DOI: 10.1364/OFC.2020.W3G.2

Automated Thermal Drift Compensation in WDM-based Silicon Photonic Multi-Socket Interconnect Systems

M. Moralis-Pegios, F. Zanetto, E. Guglielmi, V. Grimaldi, K. Fotiadis, S. Pitris, T. Alexoudi, P. De Heyn, J. Van Campenhout, D. Aguiar, G. Ferrari, M. Sampietro, A. Melloni and N. Pleros

[1]*Department of Informatics, Center for Interdisciplinary Research & Innovation, Aristotle University of Thessaloniki, Thessaloniki, Greece.*
[2]*imec, Kapeldreef 75, Leuven B-3001, Belgium*
[4]*Dipartimento di Elettronica Informazione e Bioingegneria - Politecnico di Milano, Milano, 20133 Italy.*
mmoralis@csd.auth.gr

Abstract: We present an on-chip AWGR-based interconnect system with automated thermal drift compensation along cascaded resonant structures in a dual-socket layout. Error-free operation in a 30 Gb/s data-routing scenario within a 12C temperature range is demonstrated.

1. Introduction

The compute and networking requirements, imposed by the explosive growth of Data Center (DC) traffic and the dominance of East-West communication patterns in modern DCs [1], has focused research interest on novel Multi-socket-board (MSB) interconnect technologies than can offer communication in a low energy and latency envelope. However, the currently offered solutions face a significant latency-radix trade-off. Peer-to-Peer schemes such as Intel' QPI') [2] can offer low-latency interconnection for a maximum number of 8 sockets, while switch-based layouts, like Bixby [3], can support higher radixes but suffer from critical energy and latency restrictions. Arrayed Waveguide Grating Router (AWGR) based optical architectures have shown the potential to overcome these limitations and offer any-to-any, low latency and low-energy communication among more than 8 nodes, by exploiting wavelength division multiplexing (WDM) and the cyclic wavelength routing properties of AWGRs [4]. In this context, experimentally demonstrated AWGR-based interconnects for MSB layouts have suggested dramatic energy reductions of up to 63.3% compared to QPI in a 8-socket configuration [4], while simulation-based architecture benchmarking reveals 3× energy savings, higher throughput and lower latency compared to the respective electronic MSB baseline [5].

Demarcating, however, from lab-based demonstrations towards realistic MSB deployments, the unstable DC thermal environment forms a significant challenge, as all previous demonstration rely on Silicon Photonic platforms with a rather high thermal coefficient and employ a sequence of cascaded resonant WDM structures to allow optical data modulation and multiplexing, wavelength-based routing and demultiplexing. To this end, reaping the appealing energy, throughput and latency benefits offered by AWGR-based MSB interconnect schemes necessitates an automatic monitoring and stabilization system for safeguarding resonance alignment against any environmental thermal fluctuations. Until now, research efforts resorted to multi-point light tapping [6], that effectively limits, however, the system's scalability by introducing significant optical losses in the system. Recently, the ContactLess Integrated Photonic Probe (CLIPP) system was introduced as a novel solution for non-invasive power monitoring on integrated waveguides, with its credentials having been validated in add/drop ring resonators [6] and silicon switch fabrics [8]. However, its performance has not yet been validated in multi component WDM AWGR-based layouts.

In this paper, we present for the first time to our knowledge, a CLIPP-equipped AWGR-based interconnect system that encompasses an automated feedback- control monitoring mechanism for tuning and locking the resonances of its cascaded constituent building blocks. The proposed CLIPP-assisted AWGR-based optical interconnect employs a control plane based on a CLIPP system and integrated heater control and has been evaluated in an on-chip dual-socket Silicon Photonic (SiPho) layout that comprises two Ring Modulators (RMs) followed by a 2nd order ring-based Multiplexer (MUX) module that is connected to a 16×16 DWDM AWGR optical engine subsequently connected to the 2-sockets through a similarly designed demultiplexer (DEMUX). Experimental verification of the CLIPP-assisted AWGR interconnect scheme revealed successful resonance locking of 3 different SiPho components devices in 30 ms, while the performance of the system was assessed in a 30 Gb/s NRZ routing scenario for on-chip temperature variations from 24 C to 36 C. Error-free operation was achieved across the whole 12 °C temperature range with a low power penalty of approximately 1 dB at BER 5×10^{-11}.

