IEEE Photonics Journal
An IEEE Photonics Society Publication

Open Access

A 40 Gb/s Chip-to-Chip Interconnect for 8-Socket Direct Connectivity Using Integrated Photonics

Volume 10, Number 5, September 2018

Stelios Pitris
Miltiadis Moralis-Pegios
Theonitsa Alexoudi
Joris Lambrecht
Xin Yin
Johan Bauwelinck
Yoojin Ban
Peter De Heyn
Marianna Pantouvaki
Joris Van Campenhout
Ronald Broeke
Nikos Pleros

High-speed AWGR-based chip-to-chip interconnection using integrated photonics

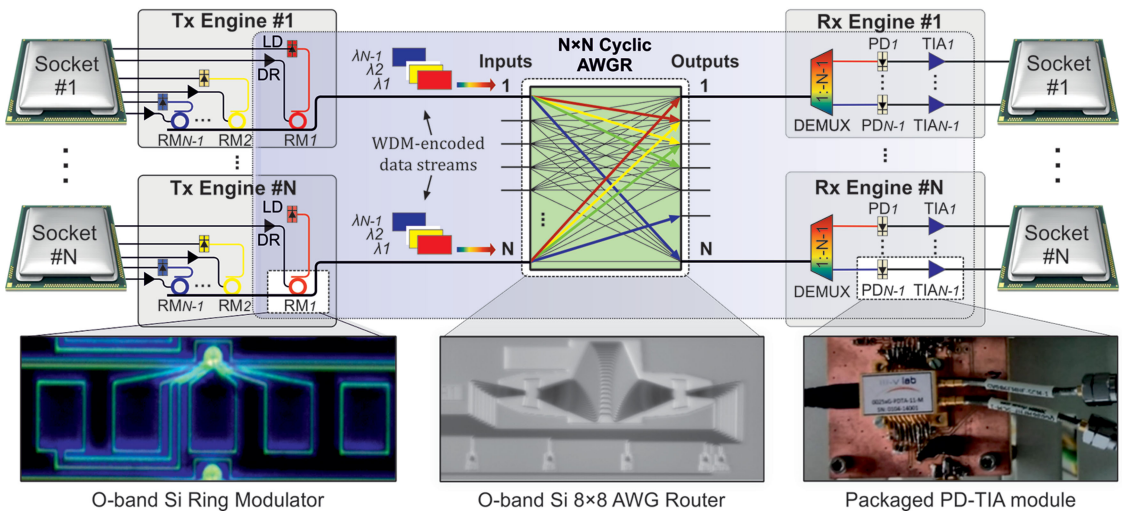

O-band Si Ring Modulator

O-band Si 8×8 AWG Router

Packaged PD-TIA module

DOI: 10.1109/JPHOT.2018.2873673
1943-0655 © 2018 IEEE

A 40 Gb/s Chip-to-Chip Interconnect for 8-Socket Direct Connectivity Using Integrated Photonics

Stelios Pitris,[1] Miltiadis Moralis-Pegios,[1] Theonitsa Alexoudi,[1] Joris Lambrecht,[2] Xin Yin,[2] Johan Bauwelinck,[2] Yoojin Ban,[3] Peter De Heyn,[3] Marianna Pantouvaki,[3] Joris Van Campenhout,[3] Ronald Broeke,[4] and Nikos Pleros [1]

[1]Department of Informatics, Center for Interdisciplinary Research and Innovation, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece
[2]IDLab, Department of Information Technology, Ghent University imec, Ghent B-9052, Belgium
[3]imec, Leuven B-3001, Belgium
[4]Bright Photonics BV, Eindhoven, 5621 AX, The Netherlands

DOI:10.1109/JPHOT.2018.2873673

Manuscript received September 6, 2018; accepted September 29, 2018. Date of publication October 4, 2018; date of current version October 18, 2018. This work was supported by the EC under the H2020-ICT-STREAMS research project (No. 688172). Corresponding author: Stelios Pitris (e-mail: skpitris@csd.auth.gr).

Abstract: We present an O-band any-to-any chip-to-chip (C2C) interconnection at 40 Gb/s suitable for up to 8-socket direct connectivity in multi-socket server boards, utilizing integrated low-energy photonics for the transceiver and routing functions. The C2C interconnect exploits an Si-based ring modulator as its transmitter and a co-packaged photodiode/transimpedance amplifier enabled receiver interconnected over an 8 $\times$ 8 Si-based arrayed waveguide grating router, allowing for a single-hop flat-topology interconnection between eight nodes. A proof-of-concept demonstration of the C2C interconnect is presented at 25 and 40 Gb/s for eight possible routing scenarios, revealing clear eye diagrams at both data rates with extinction ratios of 4.8 $\pm$ 0.3 and 4.38 $\pm$ 0.31 dB, respectively, among the eight routed signals.

Index Terms: Optical interconnections, Wavelength routing, Silicon photonics, Photonic integrated circuits.

