3.2Tb/s Heterogeneous Photonic Integrated Circuit Chip in a Co-Packaged Optics Configuration

Damien Lambert[(1)*], Jeff Rahn[(2)], Majid Sodagar[(1)], Murtaza Askari[(1)], Paveen Apiratikul[(1)], John Spann[(1)], Thang Pham[(2)], Yishen Huang[(2)], Stephen Krasulick[(1)]

[(1)]Skorpios Technologies Inc., 2706 Montopolis Dr, Austin, TX 78741, [(2)]Meta Platforms, Inc., 14 Hacker Way, Menlo Park, CA 94025
[*]Damien_Lambert@Skorpiosinc.com

Abstract: We present a low-cost, scalable 3.2Tbps heterogenous photonic integrated circuit chip assembled in a co-packaged optics configuration. The integration of III-V material directly into the Silicon-Photonic chip offers clear form-factor, density, and thermal dissipation advantages.

1. Introduction

Efficiency improvements for core switching have enabled improved scale for key switching components, enabling increases in total optical bandwidth in large scale data centers. However, the same efficiency gains have not been achieved for data transfer between the core of the switch. As a result, an increased fraction of the total system power budget gets consumed by the electrical and optical interconnects between switch ASIC cores. While optical link budgets are highly optimized, the power consumed by the electrical interconnect between ASIC core and optical module could be improved via reduction in loss in this link. This leads to interest in applying Co-Packaged Optics (CPO) as a power efficiency strategy. Additionally, CPO can reduce total power consumption per switch by 40% vs. a pluggable solution. Indeed, current form factor for pluggable optics will likely not be able to support 1.6Tbps or higher bit-rate due to high optical and electrical density requirements, thermal dissipation challenges and high-power consumption. CPO configurations are expected overcome these challenges because the switch ASIC and the opto-electrical chips are assembled in close-proximity with a very compact design [1]. This can reduce power consumption needed to drive electrical interconnects to optics, reduce latency by reducing needed signal processing, and enable compact packaging of future switch generations. An initial application has been for 51.2Tbps fabric switches in a co-packaged configuration with four 3.2Tbps transceiver modules for each of the four sides of the ASIC chip [2,3]. Meeting costs, latency, power, and bandwidth requirements introduces multiple challenges due to the high number of optical and electrical I/O connections, scalable form factor considerations, and the high thermal density. Silicon photonics offers promising solutions to overcome these challenges [6,7] because it enables fully integrated compact designs with simpler circuitry and simpler packaging. As a result, several orders of magnitude reduction in area and power consumption can be achieved. Furthermore, the combination of wavelength division multiplexing (WDM) and integrated laser source within a single silicon photonics chip can greatly reduce the number of electrical and optical I/O ports resulting in architectures benefiting from 1-5 dB higher link budgets [4]. Lastly, the CMOS compatibility of silicon photonics platforms ensures high yields and high-performance reliability at low costs, which forms a very scalable platform for multi-Tb/s architectures and sub-pJ/bit data transfer applications.

The silicon photonics platform is well suited to scale with higher modulation format complexity, such as PAM4 or QAM, to transport Ethernet data directly over DWDM [5]. The CPO integration approach is well suited for PAM4 modulation formats because it helps reducing noise in addition to saving an estimated 30% of power and 40% of packaging costs. System design favors the use of integrated laser platforms for CPO: external light sources (required when integrated lasers are not available) require high power and polarization control to deliver the light to the CPO modulator.

In this work, we present the design, fabrication, test and CPO assembly of 3.2Tbps Heterogeneous Photonic Integrated Circuit (HPIC) transceiver chips manufactured in a CMOS foundry and using a PAM4 modulation format to transmit and receive data. We further demonstrate the advantage of incorporating laser sources into the silicon chips to increase density and reduce packaging form factor and fiber count.

2. Chips architecture

We have fabricated 3.2Tbps CPO HPIC transceiver chips with a compact size, high power efficiency, high speed, and low noise. The optical chips are constructed from 4 blocks of Skorpios 800Gbps 2xFR4 transmitter and receiver pairs plus a 1x16 Spot Size Convertor (SSC) array that is designed for 51.2Tbps switch ASICs. For reliability, a redundant laser is added to each wavelength which can be turned on when on-chip monitoring detects a reduction of power. The purpose of redundant lasers is to prevent failure correlation and improve FIT of the CPO engine. As shown in Fig. 1, the chips are 30 mm long and designed to match the pitch of future generations of DSP drivers and TIAs. Given tighter constraints originating from the DSP drivers and TIA pitches, there is empty space on the HPIC, which will enable design shrinks for future production. The 3.2Tbps transceiver chip has been implemented in Skorpios proprietary HPIC technology utilizing seamless III-V integration of lasers, high-speed modulators, SOAs, and high-speed photodiodes into a wafer-scale CMOS process. The approach offers a compactness advantage beyond industry levels of integration. Furthermore, the optical outputs and inputs of each 3.2Tbps transceiver chip are designed to have on-wafer back-to-back waveguide connection with its 3.2Tbps transceiver neighbor chip within the wafer floor plan. This layout architecture coupled with Skorpios' internal laser source integration approach offers on-wafer testing capability for all components, which significantly reduces times and costs for both testing and burn-in activities.

