

High-speed (2.5 Gbps) reconfigurable inter-chip optical interconnects using opto-VLSI processors

Muhsen Aljada and Kamal E. Alameh

*Centre for MicroPhotonic Systems, Electron Science Research Institute, Edith Cowan University,
Joondalup, WA, 6027, Australia.
m.aljada@ecu.edu.au, k.alameh@ecu.edu.au*

Yong-Tak Lee and Il-Sug Chung

*Gwangju Institute of Science and Technology, Department of Information and Communication,
Gwangju 500-712, Korea.
ytleee@gist.ac.kr*

Abstract: Reconfigurable optical interconnects enable flexible and high-performance communication in multi-chip architectures to be arbitrarily adapted, leading to efficient parallel signal processing. The use of Opto-VLSI processors as beam steerers and multicasters for reconfigurable inter-chip optical interconnection is discussed. We demonstrate, as proof-of-concept, 2.5 Gbps reconfigurable optical interconnects between an 850nm vertical cavity surface emitting lasers (VCSEL) array and a photodiode (PD) array integrated onto a PCB by driving two Opto-VLSI processors with steering and multicasting digital phase holograms. The architecture is experimentally demonstrated through three scenarios showing its flexibility to perform single, multicasting, and parallel reconfigurable optical interconnects. To our knowledge, this is the first reported high-speed reconfigurable N-to-N optical interconnects architecture, which will have a significant impact on the flexibility and efficiency of large shared-memory multiprocessor machines.

©2006 Optical Society of America

OCIS codes: (200.4650) Optical interconnects; (200.2610) Free-space digital optics (090.1970) Diffractive optics

References and links

1. B.E. Lemoff, M.E. Ali, G. Panotopoulos, G.M. Flower, B. Madhavan, A.F.J. Levi, D.W. Dolfi, "MAUI: enabling fiber-to-the-Processor with parallel multiwavelength optical interconnects," *J. Lightwave Technol.* **22**, 2043 – 2054 (2004).
2. D. Fey, W. Erhard, M. Gruber, J. Jahns, H. Bartelt, G. Grimm, L. Hoppe, and S. Sinzinger, "Optical interconnects for neural and reconfigurable VLSI architectures," *Proc. IEEE*, **88**, 838 – 848 (2000).
3. L.A. Buckman Windover, J.N. Simon, S.A. Rosenau, K. Giboney, G.M. Flower, L.W. Mirkarimi, A. Grot, B. Law, C.K. Lin, A. Tandon, R.W. Gruhlke, H. Xia, G. Rankin, and D. W. Dolfi, "Parallel optical interconnects beyond > 100 Gb/s," *J. Lightwave Technol.* **22**, 2055–2063 (2004).
4. Cho. Hoyeol, P. Kapur, and K.C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," *J. Lightwave Technol.* **22**, 2021 – 2033 (2004).
5. F.-C. F. Tsai, C. J. O'Brien, N. S. Petrović, and A. D. Rakić, "Analysis of optical channel cross talk for free-space optical interconnects in the presence of higher-order transverse modes," *Appl. Opt.* **44**, 6380–6387 (2005).
6. X. Wang, F. Kiamilev, G.C. Papen, J. Ekman, P. Gui, M.J. McFadden, J.C. Deroba, M.W. Haney, and C. Kuznia, "Performance-based adaptive power optimization for digital optical interconnects," *Appl. Opt.* **44**, 6240–6252 (2005).
7. R. Bockstaele, T. Coosemans, C. Sys, L. Vanwassenhove, A. Van Hove, B. Dhoedt, I. Moerman, P. Van Daele, R. Baets, R. Annen, H. Melchior, J. Hall, P.L. Heremans, M. Brunfaut, J. Van Campenhout, "Realization and characterization of 8 x 8 resonant cavity LED arrays mounted onto CMOS drivers for POF-based interchip interconnections," *IEEE J. Sel. Top. Quant. Elect.* **5**, 224–235 (1999).

