

Influence of Transmission Impairments on the OSMOSIS HPC Optical Interconnect Architecture

Fotini Karinou, Ioannis Roudas, Kyriakos G. Vlachos, B. Roe Hemenway, and Richard R. Grzybowski

Abstract—We examine the impact of transmission impairments on the performance of the optical supercomputer interconnect architecture, initially proposed in the context of the Optical Shared MemOry Supercomputer Interconnect System (OSMOSIS) project. We study two versions of the aforementioned optical interconnect that differ in terms of the number of semiconductor optical amplifiers (SOAs) used as on-off gates. For practical reasons related to packet arbitration, the size of the crossbar switch of the optical interconnect in this study is limited to 64 ports. The switch is based on a broadcast-and-select architecture and employs DWDM in conjunction with 10 Gb/s intensity modulation/direct detection (IM/DD) per wavelength channel. We show, both by experiment and by simulation, that the minimization of the number of SOAs in the optical switch by taking advantage of the cyclic routing capability of optical arrayed waveguide multiplexers/demultiplexers leads to negligible performance deterioration compared to conventional wavelength-space switches which are prohibitive slower and do not use any inherent gain properties like in OSMOSIS.

Index Terms—Optical interconnects, semiconductor optical amplifiers (SOAs), switching.

I. INTRODUCTION

SINCE the performance of high performance computing (HPC) systems (i.e., supercomputers and computer clusters) experiences a tenfold increase every four years [1], it is expected that exascale HPC systems will be developed by 2020 [2]–[3]. Until then, to satisfy the requirements of emerging, bandwidth-demanding applications for HPC systems, it is necessary to gradually replace the inefficient, conventional electronic interconnects with optical ones. For instance, two representative examples of state-of-the-art PetaFlops (PF) supercomputers (BlueWaters [4] and POWER7-IH [5]) use optical interconnects for inter-rack communication.

The Optical Shared MemOry Supercomputer Interconnect System (OSMOSIS) project [6] proposed an optical interconnect architecture for high bandwidth, low latency, cost-effectiveness and scalability. The OSMOSIS optical interconnect uses electronics for scheduling and routing and optics for switching and transmission. Its basic building block is a two-stage, broadcast-and-select, 64×64 optical crossbar

switch fabric for synchronous, fixed-size optical cell switching. The latter is accomplished through semiconductor optical amplifiers (SOAs), acting as on-off gates [6]. The main advantage of the OSMOSIS architecture is that it performs ns-scale switching, as opposed to active optical cables, which are used for dedicated point-to-point links between pairs of nodes. In this sense, OSMOSIS is superior to the currently used optical active cable technology in terms of sharing resources. Nevertheless, the cost of the OSMOSIS architecture is still prohibitive for commercial HPC systems, due to the large number of on-off gates. The problem is exacerbated as the throughput of the interconnect must eventually grow to accommodate exascale traffic.

An economically-viable, multi-stage alternative design of the original two-stage OSMOSIS [6] crossbar switch fabric architecture, targeted at the minimization of the number of SOAs, was recently proposed [7]. More specifically, the multi-stage $N \times N$ optimum interconnect alternative design proposed in [7] can reduce the number of on-off gates from $2N\sqrt{N}$, that are required in the original two-stage OSMOSIS architecture [6], down to asymptotically $N \ln N$, where N is the number of nodes to be interconnected.

In a preliminary study [8], we experimentally investigated the performance of the multi-stage optimized crossbar switch fabric employing polarization division multiplexing (PDM) quadrature phase shift keying (QPSK) modulation and coherent intradyne detection. In a more recent publication [9], we showed that the economically-viable optimized optical switch fabric performs almost equally well to the original one when using conventional intensity modulation/direct detection (IM/DD).

Elaborating on the work of [9], in this paper, we assess the physical layer performance of the optimized 64×64 three-stage, OSMOSIS optical switch fabric and compare it to its two-stage original counterpart, both experimentally and by simulation. In particular, we evaluate the impact of SOA nonlinearities, optical filter concatenation and amplified spontaneous emission (ASE) noise accumulation, on the performance of both interconnect architectures, using 10 Gb/s IM/DD serial optical transmission. Simulation and experiment show that the optimized, cost-efficient OSMOSIS crossbar switch fabric performs almost as well as the original one, despite the fact that, in the former, optical signals travel through more concatenated SOAs.

The remainder of the paper is organized as follows: In Section II, we compare the original and the optimized crossbar

Manuscript received June 07, 2011.

F. Karinou, I. Roudas, and K. Vlachos are with the University of Patras, Patras, 26500, Greece, (F. Karinou: +30-2610-969864; fax:+30-997342; e-mail: {karinou, roudas}@ece.upatras.gr, kvlachos@ceid.upatras.gr).

B. R. Hemenway and R. R. Grzybowski are with Corning Inc., Corning, NY 14831, (email: {hemenwaybr, grzybowski}@corning.com).

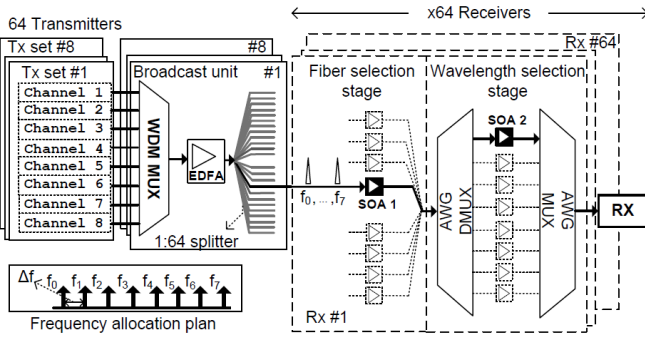


Fig. 1 Block diagram of the originally-proposed 64×64 OSMOSIS optical interconnect architecture [6]. Eight identical sets of eight equidistant carrier frequencies are multiplexed and transmitted over eight separate fibers. On each receiver card, there are two discretely-tunable stages. Eight SOAs are used per selection stage (i.e., 16 SOAs per receiver card). (Symbols: EDFA=Erbium-doped fiber amplifier, SOA=Semiconductor optical amplifier, AWG MUX/DMUX=Arrayed waveguide grating multiplexer/demultiplexer, Tx=Transmitter, Rx=Receiver).

switch designs and describe the simulation block diagrams used to evaluate their performance. Their performance is also assessed through experimental measurements using the setup described in Section III. An itemized account of the penalties due to various transmission effects is presented in Section IV. Details of the simulation models are given in Appendix A. An analytical calculation to justify the results shown in IV, is presented in Appendix B.

II. OPTICAL INTERCONNECT ARCHITECTURES

In this Section, we describe the originally-proposed [6] and the optimized [7], 64×64 OSMOSIS optical interconnect architectures and their simplified simulation block diagrams. Both architectures use fixed-wavelength transmitters and discretely-tunable, direct-detection receivers. However, they differ in the organization of the transmitters into different multiplexing hierarchies and the number of stages in the discretely-tunable receivers that perform the selection of the desired channel.

