High Capacity Switched Optical Interconnects for Low-Latency Packet Routing

K.A. Williams[1]

[1]Eindhoven University of Technology
Inter-University Research Institute COBRA on Communication Technology
Department of Electrical Engineering,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands.

High-capacity switched interconnects based on semiconductor optical amplifier switches are studied for low-latency data networking applications. Capacity is scaled through the use of wavelength multiplexed packets to show a sublinear dependence of power penalty on data capacity and sufficient power margin for scaling to 8x8 port interconnection at 100Gb/s aggregate capacity per port. Good agreement with test-bed data is achieved and a route to switched terabit/second capacity optical interconnection circuits is identified.

Introduction

Advances in high performance signal processing have been predominantly driven by improved microprocessor speeds and transistor miniaturization. However the rapid increase in the power consumption within these circuits and the delays in transmission between processors and memory incur formidable bottlenecks. Switched electronic interconnects are approaching their limits in terms of tolerable size, power consumption, crosstalk and thermal dissipation and this is leading to renewed interest in photonic techniques. However power consumption and complexity are paramount in the deployment of high-capacity photonic solutions.

The parallel transmission of multiple wavelengths has been proposed in optical interconnection [1] and exploited more recently in packet-switched optical interconnects [2,3] with a view to facilitating order of magnitude reductions in power consumption and component count. While aggregate data capacity of order 100Gb/s is achieved, there is evidence of signal impairment and little simulation study has been performed to quantify and explore the origins of power penalty or assess the potential limitations in terms of ultimate capacity and connectivity. This simulation study explores data capacity limits and connectivity for broadcast and select switch networks. Scaling in aggregate capacity is quantified in terms of power penalty for varied numbers of wavelength channels. The input power dynamic range is explored to facilitate power map assessment for increased numbers of splitter stages in larger interconnection networks. Cross-comparison is made with test bed data using commercially available components.

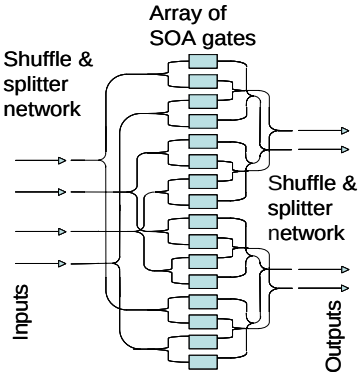

Figure 1: Example schematic of a 4 x 4 broadcast and select switch fabric

Simulation

The interconnect architecture studied is a broadcast and select fabric with cascaded splitters which fan out to an array of semiconductor optical amplifiers. Outputs are subsequently recombined at the output side. Figure 1 shows an example layout for a four input and four output (4x4) switch architecture with shuffle networks on both the input and output side.

Wavelength multiplexed sources are implemented with up to ten wavelength channels, and are modulated with 10Gb/s pseudo random bit streams with a simulation-limited pattern length of 2^7-1 and an extinction ratio of 10dB. The data streams are decorrelated in time, and spaced with a 100GHz channel separation. Multiplexing is performed in the Fourier domain before defining the time resolved input field to the SOA model. The simulated SOA output is subsequently filtered using the amplitude response of a commercial demultiplexer. Signal integrity is assessed initially in terms of Q factor before deriving power penalty.

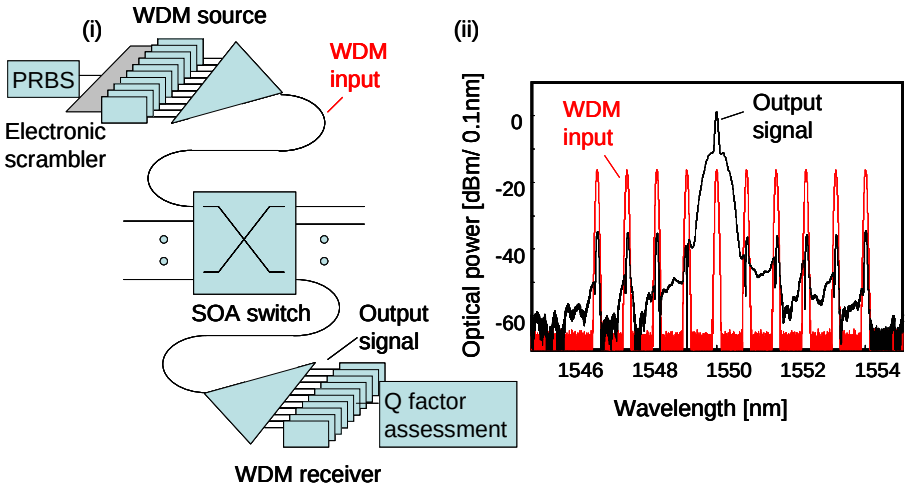

Figure 2: Assessment of the capacity of an SOA based interconnect and inset right, the simulated input spectrum and received spectrum for 100Gb/s aggregate operation