1. Introduction

The emerging High-Performance Computing applications and cloud computing are currently generating vast amounts of data traffic in the Data Centers (DC) with intra-DC traffic being expected to reach up to 71.5% of the total DC traffic by 2021 [1]. This trend is currently stretching the limits of the DC infrastructure, calling for both interconnect and computational advancements to raise the bar on efficiency and performance while keeping reasonable cost- and energy- envelopes on DC operation [2]. Towards this goal, research efforts have shifted to efficient chip-to-chip (C2C) interconnections in multi-socket server boards (MSB) seeking for a new technology toolkit for boosting computational power at a reduced energy consumption envelope [3]. MSB systems rely currently on electrically interconnected sockets and can be classified in two categories: i) "glueless" configurations, where direct point-to-point (P2P) interconnects like Intel's QuickPath Interconnect

(QPI) [4] can offer high-speed, low-latency, any-to-any C2C communication for a maximum of 4 sockets, while scaling to 8 sockets can be also achieved by employing, however, dual-hop links and, thus, increased latency setups, and ii) "glued" configurations, where scaling beyond 8-socket layouts is accomplished by exploiting active switch-based setups, such as Bixby [5] and PCI-Express switches [6], in order to interconnect multiple 4- or 8-socket QPI "islands", delivering in this way unlimited connectivity at the cost of increased energy consumption and latency.

Arrayed Waveguide Grating Router (AWGR)-based interconnections have been proposed as a promising solution to overcome this trade-off in MSB systems as they can offer low-latency, non-blocking and all-to-all optical interconnection by exploiting cyclic-routing wavelength properties and data parallelism [7]. The important advantages provided by AWGR-based interconnections in MSBs have been already indicated for 10 Gb/s data rates in the pioneering work of [8], where cyclic-accurate network simulations for different workloads demonstrated 3$\times$ energy savings, higher throughput and lower latency compared to the respective electronic MSB baseline [8]. The employment of bit-parallel transmission and bandwidth-allocation techniques in these AWGR-based architectures has been also shown to allow for a 3.5$\times$ improvement in energy efficiency [9], while broadcast-friendly transmission schemes over AWGR-based MSB layouts highlighted recently the potential for supporting cache coherency updates with additional energy gains of up to 22% [10]. Simulations of the proposed flat-topology AWGR-based MSB interconnect scheme have been restricted to the C-band regime and to 10 Gb/s line-rates [8], while proof-of-concept demonstration using C-band integrated photonic circuitry have been experimentally shown as Network-on-Chip (NoC) layouts without exceeding 0.3 Gb/s data-rate performance [11] limiting its applicability in NoC systems such as high-speed MSB architectures. Additionally, the low-loss waveguide feature [12] and the well-established short-reach interconnect infrastructure [13] of O-band remain unexploited in this domain, although O-band is a promising spectral region for such interconnection schemes.

In this paper, we experimentally demonstrate an O-band C2C interconnection scheme for up to 40 Gb/s flat-topology communication using an integrated Silicon-based Ring-Modulator (RM) transmitter (Tx), an 8 $\times$ 8 AWGR-based routing circuitry and a co-packaged InP-based photodiode (PD) and Transimpedance Amplifier (TIA) receiver (Rx). Error-free operation was achieved for 8 different routing scenarios with non-return-to-zero (NRZ) signals at 25 Gb/s yielding maximum bit error-rate (BER)-curve power variation of 1.1 dB for a 10^{-9} BER value and extinction ratio (ER) of 4.8 $\pm$ 0.3 dB among the 8 routed signals, respectively. The high-speed capabilities of the proposed scheme were also experimentally verified with successful operation at 40 Gb/s, exhibiting similar performance for all 8 routing scenarios with ER of 4.38 $\pm$ 0.31 dB. Its successful O-band operation suggests a significant potential for exploiting the low-loss optical polymer waveguide characteristics [12] and the well-established short-reach interconnect infrastructure [13] towards future realization of Electro-Optical Printed Circuit Board (EOPCB)-based MSB layouts. The energy efficiency of the proposed 40 Gb/s C2C photonic link can be estimated at 6.74 pJ/bit when assuming a 10% wall-plug efficiency for the external laser and an on-board assembled layout, leading to 58.4% reduction in energy compared to the 16.2 pJ/bit link energy efficiency of Intel QPI [14].