The fabricated HPIC chips feature 16 main active lasers and 16 redundant backup lasers remaining in standby mode. The output of each laser is split into 2 channels with a 3 dB beam splitter, providing 32 TX coherent light sources to the chip's architecture. As shown in Fig. 2 (a), the 32 light sources are modulated by 32 high-speed EAM modulators and amplified by 32 Semiconductor Optical Amplifiers (SOA). An amount of 2% from the optical output of the 32 SOAs is taped with 32 directional couplers and the power level is measured by 32 Monitoring Photo Diodes (MPD). The remaining 92% of the optical power is multiplexed into 8 tuning-free echelle grating multiplexors (Mux). Finally, the 8 optical outputs of the TX Mux are routed to the edge of the chips and the modes are expanded with 8 SSC devices. Interleaved with the TX SSC devices, the edge of the HPIC chips feature 8 RX SSC devices converting the optical modes for the inputs of 8 tuning-free echelle grating de-multiplexors (DeMux). Finally, the 32 optical outputs of the DeMux are routed to 32 Photo Diodes (PD) to perform the RX functions of the HPIC chips.

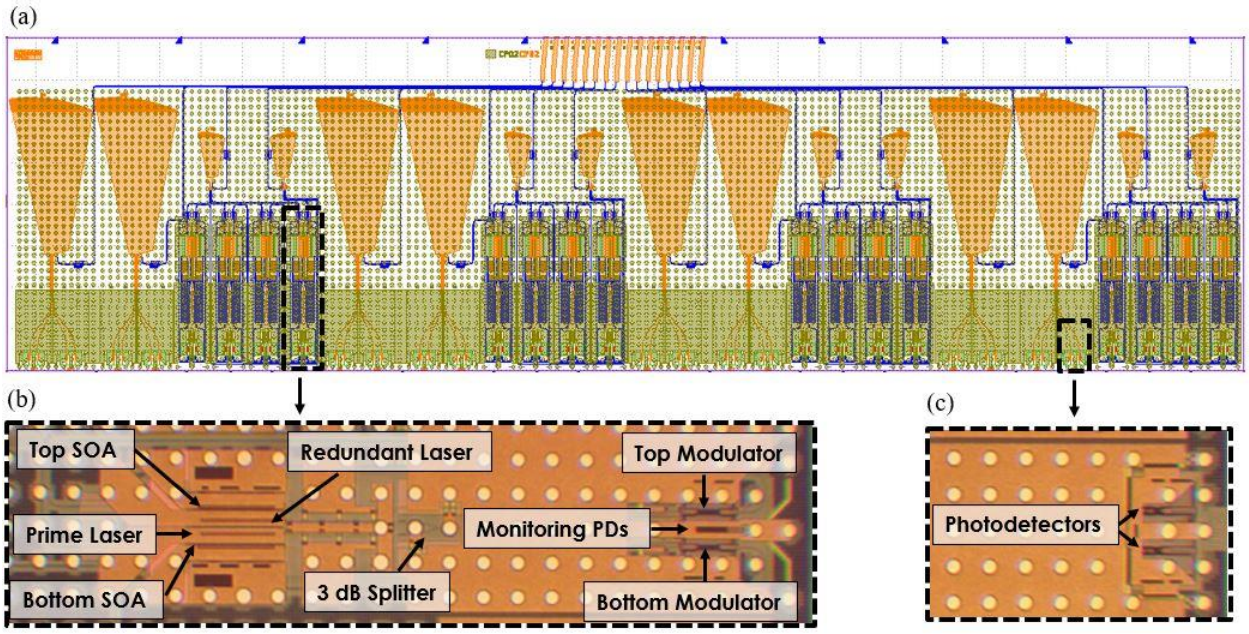

Fig. 1. (a) 3.2Tbps HPIC layout comprised of 4 blocks of 800Gbps TX interleaved with 4 blocks of 800Gbps RX, (b) Picture of 2 TX channels supporting 100Gbps each within a 800Gbps TX block, (c) Picture of 2 RX channels supporting 100Gbps each within a 800Gbps RX block.

