

A 56-Gb/s 50-mW NRZ Receiver in 28-nm CMOS

Atharav Atharav^{1b}, Member, IEEE, and Behzad Razavi^{1b}, Fellow, IEEE

Abstract—A wireline receiver consisting of a linear equalizer, a decision-feedback equalizer (DFE), a clock and data recovery (CDR) circuit, and a demultiplexer (DMUX) employs new circuit and architecture techniques that afford substantial power savings. Realized in 28-nm technology, the 56-Gb/s receiver has a bit error rate (BER) of less than 10^{-12} for a channel loss of 25 dB at 28 GHz.

Index Terms—Clock and data recovery (CDR), continuous time linear equalizer (CTLE), dual-loop decision-feedback equalizer (DFE), feedforward system, non-return-to-zero (NRZ) data, wireline receiver.

I. INTRODUCTION

THE power consumption of wireline transceivers has become increasingly more critical as higher data rates and a larger number of lanes per chip are sought. This issue is further intensified by the tradeoffs between the channel loss and the power dissipation, especially in the receive path. While PAM4 signaling is attractive for lossier channels, it has mostly dictated receiver designs incorporating analog-to-digital converters (ADCs) [1]–[4] with high power numbers. For example, the PAM4 receivers in [7], [11], and [13] draw 382, 180, and 259 mW for channel losses of 24, 16.8, and 20.8 dB, respectively. Non-return-to-zero (NRZ) receivers, on the other hand, can be realized in the analog domain, potentially consuming less power, but they must deal with a greater channel loss.

This paper introduces a 56-Gb/s NRZ receiver that draws 50 mW while exhibiting a bit error rate (BER) of less than 10^{-12} for a channel loss of 25 dB at 28 GHz and 13.5 dB at 14 GHz [5]. Such a receiver can compete with PAM4 counterparts and/or serve as part of 112-Gb/s systems that must also support 56-Gb/s NRZ reception [3], [4], [12].

Section II describes a number of design issues, and Section III presents the receiver's high-level architecture. Sections IV–VI deal with the design of the continuous time linear equalizer (CTLE), the decision-feedback equalizer (DFE), and the clock and data recovery (CDR), respectively. Section VII presents the overall receiver implementation and Section VIII summarizes the experimental results.

Manuscript received April 13, 2021; revised June 27, 2021; accepted August 24, 2021. Date of publication September 22, 2021; date of current version December 29, 2021. This article was approved by Guest Editor Amir Amirkhany. This work was supported in part by Texas Instruments Inc., in part by Realtek Semiconductor Corporation, and in part by the Taiwan Semiconductor Manufacturing Company (TSMC) University Shuttle Program. (Corresponding author: Atharav Atharav.)

Atharav Atharav was with the Department of Electrical and Computer Engineering, University of California at Los Angeles, Los Angeles, CA 90095 USA. He is now with Mediatek USA Inc., Irvine, CA 92618 USA (e-mail: atharav@ucla.edu).

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

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/JSSC.2021.3109032>.

Digital Object Identifier 10.1109/JSSC.2021.3109032

0018-9200 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

II. GENERAL CONSIDERATIONS

The twofold bandwidth efficiency of PAM4 with respect to NRZ accrues at certain costs. We briefly review the tradeoffs here.

In addition to drawing high power, ADC-based PAM4 receivers also present two other difficulties. First, they impose a stricter lower bound on the clock jitter than analog CDR circuits do. For example, a 56-GHz 7-bit ADC incurs a 3-dB signal-to-noise ratio penalty with a clock jitter of 72 fs_{rms} [8]. By comparison, a CDR jitter of 600 fs_{rms} is typically acceptable at these rates [9]. Second, the latency associated with the ADC and any digital processing that appears in the clock recovery loop can introduce jitter peaking, demanding a significant reduction in the loop bandwidth.

Analog NRZ receivers typically incorporate a CTLE, a DFE, a CDR circuit, and a demultiplexer (DMUX). A high channel loss has two critical implications for the overall design: 1) it requires a multitude of stages in the CTLE with gain boosting up to a frequency of 28 GHz. Such an approach tends to consume high power and, more importantly, faces serious bandwidth limitations as the number of stages increases; 2) the high loss translates to a large amount of jitter, even at the CTLE output, thus decreasing the phase detector (PD) gain in the CDR circuit and hence reducing its loop bandwidth, capture range, and jitter tolerance. This situation is depicted in Fig. 1(a)¹ and calls for a high boost factor in the CTLE. Otherwise, the CDR may fail to lock.

This issue is alleviated if the CDR instead senses the DFE summing junction [see Fig. 1(b)], but the CDR input capacitance substantially loads this node and degrades the DFE performance. Note that a unit interval (UI) of approximately 18 ps at 56 Gb/s poses an extremely tight loop delay for the DFE, especially in view of the clock's finite rise and fall times (≈ 5 ps).

