This article has been accepted for publication in IEEE Solid-State Circuits Letters. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/LSSC.2026.3721808

A 64-Gb/s/pin 0.25-pJ/b Single-Ended Short-Reach PAM4 Transceiver in 28-nm CMOS

Kshitiz Tyagi, *Member, IEEE,* Kunbo Ma, Joung Won Park, *Member, IEEE,* and Behzad Razavi, *Fellow, IEEE*

***Abstract*—We present a transceiver architecture for single-ended PAM4 signaling, introducing an inverter-based active-inductor terminated front-end that obviates the need for further equalizer stages, substantially reducing power and area. Realized in 28-nm CMOS technology, the 4-lane link operates at 64 Gb/s/pin with a bit error rate less than 10^{-12} while consuming only 0.25 pJ/b (0.2 pJ/b at 50 Gb/s) from a 1-V supply, and achieves bandwidth density greater than 28 Tb/s/mm^2.**

***Index Terms*—chiplets, die-to-die links, short-reach transceivers, PAM4, wireline communications.**

I. Introduction

THE push toward higher throughput rates in chiplets has motivated extensive work on signaling schemes over a single wire [1]- [4]. In principle, the rate per wire can be doubled through the use of bidirectional links [1]- [4], but the "hybrid" circuit necessary for suppressing the transmitted signal at the receiver (RX) input typically incurs several dB of loss [2]. Moreover, the hybrid's incomplete cancellation of the TX signal leads to an SNR penalty for the RX. In this brief, we instead consider employing unidirectional PAM4 signaling, and propose a compact, scalable, and low-power link architecture that operates robustly with a bit error rate (BER) $< 10^{-12}$ at 64-Gb/s.

This article is organized as follows: Section II describes the challenges associated with single-ended PAM4 signaling over interposers, Section III introduces our proposed transceiver architecture, Section IV presents the measurement results, and Section V concludes the paper.

II. Challenges with PAM4 Signaling

The principal challenge facing PAM4 in a single-ended channel arises from the driver nonlinearity. Consider, for example, the scenario depicted in Fig. 1(a), where the TX inverter-based DAC provides proper back termination and the RX incorporates a linear termination resistor, R_T. As illustrated by the simulated eye diagram in Fig. 1(b), the PAM4 data is severely distorted. The top and bottom eyes are compressed because the transistors generating the highest and lowest levels enter the triode region. Fig. 1(c) depicts this effect for $V_{RX} \approx$ 0.7 V. The arrangement in Fig. 1(a) merits two more points. First, the dc resistance of the channel degrades the voltage swing by about 140 mV. Second, V_b should be in the vicinity of $V_{DD}/2$ for optimum performance but it must also sink or source about 5 mA (for an RX signal swing of 0.5 V_{pp} and 50-ohm termination in RX), requiring its own power/area-hungry generation circuit.

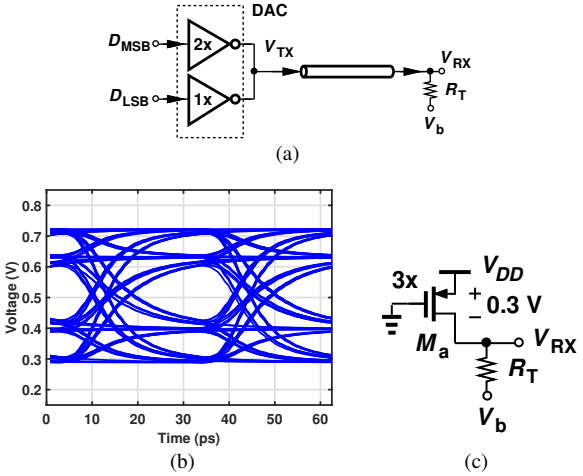

Fig. 1: (a) Channel terminated with linear resistor at the receiver, (b) simulated eye diagram, and (c) state of driver transistors when output is at maximum level.

