Performance Assessment of an Optimized Optical Supercomputer Interconnect Architecture

F. Karinou[1], I. Roudas[1], K. Vlachos[1], B. R. Hemenway[2], and R. R. Grzybowski[2]
[1]*Dept. of Electr. & Comp. Eng., Dept. of Comp. Eng. & RACTI, University of Patras, Rio 26500, Greece; E-mail: karinou@ece.upatras.gr*
[2]*Corning Inc., Corning, NY 14831, USA; E-mail: hemenwaybr@corning.com*

Abstract: We investigate the performance of an optimized optical supercomputer interconnect architecture using a minimum number of on-off gates. A 64×64 optical interconnection is demonstrated using 10 Gb/s IM/DD and 2.5 GBd coherent PDM/QPSK optical links.

OCIS codes: (200.4650) Optical interconnects; (200.6715) Switching; (250.5980) Semiconductor optical amplifiers

1. Introduction

Optical interconnects are the most promising future networking solution for High Performance Computing (HPC) systems and data centers [1]. Current state-of-the-art, commercially-available optical interconnects employ vertical-cavity surface-emitting laser (VCSEL) arrays, ribbons of multimode fibers (MMFs), and arrays of ultra-fast, high-efficiency photodiodes to create parallel, high-speed, point-to-point links between HPC system racks, while packet routing is performed electronically. The Optical Shared MemOry Supercomputer Interconnect System (OSMOSIS) project demonstrated a hybrid packet switch architecture that uses electronics for packet processing and storing and optics for switching and transmission [2]. The heart of the OSMOSIS optical interconnect is a two-stage, broadcast-and-select 64x64, optical switch fabric that makes use of semiconductor optical amplifiers (SOAs), acting as on-off gates [2]. In [3], an economically-viable, multi-stage variant of the original OSMOSIS switch fabric architecture was proposed, aiming at a minimization of SOA elements, at the expense of passive wavelength-routing components. In [4], we experimentally studied the performance of the optimized OSMOSIS 64×64 optical switch fabric, in conjunction with polarization division multiplexing (PDM) quadrature phase shift keying (QPSK) modulation and coherent intradyne detection. In this paper, for the first time, we compare the performance of the original 64×64 two-stage, broadcast-and-select, OSMOSIS optical switch fabric to its three-stage optimized counterpart, both by experiment and simulation. We use both 10 Gb/s IM/DD (intensity-modulated direct detection) and 2.5 GBd coherent PDM QPSK serial optical transmission. Simulation and experiment show that the optimized, cost-efficient architecture performs almost as well as the original one, despite the fact that, in the former, optical signals travel through three stages of concatenated SOAs. Moreover, we show that coherent PDM QPSK transmission at 2.5 GBd far outperforms 10 GBd IM/DD links, due to its resilience to SOA nonlinearities. Therefore, we argue that PDM QPSK might be a viable solution for high-spectral-efficiency serial transmission through chains of SOA on-off gates, once coherent optical technology becomes mature enough to be cost-effective for deployment in short-haul applications.

2. Architecture of the optical packet switch fabric

Fig. 1 shows the scheme of the original 64×64 OSMOSIS optical packet switch fabric (green), and its optimized counterpart (green and blue). They differ in two ways: i) the transmitters are organized using different multiplexing hierarchies and number of transmitters; and ii) wavelength selection of the desired channel is performed using a different number of selection stages (two and three, respectively) and SOAs (16 and 12 SOAs/Rx card, respectively).

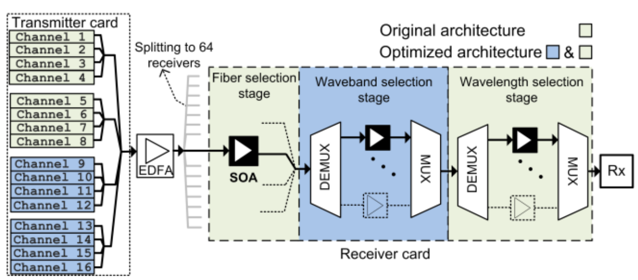

Fig. 1 Block diagram of the original (in green) and the optimized (in green and blue) 64×64 OSMOSIS optical switch fabric.