2. SiPho system demonstrator and experimental setup

To assess the performance of the CLIPP system, when employed in an AWGR-based architecture, a Silicon Photonic chip facilitating all the necessary building blocks for optical communication between two multi-processor

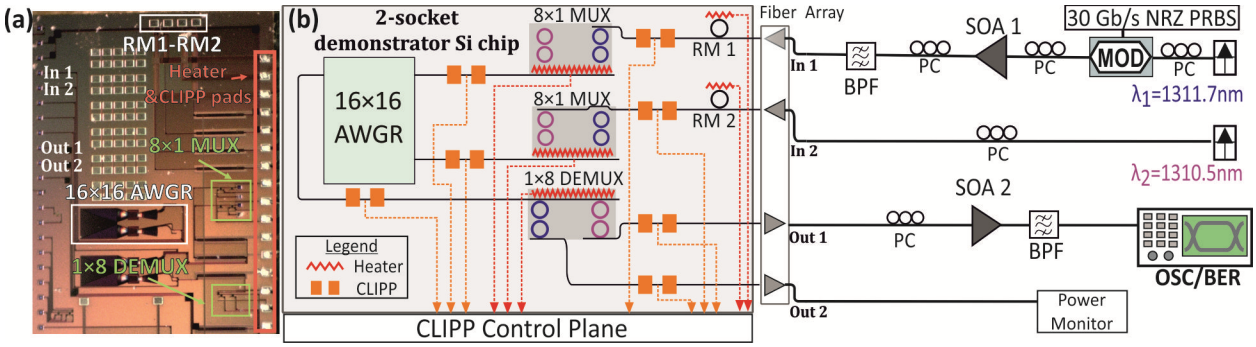

Fig. 1(a) Microscope photo of the 2-socket demonstrator chip (b) Schematic and experimental setup used for the 2-socket demonstrator.

modules (MCM), herein defined as sockets, was designed and fabricated in IMEC's ISIPP50G platform. A microscope photo of the resulting 2-socket demonstrator chip is illustrated in Fig. 1 (a). The demonstrator incorporates a 16×16 O-band DWDM AWGR, with a channel spacing of 1.10 nm and an FSR of 16.73 nm, to allow non-blocking communication between two socket cells, through wavelength-based routing. Each socket cell comprises a micro ring modulator (RM), featuring a Q factor of ~5000 and a 3-dB bandwidth of 33.8 GHz and is interconnected to the AWGR optical engine through a 1×8 MUX, based on 2nd order ring resonators, with an FSR of 9.5 nm and a channel spacing of 1.2 nm. The optical signal originating at the AWGR output is demultiplexed in an 8×1 DEMUX, identical to the one employed at the AWGR's input, with the resulting signals waveguided to the respective socket. Optical input/output access to the Si chip is achieved via TE-polarization grating couplers (GC) at the west side of the chip, denoted as In 1, In 2 and Out 1, Out 2, respectively. Integrated heaters allow thermal tuning of the RM, MUX and DEMUX resonances, while CLIPP sensors are employed after each resonant structure to monitor the optical power across the system. Finally, the CLIPP sensors and the integrated heaters are accessible through electrical DC pads at the west side of the chip and are wire-bonded to the electronic interface board.

