

A Monolithically-Integrated Optical Receiver in Standard 45-nm SOI

Michael Georgas, Jason Orcutt, Rajeev J. Ram and Vladimir Stojanović
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: mgeorgas at mit.edu

Abstract—A monolithically-integrated optical receiver for low-energy on-chip and off-chip communication is presented. The monolithic photodiode integration enables the energy-efficient and high-sensitivity sense-amplifier-based receiver design. The receiver is characterized *in situ* and shown to operate with μA -sensitivity at 3.5 Gb/s with a power consumption of 180 μW (52 fJ/bit) and area of 108 μm^2 . This work demonstrates that photonics and electronics can be jointly integrated in a standard 45-nm SOI process.

I. INTRODUCTION

In order to harness the potential of emerging many-core processor systems, the communication fabric between cores and shared off-chip memory must provide high-throughput at low power and footprint costs, overcoming chip power constraints and I/O pin limitations. Monolithically-integrated silicon-photonics offers a dense, wavelength-division multiplexed fabric with orders of magnitude better energy-efficiency and bandwidth density than electrical interconnects [1]. However, advanced process design rules place significant constraints on the integration of photonics and electronics.

The optical receiver, a necessary component in any optical link, has traditionally been designed as a discrete component for optical fiber communication. Gain and responsivity could be optimized through material selection, while packaging resulted in poor parasitic capacitance. In order to mitigate the gain-bandwidth limitation at the dominant pole of the input node, transimpedance amplifiers were implemented to lower the input resistance, R_{in} , while preserving a large transimpedance gain, R_{TIA} [2].

More recently, integrated photonics has addressed chip I/O bottlenecks through hybrid-packaged solutions [3], [4]. While more sensitive and energy-efficient than discrete receivers, the photodiode (PD) and parasitic capacitance of hybrid-packaged designs is still relatively large. A capacitance of 90 fF is reported in [3]. [5] describes a 25 fF PD capacitance (with parasitics) connected through a 20 fF microsoldier bump, leading to a receiver sensitivity of 9 μA but energy-cost of 690 fJ/bit at 5 Gb/s. This increases system laser power and dominates over-all link costs, making it less competitive with electrical solutions already at 1 pJ/bit [6].

In this paper we present an optical, sense-amplifier-based data receiver with a monolithically-integrated photodetector. Tight integration of the photodetector and the latching sense-

amplifier enables both low-energy operation and high input-sensitivity. These metrics are important, as future many-core communication fabrics will host tens of thousands of receivers per processor die, and input sensitivity maps directly through the optical link loss to the laser power requirement.

II. PHOTONIC LINKS

An example of a dense wavelength-division multiplexing (DWDM), monolithically-integrated photonic link is shown in Figure 1. A continuous wave (CW), multi- λ laser is coupled from an optical fiber onto the chip through a vertical grating coupler. The light is then routed throughout the chip along waveguides fabricated using either gate poly-silicon or the SOI body. Ring-resonant modulators, driven by integrated modulator drivers, modulate data onto a particular wavelength-channel by on-off keying. An optical clock signal could also be forwarded along with data. The modulated light is routed to another location on the die (e.g. core-to-core) or to another die (e.g. socket-to-socket). At the destination, the ring-tuning control block selects the channel to be removed from the optical bus by setting the resonance of a drop ring filter. An optical receiver, such as the one presented in this work, then converts the data back into the electrical domain by detecting the PD photocurrent.

III. RECEIVER ARCHITECTURE

The receiver architecture (Figure 2) consists of a PD connected differentially across a latching sense-amplifier (LSA), followed by a dynamic-to-static (DS) converter and an on-chip high-speed digital testing backend. The receiver operates in two clock phases, receiving one bit per clock period.

In the PD (Figure 2g,j) we make use of P+ SiGe, which is integrated in the SOI process for PMOS strain engineering and is suitable for optical absorption in the near-IR range [7]. The photodiode is extremely compact and has an estimated capacitance of 10 fF.

