A 60-Gb/s PAM4 Wireline Receiver with 2-Tap Direct Decision Feedback Equalization Employing Track-and-Regenerate Slicers in 28-nm CMOS

Kuan-Chang Chen, William Wei-Ting Kuo, and Azita Emami
California Institute of Technology, Pasadena, CA 91125

***Abstract*—This paper describes a 4-level pulse-amplitude modulation (PAM4) wireline receiver incorporating a continuous time linear equalizer (CTLE) and a 2-tap direct decision feedback equalizer (DFE). A track-and-regenerate CMOS slicer is proposed and employed in the PAM4 receiver. The reduced delay of the proposed slicer and its full-swing outputs allow the implementation of 2-tap direct decision-feedback equalization at 60-Gb/s with improved energy efficiency and area requirements. Fabricated in 28-nm CMOS technology, the PAM4 receiver achieved BER better than 1E-12 at 60-Gb/s with 1.1 pJ/b energy efficiency measured over a channel of 8.2dB loss at Nyquist rate.**

Keywords—wireline; slicer; comparator; PAM4; receiver; track-and-regenerate; equalization; decision feedback; direct feedback

I. Introduction

PAM4 signaling is an attractive solution for high-speed links with severely bandwidth-limited channels, due to its halved Nyquist frequency compared to non-return-to-zero (NRZ) modulation. However, PAM4 transceivers with multi-level signaling can be challenging to design. At the receiver side, the reduced eye-height in PAM4 sets a more stringent limit for the sensitivity of the decision circuitry. Moreover, at least 3 slicers are needed for making decisions with respect to 3 distinct thresholds [1]. Consequently, the power consumed by the slicers and the loading presented by the slicers are of greater concern in designing PAM4 receivers. Furthermore, when a decision feedback equalizer (DFE) is included in a PAM4 receiver, loop unrolling becomes considerably more power/area hungry due to the exponential increase in the number of slicers; e.g. 12 slicers to unroll the first tap of DFE [2]. In this work, a CMOS track-and-regenerate slicer is proposed and employed in a PAM4 receiver, with the aims to directly close the decision feedback loops of the first 2 taps at 60-Gb/s, offer full-swing outputs, and benefit the overall energy efficiency because of its improved performance in delay and area.

II. Receiver Architecture

A. Overall Architecture

The architecture of the PAM4 receiver is shown in Fig. 1. The analog front-end (AFE) consists of 2-stage continuous time linear equalizer (CTLE). Following the CTLE are the half-rate summers and direct 2-tap DFE. Each summer is connected to four proposed slicers in parallel, among which one slicer is responsible for the eye-monitoring, and the other three slicers

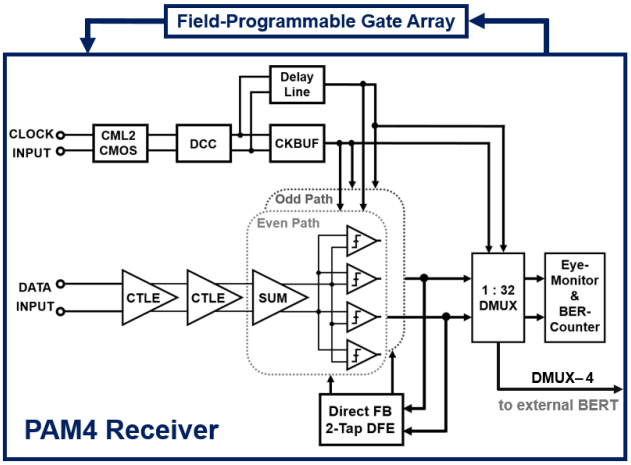

Fig. 1. The overall architecture of the PAM4 receiver.

are dedicated to recovering the analog summer outputs to their corresponding 3-bit thermometer-coded digital levels. The 3-bit thermometer outputs are first directly fed back to the summer in the other path for the first tap of DFE, and then fed back to the summer in the same path with 1-UI delay for the second tap of DFE. The outputs of the slicers are demultiplexed (1:32) for external and on-chip eye-monitor and bit-error-rate (BER) counter to evaluate the eye-opening and the BER, respectively. The clock paths are designed to take in an external pair of half-rate differential clock signals and amplify them to rail-to-rail with duty-cycle correction. Clock buffers and a digitally controllable delay line are included on the chip as the interfaces with the clocked slicers to provide clock phases for data recovery and eye-monitoring.