2. Concept & Experimental Setup

Fig. 1(a) depicts the conventional 4- (4S) and 8-socket (8S) MSB "glueless" architecture [15] currently supported by direct communication protocols between the processor sockets like Intel's QPI or HyperTransport (HT). In 4S configurations, the reference socket (blue-colored) can reach all of the neighboring sockets with 1-hop transmission (yellow-colored). In 8S topologies, each socket is directly connected to three other sockets with 1-hop, but communication with the remaining 4 sockets can be achieved via at least 2-hops-routes. Scaling the topology beyond 8S cannot be supported by "glueless" topologies and requires the use of "glued" architectures with active switches connecting different "glueless" MSB islands, leading in this way to an increased number of hops and of course higher latency in the system. Fig. 1(b) shows the respective 8S optical connectivity scheme that can be employed in MSBs by exploiting the cyclic routing capabilities of an 8 $\times$ 8 AWGR, which allows all sockets to communicate with each other within only a single-hop. Fig. 1(c) presents

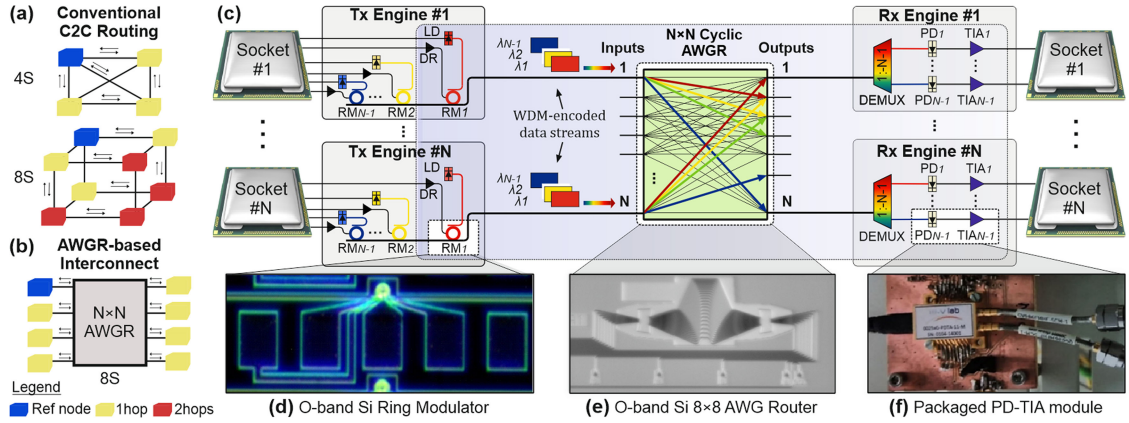

Fig. 1. (a) Conventional C2C routing in 4S and 8S MSB configurations. (b) AWGR-based routing for 8S configuration. (c) proposed optical N × N AWGR-based interconnection concept for MSB connectivity exploiting WDM-enabled Tx/Rx engines. Integrated versions of the basic building blocks employed in this demonstration corresponding to the: (d) Ring Modulator, (e) 8 × 8 cyclic-frequency AWGR, and (f) PD-TIA module (blue-highlighted areas: part of the architecture demonstrated experimentally, white-highlighted areas: basic building blocks used for the demonstration).

the corresponding physical layer C2C optical interconnect architecture in a generic N-socket layout, where an all-passive silicon photonic N × N AWGR routing element provides single-hop any-to-any communication between N processor sockets. Each socket is electrically connected to a wavelength division multiplexing (WDM)-enabled Tx optical engine equipped with N-1 laser diodes (LD), each one operating at a different wavelength. Every LD feeds a different RM to imprint the electrical data sent from the socket to each one of the N-1 wavelengths, so that the Tx engine comprises finally N-1 RMs along with their respective RM drivers (DR). All RMs are implemented on the same optical bus to produce the WDM-encoded data stream of each socket. The data stream generated by each socket enters the input port of the AWGR and gets routed via the AWGR to a destination output that is dictated by the carrier wavelength and the cyclic-frequency routing properties of the AWGR [11]. In this way, every socket can forward data to any of the remaining 7 sockets by simply modulating its electrical data onto a different wavelength via the respective RM, allowing direct single-hop communication between all sockets through passive wavelength-routing. At every Rx engine, the received WDM-encoded data stream gets demultiplexed with a 1:(N-1) optical demultiplexer (DEMUX), so that every wavelength gets received by a distinct PD. Each PD is connected to a TIA that provides the socket with the respective electrical signaling.

The main integrated transmitter, receiver and routing building blocks that were used for demonstrating experimentally the feasibility of the proposed C2C interconnect scheme comprise three discrete chips, i.e., a Si-based RM, a co-packaged PD-TIA and a Si-based 8 × 8 AWGR routing platform:

a) *Silicon micro-ring modulator:* For the Tx engine, the demonstration relied on a silicon O-band carrier-depletion micro-ring modulator. The employed RM is an all-pass ring resonator with 5 μm radius that was fabricated on imec's active platform with demonstrated 50 Gb/s modulation capabilities [16]. Fig. 1(d) shows a microscope image of the employed RM with four electrical pads to access the RM heater and the RM junction through RF probes, respectively. Access to the RM is achieved via TE-polarization grating couplers (GC). The RM was firstly characterized in the DC regime. Fig. 2(a) shows the RM transfer function shifting for a reversed applied DC voltage in the range of 0 V to −3 V. The characterization revealed 32 pm/V modulation efficiency between 0 and −1 V and a static ER of 8.45 dB for a reverse applied voltage of 3 V, respectively. The RM can be combined with a recently developed low-power driver [17], leading to an energy efficiency of 1 pJ/bit at 40 Gb/s which increases to 1.79 pJ/bit when a half free-spectral-range (FSR)-tuning is considered by utilizing the embedded heater [16] to tune the RM to its operating wavelength.