Returning to Fig. 1(a), we remark that the CTLE must drive the input capacitances of both the DFE and the CDR, a formidable challenge in view of the bandwidth requirements at this node. We also note that inserting a buffer after the CTLE does not resolve the issue because the buffer itself further limits the bandwidth.

The CDR design also poses its own difficulties. Clock generation and distribution as well as PD design generally favor a half-rate architecture at these speeds, but we also prefer to avoid CDR architectures demanding quadrature clock phases.

The foregoing observations suggest that new CTLE, DFE, and CDR techniques are necessary for a substantial reduction

¹We can choose the DFE in Fig. 1 from the topologies described in Section V. The CDR in Fig. 1 can be implemented as either traditional 2 \times oversampling or a baud-rate architecture. Each of these CDR implementations introduces substantial loading on the data path.

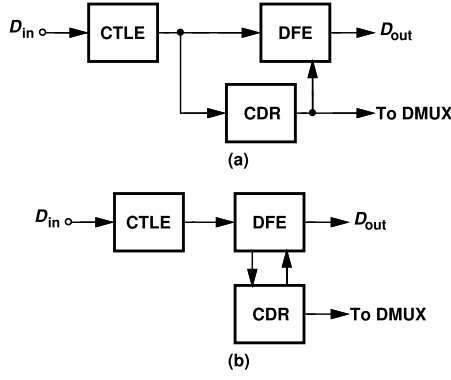


Fig. 1. RX design choices: (a) CDR fed by the CTLE, and (b) CDR fed by the DFE summing node.

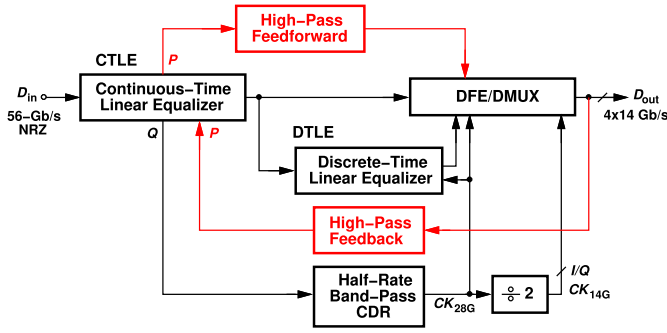


Fig. 2. Conceptual RX architecture.

of the receiver power consumption. We present such techniques in Sections III–VII.

III. NRZ RECEIVER ARCHITECTURE

The proposed receiver introduces a number of circuit and architecture techniques to ease the tradeoffs among channel loss, speed, and power consumption. Fig. 2 shows a functional diagram of the receiver so as to highlight some of the aspects on which we will focus.

The receiver data path consists of a CTLE core, a DFE core, a discrete-time linear equalizer (DTLE) [15], and a DMUX. The receiver’s performance is dramatically improved by a number of additional feedforward and feedback paths. Moreover, the half-rate CDR circuit avoids loading the main data path² and also obviates the need for a quadrature oscillator.

The receiver targets a maximum channel loss of 25 dB at 28 GHz. Accordingly, it allocates 16 dB of linear equalization to the CTLE and the DTLE. The CDR, on the other hand, is preceded with only 7.3 dB of boost. The DTLE and DFE incorporate charge-steering techniques [16] to achieve low power consumption.

The limited speed of 28-nm CMOS technology requires the use of inductive peaking at most nodes, in both the data path and the clock path, leading to considerable difficulties in the routing and distribution of signals. Thus, the design, layout,

²In this work, the CDR senses node Q , which is not in the main signal path (see Section VI).

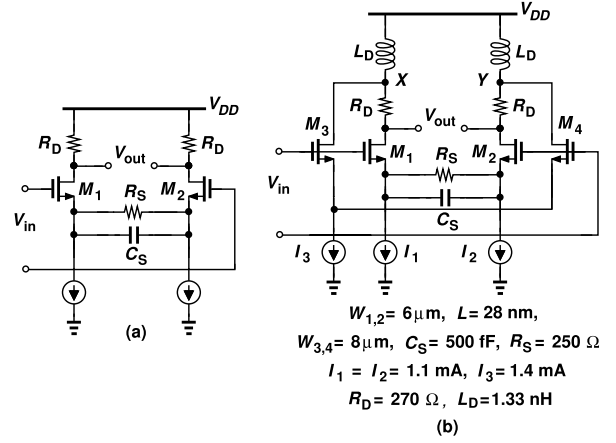


Fig. 3. (a) Basic CTLE stage, and (b) CTLE with feedforward.

and extraction steps are accompanied by electromagnetic field simulations so as to predict the bandwidths and the resonances accurately.