A. Original OSMOSIS architecture

In the original OSMOSIS architecture [6], depicted in Fig. 1, the 64 transmitters are partitioned into eight sets. The transmitters of each set, are assigned eight equidistant carrier-frequencies, f_0, \dots, f_7 . The frequency allocation plan for this scheme is shown as an inset in Fig. 1. Frequency reuse is employed among different transmitter sets. Each transmitter in Fig. 1, comprises a continuous wave (CW) laser. The output signal from each CW laser is first amplified by an erbium-doped fiber amplifier (EDFA) and then, is intensity modulated by a 10-Gb/s, non-return-to-zero (NRZ), pseudo-random Binary Sequence (PRBS), using a Mach-Zehnder modulator (MZM). All channels are synchronized in time. Each set of eight channels is wavelength division multiplexed (WDM) on a different fiber, the WDM signal is amplified by an EDFA,

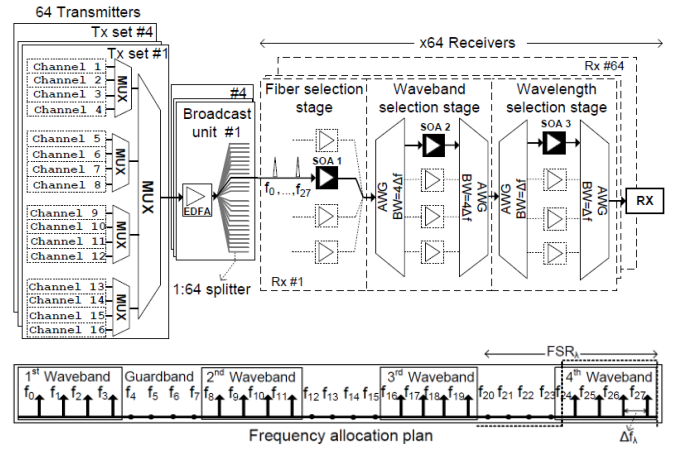


Fig. 2 Block diagram of the optimized 64×64 OSMOSIS architecture [7]. The same set of sixteen frequencies is multiplexed at the transmitter side using three multiplexing stages grouping together sets of four tributaries. Three selection stages with four SOAs per stage (i.e., 12 SOAs/Rx card in total) are needed to perform wavelength selection at the receiver. (Symbols: EDFA=Erbium-doped fiber amplifier, SOA=Semiconductor optical amplifier, AWG=Arrayed waveguide grating, BW=3-dB Bandwidth, Rx=Receiver).

and is broadcasted to all 64 receiver cards, using 1:64 star couplers. The 64 discretely-tunable receivers consist of two selection stages, choosing a single set of eight wavelengths and a single wavelength channel, respectively. Each selection stage consists of eight SOAs, acting as ON/OFF gates.

For the simulation, the channel spacing Δf is 100 GHz. More specifically, the carrier frequencies of the transmitted channels are 192.1 THz-192.8 THz (corresponding to wavelengths 1554.94-1560.60 nm). They are placed uniformly around the peak of the SOA gain, at 1557.77 nm, so they experience only a small gain variation due to the non-uniform SOA gain profile. The CW laser average power is set to -3 dBm per channel. The EDFA before the MZM has 17 dB gain to compensate for losses in the MZM and the consecutive MUX. The WDM signal reaches the second EDFA (acting as a booster amplifier) without any power variation among the eight channels. The maximum power variation of the channels at the output of the second EDFA is less than 1 dB, due to the spectral tilt in the EDFA gain profile. We assume that the SOA gain has a parabolic shape around the SOA gain peak. Uniform channel placement around the peak of the SOA's gain curve minimizes power variation. More specifically, individual channels exhibit a maximum power variation of approximately ~ 1 dB, both at the input and at the output of the first SOA, as well as at input of the second SOA. The SOAs work in the linear regime, close to their saturation point. Their input saturation power is $P_{in}^{sat} = 4$ dBm (see Fig. 9 in Appendix A). Additional ASE noise is loaded to the signal to vary the optical signal-to-noise ratio (OSNR) and estimate the error probability. After the second selection stage, the desired wavelength is filtered by an optical Gaussian filter at the entrance of the optically-pramplified receiver, with an equivalent noise bandwidth $B_0=34$ GHz [10].

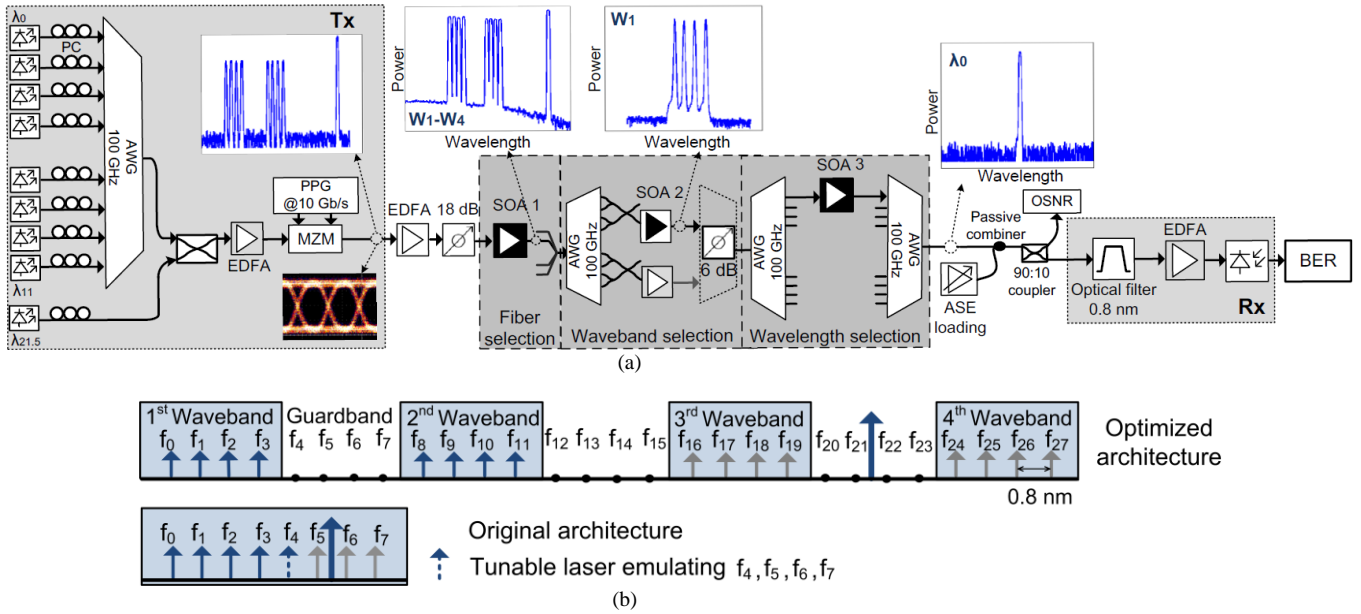


Fig. 3 (a) Experimental set-up for the study of the performance of the optimized OSMOSIS optical interconnect architecture. The black SOA boxes represent SOAs in the ON state, whereas grey SOA boxes represent SOAs that are turned OFF. Representative optical spectra are shown after successive selection stages. (b) Frequency allocation plan for the optimized and the original interconnect architectures, respectively. (Symbols: Grey arrows: original channel plan, Blue arrows: emulated channel plan, AWG=Arrayed waveguide grating, EDFA=Erbium-doped fiber amplifier, PPG=Pulse pattern generation, MZM=Mach-Zehnder modulator, SOA=Semiconductor optical amplifier, ASE=Amplified spontaneous emission noise, Rx=Receiver, and BER=Bit error rate).

B. Optimized OSMOSIS architecture

The proposed optimized 64×64 optical interconnect architecture is presented in Fig. 2. The 64 transmitters are partitioned into four sets. The transmitters of each set are assigned 16 carrier frequencies. Frequency reuse is employed among different transmitter sets. The frequency allocation plan is shown in the inset of Fig. 2. The sixteen carrier frequencies are grouped into four wavebands in sets of four, occupying an aggregate bandwidth of 2.7 THz. Guard bands facilitate waveband multiplexing/demultiplexing. The carrier frequencies span from 192.1 THz to 194.8 THz (corresponding to the wavelength range 1538.98 nm-1560.60 nm). The channel spacing within a waveband Δf is 100 GHz. The signals of each transmitter set are wavelength division multiplexed on a separate optical fiber and broadcasted to all 64 receiver cards using 1:64 star couplers. The discretely-tunable receivers use three selection stages for choosing a fiber, a waveband, and a wavelength channel, respectively. The wavelength allocation is done in a way that the periodicity of arrayed waveguide grating (AWG) MUX/DMUXs is exploited. More specifically, to reduce the number of SOAs in the proposed architecture, we use periodic MUX/DMUXs at the wavelength selection stage, with a free spectral range of $FSR=8 \Delta f$. After the third selection stage, additional ASE noise is loaded to the signal in order to assess the transmission performance. At the entrance of the optical preamplified receiver, the desired wavelength is filtered by an optical Gaussian filter, similar to the one used in the original architecture.

