Flip-Chip Integrated Silicon Photonic Bridge Chips for Sub-Picojoule Per Bit Optical Links

Hiren D. Thacker, Ying Luo, Jing Shi, Ivan Shubin, Jon Lexau°, Xuezhe Zheng, Guoliang Li, Jin Yao, Joannes Costa, Thierry Pinguet*, Attila Mekis*, Po Dong+, Shirong Liao+, Dazeng Feng+, Mehdi Asghari+, Ron Ho°, Kannan Raj, James G. Mitchell°, Ashok V. Krishnamoorthy, John E. Cunningham

Sun Labs, Oracle, San Diego, CA
°Sun Labs, Oracle, Menlo Park, CA
*Luxtera Inc., Carlsbad, CA
+Kotura Inc., Monterey Park, CA

Oracle, 9515 Towne Centre Drive, San Diego, CA 92121
Phone: (858) 526-9442, Fax: (858) 526-9176, E-mail: hiren.thacker@oracle.com

Abstract

Silicon photonics holds tremendous promise as an energy and bandwidth efficient interconnect technology for chip-to-chip and within-chip communications in high-performance computing systems. In this paper, we present a low-parasitic microsolder-based flip-chip integration method used to integrate silicon photonic modulators and photodetectors with high-speed VLSI circuits using chips fabricated on vastly different technology platforms. Both the hybrid-integrated silicon photonic transmit (Tx) and receive (Rx) components were tested to demonstrate record sub-picojoule-per-bit performance at 5 Gbps.

I. Introduction

The sustained exponential improvement in performance of semiconductor-based computing systems over the past five decades has been nothing short of incredible. However, the semiconductor industry currently faces steep challenges to continued progress. In particular, the electrical wiring between transistors on a chip and between chips has become a system-performance suffocating factor [1]. The UNIC (Ultraperformance Nanophotonic Intrachip Communication) program aims to address this problem through high-bandwidth, low-power on-chip and chip-to-chip silicon photonics. Such technologies can enable highly compact supercomputer-scale systems.

Oracle's silicon photonic technologies, including innovations in active and passive nanophotonic devices, circuits, and multi-chip packaging, offer the possibility of integrating large number of CPUs along with high-density memory chips in a single multi-chip package, and connecting them with optical communication paths of unprecedented bandwidth density, low energy, and low latency [2]. The program anticipates up to a thousand-fold reduction in energy (per bit), 10 to 100 times lower inter-chip latencies, and a dramatic reduction of physical size over today's systems.

While juxtaposition of silicon photonic devices and VLSI circuits on the same silicon substrate represents the most intimate integration of electronic and photonic technologies, achieving this requires an immense amount of sophistication in design and process integration. In the meantime, hybrid integration by flip-chip bonding is a pragmatic approach to combining best-of-breed silicon photonic devices and VLSI circuits. Such "photonic bridge chip" components may be used as units in a larger transmitter or receiver array or as drop-in communication physical layer elements in a multi-chip computing node [2]. We have successfully integrated several different configurations of silicon photonic transmitters and receivers. Silicon photonic resonant ring modulators and germanium waveguide photodetectors were built on an 8-inch 130 nm silicon-on-insulator (SOI)-CMOS platform as well as on a 6-inch SOI-Photonics platform. These were bonded to VLSI circuits built using a 12-inch 90 nm bulk CMOS technology.

In a hybrid integration scheme, the chip-to-chip interconnect plays a critical role in enabling high-speed and low-energy performance. Specifically, this interconnect ought to have extremely low electrical resistance and capacitance. To this end, we have developed an ultralow profile, small footprint microsolder bump interconnect. The geometry of the bumps keeps parasitics low while the microsolder materials enable a thermocompression bonding process. The conceptual view of silicon photonic chip and VLSI chip integration is shown in Figure 1. The hybrid integration and microsolder methods deployed here follow an approach originally reported for the integration of III-V modulators [3] and VCSELs [4] to silicon.

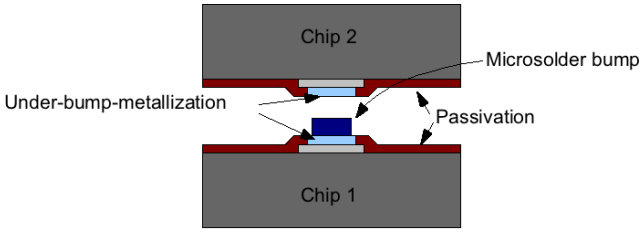

Figure 1. Schematic of a hybrid bond.

