Integrated multi-wavelength routing circuits for high capacity switched interconnects

K.A. Williams

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

Increasingly high capacity data transfer in storage area networking, high performance computing and server networks is placing extremely demanding performance requirements on next generation interconnection technology. Power consumption, physical size and latency in switch fabrics are increasing rapidly with data capacity and connectivity. Networks of tractable numbers of discrete SOAs and filters have now been studied in a number of testbeds showing scaling to Terabit/second capacity utilising wavelength striping. In this paper we study integrated circuit architectures to realise power efficiency in ultra-compact integrated routing circuits. This is expected to provide a key enabler for port-scalable interconnection. Exploiting the parallel processing of wavelength striped data in concert with new configurations of integrated photonic switching circuit, exploiting arrayed waveguide gratings and semiconductor optical amplifier technology, can lead to order of magnitude power efficiencies and capacity scaling for a given optical link latency.

Multistage switch architectures

Switching data at the photonic layer offers considerable conceptual advantages. The routing of high-capacity data may be controlled with a low-speed electronic control plane, avoiding complex, parallel, broadband electronic circuitry. Wavelength channels may be routed, split and combined in parallel without the intrinsic need for electrical power. Data can be transferred between wavelength channels without incurring delay and with low optical power penalty. However, power consumption in space and wavelength switches can be high. Customized photonic circuits and sophisticated specialized components can lead to complex physical layers requiring a high control overhead, complex packaging and large physical footprint. The scaling in terms of both connectivity and capacity remains unclear. Photonic integration is commonly cited as a route to addressing these short-comings. The reduced coupling losses, lower numbers of cooler circuits, the co-location of multiple active elements and reduced numbers of fiber pigtails can lead to reductions in power consumption, packaging materials, delays in data transmission and control complexity. Additionally, the parallel processing of multiple wavelength channels may enable further reductions in power and system complexity. Implementing large scale, high interconnectivity switches on a single photonic substrate leads to complex trade offs between packaging complexity, thermal management, electrical crosstalk and component yield. Here the scaling of interconnect networks based on multistage networks using 2×2 and 4×4 building blocks is explored, enabling a comparison of power consumption, and packaging complexity for comparable physical layer performance.

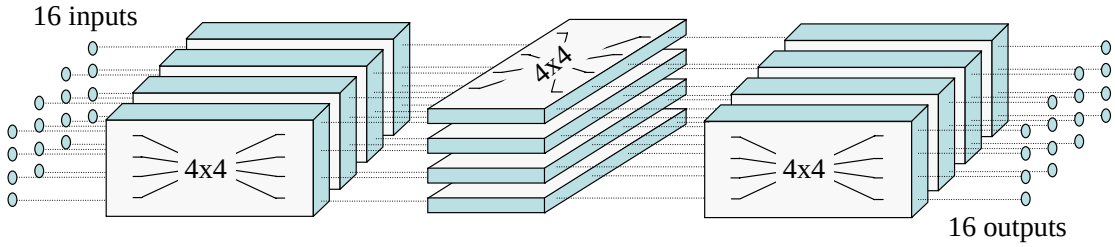

Figure 1: An example three stage switch realising 16 input and outputs.

Figure 1 shows an example implementation of a multistage switch network enabling high connectivity through the parallel interconnection of smaller more manageable building blocks. While the design rules for equivalent electronic networks are well developed, the impact of integrated SOA based circuits as the switch building blocks with their associated physical constraints is less well understood. This work explores two sizes of monolithic building block in the context of a 16 × 16 switching network for high capacity wavelength multiplexed routing.

Simulation of SOA based switch network

The SOA switch elements are simulated through a custom traveling wave code. A flat gain spectrum is implemented for the 1THz signal bandwidth so as to focus on impairments due to noise and distortion. Ten 10Gb/s channels multiplexed on a 100GHz grid enable an aggregate Terabit/s routing with the 16 × 16 configuration. The channels are offset in time and wavelength to ensure decorrelation. The simulated system is shown schematically in figure 2. Simulations are performed iteratively within a number of parameter loops as follows:

- Losses are defined at the input and output of the 2 × 2 and 4 × 4 switches
- Current is swept from 20 through to 150mA
- Input power per wavelength channel is swept from – 20dBm through to 0dBm