IV. CTLE DESIGN

Our past work indicates that the CTLE tends to be the most power-hungry block in the data path [15], [17], both because of the boost and bandwidth requirements and because of the heavy load presented to it by the (half-rate) DFE. For this reason, we wish to include no more than two stages in its design. In order to achieve sufficient boost and bandwidth, we propose several new feedforward techniques.

Consider the conventional CTLE stage shown in Fig. 3(a). The low-frequency gain and the boost factor of the circuit are approximately given by $g_{m1,2}R_D/(1 + g_{m1,2}R_S/2)$ and $1 + g_{m1,2}R_S/2$, respectively, where $g_{m1,2}$ denotes the transconductance of M_1 and M_2 . The direct tradeoff between these two design parameters severely limits the maximum achievable boost factor, especially if $g_{m1,2}R_D$ is also constrained by the supply voltage. These difficulties naturally lead to the use of multiple CTLE stages and hence bandwidth reduction.

Articulated more fundamentally, the CTLE’s role is to provide a high-pass response, which can be potentially obtained by other means. Since the CTLE stages designed for data rates of interest to us must incorporate inductors anyway, we surmise that these components can also realize a high-pass action. To this end, let us turn to the topology shown in Fig. 3(b), where undegenerated transistors M_3 and M_4 provide a feedforward path around the basic CTLE stage, but they feed the signal to the load inductors rather than to the main output. We expect the high-pass filter thus created raises the boost factor without compromising the low-frequency gain or the bandwidth.

The performance improvement resulting from feedforward can be quantified with the aid of a half circuit. Ignoring the capacitance at the output nodes for now, we have

$$\frac{V_{out}}{V_{in}} = -\frac{g_{m1,2}(R_D + L_D s)}{1 + g_{m1,2}\left(\frac{R_S}{2} \parallel \frac{1}{2C_{SS}}\right)} - g_{m3,4}L_D s \quad (1)$$




Fig. 4. Simplified CTLE model.




Fig. 5. CTLE feedforward frequency response.

where $g_{m3,4}$ denotes the transconductance of the feedforward transistors. The second term on the right-hand side represents the zero created by feedforward action and can be adjusted by selecting $g_{m3,4}$ properly.

At high frequencies, the denominator of the first term on the right-hand side approaches unity, implying that $g_{m3,4}L_Ds$ simply adds to $g_{m1,2}L_Ds$, an effect that, evidently, could be obtained by simply making L_D larger and omitting the feedforward path. The key point, however, is that the value of L_D is dictated by the capacitance at the output node, C_{out} , and cannot be arbitrarily raised in the conventional topology. Feedforward, on the other hand, affords greater flexibility in the design. With C_{out} included, we obtain

$$\frac{V_{out}}{V_{in}} = - \left[\frac{g_{m1,2}(R_D + L_Ds)}{1 + g_{m1,2} \left(\frac{R_S}{2} \parallel \frac{1}{2C_Ss} \right)} + g_{m3,4}L_Ds \right] \times \frac{1}{L_D C_{out}s^2 + R_D C_{out}s + 1}. \quad (2)$$

Notably, the relative contribution of the feedforward path remains unchanged.

To gain additional insight and deal with a more general case, let us also include the parasitic capacitance at the drains of M_3 and M_4 in Fig. 3(b) and draw the half-circuit small-signal model as shown in Fig. 4. Here, I_1 and I_3 model the drain currents of M_1 and M_3 , respectively. The transfer functions V_{out}/I_1 and V_{out}/I_3 have the same poles but not necessarily the same zeros. Of interest to us is the zero associated with V_{out}/I_3 as it provides the additional boost. To obtain this zero, we assume $I_1 = 0$ and find the value of the complex frequency, s , that yields $V_{out} = 0$. Under this condition, C_{out} and R_D carry no current, resulting in $V_X = 0$. Thus, $I_3(L_Ds \parallel (C_p s)^{-1}) = 0$ and hence $I_3 L_D s / (L_D C_p s^2 + 1) = 0$. The zero remains at the origin and, for frequencies below the circuit's poles, creates a high-pass response.

Fig. 5 sketches the frequency response of the main and feedforward paths, revealing that the latter can be so designed as to become dominant as the former's response reaches a plateau at ω_{p1} . This pole frequency is equal to $(1 + g_{m1,2}R_S/2)/(R_S C_S)$




Fig. 6. Proposed CTLE architecture.

in Fig. 3, and we wish the feedforward path's response, $g_{m3,4}L_D\omega$, to take over only for $\omega > \omega_{p1}$. That is, $g_{m3,4}L_D\omega_{p1} < g_{m1,2}R_D$, and hence

$$g_{m3,4} < \frac{g_{m1,2}R_D R_S C_S}{(1 + g_{m1,2} \frac{R_S}{2})L_D}. \quad (3)$$

These calculations³ have neglected the shunt peaking effect of L_D in the main path because it manifests itself well beyond the Nyquist frequency.