In Section II, we describe the microsolder bumping process, which includes under-bump-metallization (UBM) followed by the microsolder interconnects. The hybrid integration process, results and challenges are described in Section III. The hybrid-integrated components were assembled onto printed-circuit boards for high-speed testing and demonstration. The board-assembly process is described in Section IV. Section V details the performance of the hybrid-integrated bridge Tx and Rx components, and Section VI concludes this paper.

II. Microsolder Bump Interconnects

The microsolder bump interconnect structure comprises an under-bump-metallization (UBM) layer and a microsolder bump (Figure 2), as was previously reported in [2]. The UBM layer serves multiple functions: (i) it provides a stable, low resistance contact to the chip's I/O pads, (ii) provides a strong adhesion interface between the die bondpad and bump materials, and (iii) prevents diffusion of the bump materials into the chip. It is the foundation on which the bump and therefore the hybrid bond stand. The microsolder bumps provide much-needed vertical compliance to absorb the effects of local topological variations, chip bow, and any tool-tilt during flip-chip bonding. Additionally, the bump ensures high conductivity as thermocompression fuses metals into opposing pads during flip-chip bonding. Microsolder bumps as small as 10 µm have been successfully fabricated at pitches as low as 25 µm and can be further scaled [5]. Very high bump yield of better than 99.99% has been demonstrated with this approach [6].

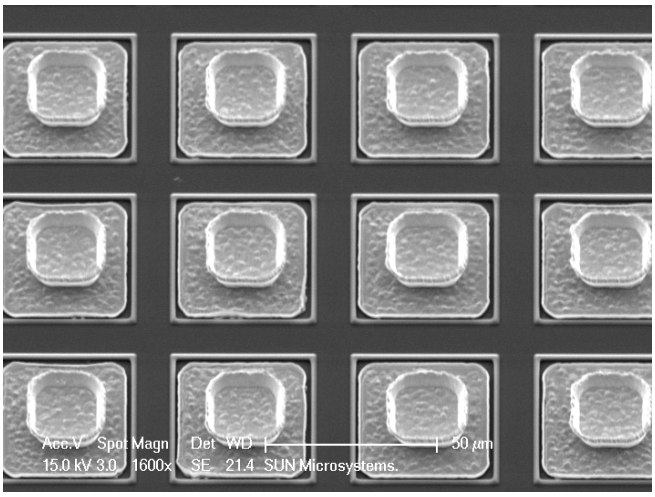

Figure 2. Scanning Electron Microscope (SEM) image of an array of microsolder bump interconnects, consisting of a UBM layer and a bump.

The microsolder bump may be fabricated on either chip or both chips in a hybrid component. In this work, as the VLSI chips were received post-dicing from the foundry, handling the tiny pieces of silicon for any extensive microfabrication was inconvenient. As such, the microsolder bumps were fabricated on the silicon photonic chips and only the UBM was built-up on the VLSI chips. The process flow for building microsolder bumps is shown in Figure 3.

Under-bump-metallization

UBM was deposited onto the chip's Al pads using either a lithography-based process or an electroless plating process.

(i) Lithography-based UBM process:

This approach uses photolithography, metal deposition and liftoff to deposit UBM on only the chip pads. Following lithography, a thin metal seed layer was deposited by physical vapor deposition in a high vacuum (10^{-6} Torr), Argon gas environment. This minimized the possibility of oxidation.

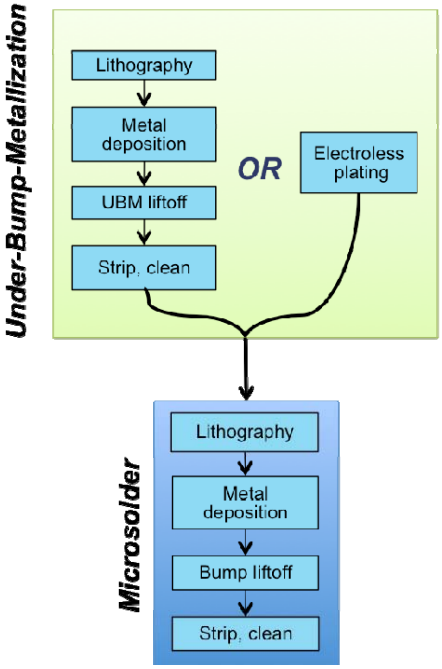

Figure 3. Microsolder process flow.

During this process, a very thin Ti and a noble metal layer were successively deposited on the sample without breaking vacuum. The Ti serves the purpose of an adhesion layer, and the cap prevents oxidation of the pads. The UBM was then made thicker by depositing additional noble metal in a subsequent operation. The metal pattern was transferred to the sample by a liftoff process.

Following liftoff, the samples were treated with a manufacturer-recommended photoresist stripper and then cleaned with standard solvents.