The LSA (Figure 2a) senses the differential photocurrent and makes a bit decision. During the reset phase ($\Phi=0$), the LSA's nodes pre-charge high. During the decision phase ($\Phi=1$), the two branches, $M_{1,3,5}$ and $M_{2,4,6}$, discharge. If an *optical-1* is received, photocurrent flows from node IN- to IN+, slowing the discharge of branch $M_{1,3,5}$ and causing it to latch *high*. Otherwise, imbalance programmed through offset




Fig. 1. An example optical link with chip-to-chip and intra-chip communication links shown. A CW laser source is coupled onto Chip A through a vertical grating. Two ring-resonant modulators imprint data onto two wavelength-channels, λ_0 and λ_1 , which propagate along the waveguide. The bus is routed over an optical fiber to Chip B. The drop rings on Chip B are each tuned to either λ_0 and λ_1 to select that channel from the bus and direct it to the correct data receiver. A second set of wavelengths, λ_2 and λ_3 carry data from Chip B to Chip A.

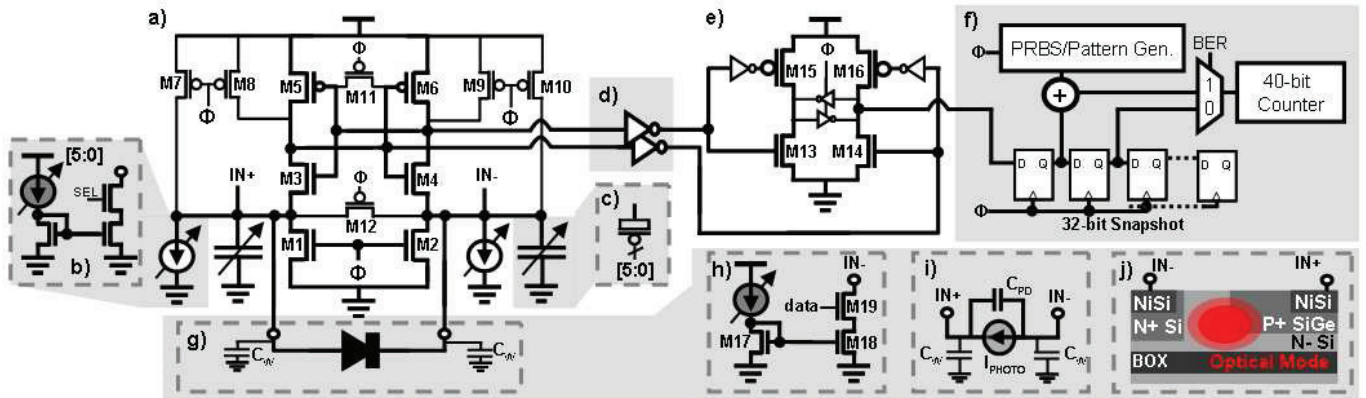


Fig. 2. Optical data receiver architecture. The LSA (a) is followed by an output buffer stage (d) and dynamic-to-static converter (e), before being fed into the digital backend infrastructure (f). The chip has receivers connected to integrated PDs (g) or electrical diode-emulation circuits (h). The simulation model is shown in (i). A cross-section of the implemented PD and optical mode is shown in (j).

compensation causes branch $M_{1,3,5}$ to latch low. The LSA transistors are sized according to [8] in order to balance speed and sensitivity. In particular, transistors $M_{3,4}$ are sized large relative to $M_{5,6}$. This lowers the trip-point voltage of the cross-coupled inverters, ensuring that they do not activate too early. Offset compensation is implemented as programmable current-steering (Figure 2b) and capacitive (Figure 2c) DACs [9], for coarse- and fine-compensation, respectively.

Figure 2h shows a diode-emulation circuit that is used to characterize the receiver's performance when decoupled from the optical devices. When the input *data* is 1, the circuit pulls current from IN-, emulating the photocurrent sourced from that node. A 0-bit sources no current. The diode-emulation circuit is driven by a pattern generator on a separate, programmable clock phase from the rest of the receiver.