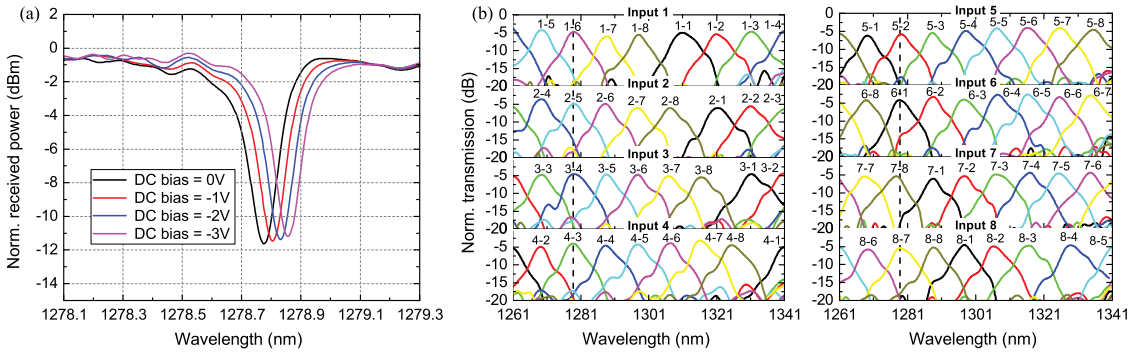

Fig. 2. Transmission spectra of: (a) the integrated RM for different DC bias voltages, (b) the integrated 8 × 8 AWGR.

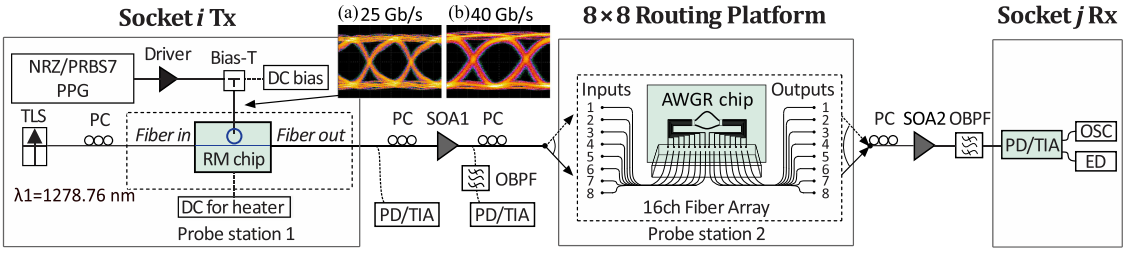

Fig. 3. Experimental setup used for the proof-of-concept demonstration employing the Si-based RM, the Si-based 8 × 8 AWGR chip and the packaged PD-TIA module. Eye diagrams of the electrical signal applied on the RM for operation at (a) 25 Gb/s (100 mV-10ps/div) and (b) 40 Gb/s (100 mV-5ps/div).

b) *Silicon 8 × 8 cyclic AWGR:* For the routing platform, the demonstration relied on an integrated silicon photonic 8 × 8 AWGR device [18]. The AWGR was designed for cyclic-frequency routing operation in the O-band with 10 nm-channel spacing and a free spectral range of 80 nm. The AWGR integration relied on imec-ePIXfab silicon photonics passives technology. Fig. 1(e) shows a microscope image of the fabricated AWGR device with part of the input/output waveguide configuration employing TE-polarization GCs. The fabricated AWGR exhibited a compact footprint of 700 × 270 μm^2. Fig. 2(b) shows the normalized spectral response of all AWGR transmission channels revealing a 5.5 nm 3-dB channel bandwidth, a maximum channel loss non-uniformity of 3.5 dB (with 2.5 dB best-case channel insertion losses) and an average channel crosstalk of 11 dB.

c) *Co-packaged PD-TIA:* For the Rx engine, a high-responsivity uni-traveling-carrier (UTC) InGaAs-InP PIN photodiode connected with a low-power transimpedance amplifier implemented in 0.13 μm SiGe BiCMOS technology was employed with demonstrated 40 Gb/s operation capabilities [19]. The integrated TIA occupied an area of 1100 × 900 μm^2 and was packaged together with the PD into a common butterfly module. Fig. 1(f) shows the single package with the fiber-pigtailed optical input and the differential RF output. The PD-TIA sensitivity for operation at 25 Gb/s and at 40 Gb/s is −10.6 dBm and −6.4 dBm, respectively, while the energy efficiency for operation at 40 Gb/s is 3.95 pJ/bit.

