Hybrid Integration of Silicon Nanophotonics with 40nm-CMOS VLSI Drivers and Receivers

Hiren D. Thacker, Ivan Shubin, Ying Luo, Joannes Costa, Jon Lexau[*], Xuezhe Zheng, Guoliang Li, Jin Yao, Jieda Li, Dinesh Patil[□], Frankie Liu[*], Ron Ho[*], Dazeng Feng[°], Mehdi Asghari[°], Thierry Pinguet[+], Kannan Raj, James G. Mitchell[*], Ashok V. Krishnamoorthy and John E. Cunningham

Oracle Labs, San Diego, CA
[*]Oracle Labs, Menlo Park, CA
[□]Rambus, previously with Oracle Labs, Menlo Park, CA
[°]Kotura Inc., Monterey Park, CA
[+]Luxtera Inc., Carlsbad, CA

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

Abstract

Oracle's scalable hybrid integration technology platform enables continuing improvements in performance and energy efficiency of photonic bridge chips by leveraging advanced CMOS technologies with maximum flexibility, which is critical for developing ultralow power high-performance photonic interconnects for future computing systems. Herein, we report on our second generation of photonic bridge chips comprising electronic drivers and receivers built in 40 nm bulk CMOS technology attached to nanophotonic devices, fabricated using SOI-photonic and 130 nm SOI-CMOS photonic technologies. Hybrid integration by flip-chip bonding is enabled by microsolder bump interconnects scaled down from our previous generation effort and fabricated on singulated dies by a novel batch processing technique based on component embedding. Generation-on-generation, the hybrid integrated Tx and Rx bridge chips achieved 2.3x and 1.7x improvement in energy efficiency, respectively, while operating at 2x the datarate (10 Gbps).

I. Introduction

Within the past decade, the semiconductor computing industry has developed multicore and multithreaded core processors to overcome the challenges and shrinking benefits of traditional technology scaling. Multichip systems built using these components will require massive amount of off-chip bandwidth and low latency chip-to-chip links at the lowest energy cost possible. Wavelength-division multiplexed (WDM) silicon photonics has the potential to provide a solution for this immense interconnect problem [1]. Within the Ultraperformance Nanophotonic Intrachip Communication (UNIC) program at Oracle, we are aggressively building a portfolio of technologies in active and passive nanophotonic devices (modulators, detectors, WDM components), circuits, and multichip packaging to achieve the vision of 15 Gbps, 300 fJ/bit photonic links between computing elements in a large array "Macrochip" [1]1.

Juxtaposition of nanophotonic devices and VLSI circuits on a common silicon substrate interconnected by low parasitic on-chip interconnects is perhaps the most intimate method for integrating electronics and photonics. Realistically, however, the design and process integration of nanophotonic devices on CMOS platforms presently lags a few technology generations behind the state-of-the-art technology. Trading-off component performance for monolithic co-integration by using lesser than the best available CMOS technology would not lead to the lowest energy per bit links. Instead, hybrid integration of best-in-breed photonic and electronic components, built on individually optimized technology platforms, is a much more pragmatic approach to achieving peak performance. In such a hybrid integrated component, the chip-to-chip interconnects must have ultralow parasitics.

Previously [2], we reported on the hybrid integration by flip-chip bonding of 130 nm SOI-CMOS photonic and SOI-photonic chips to 90 nm bulk CMOS ICs. The transmitter (Tx) and receiver (Rx) bridge chips resulting from that effort achieved ultralow-power performance of 1.6 mW [2] and 3.4 mW [4], respectively, at 5 Gbps thereby validating this approach for creating low-power integrated components. The integration was enabled by microsolder bump interconnects fabricated on one or both chips, which had an average resistance of 0.6 Ω/bump and an estimated capacitance of 20-25 fF/bump [2].

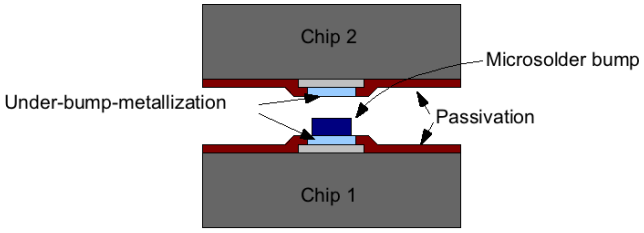

Figure 1. Schematic of a hybrid bond.

