

An Error-free 1 Tbps WDM Optical I/O Chiplet and Multi-wavelength Multi-port Laser

M. Wade[†], E. Anderson, S. Ardalan, W. Bae, B. Beheshtian, S. Buchbinder, K. Chang, P. Chao, H. Eachempatti, J. Frey, E. Jan, A. Katzin, A. Khilo, D. Kita, U. Krishnamoorthy, C. Li, H. Lu, F. Luna, C. Madden, L. Okada, M. Patel, C. Ramamurthy, M. Raval, R. Roucka, K. Robberson, M. Rust, D. Van Orden, R. Zeng, M. Zhang, V. Stojanovic, F. Sedgwick, R. Meade, N. Chan, J. Fini, B. Kim, S. Liu, C. Zhang, D. Jeong, P. Bhargava, M. Sysak, C. Sun

Ayar Labs Inc., 3351 Olcott Street, Santa Clara, CA 95054, USA

[†]mark@ayarlabs.com

Abstract: We demonstrate 128 Gbps/port ($8\text{-}\lambda \times 16\text{ Gbps}/\lambda$) natively error-free transmission across eight optical ports using a 8-port, $8\text{-}\lambda$ /port WDM remote laser source and a pair of monolithically integrated CMOS optical I/O chiplets with 4.96-5.56 pJ/bit optical Tx+Rx chiplet energy efficiency.

© 2021 The authors.

Electrical link technologies face challenges in scaling bandwidth while achieving compelling energy efficiency, bandwidth density, latency, and reach. While digital processing performance continues to scale in how much compute can be achieved per unit area of silicon, for many high-value workloads, useful compute power is limited by the ability to fuel processing cores with I/O bandwidth. For these applications, existing optical solutions that have evolved from the ethernet networking market are a mismatch in terms of power, bandwidth density, and cost. Thus, a new I/O technology is needed that can match the performance requirements for computing fabrics and have a viable path to scale to high-volumes.

In this paper, we present link characterization results from a new optical I/O architecture for computing fabrics. The architecture is comprised of an external multi-port WDM laser source (SuperNova) that is disaggregated from the bandwidth dense, CMOS optical I/O chiplet (TeraPHY) [1, 2]. SuperNova aggregates multiple wavelengths of continuous-wave light onto multiple output fibers, and the TeraPHY chiplet contains multiple full-duplex WDM Tx/Rx optical ports that contain all of the electronics circuitry (SerDes, drivers, TIA, equalization, digital control, etc.) and photonic components (waveguides, modulators, photodetectors, etc.) required for full WDM link operation. Additionally, there is a wide-parallel electrical interface on the chip to connect to host ASICs.

