High Density Silicon Photonic Integrated Circuits and Photonic Engine for Optical Co-packaged Ethernet Switch

Ling Liao[(1)], Saeed Fathololoumi, David Hui

[(1)] Intel Corporation, USA, ling.liao@intel.com

Abstract We present performance of 1.6 Tbps silicon photonic integrated circuits (SiPICs) and their key device components. These high bandwidth density SiPICs enabled fully functional photonic engines (PEs) that can be co-packaged with Ethernet switches and demonstrated IEEE standards compliant interop-link with 400 Gbps DR4 modules.

Introduction

Data center traffic is doubling every 2-3 years and is expected to surpass 20 ZB in 2021 [1-2]. To keep pace with this rapid data growth, hyper-scale data centers are doubling the bandwidth (BW) capacity of their switches and optical pluggable transceivers on a similar cadence. While pluggable optics have been able to scale BW density to keep their form factors largely unchanged from 40 Gigabit Ethernet (40GE) to 400GE, continuing this trend will be challenging. Moreover, power efficiency (pJ/bit) has largely stayed constant from generation to generation, i.e. driving up total power from ~3W at 100GE to ~12W at 400GE. Optics cost also has not scaled well and represents a growing share of the bill of materials of optical transport equipment, switches and routers with each successive BW transition [3]. A rethinking of Ethernet optics and data center connectivity is therefore needed to keep up with the ever-increasing data traffic.

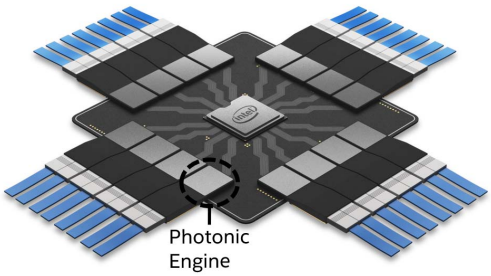

Fig.1. Schematic drawing of an optical co-packaged switch where Photonic Engines are on the same package as the switch ASIC in the center.

Co-packaged optics is considered by many in the industry to be the breakthrough technology needed to address the scaling challenges [4-5]. It aims to shrink pluggable optics that are on the face plate of switch chassis and move them inside, to be directly on the same package as the switch ASIC, as illustrated in Fig.1. This new architecture enables significant SerDes power reduction because the electrical channel between the switch die and photonic engines (PEs) can shrink in length from more than ten inches to centimetres, and it alleviates face plate physical congestion because the bulky pluggable optics are replaced by simple fiber patch panels.

Of the many design considerations for optics, BW density, power, cost and reliability are particularly important for co-package optics. In fact, reliability is uniquely important because the integrated nature of co-package optics means that any failure requires the replacement of the entire switch system and not just the failed optics. We report here the demonstration of 1.6 Tbps SiPICs that address the above requirements. We also share performance of PEs enabled by these SiPICs, including that of interop-links with DR4 pluggable transceivers.

Silicon Photonic Integrated Circuit

The SiPICs are designed to support IEEE standards compliant DR4 optical interface. Their transmitter (Tx) and receiver (Rx) channel architectures are shown in Fig.2. Different designs of the SiPICs have been demonstrated using these basic architectures, including 1.6 Tbps Tx PICs [6], Rx PICs, and fully integrated 800 Gbps Tx+Rx PICs.

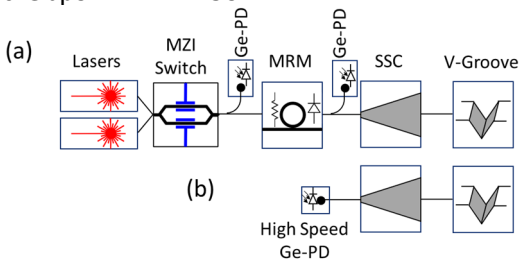

Fig.2. SiPIC architecture for each of (a) transmitter and (b) receiver channels.

