160Gb/s serial line rates in a monolithic optoelectronic multistage interconnection network

160Gb/s Serial Line Rates in a Monolithic Optoelectronic Multistage Interconnection Network

A. Albores-Mejia, K.A. Williams, F. Gomez-Agis, S. Zhang, H.J.S. Dorren, X.J.M. Leijtens, T. de Vries, Y.S. Oei, M.J.R. Heck, L.M. Augustin, R. Notzel, D.J. Robbins, M.K. Smit

Abstract— **We demonstrate very high line rate serial 160Gb/s data transmission through a semiconductor optical amplifier based multistage switching matrix. This represents both the leading edge in monolithic switching circuit complexity and the highest reported line rates through monolithically cascaded switching networks. Bit error rate studies are performed to show only modest levels of signal degradation. Power penalties of order 0.6dB and 1.2dB are observed for two stages and four stages respectively in the monolithic circuits at 160Gb/s per path.**

Index Terms— **Integrated optoelectronics, optical switches**

I. INTRODUCTION

OPTOELECTRONIC switching circuits are able to route increasingly large bandwidth data packets with bandwidth independent energy consumption and nano-second time-scale reconfigurability [1]. Large scale integration has been highlighted as a route to enable scaling to the high levels of connectivity for emerging computing applications and future interconnection networks. However stringent reductions in energy consumption, management complexity and lightspeed hardware latency are required.

The potential for energy efficiency in optoelectronics is must robustly demonstrated at the very high line rates. Recent research has highlighted near bit rate independent energy consumption as capacity is scaled through wavelength multiplexing [4][5]. Dense wavelength division multiplexing readily exploits the multi-Terahertz bandwidths of photonics, but can also lead to stringent management overheads and increased transceiver complexity. Robust wavelength registration throughout the optical network has become mandatory throughout the optical telecommunications network, but remains unattractive in cost- and energy-sensitive short reach data networking. There is therefore considerable merit in using serially multiplexed data formats. One, or a small number of widely spaced wavelength channels on non-critical wavelength grids may lead to a reduced inventory, systems miniaturisation, and management simplification. Single wavelength transmission formats remain the preferred strategy even for 100Gb/s Ethernet.

Serial transmission with relaxed wavelength registration leads to a new set of circuit level challenges and opportunities for optoelectronic switching networks. The possibility to use cooler-free circuits enables both reduced energy consumption and closer integration with high temperature electronics. Ultrabroadband interconnect networks enable capacity agnostic switching and future proofing to new transceiver technologies, which are currently undergoing considerable development.

Semiconductor optical amplifier technology is ideally tailored to offer wavelength and temperature agnostic operation [3]. Furthermore, compliance with proven integration technology offers a route to larger scale integrated circuits [4]. The use of high efficacy III-V materials may ultimately allow additional transceiver integration for optical network within chip architectures. Recent work into discrete [5-7] and monolithic [8-9] multi-stage switching circuits using semiconductor optical amplifiers is indicating routes to tens of connections, each operating at up to 100Gb/s aggregate with nanosecond time-scale reconfigurability. Capacity scaling is achieved primarily through wavelength multiplexing [7,9-10]. Higher line rate studies have been performed in the non-linear regime to realise all-optical logic, sophisticated demultiplexing functions and wavelength conversion [11-13]. However such circuits exploit different physical effects, and require auxilliary optoelectronic and optical circuitry which do not lend themselves well to large scale integration today. There is therefore currently little understanding into the scalability of monolithic multistage interconnection networks operating at high line rates.

In this work we design, fabricate and test a monolithic optoelectronic interconnection network with minimum two and maximum four stages to explore the feasibility of high serial line rate data transmission in multi-stage networks. This is believed to be the first such implementation of a multistage interconnection network on an active-passive regrown epitaxy. The performance is studied in terms of power penalty for increasing line rate and increasing number of switching stages to better understand the physical limitations in ultrahigh line rate optoelectronic integrated circuits.

