

Design of a reconfigurable optical interconnect for large-scale multiprocessor networks

I. Artundo^{*a}, W. Heirman^b, C. Debaes^a, J. Dambre^b, J. Van Campenhout^b, H. Thienpont^a

^aDept. of Applied Physics and Photonics, Vrije Universiteit Brussel, Pleinlaan 2, B-1050 Brussels, Belgium.

^bElectronics and Information Systems Dept., Ghent University, Sint Pietersnieuwstraat 41, B-9000 Ghent, Belgium.

ABSTRACT

Communication between processors and memories has always been a limiting factor in making efficient computing architectures with large processor counts. Reconfigurable interconnection networks can help in this respect, since they can adapt the interconnect to the changing communication requirements imposed by the running application, and optical technology and photonic integration allow for an easy implementation of such adaptable systems. In this paper, we present a proposed reconfigurable interconnection network in the context of distributed shared-memory multiprocessors. We show through full-system simulation of benchmark executions that the proposed system architecture can provide a significant speedup for shared-memory machines, even when physical limitations due to low-cost optical components are introduced. We propose then a reconfigurable optical interconnect implementation, making use of tunable sources and a selective broadcasting component, and we report on the first fabricated optical components of the design: refractive microlenses, fiber connectors, microprism holders and alignment plates.

Keywords: Broadcasting, Optical interconnections, Reconfigurable architectures, Shared memory systems, interconnection networks, diffractive optics

1. INTRODUCTION

Presently, electrical connections on printed circuit boards are the most common way to exchange data between different nodes (processors and memories) in multiprocessor machines. These high-speed electrical interconnection networks are running into several physical limitations such as signal attenuation, electromagnetic interference and severe crosstalk [1]. Modern interprocessor communication technologies, such as HyperTransport [2] and Sun Fireplane [3], can deliver high data throughput for the current generation of multiprocessor machines, but it is widely recognized that important technological enhancements will be required in the very near future. We can currently see an introduction of optical technologies deeper into the multiprocessor architecture to alleviate communication problems, and photonic integration already allows for on-chip interconnection networks.

In Distributed Shared-Memory (DSM) multiprocessors, all the memory of the system is distributed among its nodes. Nodes can access memory on other nodes in a software transparent way. The interconnection network is thus part of the memory hierarchy and therefore high network latencies cause a significant performance bottleneck in program execution. Reconfigurability in this aspect will allow the system to rearrange the interprocessor communication network to avoid congestion and form topologies that are best suited for the particular computing task at hand. This allows for a network topology that closely matches the traffic patterns exhibited by the running application [4].

It is the goal of this work to study how a practical reconfigurable optical network can be implemented into DSM systems with a specific focus on the issues associated to low cost optical technology. This paper is presented as follows: in Section 2 we further introduce the problematic of current interconnection networks and how they can be partially solved by the use of optics. Given the fact that wavelength reconfiguration can be easily achieved in an optical interconnect, we show in Section 3, through detailed full-system simulations, the performance improvement that can be obtained in memory access times through network topology reconfiguration. Section 4 describes a proposed implementation and the optical components needed for building such a proposed reconfigurable design, and Section 5 concludes the work.

2. RECONFIGURABLE OPTICAL INTERCONNECTS

2.1. Multiprocessor network architectures

We have focused our study on multiprocessor machines that implement a hardware-based shared-memory model. They usually have a proprietary interconnection technology yielding high throughput (tens of Gbps per processor) and very low network latency (down to a few hundred nanoseconds). Since the interprocessor communication here is largely hidden from the application developers, it is still manageable to program relatively large applications that exploit parallelism. This makes such class of machines very popular for corporate server environments. The architecture can however do little to hide communication latency, which makes the performance of such machines very vulnerable to increased network latencies.

Modern examples of this class of machines range from small, 4- or 8-way Symmetric Multiprocessor (SMP) server machines (including multi-core processors), over server mainframes with tens of processors (Sun Fire, IBM iSeries), up to supercomputers with hundreds of processors (SGI Altix, Cray XT3/4, etc.). The communication needs of these systems differ radically between them, so we will focus on the midrange servers. Here, interconnection networks have been moving away from topologies with uniform latency, such as busses, into highly non-uniform architectures where latencies between pairs of nodes can vary by a large degree. It is clear that as soon as the processor count scales up, it becomes unachievable to implement fully interconnected networks. As a result, different topologies arise, such as hyper-mesh, torus and trees, balancing high connectivity with acceptable technological complexity.

