

## Optimization of a supercomputer optical interconnect architecture

I. Roudas (1, 2), B. R. Hemenway (2), and R. R. Grzybowski (2)

1 : Dept. of Electrical & Computer Engineering, University of Patras, Rio 26500, Greece; roudas@ece.upatras.gr

2 : Corning Inc., Corning, NY 14831, USA; {hemenwaybr, grzybowski}@corning.com

**Abstract** We present a novel, optimal architecture for a  $N \times N$  supercomputer optical interconnect, composed of arrayed waveguide grating multiplexers/demultiplexers and a minimum total number, of the order of  $N \ln N$ , of semiconductor optical amplifiers used as on-off gates.

### Introduction

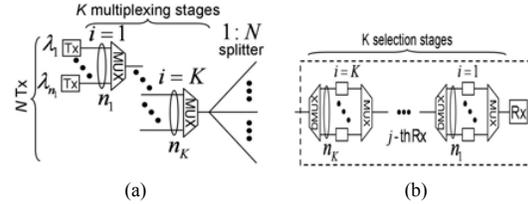
The purpose of a joint research project between Corning Inc. and IBM, called Optical Shared MemOry Supercomputer Interconnect System (OSMOSIS) [1], is to design and demonstrate a scalable, low-latency, high-throughput, optical switch suitable for future supercomputers. The viability of such supercomputer optical interconnect architectures depends critically on the drastic reduction of the cost per input/output port. This study examines how to reduce the cost of the OSMOSIS optical interconnect by optimizing the number of active components.

### OSMOSIS interconnect architecture

In the OSMOSIS project, a broadcast-and-select interconnect architecture is used to establish pairwise links between  $N$  fixed wavelength transmitters and  $N$  discretely tunable receivers in a functionally equivalent manner to a strictly non-blocking  $N \times N$  permutation switch. Channel selection at the tunable receivers can be performed by one [2] or two [3] stages of multiplexer/demultiplexer (MUX/DMUX) pairs and semiconductor optical amplifiers (SOAs) acting as on-off gates. The two-stage configuration proposed in [3] takes advantage of the periodic nature of the arrayed waveguide grating (AWG) MUX/DMUX transfer function in order to reduce the number of SOAs. A natural extension of this idea is to further increase the number of stages of the channel selector until the number of required SOAs is reduced to a minimum. We formulate this as a constrained minimization problem, where the cost function to minimize is the number of SOAs, the constraint is the number of channels, and the (discrete) variables to optimize are the number of selection stages and the number of tributaries per stage. We derive an approximate analytical solution, using the method of Lagrange multipliers. The analytical results are verified by exhaustive search of all possible realizations of the  $256 \times 256$  OSMOSIS switch fabric and by direct enumeration of the number of SOAs required in each realization.

The block diagram of the OSMOSIS architecture is shown in Fig. 1. The transmitted channels are combined using  $K$  successive wavelength division multiplexing stages, each grouping  $n_i$  tributaries (i.e., wavelengths, wavebands of different orders). All channels are broadcasted to all discretely tunable

receivers using a  $1:N$  splitter. The layout of a receiver card is shown in Fig. 1(b). It comprises  $K$  selection stages numbered in reverse order, from left to right. The AWG MUX/DMUXs of each stage  $i$  have free spectral ranges  $\Delta \nu_i = n_i \Delta \nu_{i-1}$ , with  $\Delta \nu_0 = \Delta f$ , where  $\Delta f$  is the channel spacing. By activating one SOA per stage, a single wavelength is obtained at the receiver. This selection procedure is similar to finding a path from the root to the leaves of a tree, by choosing, at each stage, one out of several branches of the same level until, finally, a leaf is reached.



**Fig. 1** OSMOSIS optical interconnect architecture (a) Transmitter multiplexing hierarchy; (b) Multistage discretely tunable receiver (Symbol: Square=SOA).

### Analytical optimization

We will show that the minimum number of SOAs, for a given number of channel selection stages  $K$ , is achieved when the number of tributaries selected in all stages is equal. If such *equipartition* exists, then the number of SOAs per receiver card is  $K \sqrt[K]{N}$  and the total number of SOAs required for the OSMOSIS architecture is  $NK \sqrt[K]{N}$ . In addition, the optimal number of channel selection stages is approximately  $\ln N$  and the minimum total number of SOAs is  $eN \ln N$ , where  $e$  is Napier's constant.

The number of SOAs per receiver card is

$$\Omega = \sum_{i=1}^K n_i \quad (1)$$

The total number of channels  $N$  is factored as

$$N = \prod_{i=1}^K n_i \quad (2)$$

We want to minimize (1) based on the constraint (2), given a fixed number of channel selection stages. Setting the total differentials of (1), (2) to zero yields:

$$\begin{aligned} d\Omega = 0 &\Rightarrow \sum_{i=1}^K dn_i = 0 \\ dN = 0 &\Rightarrow \sum_{i=1}^K (dn_i / n_i) = 0 \end{aligned} \quad (3)$$

Although  $n_i$  are positive integers, in (3) we assume they take values in a continuous range. Despite this

approximation, it will be shown, by example, that the results hold in the discrete case as well, at least for the order of magnitude of interconnect dimensions of interest in the OSMOSIS research project.