The experimental setup used for the proof-of-concept demonstration at 25 Gb/s and at 40 Gb/s employing the integrated building blocks is shown in Fig. 3. A TLS was used to produce a CW signal at $\lambda 1 = 1278.76$ nm (dotted line in Fig. 2(b)). The RM chip was optically probed with single-mode fibers through TE-polarization GCs while an RF probe was used to access the RM electrical pads. A programmable pattern generator (PPG) was used to generate an NRZ pseudo-random binary sequence (PRBS7) that was amplified by a driver amplifier and applied on the RM along with a reverse-bias DC voltage. Figs. 3(a) and (b) show the eye diagrams of the electrical signal applied on

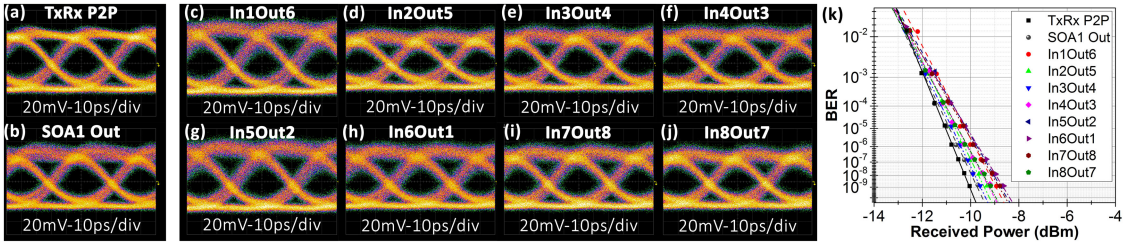

Fig. 4. Eye diagrams of the modulated signal for 25 Gb/s operation: (a) at the output of the ring modulator, (b) at the output of SOA1, (c)–(j) after routing through the respective In#iOut#j I/O ports of the AWGR and amplified by SOA2. (k) BER measurements of the modulated signal for 25 Gb/s operation at different stages of the setup.

the RM for operation at 25 Gb/s and at 40 Gb/s, respectively, exhibiting 600 mVpp and 620 mVpp, respectively. After the RM, the signal was sequentially launched into all of the 8 input ports of the AWGR. A 16-channel fiber array was used to couple the signal in and out of the respective AWGR input/output ports through the GCs. Based on the cyclic-frequency operation of the AWGR and depending on the AWGR input port where the incoming data signal at 1278.76 nm was launched, the data stream was routed each time to a different AWGR output port and received by the PD-TIA, providing in this way the 8 possible routing scenarios that correspond to 8 different input/output port combinations. To obtain the eye diagrams and the BER measurements of the signals, the PD-TIA was connected to a digital sampling oscilloscope (OSC) and to an error detector (ED), respectively. Semiconductor optical amplifiers (SOA1 & SOA2) were incorporated in the setup after the RM and the AWGR chips, respectively, to compensate for the input/output GC losses that were 9 dB for the RM chip and 9 dB for the AWGR chip, respectively. The signal quality was also monitored directly at the output of the RM and after SOA1 using an optical band-pass filter (OBPF) with 2.5 nm 3-dB bandwidth. Polarization controllers (PC) were used to maintain proper signal polarization.

3. Experimental Results

Fig. 4(a) shows the eye diagram of the 25 Gb/s modulated signal when connecting the PD-TIA directly at the RM output, corresponding to the baseline direct P2P link between the Tx and Rx without incorporating any routing functionality and revealing an ER of 5.4 dB with an amplitude modulation on the '1' level of the signal (AM) of 0.66 dB. Fig. 4(b) shows the respective eye diagram after SOA1 with an ER of 5.32 dB and an AM of 1.1 dB. Fig. 4(c)–(j) show the eye diagrams of the signal obtained after the AWGR corresponding to the 8 routing scenarios through all possible input-output port combinations denoted as In#iOut#j, exhibiting ER values of 4.8 $\pm$ 0.3 dB and AM values of 1.2 $\pm$ 0.2 dB, respectively. BER measurements were obtained at 25 Gb/s revealing error-free operation for the 8 routing combinations for a 10^{-9} BER value as shown in Fig. 4(k). The BER curves at the output of the AWGR exhibit a power variation of 1.1 dB for a 10^{-9} BER value due to different routing path insertion losses that translate to different amplification and, as such, different amplifier-induced noise levels experienced by SOA2. The RM was driven with a peak-to-peak voltage of 2.3 V_{pp}, while the applied DC bias voltage was -4 V. No DC voltage was applied to the RM heater during the experimental evaluation, but thermal stability of the chip was maintained by means of a temperature controller. The average optical power of the CW signal entering the RM was 8 dBm, leading to a -6 dBm RM output modulated signal. The optical power of the signal entering the AWGR was 6 dBm, while the output power of the routed signal at the AWGR output ranged between -8.8 dBm to -7.04 dBm for all 8 possible AWGR input/output combinations, with the differences stemming from the different AWGR channel and GC losses. The SOAs were both electrically driven at 160 mA during the experimental evaluation at 25 Gb/s.