Received data for the central 1550nm wavelength channel is presented in the right inset of figure 2 along with the wavelength multiplexed input. The output time resolved field is assessed with a receiver model [4] using an analysis of the eye diagram, and estimations are made for the Q-factor from the probability distribution functions for the ones and zero level voltages at the receiver. Bit error rate is subsequently estimated as a function of receiver sensitivity. Figure 3 shows the bit error rate performance of the switch for 100 Gb/s aggregate data transmission on ten wavelength channels.

Bit error rate performance is estimated initially for the input signal to the switch and subsequently for the amplified signals outputted from the switch fabric. A reduction in the receiver sensitivity is noted for the case of -25dB into the SOA, and further degradation to a total power penalty of 1dB is observed for the case of -20dBm input power. An increased curvature in the bit error rate dependency on received power is also observed as the error performance deviates away from noise limited sensitivity and power dependent distortion in the amplifier becomes more important. This is also

evident in the eye diagram in the right inset (ii), where a resulting broadening in the one level rail is attributable to insufficiently fast gain dynamics.

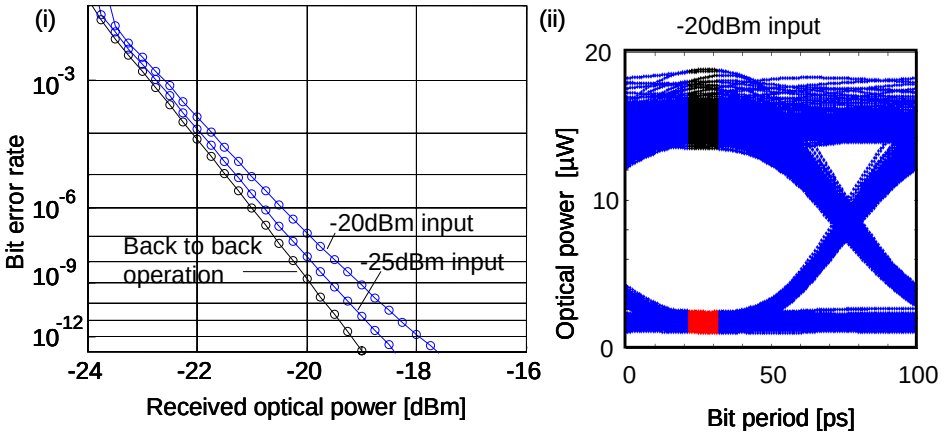

Figure 3: Simulated bit error rate characteristics for the SOA switch element operating at based interconnect and inset right, the simulated input spectrum and received spectrum for 100Gb/s aggregate operation

Capacity scaling

To explore the role of the data capacity on switch performance, the number of wavelength channels and the input power per wavelength channel into the switch is varied and the power penalty is extracted from Q-factor estimation. Figure 4 summarizes the power penalty and gain for the broader range of operating conditions.

