconnects by jointly i) selecting the ECC used for data transmission and ii) tuning the laser output power. Hence, by considering i) application requirements and power budget and ii) ECC performance and power figures, a manager configures at runtime the optical link. The configuration is two-fold. First, the most appropriate communication scheme with or without ECC (w ECC or w/o ECC in the figure) is selected and both writer and receiver are configured accordingly. Second, the optical power of the signal emitted by the laser is set in order to respect the BER requirements. This is achieved by tuning the laser current in the laser output power controller. The use of a shared, or centralized, manager to configure one emitter and one receiver ONIs is a common solution in ONoCs [11]. In these architectures, a source sends a request to the manager by specifying the destination and the communication requirements while the manager responds with the suitable configuration to apply on both source and destination sides.

Hence, adding the communication scheme (w or w/o ECC) within the requests has a limited impact on the manager cost (not detailed in this paper). The choice of the communication scheme is handled by the Operating System.

IV. ARCHITECTURE AND MODELS

This section presents the architectures we consider and the corresponding models.

A. 3D Integrated Circuits with Nanophotonic Interconnects

We assume a 3D-IC with an electrical layer and an optical layer implementing computing cores and the nanophotonic interconnect, respectively. Figure 2-a illustrates an example with four optical network interfaces (ONI) communicating using four MWSR channels (in this illustrative example, we assume a single waveguide per channel). There are thus three writers and one reader per channel: each writer is composed of N_W modulators and the reader is composed of N_W photodetector/passive MR couples.

B. MWSR Channel

Figure 2-b illustrates an MWSR channel in the nanophotonics interconnect, which is obtained by extending the point-to-point link from Figure 1 with wavelengths and writers. For this purpose, we assume N_W signal wavelengths, from λ_0 to λ_{N_W-1} , emitted by lasers and combined using a multiplexer (MUX), which can be implemented using multimode interference (MMI) coupler [12]. The un-modulated optical signals propagate on the waveguide until reaching the writer for which the channel access has been granted. The source core transmits the encoded data to the destination core using all the channel wavelengths. In our implementation, we assume an encoder/decoder pair for each wavelength, meaning that the (serial) stream of data sent by the source core is transmitted in parallel on the optical channel. Other implementations exist but are out of the scope of this paper.

As illustrated in Figure 2-c, each encoder allows transmitting data without ECC or with an ECC such as Hamming(7,4) (noted as H(7,4) in the rest of the paper). In this work, we chose Hamming for its simplicity, but other coding techniques can be used. The design of the interface is a challenging task due to the high operating frequency implied by the modulation speed. A complete description of such interface is given in Section IV-C.

The modulated signals propagate along the waveguide, cross intermediate writers and eventually reach the reader. The received data are decoded in parallel and the resulting decoded data are streamed to the destination core, as illustrated in Figure 2-d. It is worth mentioning that a high modulation speed, noted as F_{mod} , leads to similar design complexity for encoder and decoder.

As previously explained, an optical link manager ensures that the same ECC is used on source and destination sides. It is also responsible for adapting the laser output power, which is achieved by using configurable CMOS current generator. In this work, we assume a single control signal for all the laser

sources on a channel. The ECC and the laser output power are selected according to performances, power and BER requirements. The design of the manager and the implementation of the algorithms are out of the scope of this paper.

C. Electrical/Optical Interfaces

The design of interface is a challenging task due to the different signal frequencies between electrical and optical domains [13] [14]. For example, a modulation frequency F_{mod} set to 10GHz provides a throughput of 10Gbit/s on 1-bit serial data transmission. This is a typical bandwidth in the optical domain while, in the electrical domain, IPs would rather communicate at a much lower frequency, noted F_{IP} , e.g. 1GHz, but with parallel data bus of N_{data} width, such 64-bits.