Figure 1 shows a schematic of a hybrid bond. It comprises under-bump-metallization (UBM) and a microsolder bump. The UBM provides a stable, low resistance contact to the chip's I/O pads, provides a strong adhesion interface between the die bondpad and bump materials, and prevents diffusion of the bump materials into the chip. The microsolder bumps sit atop the UBM and get fused into opposing pads during flip-chip bonding, thereby creating a high conductivity chip-

to-chip connection. The hybrid integration and microsolder methods deployed here follow an approach originally reported for the integration of III-V modulators [5] and VCSELs [6] to silicon. Microsolder bumps as small as 10 µm on a 25 µm pitch [6] have been previously fabricated and can be further scaled [7]. Very high bump yield of better than 99.95% has also been demonstrated with this approach [8], [9].

This paper reports on the challenges and accomplishments in the hybrid integration of the second generation of ultralow energy Tx and Rx Si-photonic bridge chips targeted for operation at 10 Gbps. These include the fabrication of scaled down microsolder bump interconnects on and flip-chip bonding to 40 nm CMOS VLSI ICs with extreme low-K (ELK) interlayer dielectrics, which are known to be mechanically weak [10]. Section II contains a brief introduction of the chips to be integrated as well as an overview of the integration task. The processes for deposition of under-bump-metallization and microsolder are described in Sections III and IV, respectively. The latter was achieved by a batch-processing technique, wherein several singulated chips were embedded in a silicone "puck". The flip-chip bonding process and bump characterization data are summarized in Section V. Section VI details the performance of the hybrid-integrated bridge Tx and Rx components, and is followed by Conclusions.

II. Photonic Bridge Chips

Aggressive reduction in the power consumption of an optical link while also increasing channel bandwidth can be achieved not only by building more energy efficient photonic devices and electronic circuits but also by taking advantage of the raw performance gains offered by smaller geometries possible in an advanced process technology node. For this reason, the next generation of ultralow energy photonic bridge chips employed VLSI circuits built in TSMC's 40 nm bulk CMOS technology and designed to operate at 10 Gbps. Photonic bridge chips are hybrid-integrated components that embody an electrical interface for local communications and optical access for chip-to-chip or inter-chip global interconnections via either optical fibers or waveguides on another routing layer. As reported in [2], photonic bridge chips must be integrated in diving board configurations to provide access for surface-normal or edge-coupling optical I/Os.

The 40 nm VLSI chip (named "XNP"), designed for the second generation of silicon photonic bridge chips, contained a total of 106 low-power driver (Tx) and receiver (Rx) circuits arranged in four columns along the western edge and three columns along the eastern edge of the 5.2 mm × 4.5 mm chip, respectively (Figure 2). Any of these individual circuits can be selectively activated during operation. The chip also includes an on-chip pseudorandom binary sequence (PRBS) generator and checker for built-in self-test (BIST) functionality. To lower parasitic capacitance, hybrid bond pad dimensions were scaled down by roughly 20%, to 25 µm square. Details of the circuit design and chip operation may be found in [11], [12].

The XNP VLSI chip was designed to be bonding-compatible with several nanophotonic device chips built on different platforms. Figure 3 shows two such dies built in a SOI-photonic and a 130 nm SOI-CMOS photonic technology, respectively. The die in Figure 3(a) contains arrays of racetrack ring resonator modulators, while Figure 3(b) shows a die containing an array of Ge photodetectors.

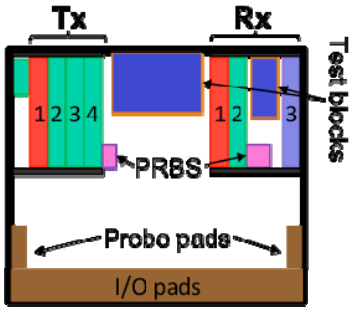

Figure 2. Floorplan of the Oracle 40nm XNP chip. The die contains low-energy driver and receiver circuits as well as built-in self-test capability.