(ii) Electroless UBM process:

UBM deposition by electroless plating is a simpler alternative to the lithography-based process. The use of electroless nickel and immersion gold as a bondpad surface finish has been widely used in the IC and PCB industries. The plating process consists of six steps, which must be carried out in quick succession. The first three are cleaning processes that rid the Al pads of any contaminants (organic and inorganic) and also remove aluminum oxide. Following cleaning, the samples were immersed in a zincation bath for deposition of a Zn seed layer, and then in autocatalytic Ni deposition and immersion Au deposition solutions. The final stack is Au/Ni/Al. Electroless plating occurs on all exposed metal surfaces on the chip. Different chip manufacturers use variants of Al-alloy (such as Al-Si, Al-Cu, etc.) as the aluminum cap layer on chip bondpads, and so the plating process was optimized depending on this material. While UBM deposition by electroless plating is a far more efficient process than the lithography-based UBM process, it requires very tight process control to obtain the desired result. Figure 4 shows images of electroless Ni/Au UBM plated on Al bondpads.

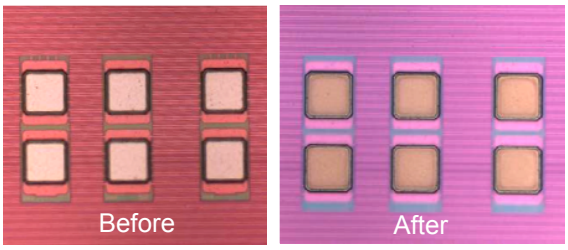

(a)

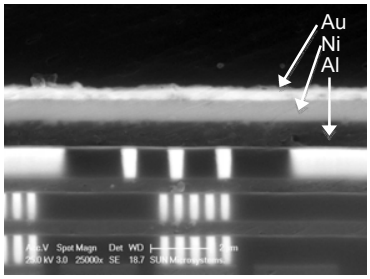

(b)

Figure 4. (a) Top-view of 30 µm wide hybrid-bonding pads with electroless Ni/Au UBM. (b) Cross-section view of pad stack-up after plating.

Microsolder

Once UBM has been deposited on the chips by vapor deposition or electroless plating, the samples are ready for microsolder bumping. This process is similar to the lithography-based UBM process. The bump pattern was first exposed onto a photoresist on the die. Following this, the microsolder alloy was deposited and the samples immersed in solvent to facilitate patterned bumping by liftoff. The samples were then cleaned with photoresist stripper and solvents. On completion of this process, the samples were ready for dicing and bonding. Figure 5 shows a top-down image of a microsolder bump interconnect.

III. Flip Chip Integrated Si-Photonic Bridge Chips

Photonic bridge chips are hybrid-integrated components that embody an electrical interface for local communications and optical access for global interconnections via either fiber or waveguides on another routing layer. Such bridge chips were attained by attachment of a VLSI chip to a silicon photonic chip. While face-to-face flip-chip bonding using microsolder bump interconnects forms a suitably low-parasitic electrical connection between the chips, doing so physically blocks access for optical I/Os to and from the photonic chip. To get around this, the photonic bridge chips were assembled in diving board configurations, such that at least one edge of the photonic chip was sufficiently exposed to provide access for surface-normal or edge-coupling optical I/Os (Figure 6).

Flip-chip Bonding Process

Flip-chip attachment of silicon photonic bridge chips was achieved via thermocompression bonding of the microsolder bumps on one chip to corresponding pads on the other. A two-way camera installed on the flip-chip bonder helped achieve an alignment accuracy of +/- 1 µm. Post-alignment a low-viscosity epoxy was dispensed on the chip surface and the chips were brought together and compressed under several pounds of loading pressure at modest temperature. For the different diving board configurations, various bond-tool heads were designed specifically to ensure that the applied pressure did not generate a moment force on the top chip, causing it to tilt relative to the bottom chip; even a slight tilt would cause high-resistance connections or worse, opens.

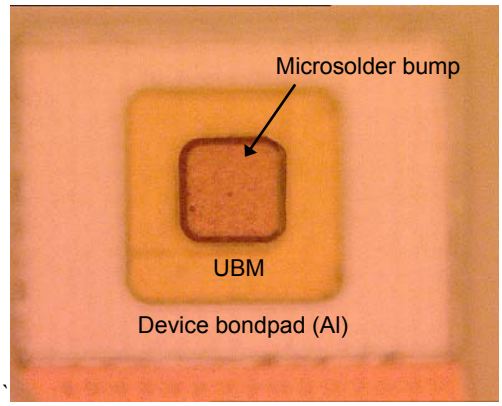

Figure 5. Top-down optical microscope image of a microsolder bump interconnect on a photonic device bondpad.

