Skew Reduction for Synchronous OE-VLSI Receiver Applications

Michael B. Venditti, *Member, IEEE*, Joshua D. Schwartz, *Student Member, IEEE*, and David V. Plant, *Member, IEEE*

***Abstract*—A dc photocurrent rejection technique for optically differential receivers is proposed to reduce the sensitivity of receiver latency to variations in average input photocurrent. This is particularly attractive in optoelectronic very large scale integration applications, where array-scale nonuniformity is significant. Using optically enabled ring oscillator experiments, we demonstrate that the technique is capable of removing between 50%–75% of the skew in large synchronously operated arrays of common-gate amplifier- and transimpedance amplifier-based optical receivers.**

***Index Terms*—Optical receiver, optically enabled ring oscillators (OEROs), optoelectronic very large scale integration (OE-VLSI), receiver latency, receiver skew, synchronous receiver arrays.**

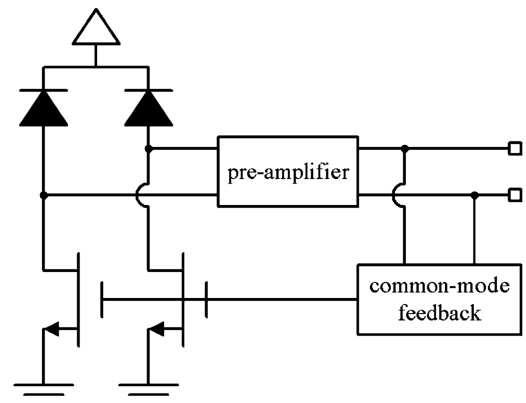

Fig. 1. Illustration of the DCPR technique.

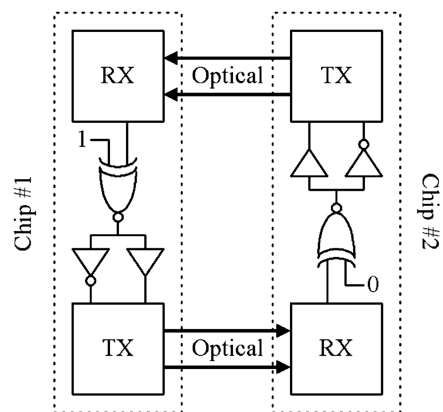

Fig. 2. Illustration of bidirectional optical link implementing the OERO.

PARALLEL synchronous digital links require tight control over latency and interchannel skew in order to obtain the highest possible performance. Reducing interconnect latency and interchannel skew in such systems allows larger regions of synchronous timing on a chip to be employed, simplifying their design [1]. In optoelectronic very large scale integration (OE-VLSI) applications, interchannel skew arises primarily from differences in on-chip electrical path lengths and differences in latency through individual receivers. Differences in on-chip electrical path lengths can be mitigated through systematic design practice [2]. Conversely, achieving uniform latency in OE-VLSI receiver arrays is complicated by basic circuit design requirements and by variables (such as the optical interconnection system, for example) that may not be known *a priori* or that may vary during system operation. The difference in latency between the receivers in an array with the maximum and minimum latencies is defined as skew.

The vast majority of previously reported OE-VLSI receiver designs are optically and electrically single-ended employing small transistor devices [3], [4]. This genre of receivers has no common mode rejection capability and, correspondingly, has a limited dynamic range [7]. Recently, an optically and electrically differential receiver design has been reported [2]. In order to contain the aggregate power dissipation of large arrays, such receivers employ small bias currents on the order of hundreds of microamperes. In contrast, traditional telecommunication receivers typically employ bias currents on the order of one to tens of miliamperes. This results in OE-VLSI receivers having an inherently narrower dynamic range and their operating points are sensitive to changes in the average input photocurrent.