III. Proposed Transceiver

With these observations in mind, we propose the conceptual front end depicted in Fig. 2(a), where the TX drives the channel with a current source to avoid its dc resistance, and the RX post-distorts the PAM4 signal. The proposed RX front end must satisfy two conditions: (a) it must robustly correct for the TX nonlinearity across PVT corners so as to maintain equal PAM4 eye heights, and (b) its input resistance, R_{in}, must create matching without the power penalty mentioned for V_b in Fig. 1(a). Both are achieved as illustrated in Fig. 2(b). Here, Inv_3, with its input and output shorted, operates as a diode-connected device and hence post-distorts the PAM4 signal generated by Inv_1 and Inv_2. This correction can be explained intuitively as follows. The incremental impedance of Inv_3, R_{in}, in fact increases as V_A departs from $V_{DD}/2$, compensating for the reduced transconductance of M_a in Fig. 1(c). The simulated R_{in} is shown in Fig. 2(c). Inv_3 terminates the line while drawing only a supply current of 1.4 mA[1].

For a symbol rate of 32 Gbaud, the bump, ESD, and RX input capacitance, denoted by C_1 in Fig. 2(b), degrades the input return loss and the bandwidth at node A. We improve

Kshitiz Tyagi and Kunbo Ma were with the Department of Electrical and Computer Engineering, University of California, Los Angeles, CA 90095, USA. Kshitiz Tyagi is now with Mediatek Inc., Irvine, CA 92618, USA (email: ktyagi30@ucla.edu).

Joung Won Park is with Samsung US Inc., San Diego, CA 92130, USA (e-mail: joung2.park@samsung.com).

Behzad Razavi is with the Department of Electrical and Computer Engineering, University of California, Los Angeles, CA 90095, USA (email: razavi@ee.ucla.edu)

[1]A generation circuit for V_b is avoided as Inv_3 is self-biased

Authorized licensed use limited to: UCLA Library. Downloaded on August 15,2026 at 17:27:50 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Solid-State Circuits Letters. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/LSSC.2026.3721808

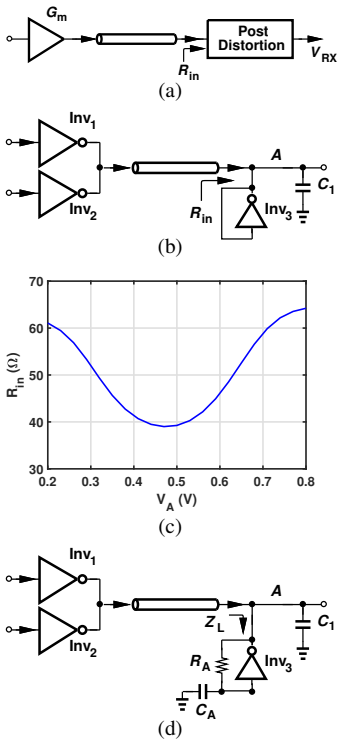

Fig. 2: (a) Conceptual view of proposed transceiver, (b) implementation using diode-connected inverter at RX, (c) input resistance of diode-connected device as a function of the RX voltage, and (d) performance enhancement using active-inductance.

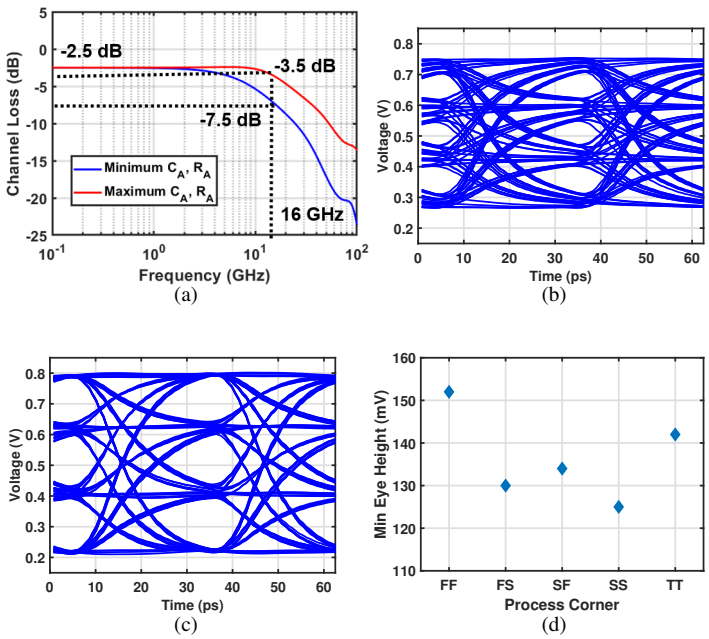

Fig. 3: (a) Channel loss with zero and maximum active inductance, (b) simulated eye diagram with zero inductance, (c) simulated eye diagram with maximum inductance, and (d) min. eye height across process corners.