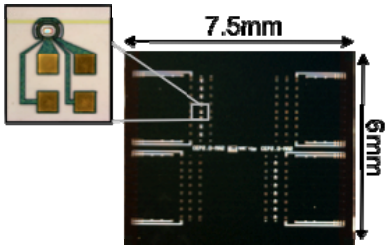

(a)

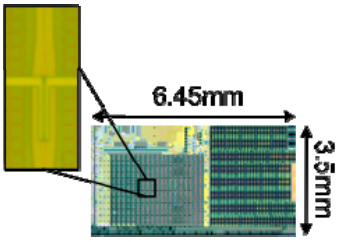

(b)

Figure 3. Examples of silicon nanophotonic dies used in building photonic Tx and Rx bridge chips. (a) SOI-photonic chip containing arrays of racetrack ring modulators. (b) 130 nm SOI-CMOS photonic chip bearing an array of Ge photodetectors.

In the present hybrid integration effort microsolder bump interconnects (i.e. UBM and bumps) were fabricated on the XNP VLSI chip while the various photonic chips were processed only to deposit UBM. As the VLSI chip is a common element in each hybrid integrated photonic bridge, it was deemed productive to develop and optimize the microsolder process on one component and avoid the overhead of designing additional photolithographic masks as well as developing and optimizing microsolder process modules for every flavor of photonic chip. The fabrication processes are discussed next.

III. Electroless Plating of UBM

Electroless ("e-less") plating is an attractive process for UBM deposition owing to its short process time and selectivity of metal deposition, which eliminates the need for lithographic patterning. In the present effort, all UBM deposition was by e-less plating of Ni/Au. The plating process module consists of six steps that must be carried out in quick succession. The first three are cleaning processes that rid the on-chip Al pads of any contaminants (organic, inorganic, and aluminum oxide) and present a clean Al surface for the build-up steps. Next, a thin layer of Zinc is deposited as a precursor for plating, followed by Ni and Au.

Figure 4 shows e-less plated UBM on hybrid bond pads of chips from a couple of different technology platforms used in hybrid integration. This included plating on chips built on a 40 nm bulk CMOS technology platform (

Figure 4(a)), which is known to have pads with higher RMS roughness and a porous dielectric stack. A cross-section SEM image of the UBM stack on a single XNP hybrid bondpad is shown in Figure 5.

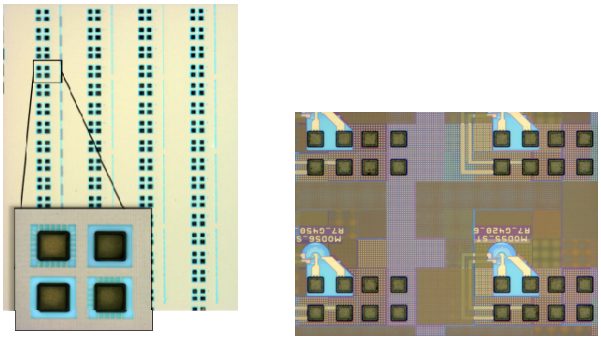

(a) (b)

Figure 4. Images of e-less plated (Ni/Au) hybrid bond pads on (a) a 40 nm XNP chip, and (b) a 130 nm SOI-CMOS photonic chip.

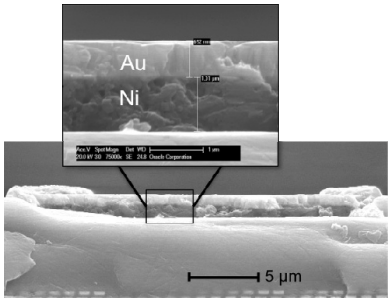

Figure 5. Cross-section SEM image of a single 25 µm hybrid bondpad on the 40 nm XNP chip. The inset shows a higher magnification image of the e-less plated Ni/Au UBM.

The UBM is targeted to be thick enough to make the metal pads coplanar with the overlapping thin-film passivation. Doing so allows the entire height of the microsolder bump to be used in the chip-to-chip connection; typically UBM thickness was between 0.7–1.7 µm. On the VLSI chips, at least 0.3 µm Au was required for the I/O pads to be process-compatible with a subsequent Al wedge wire-bonding process used to interconnect hybrids to high-speed characterization boards.

The e-less plating process is sensitive to a number of parameters, such as pH level, temperature, and concentration of the bath as well as process time, environmental conditions, and solution agitation. Fortunately, these can be managed by exercising strict process control.

Unfortunately, the quality and metallurgy of the Al pads on both electronic and photonic chips received for hybrid integration can vary; different manufacturers use variants of Al-alloys (such as Al-Si, Al-Cu, etc.) as the cap layer on chip bondpads. Therefore, the plating process had to be optimized per technology platform and in some cases, tweaked per lot of received chips. Insufficient cleaning and conditioning of the Al pads typically resulted in discontinuous UBM coverage and even corrosion Figure 6).