Large-scale optical imaging systems based on, for example, microlens arrays and fiber image guides, are commonly used with large OE-VLSI receiver arrays. In such systems, maintaining a uniform optical power throughput (and hence, a uniform average input photocurrent) across a receiver array is problematic [5]–[7]. Thus, for a large optical receiver array, individual receiver operating points will vary, along with characteristics such as gain, pole frequency, and latency, which are operating point dependent. For example, due to the dominating effect

Manuscript received September 2, 2003; revised January 22, 2004. This work was supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada and industrial partners, through the Agile All-Photonic Networks (AAPN) Research Group and through Graduate Student Fellowships. It was also supported in part by Defense Advaned Research Projects Agency (DARPA) through the Consortium for Optical and Optoelectronic Technologies in Computing (COOP) at George Mason University, and in part by the COOP-Peregrine USC Workshop and Foundry Run.

M. B. Venditti was with the Electrical and Computer Engineering Department, McGill University, Montreal, QC H3A 2A7, Canada. He is now with PMC-Sierra, Inc., Mount Royal, QC H3R 3L5, Canada (e-mail: michael_venditti@pmc-sierra.com).

J. D. Schwartz and D. V. Plant are with the Electrical and Computer Engineering Department, McGill University, Montreal, QC H3A 2A7, Canada.

Digital Object Identifier 10.1109/LPT.2004.827100

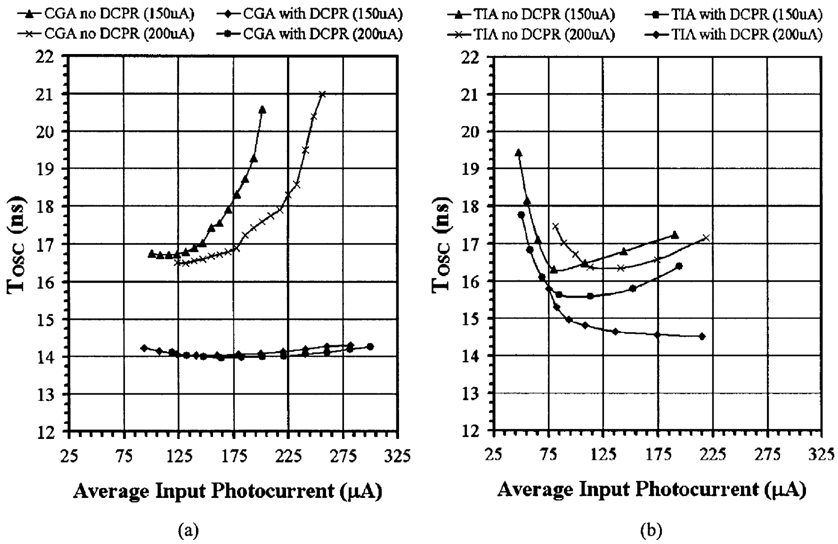

Fig. 3. Measured OERO period of oscillation versus average input photocurrent with and without DCPR circuitry for the (a) CGA- and (b) transimpedance-based receiver. The input photocurrent swing (shown in brackets in the legend) was held constant for each set of measurements.

of the photodetector (PD) capacitance C_{PD}, a preamplifier typically has a dominant pole proportional to g_m/C_{PD}, where g_m is the transconductance of the main amplifying transistor and is proportional to the dc bias current. Additionally, the resistive properties of triode-mode metal–oxide–semiconductor field-effect transistors (MOSFETs) commonly used as active feedback resistor circuit elements in transimpedance preamplifier configurations are strongly affected by an effect known as *dynamic compression*, in which the MOSFET resistance decreases as the average input photocurrent flowing through it increases [8], causing the preamplifier gain and bandwidth to be dependent on the operating point.