III. EXPERIMENTAL SETUP

The experimental setup used to measure the performance of the aforementioned, optimized 64×64 optical switch fabric is depicted in Fig. 3. Due to lack of resources, nine semiconductor lasers on the transmitters’ side emulate all 16 wavelength channels per fiber shown in Fig. 2. For the same reason, the wavelength channel distribution with reference to the SOA’s gain peak is not the same in the simulation and in the experiment. More specifically, eight DFB lasers, with carrier frequencies spaced by 100 GHz, are used to represent the first two wavebands, W_1 - W_2 . The carrier frequencies span from 192.469 THz to 193.563 THz (corresponding to the wavelength range 1548.808 nm-1557.608 nm). A ninth (tunable) semiconductor laser, with eight times the nominal average power of a single WDM channel, located at the frequency slot $f_{21.5}=192.164$ THz ($\lambda_{21.5}=1566.008$ nm), is used to represent the remaining two wavebands W_3 - W_4 (Fig. 3(a)). The nominal and the experimentally implemented frequency allocation plans are shown in Fig. 3(b), (grey and blue arrows, respectively). The CW optical signals from all nine lasers are initially combined and preamplified by an EDFA. Then, the WDM signal is modulated using a single MZM modulator by a 10-Gb/s, non-return-to-zero (NRZ), Amplitude Shift Keying (ASK), 2^7-1 pseudo-random binary sequence (PRBS). This way, time-aligned, identical wavelength channel bit sequences are generated. This corresponds to the worst-case scenario for studying the impact of cross-gain modulation (XGM) due to SOAs. Subsequently, the signal is first amplified again using a booster EDFA, and then it goes through a variable optical attenuator (VOA) with 18 dB loss, that emulates the 1:64 star

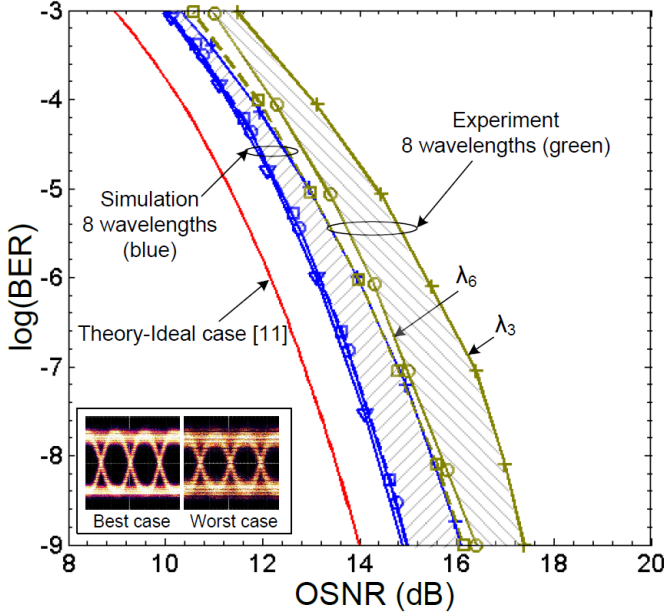


Fig. 4 BER vs OSNR (measured in a resolution bandwidth $RB=0.08$ nm) for the original architecture for all eight wavelengths (λ_0 - λ_7) investigated both by simulation and by experiment. Red solid curve: theoretical curve [11], Squares: back-to-back transmission (wavelength used: λ_6), Triangles: single channel transmission (wavelength used: λ_6), Circles: best case transmission scenario (λ_6 for simulation and experiment), Crosses: worst case transmission scenario (λ_1 and λ_3 for simulation and experiment, respectively). Blue and green color represents simulation and experimental results, respectively.

coupler of the actual OSMOSIS architecture. On the receiver side, the WDM signal passes through three cascaded selection stages that employ SOAs as on-off gates. More specifically, the WDM signal originating from any one of the four fibers can be chosen by the first selection stage, composed of a SOA, and a 4:1 combiner. The second selection stage is used to select the desired waveband. Due to lack of resources, it consists of a 100-GHz AWG MUX with interconnected arms to emulate the 400-GHz AWG required in the proposed architecture [7], a second SOA, and an attenuator of 6 dB to emulate another 400-GHz AWG MUX, not available in the lab. The use of a 100-GHz AWG with interconnected arms as a substitute for the 400-GHz AWG is not completely accurate, and results in the filtering of some out-of-channel ASE noise in each waveband before amplification. Nevertheless, we anticipate that the performance would be affected only slightly by this substitution. The same approach was employed also in [9]. Finally, the desired channel is selected via a third wavelength selection stage, which consists of a pair of conventional AWG MUX/DMUX, with 100-GHz spacing, and four SOAs. In our experiment, only the three SOAs that work in the ON state (i.e., black SOA boxes in Fig. 3(a)) are used. All SOAs have a 3-dB bandwidth of 90 nm, a small-signal gain of 15 dB, a gain peak at 1490 nm, a high input saturation power of approximately 4 dBm (see Appendix A, Fig. 9), and a low polarization dependent gain of <0.2 dB.

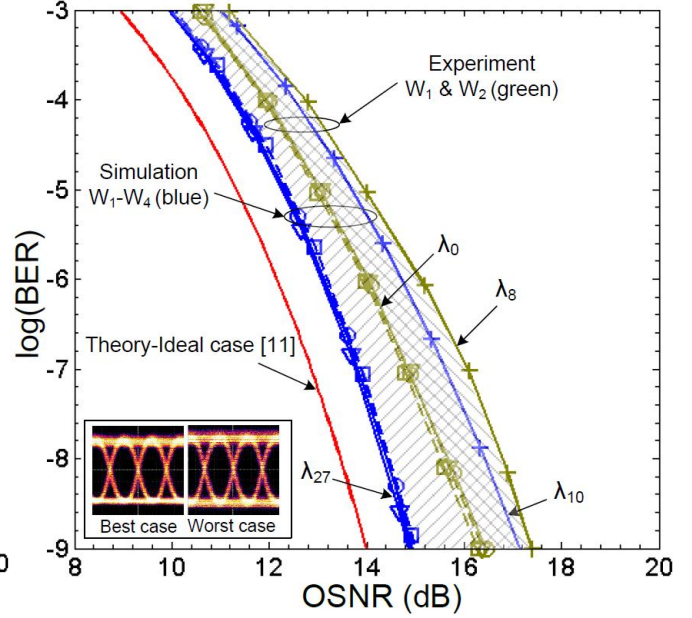


Fig. 5 BER vs OSNR (measured in a resolution $RB=0.08$ nm) for the optimized OSMOSIS architecture. Wavebands W_1 - W_4 (λ_0 - λ_{27}) were investigated both by simulation and by experiment. Red solid curve: theoretical curve [11], Squares: back-to-back transmission (wavelength used: λ_{27} and λ_0 for simulation and experiment, respectively), Triangles: single channel transmission (wavelength used: λ_{27} and λ_0 for simulation and experiment, respectively), Circles: best case transmission scenario (λ_{27} and λ_0 for simulation and experiment, respectively), Crosses: worst case transmission scenario (λ_{10} and λ_8 for simulation and experiment, respectively). Blue and green color represents simulation and experimental results, respectively.

A second experimental set-up, described in detail in [6], is used for evaluating the performance of the original OSMOSIS optical switch fabric. Its description is omitted here for brevity. Only, the implemented frequency allocation plan for that experiment is shown in Fig. 3(b). The carrier frequencies span from 192.866 THz to 193.563 THz (corresponding to the wavelength range 1548.808 nm-1554.408 nm (Fig. 3(b)). Another laser with three times the nominal average power of a single WDM channel, located at the frequency slot $f_{5,5}=193.015$ THz ($\lambda_{5,5}=1553.208$ nm), is used to represent the lasers sources not available in the lab (Fig. 3(b)).

