2.5D Heterogeneous Integration for Silicon Photonics Engines in Optical Transceivers

Radhakrishnan Nagarajan, *Fellow, IEEE*, Liang Ding, Roberto Coccioli, *Senior Member, IEEE*, Masaki Kato, Ronson Tan, Pushkraj Tumne, Mark Patterson, and Linda Liu

(Invited Paper)

***Abstract*—We present our work in the area of heterogeneous optical integration, where separately manufactured electronic components are assembled on to an active silicon photonics interposer to form a higher-level component. This process allows for the integration of components independently designed and optimized from several different technology and foundry platforms onto a common interposer. Heterogeneous integration is essential for manufacturing higher speed and performance components. Higher levels of integration also allow for closer placement of devices which minimizes the parasitic power consumed to compensate for the frequency dependent losses in the interconnect traces.**

***Index Terms*—2.5D integration, heterogeneous integration, photonic integrated circuits, silicon photonics.**

I. Introduction

PHOTONIC integration falls into two broad categories; monolithic and heterogeneous. Monolithic integration is one where the variety of optical [1] or both electrical and optical [2] components are integrated onto a common substrate, using a common foundry process. Heterogeneous integration differs from monolithic in a significant aspect that it integrates components from different foundry technology platforms, usually post fabrication, using physical attach processes [3]. Heterogeneous integration also refers to cases where different material systems, are physically bonded to a common substrate and processed together to form the final component [4]. In this paper, we will focus on the former.

The term "2.5D" heterogeneous integration refers to how the discrete chips are integrated on a common substrate. A 2D architecture is where two or more active devices are placed side-by side and interconnected via a common substrate. A 3D architecture is where two or more active devices are stacked and interconnected without the help of a common substrate, like in high bandwidth memory (HBM) [5], [6].

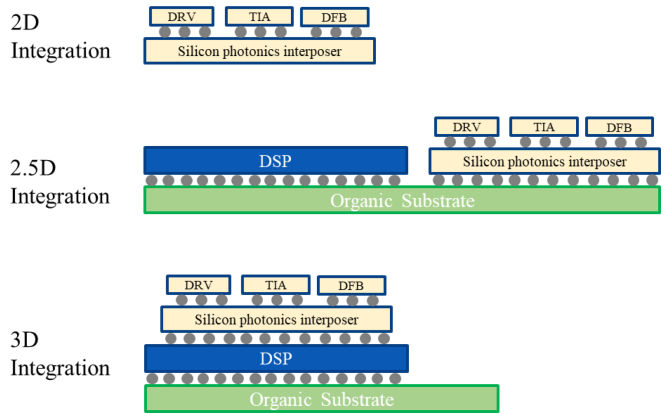

Fig. 1. 2D, 2.5D and 3D integration nomenclature as used in the paper. We discuss the 2D and 2.5D integration techniques in detail.

In this paper, we discuss a packaging technique where 2D structures, on a common silicon photonics interposer/substrate, are interconnected with other silicon devices via a package substrate. This method is referred to as the 2.5D integration.

There are differing views in the industry on what constitutes 2D, 2.5D and 3D integration [7]. We have decided to follow the IEEE guidelines given in the Heterogeneous Integration Roadmap [5]. This is illustrated in Fig. 1, where, in the case of devices discussed in this paper, the driver amplifier (DRV) for the transmitter, transimpedance amplifier (TIA) for the receiver and distributed feedback (DFB) lasers are 2D integrated on a common silicon photonics substrate. The silicon photonics substrate also has all the optical components, like the Mach Zehnder modulator (MZM), power and wavelength splitters, combiners, and high-speed photodetectors (HSPD).

The 2D optical assembly is then integrated with a digital signal processor (DSP) and other ASIC's (application specific integrated circuit) on a common high speed organic substrate, to form the final 2.5D assembly. If instead, the 2D assembly is stacked on the DSP, then it becomes a 3D assembly.

In going from a 2D to 2.5D integration, additional structures like the TSV's (Thru-Silicon-Via) may be needed in the silicon photonics assembly. This is discussed in detail in the paper.

Manuscript received 6 August 2022; revised 10 October 2022; accepted 10 October 2022. Date of publication 18 October 2022; date of current version 9 November 2022. *(Corresponding author: Radhakrishnan Nagarajan.)*

Radhakrishnan Nagarajan, Masaki Kato, and Linda Liu are with the Marvell, Santa Clara, CA 95054 USA (e-mail: radha@marvell.com; mkato@marvell.com; lindaliu@marvell.com).

Liang Ding, Roberto Coccioli, Pushkraj Tumne, and Mark Patterson are with the Marvell, Westlake Village, CA 91362 USA (e-mail: liangding@marvell.com; rcoccioli@marvell.com; ptumne@marvell.com; mpatterson@marvell.com).

Ronson Tan is with the Marvell, Tai Seng Centre, Singapore 369522 (e-mail: tronson@marvell.com).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JSTQE.2022.3214418.

Digital Object Identifier 10.1109/JSTQE.2022.3214418

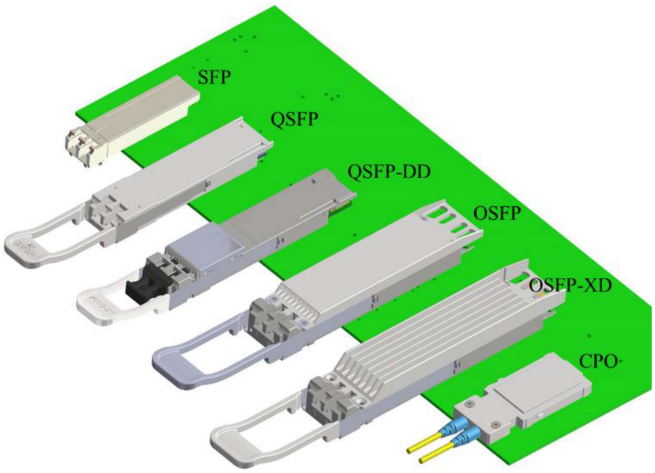

Fig. 2. Evolution of transceiver modules, with increasing data rates, over the past two decades. All, except for the CPO (which is meant to be on-board integrated), are front face pluggable modules.