Figures 2.c and 2.d illustrate the generic micro-architectures of the interfaces. The energy/performance manager (see Section III-C) selects the coding technique according to the communication requirements (e.g. task priority, BER). In addition to the mode with direct modulation (uncoded transmission), different coded modes (Hamming encoders with paths 0, to $n - 1$ in the figure) can be selected. The encoded data are outputted to a mux connected to a serializer in order to generate the stream of data at the required F_{mod} speed. The clock of the serializer register equals the modulation speed. The serializer is a well-known architecture based on a registers pipeline with a depth equal to the input size. These registers can be loaded with the input data by the use of a two-bits mux located in each register input's. Regarding the deserializer, its architecture is also based on a registers pipeline with a depth equal to the output size. These straightforward solutions allow to achieve a high throughput required by the optical modulation speeds while limiting the area (see Section V-A for synthesis results).

The receiver part, illustrated in Figure 2.d, is based on a similar micro-architecture than the emitter. It relies on decoders and on the use of a trans-impedance amplifier associated to a comparator. Obviously, the receiver path must be configured similarly to the emitter, i.e. by using a Hamming decoder with a size similar to the coder. This paper does not investigate data flow control and multi-wavelengths allocation, which have already been efficiently addressed in [13] and [14], respectively.

D. BER and Performance Models

Without coding, the SNR can be calculated from BER with

$$SNR = [\text{erfc}^{-1}(1 - 2 \cdot BER)]^2 \quad (1)$$

where erfc^{-1} is the inverse of the complementary error function. A minimum distance Hamming linear code is defined by an integer $m \geq 2$, a block length of $n = 2^m - 1$ and message length $k = n - m$. Hence the code rate of Hamming codes is $R_c = k/n$, which is the highest possible for codes with minimum distance of three and block length n . The BER of a Hamming code is defined by

$$BER = p - p(1 - p)^{n-1} \quad (2)$$




Fig. 2. ECC configuration and laser power management: detailed principle

$$\text{where } p = \frac{1}{2} \operatorname{erfc}(\sqrt{SNR}) \quad (3)$$

The energy per bit is linked to the SNR with the code rate R_c . Calculating the SNR from BER when considering Hamming codes requires to invert Equations 3 and 2. The optical signal power to be received by the photodetector (OP_{signal}) is defined with

$$SNR = \frac{\mathfrak{R} \times (OP_{signal} - OP_{crosstalk})}{i_n} \quad (4)$$

where $OP_{crosstalk}$ is the worst case crosstalk in the MWSR channel, \mathfrak{R} is the photodetector responsivity (1 A/W in this work) and i_n is the dark current ($4\mu A$). Finally, the minimum laser output power OP_{laser} is estimated using the transmission model proposed in [8]. In this model, the authors evaluate the total signal transmission in a MWSR channel by taking into account the transmission in each MR (i.e. MR in writers and reader) and the waveguide loss. We have chosen this model for its accurate estimation of the crosstalk, which is obtained by taking into account the distance between signal and MR resonant wavelengths.

Figure 3 illustrates the MR signals transmission in ON and OFF states. In the figure, λ_{MR} is the MR resonant wavelength in the OFF state and Δ_λ is the wavelength shift between ON and OFF states. The Extinction Ratio (ER) is the transmission ratio between ON and OFF states for signal at λ_{MR} . Other key characteristics of the modulators, not illustrated in the figure, are the modulation speed and the power consumption P_{MR} . In this work, these values are extracted from [15], which leads to $ER = 6.9dB$ and $P_{MR} = 1.36mW$




Fig. 3. Optical signal transmission in MR

The performance is given by the Communication Time (CT), which corresponds to the relative increase of the transmission time due to parity bits. It is thus normalized to a transmission time without ECC. For instance, when using H(7,4), 75% parity bits are added to the payload which leads to $CT = 1.75$.

E. Power Models

$P_{channel}$ is the channel power consumption per wavelength and is given by:

$$P_{channel} = P_{ENC+DEC} + P_{MR} + P_{laser}$$

with $P_{ENC+DEC}$ the encoder and decoder power, P_{laser} the laser power and P_{MR} the modulation power (all the power values are given per wavelength). $P_{ENC+DEC}$ is obtained from synthesis results and is reported in the results section. P_{laser} is obtained from the required on-chip laser output power and the laser efficiency. We consider the CMOS compatible PCM-VCSELs laser sources described in [16] and we assume

a temperature dependent lasing efficiency. In this paper, we do not take into account the thermal tuning power, which we assume to be the same for communications with and without ECC.