To get the most out of these machines, data and processes that communicate often should be clustered onto neighboring network nodes. However, this clustering problem often cannot be solved adequately when the communication pattern exhibited by the program has a structure that cannot be mapped efficiently (i.e., using only single-hop connections) on the network topology. With hard-wired topologies, performance can vary greatly for the same architecture depending on the traffic patterns generated by the applications in execution. For this reason, an efficient adaptable scheme would result in a large performance improvement by dynamically adjusting the network topology to match the specific run-time requirements.

Therefore, DSM systems are very likely candidates for the application of reconfigurable interconnection networks. Reconfigurability can adapt the interconnection network to better fit the communication needs, which depend upon the application that is running on the machine, and alleviate the large bottleneck affecting the current communication networks. It can also serve to create a backup path in case there is a failure in any of the components of the system, which is a rare situation but critical enough to take it in consideration for the high-end market.

2.2. Optical technologies in interconnects

Optics is a great candidate to introduce high throughput interconnection networks in the architecture of multiprocessor systems [5-7]. Using optical interconnects at the scale of link lengths found in multiprocessor machines (up to a few meters), the connectivity and bandwidth can be improved, whereas the design of conventional electrical interconnects is limited by the trade-off between interconnection length and bit rate. The high operating frequency of light tends to virtually eliminate frequency dependent cross-talk. The inherent voltage isolation, low power loss and consumption, and low heat dissipation are also highly demanded characteristics for high speed communication channels.

Another relatively unexploited aspect of optics is that it is easy to switch a connection path in a data transparent way, by e.g. tuning the right wavelength or adapting the path of the light beam. This way, we can construct flexible topologies that are able to respond quickly to changing traffic demands.

Of course, one must find the appropriate components for the targeted application by keeping in mind that the requirements and pricing for an optical interprocessor communication network are very different than those of telecom or local area networks. For the optical reconfiguration of short range interconnects, switching times and wavelength tuning requirements can be more relaxed as we will not be doing full packet switching but rather much slower topology adaptations exploiting long-lived dynamics in the traffic streams. When designing an interconnection system, it is a challenge to avoid complicated architectures or costly switching devices. For the leading low-cost solutions, a designer must overcome the fact that (optical) switching speed and connectivity do not come for granted, so it is necessary to find communication and reconfiguration schemes that not only overpass these limitations but try to use them for their advantage.

As laser sources, vertical cavity surface emitting lasers (VCSELs) are an attractive candidate for optical interconnects in general, primarily because of their low-cost mass production and testing on wafer-scale, low power consumption, easy array integration and coupling into optical fibers, and small form factor for easy onboard integration. Recent advances in VCSEL technology have made them widely adopted in interconnect solutions and other uses, and their reliability has been proven for high-performance applications [10-11]. Microelectromechanical structures (MEMS)-tunable VCSELs are an emerging extension to this technology, which allows for fast wavelength switching, and many technological issues have been solved in recent years, allowing their use in adaptable optical networks [12-14]. Their tuning range (a few tens of channels) and speed (between tens of ms and hundreds of ns) would be adequate for following the traffic patterns targeted in our architecture.

2.3. Optical integration and reconfigurable systems

In large optical interconnection networks, switching data directly at the photonic layer offers considerable advantages. The routing may be controlled with a lower-speed electrical control plane, avoiding complex, parallel, broadband electronics. Wavelength channels may be routed, split and combined in parallel and data can be transferred between wavelength channels without incurring in a significant delay and with low optical power penalty. However, customized photonic designs and sophisticated specialized components can lead to complex architectures, large and bulky systems, and difficult implementation in the computer system, while making escalation still not clear. Photonic integration is usually seen as a way to address these limitations. The reduced coupling losses, lower cooling needs, integration of multiple active elements and reduced numbers of optical connections can lead to reductions in fabrication price, power consumption, packaging materials, delays in data transmission and control complexity. Several demonstrators have already been built to showcase the feasibility and capacities of the new generation of integrated optical interconnects [15-17].