TABLE I
TRANSCEIVER MODULE TYPES AND CAPABILITIES

		Input Electircal Lanes			Optical Signal Rate (Gbaud)
Rate (Gbit/s)	Module Type	Number	Rate/Lane (Gbit/s)	Format	
1	SFP	1	1	NRZ	1
10	SFP	1	10	NRZ	10
40	QSFP	4	10	NRZ	4 x 10
100	QSFP-28	4	25	NRZ	4 x 25
200	QSFP-56	4	50	PAM4	4 x 25
400	QSFP-DD/OSFP	8	50	PAM4	4 x 50
800	QSFP-DD/OSFP	8	100	PAM4	8 x 50
1600	OSFP-XD	16	100	PAM4	2 x 4 x 50 or 4 x 100
3200	OSFP-XD	16	200	PAM4	2 x 4 x 100
3200	CPO	32	100	PAM4	32 x 50

The focus of the paper is heterogeneous integration on a silicon photonics substrate. There are number of commercial silicon photonics foundries currently in operation [8]. In silicon photonics, photodetectors are built from Ge grown on the Si substrate. The optical sources are either heterogeneously attached InP DFB lasers [3] or DFB lasers processed on the silicon platform, from heterogeneously attached InP epitaxial material [4], [9]. Although there are multiple approaches to package integration [6], [10], the method we have described here uses the silicon photonics device as an interposer, with TSV's.

II. Optical Transceiver Evolution

For data center interconnects, front face pluggable modules have been the mainstay of the optical transceiver market for the past 20 years. Fig. 2 shows a progression of module form factors over the years. All of them are front face pluggable except for Co-Packaged Optics (CPO).

SFP (small form factor pluggable) module, introduced in the early 2000's, and still in deployment, was used for 1 Gbit/s and 10 Gbit/s data rates [11]. The pluggable optical transceivers have progressed over time, in data rate, size and type as shown in Table I. The data rates shown in Table I do not include forward error correction (FEC) and other protocol related overhead, which increases the actual bit rates by 5% to 15%.

TABLE II
PLUGGABLE MODULE EVOLUTION IN SPEED AND POWER

	Year	Power (mW)	Rate (Gbit/s)	Energy (pJ/bit)
1G SFP	2001	1	1	1000
10G SFP	2009	2.5	10	250
40G QSFP	2012	3.5	40	88
100G QSFP (CWDM4)	2015	3.5	100	35
200G QSFP (FR4)	2018	5.5	200	28
400G QSFP-DD (FR4)	2020	10	400	25
800G OSFP (FR4)	2022	16	800	20
CPO (oif 2021.131.12)	2024	56	3200	18

As the data rates increased to 40 Gbit/s, there was a need for modules with a larger number of input electrical lanes. The SFP, which only had a single input electrical lane, morphed into a Quad Small Form Factor Pluggable (QSFP) with 4 input electrical lanes. QSFP-28 and QSFP-56 are higher bandwidth electrical variants designed to accommodate the evolution to 100 Gbit/s and 200 Gbit/s data rates [11]. For the commensurate four lane optical output, the 4 wavelengths (channels) are on a Coarse Wavelength Division Multiplexing (CWDM) grid in the O band of the optical fiber [12].

At the transition to 200 Gbit/s, the NRZ modulation format changed to PAM4, with a factor of 2 difference between the data and signal rates. Although the data rate was 50 Gbit/s per lane, the signal rate was 25 Gbaud.

At 400 Gbit/s, there was a need for 8 input electrical lanes, as there was a generational mismatch in the data rate capability of the input electrical and output optical lanes. QSFP-Double Density (QSFP-DD) [13] and Octal Small Form Factor Pluggable (OSFP) [14], module types were designed to address this need. As the electrical Serdes speeds got to 100 Gbit/s, the same module types are being used to implement 800 Gbit/s transceivers. The output optical lanes are still in multiples of 4 wavelengths.

There are pluggable transceiver proposals for data rates up to 3.2 Tbit/s. OSFP-eXtra Dense (The OSFP-XD) is a pluggable module with 16 input electrical lanes [14]. At the highest data rates, the signal rate which determines the analog bandwidth needed in these applications, has reached 100 Gbaud. At these baud rates, the electrical power needed to maintain signal integrity, limits the transmission distance over copper traces. This leads us to the CPO [10], [15], where the optical transceivers are small, highly integrated, and assembled close to the host, which in many cases is a high-speed switch, on a common substrate. This is a departure from the traditional pluggable module approach to optical transceivers.

A. Evolution of Speed and Power Consumption

Table II shows the evolution of power consumption (under worst case operating conditions) for the various module types, in the specific applications listed, with the year when they were first deployed. The power consumption includes all the optical and electrical components within the transceiver. Although the power consumption increases with data rate, the normalized energy consumption, pJ/bit, has been dropping dramatically.

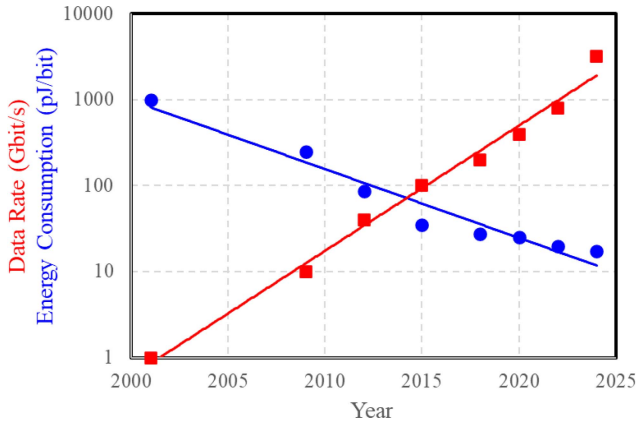

Fig. 3. Data rate and energy consumption for optical transceivers for the past two decades.

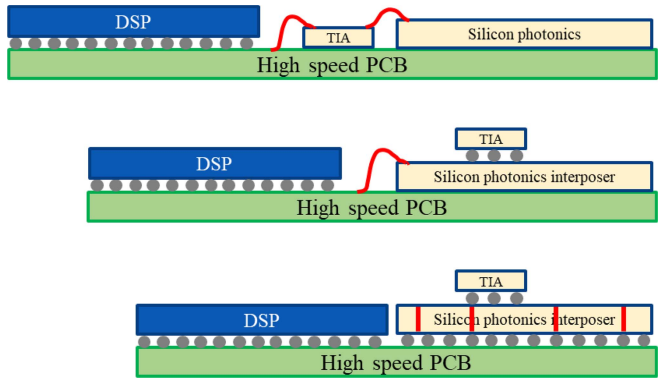

Fig. 4. Evolving approaches to packaging in optical transceivers. The wire bonds (top two diagrams) and TSV's (bottom diagram) are in red. In addition to the TSV's, there may be one or more Re-Distribution Layers (RDL) at the bottom of the silicon photonics interposer for signal routing.