V. RESULTS

A. Synthesis Results

Emitter and receiver interfaces with three possible communication modes (H(7,4), H(71,64), and uncoded) have been synthesized on 28nm FDSOI technology. The interfaces have been designed for a 64-bit bus width IP (i.e. $N_{data} = 64$) clocked at $F_{IP} = 1GHz$ and allowing to reach F_{mod} of 10Gbit/s. The synthesis results are summarized in Table I. To encode the 64-bits of the IP bus, the interface requires 16 parallel H(7,4) coders and decoders, while only one H(71,64) codec is required. The Hamming blocks are clocked at F_{IP} and compute the result in one clock cycle (outputs are clocked with registers). The SER and DES are working at F_{mod} . The critical path results show positive slacks, compared to the aimed frequencies, allowing transmissions at 10 Gbit/s.

The area is $2013\mu m^2$ and $3050\mu m^2$ for the emitter and the receiver, respectively. Static power is negligible thanks to the 28nm low leakage technology. Consumed dynamic power depends on the selected data-path and hence on the active coding blocks. Indeed, the used path is activated by an *enable* signal. The dynamic power of the transmitter ranges from $3.1\mu W$ for an uncoded transmission to $9.5\mu W$ for a transmission with H(7,4). For the receiver interface, it ranges from $4.2\mu W$ for an uncoded transmission to $10.1\mu W$ for a transmission with H(7,4). H(7,4) corresponds to the solution with the highest dynamic power consumption in the interfaces due to the highest overhead of Hamming bits. However, as this coder reaches the highest error correcting efficiency, it also allows for the highest reduction of the laser power.

In the following, we will compare the power consumption of the interface, the modulators and the laser sources. Then, we will investigate whether the improved communication quality made possible using encoded communication reduces the Laser power consumption significantly enough to overcome the increased modulation power consumption.

B. BER and Laser Power Trade-Offs

For on-chip light sources, we assume an integrated CMOS-compatible vertical-cavity surface-emitting lasers (VCSEL) [16]. The temperature of the VCSEL depends on i) the optical power to be emitted (OP_{laser}) and ii) the activity of the electrical layer. Following the methodology detailed in [8], the laser power consumption P_{laser} is estimated assuming 25% chip activity. As illustrated in Figure 4, P_{laser} linearly increases within the $0 - 500\mu W$ optical power range but follows an exponential trend for larger values. This trend is due to the efficiency of the laser which drastically drops with high temperature.

Using the model described in Section IV-D, we evaluate the optical power signal needed to reach the targeted BER according to a given coding technique. For this purpose, we

TABLE I
SYNTHESIS RESULTS OF AN INTERFACES PROVIDING NO ECC, H(7,4) AND H(71,64). $F_{IP} = 1GHz$, $N_{data} = 64bits$, AND $F_{mod} = 10Gbit/s$, TECHNOLOGY IS 28nm FDSOI.

Hardware block	Total Area (μm^2)	Critical Path (ps)	Static (nW)	Power Dynamic (μW)	Total (μW)
Transmitter					
1-bit MUX (3 to 1)	14	80	0.2	0.23	0.23
H(7,4) coders ¹	551	210	1.7	3.13	3.13
H(71,64) coder	490	350	1.6	2.51	2.52
112-bits SER, H(7,4)	433	70	6.5	6.21	6.22
71-bits SER, H(71,64)	276	70	4.1	3.24	3.25
64-bits SER, wo ECC	249	70	3.6	2.93	2.93
Total, H(7,4) com.				9.57	9.59
Total, H(71,64) com.	2013		17.7	5.99	6.01
Total, w/o ECC com.				3.16	3.18
Receiver					
64-bits MUX (3 to 1)	815	80	10.8	1.55	1.56
H(7,4) decoders	783	300	2.5	3.80	3.80
H(71,64) decoder	648	570	2.2	2.63	2.64
112-bits DESER, H(7,4)	365	60	5.5	4.75	4.75
71-bits DESER, H(71,64)	231	60	3.5	3.02	3.02
64-bits DESER, wo ECC	208	60	3.0	2.75	2.75
Total, H(7,4) com.				10.1	10.1
Total, H(71,64) com.	3050		27.5	7.21	7.23
Total, w/o ECC com.				4.29	4.32