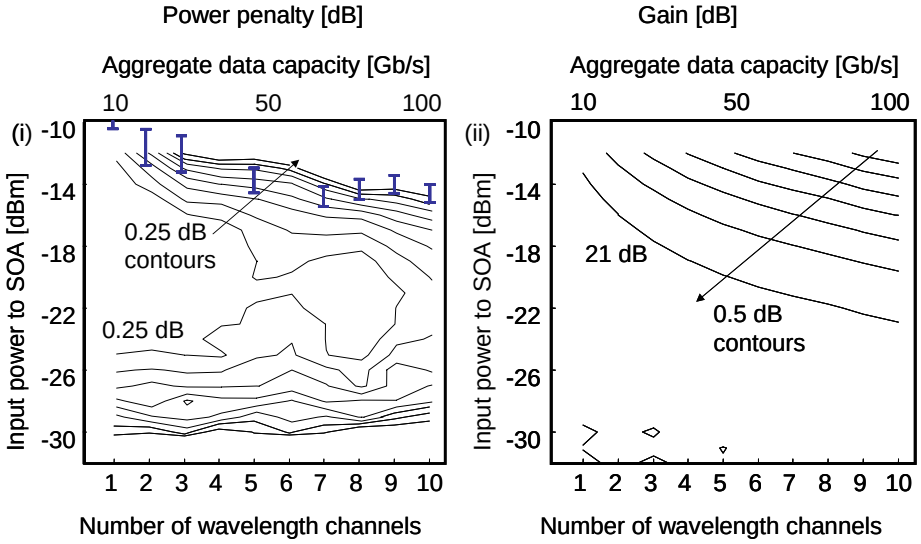

Figure 4: (i) Power penalty at bit error rate 10^{-9} and (ii) gain as a function of the number of wavelength channels and the input power per channel into the SOA switch. Also included in inset (ii) are experimental data points taken from [2] to show the experimentally measured maximum operating input power

The contour plots in figure 4(i) show the power penalty extracted from Q factor assessments as a function of both number of input wavelengths and for a range of input powers into the switch. The power penalty data indicate that power penalties of below 0.25dB are readily achieved for low numbers of wavelength channels, but as the

capacity is increased, the power range for which this is feasible drops. The right inset figure shows the gain for the measured conditions and therefore shows the correlation between power penalty and gain saturation at high operating powers. The degradation in power penalty at low power is attributable to poorer signal to noise ratio. Experimental data taken from [2] are also included in figure 4 to show good correlation with measured test-bed data. The maximum tolerated input power is observed to increase sublinearly with the number of wavelengths, indicating that further scaling in capacity is still feasible.

Connectivity scaling

The interconnection of the switch fabric presented in figure 1 and 2 indicates that the power map is determined primarily by the tolerated splitter and multiplexer losses. Assuming fixed losses of 7dB for the wavelength multiplexer and splices, and splitter losses of 3.4 $\log_2 N_p$ dB for a switch fabric of N_p input/ output pairs allows an estimate to be made for the connectivity scaling for a 100Gb/s/port switch fabric. A margin for fabrication tolerance and variations in system implementation may also be included. The data presented in figure 3 and 4 indicate a maximum input power of -18dBm into the switch, a receiver sensitivity of -20dBm and 18dB SOA gain with 1dB power penalty. The power map for these values indicates a margin of order 1dB even for an 8 x 8 architecture with each input potentially operating at 100Gb/s.

Conclusions

The scaling performance of broadcast and select switch fabrics is studied to identify the limits to performance in terms of capacity and connectivity scaling. The trends are in close agreement with test-bed data, indicating that performance is primarily limited by the design of the SOA switch implemented. Performance enhancement may be anticipated for SOAs which are further optimized for the switch architecture. However 100Gb/s/port is already achieved for an architecture which is scaleable to 8x8 interconnection, indicating a route to Terabit/second aggregate routing for a moderately sized photonic interconnect circuit.

Acknowledgements

The support a Marie Curie Chair mobility award is acknowledged. Prof Ian White at the Cambridge University Engineering Department and Dr Madeleine Glick at Intel Research Cambridge are thanked for stimulating discussions.

References

[1] M.L. Loeb and G.R. Stilwell, "High-speed data transmission on an optical fiber using a Byte-wide WDM system", IEEE Journal of Lightwave Technology, 6, 1306-1311, 1988.

[2] T. Lin, K. A. Williams, R.V. Penty, I. H. White, M. Glick and D. McAuley, "10x10Gb/s packet routing with single stage SOA switch for multi-port interconnects", IEEE Photonics Technology Letters, 18, 691-694, 2006

[3] K.A. Williams, G.F. Roberts, T. Lin, R.V. Penty, I.H. White, M. Glick and D. McAuley, "Monolithic 2x2 optical switch for wavelength multiplexed interconnects", IEEE Journal of Selected Topics in Quantum Electronics, 11, 78-85, 2005

[4] S. Gee, H. Lage, C. Park, K.A. Williams, R.V. Penty, I.H. White, and J. Barnard, "Modelled optimisation and experimental verification of a low dispersion source for long haul 2.488Gb/s systems", Hewlett Packard Journal, 48, 5, 91-106, 1997