8. V. Baukens, G. Verschaffelt, P. Tuteleers, P. Vynck, H. Ottevaere, M. Kufner, S. Kufner, I. Veretennicoff, R. Bockstaele, A. Van Hove, B. Dhoedt, R. Baets, H. Thienpont, "Performances of optical multi-chip-module interconnects: comparing guided-wave and free-space pathways," *J. Optic. Pure Appl. Optic.* **1**, 255-261 (1999).
9. R.T. Chen, L. Lin, C. Choi, Y.J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M.K. Hibbs-brenner, J. Bristow, and Y. S. Liu, "Fully embedded board-level guided-wave optoelectronic interconnects," *Proc. IEEE*, **88**, 780-793 (2000).
10. F. Mederer, R. Jager, H.J. Unold, R. Michalzik, K.J. Ebeling, S. Lehmacher, A. Neyer, and E. Griesse, "3-Gb/s data transmission with GaAs VCSEL's over PCB integrated polymer waveguides," *IEEE Photon. Technol. Lett.* **13**, 1032-1034 (2001).
11. Y. Li, J. Ai, and J. Popelek, "Board-level 2-D data-capable optical interconnect circuits using polymer fiber-image guides," *Proc. IEEE*, **88**, 794-805 (2000).
12. T. May, A.G. Kirk, D.V. Plant, J.F. Ahadian, C.G. Fonstad, K.L. Lear, K. Tatah, M.S. Robinson, and J.A. Trezza, "Interconnection of a two-dimensional array of vertical-cavity surface-emitting lasers to a receiver array by means of a fiber image guide," *Appl. Opt.* **39**, 683-689 (2000).
13. H.F. Bare, F. Haas, D.A. Honey, D. Mikolas, H.G. Craighead, G. Pugh, and R. Soave, "A simple surface-emitting LED array useful for developing free-space optical interconnect," *IEEE Photon. Technol. Lett.* **5**, 172 - 175 (1993).
14. J. Jahns, Y.H. Lee, C.A. Burrus, and J.L. Jewell, "Optical interconnects using top-surface-emitting microlasers and planar optics," *Appl. Opt.* **31**, 592-597 (1992).
15. D.V. Plant, M.B. Venditti, E. Laprise, J. Faucher, K. Razavi, M. Chateaneuf, A.G. Kirk, and J.S. Ahearn, "256-channel bidirectional optical interconnect using VCSEL's and photodiodes on CMOS," *J. Lightwave Technol.* **19**, 1093-1103 (2001).
16. D.V. Plant, and A.G. Kirk, "Optical interconnects at the chip and board level: challenges and solutions," *Proc IEEE*, **88**, 806 - 818 (2000).
17. J.J. Liu, Z. Kalayjian, B. Riely, W. Chang, G.J. Simonis, A. Apsel, and A. Andreou, "Multichannel ultrathin silicon-on-sapphire optical interconnects," *IEEE J. Sel. Top. Quant. Elect.* **9**, 380 - 386 (2003).
18. V.M. Hietala, C. Chun, J. Laskar, K.D. Choquette, K.M. Geib, A.A. Allerman, and J.J. Hindi, "Two-dimensional 8x8 photoreceiver array and VCSEL drivers for high-throughput optical data links," *IEEE J. Solid-State Circuit.* **36**, 1297-1302 (2001).
19. R. Wang, A.D. Raki, and M.L. Majewski, "Design of microchannel free-space optical interconnects based on vertical-cavity surface-emitting laser arrays," *Appl. Opt.* **41**, 3469-3478 (2002).
20. A. Goulet, N. Nieuborg, H. Thienpont, A. Kirk, P. Koczyk, P. Heremans, M. Kuijk, C. De Tandt, W. Ranson, R. Vounckx, and I. Veretennicoff, "Polarization-based reconfigurable optical interconnects in free-space optical processing modules," *IEEE Photon. Technol. Lett.* **10**, 367 - 369 (1998).
21. C.J. Henderson, D.G. Leyva, T.D. Wilkinson, "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).
22. M. Aljada, K.E. Alameh, and K. Al-Begain, "Opto-VLSI-based correlator architecture for multi-wavelength optical header recognition" *J. Lightwave Technol.* **24**, 2779-2785 (2006).
23. S. Ahderom, M. Raisi, K.E. Alameh, and K. Eshraghian, "Dynamic WDM equalizer using opto-VLSI beam processing," *IEEE Photon. Technol. Lett.* **15**, 1603-1605 (2003).
24. Z. Wang, R. Zheng, K.E. Alameh, R. Robertson, U. Mueller, and L. Bloom, "Opto-VLSI-based dynamic optical splitter," *Electron. Lett.* **40**, 1445 - 1446 (2004).
25. I.G. Manolis, T.D. Wilkinson, M.M. Redmond, and W.A. Crossland, "Reconfigurable multilevel phase holograms for Optical switches," *IEEE Photon. Technol. Lett.* **14**, 801-803 (2002).
26. H. Dammann, "Spectral characteristics of stepped-phase gratings", *Optik* **53**, 409-417 (1979).
27. R.D. Jeffery, K.E. Alameh., and M. Vasiliev, "Design of reconfigurable optical interconnects employing Opto-VLSI processors," in *Proceeding of the International Conference on Very Large Scale Integration*, 2005 (International Federation for Information Processing, Western Australia, 2005) pp. 39-44.