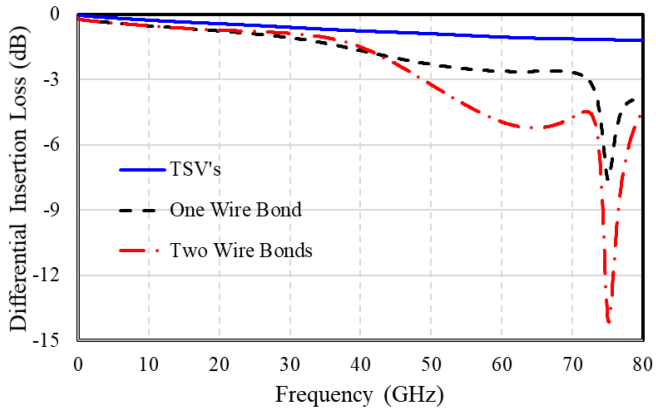

Fig. 5. Normalized electrical insertion loss, as a function of frequency, for the three packaging scenarios in Fig. 4.

As shown in Fig. 3, in the past 20 years, data rates in pluggable modules have increased by 3 orders of magnitude, while the energy consumption per bit has dropped by 2 orders of magnitude. This is a truly amazing achievement!

Using 2.5D integration, this power can be further reduced to below 10 pJ/bit. A large contributor to the total module power consumption is the host electrical interface. With integration, as the optical modules get smaller and are co-packaged with electrical host ASIC, the power at this interface can be reduced. With even tighter integration, we may not need a DSP inside the optical module, and it can be directly driven by the host ASIC. At this point the power consumption may be reduced to about 5 pJ/bit.

In the past two decades, although the energy consumption per bit has been reduced by a factor of 100 (Fig. 3), the actual power consumption in the optical module has increased by a factor 10. The signal (baud) rate, which determines the analog bandwidth, has increased by a factor 100. These two trends have put a lot of pressure on high speed and thermal packaging.

B. Evolution of Transceiver Optical Packaging

Optical transceiver packages, from different suppliers, vary considerably in design. To simplify the discussion on high-speed packaging, we have picked three cases shown in Fig. 4, where we have simplified the layout, and only shown high-speed path for the TIA.

The first row in Fig. 4 illustrates the current approach to packaging optical transceivers. The DSP is usually a ball grid array (BGA) package, or a bare die assembled on a high speed printed circuit board (PCB). The TIA is wire bonded to the PCB on one side, and wire bonded to the photonics device on the other side. These wire bond lengths are short of the order of 500 µm or less. The middle row, in Fig. 4, is what many suppliers are exploring, where a 2D optical device (integrated with the TIA) is assembled next to the DSP via a wire bond.

Although the second assembly method is better, the two, wire bond based approaches have limited analog bandwidth, and signal discontinuities, as shown in the modeling results of Fig. 5. The cleanest, from the high-speed packaging point of view, is the last row in Fig. 4, where the silicon photonics device is also an interposer with TSV's, and the assembly is completely free of wire bonds.

Ansys HFSS [16] program was used to design and simulate the performance of all three structures shown in Fig. 4. For all the wire bond cases, a double wire bond, using a 1 mil diameter wire, were used, and the wire bond length is kept at 350 µm to minimize the inductance discontinuity. As for the TSV case, we have simulated the actual via structure, with layout metal routing.

It is further worth pointing out that the modelling in Fig. 5, is for an ideal case, where we have modelled 50Ω traces separated by wire bonds. The resonances in the insertion loss behavior for the wire bond cases, due to the discontinuities and actual device characteristics, will be worse in practice (occur at lower frequencies), and will be assembly dependent. The use of TSV's allow for consistent high-speed performance of the optical transceivers.

The other aspect is thermal where the assembly needs to be cooled. Wire bonded assemblies are harder to cool, and the heat needs to be pulled out from the bottom, thru the PCB. In 2.5D integrated structures, the wire bonds are not "in the way", and the heat can be channeled to the top through package lid, in addition

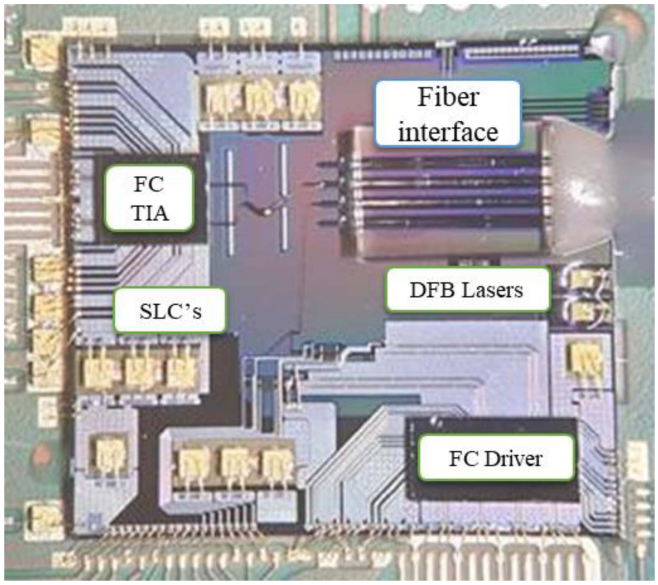

Fig. 6. 2D assembly using the silicon photonics chip as an interposer.

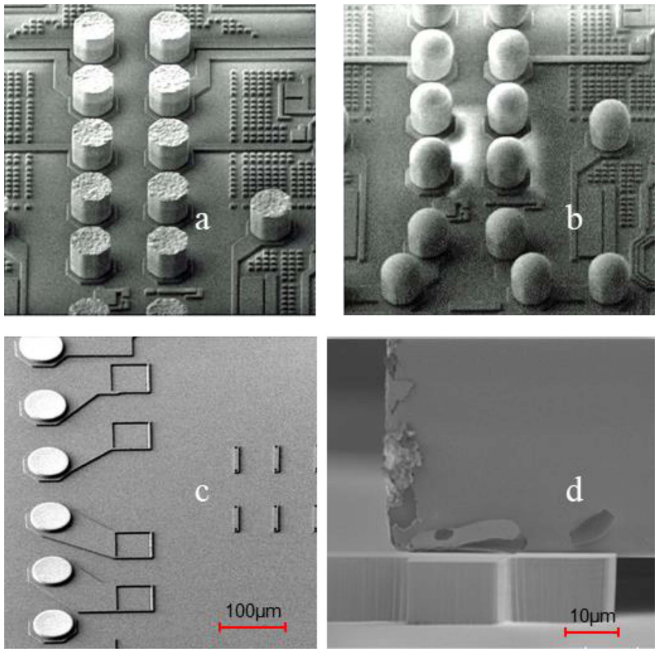

Fig. 7. The flip chip assembly process shows (a) the bumps as plated on the ASIC's, (b) the characteristic Cu bump shape after reflow, (c) the plated pads on silicon photonics, and (d) the DFB laser resting on a stopper after AuSn solder attach. The approximate scales have been included in the SEM's. They are the same for the first three SEM's, and different for the fourth.

to the PCB. This becomes important in the higher power, higher speed transceivers.