To test the feasibility of the architecture towards higher speed operation, the experimental setup of Fig. 3 was further evaluated at 40 Gb/s by simply generating a 40 Gb/s PRBS7 NRZ data signal

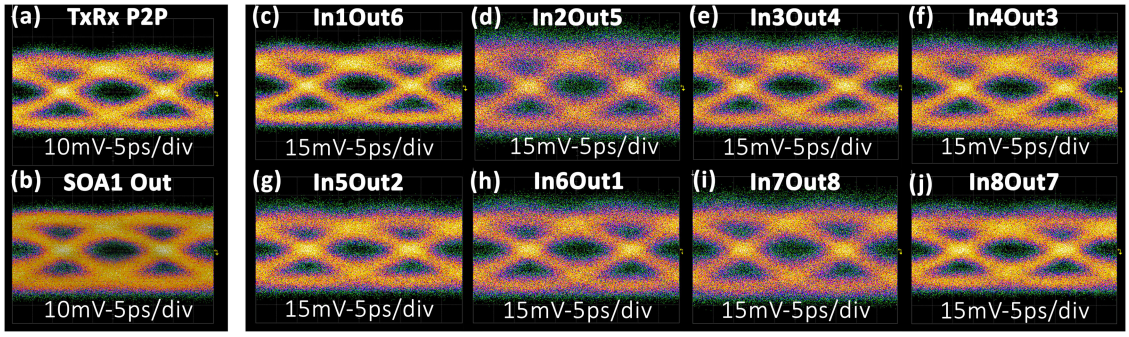

Fig. 5. Eye diagrams of the modulated signal for 40 Gb/s operation: (a) at the output of the ring modulator, (b) at the output of SOA1, (c)–(j) after routing through the respective In#iOut#j I/O ports of the AWGR and amplified by SOA2.

TABLE 1
Transmission Line Coupling & Insertion Losses for the Integrated 40 Gb/s C2C Interconnect

Metric	Value	References
LD to RM coupling losses	3 dB	[22]
RM insertion losses	1.5 dB	This work
RM to PWG coupling losses	0.5 dB	[20]
PWG to AWGR coupling losses	0.5 dB	[20]
AWGR channel insertion losses	6 dB (max)	[18]
PWG to PD-TIA coupling losses	0.5 dB	[20]

via the electrical PPG. Fig. 5(a) and (b) show the eye diagrams of the modulated signal when connecting the PD-TIA at the RM output and after SOA1 with an ER of 4.2 dB and 4.15 dB and amplitude modulation (AM) of 1 dB and 1.3 dB, respectively. Fig. 5(c)–(j) show the eye diagrams of the signal at the 8 outputs of the AWGR corresponding to the 8 routing scenarios for all possible input-output port combinations denoted as In#iOut#j, indicating successful routing at 40 Gb/s with ER values of 4.38 $\pm$ 0.31 dB and AM values of 2.3 $\pm$ 0.3 dB, respectively. The RM was electrically driven with a peak-to-peak voltage of 2.6 V_{pp}, while the applied reverse DC bias voltage was -2.5 V. The optical power of the CW signal injected at the RM input was 8 dBm, with the modulated data signal obtained at the RM output having an average optical power level of -6.3 dBm that was amplified to 10 dBm prior entering the AWGR input. The power of the signal after being routed through the 8 different AWGR port combinations was in the range of -5 dBm to -3.1 dBm. The SOAs were both electrically driven at a higher driving current of 175 mA during the evaluation at 40 Gb/s, which resulted in different gain experienced by the modulated signal by SOA1, and, thus different optical power of the signal entering the AWGR ports compared to the 25 Gb/s evaluation.

By transferring this interconnect onto a polymer hosting board comprising polymer waveguides (PWG), the high losses associated with the input/output GCs of the RM and AWGR chips can be mitigated as GCs will be replaced with low-loss adiabatic coupling structures [7] that have been shown to yield only 0.5 dB of coupling losses over the entire O-band wavelength range [20]. To this end, a potential on-board layout of the 40 Gb/s C2C interconnect will eliminate the need for SOAs in the transmission lines, turning C2C energy consumption into a parameter that depends solely on the power requirements of the RM and its respective electronic driver, the PD-TIA and the external LD that feeds the RM with the CW optical beam. Considering the employment of state-of-the-art RM drivers [17] and assuming an LD with 6.1 dBm output power and a 10% wall-plug efficiency, the energy efficiency of the proposed 40 Gb/s C2C photonic link was estimated at 6.74 pJ/bit that increases to 7.04 pJ/bit when incorporating also state-of-the-art Serialization/Deserialization (SerDes) [21], assuming transmission line devices insertion and coupling losses as shown in Table 1.

These energy efficiency values suggest a 58.4% and 56.5% improvement, respectively, compared to the 16.2 pJ/bit link energy efficiency of Intel QPI [14].