Recent attempts to implement optical reconfiguration on the interconnect include the works on a chip-to-chip (N-to-N) reconfigurable optical interconnect employing VCSEL and photodetector arrays reported in [18]. Here, the reconfiguration is done by two arrays of liquid crystal (LC) cells driven by a VLSI circuit placed on the side of a prism-like structure. They generate digital holographic diffraction gratings to steer and multicast optical beams, reconfiguring this way the switching states of the 2.5 Gbps optical links. A more limited 1.25 Gbps free space 1-to-N reconfigurable optical interconnect is reported in [19], employing VCSEL-PIN links, and binary phase gratings on a LC SLM. Beams are steered free-space over a surface of 6.4 mm² with a resolution of 50 μm. We also find in SELMOS [20] another board-level reconfigurable optical interconnect over film-waveguide 3D structures. It is composed of an embedded 3D 1024x1024 microoptical switching system based on PLZT waveguide-prism-deflector switches, and a self-organized optical network, that couples light paths between two waveguides automatically.

3. NETWORK RECONFIGURATION

3.1. Proposed architecture

Our proposed interconnection network architecture for a DSM system consists of a fixed base network connecting all the nodes (processors and local memories), arranged in a torus topology. In addition to these connections, a limited number of freely reconfigurable optical links are added periodically for a given time span (see Fig. 1) that connect nodes that are expected to have a large communication load.

A torus was chosen as the base network since it is one of the most extended models used in multiprocessor designs due to its easy implementation and high connectivity. The added optical extra links, called from now on *elinks*, can be used as direct point-to-point short-cut connections to route the traffic between processor node pairs on the network. After a certain interval of time, the elinks are freed again and can be reassigned according to other processor pairs according to new traffic measurements.

This setup, compared to the case where potentially all links in the network are available for the topology reconfiguration, is advantageous as the base network will always be available. It is therefore impossible to disconnect parts of the network, greatly reducing the complexity in the routing and reconfiguration algorithms. However, this reconfigurability can only boost system performance when some requirements are met. Indeed, as any reconfiguration scheme takes a certain time before the reconfigured links become operative again, it is crucial that reconfiguration is done wisely. Therefore, it is necessary to know in the most accurate way:

- Where to place the additional links, meaning detecting where the congestion is happening by searching which pairs of nodes are consuming more instant bandwidth.
- When is the right moment to perform the reconfiguration, obtaining the best performance for the interval of time when the extra links are assigned.




Fig. 1. Torus topology of the base interconnection network with additional reconfigurable links. The numbers correspond to the different processor nodes in the network.

The influence of these two limitations has been thoroughly studied in previous papers [21], so in the following sections we will focus on showing the performance improvement results for large size networks with the limitations stated above, and report on the components needed for the proposed system design.

3.2. System simulation

For simulating the performance of such a network architecture in a real multiprocessor system, we have based our simulation platform on the VirtuTech Simics simulator. It was configured to simulate a multiprocessor machine resembling the Sun Fire 6800 server, with 16+ UltraSPARC III processors clocked at 1GHz and running the Solaris 9 operating system. The directory-based coherency controllers and the interconnection network are custom extensions to Simics, modeling a directory-based protocol and a packet switched torus network with contention and cut-through routing. The SPLASH-2 benchmark suite was chosen as the workload. It consists of a number of scientific and technical parallelized applications and is a good representation of the real-world workload of large shared-memory machines. Average packet latency has been chosen as the low-level metric to evaluate performance improvement in the internode communications. More details on the simulation environment can be found in [21].

The system parameters taken in consideration for our simulations are the number of processor nodes in the network, p the number of extra links added to the topology, n , and the fan-out f for every node, meaning the number of outgoing optical extra links from that node. Besides, the topology will be reconfigured every certain time interval Δ_t , according to previous congestion measurements between pairs of nodes in the network. This last parameter is of paramount importance, as the switching time for the topology change (including tuning time of the VCSELs or select/route calculations) will determine how fast we can follow the traffic patterns in the communication network. This is translated into how accurately we can keep track of the evolution of network congestion, and therefore adapt the network in the best possible way, by placing the optical elinks in the right node pairs. Finally, as a way to simulate the limited tuning range of the low-cost optical sources, we will also set a limit on the connectivity for the extra links. We will effectively cluster the network in subsets of 9 nodes, a reasonable amount corresponding to available wavelength channels, allowing only for a limited set of destinations for every node.