III. 2D Integration

Fig. 6 shows a 2D assembly on a silicon photonics chip used as an interposer [3]. Integrated on this interposer (reference to Fig. 1) we have a flip chip (FC) TIA, FC Driver, DFB lasers, V-groove based fiber interface and single layer decoupling capacitors (SLC). This is assembled on a PCB and wire bonded for operation.

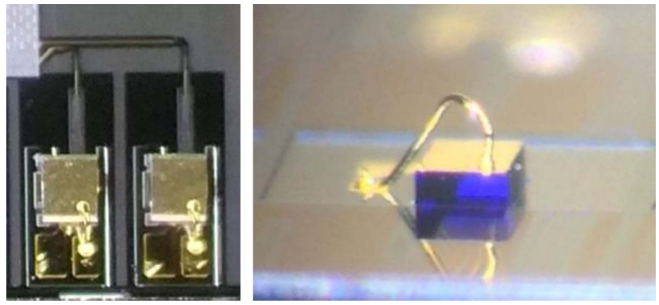

Fig. 8. Top and side views of the completed DFB laser assembly in a trench on the silicon photonics.

The silicon photonics chip is a dual channel device with edge coupled optical fibers, Mach Zehnder modulators (MZM), high speed Ge photodetectors (PD), wavelength multiplexers and demultiplexers.

A. Assembly

For the flip chip assembly process, the TIA and Driver wafers are first "bumped". As shown in Fig. 7(a) and (b), in the first step, the TIA and Driver wafers are plated with Cu (copper) pillar bumps, with SnAg (tin-silver) solder caps. The wafers then undergo a thermal reflow process, and the Cu pillar bumps get their final shape. The bumped TIA and DRV wafers are then tested and diced, and the passing devices are assembled on to the silicon photonics wafer.

The ASIC pads on the silicon photonics wafers are Cu-Ni-Au (copper-nickel-gold) plated as shown in the Fig. 7(c). The DFB pads, in a trench, are AuSn (gold-tin) solder plated.

Alignment structures are etched into a trench on the silicon photonics wafer to accurately place the DFB, in the vertical direction, with respect to the optical edge coupler [8] on the silicon photonics chip. Fig. 7(d) shows the DFB laser resting on one such vertical alignment structure. The optical coupling loss is about 2.5 dB, largely due to the mismatch between the non-optimal, asymmetric DFB mode and the symmetric edge coupler input mode.

Accurate alignment of the DFB to the silicon photonics coupler, and it's subsequent soldering is a challenge. Using a modern, precision alignment machine like the Amicra Nano [17], the DFB can be accurately placed to within ± 0.5 μm of the target position, in practice.

As shown in Fig. 8, the DFB lasers are flip chip soldered to the silicon photonics chip. There is a wire bond from the substrate side (n-side) of the DFB to a pad on the silicon.

The SLC's are attached using a conductive epoxy to the silicon, as well as the PCB, as shown in Fig. 9. The fiber interface is in a V-groove and epoxied on the silicon photonics chip as shown in Fig. 6.

B. Performance

For laboratory evaluation, the 2D assembly is attached to a high-speed PCB, as shown in Fig. 10, with high speed and DC connectors. V-grooves integrated on the silicon photonics chip

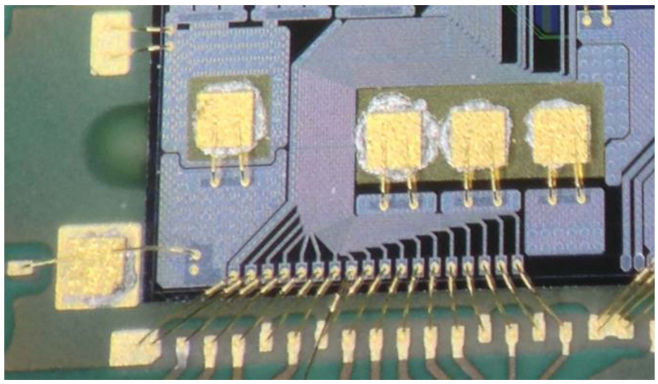

Fig. 9. SLC's and the silicon photonics chip are epoxy attached to the PCB.

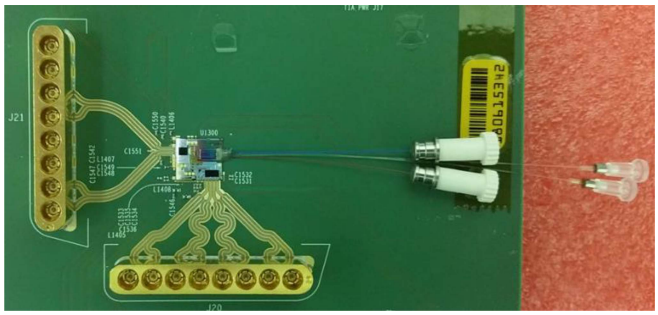

Fig. 10. Test PCB assembly of the 2D silicon photonics engine.

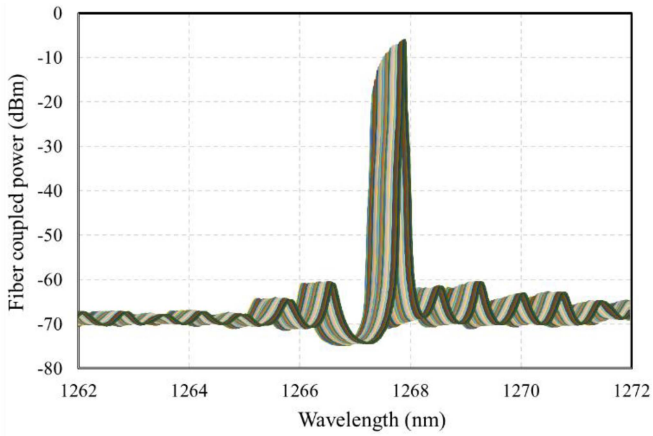

Fig. 11. Detailed optical spectrum of the 300 μm long DFB laser attached to the silicon photonics interpose.

allow the use of a passively aligned fiber array assembly for the various optical test configurations.

Fig. 11 shows the detailed optical spectrum of the DFB laser at 30 °C, with a bias current from 16 mA to 80 mA in steps of 2 mA. The detailed spectrum shows that the back reflection at the DFB-silicon photonics interface, if any, was not sufficient to cause any spectral impairments. By measuring the wavelength shift with temperature ($d\lambda/dT$) and with electrical power ($d\lambda/dP$), we calculated the thermal resistance (dT/dP) of the 300 μm long DFB laser attached to the silicon photonics interposer to be 60±2 °C/W.

Fig. 12. PAM ASIC test board used to drive the test PCB in Fig. 10.