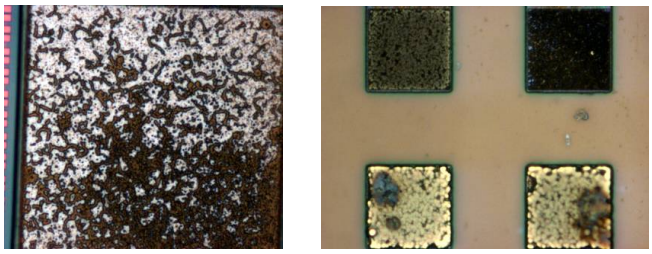

(a) (b)

Figure 6. (a) Image of an Al pad with discontinuous e-less plating. (b) Image of plated pads with corrosion from residual contamination.

illustrates another phenomenon observed when plating the 40 nm XNP chips. Pads on the same chip, even those sitting next to each other, plated differently; a distinct difference was observed not only in the coloration of pads but also in plated thickness. It was further noted that the variation was dependent on the circuit nodes to which pads were connected. From our experiments and literature review [13], it was concluded that exposure of the chips to ambient light before and during plating can alter the plating electrochemical reactions, which affects the quality and quantity of the plating deposition. Consequently, plating in near-dark conditions was implemented to achieve a uniform high-quality UBM layer across all pads on a die (see Figure 4 (a).

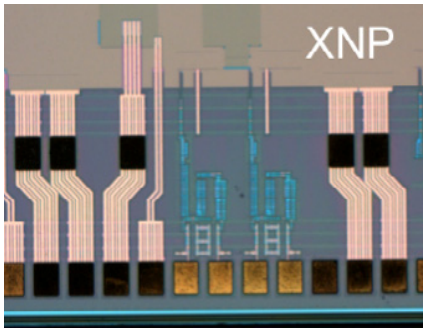

Figure 7. Photo of plated pads on the XNP chip illustrating dependence of the e-less plating reactions on the circuit node to which the pads to be plated were connected.

IV. Batch Processing of Microsolder Bumps

Microsolder bumps are fabricated atop the UBM on bondpads dedicated for hybrid bonding. The process flow is based on lithography, bump deposition, and liftoff. As such, the overall microsolder performance and yield are inherently dependent on the quality and uniformity of the photoresist film coating on a chip's surface.

A photograph of the XNP chip, only 5.2 mm × 4.5 mm in size, is shown in Figure 8(a). These chips were manufactured on a foundry shuttle run, which is a cost effective solution for small footprint and low-volume chip fabrication; however, shuttle-run chips are delivered in singulated form instead of full wafers. Post processing a small-footprint chip such as XNP to form microsolder bumps is a challenging task because of handling alone. Here, this task is made even tougher due to circuits and hybrid bondpads placed only 100 µm away from chip's diced edge (Figure 8(b)).

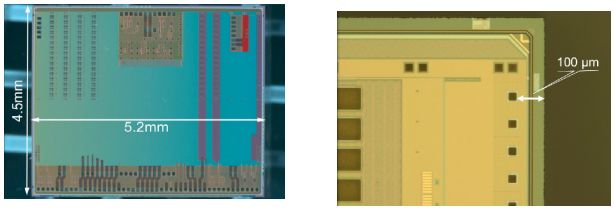

(a) (b)

Figure 8. (a) Top view of the 40 nm XNP VLSI chip. (b) A higher magnification image of the XNP chip's top right corner. In this region, the microsolder bump sites are laid out as close as 100 µm from the diced edge of the chip.

In a photolithography process sequence, deposition of photoresist by a standard spin-on process typically yields an edge "bead" that is much thicker than the target photoresist thickness. Figure 9(a) shows the result of spinning photoresist onto a single XNP chip using typical microsolder process parameters. As is evident from the surface profile plot in Figure 9(b) and Figure 9(c), the photoresist around the edge is nearly 40 µm thicker than at the center; additionally, this undesirable topology extends more than 1 mm laterally towards the chip-center, covering a majority of the microsolder sites. Owing to this varying topology, fine, accurate and uniform resist patterning, to the order of microsolder dimensions (15-20 µm), would not be possible with either projection or contact lithography. In the case of the latter, which is used in our post-CMOS back-end processing, the photomask could not even be brought into required proximity with the resist coated chip surface.