Fig. 4. P_{laser} is estimated from the output optical power OP_{laser} , the chip activity and the temperature-dependant laser efficiency.

assume an MWSR architecture with 12 ONIs, 16 wavelengths, 6 cm waveguide length, 0.274dB/cm waveguide loss [17] and $ER = 6.9dB$ [15]. We assume transmissions with H(7,4), H(71,64) and without ECC, and targeted BER values range from 10^{-3} to 10^{-12} , Figure 5. As expected, the transmission without codes leads to the higher laser power consumption. For instance, each laser source consumes $14.3mW$ for $BER = 10^{-11}$. In case H(71,64) is used, P_{laser} drops to $7.12mW$ and it further decreases to $6.64mW$ for H(7,4). This significant reduction in the laser power is due to the added redundancy bits, which improve the optical channel robustness. Since additional errors in the transmission are acceptable, the optical power emitted by the light sources can be reduced. Interestingly, while targeting a 10^{12} BER without ECC is not possible since it exceeds the maximum optical power deliverable by the laser (i.e. $700\mu W$), reaching this BER is possible using H(71,64) and H(7,4) ($P_{laser} = 7.1mW$ and $7.6mW$ respectively). In the following, we will investigate the energy/bit figures and the communication time overheads due to the transmission of the redundancy bits .




Fig. 5. P_{laser} for a given MWSR channels is calculated from the targeted BER and the selected ECC code.

C. Performance and Power Trade-off

Figure 6.a summarizes the channel power consumption per wavelength for $BER = 10^{-11}$. For communication without ECC, the laser sources cost for 92% of the total power. Hence, the laser power reductions achieved using H(71,64) and H(7,4) lead to significant improvement of the total power (-45% and -49% respectively). Regarding the energy-per-bit figures, since H(7,4) is the less energy efficient solution due to the large number of redundancy bits to transmit. In this study, H(71,64) is the most energy efficient solution with 3.76pJ/bit wrt. 3.92pJ/bit for a communication w/o ECC (i.e. 4% improvement). Furthermore, H(71,64) also demonstrates the lower power consumption thanks to the nearly 50% reduction: for the 16-wavelengths channel we consider, the power consumption drops from 251mW to 136mw for each waveguide. Assuming 16 waveguides per MWSR channel and 12 ONIs, the total power saving reaches 22W for the whole interconnect. To summarize, our method allows significant reduction in the power while maintaining the energy/bit figures: this contributes to solve the power wall and temperature (hot spot) issues in nanophotonic interconnects. It is worth mentioning that additional power saving can be achieved by turning off the lasers when no activity is detected [9]. Reducing the power also comes at a price of larger communication time since the size of the data stream proportionally increases with the added redundancy bits. Figure 6.b illustrates the power/performance trade-off for BER ranging from 10^{-6} to 10^{-12} . For a given BER, all the coding techniques belong to the Pareto front, thus demonstrating the potential of the proposed approaches to overcome performance and power issues in nanophotonic interconnects.

VI. CONCLUSION

In this paper, we investigated the use of Error Correcting Code (ECC) to reduce the laser output power according to the communication quality requirement. We have shown that using ECC not only allows reducing the laser power consumption, but also allows to target BER not reachable without coding techniques and without compromising energy per bit. Since lasers significantly contribute to the optical channel power, global power reduction of the optical interconnect is also possible under specific scenario. This trend has been observed especially for H(71,64) ECC. In our future work, we will