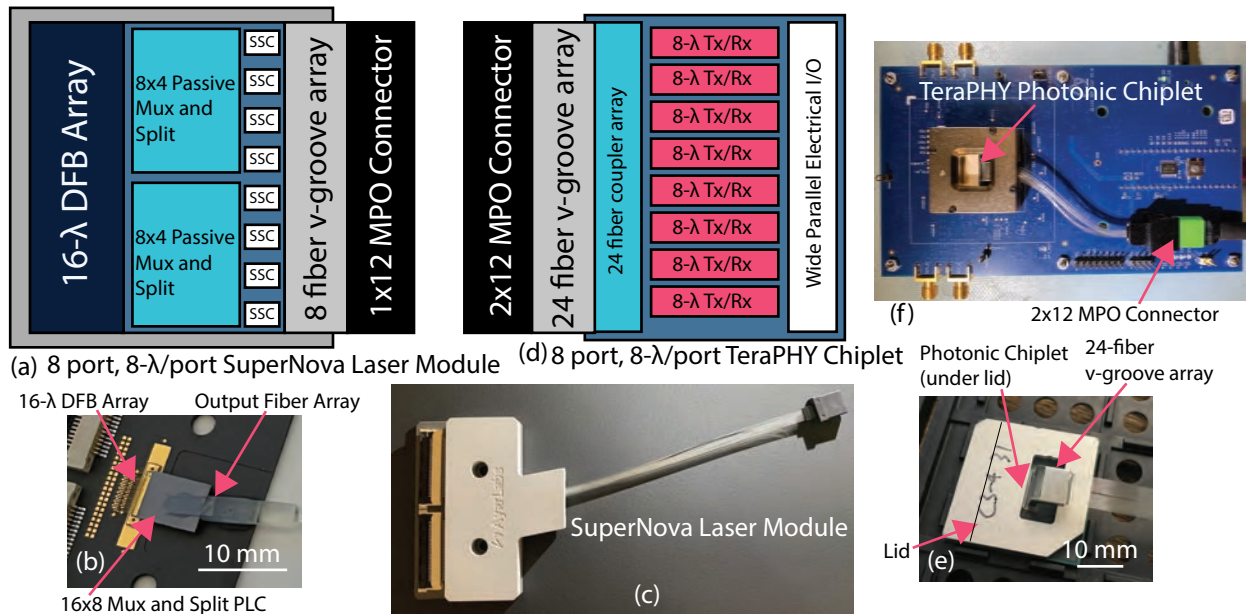


Fig. 1: (a) Block diagram of the multi-port, multi-wavelength laser module. (b) Picture of the assembled laser module without lid and (c) with lid. (d) Block diagram of the optical I/O CMOS chiplet. (e) Picture of the assembled optical I/O chiplet package. (f) Picture of the optical chiplet test board.

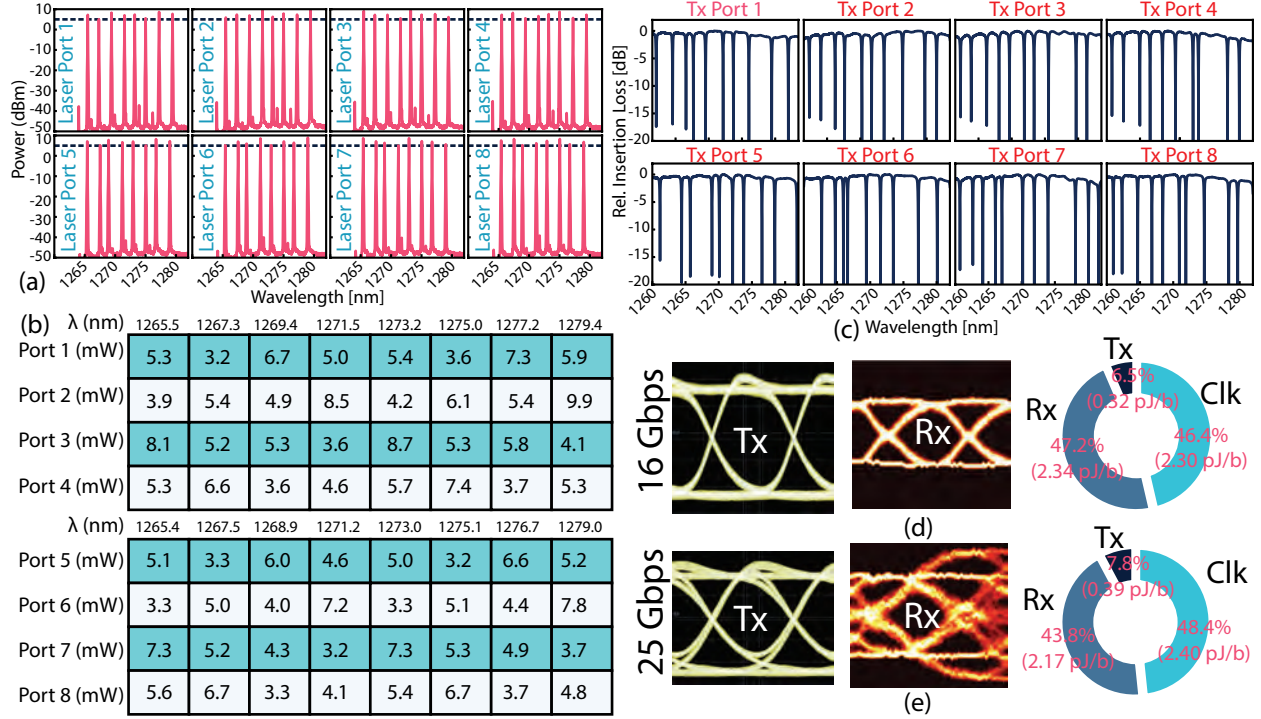


Fig. 2: (a) Measured optical power per port. The dashed line is 5dBm. (b) Tabulated laser wavelengths and power across all ports. (c) Measured relative insertion loss through the transmitter. (d) 16 Gbps and (e) 25 Gbps characterization. The 25 Gbps Rx eye is after on-chip decision feedback equalization (DFE) integrated in each receiver.