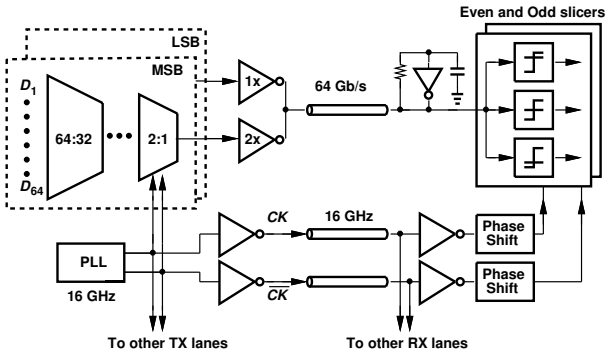

Fig. 4: Overall architecture of the proposed transceiver.

both by transforming Inv_3 to an active inductor [5], as shown in Fig. 2(d). We have $Z_L \approx (R_A C_A s + 1)/(C_A s + G_{m3})$, where G_{m3} is the transconductance of Inv_3. As shown in [6], this leads to an equivalent inductance of about $R_A C_A/G_{m3}$. The dimensions and bias current of Inv_3 determine the low-frequency match, and C_A and R_A set the effective inductance to optimize the high-frequency performance in view of C_1. For C_1 = 80 fF, we plot the effective channel loss in Fig. 3(a) with C_A and R_A set to their minimum and maximum values, noting an improvement of 4 dB at 16 GHz (the dc loss is 2.5 dB due to the channel resistance). We also plot the received eye in Figs. 3(b) and (c) for the two cases, observing the significant role played by the active inductor in improving the performance. No further equalization is therefore necessary. We also note that, compared to the conventional current-mode circuits, e.g., continuous-time linear equalizers, the inverter-based design is more compact, draws less power, provides higher linearity, and scales better with process technology [5].

The robustness of the proposed front end is evaluated by monitoring the worst-case eye height across process corners. Shown in Fig. 3(d), the results suggest a variation from 125 mV to 152 mV. Including supply and temperature variation, the minimum eye height is found to be 110 mV, which provides ample margin for BER $< 10^{-12}$. The overall link architecture is depicted in Fig. 4. The TX serializer is driven by a half-rate type-I PLL. The 16-GHz differential clock is also forwarded to the RX. Two sets of slicers detect the received PAM4 signal on the rising and falling edges of the clock. The prototype consists of four data lanes and one clock lane.

Another key concept employed in this work is to minimize the complexity and hence area of transceiver circuits, leading to two critical advantages. First, the bandwidth density (in Tb/s/mm^2) is boosted by a factor of 5.3 compared to the prior art [4]. Second, more importantly, the compact floor plan and short interconnects reduce the clock network power consumption substantially. This power amounts to 0.08 pJ/b in [2] but only 0.03 pJ/b in our work.

Our complexity and power reduction strategy consists of four principles. First, the two 64-to-1 serializers in the TX incorporate quadrature clocks to obviate the need for any latches [7]. This approach saves power in both the data and clock paths. Second, the RX employs simple programmable delay lines, rather than phase interpolators, to adjust the

Authorized licensed use limited to: UCLA Library. Downloaded on August 15,2026 at 17:27:50 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Solid-State Circuits Letters. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/LSSC.2026.3721808

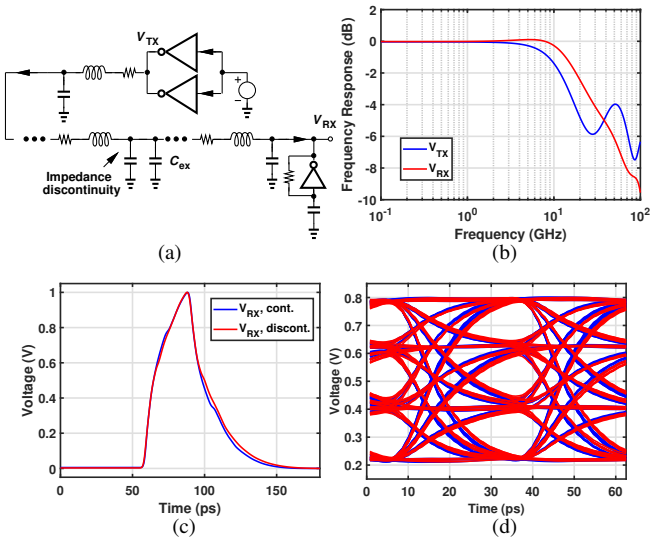

Fig. 5: (a) Channel model including impedance discontinuity, (b) simulated normalized frequency response at TX output and RX input,(c) simulated normalized pulse response at RX input, and (d) simulated RX eye diagram with (red) and without (blue) inclusion of the discontinuity, revealing less than 4% degradation in the heights.