To provide qualitative analysis of the impact of parasitic capacitances and operation frequency on the receiver sensitivity, an equivalent model of the LSA is shown in Figure 3.

Figure 3a shows the input nodes at the end of LSA reset ($t = 0$), pre-charged *high*. I_{cm} models $M_{1,2}$ pulling down on the input nodes until cross-coupled inverters M_{3-6} turn on. C_w represents the wiring capacitance from the PD to the receiver. The model divides the decision phase into two steps: integration, and evaluation (Figure 3b). During the integration

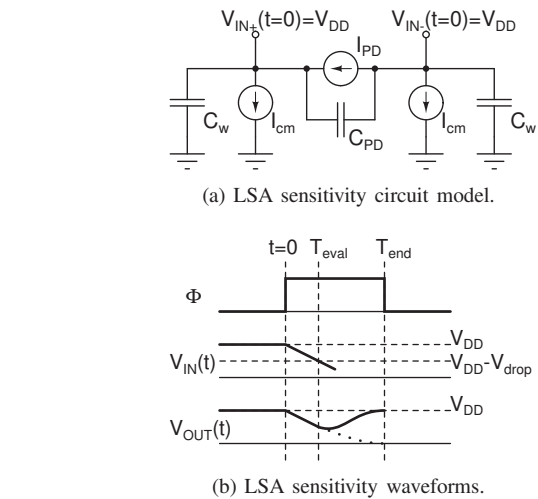


Fig. 3. LSA sensitivity model.

phase of duration T_{eval} (Equation 1), the photocurrent is integrated across $C_{int} = C_{PD} + C_w/2$, resulting in a voltage difference, V_{diff} , at the onset of evaluation (Equation 3).

$$T_{eval} = \frac{C_w V_{drop}}{I_{cm}} \quad (1)$$


$$V_{diff} = \frac{I_{PD}T_{eval}}{C_{int}} \quad (2)$$

$$= \frac{I_{PD}C_wV_{drop}}{C_{int}I_{cm}} \quad (3)$$


$$V_{out} = AV_{diff}e^{\frac{T_{end}-T_{eval}}{\tau}} \quad (4)$$

$$I_{PD} = \frac{V_{out}C_{int}I_{cm}}{AC_wV_{drop}}e^{-\frac{T_{end}}{\tau}}e^{C_w\frac{V_{drop}}{I_{cm}\tau}} \quad (5)$$

During the evaluation phase, V_{diff} regenerates exponentially until T_{end} according to Equation 4. Re-arranging, the current-sensitivity of the receiver can be expressed by Equation 5.



(a) Effect of C_w on receiver sensitivity. $C_{PD} = 10fF$.



(b) Effect of C_{PD} on receiver sensitivity. $C_w = 2.5fF$.

Fig. 4. LSA sensitivity model.

Figure 4a shows through extracted simulation that for high data-rates where the exponential is not completely settling, wire capacitance, C_w , delays the onset of evaluation, shortening the evaluation time and therefore demanding exponentially more input photocurrent. Figure 4b shows that PD capacitance, C_{PD} , reduces V_{diff} linearly, demanding only proportionally more photocurrent (Figure 4b). As our PD was implemented on the same die as the receiver, the low-metal-layer routing between the PD and receiver results in a small C_w , assumed here to be ≈ 2.5 fF. The proposed topology may suffer in scenarios where a second die provides the optical transport layer, necessitating through-silicon vias (TSV) or microsolder bumps where C_w may increase above 20 fF [10]. Figure 4a

shows that for C_w in this range and data rates above 4 Gbps, the sensitivity becomes prohibitively poor.

The output of the LSA is buffered (Figure 2d) to isolate the LSA decision nodes from the data-dependent capacitance looking into the DS (Figure 2e).