Figure 1(a) shows the main functional blocks of the laser source, and Fig. 1(b),(c) shows a picture of the assembled SuperNova laser without and with the lid. The laser combines a 16- λ DFB laser array with a passive planar lightwave circuit (PLC) that does the multiplexing and splitting to an 8-fiber output array, with each fiber containing 8 wavelengths. The fiber array connects to an off-the-shelf 1x12 MPO connector. Figure 1(d) shows a simplified functional block diagram of the TeraPHY optical I/O chiplet. Figure 1(e) shows a picture of the assembled TeraPHY chiplet package. Figure 1(f) shows a picture of the test CMOS optical I/O chiplet test board.

Figure 2 shows a summary of characterization data for the SuperNova laser and the TeraPHY optical I/O chiplet. The output optical spectra for all eight optical ports is shown in Fig. 2(a) and is tabulated in Fig. 2(b). The laser spectra are taken at $\approx 30^\circ$ C case temperature. The average optical power per laser wavelength is ≈ 5.3 mW across all wavelengths and all output ports. The variability in the output powers comes mainly from variability in the optical insertion losses through PLC.

Figure 2(c),(d),(e) shows a summary of key characterization data from the TeraPHY optical I/O chiplet. The untuned optical spectra for all eight transmitter ports are shown in Fig. 2(c) showing eight microring resonances per port. Note that the spectra shows an untuned channel spacing. Each microring is thermally tuned upon link initialization to match the incoming laser grid. The nominal microring resonance spacing is designed to span the full free spectral range (FSR) of one microring, which in this design is ≈ 3.2 THz. Each optical channel can be configured to run at 16 to 25 Gbps per wavelength, as shown in the transmitter and receiver optical eye diagrams in Fig. 2(d),(e). The Tx/Rx characterization data is acquired using built-in self-test (BIST) modes of operation that use on-chip pseudo-random bit sequence bit-error rate checkers. The energy efficiency breakdown for the optical macros is shown in Figs. 2(d),(e). To characterize the energy efficiency, all eight optical ports are enabled, and we measure the current draw of all voltage rails to arrive at a measured total power. We then divide the measured power by the aggregate data rate to arrive at pJ/bit. We measure 4.96 pJ/bit for both 16 Gbps and 25 Gbps operation per optical channel with small differences in the power breakdown. The power is dominated by the receiver and high-speed clock distribution. Although the thermal tuning power was negligible in the direct measurement, the thermal tuning power will vary depending on the operating conditions when co-packaged with a host ASIC. We estimate that the maximum thermal tuning power that would ever be required is ≈ 5 mW per optical channel. This results in an additional maximum 0.4-0.6 pJ/bit for 25 Gbps and 16 Gbps, respectively, which brings the estimated range of energy efficiencies to 4.9-5.5 pJ/bit.