1. Introduction

In large shared-memory multiprocessor machines, hundreds of CPUs and terabytes of memory and integrated storage devices, requiring ultrawide bandwidths and expected to carry as much as a few Tb/s between processors by the end of this decade [1]. The imbalance between satisfying on-chip computing power and insufficient off-chip short-distance communication performance has lead to an interconnect crisis in microelectronics [2]. Despite their reliability, simplicity, and low cost, electrical interconnects have fundamental technical challenges including reduced power requirements, reduced transmission latency, and greater interconnect density, which prevent them from being used in high-capacity interconnects. On

the other hand, optical interconnects have the potential to not only solve many of these problems but also add benefits of extended reach and immunity to electromagnetic interference (EMI) [3].

Optical interconnects [4]-[8] have recently emerged as viable alternatives to high-speed data buses in electronics computers and signal processors, and at the end of this decade optical interconnects will be expected to carry much of the board-to-board and processor-to-processor bandwidth burden in high-end computer systems [1]. Among the optical interconnect structures proposed for high-speed data communications are polymer waveguides [9] [10], fibre image guides [11], [12], fibre ribbons [1] [3], and free space optical interconnects using lens and mirror systems [13]-[19].

Free-space optical interconnects, in particular, offer promising solutions to achieve high-bandwidth, low-power-consumption data communication links. In addition, the use of free-space enables dense and reconfigurable optical interconnects to be realised simultaneously.

Reconfigurable free-space optical interconnect modules employing polarization-selective diffractive optical elements in combination with a liquid crystal based polarization controller have previously been reported by Goulet, et al.[20]. However, these modules result in high optical losses especially when the number of output ports increases. Recently, a 1.25Gb/s free space 1-to-N optical interconnect structure employing a single ferroelectric liquid-crystal spatial light modulator in conjunction with a free space optical polarizing beam splitter, half-wave plates and collimating lenses has been reported [21].

In this paper, we propose and demonstrate the concept of chip-to-chip high-speed reconfigurable optical interconnects employing VCSEL and photodetector arrays in conjunction with two Opto-VLSI processors driven by digital phase holograms that reconfigure the switching states of the optical links. To our knowledge, this is the first high-speed Opto-VLSI-based reconfigurable N-to-N optical interconnect architecture reported to date.

The paper is organised as follows: In Section 2 we review the steering and multicasting capabilities of Opto-VLSI processors. Section 3 presents the reconfigurable optical interconnect architecture, Section 4 presents the optical component design specifications for the architecture. Experimental setups, results and discussions are reported in Section 5.

2. Opto-VLSI processor

An Opto-VLSI processor is an electronically controlled, software-configured, polarization independent, motionless beam steerer, comprising an array of liquid crystal (LC) cells driven by a Very-Large-Scale-Integrated (VLSI) circuit, that generates digital holographic diffraction gratings to steer and/or shape optical beams [22-24], as illustrated in Fig. 1. Each pixel is assigned a few memory elements that store a digital value, and a multiplexer that selects one of the input voltages and applies it to the aluminum mirror plate.