3. Test results and CPO Assembly

Prior to performing CPO assembly, the 3.2Tbps CPO HPIC transceiver chips were tested for TX RF performance and RX responsivity. For RF testing, we employed a PAM4 modulation format using the SSPRG pattern. As seen in Fig. 2(b), all 32 channels of the HPIC chips provided 100Gbps modulation speed meeting the FR4 specifications at 1.3V Vpp. Specifically, we observed Extinction Ratio (ER) greater than 3.5 dB, outer OMA levels greater than -0.3 dBm, TDECQ lower than 3 dB and OMA-TDECQ greater than -1.6 dBm for all channels. After RF testing, the 3.2Tbps HPIC chips were flip-chip bonded to a low temperature co-fired ceramic (LTCC) electrical fan-out IC. Along with each HPIC chip, 2 Marvell DSP chips and 4 Macom TIA chips were also flip-chip bonded to the same LTCC IC mechanical support. At the time of design and fabrication of the 3.2Tbps CPO HPIC transceiver chips, Skorpios silicon photonic platform could accommodate the channel pitch density required for 51.2Tbps fabric switches, however, both DSP and TIA chips available commercially could only accommodate half the pitch density. As such, we have chosen to demonstrate the concept of CPO integration capability by connecting electrically and utilizing every other channel from the 3.2Tbps CPO HPIC transceiver chips. Thus, the 2 DSP chips and 4 TIA chips can support a PAM4 operation for 16 of the 32 TX channels of each HPIC chip. While the LTCC IC is designed to emulate a socket mount configuration onto a 51.2Tbps ASIC switch, it was soldered onto an evaluation board to enable traffic testing. Finally, 8 fibers were butt-coupled to the edges of the chips to complete the assembly.

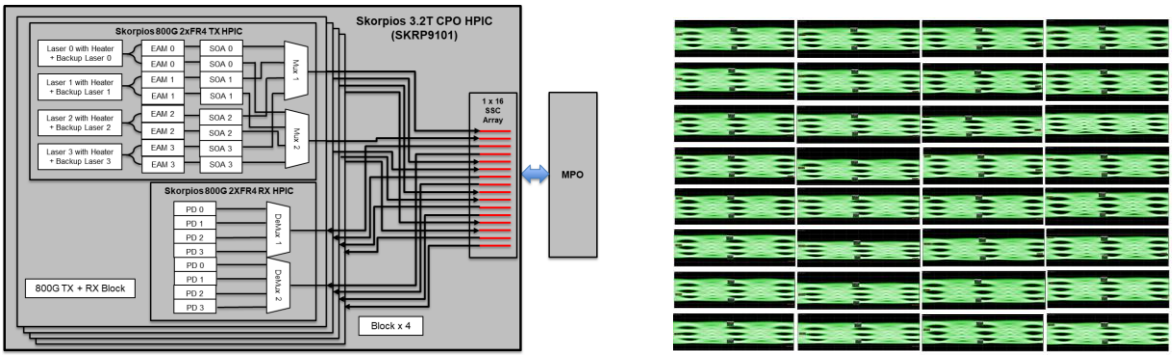

Fig. 2. (a) Left: Architecture of the 3.2Tbps CPO HPIC, (b) Right: Eye diagram for 32 PAM4 channels.

4. Conclusion

We have successfully demonstrated the capability to integrate 3.2Tb/s CPO HPIC chips compatible with the assembly and operation of 51.2Tbps fabric switches. Furthermore, Skorpios' silicon photonics approach to monolithically integrate high-efficiency III-V material into SOI wafers was demonstrated to be a highly scalable platform supporting multi-Tbps applications for next generations of Hyperscale datacenters.

5. References

[1] C. Minkenberg et al, "Co-packaged datacenter optics: Opportunities and challenges," IET Optoelectron, 2021;15:77–91 (2021).
[2] Consortium for On-Board Optics, "Design Considerations of Optical Connectivity in a Co-Packaged or On-Board Optics Switch," *www.onboardoptics.org*, Cobo (2022).
[3] CPO document, "3.2T-Copackaged-Optics-ModulePRD-1.0," V1.0, 2021 URL: http://www.copackagedoptics.com/wpcontent/uploads/2021/02/JDF-3.2-Tb_s-CopackagedOptics-Module-PRD-1.0.pdf.
[4] B. Buscaino et al, "External vs. Integrated Light Sources for Intra-Data Center Co-Packaged Optical Interfaces," IEEE J. Lightwave Technology, vol. 39, no. 7, pp. 1984 – 1996, 2021.
[5] R. Nagarajan, et al, "Silicon Photonics-Based 100 Gbit/s, PAM4, DWDM Data Center Interconnects," J. Opt. Comm. and Networking, vol. 10, no. 7, pp. B25-B36, 2018.
[6] A. V. Krishnamoorthy et al, "Progress in low-power switched optical interconnects," IEEE J. Sel. Topics Quantum Electron., vol. 17, no. 2, pp. 357–376, 2011.
[7] G. Li et al, "A CMOS Wafer-Scale, Monolithically Integrated WDM Platform for TB/s Optical Interconnects," Optical Fiber Communication Conference (Optical Society of America, 2014), Th1C.2, 2011.