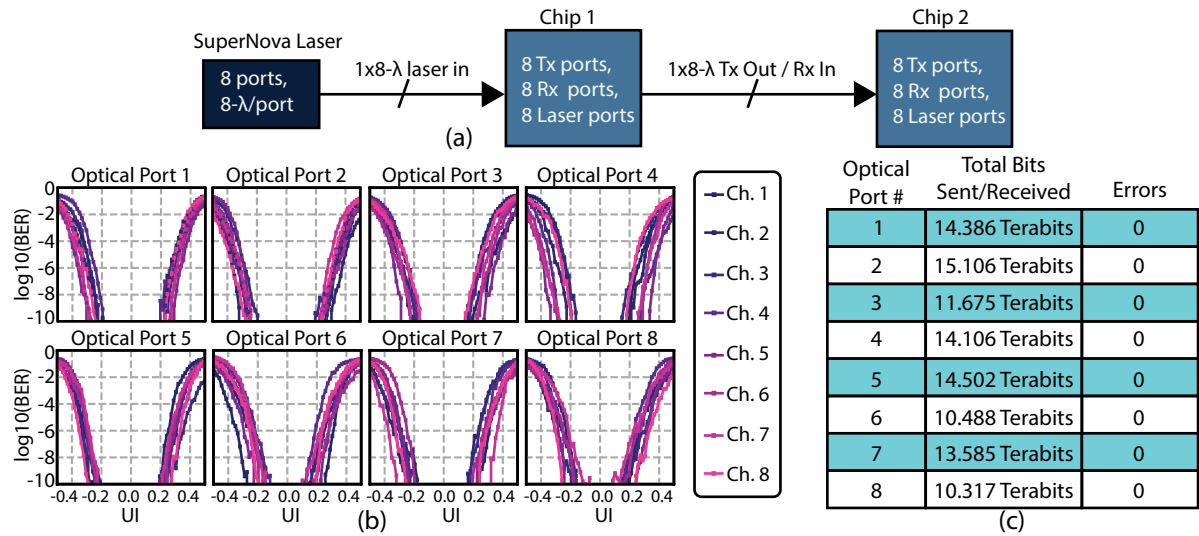


Fig. 3: (a) Experimental setup for characterization. One port at a time was characterized. (b) Measured BER bathtub curves for all channels, all ports. (c) Table summarizing error-free operation. A total of 104.167 Terabits was sent and received across all ports.

Figure 3 summarizes the full link results. An isolator is used on the output of the SuperNova laser, and all fibers are polarization maintaining, single-mode fibers. Each optical port is characterized individually with all eight WDM channels operational such that there is 128 Gbps of simultaneous bandwidth. During initial link bring-up, the microrings on the Tx and Rx undergo automated optimization that establishes thermal lock to the incoming wavelength grid. Additionally, since the absolute frequency microring resonance frequency is sensitive to manufacturing variation, the initialization routine also establishes Tx microring to Rx microring mapping. Figure 3(b) shows measured bathtub curves for all eight optical ports. For each plot, there are $1e12$ bits sent per channel such that the center region of the plots represents error free transmission at an aggregate 1 Terabit ($1e12$ bits) across all channels. All channels operate simultaneously during the measurements for each BER bathtub curve. Figure 3(c) tabulates the measured results and confirms error-free operation. In addition to the link characterization data shown in Fig. 3, we also confirmed simultaneous (all ports, all channels) 1.024 Tbps error-free operation using a reference benchtop laser with 4.5 mW average power per wavelength.

In conclusion, we have demonstrated error-free communication at 128/Gbps per optical port using an $8-\lambda$ microring-based WDM architecture. We measured the optical Tx/Rx chiplet energy efficiency to be 4.96 pJ/bit with an estimated additional < 0.6 pJ/bit from thermal tuning depending on the data rate and the environmental operating conditions. Near-term future work will focus on quantifying the laser power contribution, enabling simultaneous operation of all optical ports to match the reference laser performance, increasing the data rate per optical channel during multi-port operation, and enabling full-duplex links. We expect further improvements in optical link budget to come from improved consistency of insertion loss across the 24 fiber v-groove array connected to TeraPHY. Additionally, better optimization of the assembly of the SuperNova module is expected to obviate the need of an external isolator. We believe this demonstration shows the viability of microring-based WDM optical I/O architectures with remote laser sources to meet the emerging energy efficiency, bandwidth density, and latency (FEC-free) requirements of multi-Tbps computing fabrics.

Acknowledgments: We would like to acknowledge Dr. Gordon Keeler and the DARPA PIPES program. This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). 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 or the U.S. Government.

References

1. M. Wade *et al.* "TeraPHY: A Chiplet Technology for Low-Power, High-Bandwidth In-Package Optical I/O," in IEEE Micro, vol. 40, no. 2, pp. 63-71, 1 March-April 2020
2. C. Sun, *et al.*, "TeraPHY: An O-Band WDM Electro-Optic Platform for Low-Power, Terabit/s Optical I/O," VLSI Technology Symposium, 2020