3. Experimental setup

The laboratory implementation of the optimized 64×64 OSMOSIS optical packet switch fabric is shown in Fig. 2. Due to lack of resources, at the transmitter's side, only nine lasers are used to emulate all 16 transmitted channels. Their wavelength allocation is shown in Fig. 2a. We populate the wavelength slots of the first two wavebands W_1-W_2, 0-3 and 8-11, respectively. We represent the remaining two wavebands W_3-W_4, i.e., λ_{16}-$\lambda_{27,}$ by a single light source placed at $\lambda_{21.5}$, having eight times the nominal power of a single channel (Fig. 2a). Wavelength slots shown in white serve as guard bands, in order to facilitate wavelength/waveband multiplexing/demultiplexing. The first eight

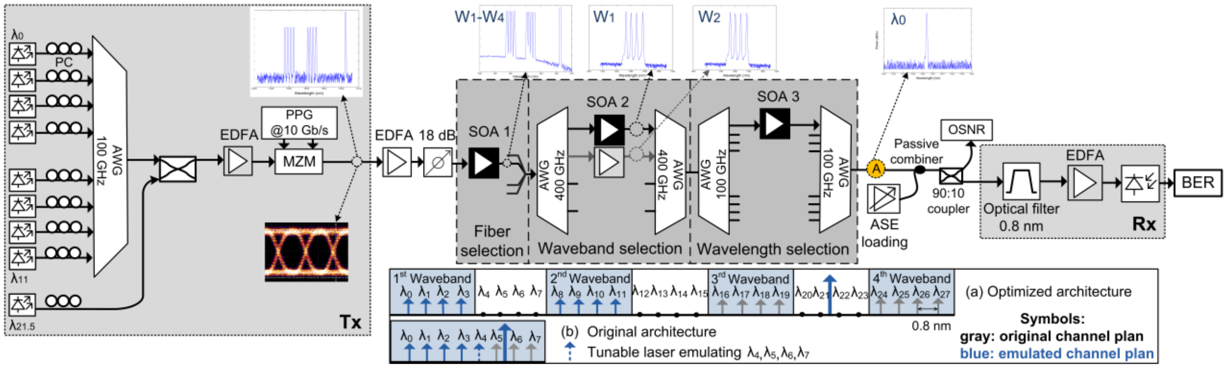

Fig. 2 Experimental set-up for optimized architecture and implemented wavelength plan for (a) optimized and (b) original architecture.

wavelength channels are multiplexed with an arrayed waveguide grating (AWG) and then combined with the $\lambda_{21.5}$ through a coupler.The WDM signal is preamplified by an erbium-doped fiber amplifier (EDFA) and modulated using a Mach-Zehnder modulator (MZM), fed with a 10 Gb/s pseudo-random binary sequence (PRBS), to form eight identical non-return-to-zero Amplitude Shift Keying (NRZ-ASK) WDM signals. The latter is a worst-case scenario for studying the impact of self-gain modulation (SGM) and cross-gain modulation (XGM) due to SOAs. The signal is then boosted using a second EDFA. The 1:64 star coupler of the actual OSMOSIS architecture is emulated by a variable optical attenuator set to 18 dB loss. Subsequently, the signals pass through three consecutive wavelength selection stages that employ SOAs as on-off gates. The SOAs under study have a 3-dB bandwidth of 90 nm, a small-signal gain of 15 dB, a gain peak at 1490 nm, and a low polarization dependent gain (PDG<0.2 dB). The first SOA is used to select one out of four fibers and is followed by a 4:1 combiner. The waveband selection stage consists of a pair of 400-GHz bandwidth AWGs interconnected with four SOAs acting as on-off gates. The selection of the desired channel is performed via a third wavelength selection stage, which is comprised of a pair of conventional AWG MUX/DMUX, with 100 GHz spacing, and a single SOA. The experimental setup used for measuring the performance of the original OSMOSIS optical switch fabric isn't shown here due to space limitations but it is similar to the one described in [2], whereas the wavelength allocation is shown in Fig. 2b. To assess the error probability of the two architectures, amplified spontaneous (ASE) noise is loaded after the last selection stage, in order to adjust the delivered OSNR to the receiver. Then, the selected wavelength is filtered by a 0.8 nm bandwidth optical filter and detected by a p-i-n diode.