An alternate technique was therefore developed to address the resist planarity issue on a small-footprint die. This process borrows from technology known in the packaging industry as component embedding [14]. With this method, several different chips can be embedded into an intermediate medium atop an organic substrate or a wafer. The embedded chips are planar with the intermediate material, and thus, can be processed in parallel, irrespective of their individual dimensions, using standard wafer-scale microfabrication techniques.

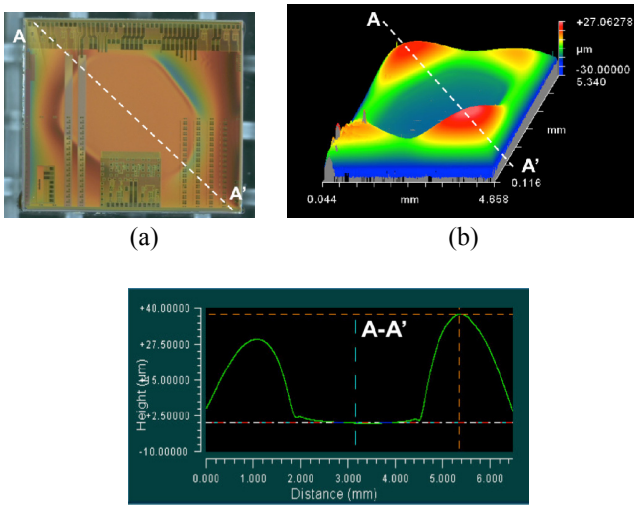

(a) (b)

(c)

Figure 9. Photoresist topology on the XNP chip from a typical spin-on process. (a) Top view, (b) 3D profile, and (c) surface profile across section A-A'.

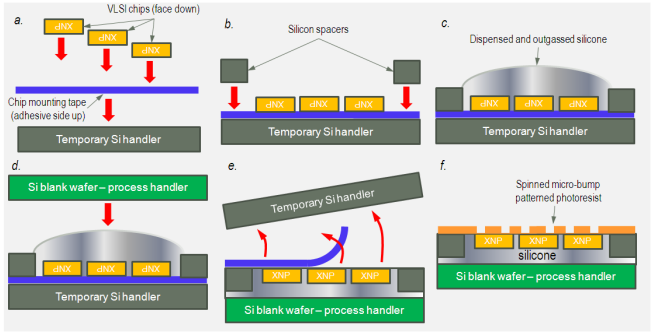

Figure 10. Process flow for batching XNP chips for microsolder processing.

The process flow for embedding or "batching" multiple chips into a "puck" is schematically depicted in Figre 10. A temporary tacky tape is first spread flat across a temporary Si handler wafer with the adhesive side face-up (Figure 10(a)). It is ensured that the tape and tape-wafer interface are free of any contamination or air pockets. Next, XNP chips with UBM are placed face-down in the desired array configuration, with approximately 1 mm separation between them, onto the adhesive side of the tape. Dummy silicon spacers are distributed and secured in a similar manner around the periphery of the array (Figure 10(b). After confirming that the chips are in full contact with the mounting tape, a RTV silicone compound from Dow Corning is poured from above to completely immerse the chip array (Figure 10(c)). The silicone material is outgassed in a vacuum chamber and a blank Si "process handler" wafer is placed atop the silicone to flatten it out and force the material into the gaps between the chips and spacers (Figure 10(d)). The assembly is then placed in a 65ºC oven for 24 hours to allow the silicone to cure fully. Next, the temporary handler wafer is easily removed and the chip mounting tape is peeled off. This tape does not leave any residue and does not affect the on-chip circuits or metal pads.

Any excess silicone around the process handler wafer edge is cut off and the resulting wafer-sized puck is ready for microsolder processing (Figure 10(f)).

Several experiments were conducted to identify the most optimal distribution of XNP and spacer chips. Process-worthy pucks containing chips arranged in 1×2, 1×3, 2×2, and 3×3 arrays, with or without the spacer chips nearby, have been demonstrated. The most process-optimized and frequently used configuration was a 1×3 chip array on a 2" diameter silicon process handler wafer with four spacers cleaved and situated to resemble an original 2" substrate (Figure 11(a)).