The bits stored in the DS are fed into the on-chip digital test backend for *in situ* processing (Figure 2). The backend, consisting of synthesized PRBS and pattern generators, snapshots, and counters, gathers bit-error-rate and receiver decision threshold data and exports only statistics off-chip.

IV. MEASURED RESULTS

The monolithically-integrated data receiver was fabricated in a standard 45-nm SOI process, as a part of a flexible electronic-photonics test vehicle.

Figure 5 shows two DC photocurrents generated by a 1310-nm wavelength laser, coupled into the chip through a vertical coupler and horizontal waveguide made with front-end body Si. The receiver's threshold is swept using the offset circuitry (Figure 2b,c) while recording the output decision statistics. Photocurrent values were de-embedded through simulation. Though the receiver was able to detect photocurrent from the PD, a foundry error in the SiGe mask definition limited achievable PD bandwidth.

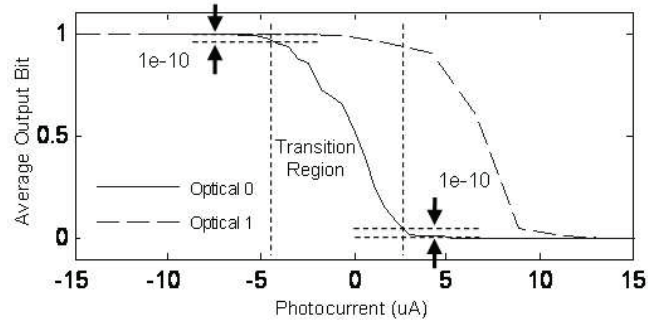



Fig. 5. The measurement shows the receiver's ability to distinguish between a DC *optical-1* and *optical-0*.


Figure 6a shows the receiver's sensitivity vs frequency for different supply voltages. Sensitivity is measured on a PD-connected receiver (Figure 2g,j) as the width of the transition region (Figure 5) of an optical-0. As clock frequency increases, sensitivity degrades exponentially as predicted by our model due to the decrease in T_{end} .

Figure 6b shows the energy-cost of the receiver. The linearity emphasizes the digital design, with power following $P_{digital} = fCV_{DD}^2$, keeping the receiver energy-cost $\approx 50fJ/b$ across a range of frequencies.

Figure 7 shows the bit-error-rate eye diagram of the receiver when configured with a PD-emulation circuit (Figure 2h). Clock phase and receiver threshold were swept for a 31-bit PRBS data sequence at 3.5 Gb/s and a supply of 1.1 V, and error statistic were gathered *in situ* using the digital backend. Clock rates above 3.7 GHz caused the digital testing backend to fail.



(a) Sensitivity.



(b) Energy cost.

Fig. 6. Measured data receiver performance.

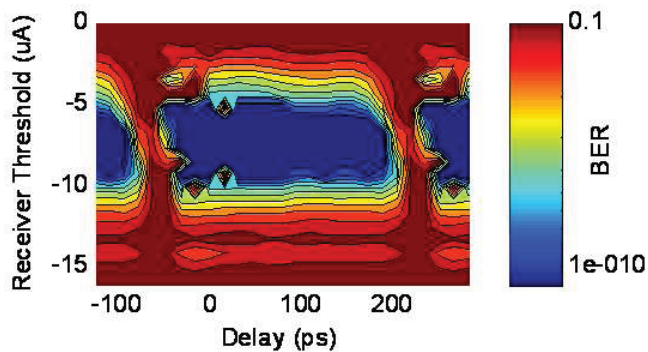


Fig. 7. Electrical *in situ* eye-diagram at 3.5 Gb/s.

A die photo is shown in Figure 8. The chip contains 72 test cells that implement combinations of optical modulators and receivers. Each receiver has a circuit area of $108 \mu\text{m}^2$ and PD area of $416 \mu\text{m}^2$.