B. CTLE and Summer Architecture

The CTLE adopts the topology of RC source-degenerated differential amplifier with digital programmability to enable/disable the peaking and adjust the peaking frequency, as shown in Fig. 2(a). The architecture of the resistively loaded current-mode logic (CML) summers is shown in Fig. 2(b). Depending on the resolved previous two symbols; that is, the corresponding six thermometer-coded digital signals (in differential fashion), the six tail currents are respectively steered to one of the two load resistors to perform DFE summation. All these tail currents are summed and mirrored to a common-mode restoration block, which injects currents evenly from the supply

978-1-7281-6031-3/20/$31.00 ©2020 IEEE

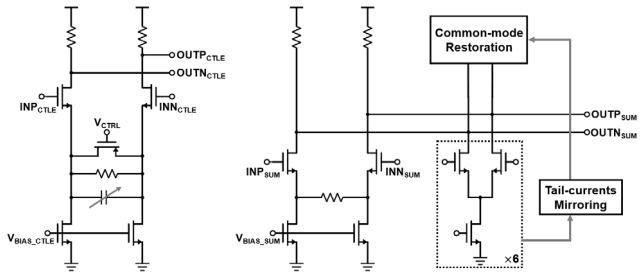

(a) (b)

Fig. 2. (a) The schematic of the CTLE. (b) The architecture of the summer.

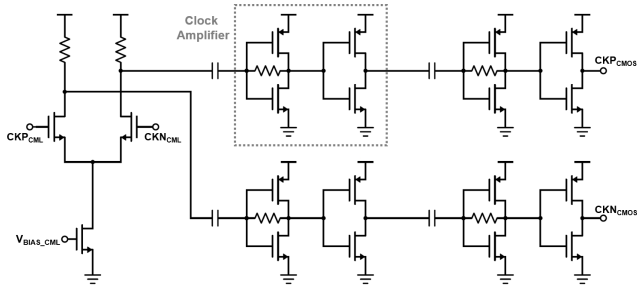

Fig. 3. The schematic of the CML-to-CMOS amplifier.

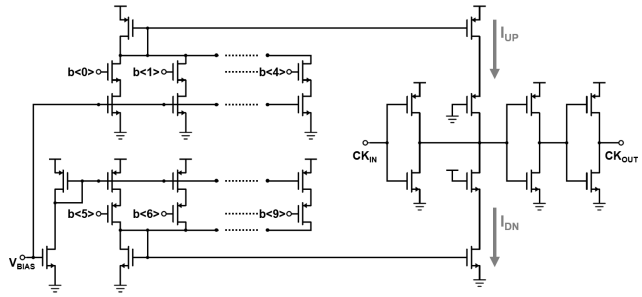

Fig. 4. The schematic of the duty-cycle correction circuitry. Single-ended part is shown for clarity.

to the summing nodes to maintain the common-mode level, irrespective of the setting of DFE; i.e., the amounts of the tail currents.

C. Clock Path

Fig. 3 shows the schematic of the CML-to-CMOS amplifier. The use of AC coupling capacitor and an inverter with its input and output connected via a resistor ensures the DC value of the clock signal is biased to around half of the supply. Simulation results show the outputs of the CML-to-CMOS amplifier are rail-to-rail when the peak-to-peak inputs are larger than 40 mV for 15 GHz clock signals. Fig. 4 presents the duty-cycle correction circuitry. The duty-cycle is adjusted by varying the currents I_{UP} and I_{DN}, which are digitally controllable by 10 bits, b<9:0>. In addition, the value of V_{BIAS} in Fig. 4 is also digitally tunable to accommodate both large duty-cycle distortion (e.g. ±20%) and fine-tuning of the duty-cycle.