1550-4794/09 $26.00 © 2009 IEEE
DOI 10.1109/HOTI.2009.10

II. Optoelectronic Circuit Implementation

The four input, four output multistage optoelectronic switching circuit is implemented on an active-passive regrown InGaAsP/InP epitaxial wafer. The circuit is described separately in terms of the fabrication approach, the cell structure used, and the interconnection circuit.

A. *Fabrication*

The circuit is implemented on a multi-project wafer (MPW) enabled by the JePPIX platform [14]. The epitaxial material uses an InGaAsP/InP active layer designed for 1550nm operation. A photolithographically defined etch and regrowth provides active islands periodically across the wafer with low loss passive heterostructures in the rest of the topology. The circuits are processed with a two step reactive ion etch to define both shallow low-loss waveguides and deeper waveguides for selected optical components and functions. Planarisation and gold evaporation and plating complete the processing. Circuits are subsequently cleaved to die. No facet coating has been applied to the circuit prior to this study leading to residual reflectivities of order 34%. The circuit reported in this work includes twelve active islands on a pitch of 250 microns, within an area of approximately 15mm^2. Figure 1 shows a microscope photograph of the circuit. The cleaved facet which includes all the input and output waveguides is visible at the lower boundary of the circuit. Features on the front facet are a reflection of the underlying surface. Some of the passive waveguides are discernable in the top darker region while the active elements are positioned underneath the twelve electrodes in the centre of the circuit. The circuit is bonded to an AlN carrier. The carrier is gold plated and patterned for ease of wire bonding and electronic connection.

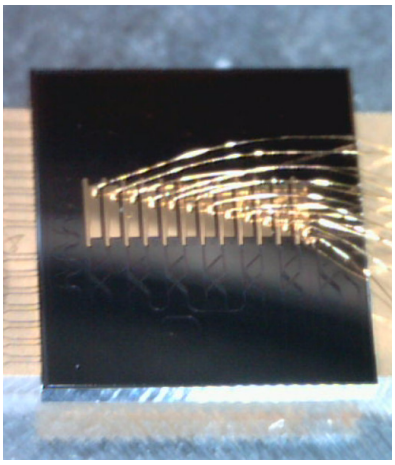

Fig. 1. Multistage optoelectronic switching circuit. The InGaAsP/InP chip is die bonded to a patterned AlN carrier and amplifier gates are wire bonded for separate electronic addressing. Crossbars are aligned vertically.

B. *Cross-bar switch cell*

The architecture is based upon six interconnected crossbar cells with two input and two outputs. The cell is derived from the broadcast and select design widely implemented in semiconductor optical amplifier based switches. This generic form is shown in figure 2(i). The signals from both inputs are split using multimode interference couplers and separately gated prior to the shuffling of two paths and the subsequent recombination. Noting that the electrode numbers scale poorly with increasing numbers of connections, we simplify the electronic implementation. Increased overall connectivity is therefore practicable in such multi-stage interconnection networks with additional numbers of switching cells [1].

Figure 2(ii) shows the concept for simplified electronic wiring. Two waveguide connections on the left-side splitters are exchanged to allow the upper and lower electrodes to be paired. Operation is limited to cross and bar state operation, but the control complexity is reduced. The paired gates can now be biased with the same electrode, and importantly, can be implemented within the same active island [15]. Only two active islands are now used for a crossbar. This flexibility at the optical plane allows considerable simplification at the opto-electronic and electronic plane. The electrode pair requires complementary binary drivers, and this therefore simplifies electronic control significantly. The full circuit contains six crosspoint elements, and twenty-four amplifier gates, but only six logical states are required to define the full switch configuration.