IV. SIMULATION AND EXPERIMENTAL RESULTS

In this Section, we evaluate the performance of both optical interconnect architectures using the raw BER as a criterion. In practice, the raw BER is reduced by using additional forward error correction (FEC) coding and/or automatic repeat re-request (ARQ) protocol. We investigate the impact of SOA nonlinear effects, such as self-gain modulation (SGM), XGM, and four wave mixing (FWM), optical bandwidth narrowing due to AWG concatenation, and ASE noise accumulation on the performance of the optical interconnect. We show that the optimized architecture performs well, almost equally to the original one, despite the stricter limitations imposed by the additional stage of SOAs.

We define here three reference systems used in the following subsections for performance comparison: i) by the term 'ideal system' we refer to the one described in [11], where optical signal is assumed to be distortionless and the ASE noise is filtered at the IM/DD receiver by a brickwall optical BPF and an integrate-and-dump LPF (in the absence of a polarizer); ii) the term 'back-to-back' refers to a hypothetical scenario where the selection stages of the receiver (i.e., grey boxes in Fig. 3(a)) are omitted; and iii) the term 'single-channel transmission' refers to the hypothetical case where only one wavelength, i.e., the one that experiences the best or worst performance, is transmitted through the OSMOSIS architecture. A commercially-available software tool (VPI TransmissionMaker) enhanced with custom-made modules in Matlab for OSNR measurements, was used for carrying out the simulations shown [12]. Moreover, Mathematica software was also used to validate the SOA model.

A. Overview

The BER is assessed as a function of the received OSNR for each of the transmitted wavelengths, both by experiment and simulation. The final results are shown in Fig. 4 and Fig. 5, for the original and the optimized architecture, respectively. The hatched areas between two consecutive curves in both figures, represent families of closely spaced BER curves corresponding to different wavelength channels.

In Fig. 4, measurements for the original architecture show that the spread among all eight channels is approximately 1 dB, both in simulation and experiment.

Experimental results in Fig. 5 reveal that there is a spread of 1.2 dB, in terms of required OSNR, among the eight channels of waveband W_1 and W_2 for $BER = 10^{-9}$. In close agreement with the experiment, simulation results indicate that there is a OSNR spread of approximately 1.45 dB among the eight channels of wavebands W_1 and W_2 . However, simulation also shows that there is a spread of approximately 2 dB between all 16 simulated channels of the optimized architecture for $BER = 10^{-9}$. The BER curve for an ideal system [11] is also shown for comparison. For qualitative comparison, experimental eye diagrams are also shown as insets in Fig. 4-5, for the best and the worst channels, for both architectures.

Finally, the performance of both architectures is also tested without in-line SOAs, in a back-to-back configuration (squares in Fig. 4 and Fig. 5), as well as for single-channel transmission, after the three selection stages of the optimized configuration (triangles in Fig. 5). We observe that the penalty difference is negligible among the back-to-back, the single-channel, and the best case transmission scenario (i.e., λ_{27}) for the optimized architecture (Fig. 5). Similarly, the curves corresponding to the back-to-back and the best case transmission scenario (i.e., λ_6) for the original architecture (Fig. 4) are indistinguishable.

It is worth noting that the measured performance, in both architectures, is worse than the one predicted by simulation by approximately 1 dB (e.g., compare the back-to-back cases for the simulation and the experiment, respectively). This small

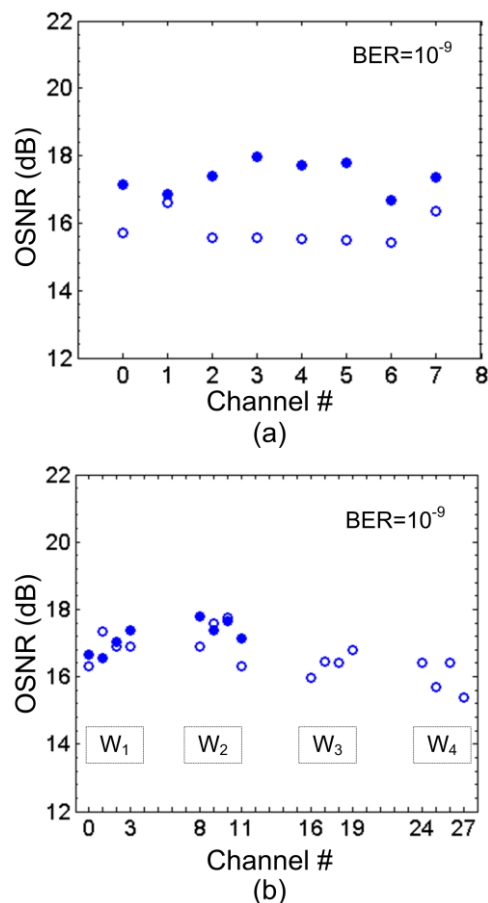


Fig. 6 OSNR (measured in a resolution bandwidth $RB=0.08$ nm) required for $BER=10^{-9}$ for all channels: (a) original and (b) optimized optical interconnect architecture. (Symbols: W_1 - W_4 : wavebands used in the optimized architecture; Open and filled circles: simulated and experimental results, respectively).

difference is attributed to the following parameter mismatch between simulation and experiment: i) the use of narrower optical filters in the simulation; ii) the use of the same data pattern for modulating all wavelength channels in the experiment, in contrast to the simulation, where each laser is independently modulated with a different bit sequence (a 4 dB degradation is observed using the same data patterns); and iii) the omission of polarization dependent gain (PDG) from the SOA simulation model [13]. Despite this small discrepancy, from Fig. 4-5, we can safely conclude that the performance of the optimized architecture is slightly worse than the performance of the original one for the worst channel (whereas it is identical for best channel). We point out here, that the experimental results shown in Fig. 4 and Fig. 5 are optimized compared to the ones reports in [9] by 1.4 dB, a penalty found to be due to the low extinction ratio (ER) that holds in the experiment [14]. The ER in [9] was 8 dB, while in the present work the ER was optimized by 5 dB to avoid penalty due to a reduced modulation ER.

In Fig. 6, the OSNR (measured in a resolution bandwidth of 0.08 nm and expressed in dB) required for a $BER = 10^{-9}$ is shown for each of the transmitted wavelengths of both architectures. The OSNR variation among channels in both simulation and experiment is approximately the same, about 1 dB. In Fig. 6(b) the OSNR values experience a degradation




Fig. 7 BER vs OSNR (measured in RB=0.08 nm) (a) for the worst channel of the original OSMOSIS architecture and (b) for the worst channel of the optimized OSMOSIS architecture. (Symbols: ideal system [11]: red line; Back-to-back: blue circles; Single-channel transmission through the actual interconnect: red squares; WDM transmission through the actual interconnect after the substitution of SOAs and EDFAs by ideal, flat-gain amplifiers: black crosses; WDM transmission through the actual interconnect i.e., a concatenation of two and three SOAs, for the original and the optimized scheme, respectively: dash-dotted curve with stars.

from longer to shorter wavelengths (i.e., a higher OSNR value is needed for the same BER moving from longer to shorter wavelengths). This is partially attributed to the noise figure (NF) of the SOAs, which increases for shorter wavelengths [17].

No FWM [18]–[19] products were observed in the recorded spectrum for the specific operating conditions and the wavelengths under test, in the experiment. This indicates that FWM is not a major factor to the observed signal.

From the above, we conclude that the optimized architecture provides a good trade-off between SOA count reduction and performance degradation. Its relatively small penalty in required OSNR compared to the original architecture justifies its use.

In the following subsections *B* to *D*, using simulation, we quantify the contribution of different transmission effects in the performance degradation of both architectures.