For the 1.6 Tbps Tx SiPIC, 32 hybrid silicon distributed feedback (DFB) lasers are integrated to support a pair of redundant lasers for each of the 16 optical channels. Our hybrid silicon laser technology bonds III-V gain materials to the SiPh wafer, after which arrays of lasers are fabricated using CMOS processes [7]. When the size of the bonded III-V gain material is unchanged, making one or multiple lasers has nominally the same cost. To keep the size of the gain material the same, the redundant lasers are placed at half

pitch. This redundant integrated laser design has two key benefits compared to remote laser approaches. First, it removes the complexity, cost, and added optical loss associated with the laser-to-SiPIC fiber assembly. Second, full laser redundancy greatly improves failure in time (FIT) rates of laser arrays, as shown in Fig.3. Our laser has a FIT/laser of ~2, but a FIT/laser of 10 is also included for comparison. For a 51.2 Tbps switch where 512 lasers are needed to support 512 106 Gbps PAM4 channels, laser redundancy can improve laser related system FIT to ~0.1.

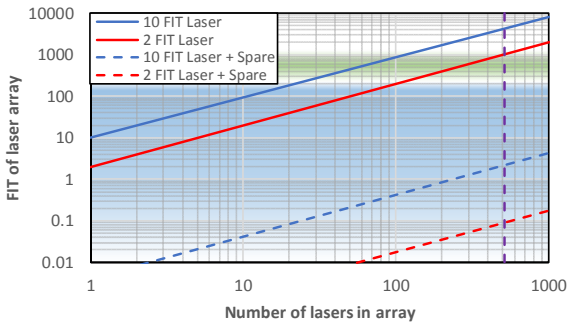

Fig.3. Calculated FIT rates of laser arrays as a function of the number of lasers in the array, for designs with and without laser redundancy [8]. For a 51.2 Tbps co-packaged switch using lasers that each has a FIT of 2, laser redundancy improves laser-related system FIT by 4 orders of magnitude to ~0.1.

Following the 32 hybrid Si lasers on the SiPIC are Mach-Zehnder interferometer (MZI) switches with phase tuners to select the operating laser from each redundant laser pair. The laser light is then modulated using micro-ring modulators (MRM) with integrated heaters and Ge-photodetectors (Ge-PDs) that are part of a dynamic control loop to optimize and maintain the MRM bias point insertion loss (IL). Modulated light is then expanded to a 9 µm mode field diameter using integrated spot size converters (SSCs). V-grooves are formed on-die to allow for low cost passive alignment of SMF28 fibers to the SiPIC. For the Rx, the incoming light is coupled on die using polarization insensitive SSCs and directed to the integrated high speed Ge-PDs.

The integrated hybrid lasers, as measured on multiple wafers from multiple process runs, have wavelengths of 1311±4.5 nm over 70 °C range, side mode suppression ratio (SMSR) >45 dB, and relative intensity noise (RIN) <-145 dB/Hz. For reference, IEEE 802.3-bs DR4 specifications for wavelength, SMSR and RIN are 1311±6.5 nm, >30 dB, and <-136 dB/Hz, respectively. Optical power in the silicon waveguide of a representative hybrid laser is measured to be > 20 mW at 80 °C with < 100 mA injection current.

The MRMs after the MZI-based laser selector switches have 10 µm radius and are designed for critical coupling [9]. They have 6.7 nm free spectral range, > 15 dB DC extinction ratio (ER), and ~5.5k quality factor (Q). This ring Q allows for good balancing of drive voltage swing, photon lifetime, and process tolerance. The high-speed phase modulator of the MRM is based on a reverse bias PN diode and has a measured phase efficiency of 0.52 V*cm and BW of 38 GHz at -1.5 V DC. Fig.4. shows shift of the MRM spectra and resonance as a function of applied voltage. When operating the MRM at 6 dB IL, a 2 V_{pp} swing centered at -1 V_{dc} yields ~5 dB ER.

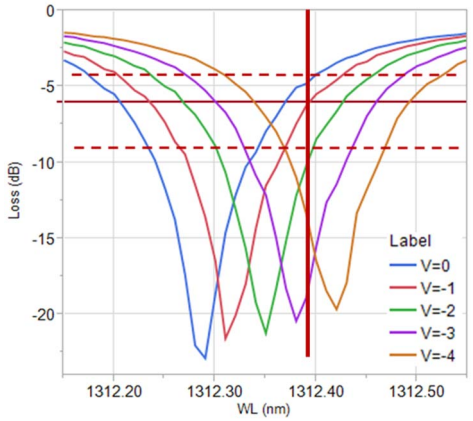

Fig.4. Transmission spectra of MRM, showing resonance shift with bias and 5dB ER with 2V swing.