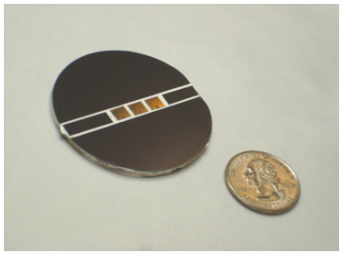

(a)

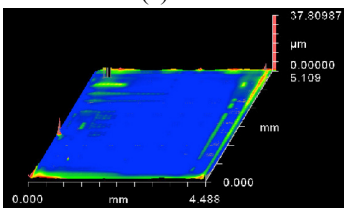

(b)

Figure 11. (a) A completed puck of three VLSI chips batched together for microsolder bumping. (b) 3D profile of photoresist topology on a chip embedded in the puck. The resist thickness variation was measured to be less than 2 µm.

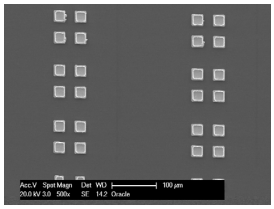

(a)

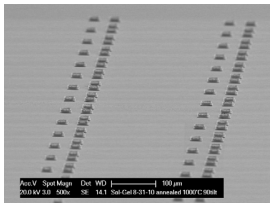

(b)

Figure 12. SEM images of arrayed microsolder interconnects on a 40 nm XNP VLSI chip. The bumps are 20 µm square on a 48.6 µm pitch. Each bump consists of an e-less plated UBM layer and a microsolder bump fabricated using the described batch process.

This chip-batching process was successful in mitigating the edge bead problem. Figure 11(b) shows a 3D profiler scan of the typical photoresist topology across a single chip in the puck from Figure 11(a). The resulting resist surface variation across the chip is less than 2 µm, which meets the stringent requirements for microsolder fabrication. With optimization of the photoresist baking profiles and proximity aligner operation, the multi-chip pucks were readily processed to fabricate arrays of microsolder bumps across the XNP chips with high yield and reproducibility (Figure 12), including on the bondpad column at the very edge of the chip. After microsolder processing, the chips can be readily released from the puck without any residual silicone adhering to its surface.

This is the first time that our microsolder bump interconnects have been process integrated with 40 nm CMOS ICs.

V. Flip Chip Bonding

The processed chips were bonded using a thermal compression bonding process with an alignment accuracy of +/- 1 µm. Two exemplary hybrid integrated components are shown in Figure 13 and Figure 14. The former consists of an XNP chip bonded to an array of racetrack ring modulators fabricated in an SOI-photonic technology to yield a hybrid integrated Tx, and the latter consists of an XNP chip bonded to an array of Ge photodetectors fabricated on a 130 nm SOI-CMOS photonic chip to yield a hybrid integrated Rx.

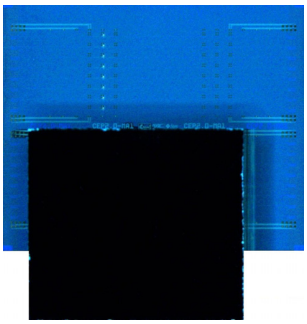

Figure 13. Photograph of a hybrid integrated transmitter bridge chip consisting of electronic driver circuits in 40 nm CMOS (face down) bonded to racetrack ring modulators in SOI-photonic technology (face up). This component employed edge-coupled optical I/Os.

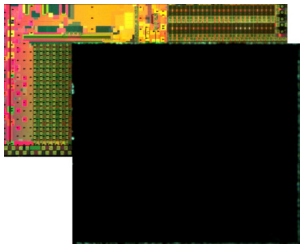

Figure 14. Photograph of a hybrid integrated receiver bridge chip consisting of electronic receiver circuits in 40 nm CMOS (face down) bonded to Ge photodetectors in 130 nm SOI-CMOS photonic technology (face-up). This component employed surface-normal grating couplers as optical I/Os.

The XNP and various silicon nanophotonic chips used in hybrid integration were co-designed to include a number of test structures at matching locations to allow the measurement of bond planarity and bump resistance after flip-chip bonding. Using these, the average microsolder bump resistance was measured to be only 0.37 Ω/bump. This is a nearly 2x reduction in resistance relative to measurements from our hybrid integrated bridge chips reported last year [2].

From simulations, we estimate the total bump parasitic capacitance to be between 15-21 fF. Therefore, the interconnect bandwidth of a microsolder bump interconnect, assuming a lumped RC model, far exceeds 10 THz and should easily support digital data rates exceeding 40 Gbps.