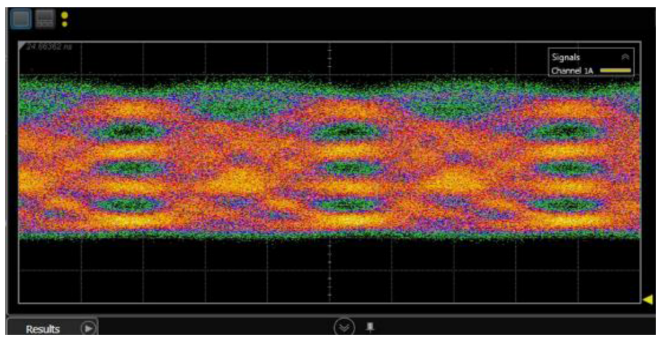

Fig. 13. PAM eye diagram from the 2D light engine at 28.125Gbaud.

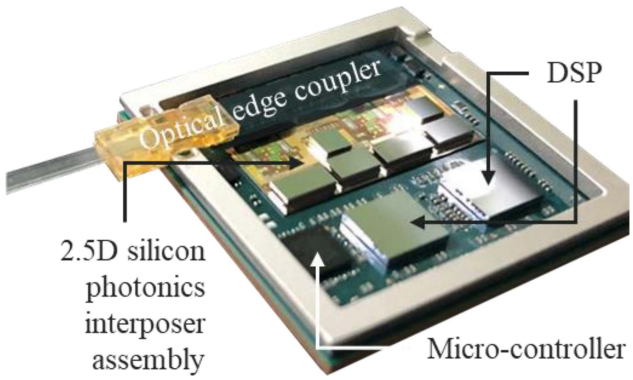

Fig. 14. 2.5D assembly of the 2D light engine on a high-speed, organic substrate with a PAM4 DSP and other ASIC's.

A PAM4 DSP test board, shown in Fig. 12, was used to drive the 2D silicon photonics light engine via RF breakout cables. The PAM4 output optical eye diagram from the light engine, for 28.125 Gbaud, is shown in Fig. 13.

The overall performance was impacted by the long RF cables used to connect the PAM4 test board to the light engine PCB, and the wire bonds from the light engine to the PCB. A fully integrated version, discussed in the next section, does not suffer the same penalty.

III. 2.5D Integration

Fig. 14 shows a 2.5D assembly of the 2D light engine on high-speed substrate, with the high-speed DSP's and controller ASIC's. This assembly does not use any wire bonds, except for the DFB laser. As shown in the block diagram in Fig. 1, the silicon photonics interposer has TSV's and backside solder

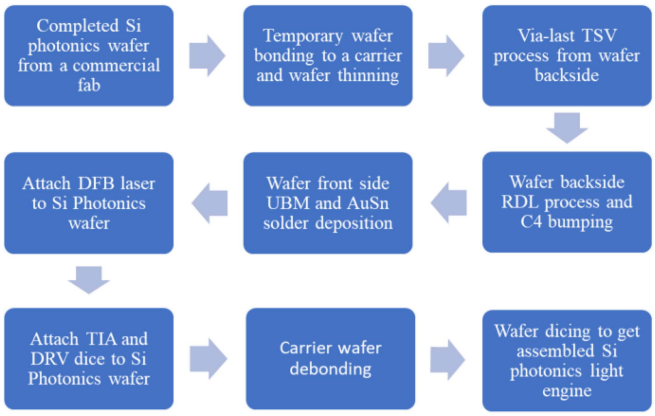

Fig. 15. Wafer level light engine assembly process.

bumps. This is assembled on an organic substrate using a mass reflow process.

Using the existing semiconductor manufacturing ecosystem, we can fabricate TSVs and backside C4 bumps on Si photonics wafers. These are the key enabling features for realizing the compact 2.5D integration.

The Si photonics wafers in this work were fabricated in a commercial Si photonics foundry followed by a via-last TSV process [18], [19], [20], [21], [22], [23], C4 bumping [24], die attach and wafer dicing done at an assembly vendor.

A. Via-Last Process

2D silicon photonics interposer is manufactured at the wafer-level [20]. Fig. 15 shows the details of the wafer-level, light engine assembly flow.

The completed silicon photonics wafer is first thinned and the TSV's are formed using a "via-last" process. Next a Re-Distribution Layer (RDL) [22] and C4 bumps are fabricated on the bottom side of the wafer. The wafer is then ready for front side die attach.

The via-last process, from the back side of the wafer is used to fabricate the TSVs in a silicon photonics wafer. It is more supply-chain friendly compared to the via-first and via-middle processes. Although silicon photonics foundries have rapidly grown in recent years, few of them provide an integrated TSV process. On the other hand, many Outsourced Semiconductor Assembly and Test (OSAT) companies have established TSV process that can be easily used with a silicon photonics wafer. Most of the OSAT TSV processes rely on via-last integration scheme. We believe that CMOS foundry silicon photonics process followed by OSAT-based via-last process is a supply-chain friendly solution for enabling 2.5D integration.

A typical via-last process is shown as in Fig. 16. At the OSAT the silicon photonics wafer is first attached onto a carrier wafer. Grinding and polishing is used to thin the silicon photonics wafer down to a thickness of 50 µm to 200 µm, based on the OSAT's TSV capability. Then, the TSV is etched, stopping on designated metal pads embedded on the front side of the wafer. Following the TSV etching, a dielectric layer is applied into the via for electrical isolation and surface passivation. Next a RDL

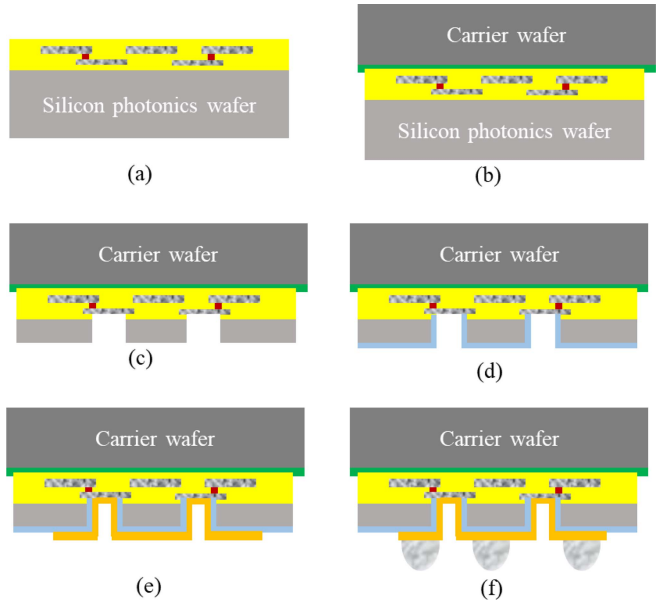

Fig. 16. Via-last TSV process flow for the 2.5D silicon photonics interposer. (a) Silicon photonics start wafer. (b) Bonding to a carrier. (c) Wafer thinning and TSV etching. (d) Dielectric deposition. (e) RDL formation. (f) Back side C4 or Cu pillar bumping.

is formed by a plating process. C4 or Cu pillar bumping process is then done to form the interconnect on the back side of the wafer.