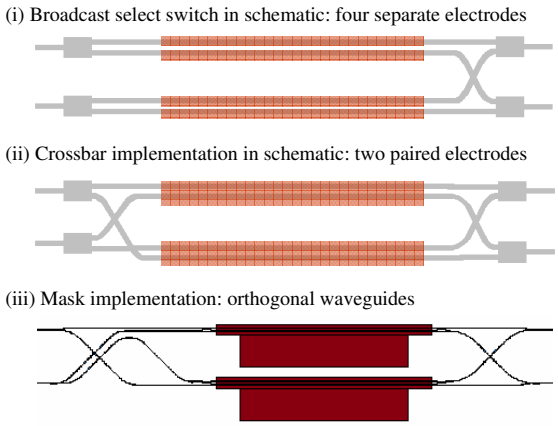

Fig. 2. The base 2x2 switch cell with reduced complexity electrodes.
(i) Schematic showing generic broadcast and select waveguides with individual electrode addressing
(ii) Schematic showing crossbar implementation in this work. The optical wiring is tailored to enable two pairs of electrodes for cross and bar state
(iii) Mask file data for the waveguide layer and the electrode layer whereby the waveguides are arranged for low-crosstalk orthogonal crossings

The final image in figure 2(iii) shows the mask layer implementation for the waveguide layer and the metal layer for the contacting to the gate pairs. A combination of shallow and deep etch waveguides are performed to suit the function required. For example, shallow waveguides are used to ensure improved crosstalk performance waveguide crossings and

facilitate minimum leakage electro-optic components. Deep etch waveguides facilitate tight waveguide bends of order 100 microns and fabrication-tolerant splitter and combiner performance. Mode matching between the waveguide designs is performed using adiabatic tapers. Six crossbar elements are repeated along one dimension to map onto the available area in the predefined active-passive regrown wafer.

C. *Interconnection*

This study is restricted to architectures which concatenate multiple stages of two input two output crossbar elements. These may include, amongst others, the N-stage planar architecture from Spanke and Benes [16] as shown in figure 3 through to the more recently described optoelectronic omega network [17]. The schematic in figure 3 shows sufficient numbers of crossbar stages to enable rearrangeably non-blocking operation.

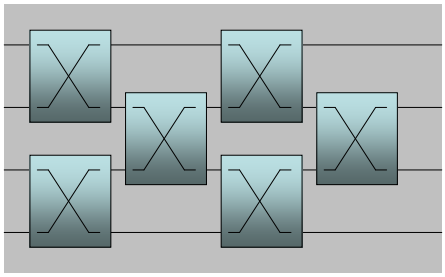

Fig. 3. Multistage switch architecture example enabling four inputs on the left, four outputs on the right. The architecture has a minimum number of two stages, and a maximum number of four stages of crossbar switch cells.

The circuit which has been realised is based on the N-stage planar network schematically drawn in figure 3. Full function for the fabricated circuit is constrained by the inadvertent exchange of two internal nodes and an electrical short circuit at one of the twelve electrodes. Importantly however, the ten correctly functioning paths include both the shortest path and the longest path through the architecture. This allows us to estimate both best and worst case circuit performance for the intended four input four output multistage network.

The optical 'wiring' between crossbar stages incorporates a rich mix of optical waveguide components. Combinations of tight deep etched bends and orthogonal waveguide crossings enable the interconnection of the crossbar elements. The input and output waveguides are all placed on the same facet for ease of subsequent packaging by means of fiber arrays. The layout is primarily restrained by the regrown wafer topology and the spacing of electronic contact pads rather than the dimensions of any given optical or optoelectronic element. Figure 4 shows the implemented layout in this work. Superimposed on the circuit are the paths for two stage routing and four stage routing assessment.

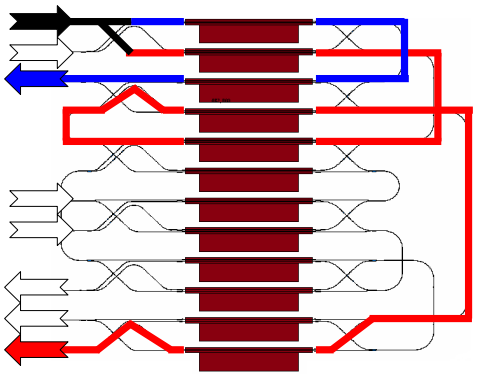

Fig. 4. Interconnection diagram for the fabricated optoelectronic circuit. The examples of the shortest and longest paths as studied in this work are highlighted.