Fig. 1(b) shows the experimental setup employed for the characterization and the 30 Gb/s routing scenario of the 2-socket demonstrator chip. A bit pattern generator was used to generate an NRZ pseudo-random binary sequence (PRBS7-1) at 30 Gb/s, with a peak amplitude of 300 mV, that was amplified by a commercial high-speed RF driver amplifier to drive a 30 GHz O-band MZM modulator. Light from a tunable laser source (TLS) (λ1=1311.7 nm), was injected to the MZM modulator, with the resulting modulated signal after amplification in a semiconductor optical amplifier (SOA) and filtering in a 3 nm 3-db bandwidth optical band pass filter (OBPF), coupled to the In1 port of the Si chip, through a GC. A distributed feedback laser (DFB) was used to generate a continuous wave beam (λ2=1310.5 nm), similarly coupled to the chip In2 port. The optical signal originating from the modulated input and routed to Out 1 port was coupled out of the chip through a GC and after amplification in an SOA and filtering in a 0.5-nm 3-db bandwidth OBPF, was monitored at a sampling scope and evaluated at a Bit Error Rate Tester (BERT). The optical signal originating from the CW input was similarly coupled out of the chip and monitored in an optical power meter. The average optical power of the modulated and CW signals injected were 10 dBm and 8 dBm respectively, while the average optical power of the signals emerging at output ports Out1 and Out2 were -12.6 dBm and -14.2 dBm respectively. Breaking down the optical losses, the two GC's imposed 10 dB at1310 nm, the two MUX 5 dB and the AWGR 7-8 dB, resulting in total optical losses of 22.6 dB and 22.2 dB for the two optical signals. The two SOAs were electrical driven at 150 mA and 200 mA, providing a gain of 10 dB and 18 dB respectively.

The photonic chip was finally mounted on a compact interface board, depicted in Fig. 2(a), to facilitate optical coupling and electrical access to the CLIPP sensors and integrated heaters. The board is subsequently connected to the FPGA-based control platform that reads all CLIPPs in parallel, drives all heaters in parallel and operates the

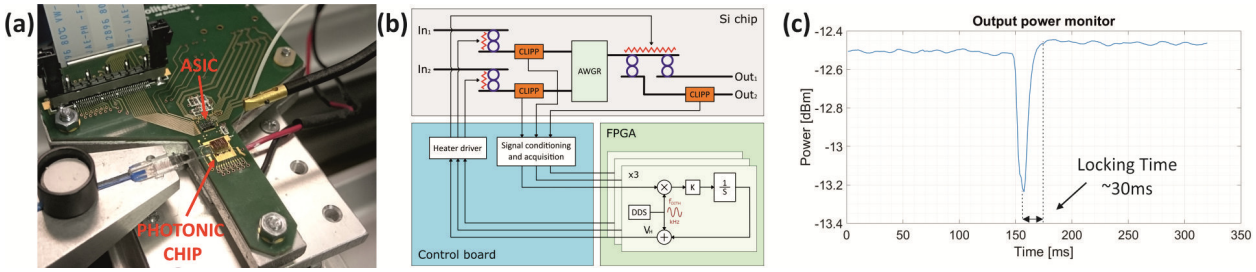

Fig.2 (a) Electronic board for interconnection to the SiPho chip (b) Schematic of the control architecture (c) Response to a 150 pm wavelength shift.

most suitable strategy in real time to perform the closed-loop control of the system, providing reconfigurability and flexibility to the system. Fig. 2 (b) depicts the locking procedure of the control system, incorporating the CLIPPs sensor and the integrated heater control. To lock the working point of the ring based MUXes and DEMUXes and counteract unavoidable thermal fluctuations that otherwise would impair the system functionality, a real time locking was implemented based on the dithering technique. According to it a small modulation signal, with 5mV peak power and frequency of 6 kHz, is applied to the heater of each device in order to extract the derivative of its transfer function. This information, used as the error signal of the control loop, is minimized by an integral controller that drives the heater to keep the devices at resonance. In this way, a robust and power independent feedback loop is implemented. The bandwidth of the loop has been set to about 20 Hz to minimize noise yet allow a fast recovery time within less than 30 ms as illustrated in Fig.2 (c).