VI. Hybrid Integrated Bridge Chip Performance

The hybrid bonded chip assemblies were die attached and wire bonded to a printed circuit board (PCB) for performance characterization (Figure 15). To mimic an inter-chip communications application, we used the on-chip pseudorandom binary sequence (PRBS) generator as the data source to drive the nanophotonic ring modulators, and the on-chip PRBS checker to measure the bit-error-rate (BER) for the data received by the receiver. External clock sources were used to clock both transmitter and receiver.

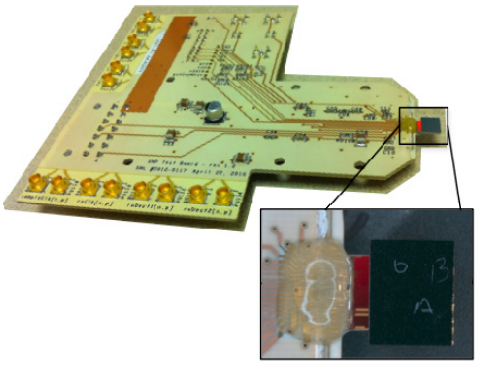

Figure 15. Hybrid integrated photonic bridge chip wirebonded to a high-speed board for characterization.

The characterization results of an exemplary Tx bridge chip and an exemplary Rx bridge chip are described below. Detailed descriptions of the driver and receiver design, on-chip BIST functions, test setup, and component characterization can be found in [11] and [12].

Transmitter (Tx) Performance

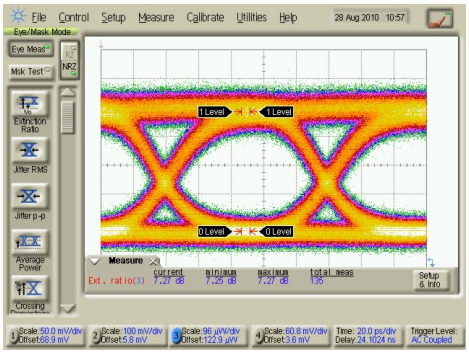

Figure 16. Optical "eye" diagram at 10 Gbps for a hybrid integrated transmitter component, consisting of 2 V cascode drivers in 40 nm CMOS driving racetrack ring modulators in a SOI-photonic technology [12].

Figure 16 shows a measured optical "eye" diagram for 10 Gbps data transmission from the hybrid integrated silicon photonic transmitter using a racetrack ring modulator (Figure 13); the eye is open with >7 dB extinction ratio. The transmitter was measured to have a BER better than 10^{-12} and consumed a mere 1.35 mW, excluding the laser.

Receiver (Rx) Performance

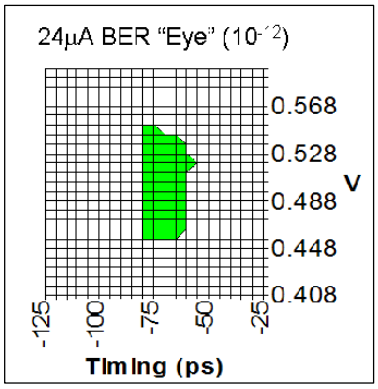

Figure 17. Performance of a hybrid integrated silicon photonic receiver at 10 Gbps. The plot shows the measured BER "eye" at an average input photocurrent of 24 µA [12].

The hybrid-integrated receiver shown in Figure 14 was wirebonded to the high-speed board and tested on a station with a lensed fiber probe for the optical input. Figure 17 plots the virtual "eye" opening at 10 Gbps as measured by the on-chip BIST circuits. The green area corresponds to a BER under 10^{-12}. This hybrid integrated receiver component was measured to consume only 3.95 mW.

Conclusions

In this paper, we have presented a second generation of silicon-photonic bridge chips comprising a 40 nm bulk CMOS VLSI chip hybrid integrated with silicon nanophotonic devices. Exemplary Tx and Rx bridge chips built in this configuration achieved ultralow energy performance of 1.35 mW and 3.95 mW, respectively, while operating at 10 Gbps. To enable this integration, microsolder bumps were process-integrated onto 40 nm technology CMOS chips with an ELK ILD stack. This was achieved by implementing an e-less Ni/Au UBM process and a novel batch processing method, based on component embedding, for microsolder patterning and deposition. This batch processing method allows parallel processing of multiple chips of varying dimensions while being inherently ambivalent to the process technologies used to build them. With continuing improvements in CMOS technology, silicon photonic devices, and microsolder scaling we expect to develop even more efficient and higher bandwidth photonic transmitters and receivers for use in high performance inter/intra-chip WDM photonic links.