B. Penalty due to SGM

As its name indicates, SGM is gain modulation due to instantaneous channel power variation at the SOA input [20]. In this subsection, we focus on the penalty due to SOA SGM in both optical interconnect architectures under study. To distinguish between the penalty due to SGM and the total system penalty, we perform simulations using a single channel, i.e., the worst and the best channel, in both architectures.

The results for the worst channel are shown in Fig. 7 (red squares). The curves for the ideal system (red line), the back-to-back case (blue circles), and the case of WDM transmission through the interconnect (dash-dotted curve with stars), are also included for comparison. An additional curve corresponding to the WDM case, obtained by substituting all SOAs and EDFAs by ideal, flat-gain amplifiers, is also shown (black crosses).

Comparing the back-to-back case with the single-channel transmission, we conclude that the penalty is negligible in the original architecture (Fig. 7(a)) and 1.5 dB in the optimized one (Fig. 7(b)) at an error probability of 10^{-9} . In contrast, for the best channels, the penalty is negligible in both architectures (not shown here to avoid clutter). The difference of 1.5 dB in performance between the worst and the best channel of the optimized interconnect is explained as follows: the worst channel reaches higher power levels, due to its allocation closer to the gain peak of the SOAs, and therefore, is more affected by SGM. As expected, SGM is more severe in the optimized architecture than in the original one because the signal passes through one more SOA in the former case.

C. Penalty due to XGM

SOAs are subject to XGM [21] that results into data pattern-dependent crosstalk among WDM signals when ASK is used [22]–[23]. The use of quasi-constant envelope modulation formats, such as return-to-zero (RZ) differential phase-shift keying (DPSK) [15], [22], RZ differential quadrature phase-shift keying (RZ-DQPSK) [16], or PDM-QPSK [8]–[9], has been proposed to counteract this effect. We can assess the penalty due to XGM from Fig. 7, by comparing the results for WDM transmission through the optical interconnect after substitution of SOAs and EDFAs by ideal, flat-gain amplifiers.

The penalty due to XGM is negligible for the best channel in both architectures. Plots of the best channel performance are not shown in Fig. 7 to avoid clutter. In contrast, the XGM penalty is equal to 1 dB, for the worst channel of the original architecture (Fig. 7(a)), and 2 dB, for the worst channel of the optimized architecture (Fig. 7 (b)), respectively.

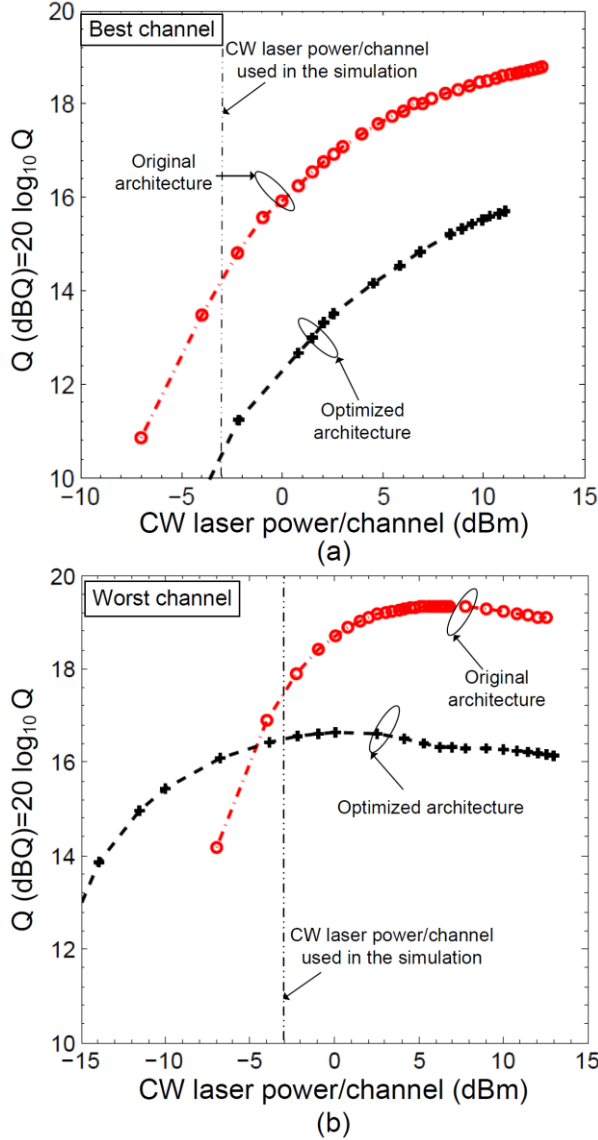


Fig. 8 Q-factor vs CW laser power (a) for the best and (b) for the worst channel, of both the original and the optimized architecture. (Symbols: Red dashed line with circles: WDM transmission in the original architecture, Black dashed line with crosses: WDM transmission in the optimized architecture).

In the original architecture, the WDM signal bandwidth is smaller so channels are less affected by the SOA gain non-uniformity. More specifically, the eight channels are located uniformly around the SOA's gain peak and occupy a bandwidth of 800 GHz (inset in Fig. 1). In particular, the eight channels exit the first SOA with negligible difference in power. On the other hand, in the optimized architecture, the 16 channels occupy a bandwidth of 2.7 THz (inset of Fig. 2) around the SOA's gain peak, and, consequently, they experience a larger power variation. In addition, power variation as a function of wavelength, due to the gain non-uniformity of the EDFAs and the SOAs, and the insertion loss non-uniformity of the cascaded AWGs exacerbates the power variation among channels. This leads to a different behavior among channels with respect to XGM and to higher penalties compared to the original architecture.

Fig. 8(a),(b) show the Q-factor as a function of transmitted power per channel, for the best (λ_6 and λ_{27}) and the worst (λ_1 and λ_{10}) channels of the original and the optimized architecture.

In both architectures, the performance reaches an optimum and remains at a constant level. The ceiling of the Q-factor for higher power levels is analytically explained in Appendix B.

D. Penalty due to the concatenation of optical MUX/DMUXs

AWG MUX/DMUX concatenation leads to narrowing of the optical bandwidth, which, in turn, results in signal attenuation and distortion [10]. In this subsection, we evaluate, by simulation, the penalty due to the narrowing of the optical bandwidth of the aggregate transfer function of the cascaded AWGs, in both interconnect architectures under study.

We consider conventional AWGs with Gaussian amplitude transfer function and linear phase transfer function. To focus on filter-induced distortion exclusively, we substitute all EDFAs and SOAs, in Fig. 1 and Fig. 2, with ideal (flat-gain) amplifiers. In this way, we neglect all transmission effects related to optical amplifiers.

Plots of BER as a function of OSNR are shown in Fig. 7, for the worst case transmission scenario, for both architectures. Comparing the curve for the back-to-back case to the one for the ideal flat-gain EDFAs, we conclude that this penalty is negligible in both cases. These results can be easily explained by the following analysis.

The Gaussian transfer function of each of the four AWGs is given by the formula

$$H_i(f) = A_i e^{-\frac{1}{2} \left(\frac{f-f_{0i}}{f_{ci}} \right)^2} \quad (1)$$

where, A_i is the amplitude, f_{0i} is the center frequency and f_{ci} is the cut-off frequency of the i -th AWG ($i = 1, \dots, 4$).