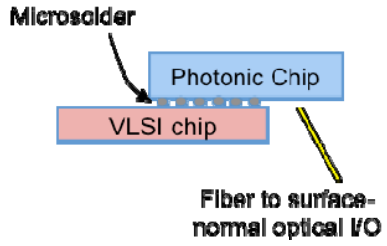

(a)

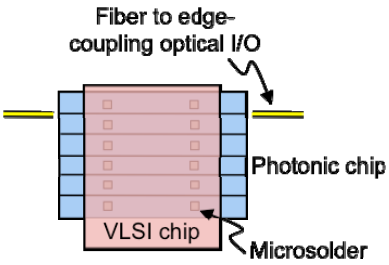

(b)

Figure 6. Schematics of flip-chip integrated silicon photonic bridge chips in diving board configurations.

Bonding Results

Figure 7 shows a die-photo of the 90 nm bulk CMOS IC that was designed and built for use in the photonic bridges. A total of 330 low-power driver and receiver circuits were arranged along a single row at the north end and along four columns on either side of the 5.25 mm x 6.3 mm chip; high-speed electrical input/output pads were located at the south end. The shaded area indicates unrelated test circuits. Having the driver and receiver circuits laid out in this configuration allowed for flexibility in using the same VLSI chip for

bonding to photonic chips with surface-normal as well as edge-coupling optical I/Os. For flip-chip bonding verification, test-pad pairs were placed at the top of each column of circuits. Additional details on the chip design are described in [7].

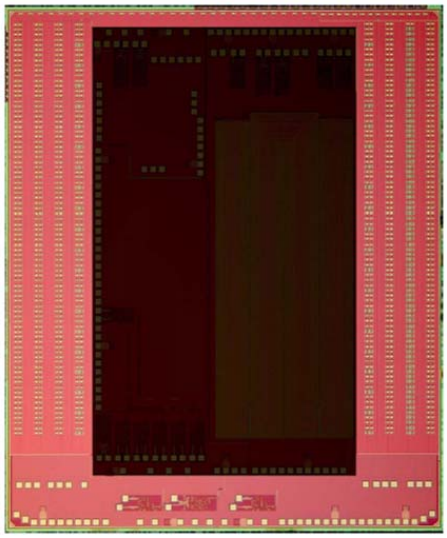

Figure 7. Die shot of the 90 nm bulk CMOS technology VLSI chip that was used in the photonic bridge components.

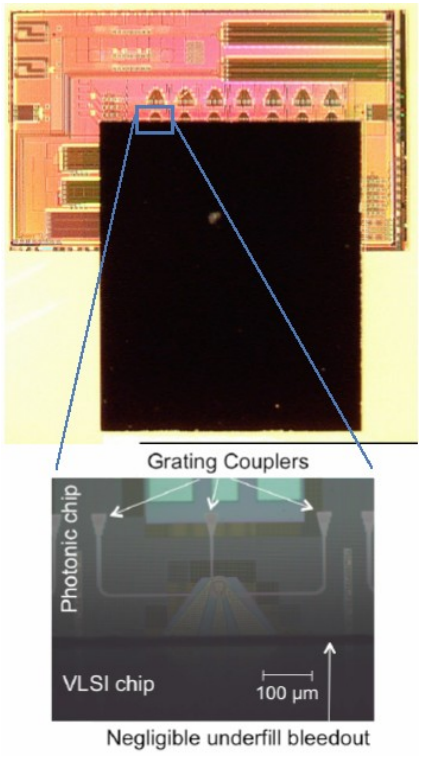

Figure 8. Image showing the 90 nm bulk CMOS VLSI chip flip-chip assembled to a 130 nm SOI-CMOS photonic chip. The inset image shows the photonic chip extending from underneath the VLSI chip. Here the underfill edge bead was controlled to be within a few microns of the chip edge.

The various photonic chips used in bridge chip construction included resonant microring modulator chips and germanium waveguide photodetector chips built in both a 130 nm CMOS process technology as well as an SOI-Photonics process technology.

Figure 8 shows an image of the VLSI chip flip-chip bonded to a Si photonic chip fabricated at Luxtera and containing an array of microring modulator devices. These modulators used grating-coupler terminated waveguides for optical input and output. The inset image shows the area of the photonic chip that extends beyond the VLSI chip and contains the optical waveguides and grating couplers. Arrangements such as this, where the grating couplers ended up only a short distance away from the VLSI chip edge, required tight control over the extent of the edge-bead from the underfill epoxy used in the flip-chip bonding process. Any epoxy on the grating couplers would degrade output efficiency. As seen on this image, the edge-bead extended less than 20 µm from the VLSI chip edge.