As a result, in figure 2 a summary is shown of the average packet latency for different network architectures. First, a non-reconfigurable torus-only network is shown. We add then the extra links in the following network simulations, always with $n = p$. The second set of bars plots the latency for an idealized, $f = 4$, $\Delta_t = 100 \mu\text{s}$ network. Next, the fan-out is restricted to a more realistic $f = 2$, then $f = 1$ along with the limited 1-to-9 connectivity is simulated. Finally, the reconfiguration interval is lengthened to $\Delta_t = 1 \text{ ms}$ for this last configuration. The average latency can be seen to drop when introducing a reconfigurable network. Some of this gain has to be relinquished when implementation constraints

(fan-out limitation, limited connectivity, slower tuning) are introduced, but even then a very visible performance improvement can be obtained. A torus has $4 \cdot p$ unidirectional links. The n elinks therefore represent an additional total bandwidth of 25%. The "global" measurements in the figure show a torus network in which each link has a bandwidth that is 25% higher than the link bandwidth in the previous situation. The total network bandwidth is therefore the same as in the $n = p$ reconfigurable cases. As conclusion, memory access times can be reduced up to 40% of the default network case, meaning that even with the associated physical restrictions due to a low-cost implementation, reconfigurability in the network can bring a significant performance improvement, especially over large size networks.




Fig. 2. Average packet latency reduction for several network architectures, depending on the fan out, the topology and the reconfiguration interval length.

4. SELECTIVE OPTICAL BROADCAST IMPLEMENTATION

4.1. System design

In a practical implementation of a tunable VCSEL, the number of available wavelengths is generally very limited, normally having a trade-off between tuning speed and the channel count. The restriction on the number of channels prohibits the use of a broadcast-and-select scheme in which all processing nodes (say over 64 nodes) are connected together via a single star-coupling element, besides power considerations. We therefore propose a selective optical broadcasting (SOB) design which broadcasts each channel to only a limited number of outputs, dictated by the number of wavelength channels on the tunable VCSEL, but can scale this way to larger configurations by clustering the network.

On previous years, an optical module based on this concept has been prototyped at our labs [23], where the initial design was very similar, still including the diffractive microlenses and a microprism for the end-point interconnect. Similar designs have been proposed already [18], with the optical interconnection designed and tested by McFadden et Al. [22] as one of the closest ones. In the latter, an optical setup is built comprising microlenses, etched compound slope microprisms and a field-scale curved macromirror for the in-plane reflection.

As a logical extension of this work, a redesign of our previous interconnect module with extended alignment features, including fiber connectors, is then presented as follows. The proposed network is composed of tunable VCSEL sources, single mode/multimode optical fibers, two sets of diffractive microlenses and a selective optical broadcast (SOB) element with the corresponding alignment features (see Fig. 3 and 4). The single mode fibers coming from the tunable sources in each processing node are bundled into a fiber array at the ingress of the SOB. Two plates are used to keep alignment of the fibers, by guiding the fiber cores through microholes until the exit plane. This design allows for an escalation of 20x10 optical channels available for reconfiguration.

The SOB component, thanks to the input set of diffractive microlenses with a 3x3 splitting pattern, will then fan-out the signal to an array of spots at the output, and light will be focused back to the output multimode fibers via microlenses too. Reception is done again at processor side with resonant cavity with alternative wavelengths. This way, each processor is capable of connecting to a restricted set of 9 different nodes, and by tuning the right wavelength, one can address the proper destination from the subset. It is important to note that the mapping of the source node connections and the receiving nodes on the broadcasting component is now critical, because it directly determines the possible addressable nodes for every transmitting processor.




Fig. 3. Overall system diagram, showing how alignment and broadcasting is achieved with the right angle microprism.




Fig. 4. SOB system 3D diagram, front and side view. Two perpendicular holder plates keep the right-angle prim fixed and aligned with the lens plate.

