FPGA-Programmable 512 Gbps 2.5D DWDM Photonic Network Interface Card

James Robinson,[1,*] Maarten Hattink,[1] Liang Yuan Dai,[1] Max Haimowitz,[1] Mao Li,[1] Daniel Jang,[1] Padraic Morrissey,[2] Arun Kumar Malik,[2] Peter O'Brien,[2] Seth Robertson,[3] Mingoo Seok,[1] and Keren Bergman[1]

[1] *Dept. of Electrical Engineering, Columbia University in the City of New York, New York, New York, 10027, USA*
[2] *Photonics Packaging Group, Tyndall National Institute, University College Cork, Cork, T12R5CP, Ireland*
[3] *Peraton Labs, Basking Ridge, NJ, 07920, USA*

[*] *james.robinson@columbia.edu*

Abstract: We present a 2.5D integrated, FPGA-programmable silicon photonic network interface card implementing a bidirectional 512Gbps dense wavelength-division multiplexed link. We demonstrate sixteen transmitter and eight of sixteen receiver channels operating error-free at 16 Gbps each.

1. Introduction

Silicon photonic technologies promise dramatic improvements to network performance in data-center and super-computing environments, including improvements to training speeds for machine learning and artificial intelligence applications [1] [2]. Photonic devices have seen major improvements in performance over the past two decades [3] [4], but there remains a need for hardware to conduct system-level experiments to perform network-scale validation of existing simulations. We report on our Optical Network Interface Card (ONIC), a 2.5D integrated, FPGA configurable, 512 Gbps dense wavelength-division multiplexed (DWDM) transceiver capable of addressing this need, as well as demonstrate a representative examination of its bit error rate (BER) characteristics.

2. ONIC and Package Description

The ONIC consists of a custom PCB integrated into a larger mechanical module, depicted in Fig. 1a, which also contains an FPGA development board, power, and cooling. The module provides fiber interfaces to connect to an optical fabric and can be connected to a PCIe interface via an extension cable. The PCB itself consists of a 12-layer PCB hosting a photonic integrated circuit (PIC), shown in Fig. 1b, and four electronic integrated circuits (EICs). All five custom ICs are flip-chip soldered on a 17-layer BGA ceramic interposer, arranged as shown in Fig. 1c. A cross section of the full 2.5D integration is shown in Fig. 1d. The PCB provides high-speed data paths to the interposer, low-speed thermal control circuitry, and voltage regulators.

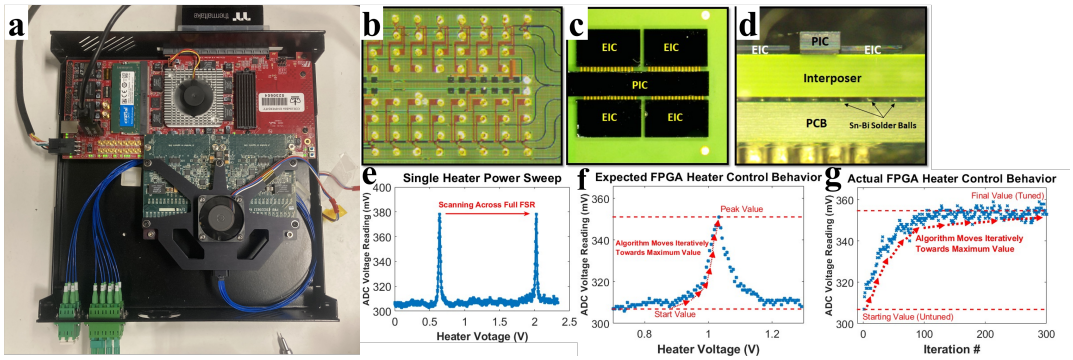

Fig. 1. a) ONIC package, including the custom PCB (green), FPGA (red), PCIe connector (top), and optical ports (bottom left). b) PIC with solder balls. c) PIC and EICs flip-chipped onto the interposer. d) Side view of the 2.5D integration. e) FPGA-controlled resonator heater control. f) Expected automated power peak-finding behavior. g) Actual automated peak-finding and peak-tracking behavior.

The FPGA serves multiple purposes: providing an interface between standard PCIe-based hardware and the optical network, programming and configuring the PIC and EICs, implementing the PIC's thermal control algorithms, provision of internal test data sequences, and potentially offloading some compute tasks. The PIC supports a bidirectional bandwidth of 512 Gbps through two independently routable 256 Gbps optical links. Each link is composed of 16 wavelength channels, implemented through a cascaded array of disk modulators, ring filters for demultiplexing, and integrated photodiodes. Each disk modulator's radius is designed for a specific wavelength to minimize required thermal tuning and closely match resonances to pre-selected laser wavelengths. The modulators exhibit a high extinction ratio of approximately 21 dB and a total insertion loss of 5.5dB for all 16 modulators per bus. The four EICs, each servicing one half-link, contain programmable analog front ends, transimpedance amplifiers, DC current cancellation loops, amplification stages, and clock and data recovery circuits.