Acknowledgments

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

The authors also thank Dr. Jag Shah of DARPA for his inspiration and support of this program. This material is based upon work supported, in part, by DARPA under Agreement No. HR0011-08-09-000. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Approved for public release. Distribution unlimited.

References

1. 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," *Proc. IEEE*, vol. 97, pp. 1337-1361, Jul. 2009.
2. H.D. Thacker, Y. Luo, J. Shi, I. Shubin, J. Lexau, X. Zheng, G. Li, J. Yao, J. Costa, T. Pinguet, A. Mekis, P. Dong, S. Liao, D. Feng, M. Asghari, R. Ho, K. Raj, J.G. Mitchell, A.V. Krishnamoorthy, and J.E. Cunningham, "Flip-Chip Integrated Silicon Photonic Bridge Chips for Sub-Picojoule Per Bit Optical Links," in *Proc. Electron. Compon. and Technol. Conf.*,Las Vegas, NV, 2010, pp. 240-246.
3. 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 *Proc. OFC/NFOEC*, San Diego, CA, 2010.
4. 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.
5. K.W. Goossen, J.E. Cunningham, and W.Y. Jan, "GaAs 850nm Modulators Solder-Bonded to Silicon," *IEEE Photonics Technol. Lett.,* vol. 5, p. 776-778, Jul. 1993.
6. 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, and L.A. D'Asaro, "Vertical-Cavity Surface-Emitting Lasers Flip-Chip Bonded to Gigabit-per-Second CMOS Circuits," *IEEE Photonics Technol. Lett.*, vol. 11, pp. 128-131, Jan. 1999.
7. J.E. Cunninghmam, A.V. Krishnamoorthy, I. Shubin, J.G. Mitchell, and X. Zheng, "Aligning Chips Face to Face for Dense Capacitive Communication," *Coupled Data Techniques,* Springer, R. Ho and R. Drost Ed., 2010.
8. 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 J. Sel. Topics Quantum Electron.*, vol. 2, pp. 77-84, Apr. 1996.
9. A.V. Krishnamoorthy, and K.W. Goossen, "Optoelectronic-VLSI: Photonics Integrated with VLSI Circuits," *IEEE J. Sel. Topics Quantum Electron.*, vol. 4, pp. 899-912, Nov. 1998.
10. C.C. Hsia, "The Quest of Porous ELK Materials for High Performance Logic Technologies," *Microelectronic Engineering,* vol. 83., iss. 11-12, pp. 2055-2058, 2006.
11. F. Liu, D. Patil, J. Lexau, P. Amberg, M. Dayringer, J. Gainsley, H.F. Moghadam, X. Zheng, J. Cunningham, A. Krishnamoorthy, E. Alon, R. Ho, "10 Gbps, 530 fJ/b Optical Transceiver Circuits in 40nm CMOS," submitted, *Proc. IEEE Symp. VLSI Circuits*, 2011.
12. X. Zheng, D. Patil, J. Lexau, F. Liu, G. Li, H. Thacker, Y. Luo, I. Shubin, J. Li, J. Yao, P. Dong, D. Feng, M. Asghari, T. Pinguet, A. Mekis, P. Amberg, M. Dayringer, J. Gainsley, H.F. Moghadam, E. Alon, K. Raj, R. Ho, J. Cunningham, and A. Krishnamoorthy, "Ultra-Efficient 10Gb/s Hybrid Integrated Silicon Photonic Transmitter and Receiver," *Optics Express*, to be published.
13. A.J.G. Strandjord, M. Johnson, H. Lu, D. Lawhead, R. Hanson, and R. Yassie, "Electroless Nickel-Gold Reliability UBM, Flipchip, and WLCSP, (Part I of III)," *Proc. of Int. Microelectron. Packag. Soc.*, Oct. 2006.
14. M. Brunnbauer, E. Furgut, G. Beer, T. Meyer, H. Hedler, J. Belonio, E. Nomura, K. Kiuchi, and K. Kobayashi, "An Embedded Device Technology based on a Molded Reconfigured Wafer," in *Proc. Electron. Compon. and Technol. Conf.*, San Diego, CA, 2006, pp. 547-551.