The Ge-PDs are used both as power monitors for Tx MRM bias point IL control as well as Rx high speed detection. They are vertical PIN diodes formed by growing germanium on top of silicon waveguides. At 30°C and -2.6 V, the Ge-PDs have dark current of <100 nA. Responsivity saturates at -1 V, as shown in Fig.5, reaching ~0.9 A/W. High speed measurements with 50 Ω RF termination shows a typical BW of 39 GHz at -2 V bias.

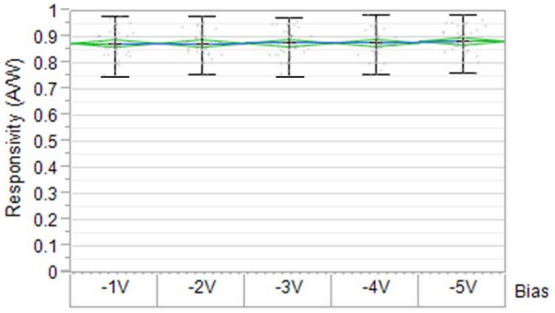

Fig.5. Ge-PD responsivity at room temperature.

Tx and Rx SSCs to enable efficient SMF28 fiber coupling are based on a SiN inverse taper design. They have polarization insensitive performance with ~0.8 dB taper loss, ~0.5 dB active alignment coupling loss, and 0.4 dB ripple over O-band. Passive optical alignment using integrated V-groove arrays shows mean coupling loss of ~1.16 dB. The ~0.7 dB coupling loss delta between passive and active alignment can be reduced with index matching epoxy and improved V-groove targeting.

Photonic Engine

The PE architecture implemented several innovations to scale BW density. These included the use of compact bare ICs such as our SiPICs, advanced 3D chip packaging, and fine-pitch LGA connector to interface with the switch package. Also, dc-dc voltage converters are moved outside of the PE, so the needed voltage rails come in through the connector from the host platform. The PE uses commercially available SerDes dies with high-voltage swing driver outputs to directly drive the MRMs, and TIA ICs are placed between the SerDes die and the Rx Ge-PDs. As shown in Fig.6, all high-power dies are placed on the top side of the PE organic substrate to allow for efficient heat extraction. Low power consuming ICs (controller IC and flash die) are bonded to the bottom side of the substrate. The PE is 25 mm x 22 mm, roughly the size of a postage stamp.

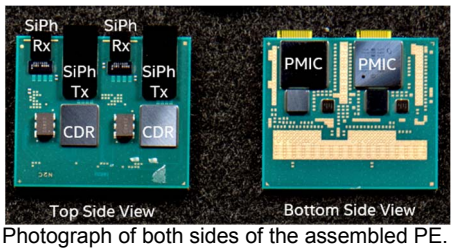

Fig.6. Photograph of both sides of the assembled PE.

Both the SerDes and TIA ICs can only support one 400 Gbps DR4 port. So, the two sets of ICs integrated in the PE can only enable operation of two 400 Gbps DR4 optical links, even though each SiPh Tx die is designed for and capable of supporting four 400 Gbps DR4 (1.6 Tbps). Fig.7 shows the eight PE Tx eye diagrams after the TDECQ equalizer, measured using a Keysight DCA-M system and SerDes' internal PRBS13Q generated pattern at 106.25 Gbps. The measured performance is in good agreement with simulations, showing TDECQ of 1.7-2.1 dB, ER >3.5 dB and RLM >0.9, which are compliant with IEEE802.3-bs DR4 specifications.

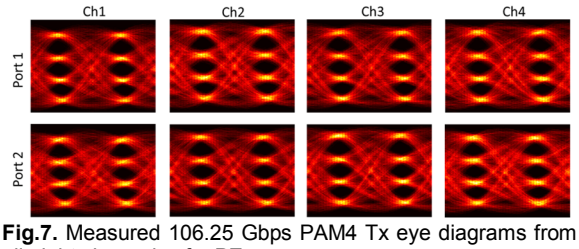

Fig.7. Measured 106.25 Gbps PAM4 Tx eye diagrams from all eight channels of a PE.

Performance of the PE Rx was measured using a commercial DR4 Tx as the reference source. The Tx signal has 2.76 dB TDECQ, 0.78 dB Ceq, and 4.4 dB ER and is a PRBS31 data pattern. Fig.8 shows the Rx sensitivity waterfall curves for all eight Rx channels, and they all have better than -4.5 dBm sensitivity at 2.4E-4 error rate, again more than meeting IEEE802.3-bs 400 Gbps DR4 specifications.