Four inputs and four outputs are defined on the left side of the figure. These connect directly to the vertically stacked crossbar elements which comprise two electrodes for the cross and bar states respectively. These are interconnected by means of deep etch curved waveguides with minimum bend radius of 100 microns. Waveguide crossings are implemented orthogonally in shallow etched waveguides for minimum crosstalk and loss.

III. Experimental Assessment

The switching circuit is epoxy bonded to a gold plated, custom patterned, AlN substrate. The twelve electrodes are wire bonded to coplanar tracks and connected to individually controllable current sources. Preliminary measurements assess the electrical properties of the amplifier gates. Maximum resistance values of 12 Ω are measured with one short circuit.

The chip-submount assembly is subsequently fixed by vacuum clamp to a temperature stabilised positioning system operated at a temperature of 22 degrees. Two lensed fibers are scanned across the facet to assess the switching paths individually. Fibre to chip coupling loss for the non-optimised reflective interface is estimated to be of order 6dB from scaled photocurrent measurements. Ten paths are verified for connection. Operating currents are varied up to 150mA per switch gate pair, at which point facet reflections lead to oscillations. There is no discernable evidence of intracircuit reflection.

The optical gain spectrum is centred at 1565nm. The cascading of multiple stages of SOA leads to a small narrowing of the 3dB bandwidth to 20nm after four stages, corresponding to potentially useful bandwidth of 2.6THz per path for the chosen epitaxy. The measurements performed within this study use a centre wavelength of 1550nm.

IV. High Speed Serial Data Routing

The experimental assessment for the switch circuit is performed at both 40Gb/s and 160Gb/s over paths representing the best and worst cases through an N-stage planar network. These are highlighted in figure 3 as the shortest two stage path and the longest four stage path respectively.

A mode-locked fiber laser is used to generate a 40GHz optical impulse train. The pulses are electrically modulated at the Mach-Zehnder modulator. A four-fold interleaver replicates decorrelated copies of the optical data to generate the time division multiplexed 160Gb/s serial data. The interleaver is tailored to take as input 2^7-1 sequences for each tributary channel and to generate 2^7-1 sequences at the multiplexed line rate (see appendix). The time multiplexed optical data is amplified and filtered with a 5nm bandwidth filter. The polarisation state is aligned for maximum gain and extinction ratio. Light is input and output from the circuit under test using lensed single mode fibres. The return path to the error test equipment includes a filtered Erbium pre-amplifier and an optical demultiplexer from 160Gb/s to 40Gb/s. The experimental assessment of the switch circuit is performed in a separate laboratory to the 160Gb/s error rate assessment system. Over 100 metres of interconnecting standard single mode fibre is therefore also deployed in the experimental assessment. An additional 8 meters and 12 meters of dispersion compensating fiber accommodate the incurred chromatic dispersion after each fiber span.

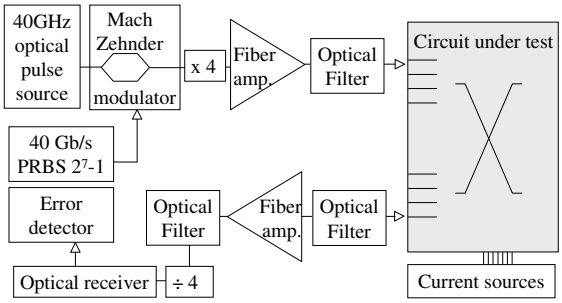

Filter bandwidths 5nm; polarisation control not shown

Fig. 5. Experimental arrangement for the assessment of the circuits at up to 160Gb/s serial line rates.

A. 40Gb/s Line rate

Measurements for 40Gb/s serial data dispense with the time division multiplexer (x4) and the time division demultiplexer (/4) in figure 4 but the experimental arrangement and filter bandwidths are otherwise identical. An aggregate optical power of +7dBm in-fibre is injected into the circuit. Measurements of the error rate performance are made before and after the insertion of the two circuit paths.