B. Assembly

A chip-to-wafer (also called the die-to-wafer) bonding process is used to assemble the DFB, TIA and DRV on the top side of the silicon photonics wafer.

The TIA and DRV wafers are tested, diced into individual dies and the passing dies are sorted for assembly on to the silicon photonics wafer. As in the 2D process, the TIA and DRV wafers have Cu bumps on them.

The completed wafer, after the assembly process, is shown in Fig. 17. The fully assembled wafer is first tested, diced and the yielded light engine dies go on to the substrate assembly step. The inset in Fig. 17 shows an individual light engine which has been diced from a completed wafer.

The light engine interposer assembly is then attached to the organic substrate using mass reflow. After the reflow step, the C4 (controlled collapse chip connection) bumps on the rear of the light engine form the electrical connection to the substrate.

A planar lightwave circuit (PLC) fiber array assembly (as shown in Fig. 14) was used as the optical signal edge coupler. This, prefabricated, multi-fiber assembly, is actively aligned to the silicon photonics chip and epoxied to the substrate. The optical coupling is less than 1 dB.

Fig. 18 shows the cross-section of the silicon photonics interposer on a multi-layer organic substrate assembly of Fig. 14. The assembly in Fig. 18 consists of TIA, DRV and DFB on a silicon photonics interposer. Fig. 18 shows the TSV's and C4 bumps on the interposer and the details of the multi-layer organic substrate.

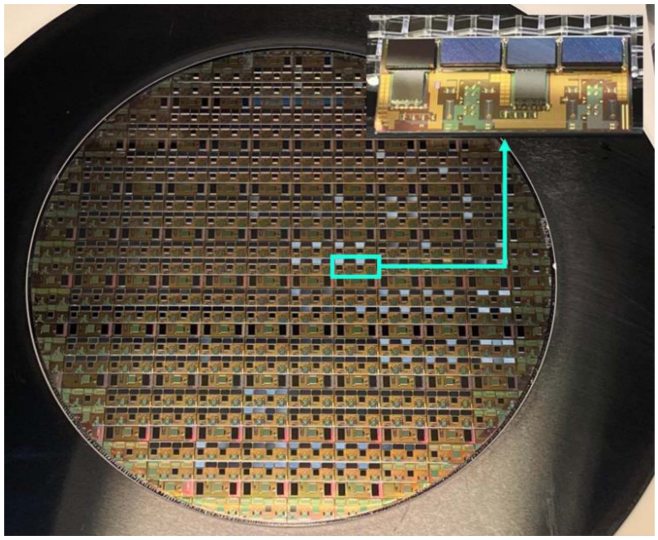

Fig. 17. 200 mm/300 mm wafer level assembly. The diced light engine is shown as the inset.

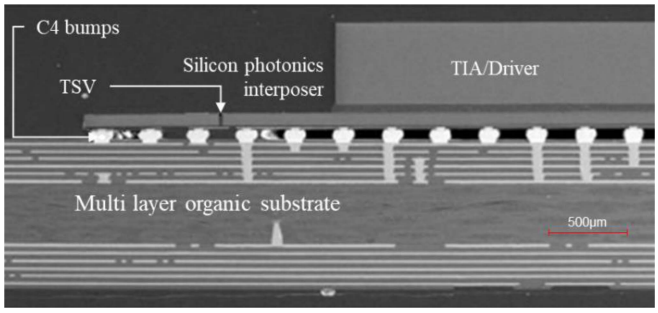

Fig. 18. Cross section of the 2.5D assembly using the silicon photonics as the interposer. An approximate scale has been included in the SEM.

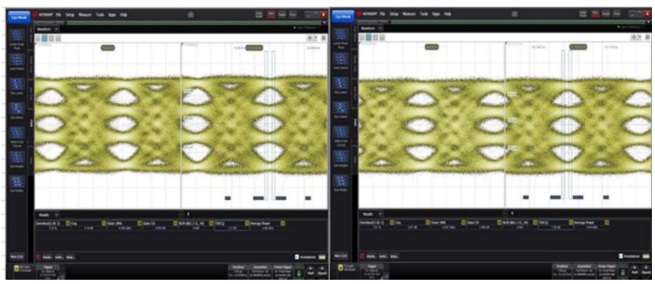

Fig. 19. PAM eye diagram at 26.5625 Gbaud (53.125 Gbit/s). Eye diagrams for 4 channels are shown.

This is the power of 2.5D heterogeneous integration. The electronic ASIC's (CMOS and Bi-CMOS) and photonics can be from different semiconductor technologies, wafer sizes (200 mm or 300 mm) and foundries. They can all be integrated onto a common platform.

C. Performance

Fig. 19 shows optical transmit eye diagrams for 26.5625 Gbaud PAM4 (53.125 Gbit/s equivalent data rate). Fig. 20 shows

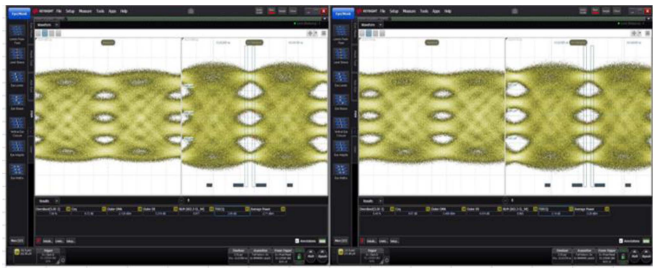

Fig. 20. PAM eye diagram at 53.125 Gbaud (106.25 Gbit/s). Eye diagrams for 4 channels are shown.

IV. System Level Applications

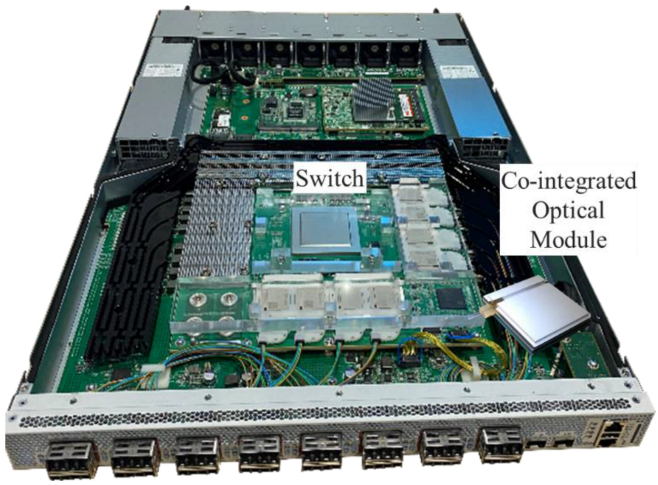

Fig. 21. Switch (Teralynx7^{TM}) integrated with 16 2.5D integrated optical modules shown in Fig. 14.

the same module operating a twice the baud rate, for a data rate of 106.25 Gbit/s.