V. CONCLUSION

Integrated photonics has emerged as an I/O technology that can meet the throughput demands of future many-core processors. In this work, the monolithic integration of the photodetector enables the design of a fully-digital, low-energy receiver with high input sensitivity. A qualitative model of the receiver provides intuition for the impact of different PD integration scenarios on the receiver's sensitivity performance.




Fig. 8. Die, backside photonic cell and receiver photos.

The sense-amplifier-based latching receiver is shown to operate with μA -sensitivity at 3.5 Gb/s with an energy-efficiency of 52 fJ/b. This work demonstrates the first monolithic electronic-photonics integration in a sub-100-nm standard SOI process.

ACKNOWLEDGMENT

The authors would like to thank the Integrated Photonics teams at MIT and UC Boulder. The work was supported in part by DARPA, NSF, FCRP ICF, Trusted Foundry, Intel, MIT CICS, Santec and NSERC.

REFERENCES

- [1] C. Batten, A. Joshi, J. Orcutt, C. Holzwarth, M. Popovic, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic, "Building manycore processor-to-DRAM networks with monolithic CMOS silicon photonics," *Micro, IEEE*, vol. PP, no. 99, p. 1, 2009.
- [2] C. Kromer, G. Sialm, T. Morf, M. Schmatz, F. Ellinger, D. Erni, and H. Jackel, "A low-power 20-GHz 52-dB-Ohm transimpedance amplifier in 80-nm CMOS," *JSSC*, vol. 39, no. 6, pp. 885–894, June 2004.
- [3] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang, "Optical I/O technology for tera-scale computing," in *ISSCC*, feb. 2009, pp. 468–469,469a.
- [4] A. Ryljakov, C. Schow, B. Lee, W. Green, J. Van Campenhout, M. Yang, F. Doany, S. Assefa, C. Jahnes, J. Kash, and Y. Vlasov, "A 3.9ns 8.9mw 4x4 silicon photonic switch hybrid integrated with cmos driver," in *ISSCC*, feb. 2011, pp. 222–224.
- [5] G. Li, X. Zheng, J. Lexau, Y. Luo, H. Thacker, T. Pinguet, P. Dong, D. Feng, S. Liao, R. Shafiqi, and et al., "Ultralow-power silicon photonic interconnect for high-performance computing systems," *Proceedings of SPIE*, vol. 7607, no. Optoelectronic Interconnects and Component Integration IX, pp. 760 703–760 703–15, 2010.
- [6] K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, E. Suzuki, N. Masuda, T. Takemoto, F. Yuki, and T. Saito, "A 12.3-mw 12.5-gb/s complete transceiver in 65-nm cmos process," *JSSC*, vol. 45, no. 12, pp. 2838–2849, dec. 2010.
- [7] J. Polleux and C. Rumelhard, "Optical absorption coefficient determination and physical modelling of strained SiGe/Si photodetectors," in *High Performance Electron Devices for Microwave and Optoelectronic Applications, 8th IEEE International Symposium on*, 2000, pp. 167–172.
- [8] S. Mukhopadhyay, R. Joshi, K. Kim, and C.-T. Chuang, "Variability analysis for sub-100nm PD/SOI sense-amplifier," *ISQED*, pp. 488–491, March 2008.
- [9] K.-L. Wong and C.-K. Yang, "Offset compensation in comparators with minimum input-referred supply noise," *JSSC*, vol. 39, no. 5, pp. 837–840, may 2004.
- [10] J.-S. Kim, C. S. Oh, H. Lee, D. Lee, H.-R. Hwang, S. Hwang, B. Na, J. Moon, J.-G. Kim, H. Park, J.-W. Ryu, K. Park, S.-K. Kang, S.-Y. Kim, H. Kim, J.-M. Bang, H. Cho, M. Jang, C. Han, J.-B. Lee, K. Kyung, J.-S. Choi, and Y.-H. Jun, "A 1.2V 12.8GB/s 2Gb mobile wide-I/O DRAM with 4x128 I/Os using TSV-based stacking," in *ISSCC*, feb. 2011, pp. 496–498.