4. Simulation and experimental results

We evaluate the bit error rate (BER) as a function of the OSNR for each of the transmitted wavelengths for both interconnect architectures, both by experiment and simulation. We want to evaluate the impact of the ASE accumulation and SOA nonlinear effects. Simulation results in Fig. 3a reveal that there is a spread of approximately 2 dB between the 16 simulated channels of the optimized architecture. Channel 1 and Channel 10 exhibit the best and the worst performance, respectively. Eye diagrams for the wavelengths investigated in the experiment are shown for qualitative comparison. In the experiment, in Fig. 3b, the spread between implemented channels is 1.2 dB, for the 8 channels that are studied. The delivered OSNR is enough to achieve error free operation, although this is not shown in the graph. The eye diagrams for the best and the worst case are also shown. In Fig. 3a, the accuracy of the semi-analytical method used for the calculation of error probability for an IM/DD Rx with realistic filter (Simulation-Ideal case) was compared to the analysis of [5] for an IM/DD receiver with a brickwall optical BPF and an integrate-and-dump LPF. The latter takes into account the non-Gaussian noise statistics at the receiver but neglects any signal distortion due to filter-induced intersymbol interference (Theory-Ideal case). These two cases were proven to be in good agreement as shown in Fig. 3. The OSNR penalty for the simulated best and worst channels compared to the IM/DD receiver of [5] is 1 dB and 3 dB, respectively. The OSNR

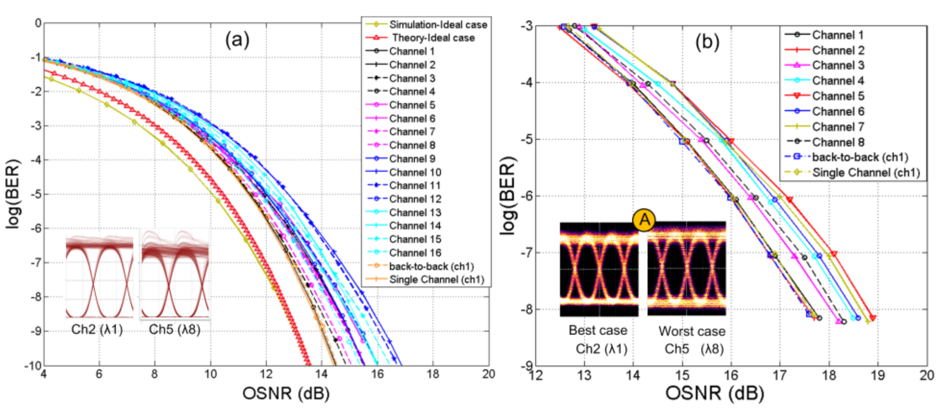

Fig. 3 BER vs OSNR (measured at RB 0.07 nm) for the optimized OSMOSIS architecture: (a) Simulation and (b) Experimental results