We have started with the prototyping of the micro-optical assembly. Up to now, several components of the design have been fabricated and characterized, by making use of the Deep Proton Writing (DPW) technology [24], a micro-optical rapid prototyping technology that is being developed at the VUB. DPW consists of the irradiation of polymethyl metacrylate (PMMA) with a pencil-like proton beam followed by a selective etching or swelling of the irradiated zones. Selective etching results in high quality optical surfaces, or micro-hole arrays in case of proton beam point irradiations. We can also swell the point irradiations, resulting in large arrays of microlenses with dedicated focal numbers. In the following subsections we detail the fabricated components up to date.

4.2. Microlenses and diffractive grating

The first component is a lens plate with a 5x5 array of 140 μm spherical refractive microlenses with a targeted 350 μm focal length and a sag height of 15.16 μm (see Fig. 5a). This lens plate is intended to be used preliminary to test the point-to-point communication of the assembly, as the actual lens plate with the diffractive capabilities is being fabricated in the moment. However, calibration of the technology for the precise lens fabrication has proven to be more difficult than expected. Averaged standard deviation for lens height in 7 different irradiations is 2.01 μm , giving an excessive variability in focal length for the same repeated process. Further tests are being made to improve on the quality and accuracy of the DPW lens swelling technique.




Fig. 5. (a) Microscope picture of a 5x5 array of refractive lenses. (b) 3x3 splitting diffractive pattern.

For the actual diffractive design, simulations have been carried out first with LightTrans VirtualLab v1.3 for modelling a 3x3 splitting pattern over 1000 iterations. We obtained a basic cell (see Fig. 5b) with a diffraction efficiency of 89.5% and a Signal-to-Noise ratio (SNR) of 137 dB, as the latter will be an important factor to consider for the light propagation inside the prism due to low-power emission and internal scattering. This basic cell pattern is then replicated all over the fresnel-like microlens to form the diffractive design, which will be fabricated through e-beam written masks over fused silica.

Propagation of light was simulated along the system, obtaining a 74.6% of propagated optical power at detector with a worst-case point-to-point SNR of 20.37 dB. Analyzing single received spots on the detector side, there is a peak variation of 4.7% in received power as the diffractive pattern gets a different output depending on the wavelength (see Fig. 6). Note that peak power is slightly displaced from central design wavelength due to sampling errors on the diffractive patter and approximations on the simulation engine.

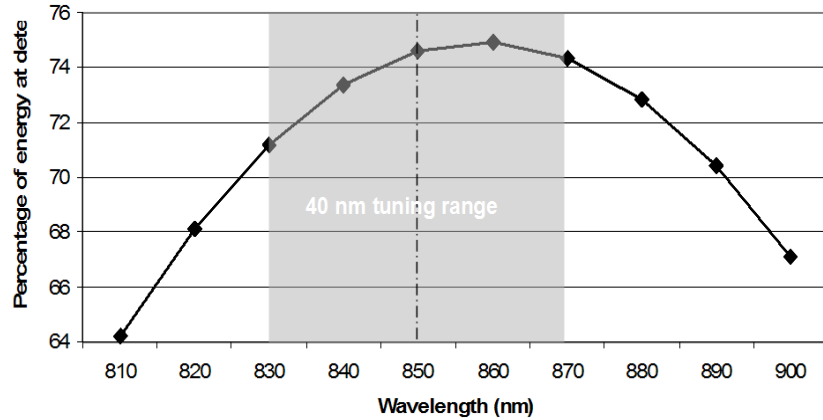




Fig. 6. Power variation at detector side with wavelength. Tuning range interval represents common tuning range for low-cost tunable VCSELs.

4.3. Fiber alignment plates

The second component (see Fig. 7a) that was fabricated so far is the fiber prealignment plate, a 500 μm thick PMMA rectangular plate measuring 20x7 mm, with two arrays of 5x5 irradiated microholes of 125 μm diameter for the fibers and two irradiated circular holes of 700 μm for the MT ferrule micropins. The rectangular plate outline has been micromilled. We measured the etched microholes and calculated their circularity by employing a circle as the assessment feature, obtaining a 0,016 mean error over a perfect shape. Perfect cylindrical shape was not obtained, showing a conical shape due to the proton beam scattering on the PMMA; diameter of the holes back side was 165 μm , making fiber insertion easier. The insertion still needs to be done manually, by removing first the plastic coating from the fiber core before passing through, but a translation-stage insertion setup is planned in the future.