We combine (3) using the Lagrange multiplier  $\lambda$  [4]

$$\sum_{i=1}^K dn_i - \lambda \sum_{i=1}^K (dn_i / n_i) = 0 \quad (4)$$

From (4) we obtain  $n_i = \lambda$  for  $i = 1, \dots, K$ . The value of the Lagrange multiplier can then be calculated by substituting into (2)

$$N = \prod_{i=1}^K n_i = \lambda^K \Leftrightarrow \lambda = \sqrt[K]{N} \quad (5)$$

Thus, for  $K$  channel selection stages, the minimum number of SOAs per receiver card is  $\Omega_{\min|K} = K \sqrt[K]{N}$  and the minimum total number of SOAs required for the  $N \times N$  interconnection is  $\Omega_{\text{tot}|K}^{\min} = NK \sqrt[K]{N}$ .

The above results are derived assuming a given number of channel selection stages  $K$ . It is instructive to optimize the number of channel selection stages in order to achieve a global minimum of the number of SOAs. An approximate value can be obtained assuming  $K$  lies in a continuous range and taking the derivative of  $\Omega_{\min|K}$  with respect to  $K$ . It is straightforward to show that  $K_{\text{opt}} = \ln N$ , which results in a global minimum number of SOAs per receiver card  $\Omega_{\min} = e \ln N$  and a minimum total number of SOAs equal to  $eN \ln N$ . The latter formulae should be considered as lower bounds to the number of SOAs required per receiver card and per interconnect, respectively.

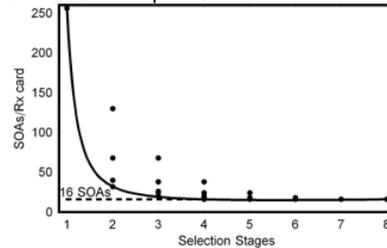
To compare the merit of receiver architectures with different numbers of input/output ports  $N$ , we introduce an *optimality factor* defined as  $F = \Omega_{\min} / \Omega$ , where  $\Omega_{\min}, \Omega$  are the theoretically global minimum (albeit non-attainable, due to the discreteness of  $n_i, K$ ) and the actual number of SOAs per receiver card given by direct enumeration, respectively. Obviously,  $0 < F < 1$ .

### Numerical analysis

To validate the analytical model, in Fig. 2 we compare the number of SOAs in all possible realizations of the receiver cards used in the 256×256 OSMOSIS switch fabric given by direct enumeration (points) and by the theoretical formula for  $\Omega_{\min|K}$  (solid curve). The analytical model is accurate when equipartition of tributaries occurs ( $K = 1, 2, 4, 8$ ). It is worth noting that the minimum number of 16 SOAs is achieved for equipartitions ( $K = 4, 8$ ) and quasi-equipartitions of tributaries ( $K = 5-7$ ), since decomposition of a single selection stage of four tributaries into two selection stages of two tributaries each does not reduce the number of SOAs. The theoretical optima are  $K_{\text{opt}} = 5.5$ ,  $\Omega_{\min} = 15.1$ , respectively.

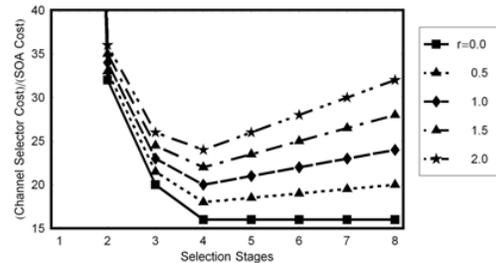
The optimality factor is used to compare the merit of the most economical realizations of OSMOSIS optical interconnects with dimensions up to 2000×2000. We

show that it is possible to approach the optimum when the number of channels can be factored into a product of small primes, e.g., architectures reach within 0.5% of the optimum for  $N = 3^m$  and within 6% of the optimum for  $N = 2^m$ . For  $N = 16, 64, 256, 1024$  the tributaries can be organized into tetrads and this jointly minimizes the numbers of SOAs and the required MUX/DMUX pairs.



**Fig. 2** Number of SOAs required per receiver card for 256 channels partitioned into tributaries of  $K$  selection stages. (Symbols: Curve: theoretical model, Points: enumeration).

This is better illustrated in Fig. 3 for the 256×256 OSMOSIS switch fabric. The cost function (1) is modified to take into account the number of MUX/DMUX pairs and the relative cost  $r = A/B$  between a MUX/DMUX pair  $A$  and an SOA  $B$ . The receiver's channel selector cost is obtained by direct enumeration and is normalized to the per unit SOA cost  $B$ . For  $r = 0$ , the MUX/DMUX pairs have negligible cost and the points coincide with the most cost-efficient realizations shown in Fig. 2. For  $r > 0$ , it is observed that the most cost-efficient 256-channel tunable receiver should have four selection stages.



**Fig. 3** Normalized comprehensive cost function (per unit SOA cost) for  $N = 256$  using direct enumeration.

In conclusion, the present theoretical study minimizes the cost of the OSMOSIS optical interconnect by exploiting the periodicity of the transfer functions of the AWG MUX/DMUXs, in order to reduce the number of SOAs. The feasibility of the proposed architectures ultimately depends on the transmission impairments induced by the optical components.

### References

- [1] R. Hemenway et al., *J. Opt. Netw.*, pp. 900–913, 2004.
- [2] M. Zirngibl et al., *Phot. Technol. Lett.*, pp. 513–515, 1994.
- [3] N. Kikuchi et al., *Electron. Lett.*, pp. 312–314, 2003.
- [4] G. Arfken, *Mathematical Methods for Physicists*, 3d ed., Orlando, FL: Academic Press, pp. 945–950, 1985.