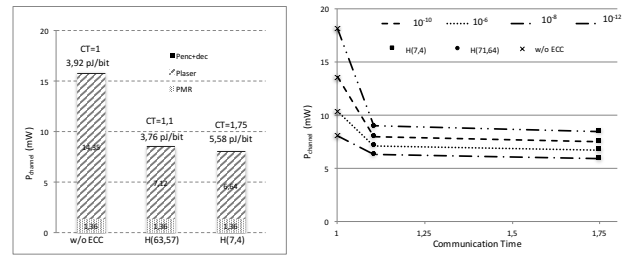


Fig. 6. a. (left) Power contribution in an MWSR channel for $BER = 10^{-11}$. b. (right) Power and performance trade-off wrt. BER and ECC.

further investigate the achievable energy efficiency improvement by simulating the execution of standard benchmark applications on nanophotonic interconnects.

REFERENCES

- [1] C. Sun *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, pp. 534–538, 12 2015.
- [2] Vantrease, D. *et al.*, "Corona: System implications of emerging nanophotonic technology," in *Computer Architecture, 2008. ISCA '08. 35th International Symposium on*, June 2008, pp. 153–164.
- [3] Goldsmith, A. J. *et al.*, "Design challenges for energy-constrained ad hoc wireless networks," *Wireless Commun.*, vol. 9, no. 4, Aug. 2002.
- [4] S. Chouhan *et al.*, "System-level design space exploration methodology for energy-efficient sensor node configurations: An experimental validation," *IEEE TCAD*, vol. 31, no. 4, pp. 586–596, April 2012.
- [5] Alam, Muhammad *et al.*, "On the energy savings of adaptive transmit power for wireless sensor networks radio transceivers," in *Proc. ARCS. VDE*, 2013, pp. 1–6.
- [6] A. Mineo *et al.*, "A closed loop transmitting power self-calibration scheme for energy efficient winoc architectures," in *DATE*, March 2015, pp. 513–518.
- [7] R. Wu *et al.*, "Variation-aware adaptive tuning for nanophotonic interconnects," in *Computer-Aided Design (ICCAD), 2015 IEEE/ACM International Conference on*, Nov 2015, pp. 487–493.
- [8] H. Li *et al.*, "Towards maximum energy efficiency in nanophotonic interconnects with thermal-aware on-chip laser tuning," *IEEE Transactions on Emerging Topics in Computing*, vol. PP, no. 99, pp. 1–1, 2016.
- [9] Demir, Yigit *et al.*, "Towards energy-efficient photonic interconnects," pp. 93 680T–93 680T–12, 2015.
- [10] C. Li *et al.*, "An energy-efficient silicon microring resonator-based photonic transmitter," *IEEE Design Test*, vol. 31, no. 5, 2014.
- [11] Wu, Xiaowen *et al.*, "Suor: Sectioned unidirectional optical ring for chip multiprocessor," *J. Emerg. Technol. Comput. Syst.*, vol. 10, no. 4, pp. 29:1–29:25, Jun. 2014.
- [12] F. Mandorlo *et al.*, "Controlled multi-wavelength emission in full cmos compatible micro-lasers for on chip interconnections," *Journal of Lightwave Technology*, vol. 30, no. 19, pp. 3073–3080, Oct 2012.
- [13] Z. Wang *et al.*, "A holistic modeling and analysis of optical-electrical interfaces for inter/intra-chip interconnects," *IEEE TVLSI*, vol. 24, no. 7, pp. 2462–2474, July 2016.
- [14] Ortin-Onon *et al.*, "Capturing the sensitivity of optical network quality metrics to its network interface parameters," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 15, pp. 2504–2517, 2014.
- [15] Michal Rakowski *et al.*, "Low-power, low-penalty, flip-chip integrated, 10gb/s ring-based 1v cmos photonics transmitter," in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2013*. Optical Society of America, 2013, p. OM2H.5.
- [16] C. Sciancalepore *et al.*, "Thermal, modal, and polarization features of double photonic crystal vertical-cavity surface-emitting lasers," *IEEE Photonics Journal*, vol. 4, no. 2, pp. 399–410, April 2012.
- [17] P. Dong *et al.*, "Low loss silicon waveguides for application of optical interconnects," in *IEEE Photonics Society Summer Topicals 2010*, July 2010, pp. 191–192.