(a)



(b)

Fig. 7. (a) Fiber prealignment plate. (b) Fiber alignment plate.

The third component (see Fig. 7b) corresponds to the fiber alignment plate, a 30x30 mm 500 μm thick PMMA plate with two arrays of 5x5 irradiated microholes of 125 μm diameter for the fibers, two irradiated circular holes of 700 μm diameter for the micropins, and three 382 μm diameter irradiated circular holes to hold the microspheres that will align the fiber plate with the lens plate. The hardened steel microspheres are 508 μm in diameter and they are positioned in a 16 mm side equilateral triangle for stable equilibrium.

4.4. Right-angle microprism and prism holders

The fourth and fifth component fabricated up to date are two identical PMMA plates, 500 μm thick too, that will hold and align the prism above the lens plate. The prism is a BK7-glass right angle prism, 5 mm side and 7.1 mm hypotenuse, ± 3 arcmin tolerance on the right angle, a surface flatness of $\lambda/10$ and a 40-20 scratch-dig surface finish. The plates consist of a DPW-irradiated right angle for maximum precision on the feature alignment, and two irradiated rectangular holes of 1005x500 μm with rounded corners to fit the holders on the lens plate. The rest of the outline, including the handle in the right side, has been micromilled with a 500 μm drill as the surface quality here is not critical. These components are depicted below in Figure 8.




Fig. 8. (a) Right-angle microprism (b) Prism holder plate with DPW-irradiated right angle

5. CONCLUSIONS

In this paper, we have addressed the problem of communication bottleneck in multiprocessor interconnection networks due to high internode latencies, and we have proposed the use of optical technologies to alleviate congestion by implementing a reconfigurable network. In this network, extra optical links are placed temporarily where high load is expected in the network, obtaining memory access reductions of up to 40% for several benchmark applications. A practical low cost implementation has been proposed, a system where light is emitted from the processor nodes through tunable VCSELs, guided with optical fibers to a optical broadcast component, and selectively received by a set of 9 other nodes.

Several components of this design have been already fabricated, and we reported on the fiber prealignment and alignment plates, the 5x5 arrays of refractive microlenses, the 3x3 diffraction grating and the holder structures for the right angle microprism. Future work includes fabricating the rest of the components and assembling the whole system in a testbed.

Given the advances in integrated on-chip parallel optical technologies, our system could be easily scaled down to make a more compact adaptive interconnect with a higher channel density, and the same principles could be applied to implement an integrated reconfigurable network.

ACKNOWLEDGEMENTS

This work is financed by the EC 6th FP Network of Excellence on Micro-Optics (NEMO), the FWO, IWT-GBOU, VUB-GOA, the Interuniversity Attraction Poles Program photonics@be (IAP-Phase VI), and the OZR of the VUB.