III. Slicer Design

Slicers are widely employed in mixed-signal circuits and systems, including analog-to-digital converters (ADCs), adaptive configuration loops, and data receivers. Prevalent topologies of slicers consist of CML [3] and StrongArm [2, 3], as respectively shown in Fig. 5(a) and Fig. 5(b). As the output swing magnitude of the CML slicer depends on the product of the tail current and the load resistance ($I_{TAIL} \times R_L$), the power consumption of the CML slicer is significantly larger than those of StrongArm [3]. By contrast, StrongArm generates rail-to-rail full swings with the cross-coupled CMOS pairs, provided that the regeneration time is sufficient. However, these full-swing signals are grown from approximately zero due to the resetting mechanism of the StrongArm. For high data-rate operations, the time required for the output signals of the StrongArm to grow from approximately zero to the level which can be identified as digital outputs may not be sufficient. A CMOS track-and-regenerate slicer is therefore proposed and designed to improve the clock-to-Q delay and output swing. When the DFE is implemented with the proposed slicer, digital-level outputs are directly available and the settling time specification of the summer is relaxed in consequence of the reduced slicer delay, leading to a DFE design operating at high data rates with improved power efficiency.

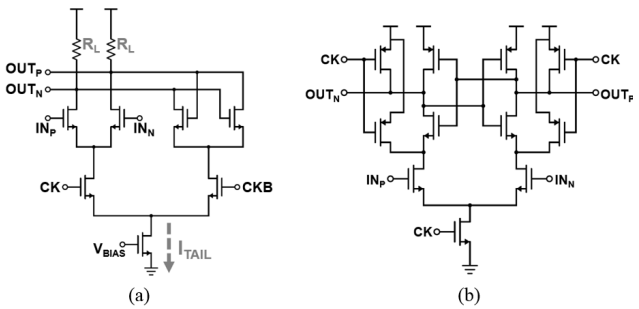

(a) (b)

Fig. 5. Prevalent slicer topologies. (a) CML slicer. (b) StrongArm slicer.

Fig. 6 shows the schematics and operations of the proposed slicer. When CK is logic low and CKB is logic high, M1—M8 and M11—M14 perform the tracking function with M9, M10, M17, and M18 turned off, and they overwrite the latch outputs (i.e. OUTP and OUTN). M15 and M16 are always on and conduct relatively weak currents to avoid the cross-coupled pairs (M19—M22) recover from being completely off. In the other half of clock cycle; that is, when CK is logic high and CKB is logic low, the cross-coupled pair conducts significantly more currents with M17 and M18 turned on, enabling positive feedback to regenerate its differential output. It is the non-resetting feature that allows the regeneration process to start with higher signal levels and consequently reduce the delay. The slicer offset can be compensated by setting TH_P and TH_N correspondingly with on-chip voltage digital-to-analog converters (VDACs) for a given threshold level.

Fig. 7 illustrates the features of the conventional "reset-and-regenerate" StrongArm and the proposed CMOS "track-and-regenerate" slicer, and comparisons between them. The input signals to the slicers are shown in Fig. 7 (a), representing a worst case pattern when a weak negative symbol, i.e. (MSB, LSB) = (-1, +1), is between a long sequence of strong positive symbols, i.e. (MSB, LSB) = (+1, +1). Using the simulated waveforms shown in Fig. 7(b) and Fig. 7(c), a few conclusions can be drawn. First, in contrast to the conventional CML slicer, the