Fig. 1. Opto-VLSI processor and a typical 8-phase LC cell structure design.

Figure 1 also shows a typical layout and a cell design of a multi-phase Opto-VLSI processor. Indium-Tin Oxide (ITO) is used as the transparent electrode, and evaporated aluminum is used as the reflective electrode. By incorporating a thin quarter-wave plate (QWP) layer between the liquid crystal and the VLSI backplane, a polarization-insensitive Opto-VLSI processor can be realized, allowing optical beam steering with polarization-dependent loss as low as 0.5 dB, as demonstrated by Manolis et al. [25]. The ITO layer is generally grounded and a voltage is applied at the reflective electrode by the VLSI circuit below the LC layer to generate stepped blazed gratings for optical beam steering.

Figure 2 illustrates the steering and multicasting capabilities of an Opto-VLSI processor of pixel size d , driven by digital phase holograms (Fig. 2(a)). A blazed grating of arbitrary pitch can be generated by digitally driving a block of LC pixels with a linear phase hologram whose pitch can be controlled by changing the voltage applied to each pixel, thus realising beam steering. For a blazed grating of pitch $q \times d$, the steering angle Θ is proportional to the wavelength (λ) of the light and inversely proportional to $q \times d$, as shown in Fig. 2(b). On the other hand, if a multicasting phase hologram is synthesised, multiple beams can be generated, whose intensities can be controlled by reconfiguring the phase hologram.

For a small incidence angle, the maximum steering angle of the Opto-VLSI processor is given by [23]:

$$\theta_{\max} \approx \frac{\lambda}{M \cdot d} \text{ (Radians)} \quad (1)$$

Where M is the number of phase levels, d is the pixel size, and λ is the wavelength of the incident beam. For example, a 4-phase Opto-VLSI processor having a pixel size of 5 microns can steer a 1550 nm laser beam by a maximum angle of around $\pm 4^\circ$. The maximum diffraction efficiency of an Opto-VLSI processor is given by [26]:

$$\eta = \text{sinc}^2 \left(\frac{\pi n}{M} \right) \quad (2)$$

where $n = gM + 1$ is the diffraction order ($n = 1$ is the desired order), and g is an integer. Thus an Opto-VLSI processor with binary phase levels can have a maximum diffraction efficiency of 40.5%, while a four phase levels allow for efficiency up to 81%. The higher diffraction orders (which correspond to the cases $g \neq 0$) are usually unwanted crosstalk, which must be attenuated or properly routed outside the output ports to maintain a high signal-to-crosstalk performance.




Fig. 2. (a) Phase holograms driving various pixel blocks, (b) Principle of beam steering and multicasting using an Opto-VLSI processor.

3. Optical interconnect architecture

Figure 3 illustrates the concept for the reconfigurable optical interconnect architecture which employs VCSEL and photodetector (PD) arrays in conjunction with two Opto-VLSI processors attached on an optical substrate, which integrates microlens arrays that collimate and focus the different optical beams.

The active window of the first Opto-VLSI processor is partitioned into different pixel blocks, each allocated to a VCSEL element. On the other hand, the pixel blocks of the second Opto-VLSI processor are assigned to the various photodetector elements. For an optical interconnect between VCSEL # i to PD # j , the i^{th} pixel block of the first Opto-VLSI processor is driven by digital phase hologram that steers the i^{th} VCSEL beam along the j^{th} pixel block of the second Opto-VLSI processor for final steering along the j^{th} PD element.