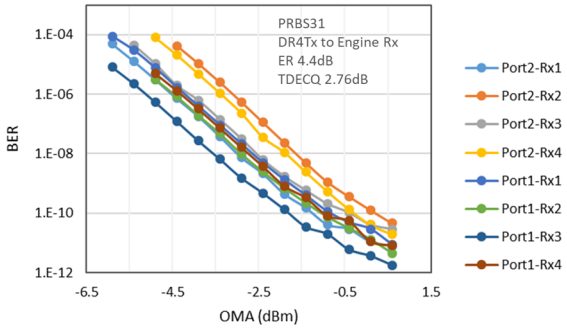

Fig.8. Sensitivity curves for all eight channels of PE Rx.

To demonstrate interoperability, a PE and DR4 module are optically connected in full loop-back mode. An Ethernet traffic generator with a PRBS31 payload is used to source the DR4 electrical data and analyze the return data. Fig.9 shows the FEC statistics collected over 4 days for the entire loop (DR4 to PE and back). The highest number of errors corrected within an Ethernet frame is six. Since KP FEC can correct up to 15 error frames, the data shows that there is significant margin to maintain the required BER target. The PEs, using commercial LR SerDes, have 19.2 pJ/bit power efficiency, compared to ~30 pJ/bit of pluggables. We see path to 13.5 pJ/bit with improved design and XSR SerDes.

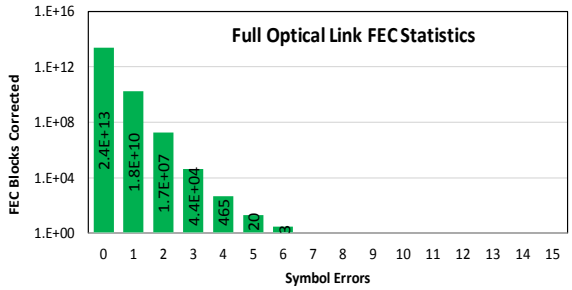

Fig.9. FEC statistics over 4 days of continuous running the link between PE and DR4.

Conclusion

We demonstrated 1.6 Tbps SiPICs that integrate hybrid Si laser arrays to eliminate the need for laser assembly, MRMs to improve BW density, and SSCs along with V-grooves to enable passive fiber alignment. We also realized laser redundancy to improve system reliability and integrated Ge-PDs to enable both MRM bias control and Rx high speed detection. Using the SiPICs, we demonstrated compact PEs that met DR4 specifications and successfully supported post-FEC error-free inter-op links. These PEs have also been co-packaged with Ethernet switches and demonstrated full functionality [10].

Acknowledgements

We wish to thank all our Intel colleagues in design, fabrication, assembly, test and reliability who made this work possible.

References

[1] Cisco Global Cloud Index, 2016-2021

[2] F. Chang, *Datacenter Connectivity Technologies: Principles and Practice*, River Publishers, 2018, Chap1.

[3] *The State of The Optical Communications Industry: 2020*, LightCounting, May 2020

[4] C. Minkenberg *et al*., "Network Architecture in the Era of Integrated Optics", *J. Opt. Commun. Netw*., 11, pp. A72-A83, 2019.

[5] B. Wheeler, “Prepare for Copackaged Optics, Limits on Serdes Speed Will Drive Transition to Optical I/O”, Microprocessor Report, Linley Group, May 2020

[6] S. Fathololoumi et al., “1.6Tbps Silicon Photonics Integrated Circuit for Co-Packaged Optical-IO Switch Applications”, Optical Fiber Communications Conference (OFC), San Diego, CA, 2020.

[7] R. Jones *et al*., "Heterogeneously Integrated InP/Silicon Photonics: Fabricating Fully Functional Transceivers", *IEEE Nanotechnology Mag*., vol. 13, no. 2, pp. 17-26, 2019.

[8] R. Jones, “Lasers for Co-Packaged Optics”, OIF Workshop: Co-packaged Optics - Why, What and How, Oct 2020.

[9] J. Sun, *et al*., "A 128 Gb/s PAM4 Silicon Microring Modulator with Integrated Thermo-Optic Resonance Tuning," *J. Lightwave Technol.,* 37, pp. 110-115, 2019.

[10] Intel demonstrates industry-first co-packaged optics ethernet switch: https://newsroom.intel.com/news/intel-demonstrates-industry-first-co-packaged-optics-ethernet-switch/#gs.j0puly