Prior to implementing the loop, the baseline performance of the thermal noise limited receiver model is evaluated and decorrelated wavelength division multiplexed input fields are generated. Within the innermost simulator loop, the following sequence is implemented:

- Input the WDM field sequentially into a recirculating loop
- Attenuate as per switch architecture (3.5dB/splitter and 2dB/optical interface)
- Solve the travelling wave rate equations for the SOA for a 2^7-1 bit sequence
- Attenuate as per switch architecture (3.5dB/splitter and 2dB/optical interface)
- Demultiplex the output field for the wavelength channel under test
- Evaluate the power penalty from Q factor analysis for the receiver model output
- Record the power penalty, gain and optical powers
- Define the next input from the current output for the next loop iteration

No wavelength filtering is done within the switch fabric itself and therefore within the loop and this is observed to lead to a degradation in the optical signal to noise performance. The modelled SOA is of a highly conservative design, with 4 quantum wells, 1.3% mode overlap per well, and a length of 800μm, and therefore distortion is also evident at the higher input powers.

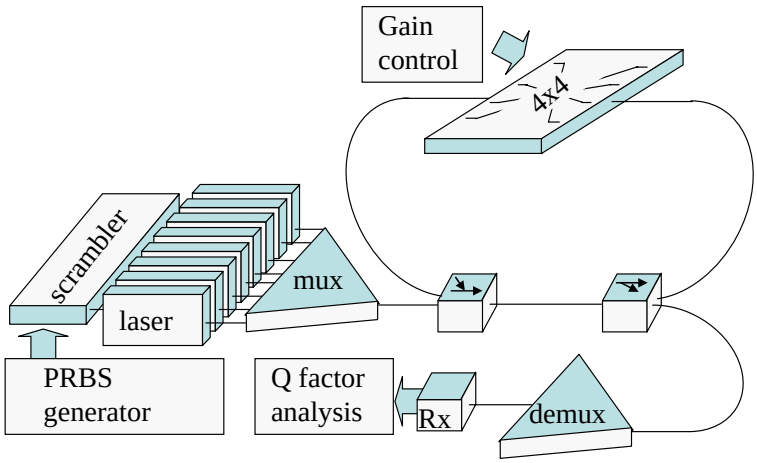

Figure 2: Schematic layout for the simulated physical layer, whereby the cascaded performance of a 4 × 4 switch block is simulated for wavelength multiplexed payloads.

The model implementation is shown schematically in figure 2 in terms of a recirculating signal through the same switch model operating under otherwise identical conditions. Two monolithic switch elements are simulated with 2 × 2 input outputs and 4 × 4 input outputs in configurations to enable 16 × 16 interconnection, and more. For ease of comparison, the power penalty is plotted in figure 3 as a function of loop number for the two architectures. A reduction in power penalty is observed for a given loop iteration for the smaller switch element due to the lower required operating current of 55mA. This contrasts to the value of 110mA used for the 4 × 4 element.

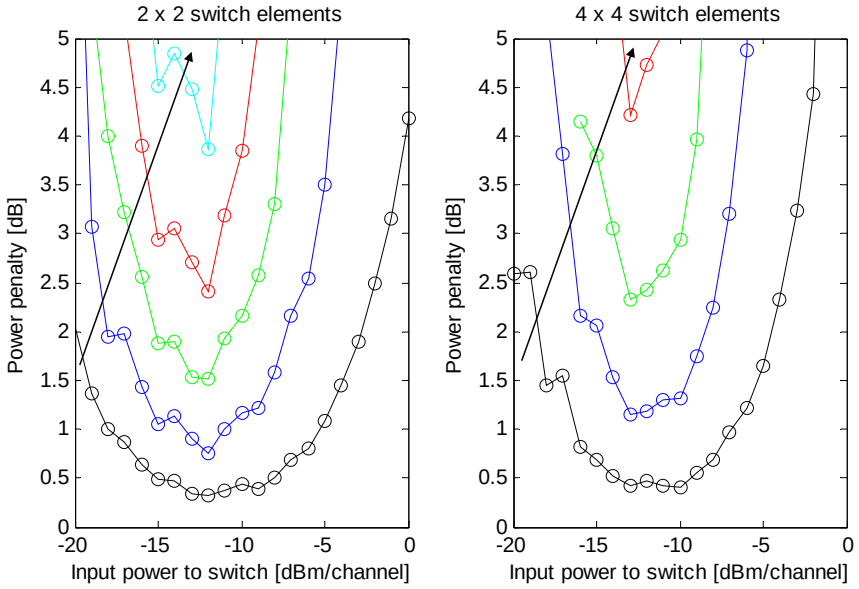

Figure 3: Power penalty as a function of input power for switch fabrics utilising 2 × 2 (left) and 4 × 4 (right) monolithic elements. The data sets indicate degrading power penalty as increasing numbers of loops are simulated.