JWA86.pdf

penalty for the experimental best and worst channel, compared to the IM/DD receiver of [5], is 5 dB and 6.2 dB, respectively. The experimental results are by 4 dB worse than the ones provided in the simulation. This is attributed to two different mechanisms: i) the much higher extinction ratio at the receiver used in simulation [6] and ii) the narrower optical filtering performed in the simulations. In Fig. 3, the different behavior of the transmitted channels in terms of performance (i.e., spread of the curves in both cases), is attributed to the power variation among the channels, as each one experiences a different gain from the EDFAs and the SOAs and different insertion losses, while passing through the cascaded AWGs. These result both in unequal accumulation of ASE noise for each channel, as well as channel power variation at the receiver. The performance of the architecture is also tested without in-line SOAs (back-to-back case) and in single channel transmission after three stages, as shown in Fig. 3a and Fig. 3b. We observe that there is no penalty between the back-to-back, the single channel transmission, and the best case scenario transmission, for both simulation and experiment. A similar study is also carried out for the original architecture. In Fig. 4, we summarize the results for the best and the worst channels, in terms of performance, for both architectures, both in the simulation and in the experiment. In the simulation, the OSNR penalty compared to the matched-filter receiver for the best channel, in both architectures, is approximately 1 dB. In contrast, for the worst channel, the OSNR penalty is 2 dB in the original interconnect architecture, as opposed to 3 dB for its optimized counterpart. Therefore, we 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. In the experiment, we observe that the optimized architecture is slightly better for both the best and the worst channel. Experimental measurements are in fairly good agreement with the simulation results (Fig. 4), for all 8 wavelengths studied both by simulation and experimentally. Finally, our experiment confirms that the best performance is approximately the same for both architectures, although the optimized configuration is slightly superior to the initial one. For a qualitative comparison, eye diagrams are also shown for each case. Finally, we experimentally study the performance of the optimized optical interconnect for the hypothetical, futuristic scenario when PDM-QPSK and coherent detection may be used. The experimental configuration of the optical PDM/QPSK interconnection was described and studied in [4]. Here, the performance of the optical PDM-QPSK interconnection is improved, compared to the preliminary measurements presented in [4], by using an ECL laser with a 3-dB linewidth of 200 kHz. In Fig. 4, on the right, it is shown that, in the case of IM/DD modulation, the OSNR penalty between the back-to-back case and after the three selection stages is about 1 dB. This is attributed to the nonlinearities, such as XGM, that greatly affect the ASK format, as well as to ASE noise accumulation as the signal passes through the concatenated SOAs. On the contrary, the PDM-QPSK format is more robust and outperforms IM/DD by far, approximately by 7 dB for a BER=10^{-3}, after three selection stages. Representative PDM-QPSK constellation diagrams are also shown for OSNR=10 dB.

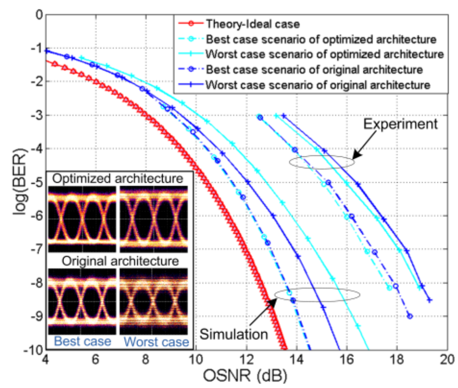

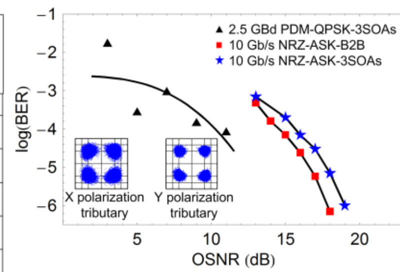

Fig. 4 Left: Performance comparison of the original and the proposed optimized architecture. Right: Experimental comparison of coherent PDM-QPSK and IM/DD for the optimized architecture.

5. Conclusion

In conclusion, we studied for the first time, both theoretically and experimentally, the performance of an optimized optical interconnect architecture using a minimum number of SOA-based on-off gates. Comparison with the original configuration, which was proposed in the context of the OSMOSIS project, revealed that the economically-viable new interconnect configuration performs almost equally well to the more compact, but much more SOA consuming, original one. We also showed that coherent optical PDM-QPSK interconnections largely outperform their conventional IM/DD counterparts, mainly due to their resilience to SOA nonlinearities.

6. References

[1] P. Pepeljugoski et al.,“Data Center and High Performance Computing Interconnects for 100 Gb/s and Beyond”, Proc. OFC'07, OMR4 (2007).
[2] R. Hemenway et al., “Optical packet-switched interconnect for supercomputer applications,” OSA JON **3**, 900-913 (2004).
[3] I. Roudas et al., “Optimization of a supercomputer optical interconnect architecture,” Proc. IEEE LEOS 2007, ThG3 (2007).
[4] F. Karinou et al., “Wavelength-space permutation switch with coherent PDM QPSK transmission for supercomputer optical interconnects”, Proc. OFC'10, JWA62 (2010).
[5] A. Humblet and M. Azizoglu, “On the bit error rate of lightwave systems with optical amplifiers”, IEEE JLT **9**, 1576-1582 (1991).
[6] G. P. Agrawal, *Fiber Optic Communication Systems*, New York, Wiley (2002).