In a practical case, where multiple sockets communicate simultaneously with other sockets, the AWGR channel crosstalk can be a limiting factor affecting the entire system performance. Given that the AWGR device used in our experiments was one of the few AWGR structures reported to operate in O-band [23], [24] and was designed for CWDM operation, making it rather challenging to achieve a low channel crosstalk, the system performance of the proposed AWGR-interconnect when operating in all-to-all communication scenaria could be improved via one of the following two approaches: a) To replace the employed AWGR module by a redesigned DWDM O-band AWGR with low-crosstalk characteristics, taking advantage of the improved crosstalk properties of DWDM structures compared to CWDM design, as this has been confirmed by respective C-band AWGR designs where DWDM layouts report on low-crosstalk values in the range of 16.8–25 dB [25]–[29], b) To enrich the system architecture by retaining the CWDM AWGR and equipping every socket with a set of seven unique wavelengths, while adding a DWDM AWG DEMUX at every AWGR output. In this way, all seven wavelengths reaching the same CWDM AWGR output port from seven different originating sockets can be adjusted to have slightly different wavelength values still residing in the same AWGR pass-band, so that finally the DWDM AWG DEMUX can separate them successfully to the destined sockets without suffering from any crosstalk issues. This approach of combining Coarse with Dense WDM AWG and AWGR structures for overcoming the crosstalk issues of the CWDM component has been already applied successfully in relevant wavelength routing experiments [30].

4. Conclusion

We have demonstrated an O-band C2C interconnection scheme for up to 40 Gb/s 8-socket direct communication in MSBs, exploiting integrated photonics for its transmitter, receiver and routing building blocks. Proof-of-concept validation employing an O-band carrier-depletion Si-based RM, a Si-based 8 $\times$ 8 cyclic AWGR and a packaged PD-TIA module was successfully demonstrated at 25 Gb/s and 40 Gb/s for 8 different routing scenarios, revealing in all cases clearly open eye diagrams. The 40 Gb/s line-rate capabilities together with the single-hop connectivity perspective among 8 nodes indicate a significant potential for reducing both energy and latency in “glueless” MSB setups when transferring the C2C interconnect architecture onto a 1300 nm optical polymer-waveguide hosting board, suggesting a potential energy gain of $>$58% compared to state-of-the-art QPI interconnects.

Acknowledgment

The PD-TIA was developed in part by the EU-funded FP7-ICT-DISCUS research project (No. 318137) and in part by the EU-funded FP7-ICT-MIRAGE research project (No. 318228).

References

[1] Cisco, “Cisco global cloud index: Forecast and methodology, 2016–2021,” San Jose, CA, USA, White Paper, Feb. 1, 2018. [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html
[2] T. Chen, X. Gao, and G. Chen, “The features, hardware, and architectures of data center networks: A survey,” *J. Parallel Distrib. Comput.*, vol. 96, pp. 45–74, 2016.
[3] J. Gripp, J. E. Simsarian, J. D. Le Grange, P. Bernasconi, and D. T. Neilson, “Photonic terabit routers: The IRIS project,” in *Proc. Conf. Opt. Fiber Commun., Collocated Nat. Fiber Opt. Engineers Conf.*, San Diego, CA, USA, 2010, pp. 1–3.
[4] R. Maddox, *Weaving High Performance Multiprocessor Fabric: Architectural Insights to the Intel QuickPath Interconnect.* Hillsboro, OR, USA: Intel Press, 2009.
[5] T. Wicki and J. Schulz, “Bixby: The scalability and coherence directory ASIC in Oracle’s highly scalable enterprise systems,” in *Proc. IEEE Hot Chips 25 Symp.*, Stanford, CA, USA, 2013, pp. 1–34.
[6] J. Ajanovic, “PCI express 3.0 overview,” in *Proc. IEEE Hot Chips 21 Symp.*, Stanford, CA, USA, 2009, pp. 1–61.
[7] G. T. Kanellos and N. Pleros, “WDM mid-board optics for chip-to-chip wavelength routing interconnects in the H2020 ICT-STREAMS,” *Proc. SPIE*, vol. 10109, Feb. 20, 2017, Art. no. 101090D.