3. Results

Several ONIC prototypes were assembled to validate the high speed transmitter and receiver capabilities, as well as the thermal control capabilities. The thermal control capabilities, were demonstrated with the main FPGA-ONIC package, Fig. 1a. First, broad-range thermal tuning of greater than one free spectral range was demonstrated, Fig. 1e, showing the ability to dynamically reconfigure the photonic resonators to various wavelengths as demanded by a given compute node or task, and enabling broad-temperature range thermal control. Next, in Fig. 1f, the expected performance of the peak-power finding algorithm is shown on an actual measured heater sweep, illustrating the basic mechanics of the resonator finding an assigned optical wavelength and maximizing the power observed. Last, Fig. 1g demonstrates an actual automated example of this FPGA-programmable peak-finding and peak-tracking behavior, in this case showing the power increasing from a low initialization value to a higher value, which is then maintained over time.

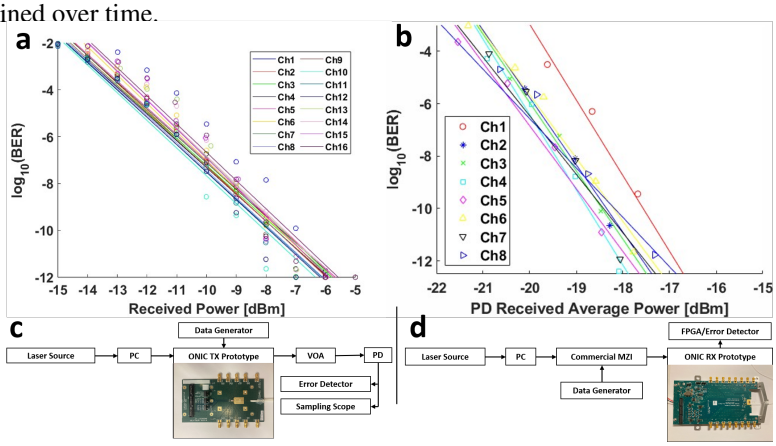

Fig. 2. a) BER performance of the modulators. b) BER performance of the receivers. c) System schematic for the modulator experiments. d) System schematic for the receiver experiments.

The transmitter and receiver high speed capabilities were validated with dedicated transmitter and receiver test boards, with the test setups shown in Fig. 2c-d including images of these prototypes. Sixteen wavelengths in the C-band, matching the ONIC, were selected and used to demonstrate the BERs of the modulators. For the receiver, eight of the 16 wavelengths were used to validate the filtered signals. To demonstrate modulator performance, 16 Gbps signals were modulated onto one of these sixteen wavelengths using the prototype board and received using a Thorlabs RXM40AF photodiode/TIA module, with data generation and BERT capability provided by an MP1900A PPG. For the receiver experiments, a Thorlabs MX35E MZI optical transmitter, with the aforementioned 16 Gbps PPG signal, was used to generate a modulated optical signal, which was then received using the ring filters and photodiodes integrated into the prototype package, a high-speed TIA 2.5D integrated with the receiver, and an FPGA for BERT. The modulators operated error free across all wavelengths at all received power levels at or above -5dBm and good performance at lower power, as shown in Fig.2a. The receivers operated error free on all wavelengths at or above -17dBm received power, as shown in Fig. 2b.

4. Conclusion

We present a 2.5D integrated, FPGA programmable silicon photonic network interface card implementing a bidirectional 512 Gbps dense wavelength-division multiplexed link. We demonstrate sixteen transmitter and eight receiver channels operating error-free at 16 Gbps per wavelength. These results provide an important step in the development of the hardware necessary for large-scale practical network experiments using high-bandwidth integrated silicon photonic links. In particular, the presented system offers powerful utility in the verification of existing simulation work that suggests substantial performance improvements of data-center and super-computer systems when training large machine learning or artificial intelligence models.

References

1. F. Douglis *et al.*, "FLEET—Fast Lanes for Expedited Execution at 10 Terabits: Program Overview," IEEE Internet Comput. **25**, 79–87 (2021).
2. M. Wade *et al.*, "Teraphy: a chiplet technology for low-power, high-bandwidth in-package optical i/o," IEEE Micro **40**, 63–71 (2020).
3. S. Y. Siew *et al.*, "Review of silicon photonics technology and platform development," J. Light. Technol. **39**, 4374–4389 (2021).
4. N. Abrams *et al.*, "Silicon photonic 2.5d multi-chip module transceiver for high-performance data center," J. Light. Technol. **38**, 3346–3357 (2020).

Acknowledgements: This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA) under the FastNICs program and by Department of Energy (DOE) Small Business Innovation Research (SBIR) ASCR Program under contract DE-SC0017182. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense, DOE, or the U.S. Government.