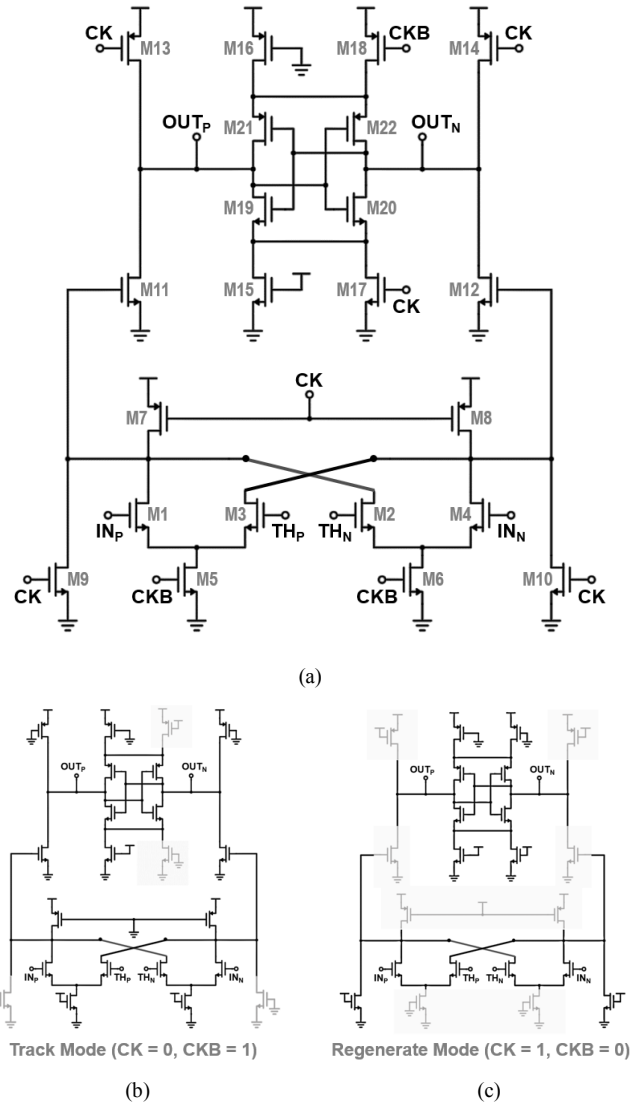

Fig. 6. (a) The circuit schematic of the proposed CMOS track-and-regenerate slicer. (b) The proposed slicer in track mode (CK = 0, CKB = 1). (c) The proposed slicer in regenerate mode (CK = 1, CKB = 0).

proposed slicer offers rail-to-rail output swings and thus direct availability of digital-level outputs. Second, compared to the StrongArm slicer, instead of resetting the latch, the proposed slicer tracks the input signals like the CML slicer does, helping to reduce the required regeneration time. Consequently, the proposed slicer improves the delay and also the output swing over the StrongArm. With the sizes of the input transistors and the output cross-coupled pairs designed to be identical, the worst-case clock-to-Q delay (with switching points defined as ± 450mV) is simulated to be 30.96ps for the StrongArm, whereas it reduces to 15.34ps for the proposed slicer. Moreover, the proposed slicer is less sensitive to the change in power supply. A voltage drop of 50mV from a 900mV supply hinders the StrongArm from resolving the relatively weak negative input to digital-level, while the penalty for the proposed slicer is only 2.36ps of increase in delay.

Fig. 8 presents the superior output swing and input sensitivity of the proposed slicer in comparison to the

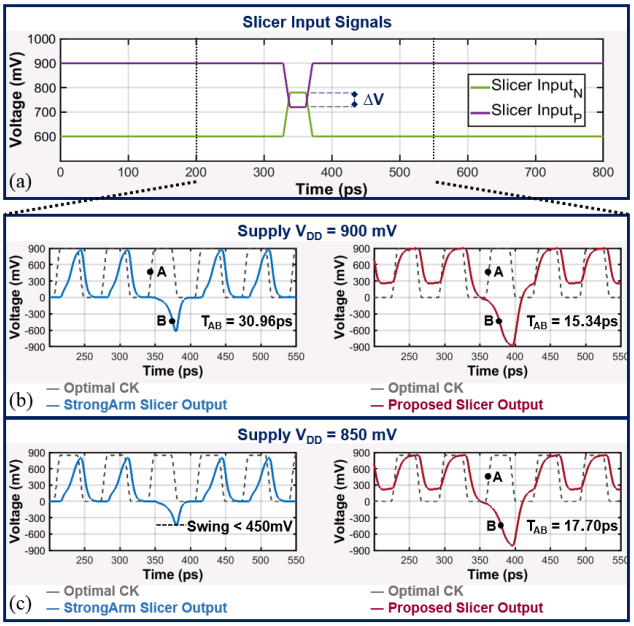

Fig. 7. Simulations and comparisons of the clock-to-Q delay between the reset-and-regenerate StrongArm and the proposed CMOS track-and-regenerate slicer. (a) The input signals to the slicers. (b) The optimal clock signals and the resulting output waveforms of the slicers with 900 mV supply. (c) The optimal clock signals and the resulting output waveforms of the slicers with 850 mV supply.