Figure 9 shows an image of the VLSI chip flip-chip bonded to a silicon photonic chip fabricated at Kotura also containing an array of microring modulator devices. These devices utilized in-plane edge coupling waveguides, seen as horizontal lines in the image, for optical input and output.

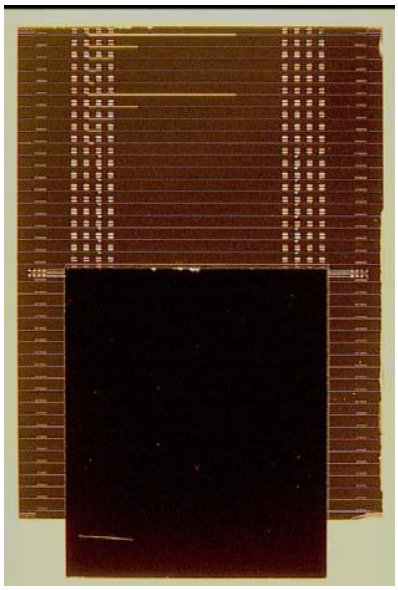

Figure 9. Image of the 90 nm bulk CMOS VLSI chip flip-chip assembled to an SOI-Photonic chip bearing an array of microring modulators with edge-coupling waveguides for optical input and output.

Bonded Chip Verification and Bump Resistance Measurements

The first row of pads in all the vertical columns on either side of the VLSI chip were shorted in pairs, *i.e.* a total of 4 pair of pads on either side. Each pad pair formed one half of a test structure that was added on the VLSI chips to enable measurement of bump resistance after flip-chip bonding. The photonic chips were designed to have corresponding hybrid bond test pads with pairs of traces connecting each pad to probe pads near the photonic chip edge. When two chips were correctly aligned and flip-chip bonded, a 4-point resistance

measurement for each bump pair was enabled as a metric for bonding verification (Figure 10).

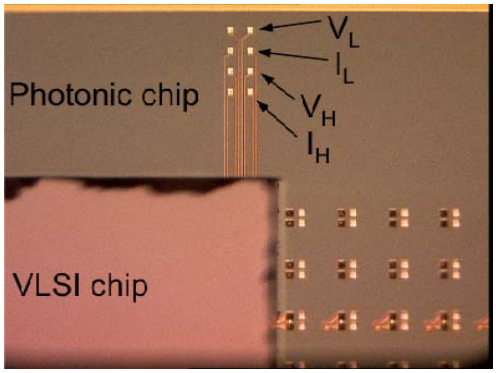

Figure 10. Image showing bump resistance measurement and flip-chip bonding verification site on a photonic bridge chip.

Using these resistance test structures, we measured an average bump resistance of 0.6 ohm across several hybrid bridge-chips; the minimum measured bump resistance was 0.05 ohm. We expect the bump capacitance to be on the order of 20-25 fF.

Considerations for Hybrid Integration of Edge-Coupling Photonic Chips

Hybrid integration of photonic chips having edge-coupling waveguides requires additional preparation prior to use. Specifically, the input and output waveguide facets need to be polished and coated with an anti-reflection (AR) material for optimal fiber-to-waveguide optical coupling. Typically, these processes are carried out on singulated die, but doing so would require handling the tiny singulated chips for microsolder post-processing, while also preserving the quality of the waveguide facets – a tall order. As such, the adopted process-of-record was to carry out waveguide polishing and AR-coating (ARC) after flip-chip integration. This process itself is not without risk; polishing is a harsh mechanical process, and components must be held with a significant amount of force to keep them stable while the polishing wheel is in operation. In this case, the component is a flip-chip bonded hybrid, which would not tolerate significant surface-normal or shear forces over the bond-zone. In addition, standard AR-coating processes require a 200°C environment, which is beyond the tolerance limit of the underfill used.

Two processing innovations helped overcome these hurdles. First, for the polishing process, a small piece of the polishing gasket material was taped down to the exposed area of the edge-coupling photonic chip in the hybrid. Using this as both a spacer and load-bearing segment, the hybrid was mounted onto a polishing fixture to carry out the process. With such fixturing, no load is applied directly on the VLSI chip and the bond-zone, which contains the microsolder bump connections. Second, for the AR-coating, a room temperature process was developed and qualified for use on the hybrid chips. Additionally, a temporary protective cover was placed over the exposed wirebond pads on the VLSI chip.

Table 1 shows measurements of bump resistance at different stages of the hybrid-integration process development for edge-coupling photonic chips: after flip-chip bonding, after polishing, and after ARC. No significant change in resistance was measured and the bumps continued to maintain their superior low-resistance characteristic.