Assuming that all center frequencies are perfectly aligned to the channel carrier frequency f_0 , the aggregate transfer function $H_{tot}(f)$ of the four concatenated AWGs is

$$H_{tot}(f) = \prod_{i=1}^4 H_i(f) = A e^{-\frac{1}{2} \left(\frac{f-f_0}{f_c} \right)^2} \quad (2)$$

where f_c is the overall cut-off frequency

$$f_c = \left(\sum_{i=1}^4 f_{ci}^{-2} \right)^{-\frac{1}{2}} \quad (3)$$

By substituting $f_{c1} = f_{c2} = 200$ GHz and $f_{c3} = f_{c4} = 20$ GHz, for the cut-off frequencies of the AWG pairs of the discretely-tunable Rx in the OSMOSIS optimized interconnect, we find that $f_c \approx 14$ GHz. The equivalent noise bandwidth of the

TABLE I
EFFECTIVE SIMULATION PARAMETERS

Parameter	Unit	Description	Value
Γ		Confinement factor	0.13
α_0	m^2	Differential gain coefficient	5.3×10^{-20}
N_{tr}	m^{-3}	Carrier density at transparency	1.56×10^{24}
α_s	m^{-1}	Internal loss	1725
L	m	Device length	10^{-3}
w	m	Active layer width	3×10^{-6}
d	m	Active layer thickness	5×10^{-7}
A	s^{-1}	Linear recombination coefficient	3×10^8
B	$m^3 s^{-1}$	Bimolecular recombination coefficient	2×10^{-16}
C	$m^6 s^{-1}$	Auger recombination coefficient	6×10^{-41}

aggregate transfer function is $B_0 = \sqrt{\pi}f_c \approx 25$ GHz [10]. Given that the 95% of the power of an ideal 10 Gb/s NRZ ASK signal occupies 30 GHz [24] we conclude that the signal distortion due to the bandwidth narrowing arising from MUX/DMUX concatenation would be negligible, as indicated by the simulation results shown in Fig. 7(b).

The above analysis does not take into account the increased filtering of AWG MUX/DMUXs in the presence of spectral broadening due to SPM and XPM in SOAs. This is taken indirectly into account in Fig. 7(a), (b) (dash-dotted curves with stars).

V. SUMMARY

In this paper, we assessed the performance of a recently-proposed, economically-viable, broadcast-and-select, 64×64 , optical interconnect architecture [7] using a minimum number of SOA-based on-off gates. We compared its performance to its originally proposed counterpart [6], both by experiment and by simulation. We showed that both interconnect architectures perform almost equally well, while the optimized one is advantageous in terms of the number of SOAs required to perform permutation switching.

To examine the impact of SOA nonlinearities and ASE noise accumulation on the performance of the optical interconnect architectures under consideration, each effect was studied separately. The proposed optimized architecture proved to be tolerant in SOA's nonlinearities. SGM and XGM worsen the performance of the optimized architecture in terms of required OSNR for error free operation only by about 1 dB for the worst case scenario, when compared to the corresponding performance of the original configuration. FWM and the bandwidth narrowing imparted by the concatenation of AWGs in the architecture proved to have




Fig. 9 Fitting experimental data for the SOA gain (blue) and for the output power (green), as a function of input power using the TLM simulation model. (Symbols: P_{in} =Input SOA power, P_{out} =Output SOA power, Solid line: Simulation, Crosses: Experimental values).

negligible impact on the degradation of the performance of the optimized interconnect. ASE noise accumulation and WDM channel power variation due to the wavelength dependent gain of the SOAs affect the performance of the optical interconnect too, resulting both in unequal accumulation of ASE noise for each channel, as well as in channel power variations at the receiver. Finally, we experimentally studied the performance of both optical interconnect architectures. Experimental and theoretical results were in good qualitative agreement.

APPENDIX A

In this Appendix, we briefly describe: i) the SOA model used in the simulation and its simulation parameters, and ii) the formula used for the error probability for the ideal system.

A. SOA model validation

For the simulation of the SOAs, the Transmission Line Model (TLM) was used [13]. In the simulation, we assumed that the SOAs in the simulation have 1-mm length and a 90-nm 3-dB gain bandwidth. The other parameters used in the SOA model are extracted by least squares fitting of experimental data for the SOA gain and output power as a function of input SOA power (Fig. 9) [25]. The most important SOA parameters are shown in Table 1.

B. Error probability evaluation of the ideal system

As a reference, in Fig. 4, 5 and 7 we plot the error probability of an IM/DD receiver with a brickwall optical BPF and an integrate-and-dump LPF, in the absence of a polarizer [11]. As a sanity check, the accuracy of the semi-analytical method used in simulation [12], [26] for the calculation of error probability, was compared to the accurate analytical

method of [11]. The two methods were found to be in good agreement.

The analytical model of [11] takes into account the non-Gaussian ASE noise statistics at the receiver but neglects any signal distortion due to filter-induced intersymbol interference. The analytical expression of the error probability P_e for the receiver of [11], in the absence of a polarizer, is:

$$P_e = \frac{1}{2} \left(1 + 2\bar{N}_{ph}d \right) e^{-2\bar{N}_{ph}d} + \frac{1}{2} \left[1 - Q_2 \left(2\sqrt{\bar{N}_{ph}}, 2\sqrt{\bar{N}_{ph}d} \right) \right] \quad (4)$$

where \bar{N}_{ph} is the average number of photons at the receiver. The latter is related to the OSNR, calculated in both polarizations, through the equation, $OSNR = P_s / (2h\nu\Delta\nu nsp) = \bar{N}_{ph} / 2$. The last equality holds for a resolution bandwidth $\Delta\nu = 2B_e$, a LPF equivalent noise bandwidth $B_e = 1/(2T_b)$ where T_b is the bit period, a spontaneous emission factor $nsp = 1$, and an infinite extinction ratio. In (4), d is the normalized optimum threshold, $0 \leq d \leq 1$, and Q_m is the generalized Marcum's function [27]

$$Q_m(a, b) = \int_b^\infty x \left(\frac{x}{a} \right)^{m-1} e^{-\frac{x^2+a^2}{2}} I_{m-1}(ax) dx \quad (5)$$

where $I_m(x)$ is the m -th order modified Bessel function of the first kind.

The deterministic semi-analytical method [26] for error probability estimation is used in the simulations. When the deterministic approach is used, the BER is calculated from the deterministic signal and statistical properties of the optical, thermal and shot noises. The module finds the exact (nongaussian) moment generating function (MGF) of the detected signal, taking into account the optical noise spectral shape, relations between the signal and noise polarization states, thermal and shot noises of the receiver and correlations due to postdetection filtering. The bit error rate is then calculated from the MGF using the saddle-point approximation technique [12].

APPENDIX B

In this Appendix, using small-signal analysis, we interpret the Q-factor ceiling observed in Fig. 8, when the launched power per channel is high. We assume N amplitude shift keying (ASK) signals of equal average power at the SOA input. Let G be the instantaneous SOA gain, G_s the SOA small-signal gain, P_{tot} the instantaneous total input power, and P_s the SOA input saturation power. From the SOA rate equation for the optical power [28], assuming zero internal loss and constant carrier concentration, yields

$$G = G_s e^{-\frac{(G-1)P_{tot}}{P_s}} \quad (6)$$

Assume that the total input power to the SOA is $P_{tot} = P_0 + \delta P$, where P_0 is the sum of average powers of all channels and δP is a small perturbation ($\delta P \ll P_0$) due to the sum of the instantaneous power variation of all channels. Substituting in (6), the amplifier gain is

$$G_0 + \delta G = G_s e^{-\frac{(G_0 + \delta G - 1)(P_0 + \delta P)}{P_s}} \quad (7)$$

Expanding the terms in parentheses and ignoring the second-order term $\delta G \delta P$ yields

$$\delta G = -\frac{G_0(G_0 - 1)\delta P}{P_s + G_0 P_0} \quad (8)$$

The instantaneous power per channel at the output of the SOA amplifier is

$$P^{out} + \delta P_{out} = (G_0 + \delta G)(\bar{P} + \delta P_i) \quad (9)$$

where \bar{P} is the average input power per channel and δP_i is the instantaneous power variation per channel ($\delta P_i \ll \bar{P}$). Given that $P^{out} = G_0 \bar{P}$ and neglecting $\delta G \delta P_i$, we get

$$\delta P_{out} = \bar{P} \delta G + G_0 \delta P_i \stackrel{(8)}{=} \left[-\frac{G_0(G_0 - 1)}{P_s + G_0 P_0} \right] \bar{P} \delta P + G_0 \delta P_i \quad (10)$$

Because δP , δP_i are independent random variables with zero mean value and variances $\sigma_{\delta P}^2$ and $\sigma_{\delta P_i}^2$, respectively, δP_{out} is also a random variable with zero mean value and variance

$$\sigma_{\delta P_{out}}^2 = \left[\frac{G_0(G_0 - 1)}{P_s + G_0 P_0} \right]^2 \bar{P}^2 \sigma_{\delta P}^2 + G_0^2 \sigma_{\delta P_i}^2 \quad (11)$$

For each transmitted channel the instantaneous input power, P_i can be explicitly written as

$$P_i = \bar{P} + b_i \Delta \quad (12)$$

where b_i is a binary random variable taking values in the set $\{-1, 1\}$, and Δ is the fraction of power added to or subtracted from the average input power to yield a logical ONE or ZERO.