The DSP's in the package have internal pseudo-random bit sequence (PRBS) generators. We use that feature to generate the optical eye diagrams and generally test the modules. The quality of the optical eye diagrams both at 53.125 Gbit/s and 106.25 Gbit/s are very good.

The TDECQ (transmitter and dispersion eye closure quaternary) [25], [26] for the 53.125 Gbit/s mode (Fig. 19) is less than 1.0 dB and for the 106.25 Gbit/s mode (Fig. 20) is in the range of 2.0 dB. Since the light engine was optimized for 100 Gbit/s operation, the performance at 50 Gbit/s was very good. The TDECQ performance for both applications is well below the IEEE specifications [27], [28].

We are working on improving the TDECQ performance of the 106.25 Gbit/s mode, by further optimizing the equalization at the DSP-light engine interface.

IV. System Level Applications

Fig. 21 shows 16 of these 2.5D optical engine modules (one of them, a lidded version of Fig. 14, is shown as an inset in Fig. 21) assembled on a common high-speed PCB with a 12.8 Tbit/s Teralynx7^{TM} host switch ASIC [29]. This PCB is attached, via mezzanine connector, to the main board which has the power

supplies, cooling fans, CPU, memory, etc., needed for operation in a chassis.

The switch carried traffic in a recent live demo [30].

V. Conclusion

Heterogeneous integration is an enabler for realizing higher speed optical components. Eliminating wire bonds and other sources of signal discontinuities are the key for high speed and improved thermal performance.

There are several approaches to heterogeneous integration, and we have explored one in this paper. The method discussed here uses the silicon photonics chip as the interposer for the 2D integration. This interposer, with TSV's and backside solder bumps, is assembled on to a high-speed substrate for the final 2.5D assembly.

This approach may be extended to include multiple 2D structures on a common substrate to manufacture an optical component with extended functionality or to enable the optical component to be integrated with a much larger ASIC, like a switch, to realize a CPO on a common substrate. These CPO's can then be used to form very compact systems with co-located switch and optical modules.

Acknowledgment

The authors would like to thank the exceptional engineering team at Marvell which contributed enormously to this work.

References

[1] R. Nagarajan, C. Doerr, and F. Kish, "Semiconductor photonic integrated circuit transmitters and receivers," in *Optical Fiber Telecommunications*, vol. VI A, I. Kaminow, T. Li, and A. Willner Eds. New York, NY, USA: Elsevier, 2013, pp. 25–98.

[2] M. Rakowski et al., "45nm CMOS - Silicon Photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2020, pp. 1–3.

[3] R. Nagarajan, L. Ding, M. Kato, and R. Tan, "2.5D Heterogenous silicon photonics light engine with integrated DFB lasers and electronics," in *Proc. OCP Future Technol. Symp.*, 2020, Art. no. FTS116s2.

[4] H. Yu et al., "400Gbps Fully integrated DR4 silicon photonics transmitter for data center applications," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2020, pp. 1–3.

[5] "Chapter 22: Interconnects for 2D and 3D architectures," *IEEE Heterogeneous Integration Roadmap*, 2021, [Online]. Available: http://eps.ieee.org/hir

[6] J. Lau, *Semiconductor Advanced Packaging*. Singapore: Springer, 2021.

[7] "About 2.5D technology," [Online]. Available: https://nhanced-semi.com/technology/about-2-5d-technology/

[8] S. Y. Siew et al., "Review of silicon photonics technology and platform development," *J. Lightw. Technol.*, vol. 39, no. 13, pp. 4374–4389, Jul. 2021.

[9] C. Xiang et al., "High-performance silicon photonics using heterogeneous integration," *IEEE J. Sel. Top. Quantum Electron.*, vol. 28, no. 3, May–Jun. 2022, Art. no. 8200515.

[10] R. Mahajan et al., "Co-packaged photonics for high performance computing: Status, challenges and opportunities," *J. Lightw. Technol.*, vol. 40, no. 2, pp. 379–392, Jan. 2022.

[11] SFP Multi-Source Agreement, 2006, [Online]. Available: http://www.snia.org/sff/specifications

[12] CWDM4 Multi-Source Agreement, 2014, [Online]. Available: http://www.cwdm4-msa.org/

[13] QSFP-DD Multi-Source Agreement, 2016, [Online]. Available: http://www.qsfp-dd.com/

[14] OSFP Multi-Source Agreement, 2017, [Online]. Available: https://osfpmsa.org/

[15] "Implementation agreement for a 3.2Tb/s co-packaged (CPO) module," in *Proc. Opt. Internetworking Forum*, 2022, pp. 1–28.

[16] Ansys HFSS, 2021, [Online]. Available: https://www.ansys.com/products/electronics/ansys-hfss

[17] NANO - Die Bonder and Flip Chip Bonder, 2018, [Online]. Available: https://amicra.semi.asmpt.com/en/products/die-flip-chip-bonder/nano-die-bonder-and-flip-chip-bonder/

[18] X. Zhang et al., "Heterogeneous 2.5D integration on through silicon interposer," *Appl. Phys. Rev.*, vol. 2, 2015, Art. no. 021308.

[19] X. Jing, F. Dai, W. Zhang, and L. Cao, "Via last TSV process for wafer level packaging," in *Proc. IEEE 17th Int. Conf. Electron. Packag. Technol.*, 2016, pp. 1216–1218.

[20] L. Ding et al., "Demonstration of OSAT compatible 300 mm through Si interposer," in *Proc. IEEE 15th Int. Conf. Electron. Packag. Technol.*, 2013, pp. 430–434.

[21] M. J. Kim et al., "Novel 2.5D RDL interposer packaging: A key enabler for the new era of heterogenous chip integration," in *Proc. IEEE 71st Electron. Comp. Technol. Conf.*, 2021, pp. 321–326.

[22] T. Zhou, T. Hang, S. Ma, and M. Xiang, "Design and development of encapsulation process for CIS-TSV wafer level package," in *Proc. IEEE 21st Int. Conf. Electron. Packag. Technol.*, 2020, pp. 1–3.

[23] Y. Yang et al., "Though-Si-via (TSV) keep-out-zone (KOZ) in SOI photonics interposer: A study of the impact of TSV-induced stress on Si ring resonators," *IEEE Photon. J.*, vol. 5, no. 6, Dec. 2013, Art. no. 2700611.