In addition, the polished and AR-coated waveguides were measured to have a -14dB reflection.

Table I. Microsolder bump resistance measurements at various stages of the hybrid-integration process for edge-coupling photonic chips

Sample #	R_{bump} after flip-chip	R_{bump} after polishing	R_{bump} after polishing & ARC
1	0.4 Ω	0.8 Ω	N/A
	0.5 Ω	0.8 Ω	
2	0.2 Ω	0.2 Ω	N/A
	1.1 Ω	1.2 Ω	
3	0.9 Ω	N/A	1.0 Ω
	1.6 Ω		1.9 Ω
4	0.6 Ω	N/A	0.7 Ω

IV. Board Assembly

For testability and demonstration, the hybrid-bonded bridge chips were assembled in a chip-on-board configuration following flip-chip attachment. Physically, the high-speed test boards were designed to have a small ledge extending from one edge of the board where the bridge chip, with the VLSI chip facing up, was die-attached using a conductive epoxy. The electrical I/Os of the VLSI chip were then wire-bonded to the test board using an Al wedge-bonding process. An example of a board-assembled bridge-chip is shown in Figure 11.

V. Photonic Bridge-Chip Performance

The integrated photonic-bridge transmitters and receivers were tested in a clocked digital link environment.

Photonic Tx

Input to the modulator drivers came from a second board through a flex connector, and additional traces after the connector on the test board. Because resonator based devices could be sensitive to ambient temperature variations, a copper heat sink was directly attached to the modulator chip, and the ambient airflow to the transmitter was controlled. This simple passive thermal management accomplished very stable operation. Performance results of two significant variants of photonic Tx bridge-chips are summarized below:

(i) 130 nm CMOS-photonic ring modulator (15 μm ring diameter) integrated with a 90 nm bulk CMOS driver circuit: The integrated transmitter achieved error-free operation with an ultra-low power consumption directly measured below 400 fJ/bit at a datarate of 5 Gbps. The transmitter achieved stable error-free performance over 3.5 days, and a BER better than 10^{-15}. Further performance details are available in in [8]. The output "eye" diagram is shown in Figure 12.

(ii) SOI-photonic ring modulator (15 μm ring diameter) integrated with a 90 nm bulk CMOS driver circuit: This

transmitter also achieved error-free transmission with a BER better than 10^{-13} without active tuning or temperature control. The total transmitter power consumption was measured to be 320 fJ/bit. Further performance details are available in [9]. The output "eye" diagram is shown in Figure 13.

In these measurements, the energy-per-bit includes the energy consumed by the photonic device, the driver circuit as well as any circuit parasitics including the flip-chip bond-pads. In the eye diagrams shown in Figures 12 and 13, the double falling edge is due to a timing error related to non-optimal clock loading, easily corrected with timing and layout optimization.

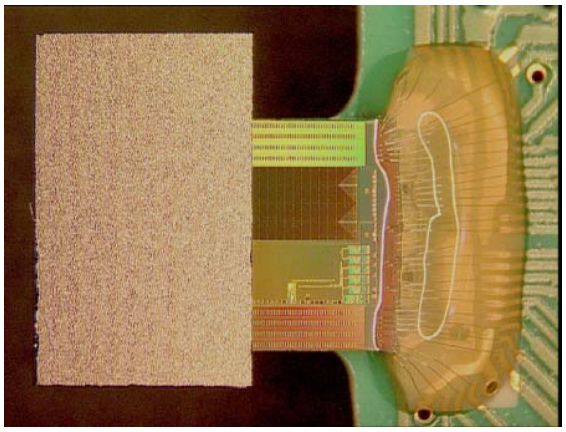

(a)

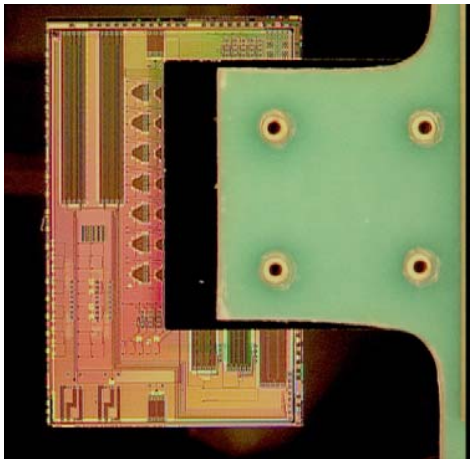

(b)

Figure 11. Images of a photonic transmitter bridge chip assembled onto a test board. (a) View from the topside with the VLSI chip face-up, (b) View from the bottom-side with the photonic chip face-up. The mounting ledge on the PCB is also seen here.