4. Component specifications and circuit designs

The PCB driving the VCSEL and PD array was designed to prove the capability of the Opto-VLSI processors to reconfigure high-speed optical interconnects. The VCSEL circuitry is shown in Fig. 4(a). It comprise of a single mode (SM) 850nm 1x12 VCSEL array and a 3.2 Gbps VCSEL driver (LTC5100) having a unique output stage confining the modulation current to the ground, thus isolating the high-speed signal from the power supply to minimise the RF interference. The VCSEL array has a threshold current of 3mA, maximum output power of 1.2 mW and 250 μ m element spacing. Only three VCSEL elements were bonded as shown in Fig. 4 (b). The circuit design for a VCSEL element is shown in Fig. 4(c), where the data inputs are AC coupled, eliminating the need for external capacitors. Note that the VCSEL driver, LTC5100, has a precisely controlled 50 Ω output that is DC coupled to the laser enabling arbitrary placement of the VCSEL chip. The process control of the LTC5100 uses an I²C serial interface.




Fig. 3. Reconfigurable optical interconnect architecture.

The PCB driving the PD array is shown in Fig. 5(a). The PD array is a 5 Gbps GaAs PIN photodiode of responsivity 0.5 A/W. Figure 5(b) is a photomicrograph showing the PD array bonded to the PCB. Each PD element is connected to a 3.2Gbps MAX3725 transimpedance amplifier, which preamplifies the detected photocurrent, followed by a 3.5Gbps MAX3748A multirate limiting amplifier, as shown in Fig. 5(c).




Fig. 4. (a) PCB of the VCSEL circuitry, (b) VCSEL array after flip-chip bonding, and (c) transmitter circuit design.




Fig. 5. (a) PCB of the PD circuitry, (b) PD array after flip-chip bonding, (c) Receiver circuit design.

5. Experimental setups and results

In order to prove the reconfigurability of the proposed architecture, we carried out three experimental scenarios using a 1x4096-pixel, 256-phase Opto-VLSI processor of 1.8-micron pitch. The first scenario was intended to demonstrate a reconfigurable single optical interconnect between a VCSEL element and any PD element. The second scenario focused on demonstrating the multicasting capability of the proposed reconfigurable optical interconnect architecture. Finally, the third scenario shows the capability of the proposed architecture to establish parallel optical interconnects

5.1 First Scenario: A single reconfigurable optical interconnect

The experimental set up of the first scenario is shown in Fig. 6, where the Opto-VLSI processors were reconfigured to establish optical interconnects between a VCSEL element modulated with a 2.5Gbps clock signal and the three PD elements. For each optical interconnect, optimised phase holograms are generated on the first Opto-VLSI processor to steer the VCSEL beam to different pixel blocks on the second Opto-VLSI processor. Each pixel block is driven by an optimised phase hologram that steers the VCSEL beam towards the targeted PD element.




Fig. 6. First interconnect scenario: establishing reconfigurable single optical interconnects.

Figures 7(a)-(c) show the optimum steering holograms which drove the first Opto-VLSI processor to establish optical interconnects from the VCSEL element to the first, second, and third PD element, respectively. Figure 7(d) shows the phase holograms driving the second Opto-VLSI processor, which steer the VCSEL beams to the various PD elements. It is noted that the steering phase holograms are blazed gratings with different pitches which are inversely proportional to the steering angle, as evident from Eq. (1). A computer algorithm based on real-time optimisation was developed to generate the optimum phase holograms that maximise the signal at the output port of the intended optical interconnects.




Fig. 7. Optimum steering holograms of the Opto-VLSI processors for establishing the different optical interconnects from the VCSEL element to any PD element.

Figure 8 shows the output waveforms and eye diagrams for each interconnect. It is noticed that the output signals are mainly contaminated with photoreceiver thermal noise, however, the output eye diagram opening is large and clear, implying that the signal is adequate for further signal processing .




Fig. 8. Output signals and eye diagrams for optical interconnect between the VCSEL element and the (a) first (b) second and (c) third PD element.

5.2 Second Scenario: Multicasting optical interconnect

Figure 9 shows the second interconnect scenario, where a multicasting is established between a VCSEL element and two PD elements. A multicasting hologram is generated on the first Opto-VLSI processor to equally split the VCSEL beam into two beams, and route them toward the pixel blocks (on the second Opto-VLSI processor) assigned to the first and second PD elements. These pixel blocks are driven by phase holograms that steer the incoming beams to the first and second PD elements. Figure 10 shows the multicasting phase hologram driving the first Opto-VLSI processor as well as the steering phase holograms of the second Opto-VLSI processor, for the multicasting scenario. Figure 11 shows the detected optical signals and eye diagrams at the output ports of both optical interconnect. It is obvious that the output signals have higher noise than those of the first scenario. However, the output eye diagram opening is adequate.