phase of the forwarded clock[2]. Third, the exclusive use of C^{2}MOS logic throughout the transceiver greatly reduces the circuits' footprints. Fourth, the MSB and LSB line drivers rely on the inverters' on-resistance to provide proper back termination, avoiding resistors in series with the transistors. This PVT-dependent resistance proves adequate for suppressing secondary reflections because the highly resistive channel "smoothes" out the frequency response in the presence of impedance discontinuities. As shown in Fig. 5, a discontinuity that increases the first postcursor by 20%, and causes a dip, affects the overall response and eye diagram negligibly. Figure 6 plots the simulated TX and RX return loss, displaying moderate matching up to the Nyquist frequency. We make two remarks. First, the active inductor, in fact, improves the input match in the presence of the RX input capacitance. According to simulations, $|S_{11}|$ would only be about -2 dB at 16 GHz without the active inductor. Second, an $|S_{11}|$ of -7.5 dB is tolerable here because the high resistance of the channel ($\approx 38\Omega$) tends to suppress secondary reflections.

The RX shown in Fig. 7 merits two more remarks. First, the slicers are based on the double-tail topology [8], with the middle and top eye slicers using NMOS input pairs and the bottom one, a PMOS pair. They exhibit a sensitivity of 16 mV at 16 GHz and an input-referred noise of 1 mV_{rms}. For a BER $< 10^{-12}$ and a minimum eye height of 110 mV, we obtain a margin of 70 mV for supply noise and crosstalk. Second, for ease of testing, we recombine the slicer outputs by a current-mode DAC and deliver the resulting PAM4 signal to off-chip instrumentation.

[2]the delay is adjustable over a 1-UI range, and the optimal code is selected at startup through foreground calibration

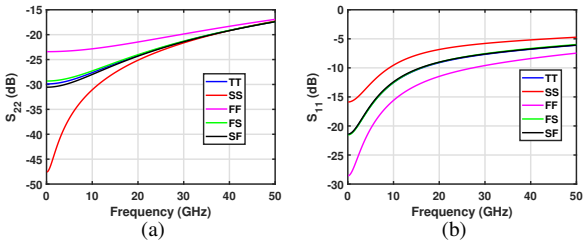

Fig. 6: Simulated return loss vs frequency across process corners at (a) the TX, and (b) the RX.

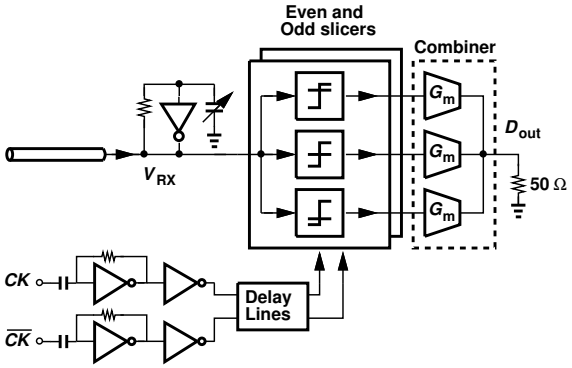

Fig. 7: Receiver block diagram including output combining stage.

Due to chip area limitations, we emulate the interposer channel by means of a distributed on-chip RC network[3] [Fig. 8(a)]. Fig. 8(b) shows the loss of this network and compares it with the loss of a reference interposer channel, showing close agreement up to twice the Nyquist frequency. This structure is used in both the data and clock paths.