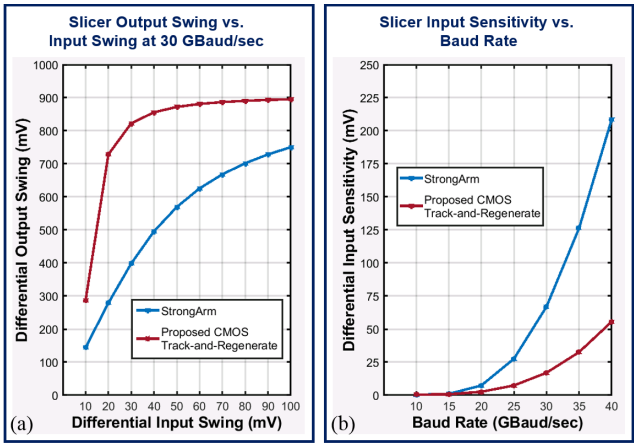

Fig. 8. (a) Simulated slicer output swings with distinct input swings at 30 GBaud/sec (b) Simulated slicer input sensitivity at different baud rates.

StrongArm. Using the input pattern in Fig. 7 (a), but with ΔV swept from 10mV to 100mV instead, Fig. 8 (a) shows that the proposed slicer outperforms the StrongArm, and recovers the input signal to a stronger output. Next, by defining the input sensitivity as the minimum required differential input swing, i.e. ΔV in Fig. 7(a), for the output swing to be larger than 650mV, the sensitivity performance at different baud rates are simulated and summarized in Fig. 8 (b). The input differential pairs and the output cross-coupled pairs in both slicers are designed to be identical for a fair comparison, and they present similar area and loading to the summer circuitries. The proposed track-and-regenerate slicer offers higher gain thanks to its non-resetting feature when the allocated regeneration time becomes stringent.

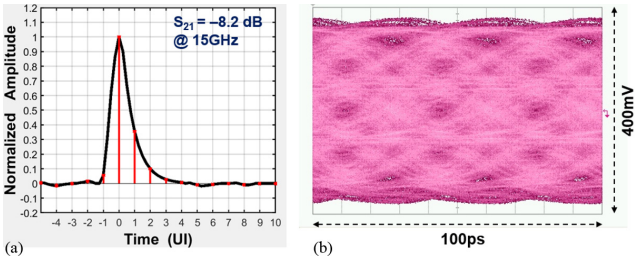

Fig. 9. (a) Measured 30-GBaud/s pulse response at the receiver input. (b) Measured receiver (single-ended) input eyes at 60-Gb/s PAM4.

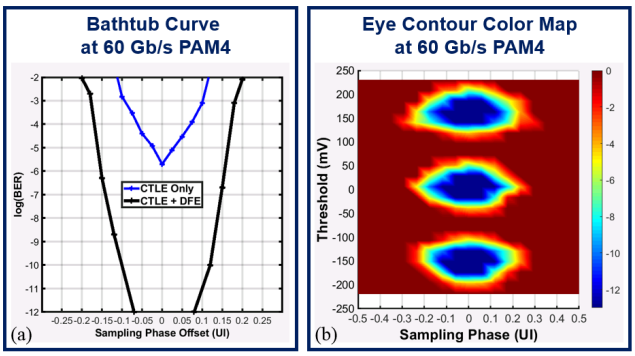

Fig. 10. (a) Measured bathtub curve at 60-Gb/s PAM4. (b) Measured eye-contour at 60-Gb/s PAM4.

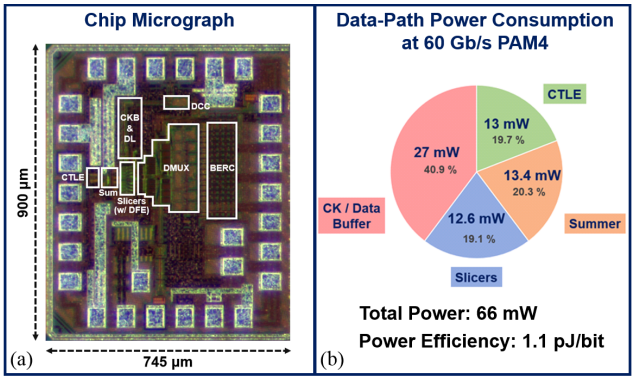

Fig. 11. (a) Chip micrograph with key building blocks highlighted, including CTLE, summer (Sum), slicers with DFE logics, de-multiplexer (DMUX), bit-error-rate counter (BERC), duty-cycle correction (DCC), clock buffers (CKB), and delay line (DL). (b) Measured receiver data-path power consumption at 60-Gb/s.