Photonic RX

The receiver was tested with a commercial lightwave transmitter driven by 2^{31}-1 PRBS data from a pattern generator at 5 Gbps. The high speed digital output from the bridge chip passed through short wire-bonds, then traveled over several inches of trace leading to a flex connector, then additional traces on a second test board and finally to a BER tester through high quality RF cables.

A photonic Rx bridge chip comprising of a 130 nm CMOS-photonic germanium waveguide detector bonded to a 90 nm bulk CMOS receiver circuit achieved a sensitivity of -18.9 dBm at 5Gbps for a BER of 10^{-12} (see Figure 14). The Rx bridge-chip was measured to have an ultralow power consumption of 690 fJ/bit. Additional performance details are available in [10].

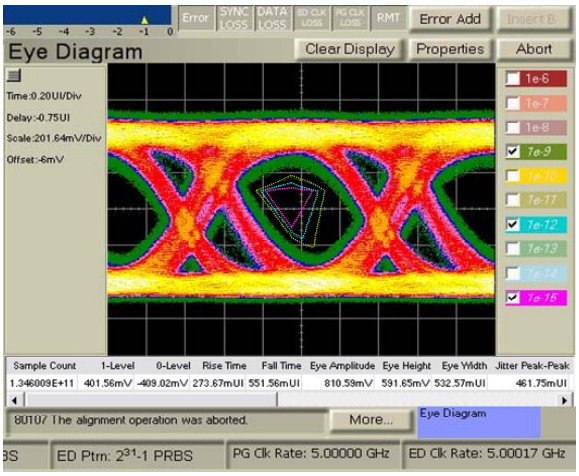

Figure 12. "Eye" diagram of photonic Tx output comprising of a 130 nm CMOS-photonic modulator driven by a 90 nm bulk CMOS driver [8].

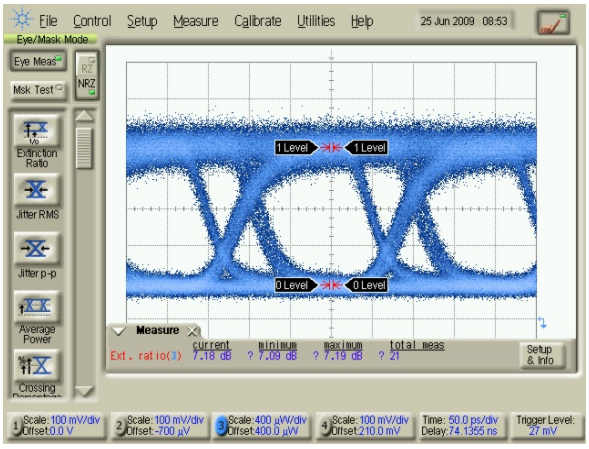

Figure 13. Output "Eye" diagram of a photonic Tx bridge chip comprising of a SOI-photonic ring modulator driven by a 90 nm bulk CMOS driver [9].

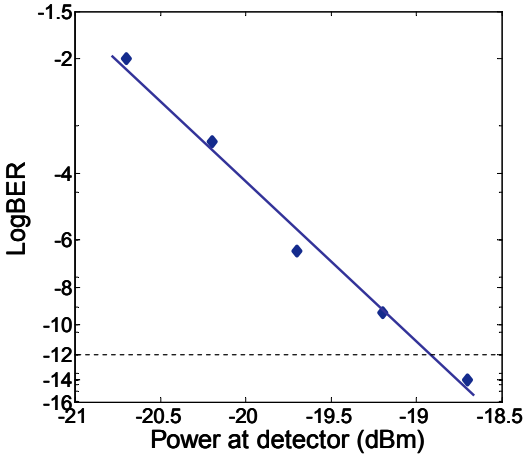

Figure 14. Photonic Rx bridge chip sensitivity measurement results [10].

VI. Conclusions & Future Directions

Hybrid integration using ultralow parasitic microsolder interconnections is an enabling packaging solution that brings together best-in-class technologies for high-speed and low-power photonic transmitters and receivers. The microsolder bump interconnects achieved in this work have a measured average resistance of 0.6 Ω/bump and an estimated capacitance of 20-25 fF/bump. These have been built at pitches as low as 40 µm and have been proven to a 10 µm bump size. We have successfully commissioned this microsolder and hybrid integration technology to build ultralow power si-photonic bridge-chip transmitters and receivers. The successful demonstration of bridge chips with a transmitted energy/bit as low as 320 fJ/bit and received energy/bit as low as 690 fJ/bit are significant steps towards a complete sub-picojoule/bit link that is expected to be critical for future inter-chip and intra-chip interconnect applications.