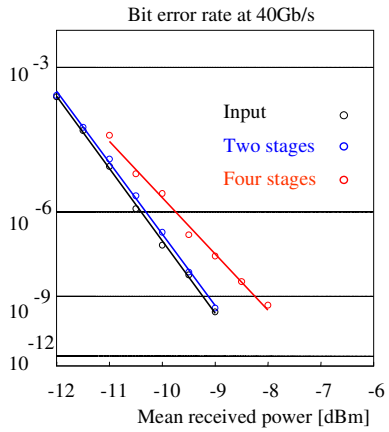

Fig. 6. Bit error rate performance at 40Gb/s serial line rate.

Data in figure 6 are measured for switch electrode currents of 120 mA for each electrode. Negligible power penalty is observed for two switch stages with near-complete overlap between the back to back and short switch path error rate dependence on receiver power. A small deviation in slope is perceivable for the four stage analysis leading to a power penalty of order 1dB. The excess penalty from four stages may be attributable to a build up in amplified spontaneous emission noise. All traces indicate error rates to below 10^{-9} with no evidence of error floors.

B. 160Gb/s Line rates

For the 160Gb/s data routing experiments, all four tributary channels are aggregated in time and routed through the switch circuit. They are subsequently assessed independently to study the error rate performance and power penalty.

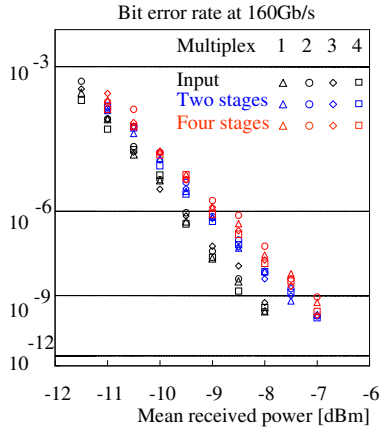

Fig. 7. Bit error rate performance at 160Gb/s. The channels are optically time multiplexed from four 40Gb/s tributaries which are denoted by triangles, circles, diamonds and squares respectively.

The power penalty for the two stages is measured to be of the order of 0.6dB for switch operating currents of 120 mA and an in-fibre input power of +7 dBm. This increases to 1.2dB for the four stage path operating at bias currents of between 95mA and 130mA. Data are shown in figure 7.

Time resolved representations are recorded with the eye diagram traces in figure 8. An Agilent 86119A ultra-wideband optical sampler front end is used in combination with an oscilloscope to view the 160Gb/s data. A clear opening between the ones level and the zeros level for the return-to-zero data format is observed, indicative of the low signal degradation observed.

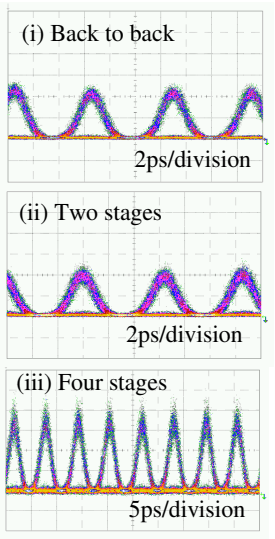

Fig. 8. Eye diagrams for the data modulated 160Gb/s data streams (i) without the switching network, (ii) passing two stages of the switching network, and (iii) after four stages of the switching network.

V. Discussion

The design, implementation and the demonstration of the first four stage switching network and the subsequent low-penalty transmission of serial 160Gb/s represent important advances in their own right. It is none-the-less of interest to understand the limitations which may result from the proposed technology.

The logarithmic relationship between the bit error rate and the receiver power is indicative of noise limited performance. No evidence of optical power dependent degradation which may be attributable to patterning or nonlinearity is therefore observed. This is an important finding, and is indicative of further performance enhancement with an optimised power map within the circuit.