[3]this representation is effective since typical interposer channels are RC-dominant [2], [9]

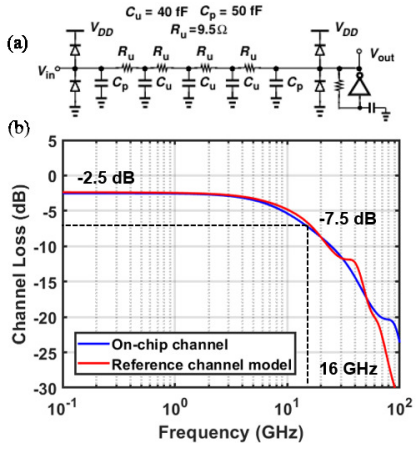

Fig. 8: (a) On-chip channel implementation using passives, with the RX termination included, and (b) simulated loss characteristic compared with reference interposer link

.

Authorized licensed use limited to: UCLA Library. Downloaded on August 15,2026 at 17:27:50 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Solid-State Circuits Letters. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/LSSC.2026.3721808

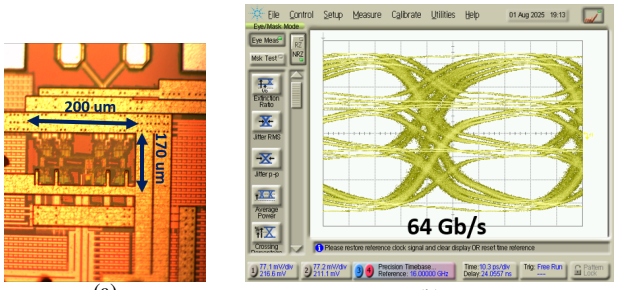

(a) (b)

Fig. 9: (a) Die micrograph, and (b) measured output data eye diagram, at half the data rate due to 1:2 deserialization.

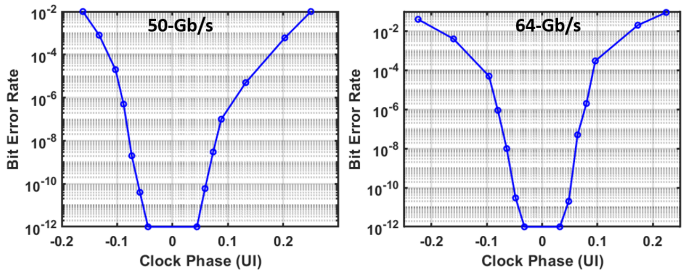

Fig. 10: Measured bathtub curves at 50- and 64-Gb/s.

IV. Measurement Results

The prototype has been fabricated in TSMC's 28-nm technology [Fig. 9(a)]. The total active area is 0.009 mm^2(each TX occupies 24 μm $\times$ 35 μm, RX occupies 30 μm $\times$ 35 μm, and the PLL occupies 40 μm $\times$ 40 μm), yielding a bandwidth density of 28 Tb/s/mm^2. After amortizing the PLL power (2.5 mW per lane), we obtain a per-lane value of 16 mW at 64 Gb/s, and hence an efficiency of 0.25 pJ/b (The TX and RX consume 6 mW and 7.5 mW, respectively).

Fig. 9(b) shows the 64-Gb/s PAM4 eye at RX reconstruction DAC output[4]. Using Keysight's M8040A, we obtain the bathtub curves depicted in Fig. 10, which suggest a horizontal opening of 0.07 UI at 64 Gb/s and 0.09 UI at 50 Gb/s (the clock phase is swept using the on-chip delay lines in steps of roughly 0.6 ps). The phase is first varied toward the left and right extremes, where the BER rises sharply, and then swept in fine steps in between. Figure 11 shows the measured PLL phase noise at the 1-GHz output, yielding an rms jitter of 260 fs. Table I summarizes the results, indicating efficiency and bandwidth density improvements of 4.8X and 5.3X, respectively, with respect to [3], [4].

V. Conclusion

We present a single-ended 64-Gb/s PAM4 transceiver architecture that is suitable for die-to-die interconnects. A front-end design employing active-inductor peaking and termination obviates the need for equalizer stages in the receiver, saving considerable power and area and mitigating signal integrity degradation due to noise and nonlinearity. Measurement results of the 4-lane prototype, fabricated in 28-nm CMOS technology, reveal improvements of 4.8X in the energy efficiency and 5.3X in the bandwidth density over the prior art.