If r is the extinction ratio, defined as $r = \frac{P_{ZERO}}{P_{ONE}}$ [23], where

P_{ZERO} , P_{ONE} are the instant powers for the ZEROS and ONES, respectively, Δ can be expressed as

$$\Delta = \frac{\bar{P}}{1+r} (1-r) \quad (13)$$

The variance of δP_i is

$$\sigma_{\delta P_i}^2 = \bar{P}^2 \left(\frac{1-r}{1+r} \right)^2 \quad (14)$$

Due to channel independence, the variance of the total instant power variation δP is

$$\sigma_{\delta P}^2 = N \sigma_{\delta P_i}^2 \quad (15)$$

At the limit $\bar{P} \rightarrow \infty$, the SOA ASE noise is negligible and the OSNR at the output of the SOA is calculated by the formula

$$\text{OSNR} = \frac{P^{\text{out}}}{\sigma_{\delta P_{\text{out}}}} \quad (16)$$

or, equivalently,

$$\text{OSNR} \stackrel{(10)-(15)}{=} \frac{G_0}{\left\{ \left[\frac{G_0(G_0-1)}{G_0 N \bar{P} + P_s} \right]^2 N \bar{P}^2 + G_0^2 \right\}^{\frac{1}{2}} \left(\frac{1-r}{1+r} \right)} \quad (17)$$

where $G_0 = G_0(P_0) = G_0(N\bar{P})$.

It is observed that as $\bar{P} \rightarrow \infty$, the OSNR reaches a ceiling. The same is true for the Q-factor given by [14]

$$Q = \frac{\mu_1 - \mu_0}{\sigma_1 + \sigma_0} \quad (18)$$

where

$$\begin{aligned} \mu_1 &= (G_0 + \delta G)(\bar{P} + \delta P_i) \\ \mu_0 &= (G_0 - \delta G)(\bar{P} - \delta P_i) \\ \sigma_1^2 &= 4(G_0 + \delta G)(\bar{P} + \delta P_i) h\nu \Delta\nu \cdot n_{sp} (G_0 + \delta G - 1) \\ \sigma_0^2 &= 4(G_0 - \delta G)(\bar{P} - \delta P_i) h\nu \Delta\nu \cdot n_{sp} (G_0 - \delta G - 1) \end{aligned} \quad (19)$$

We set

$$G_0 \pm \delta G - 1 \simeq G_0 \pm \delta G \quad (20)$$

which is valid for $G_0 \gg 1$, we use the 2nd order Taylor series to get

$$(\bar{P} \pm \delta P_i)^{\frac{1}{2}} = \bar{P} \pm \frac{1}{2} \delta P_i \quad (21)$$

and we assume that for XGM, the average error probability is equal to the error probability of the innermost traces in the eye diagram corresponding to all ONES and ZEROS [23]. By substitution of (19)-(21) in (18) we get

$$Q \simeq \frac{1}{2\sqrt{a}} \left(\frac{1-r}{1+r} \right) \frac{N\bar{P} + P_s}{G_0(N\bar{P})N\bar{P} + P_s} \quad (22)$$

where $a = h\nu \Delta\nu \cdot n_{sp}$. As $\bar{P} \rightarrow \infty$, Q reaches a ceiling similar to the one observed in Fig. 8 for large-signal modulation which is

$$Q = \frac{1}{2\sqrt{a}} \left(\frac{1-r}{1+r} \right) \quad (23)$$

It is worth noting here that the curves in Fig. 8 do not assume the bell shape shown in [29] (Chapter 14, Fig. 15), i.e., the optical interconnect performance does not degrade at high SOA input power levels but reaches a ceiling. This apparent contradiction is due to the fact that, in our simulations, the input power to the photodiode is unlimited, whereas, in practice, there is a maximum optical power that can be received by the photodiode. In [29], there is an attenuator before the photodiode that keeps the optical power below a certain level.

REFERENCES

- [1] <http://www.top500.org>.
- [2] A. F. Benner, "Cost-effective optics: Enabling the Exascale roadmap," 17th Symposium on High Performance Interconnects, pp. 133–137, New York, NY, Aug. 2009.
- [3] P. Pepeljugoski, J. Kash, F. Doany, D. Kuchta, L. Schares, C. Schow, M. Taubenblatt, B. J. Offrein, and A. Benner, "Low power and high density optical interconnects for future supercomputers," in *Proc. Optical Fiber Communications (OFC/NFOEC)*, San Diego, CA, Mar. 2010, paper OThX2.
- [4] P. Pepeljugoski, J. Kash, F. Doany, D. Kuchta, L. Schares, C. Schow, M. Taubenblatt, B. J. Offrein, and A. Benner, "Towards exaflop servers and supercomputers: The roadmap for lower power and higher density optical interconnects," in *Proc. 36th Eur. Conf. Optical Communication (ECOC)*, Torino, Italy, Mar. 2010, paper Tu4G3.
- [5] A. F. Benner, D. M. Kuchta, P. K. Pepeljugoski, R. A. Budd, G. Hougham, B. V. Fasano, K. Marston, H. Bagheri, E. J. Seminario, H. Xu, D. Meadowcroft, M. H. Fields, L. McColloch, M. Robinson, F. W. Miller, R. Kaneshiro, R. Granger, D. Childers, and E. Childers, "Optics for high-performance servers and supercomputers," in *Proc. Optical Fiber Communications (OFC/NFOEC)*, San Diego, CA, Mar. 2010, paper OTuH1.
- [6] R. Hemenway, R. R. Grzybowski, C. Minkenbergh, and R. Luijten, "Optical packet-switched interconnect for supercomputer applications," *Journal of Optical Networking*, vol. 3, no. 11, pp. 900–913, Dec. 2004.
- [7] I. Roudas, B. R. Hemenway, and R. R. Grzybowski, "Optimization of a supercomputer optical interconnect architecture," in *Proc. IEEE/OSA LEOS Annual Meeting*, Lake Buena Vista, FL, Oct. 2007, paper ThG3.
- [8] F. Karinou, I. Roudas, K. Vlachos, C. S. Petrou, A. Vgenis, and B. R. Hemenway, "Wavelength-space permutation switch with coherent PDM QPSK transmission for supercomputer optical interconnects," in *Proc. Optical Fiber Communications (OFC/NFOEC)*, San Diego, CA, Mar. 2010, paper JWA62.
- [9] F. Karinou, I. Roudas, K. Vlachos, B. R. Hemenway, and R. R. Grzybowski, "Performance assessment of an optimized optical supercomputer interconnect architecture," in *Proc. Optical Fiber Communications (OFC/NFOEC)*, Los Angeles, CA, Mar. 2011, paper JWA086.
- [10] I. Roudas, N. Antoniadis, T. Otani, T. E. Stern and R. E. Wagner, "Accurate modeling of optical multiplexer/demultiplexer concatenation in transparent multiwavelength optical networks," *J. Lightw. Technol.*, vol. 20, no. 6, pp. 218–228, June 2002.