[8] P. Grani, R. Proietti, S. Cheung, and S. J. Ben Yoo, "Flat-topology high-throughput compute node with AWGR-based optical-interconnects," *IEEE J. Lightw. Technol.*, vol. 34, no. 12, pp. 2959–2968, Jun. 2016.
[9] P. Grani, G. Liu, R. Proietti, and S. J. Ben Yoo, "Bit-parallel all-to-all and flexible AWGR-based optical interconnects," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, Los Angeles, CA, USA, 2017, pp. 1–3.
[10] S. Pitris *et al.*, "O-band energy-efficient broadcast-friendly interconnection scheme with SiPho Mach–Zehnder modulator (MZM) & arrayed waveguide grating router (AWGR)," in *Proc. Opt. Fiber Commun. Conf. Expo.*, San Diego, CA, USA, 2018, pp. 1–3.
[11] R. Yu *et al.*, "A scalable silicon photonic chip-scale optical switch for high performance computing systems," *Opt. Exp.*, vol. 21, no. 26, pp. 32655–32667, Dec. 2013.
[12] A. Sugama, K. Kawaguchi, M. Nishizawa, H. Muranaka, and Y. Arakawa, "Development of high-density single-mode polymer waveguides with low crosstalk for chip-to-chip optical interconnection," *Opt. Exp.*, vol. 21, no. 20, pp. 24231–24239, Oct. 2013.
[13] *IEEE Standard for Ethernet: Physical Layer Specifications and Management Parameters for 40 Gb/s and 100 Gb/s Operation Over Fiber Optic Cables*, IEEE Standard P802.3bm/D3.23.3 (Amendment of IEEE Standard 802.3-2012), Nov. 21, 2014, pp. 1–16.
[14] C. Gough, I. Steiner, and W. Saunders, *Energy Efficient Servers: Blueprints for Data Center Optimization*, 1st ed. Berkeley, CA, USA: Apress, 2015.
[15] D. Ziakas, A. Baum, R. Maddox, and R. Safranek, "Intel quickpath interconnect architectural features supporting scalable system architectures," in *Proc. 18th IEEE Symp. High Perform. Interconnects*, 2010, pp. 1–6.
[16] M. Pantouvaki *et al.*, "Active components for 50 Gb/s NRZ-OOK optical interconnects in a silicon photonics platform," *IEEE J. Lightw. Technol.*, vol. 35, no. 4, pp. 631–638, Feb. 2017.
[17] H. Ramon *et al.*, "Low-power 56 Gb/s NRZ microring modulator driver in 28 nm FDSOI CMOS," *IEEE Photon. Technol. Lett.*, vol. 30, no. 5, pp. 467–470, Mar. 2018.
[18] S. Pitris *et al.*, "Silicon photonic 8 × 8 cyclic arrayed waveguide grating router for O-band on-chip communication," *Opt. Exp.*, vol. 26, no. 5, pp. 6276–6284, Mar. 2018.
[19] B. Moeneclaey *et al.*, "A 40-Gb/s transimpedance amplifier for optical links," *IEEE Photon. Technol. Lett.*, vol. 27, no. 13, pp. 1375–1378, Jul. 2015.
[20] R. Dangel *et al.*, "Polymer waveguides enabling scalable low-loss adiabatic optical coupling for silicon photonics," *IEEE J. Sel. Topics Quantum Electron.*, vol. 24, no. 4, pp. 1–11, Jul. 2018.
[21] V. Stojanović *et al.*, "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes [Invited]," *Opt. Exp.*, vol. 26, no. 10, pp. 13106–13121, May 2018.
[22] M. Seifried *et al.*, "Monolithically integrated CMOS-compatible III-V on silicon lasers," *IEEE J. Sel. Topics Quantum Electron.*, vol. 24, no. 6, pp. 1–9, Nov. 2018.
[23] N. A. Idris and H. Tsuda, "6.4-THz-Spacing, 10-channel cyclic arrayed waveguide grating for T- and O-band coarse WDM," *IEICE Electron. Exp.*, vol. 13, no. 7, pp. 1–7, Apr. 2016.
[24] Y. Wu, T. Lang, and J. He, "Horseshoe-shaped 16 × 16 arrayed waveguide grating router based on SOI platform," in *Proc. Asia Commun. Photon. Conf.*, 2017, Art. no. S4J.3.
[25] J. Wang *et al.*, "Low-loss and low-crosstalk 8 × 8 silicon nanowire AWG routers fabricated with CMOS technology," *Opt. Exp.*, vol. 22, no. 8, pp. 9395–9403, Apr. 2014.
[26] S. Takenobu *et al.*, "All-polymer 8 × 8 AWG wavelength router using ultra low loss polymer optical waveguide material (CYTOP)," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Engineers Conf.*, 2008, pp. 1–3.
[27] S. Pathak, M. Vanslembrouck, P. Dumon, D. Van Thourhout, and W. Bogaerts, "Compact 16 × 16 channels routers based on silicon-on-insulator AWGs," in *Proc. Annu. Symp. IEEE Photon. Benelux Chapter*, 2011, pp. 101–104.
[28] G. Chen, J. Zou, T. Lang, and J. He, "Compact 4 × 4 1250 GHz silicon arrayed waveguide grating router for optical interconnects," *Proc. SPIE*, vol. 9367, Feb. 27, 2015, Art. no. 936717.
[29] G. Song, J. Zou, and J. He, "Ultra-compact silicon-arrayed waveguide grating routers for optical interconnect systems," *Chin. Opt. Lett.*, vol. 15, no. 3, p. 030603, 2017.
[30] N. A. Idris *et al.*, "Full-mesh T- and O-band wavelength router based on arrayed waveguide gratings," *Opt. Exp.*, vol. 24, no. 1, pp. 672–686, Jan. 2016.