[4]the non-linearity of this eye is acceptable as long as the minimum height meets the measuring instrument sensitivity

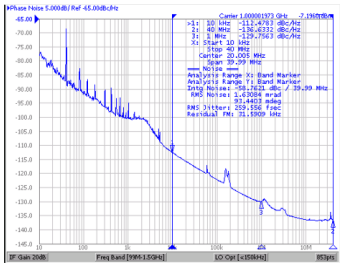

Fig. 11: Measured PLL phase noise at 1 GHz.

TABLE I: Performance summary and comparison

Reference	ISSCC'26 [1]	JSSC'23 [2]	ISSCC'25 [3]	ISSCC'26 [4]	This Work
Data Rate (Gb/s)	56	50	64	64	50, 64
Modulation	Bidirectional	Bidirectional	Bidirectional	Bidirectional	PAM4
BER	$<10^{-12}$	$<10^{-12}$	$<10^{-16}$	$<10^{-12}$	$<10^{-12}$
Channel Loss @ Nyquist (dB)	6.5	5.4	2.8	6	5
Bandw. Density (Tb/s/mm^2)	16.5	4.45	4.88	5.33	28.05
Power Efficiency (pJ/bit)	0.292	0.297	1.21	1.5	0.2 (50 Gb/s) 0.25 (64 Gb/s)
Technology	28–nm CMOS	5–nm FinFET	28–nm CMOS	28–nm CMOS	28–nm CMOS

Acknowledgment

We thank Samsung Semiconductor US Inc. for funding this work, and TSMC for chip fabrication.

References

[1] K. Kim, Y. Lee, D. Na, H. -J. Park, J. Song and W. -S. Choi, "A 0.292pJ/b 56Gb/s/wire Capacitively Driven Simultaneous Bidirectional Transceiver with PVT/Mismatch Tracking for XSR and D2D Interfaces in 28nm CMOS," IEEE ISSCC, San Francisco, CA, USA, 2026, pp. 148-150.

[2] Y. Nishi et al., "A 0.297-pJ/Bit 50.4-Gb/s/Wire Inverter-Based Short-Reach Simultaneous Bi-Directional Transceiver for Die-to-Die Interface in 5-nm CMOS," in IEEE Journal of Solid-State Circuits, vol. 58, no. 4, pp. 1062-1073, April 2023.

[3] Z. Wang et al., "A 64Gb/s/wire 10.5Tb/s/mm/Layer Single-Ended Simultaneous Bi-Directional Transceiver with Echo and Crosstalk Cancellation for a Die-to-Die Interface in 28nm CMOS," IEEE ISSCC, San Francisco, CA, USA, 2025, pp. 588-590.

[4] X. Cheng et al., "A 72Gb/s/pin Single-Ended Simultaneous Bi-Directional Transceiver with C-Peaking Leakage Cancellation and Dual-Loop Hybrid Impedance Calibration for Chiplet Interfaces," IEEE ISSCC, San Francisco, CA, USA, 2026, pp. 150-152.

[5] J. Im et al., "A 112-Gb/s PAM-4 Long-Reach Wireline Transceiver Using a 36-Way Time-Interleaved SAR ADC and Inverter-Based RX Analog Front-End in 7-nm FinFET," in IEEE Journal of Solid-State Circuits, vol. 56, no. 1, pp. 7-18, Jan. 2021.

[6] B. Razavi, "The Active Inductor [A Circuit for All Seasons]," in IEEE Solid-State Circuits Magazine, vol. 12, no. 2, pp. 7-11, Spring 2020.

[7] Y. Chang, A. Manian, L. Kong and B. Razavi, "An 80-Gb/s 44-mW Wireline PAM4 Transmitter," in IEEE Journal of Solid-State Circuits, vol. 53, no. 8, pp. 2214-2226, Aug. 2018.

[8] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl and B. Nauta, "A Double-Tail Latch-Type Voltage Sense Amplifier with 18ps Setup+Hold Time," IEEE ISSCC, San Francisco, CA, USA, 2007, pp. 314-605.

[9] D. Das Sharma, G. Pasdast, Z. Qian and K. Aygun, "Universal Chiplet Interconnect Express (UCIe): An Open Industry Standard for Innovations With Chiplets at Package Level," in IEEE Transactions on Components, Packaging and Manufacturing Technology, vol. 12, no. 9, pp. 1423-1431, Sept. 2022.

Authorized licensed use limited to: UCLA Library. Downloaded on August 15,2026 at 17:27:50 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.