Acknowledgments

The authors would like to thank Mr. John Simons and Mr. Somphet Siphengphone for their assistance with the microfabrication.

This material is based upon work supported, in part, by DARPA under Agreement No. HR0011-08-09-0001. The authors thank Dr. Jag Shah of DARPA MTO for his inspiration and support of this program. The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense. Approved for Public Release. Distribution Unlimited.

References

1. J.D. Meindl, R. Venkatesan, J.A. Davis, J. Joyner, A. Naeemi, P. Zarkesh-Ha, M. Bakir, T. Mule, P.A. Kohl, and K.P. Martin, "Interconnecting Device Opportunities for Gigascale Integration," in *Proceedings of IEEE International Electron Devices Meeting*, Dec. 2001.
2. A.V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J.E. Cunningham, "Computer Systems Based on Silicon Photonic Interconnects," *Proceedings of the IEEE*, vol. 97, no. 7, pp. 1337-1361, Jul. 2009.
3. K.W. Goossen, J.E. Cunningham, and W.Y. Jan, "GaAs 850nm Modulators Solder-Bonded to Silicon," *IEEE Photon. Technol. Lett.,* vol. 5, p. 716, 1993.
4. A.V. Krishnamoorthy, L.M.F. Chirovsky, W.S. Hobson, R.E. Leibenguth, S.P. Hui, G.J. Zydzik, K.W. Goossen, J.D. Wynn, B.J. Tseng, J. Lopata, J.A. Walker, J.E. Cunningham, L.A. D'Asaro, "Vertical-Cavity Surface-Emitting Lasers Flip-Chip Bonded to Gigabit-per-Second CMOS Circuits," *IEEE Photon. Technol. Lett.*, vol. 11, no. 1, pp. 128-131, Jan 1999.
5. J.E. Cunninghmam, A.V. Krishnamoorthy, I. Shubin, J.G. Mitchell, and X. Zheng "Aligning Chips Face to Face for Dense Capacitive Communication," in *Coupled Data Techniques,* Springer, R. Ho and R. Drost Ed., 2010, in Print.
6. A.L. Lentine, K.W. Goossen, J.A. Walker, L.M.F. Chirovsky, L.A. D'Asaro, S.P. Hui, B.J. Tseng, R.E. Leibenguth, J.E. Cunningham, W.Y. Jan, J.-M, Kuo, D.W. Dahringer, D.P. Kossives, D.D. Bacon, G.Livescu, R.L. Morrison, R.A. Novotny, and D.B. Buchholz, "High-Speed Optoelectronic VLSI Switching Chip With >4000 Optical I/O Based on Flip-Chip Bonding of MQW Modulators and Detectors to Silicon CMOS," *IEEE Journ. Select. Topics in Quantum Electron.,* vol. 2, iss.1, pp. 77-84, 1996.
7. R. Ho, J. Lexau, F. Liu, D. Patil, R. Hopkins, E. Alon, N. Pinckney, P. Amberg, X. Zheng, J.E. Cunningham, and A.V. Krishnamoorthy, "Circuits for Silicon Photonics on a "Macrochip"," in *Proceedings of IEEE Asian Solid State Circuits Conference*, Nov. 2009.
8. X. Zheng, J. Lexau, Y. Luo, H.D. Thacker, T. Pinguet, A. Mekis, G. Li, J. Shi, P. Amberg, N. Pinckney, K. Raj, R. Ho, J.E. Cunningham, and A.V. Krishnamoorthy, "Ultra-low-energy All-CMOS Modulator Integrated with Driver," *Optics Express*, vol. 18, iss. 3, pp. 3059-3070, 2010.
9. G. Li, X. Zheng, J. Lexau, Y. Luo, H.D. Thacker, P. Dong, S. Liao, M. Asghari, J. Yao, J. Shi, P. Amberg, N. Pinckney, K. Raj, R. Ho, J.E. Cunningham, and A.V. Krishnamoorthy, "Ultralow-Power High-Performance Si Photonic Transmitter," in *Proceedings of OFC/NFOEC*, Mar 2010.
10. X. Zheng, F. Liu, D. Patil, H.D. Thacker, Y. Luo, T. Pinguet, A. Mekis, J. Yao, G. Li, J. Shi, K. Raj, J. Lexau, E. Alon, R. Ho, J.E. Cunningham, A.V. Krishnamoorthy, "A Sub-Picojoule-Per-Bit CMOS Photonic Receiver for Densely Integrated Systems," *Optics Express*, vol. 18, iss. 1, pp. 204-211, Feb. 2010.