[24] S. Ray, K. Beckham, and R. Master, "Flip-chip interconnection technology for advanced thermal conduction modules," in *Proc. IEEE 41st Electron. Comp. Technol. Conf.*, 1991, pp. 772–778.

[25] J. King, D. Leyba, and G. LeCheminant, "TDECQ (transmitter dispersion eye closure quaternary) replaces historic eye-mask and TDP test for 400 Gb/s PAM4 optical transmitters," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2017, pp. 1–3.

[26] S. Echeverri-Chacón et al., "Transmitter and dispersion eye closure quaternary (TDECQ) and its sensitivity to impairments in PAM4 waveforms," *J. Lightw. Technol.*, vol. 37, no. 3, pp. 852–860, Feb. 2019.

[27] *IEEE Standard for Ethernet - Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation*, *in IEEE Std 802.3bs-2017 (Amendment to IEEE 802.3-2015 as amended by IEEE's 802.3bw-2015, 802.3by-2016, 802.3bq-2016, 802.3bp-2016, 802.3br-2016, 802.3bn-2016, 802.3bz-2016, 802.3bu-2016, 802.3bv-2017, and IEEE 802.3-2015/Cor1-2017)*, 2017.

[28] *IEEE Standard for Ethernet - Amendment 11: Physical Layers and Management Parameters for 100 Gb/s and 400 Gb/s Operation over Single-Mode Fiber at 100 Gb/s per Wavelength*, *in IEEE Std 802.3cu-2021 (Amendment to IEEE Std 802.3-2018 and its approved amendments)*, 2021.

[29] "Marvell Teralynx 7 Data Center Ethernet Switch, 2021, [Online]. Available: https://www.marvell.com/content/dam/marvell/en/public-collateral/switching/marvell-teralynx-7-product-brief.pdf

[30] "Marvell unveils co-packaged optics technology platform at OFC 2022," 2022, [Online]. Available: https://www.marvell.com/company/newsroom/marvell-unveils-co-packaged-optics-technology-platform-at-ofc-2022.html

Radhakrishnan Nagarajan (Fellow, IEEE) received the B.Eng. degree (1st Class Hons.) in electrical engineering from the National University of Singapore, Singapore, the M.Eng. degree in electronic engineering from the University of Tokyo, Tokyo, Japan, and the Ph.D. degree in electrical engineering from the University of California, Santa Barbara, CA, USA.

He is currently the SVP and CTO with Marvell, Optical Connectivity Group. He was a Fellow with Infinera, where he led the development of multiple generations of InP based large scale photonics integrated circuits. Prior to that, he was with SDL, which was acquired by JDS Uniphase, where he developed high power pump laser modules for EDFA applications. He has authored and coauthored more than 190 publications in journals and conferences, and five book chapters in the areas of high-speed optical components and photonics integration. He has been awarded 221 U.S. patents.

Dr. Nagarajan is a Fellow of the Optical Society (Optica) and a Fellow of the Institution of Engineering and Technology. He was the recipient of the 2022 IPRM Award and 2006 IEEE LEOS Aron Kressel Award and for his breakthrough contributions to the commercialization of Large Scale Photonic Integrated Circuits.

Liang Ding received the B.S. degree in optical and optoelectronic physics from the Huazhong University of Science and Technology, Wuhan, China, and the Ph.D. degree from the Division of Microelectronics, School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore.

He is currently a Senior Principal Engineer with Marvell Semiconductors working on packaging technology development for Marvell's optics product line. Prior to Marvell, he was a Research Scientist with the Institute of Microelectronics, Agency for Science, Technology and Research, Singapore, from 2008 to 2014. He has authored and coauthored more than 100 papers in journals and conferences in Si photonics, advanced packaging, and Si nanocrystal electronics and optoelectronic devices.

Dr. Ding was the recipient of the prestigious Singapore Millennium Foundation Ph.D. Fellowship in 2007. In 2008, he was also the recipient of the Chinese Government Award for Outstanding Self-financed Student Abroad.

Roberto Coccioli (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Florence, Florence, Italy, in 1995.

After Postdoctoral studies with UCLA in the field of computational methods for electromagnetics applied to microwave and optical devices, he joined Conexant Systems working on the area of microelectronic packaging. He joined Inphi in May 2001, later acquired by Marvell in 2021, and his current role is the Sr. Director of package development leading a team of engineers responsible for all aspects of package design for high-data rate serdes and silicon photonic devices, including signal and power integrity, thermal management, chip-package mechanical interaction, and manufacturability.

Roberto Coccioli holds 30 US patents in the field of microelectronic and silicon photonic packaging.

Masaki Kato received the B.Eng. degree in electrical engineering, the M.Eng. degree in electronic engineering, and the Ph.D. degree in electrical engineering from the University of Tokyo, Tokyo, Japan. He is currently the AVP of Optical Engineering at Marvell leading the development of Silicon Photonics circuits. He was a Distinguished Engineer with Infinera, where he worked on InP-based large scale photonics integrated circuits. He has authored and coauthored more than 80 publications in journals and conferences. He has been awarded 94 U.S. patents.

Ronson Tan received the B.S. (highest Hons.), M.S., and Ph.D. (1993) degrees in electrical engineering from the Georgia Institute of Technology in Atlanta, Atlanta, GA, USA. Over the years, he has been involved in the development of multiple optical transceiver platforms. He is currently the Senior Principal Engineer at Marvell and was with JDS Uniphase. Prior to JDS Uniphase, he cofounded E2O Communications which was later acquired. Prior to that, he was with HP's Optical Communications Division.

Pushkraj Tumne received the B.Eng. degree in mechanical engineering from the University of Mumbai, Mumbai, India, the M.Eng. and Ph.D. degrees in industrial and systems engineering from the State University of New York, Binghamton, NY, USA. He is currently a Principal Packaging Engineer with Marvell. He has more than ten years of IC package technology development experience. His area of expertise is in research and technology integration of complex, hybrid MCM packages, and 2.5D silicon photonics modules.

Mark Patterson received the B.S. degree in mechanical engineering from the University of California, San Diego, CA, USA, in 1998, and the M.S. degree in mechanical engineering from San Diego State University, San Diego, CA, USA, in 2009. He is currently Principal Packaging Engineer with Marvell Semiconductor working in the area of mechanical and thermal design of optoelectronic packaging and systems. He developed packaging solutions for various telecom and high-speed digital applications with Applied Micro Circuits, now MACOM. He has been issued seven U.S. patents.

Linda Liu received the B.S. degree in electrical engineering from San Jose State University, San Jose, CA, USA. She is also the Sr. Principal Engineer with Marvell, Optical Connectivity Group. She was a Distinguished Technical Staff with Lumentum.