

Temperature Tolerant On-Chip WDM Silicon Photonic Transmitter and AWGR-based Routing Interconnects

I. Roumpos¹, T. Chrysostomidis¹, V. Grimaldi², F. Zanetto², F. Toso², P. De Heyn³, Y. Ban³, J. Van Campenhout³, G. Ferrari², M. Sampietro², F. Morichetti², A. Melloni², K. Vysokinos¹, T. Alexoudi⁴, N. Pleros⁴, and M. Moralis-Pegios⁴

¹Department of Physics, Center for Interdisciplinary Research & Innovation, Aristotle University of Thessaloniki, Thessaloniki, Greece.

² Department of Electronics, Information and Bioengineering- Politecnico di Milano, Milano, 20133 Italy.

³imec, Kapeldreef 75, Leuven B-3001, Belgium.

⁴Department of Informatics, Center for Interdisciplinary Research & Innovation, Aristotle University of Thessaloniki, Thessaloniki, Greece.
ioroumpo@auth.gr

Abstract: We demonstrate automated thermal drift compensation in a two-socket AWGR interconnect, incorporating a ring-modulator transmitter. Stable operation with an average $Q=5.8$ over a range of 9°C is achieved for 25 Gb/s on-chip modulated data.

1. Introduction

The relentless growth of data center (DC) traffic [1], fueled by the proliferation of cloud computing and AI workloads, is putting tremendous strain on the required processing speed and power-envelope of the widely employed Multi-Socket Boards (MSB). In this context, optical architectures employing Arrayed Waveguide Grating Router (AWGR) have been gathering pace during the last years, as a potential technological candidate that can break through the radix and power consumption limitations of current deployments, by leveraging wavelength-based routing to allow for high-port count and low latency interconnectivity [2]-[5]. Previous demonstrations have already showcased the benefits of such architectures, allowing interconnection of up to 8 sockets with 4.2 pJ/bit energy efficiency [5] and revealing a 5x reduction of execution time in multi-GPU systems [2].

Despite these impressive credentials, monolithic deployment of AWGR-based interconnection architectures necessitates the co-integration of several supporting photonic components, such as optical data modulators, multiplexers (MUX) and demultiplexers (DEMUX), inevitably imposing the challenge of tuning their transfer functions in the optimum operating point. This is exaggerated when considering architectures based on the highly power efficient and low footprint photonic ring resonators (RR) [6], that are yet highly sensitive to the temperature fluctuations inherent in a DC environment. In this context, a ContactLess Integrated Photonic Probe (CLIPP) -based system [7] was previously utilized for stabilizing the operating point of three constituent RR-based MUX and DEMUX blocks of an AWGR-based photonic system, validating its functionality in safeguarding the photonic building blocks optimum operating point over a MUX-AWGR-DEMUX resonance chain. Incorporating, however, Ring Resonator Modulator (RRM) -based transmitter function in the resonant element chain in order to complete an on-chip transmitter/routing interconnect system for multi-socket applications [2]-[5] requires a more sophisticated approach, as the locking near the optimum operating point of the RRM cannot be sustained through a simple optical power monitoring mechanism [8].

In this paper we present to the first time to our knowledge a temperature tolerant WDM Silicon Photonic transmitter and AWGR-based routing interconnect system by enforcing a dual control CLIPP-based feedback circuitry over a chain of RRM-MUX-AWGR resonant elements. The control system safeguards the working point of each structure against thermal and wavelength fluctuations, utilizing a novel strategy that exploits the 2nd derivative of the RRM transfer function to lock the RRM to a near-optimum operating point. The functionality of the interconnection scheme was evaluated in a 25 Gb/s two-socket modulation and interconnection routing scenario, during a 15-min long thermal stress test, with imposed temperature variations ranging from 27 °C to 36 °C. The quality of the interconnection was validated through eye-diagrams measurements, revealing an average Q factor of 5.8, that correspond to a theoretical bit-error-rate of 1.82×10^{-9} .