Fig. 9. Second interconnect scenario: establishing a multicasting optical interconnects between a VCSEL element and two PD elements.




Fig. 10. (a) Multicasting hologram generated on the first Opto-VLSI processor to split the VCSEL beam toward the pixel blocks assigned to the first and the second PD elements on the second Opto-VLSI processor, (b) the steering holograms generated on the second Opto-VLSI processor to steer the optical beams to the first and the second PD elements.




Fig. 11. Output signals and the eye diagrams at the output port of (a) the first PD element, and (b) the second PD element.

5.3 Third scenario: parallel reconfigurable optical interconnects

Figure 12 shows the third interconnect scenario, where parallel reconfigurable optical interconnects are simultaneously established between the VCSEL elements and the PD elements. Both Opto-VLSI processors are loaded with proper steering holograms to steer the beams of VCSEL #1, #2, and #3 to PD elements #3, #2, and #1, respectively. Fig. 13 shows the steering holograms generated on both Opto-VLSI processors. Figure 14 shows the detected optical signals and eye diagrams at the output ports. The high output noise is attributed to the crosstalk induced by the high diffraction orders of the Opto-VLSI processors. However, based on our previous work, by optimally positioning the VCSEL and PD elements, the zeroth and higher order diffraction beams of the Opto-VLSI processors can be routed outside the active areas of the PD elements, and thus the crosstalk can be minimised.




Fig. 12. Third interconnect scenario: establishing parallel optical interconnects.




Fig. 13. Steering hologram generated for parallel optical interconnects: (a) on the first Opto-VLSI processor (b) on the second Opto-VLSI processor.




Fig. 14. Output signals and the eye diagrams at the various output ports for the case of parallel optical interconnect.

It is also important to note that the scalability of the demonstrated reconfigurable optical interconnects architecture depends on the size of the active window of the Opto-VLSI processor, which can practically be as large as 20×20 mm. By allocating 128×128 pixels for each optical interconnect and using a pixel size of $5 \mu\text{m} \times 5 \mu\text{m}$, 32×32 optical interconnects can be realised.

Note also that the discrete nature of the phase holograms generated by the Opto-VLSI processors give rise to high-order diffraction beams which fall within the active areas of the various PD elements, thus resulting in crosstalk, as reported in [27]. In order to minimise the crosstalk, the spacing between the VCSELs was intentionally made non-uniform, as seen in Figs 4(b) and 5(b). This enabled the unwanted high-order beams to be routed outside the active areas of the PD elements.

6. Conclusion

We have demonstrated a dynamic reconfigurable optical interconnects using two Opto-VLSI processors. The architecture was constructed using 850nm VCSEL array, photodiode (PD) array, and two Opto-VLSI processors. Each Opto-VLSI processor is driven with proper steering holograms. The proof of the concept has been demonstrated for different scenarios including single, multicasting, and parallel reconfigurable optical interconnects at 2.5Gbps. The performance of the architecture was evaluated by measuring the output signals and eye diagrams for different optical interconnect scenarios, and results have shown that in order to reconfigure an optical interconnect and maintain an adequate eye diagram opening, the VCSEL and PD elements must appropriately be placed and the Opto-VLSI processors must be driven with optimal steering and multicasting holograms that minimise the high order diffraction, and hence the crosstalk. It is worthwhile noting that the issues of latency and

limited bandwidth encountered in electronic multichip communication links are significantly suppressed using optical interconnects. The proposed inter-chip reconfigurable optical interconnects architecture has applications in processor-memory and processor-processor communication.

Acknowledgments

This work is supported by the Office of Science and Innovation, Government of Western Australia. The authors would like to extend their thanks to Dr Roger Jeffery for his help in developing the VCSEL and PD drivers and to Prof. Adam Osseiran for useful discussions.