REFERENCES

- [1] Miller, D. A. B., "Rationale and challenges for optical interconnects to electronic chips," *Proceedings of the IEEE* **88**, 728-749 (2000).
- [2] HyperTransport consortium, "HyperTransport IO Technology Comparison with Other IO Technologies," <http://www.hypertransport.org> (2004).
- [3] Charlesworth, A., "The sun fireplane system interconnect," *Proc. of the ACM/IEEE conference on Supercomputing*, pp. 7 (2001).
- [4] Dvorak, V., "Reconfigurability of the interconnect architecture for chip multiprocessors," *Proc. of the 4th international symposium on Information and communication technologies*, 136-141 (2005).
- [5] Mohammed, E., et Al., "Optical Interconnect System Integration for Ultra-Short-Reach Applications," *Intel Technology Journal* **8(2)**, 115-27 (2004).
- [6] Collet, J., Litaize, D., Campenhout, J. V., Desmulliez, M., Jesshope, C., Thienpont, H., Goodman, J., Louri, A., "Architectural approach to the role of optics in monoprocessor and multiprocessor machines," *Applied Optics* **39**, 671-682 (2000).
- [7] Benner, A. F., et Al., "Exploitation of optical interconnects in future server architectures," *IBM Journal of Research and Development* **49**, 755-775 (2005).
- [8] Tissot, Y., Russell, G. A., Symington, K. J., Snowdon, J. F., "Optimization of Reconfigurable Optically Interconnected Systems for Parallel Computing," *Journal of Parallel and Distributed Computing* **66**, 238-247 (2006).
- [9] Chen, G., Chen, H., et Al., "Predictions of CMOS compatible on-chip optical interconnect," *Integration, the VLSI Journal* **40(4)**, 434-446 (2007).
- [10] Koyama, F., "Recent Advances of VCSEL Photonics," *Journal of Lightwave Technology* **24**, 4502-4513 (2006).
- [11] Cunningham, J. E., McElfresh, D. K., Lopez, L. D., Vacar, D., Krishnamoorthy, A. V., "Scaling vertical-cavity surface-emitting laser reliability for petascale systems," *Applied Optics* **45**, 6342-6348 (2006).
- [12] Maute, M., Kögel, B., Böhm, G., Meissner P., Amann, M. C., "MEMS-Tunable 1.55- μm VCSEL With Extended Tuning Range incorporating a Buried Tunnel Junction," *IEEE Photonic Technology Letters* **18**, 688-690 (2006).
- [13] Huang, M. C. Y., Bun Cheng, K., Zhou, Y., Pisano, A. P., Chang-Hasnain, C. J., "Monolithic Integrated Piezoelectric MEMS-Tunable VCSEL," *IEEE Journal of Selected Topics in Quantum Electronics* **13**, 374-380 (2007).
- [14] Huang, M. C. Y., Zhou Y., Chang-Hasnain, C. J., "A nanoelectromechanical tunable laser," *Nature Photonics* **2**, 180 (2008).
- [15] Shacham, A., "Building Ultralow-Latency Interconnection Networks Using Photonic Integration," *IEEE Micro* **27**, 6-20 (2007).
- [16] Nagarajan, R., et Al., "Large-scale photonic integrated circuits," *IEEE Journal of Selected Topics in Quantum Electronics* **11**, 50-65 (2005).
- [17] Kimerling, L. C., et Al., "Electronic-photonic integrated circuits on the CMOS platform," *Proc. SPIE* **6125**, (2006).
- [18] Aljada, M., Alameh, K. E., "High-speed (2.5 Gbps) reconfigurable inter-chip optical interconnects using opto-VLSI processors," *Optics Express* **14**, 6823-6836 (2006).
- [19] Henderson, C. J., Leyva, D. G., Wilkinson, T. D., "Free space adaptive optical interconnect at 1.25 Gb/s, with beam steering using a ferroelectric liquid-crystal SLM," *J. Lightwave Technol.* **24**, 1989-1997 (2006).
- [20] Yoshimura, T., Ojima, M., Arai, Y., Asama, K., "Three-dimensional self-organized microoptoelectronic systems for board-level reconfigurable optical interconnects - Performance, Modeling and simulation" *IEEE Journal of Selected Topics in Quantum Electronics* **9(2)**, 492-511 (2003).
- [21] Heirman, W., Dambre, J., Artundo, I., Debaes, C., Thienpont, H., Stroobandt, D., Van Campenhout, J., "Predicting the performance of reconfigurable optical interconnects in distributed shared-memory systems," *Photonic Network Communications* **15**, 25-40 (2008).
- [22] McFadden, M. J., et Al., "Multiscale free-space optical interconnects for intrachip global communication: motivation, analysis, and experimental validation," *Appl. Opt.* **45**, 6358-6366 (2006).
- [23] Artundo, I., Desmet, L., Heirman, W., Debaes, C., Dambre, J., Van Campenhout, J., Thienpont, H., "Selective optical broadcasting in reconfigurable multiprocessor interconnects," *Proc. of SPIE Photonics Europe* **6185**, (2006).
- [24] Debaes, C., Van Erps, J., Vervaeke, M., Volckaerts, B., Ottevaere, H., Gomez, V., Vynck, P., Desmet, L., Krajewski, R., Ishii, Y., Hermanne, A., Thienpont, H., "Deep Proton Writing: a rapid prototyping polymer microfabrication tool for micro-optical modules", *New Journal of Physics, Focus on nanotechnology* **8**, 270 (2006).