2. Experimental setup, system architecture and control strategy

The performance of the two-socket, dual control AWGR interconnection system, was evaluated in a previously demonstrated SiPho chip [7], depicted in Fig. 1(a). The chip includes all the necessary building blocks for interconnecting two-processor modules, hereafter defined as sockets. At the core of the interconnection system lies a 16×16 AWGR-based routing engine, with a channel spacing of 1.10 nm and an FSR of 16.73 nm, that allows time-of-flight, non-blocking connectivity between the two sockets. Each socket is equipped with a RRM with a 3-dB




Fig. 1(a) Microscope photo of the SiPho chip and experimental setup (b) schematic of the photonic system and control platform

bandwidth of 33.8 GHz and a Q factor of 5000, that is followed by a 1x8 MUX, consisting of two 2nd order RR with a channel spacing of 1.2 nm and FSR of 9.5 nm. Interconnection to different output sockets is achieved through the AWGR-based router by controlling the transmission wavelength, with the signal traversing the AWGR emerging at a different output for each transmission wavelength. Optical access to the SiPho chip is achieved through PDK-based Grating Couplers (GCs) at the input and output of the communication system, each introducing approximately 4.5 dB of optical loss in the 1310 nm transmission window.

The experimental system employed for the evaluation of the 25 Gb/s two-socket dual control AWGR-based system is illustrated in Fig. 1(a). Two light beams at $\lambda_1=1310.3$ nm and $\lambda_2=1311.3$ nm, originating from a tunable laser source and a DFB laser respectively, were injected at the input optical ports of Socket 1(In1) and 2 (In2). Each CW signal traverses a RRM, that is responsible of superimposing the electrical data to the optical signal. For this demonstration only the signal from the 1st socket was modulated on-chip, due to design limitations that constrained access to the RF pads of the 2nd socket's RRM. An Arbitrary Waveform Generator (AWG) was used to generate a 25 Gb/s NRZ pseudo-random binary sequence (PRBS-9) with a pattern length of 512 bit and a peak-to-peak amplitude of 550 mV. The electrical data stream was amplified by an RF amplifier, combined with a DC voltage of -2.5 V via a Bias-T and finally applied to RRM1 by a 40 GHz GS probe tip. Both the CW and the modulated data signals emerging at Output ports 1 and 2, were coupled out of the chip via GCs. The CW signal originating from Socket 2 at Out2, was evaluated in terms of power fluctuations, with an optical power meter. The modulated signal originating from Socket 1 at Out1 after amplification in a PDFA and filtering in an optical bandpass filter (OBPF) with a 3-dB bandwidth of 1nm, was injected in a 70 GHz photodiode and monitored by a sampling scope (OSC). The two optical paths introduce an optical loss of ~ 19 dB and ~ 25 dB respectively, that can be broken down to: (i) a common part of 19 dB with 9 dB from Input/Output GCs, 2.5 dB by the 1x8 MUX and 7.5 dB from the AWGR (ii) an extra loss introduced by the RF modulation of RRM1, ~ 6 dB during locking. The schematic diagram, in Fig. 1(b), illustrates the electrical interface points of the CLIPP-based control circuit. Each resonance-based structure is equipped with an integrated heater, that enables thermal resonance tuning, while CLIPP sensors are placed after each component to monitor the traversing optical power. Heaters and CLIPP sensors are subsequently connected to the FPGA-based control device that acquires the signal from all the sensors and drives all the heaters. The board is equipped with a thermoelectric cooler (TEC) and a thermistor, allowing real time monitoring and control of the average temperature of the chip.

Real time resonance locking of both the RRM1 modulator and the two 1x8 MUXes, was accomplished through a dual control strategy that enables both maximum-output power operation for the MUXes, and near-optimum operating point for the RRM. Fig. 2(a) depicts the transfer function of the 1x8 MUX of Socket1, derived by sweeping its




Fig. 2(a) MUX CLIPP signal with dithering signal vs heater power (b) Transfer function of MUX locked at the AWGR response, (c) RRM1 CLIPP signal and second harmonic amplitude as function of heater power (d) Transmission penalty at RRM1 as function of wavelength shifting.