3. Experimental assessment of system demonstrator

The performance of the proposed optical interconnection scheme, incorporating the CLIPP sensor-based feedback control, was assessed in a 30 Gb/s routing scenario with both optical paths locked, that encompassed varying thermal conditions, emulating the temperature instabilities imposed on an optical chip in a DC environment. To this end, the optical chip was mounted on a thermo-electric cooler that allowed fine thermal control of the Si chip, while the quality of the optical link was monitored through continuous Bit Error Rate (BER) measurements in a Real Time Oscilloscope (RTO), acquired optical eye diagrams and BER curves obtained at a BERT at certain temperature points. Fig.3 (a) depicts the imposed variations on the chip's temperature over a 20-minute long thermal stability test, featuring values from 24 °C to 36 °C, with a time step of 1 second. Fig. 3 (b) illustrates the variations of the MUX heater voltage, located at the optical path of the modulated optical signal prior entering the AWGR, that arise from the feedback control of the CLIPP based system. It should be noted, that the heater voltage variations, are mirroring the temperature variations, denoting the feedback system's response precision. Fig. 3 (c), illustrates the BER evolution during the 20-min experiment, revealing almost zero detected errors for a temperature variation of 12 C, with an instantaneous loss of error-free operation. The obtained BER measurements at the BERT along with the captured eye diagrams are illustrated in Fig. 3 (d). Error-free operation at 5×10^{-11} error-rate for all cases is reported across different temperature operating points, with a power penalty variation of 0.3 dB and a maximum power penalty of ~1 dB when compared with the input signal. Moreover, the effect of locking the second optical signal was also assessed through BER measurements, revealing an almost negligible power penalty of ~0.4 dB.

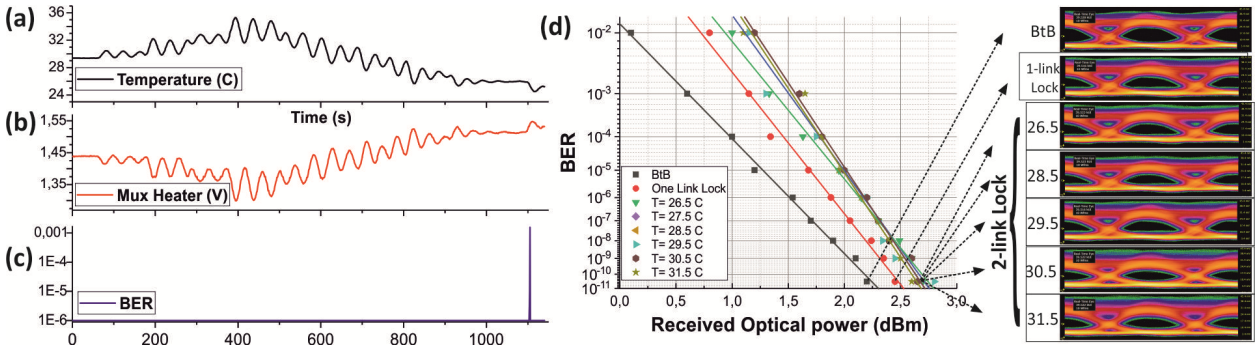

Fig. 3 (a) On-chip temperature variation (b) Mux heater voltage (c) BER measured at an RTO, obtained during 20-min thermal stability Test (d) BER curves and acquired optical eye diagrams

Acknowledgements

This work was supported by the EC H2020 ICT-STREAMS (688172) and NEBULA (871658) projects.

References

1. Cisco, "Cisco Global Cloud Index: Forecast and Methodology, 2015–2020," White Pap. (2015).
2. R. Maddox, Weaving High Performance Multiprocessor Fabric: Architectural Insights to the Intel QuickPath Interconnect , Intel Press, (2009).
3. T. Wicki, et al. "Bixby: The scalability and coherence directory ASIC in Oracle's highly scalable enterprise systems," Hot Chips ,pp 1-34, (2013).
4. S. Pitris et al., "A 40 Gb/s Chip-to-Chip Interconnect for 8-Socket Direct Connectivity Using Integrated Photonics", Photon. Journal,10(2018)
5. P. Grani et al.,"Flat-Topology High-Throughput Compute Node With AWGR-Based Optical-Interconnects," JLT, 34(12), 2959-2968, (2016)
6. A. Gazman et al.,"Automated Thermal Stabilization of Cascaded Silicon Photonic Ring Resonators for Reconfigurable WDM Applications," ECOC, pp.. 1-3 (2017).
7. D Aguiar et al. "Automatic tuning of microring-based hitless reconfigurable add-drop filters", OFC (2018).
8. A.Annoni et al. "Automated routing and control of silicon photonic switch fabrics", IEEE JSTQE **22**(6) pp.169-176 (2016).
9. S. Grillanda et al. "Non-invasive monitoring and control in silicon photonics using CMOS integrated electronics", Optica, 1(3), 129-136 (2014)