Increasing the splitting losses between each stage would allow the amplifiers to operate in a lower noise figure regime and enable a rapid increase in connectivity. This enables a more rapid scaling of connectivity with the number of stages, providing a promising route to tens of connections. It is noted that the circuit real estate for large interconnection circuits may be predominantly defined by the constraints in active-passive regrowth technology. Additionally the single column of active islands is a restriction imposed only by the multi-project wafer processing, and does not constrain future design. The imposed pitch of active elements is expected to facilitate sufficient thermal isolation and ease of electronic connection.

The achieved net gain which can be accessed within the circuit has been modest due to 34% reflections at both the input and output waveguides. This leads the circuit to oscillate at relatively modest currents and has constrained the values used in this work. Anti-reflection coatings can be readily applied to reduce reflections to well below 1% and thereby suppress spectrally dependent gain fluctuations. This is expected to facilitate operation with higher gain.

Benchmarking with alternative technologies may be most readily considered in terms of the energy consumption per switched bit. Assuming an operating current of 0.13A maximum for each crossbar, a worst case series resistance of 12Ω and a diode voltage at 1550nm wavelength to be 0.8V, an intrinsic power consumption per crossbar of 0.3W can be reached. For an N-stage planar circuit with four inputs operating with 160Gb/s per path, and six crossbars consuming 0.3W we can estimate an energy efficiency of under 3pJ/bit. The potential for even further energy savings will evidently arise from trading power penalty and loss for additional optical multiplexing at the transceiver. The bandwidth of the optoelectronic circuit is of order tens of nanometers, while only 5nm of bandwidth is currently exploited.

VI. Conclusion

The excellent power penalty performance of 1.2dB at 160Gb/s over four stages of crossbar switches in a monolithic interconnection network represents an important milestone in high-bandwidth monolithic circuit integration and offers a highly promising route to large scale monolithic interconnection circuits operating at ultrahigh line rates.

Appendix I:

Optical time division multiplexing

Generating true pseudorandom sequences at ultrahigh bit rates using bit interleaved sequences of the same sequence is facilitated by the correct choice of time delays when combining the sequences.

In this work, an electrical binary sequence length of 2^7-1 is generated at 40Gb/s with a pulse pattern generator. The clock rate is quadrupled to 160Gb/s using two serial Mach-Zehnder fiber interferometers: a commercially supplied optical clock multiplier from Pritel. In one arm of each Mach-Zehnder interferometer there is a fine adjust delay for bit synchronisation and power equalisation. The second arm includes a half pattern length delay line. This careful choice of delay results in the generation of 2^7-1 patterns at twice the

repetition frequency. The operation is repeated to transform the 80Gb/s 2^7-1 PRBS to a sequence at 160Gb/s with the appropriate delay in the second Mach Zehnder interferometer stage. For the bit at time interval t, the pseudo-random bit sequence at the doubled line rate may be derived as follows:

$$\text{PRBS}_{80\text{Gb/s}}(\tfrac{1}{2}\,t) = \begin{cases} \text{PRBS}_{40\text{Gb/s}}(t) & t \text{ odd} \\ \text{PRBS}_{40\text{Gb/s}}(t + \tfrac{1}{2}\,2^7) & t \text{ even} \end{cases}$$

The operation is repeated to transform the 80Gb/s 2^7-1 PRBS to a sequence at 160Gb/s with half the delay in the second Mach Zehnder interferometer stage. The specific experimental arrangement used is shown in fig. 9 below. Polarisation maintaining components are used throughout the interleaver.

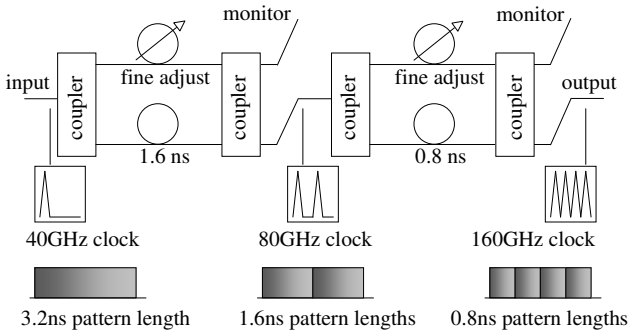

Fig. 9. Experimental arrangment for generating pseudo random bit sequences at 2^7-1 at 160Gb/s using optical time division multiplexing. This apparatus is an explicit description of the functional block shown as x4 in figure 5.