integrated heater, along with the amplitude of the output dithering signal measured by the CLIPP sensor [7]. Locking the resonance of the MUX to the zero of dithering signal, results in the tuning of the MUXs resonance to the one of the AWGR, with the two overlapping transfer functions depicted in Fig. 2(b). The same control strategy cannot be applied to the RRM, as in this case, the optimum operating point is not located at a minimum or maximum of the transfer function [6]. Instead, we developed a novel control strategy that relies on the use of the 2nd derivative of the CLIPP signal, to lock RRM1 to a near-optimum operating point. Fig. 2(c) shows the RRM1's transfer function along with its 2nd derivative, acquired by synchronously demodulating the CLIPP readout at twice the dithering frequency. Locking the RRM1 to the zero of this signal, results in an almost optimum operating point, as illustrated in Fig. 2(d), where we plot the RRM1 transmission penalty and highlight the position of the locking point. Moreover, as it can be seen in Fig. 2(d), the use of a 2 mV dithering does not interfere with the transfer function of RRM1, as the transmission penalty curves in the 0 mV case that represents the un-controlled status, and the 2 mV referring to the controlled status are identical. Finally, it should be noted, that this technique does not require any calibration and it is independent of fluctuations of the input power, two relevant advantages especially in complex systems with multiple devices.

3. 25 Gbps routing scenario

The performance of two-socket, dual-control strategy interconnection scheme was evaluated in a 25 Gb/s routing scenario. In order to demonstrate the closed-loop controller ability to compensate the thermal drifts, an external temperature fluctuation was applied on the chip through the on-board TEC. Fig. 3(a) depicts the temperature variation on the photonic system during a 15-minute-long thermal stress test, with an acquisition step of 1 sec, and a temperature range of 9 °C (27 °C to 36 °C). Fig.3 (b) illustrates the time evolution of RRM1's heater and the two MUXes heaters voltages, during the testing time window, revealing the system's response to the temperature variation. Fig.3 (c) depicts the monitored optical power of the CW signal at Out2 of the AWGR revealing a small deviation of ~ 2 dB originating from the transfer function of the AWGR that could not be tuned during the experiment. Fig.3 (d) shows the optical power of the modulated signal at AWGR out1 with the corresponding Q factor values, extracted from the captured eye diagrams, in each recorded temperature. As it can be observed from the graphs at Fig. 3(c) and (d) the power is decreasing at the edge temperatures, again due to the transfer function of the AWGR. The measured Q values ranged from 4.6 to 6.8, with an average value of 5.8, validating the correct functionality of the locking mechanism. Fig. 3(e) illustrates the captured eye diagrams for different set temperatures during the thermal stress test, further confirming the high precision of the locking strategy, even when incorporating two MUXes and an RRM.




Fig. 3 (a) Temperature variations on the photonic system (b) Voltage variations of ring modulator (RRM1) and two MUXes heaters (c) Monitored optical power at AWGR Out2 (d) Monitored optical power at AWGR Out1 with the corresponding Q factor values at each temperature (e) Acquired eye diagrams at five different temperatures

Acknowledgments

This work was supported by the EC H2020 NEBULA (871658) project.

References

- [1] Cisco, "Cisco Annual Internet Report (2018–2023)," White paper (2020).
- [2] M. Moralis-Pegios *et al.*, "Silicon circuits for chip-to-chip communications in multi-socket server board interconnects," *IET OPTOELECTRONICS*, vol. 15, (2021).
- [3] T. Alexoudi *et al.*, "Optics in Computing: From Photonic Network-on-Chip to Chip-to-Chip Interconnects and Disintegrated Architectures," in *JLT*, vol. 37, (2019).
- [4] M. Fariborz *et al.*, "Silicon Photonic Flex-LIONS for Reconfigurable Multi-GPU Systems," in *JLT*, vol. 39, (2021).
- [5] S. Pitris *et al.*, "A 40 Gb/s Chip-to-Chip Interconnect for 8-Socket Direct Connectivity Using Integrated Photonics," *Phot. Journal*, 10(2018).
- [6] M. Pantouvaki, *et al.*, "Active Components for 50 Gb/s NRZ-OOK Optical Interconnects in a Silicon Photonics Platform," in *JLT* 35 (2017).
- [7] F. Zanetto *et al.*, "WDM-Based Silicon Photonic Multi-Socket Interconnect Architecture With Automated Wavelength and Thermal Drift Compensation," in *JLT*, vol. 38 (2020).
- [8] H. Li *et al.*, "A 3D-Integrated Silicon Photonic Microring-Based 112-Gb/s PAM-4 Transmitter With Nonlinear Equalization and Thermal Control," in *IEEE JSSC*, vol. 56 (2021).
- [9] M. Moralis-Pegios *et al.*, "52 km-Long Transmission Link Using a 50 Gb/s O-Band Silicon Microring Modulator Co-Packaged With a 1V-CMOS Driver," in *Phot. Journal*, vol. 11(2019).