IV. Experimental Results

The PAM4 receiver chip was fabricated in 28-nm CMOS, and was fully characterized via measurements. It achieves BER better than 1E-12 at 60-Gb/s when tested with PRBS-7, 9, 31 patterns over the channel characterized by its pulse response as shown in Fig. 9(a). The resulting receiver input eyes are completely closed as can be seen in Fig. 9(b). The channel loss at 15GHz is measured to be 8.2dB, composed of the cable loss and PCB-trace loss. The bathtub curves with DFE switched on/off are measured and plotted as Fig. 10(a), showing the effectiveness of the PAM4-DFE at 60-Gb/s. The horizontal opening for BER = 1E-12 is 0.15UI. Fig. 10(b) is the measured PAM4 eye contour after equalization. The chip micrograph with key building blocks highlighted is presented in Fig. 11(a), and the receiver data-path power consumption along with its breakdown at 60-Gb/s is shown in Fig. 11(b).

TABLE I.

	This Work	JSSC'17 [2]	JSSC'17 [4]	ISSCC'18 [5]	VLSI'19 [6]
Technology	28-nm CMOS	16-nm FinFET	65-nm CMOS	16-nm FinFET	40-nm CMOS
Data Rate	60 Gb/s	40-56 Gb/s	60 Gb/s	19-56 Gb/s	52 Gb/s
Modulation	PAM4	PAM4	NRZ	PAM4	PAM4
Equalization	CTLE 2-Tap DFE	CTLE 10-Tap DFE	CTLE 2-Tap FFE 3-Tap DFE	CTLE	CTLE FFE
Power	66 mW	230 mW @ 56G	136 mW	360 mW @ 56G	48mW
Power Efficiency	1.1 pJ/b	4.11 pJ/b	2.26 pJ/b	6.4 pJ/b	0.92 pJ/b
Channel	8.2 dB @ 15G	10 dB @ 14G	21 dB @ 30G	7.4 dB @ 14G	7.3 dB @ 13G

V. Conclusion

A CMOS track-and-regenerate slicer is proposed and designed to improve the clock-to-Q delay and output swing over the conventional CML slicer and the StrongArm. A PAM4 receiver employing the proposed slicer demonstrates an energy-efficient direct PAM4-DFE design, which benefits from the relaxed settling time constraint due to the reduced slicer delay, and also the immediate availability of full-swing signals at the slicer outputs. The prototype fabricated in 28-nm CMOS achieves power efficiency of 1.1 pJ/b at 60-Gb/s over a channel of 8.2dB loss at Nyquist. Table I. summarizes the receiver performance and compares it with prior arts.

Acknowledgements

The authors would like to thank D. A. Nelson, Caltech MICS lab members and Caltech CHIC lab.

References

[1] J. Lee, P. Chiang, P. Peng, L. Chen and C. Weng, "Design of 56 Gb/s NRZ and PAM4 SerDes Transceivers in CMOS Technologies," in IEEE Journal of Solid-State Circuits, vol. 50, no. 9, pp. 2061-2073, Sept. 2015.

[2] J. Im et al., "A 40-to-56 Gb/s PAM-4 Receiver With Ten-Tap Direct Decision-Feedback Equalization in 16-nm FinFET," in IEEE Journal of Solid-State Circuits, vol. 52, no. 12, pp. 3486-3502, Dec. 2017.

[3] T. Toifl et al., "A 22-gb/s PAM-4 receiver in 90-nm CMOS SOI technology," in IEEE Journal of Solid-State Circuits, vol. 41, no. 4, pp. 954-965, April 2006.

[4] J. Han, N. Sutardja, Y. Lu and E. Alon, "Design Techniques for a 60-Gb/s 288-mW NRZ Transceiver With Adaptive Equalization and Baud-Rate Clock and Data Recovery in 65-nm CMOS Technology," in IEEE Journal of Solid-State Circuits, vol. 52, no. 12, pp. 3474-3485, Dec. 2017.

[5] P. Upadhyaya et al., "A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with a configurable ADC in 16nm FinFET," 2018 IEEE International Solid - State Circuits Conference - (ISSCC), San Francisco, CA, 2018, pp. 108-110.

[6] C. Wang, G. Zhu, Z. Zhang and C. P. Yue, "A 52-Gb/s Sub-1pJ/bit PAM4 Receiver in 40-nm CMOS for Low-Power Interconnects," 2019 Symposium on VLSI Circuits, Kyoto, Japan, 2019, pp. C274-C275.