Optical time division demultiplexing is performed by optically gating the received 160Gb/s data stream with an electroabsorption modulator. The 40GHz data clock is applied in electrical form to the modulator. The phase delay between the received optical input and the applied electrical input determines which demultiplexed 40Gb/s tributary is selected for subsequent error rate analysis. Optically de/multiplexing of distinct data sequences with relaxed timing tolerance is feasible with parallel optical modulators.

ACKNOWLEDGMENT

The authors wish to thank the Photonic Integration Platform JePPIX (coordinator D.J. Robbins) for provision of the device. The European Commission Framework 7 Network of Excellence program BONE: Building the Future Optical Network in Europe and the Mexican National Science and Technology Council (CONACyT) are also acknowledged.

REFERENCES

[1] White I.H. et al., "Scalable optical switches for computing applications", OSA Journal of Optical Networking, Invited paper, 8, 2, 215–224 (2009)
[2] Williams, K.A., "Physical layer modelling of semi-conductor optical amplifier based Terabit/second switch fabrics", NUSOD, Post-deadline paper ThPD5 (2008)
[3] Wang. H. et al., "Temperature independent optical amplification in uncooled quantum dot optical amplifiers", Optical Fiber Communications conference, 2008
[4] Nagarajan, R. and Smit, M. "Photonic integration", IEEE LEOS News, 21, 3, 4-10, June 2007
[5] T. Lin et al., "Capacity scaling in a multi-host wavelength striped SOA-based switch fabric", IEEE/OSA Journal of Lightwave Technology, 25, 3, 655-663 (2007)
[6] R. Luijtens and R. Grzybowski, "The OSMOSIS optical packet switch for supercomputers", Optical Fiber Communications Conference, OtuF3, San Diego (2009)
[7] O. Liboiron-Ladoucceur et al., "The data vortex optical packet switched interconnection network", IEEE/OSA Journal of Lightwave Technology, 26, 13, 1777-1789 (2008)
[8] A. Albores-Mejia et al., "Scalable quantum dot optical switch matrix in the 1.55 micron wavelength range", Photonics in Switching, Paper D-06-4 (2008)
[9] H. Wang et al., "Lossless multistage SOA switch fabrics using high capacity monolithic 4x4 SOA circuits", Optical Fiber Communications Conference, OWQ2 (2009)
[10] K.A. Williams, "Integrated semiconductor optical amplifier-based switch fabrics for high capacity interconnects", OSA Journal of Optical Networking, 6, 2, 189-199 (2007)
[11] Kristian E. Stubkjaer, "Semiconductor optical amplifier-based all-optical gates for high-speed optical processing", IEEE Journal on Selected Topics in Quantum Electronics, 6, 6, 1428-1425 (2000)
[12] M. Scaffardi et al., "Photonic combinatorial network for contention management in 160 Gb/s interconnection networks based on all-optical 2×2 switching elements", IEEE Journal of Selected Topics in Quantum Electronics, 13, 5, 1531-1539 (2007)
[13] Liu, Y. et al., "Ultrafast all-optical wavelength routing of data packets utilizing an SOA-based wavelength converter and a monolithically integrated optical Flip–Flop", IEEE Journal of Selected Topics in Quantum Electronics, 14, 3, 801 - 807 (2008)
[14] www.jeppix.eu
[15] A. Albores-Mejia et al., "Integrated 2x2 quantum dot optical crossbar switch in 1.55 μm wavelength range", IET Electronics Letters, 45, 6, 313-314 (2009)
[16] R. Spanke and V.E. Benes, "N-stage planar optical permutation network", OSA Applied Optics, 26, 7, 1226-1229 (1987)
[17] A. Shacham et al., "Experimental demonstration of a complete SPINet optical packet switched interconnection network", Optical Fiber Communications Conference, Paper OThF7 (2007)