Previous work has shown that short optical pulses (i.e., a return-to-zero signaling format) can be used to remove skew and timing jitter in optical links that employ multiple quantum well reflection modulators [9]. Optically single-ended receiver designs lack the means to accommodate a wide range of average input photocurrents, and are particularly susceptible to skew and timing jitter in large arrays [7]. This work presents a technique for dynamic range enhancement and skew reduction in optically differential receiver arrays that employs conventional nonreturn-to-zero signaling based on a dc photocurrent rejection (DCPR) technique. The DCPR technique implemented was adapted from an optically single-ended infrared application to reject ambient light [10] for use in an optically differential configuration, and is illustrated in Fig. 1. An n-channel metal–oxide semiconductor transistor in each input branch of the preamplifier shunts the average input photocurrent from the PDs under the control of a common-mode feedback circuit to maintain a stable operating point for the preamplifier. With the requirement of having differential optical inputs, the modified DCPR technique has a simple and space-efficient implementation and is, thus, suitable for OE-VLSI applications where physical space is typically constrained. The advantages of the DCPR technique were demonstrated using a noninvasive experimental technique based on optically enabled ring oscillators (OEROs) [11].

To evaluate the benefits of the DCPR technique, a test chip was designed and fabricated in a 0.5-μm silicon-on-insulator complimentary metal–oxide–semiconductor process technology from Peregrine Semiconductor. Two fully differential receiver preamplifier designs were considered, one based on a common-gate amplifier (CGA) configuration, and the other on a transimpedance configuration. Two receivers for each preamplifier design (one with and one without DCPR circuitry) were implemented. The test chip was designed to allow an OERO to be formed with a bidirectional optical link between two test chips, as illustrated in Fig. 2. Each side of the link comprised a receiver connected to a transmitter through an XNOR gate that drives a buffer and an inverter. One XNOR gate input was connected to the receiver output, and the other to a control signal that could be set logically high (low) to set the logic function of the XNOR gate to be a buffer (inverter). The XNOR gate control signal for the test chip on either side of the link were set to opposite logic polarities to achieve a logical inversion around the ring. This permitted oscillation to occur when the ring was closed by connecting the transmitter outputs of one chip to the receiver inputs of the other chip, and vice versa. The period of oscillation of the OERO was determined by the latency of the circuit stages along the ring and the optical time of flight between the two chips.

The test chips were packaged in p-i-n grid-array packages along with 1×4 bars of vertical-cavity surface-emitting lasers and PDs, and the packaged chips were optically aligned using conventional optomechanical hardware. One of the optical signals from the setup was tapped from the ring using a beam splitter and converted to an electrical signal using an external detector. Its period was then measured using a digitizing oscilloscope. Additional information on packaging and optical alignment can be found in [11].

Simulations of the OERO using SPICE verified that there was a range of average input photocurrents (corresponding to the dynamic range of each receiver) for which changes in the period of oscillation $(\Delta T_{\mathrm{OSC}})$ corresponded directly to a change in receiver latency (ΔL_{RX}), with the latency in other circuit elements (XNOR gate, buffer, inverter, and transmitter) not materially changing. This, coupled with the fact that the optical time

TABLE I
REDUCTION IN SKEW OFFERED BY DCPR TECHNIQUE

Receiver Design	Input Photocurrent Swing	Average Input Photocurrent Range	Skew		
			No DCPR	With DCPR	Reduction
CGA	150 μA	100 μA to 150 μA	160 ps	55 ps	66%
	200 μA	125 μA to 175 μA	190 ps	80 ps	58%
TIA	150 μA	75 μA to 100 μA	80 ps	20 ps	75%
	200 μA	100 μA to 175 μA	180 ps	90 ps	50%

of flight between chips is invariant with the average input photocurrent, allows ΔL_{RX} to be written as

$$\Delta L_{\mathrm{RX}} \approx \frac{\Delta T_{\mathrm{OSC}}}{2}. \quad (1)$$

For any given pair of average input photocurrents $(I_{\mathrm{PH1}}, I_{\mathrm{PH2}})$ falling within the dynamic range of the receiver, the corresponding ΔL_{RX} can be taken as the worst-case skew between two receivers in an array for which the nonuniformity of the optical system power throughput results in a distribution of average input photocurrents for which the maximum and minimum correspond to I_{PH1} and I_{PH2}.