- [11] A. Humblet and M. Azizoglu, "On the bit error rate of lightwave systems with optical amplifiers," *J. Lightw. Technol.*, vol. 9, no. 11, pp. 1576–1582, Nov. 1991.
- [12] VPItransmissionMaker™ Optical Systems User's Manual.
- [13] A. J. Lowery, "Transmission-line modeling of semiconductor laser: The transmission-line laser model," *Int. J. Num. Model.*, vol. 2, pp. 249–265, 1989.
- [14] G. P. Agrawal, *Fiber-Optic Communication Systems*. Wiley, New York (2002).
- [15] Z. Li, Y. Dong, J. Mo, Y. Wang and C. Lu, "1050-km WDM transmission of 8×10.709 Gb/s DPSK signal using cascaded in-line semiconductor optical amplifier," *IEEE Photon. Technol. Lett.*, vol. 16, no. 7, pp. 1760–1762, Jul. 2004.
- [16] P. S. Cho, Y. Achiam, G. Levy-Yurista, M. Margalit, Y. Gross, and J. B. Khurgin, "Investigation of SOA nonlinearities on the amplification of DWDM channels with spectral efficiency up to 2.5 b/s/Hz," *IEEE Photon. Technol. Lett.*, vol. 16, no. 3, pp. 918–920, Mar. 2004.
- [17] M. J. Connelly, "Wideband semiconductor optical amplifier steady-state numerical model," *IEEE J. Quantum Electron.*, vol. 37, no. 3, pp. 439–447, Mar. 2001.
- [18] G. P. Agrawal, "Four-wave mixing and phase conjugation in semiconductor media," *Opt. Lett.*, vol. 12, no. 4, pp. 260–262 Apr. 1987.
- [19] D. F. Geraghty, R. B. Lee, M. Verdiell, M. Ziari, A. Mathur, and K. J. Vahala, "Wavelength conversion for WDM communication systems using four-wave-mixing in semiconductor optical amplifiers," *IEEE J. Select. Topics Quantum Electron.*, vol. 3, no. 5, pp. 1146–1155, Oct. 1997.
- [20] T. Durhuus, B. Mikkelsen, C. Joergensen, S. L. Danielsen, and K. E. Stubkjaer, "All-optical wavelength conversion by semiconductor optical amplifiers," *J. Lightw. Technol.*, vol. 14, no. 6, pp. 942–954, June 1996.
- [21] K. Inoue, "Crosstalk and its power penalty in multichannel transmission due to gain saturation in a semiconductor laser amplifier," *J. Lightw. Technol.*, vol. 7, no. 7, pp. 1118–1124, July 1989.
- [22] P. S. Cho and J. B. Khurgin, "Suppression of cross-gain modulation in SOAs using RZ-DPSK modulation format," *IEEE Photon. Technol. Lett.*, vol. 16, no. 1, pp. 162–164, Jan. 2004.
- [23] R. Ramaswami and P. A. Humblet, "Amplifier induced crosstalk in multichannel optical networks," *J. Lightw. Technol.*, vol. 15, no. 1, pp. 1882–1896, Jan. 2003.
- [24] K. S. Shanmugan, *Digital and Analogue Communication Systems*, Wiley, New York (1980).
- [25] M. Sauer, R. Hemenway, R. Grzybowski, D. Peters, J. Dickens, and R. Karfelt, "A scalable optical interconnect for low-latency cell switching in high-performance computing systems," in *Proc. SPIE Photonics West*, paper 6124–22, San Jose, CA, Jan. 2006.
- [26] M. Jeruchim, "Techniques for estimating the bit error rate in the simulation of digital communication systems," *IEEE J. Sel. Areas. Commun.*, vol. SAC2, pp. 153–170, Jan. 1984.
- [27] J. Proakis, *Digital Communications*, 4th ed. New York: McGraw-Hill, 2000, pp. 177–178.
- [28] G. P. Agrawal and N. A. Olsson, "Self-phase modulation and spectral broadening of optical pulses in semiconductor laser amplifiers," *IEEE J. Quantum Electron.*, vol. 25, no. 11, pp. 2297–2306, Nov. 1989.
- [29] L. H. Spiekman, "Semiconductor optical amplifiers," in *Optical Fiber Telecommunication*, Academic Press, San Diego, CA, 2002, vol. IV-A, ch14, pp. 699–731.

Fotini Karinou was born in Arta, Greece, in 1983. She received her Diploma in Electrical and Computer Engineering, with specialization in Telecommunications and Information Theory, from the University of Patras, Rio, Greece, in 2007. From October 2006 to May 2007, she was an exchange student at the Institute for Quantum Optics and Quantum Information (IQOQI), at Vienna University of Technology, participating in the initial attempts for the realization of a commercial prototype of a quantum key distribution (QKD) system based on polarization entanglement. She has been involved in EU FP6-project SECOQC and EU FP7-project ICT-BONE. She is currently working towards the PhD degree at the University of Patras. Her research interests lie in the fields of optical interconnects and advanced modulation formats.

Ms Karinou has been a member of the Technical Chamber of Greece and an OSA student member since 2009.

Ioannis Roudas, (M'00) received his B.S. in Physics and an M.S. in Electronics and Radio-engineering from the University of Athens, Greece in 1988 and 1990 respectively, and an M.S. and a Ph.D. degree in coherent optical communication systems from the Ecole Nationale Supérieure des Télécommunications, Paris, France in 1991 and 1995, respectively. During 1995-1998 he worked at the Optical Networking Research Department at Bell Communications Research (Bellcore), Red Bank, NJ. During 1999-2002, he was with the Photonic Modeling and Process Engineering Department at Corning Inc., Somerset, NJ. Since 2003, he has been with the Department of Electrical and Computer Engineering at the University of Patras, Greece, as an Associate Professor of Optical Communications. He has taught, as an adjunct professor, at Columbia University, the City University of New York, and the Hellenic Open University. His current research focuses on spectrally-efficient, multilevel modulation formats, adaptive electronic equalization, and optical interconnects.

He is the author or co-author of more than 65 papers in scientific journals and international conferences and holds one patent. He serves as a reviewer for the Journal of Lightwave Technology, IEEE Photonics Technology Letters, IEEE Journal of Quantum Electronics, Optics Express, and Optics Communications. He is a member of the Optical Communications Subcommittee for the IEEE/LEOS Annual meetings in 2009-2011.

Kyriakos G. Vlachos, received his Dipl.-Ing. degree in electrical and computer engineering from the National University of Athens (NTUA), Greece, in 1998 and his Ph.D. in electrical and computer engineering, also from NTUA, in 2001. From 1997 to 2001 he was a Senior Research Associate in the Photonics Communications Research Laboratory (ICCS/NTUA). In 2001, he joined Bell Laboratories, Lucent Technologies, working in the Applied Photonics Group. Since 2003, he had been also a member of the Computer Engineering Laboratory of the Technical University of Delft, and since 2005 he has been a Faculty Member of the Computer Engineering and Informatics Department of the University of Patras, Greece. His research interests are in the areas of architectures and technologies for broadband, high-speed networks. Prof. Vlachos has participated in various research projects funded by the European Commission (IST-STOLAS, IST-PRO3, ESPRIT-DOALL, ephoton/ONE+, ICT-BONE, and ICT-DICONET). He is a member of IEEE and the Technical Chamber of Greece. He is the (co)author of more than 120 journal and conference publications and also holds two international patents.

B. Roe Hemenway completed his doctorate at Stanford University in 1990 on the subject of integrated silicon free-carrier light modulators. Previously he worked at AT&T Bell Labs (Holmdel) on high-speed lasers and detectors. During 1990-1999 he worked on all-optical wavelength routed networks, high power erbium doped fiber amplifiers, optical frequency conversion and spacecraft optical networks at MIT Lincoln Laboratory. Most recently he has been with Corning Incorporated leading network equipment, optical packet switching and more generally optical physics research. His present interests include intellectual property development and low-power high speed optoelectronics for computer interconnects

Richard R. Grzybowski, photograph and biography not available at the time of publication