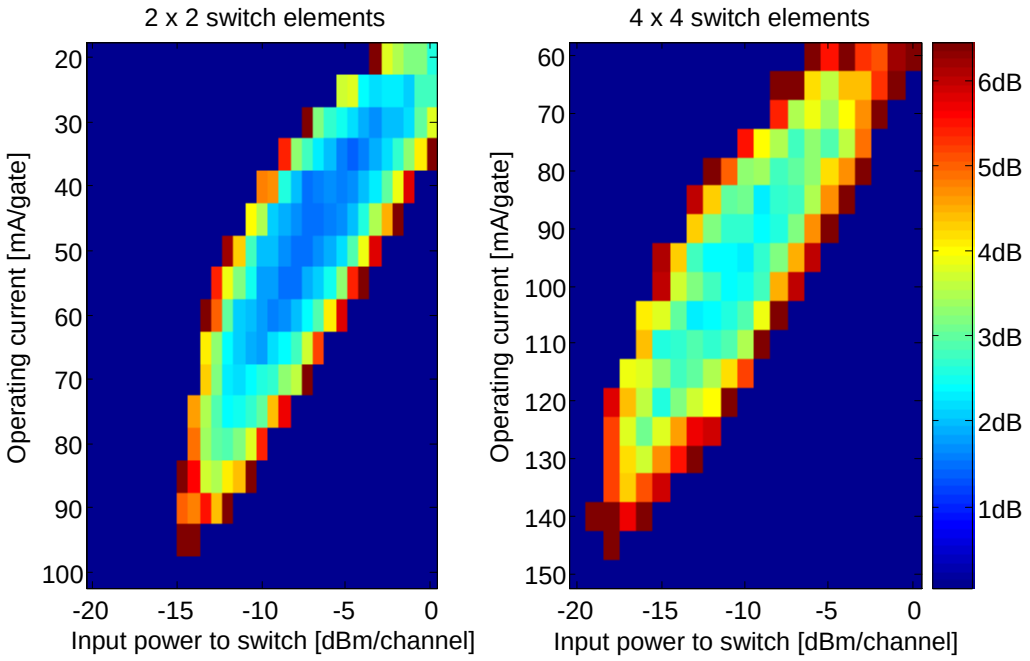

Figure 4: Colourmap images for power penalty after three stages indicating the operating ranges for 2 × 2 and 4 × 4 based switch elements as a function of both input current and input optical power.

Signal degradation is observed for both low input powers and high input powers due to amplified spontaneous emission and optically induced distortion respectively. Four cascades of 2 × 2 switches are readily achieved while only three 4 × 4 switches are feasible for the same power penalty. In both cases a 16 × 16 interconnection network would be feasible using either thirty two 2 × 2 switches or twelve 4 × 4 switches, amongst other possible combinations. While the input power dynamic ranges are comparable, the 4 × 4 implementation requires 75% of the optical interfaces required for the 2 × 2 implementation and less than 40% of the number of monolithic circuits. These parameters having a direct impact on the packaging complexity which would be expected to dominate implementation costs. While the number of on-state gates is also reduced to 75% for the 4 × 4 implementation, the operating current per gate is doubled, and this leads to an increase in power consumption of order 50% to approximately 0.15W/gate. Considering an aggregate switch capacity of 16 inputs × 10λ × 10Gb/s, this corresponds to a power consumption of 2.8mW/Gb/s and 4.5mW/Gb/s respectively, excluding cooling and driver electronics. A trade off between packaging complexity and power consumption for comparable physical layer performance is therefore evident when addressing the optimum levels of monolithic complexity in photonic switch circuits.

Acknowlegdements

The European Commission is acknowledged for personal financial support through the Marie Curie Chair programme. The work is also performed within the framework of the ePhoton/ONe Network of Excellence.