For each set of experimental measurements using the OERO setup, a constant optical power swing was maintained, and the average transmitted power was swept while recording ΔT_{OSC}. The results are presented in Fig. 3. For ranges of average input photocurrent corresponding approximately to the dynamic range of the receiver designs without DCPR for a given input photocurrent swing, the reduction in skew offered by the DCPR technique is significant, as presented by the results in Table I. For the CGA-based design, a reduction in skew between 58% and 66% was obtained. For the transimpedance-based design, a reduction in skew between 50% and 75% was obtained.

The results of Fig. 3 also clearly indicate that the dynamic ranges of the receiver designs are increased significantly, particularly for the CGA-based receiver and for the transimpedance-based receiver for larger input photocurrent swings.

In conclusion, the stabilization of the preamplifier operating point through the use of the DCPR technique results in a reduced dependence of the receiver latency on the average input photocurrent. This improves the dynamic range of the receiver and results in significantly less skew in synchronous OE-VLSI applications involving large receiver arrays.

REFERENCES

[1] G. A. Keeler, D. Agarwal, C. Debaes, B. E. Nelson, N. C. Helman, H. Thienpont, and D. A. B. Miller, "Optical pump-probe measurements of the latency of silicon CMOS optical interconnects," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 1214–1216, Aug. 2002.

[2] M. B. Venditti, E. Laprise, J. Faucher, P.-O. Laprise, J. E. Lugo, and D. V. Plant, "Design, test, and verification of an optoelectronic-VLSI chip with 540-element receiver/transmitter arrays using differential optical signaling," *IEEE J. Select. Topics Quantum Electron.*, vol. 9, pp. 361–379, Mar./Apr. 2003.

[3] A. C. Walker *et al.*, "Design and construction of an optoelectronic crossbar switch containing a terabit per second free-space optical interconnect," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 236–249, Mar./Apr. 1999.

[4] A. L. Lentine *et al.*, "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. Select. Topics Quantum Electron.*, vol. 2, pp. 77–84, Apr. 1996.

[5] M. Châteauneuf, A. G. Kirk, D. V. Plant, T. Yamamoto, and J. D. Ahearn, "512-channel vertical-cavity surface-emitting laser based free-space optical link," *Appl. Opt.*, vol. 41, pp. 5552–5561, 2002.

[6] T. Maj, A. G. Kirk, D. V. Plant, J. F. Fonstad, K. L. Lear, K. Tatah, M. S. Robinson, and J. A. Trezza, "Interconnection of a two-dimensional array of vertical cavity surface emitting lasers to a receiver array via a fiber image guide," *Appl. Opt.*, vol. 39, pp. 683–689, 2000.

[7] M. B. Venditti and D. V. Plant, "On the design of large receiver and transmitter arrays for OE-VLSI applications," *J. Lightwave Technol.*, vol. 21, pp. 3406–3416, Dec. 2003.

[8] M. Ingels, G. Van der Pals, J. Crols, and M. Steyaert, "A CMOS 18 THz-Ω 240 Mb/s transimpedance amplifier and 155 Mb/s LED-driver for low-cost optical fiber links," *IEEE J. Solid-State Circuits*, vol. 29, pp. 552–1559, Dec. 1994.

[9] G. A. Keeler, B. E. Nelson, D. Agarwal, and D. A. B. Miller, "Skew and jitter removal using short optical pulses for optical interconnection," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 714–716, June 2000.

[10] K. Phang and D. A. Johns, "A CMOS optical preamplifier for wireless infrared communications," *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 852–859, July 1999.

[11] J. D. Schwartz, M. B. Venditti, and D. V. Plant, "Experimental techniques using optically enabled ring oscillators," *Appl. Opt.*, vol. 43, pp. 2456–2461, Apr. 2004.