Let us extend the feedforward concept to our entire two-stage CTLE. Shown in Fig. 6, the proposed circuit employs G_{mf1} , G_{mf2} , and G_{mf3} to carry each interface's signal forward. The core stages have the same topology as that in Fig. 3(a), and G_{mf1} and G_{mf2} are chosen according to (3). Similarly, $G_{mf3}L_D2\omega$ begins to lift the gain as the rest of the circuit approaches a flat response. In this design, the two main stages are identical to the arrangement shown in Fig. 3(b). The third path, G_{mf3} , has $W = 4 \mu\text{m}$ and a tail current of 1.5 mA.

The performance of the CTLE is studied as follows. Due to the significant effect of layout parasitics at these frequencies, we perform simulations only based on the extracted layout.⁴ The inductors are included as RLC models obtained from Ansoft's HFSS. Figs. 7–9 plot the simulated CTLE ac response,⁵ group delay and single-bit response (SBR), respectively for four cases: with no feedforward, with only G_{mf1} , with G_{mf1} and G_{mf2} , and with G_{mf1} , G_{mf2} , and G_{mf3} . We observe that the boost factor rises, but the corresponding frequency falls. Nevertheless, the ultimate goal is to counteract the channel loss and, as shown in Fig. 10, the channel-CTLE cascade indeed provides a much flatter response in the presence of the feedforward paths. Fig. 11 plots the simulated eye diagrams at different nodes in the proposed CTLE architecture (see Fig. 6).

While the frequency response improvement due to G_{mf3} in Figs. 7 and 10 appears marginal, this path still opens the eye

³In this work, $g_{m3,4} = 7 \text{ ms}$, which satisfies the upper bound of $g_{m3,4} < 87 \text{ ms}$ in (3) by a considerable margin. In other words, the current consumption of the feedforward stage and the loading on the preceding stage dictate the choice of $g_{m3,4}$ in this design.

⁴The CTLE performance simulations (see Figs. 7–11) include the loading due to high-frequency feedforward and feedback DFE paths at node P (see Section V) and the CDR's PD at node Q (see Section VI).

⁵The simulated CTLE ac response is not normalized, i.e., the CTLE gain at low frequencies is indeed equal to 0 dB.




Fig. 7. Simulated ac response (a) conventional CTLE, (b) proposed CTLE with G_{mf1} , (c) proposed CTLE with G_{mf1} and G_{mf2} , and (d) proposed CTLE with G_{mf1} , G_{mf2} , and G_{mf3} .




Fig. 8. Simulated group delay (a) conventional CTLE, (b) proposed CTLE with G_{mf1} , (c) proposed CTLE with G_{mf1} and G_{mf2} , and (d) proposed CTLE with G_{mf1} , G_{mf2} , and G_{mf3} .




Fig. 9. Simulated SBR (a) conventional CTLE, (b) proposed CTLE with G_{mf1} , (c) proposed CTLE with G_{mf1} and G_{mf2} , and (d) proposed CTLE with G_{mf1} , G_{mf2} , and G_{mf3} .

height by 10 mV and the eye width by 0.5 ps (see Section V), a reasonable result for a cost of 1.5 mA.

V. DFE DESIGN

The 56-Gb/s NRZ data stream at the CTLE output suffers from significant eye closure and jitter. In fact, with a channel loss of 25 dB at Nyquist, the CTLE output eye is still

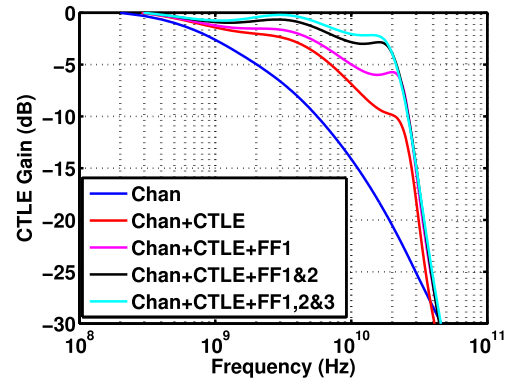


Fig. 10. Frequency response of 25-dB lossy channel cascaded with CTLE (FF $_j$ corresponds to G_{mfj} in Fig. 6).




Fig. 11. Simulated eye diagram in proposed CTLE architecture (see Fig. 6) at (a) stage 1 output, (b) stage 2 output, (c) node Q , and (d) node P .

completely closed. The DFE thus bears a heavy burden of equalization. We propose a number of techniques that enable the DFE to operate at these speeds while drawing 4 mW. We first introduce the concepts and then quantify their efficacy by simulations.

Full-rate and half-rate DFE architectures typically face an extremely tight timing constraint; in Fig. 12(a), we have $t_{CK-Q} + t_{setup} + t_{FB} < 1$ UI, where the three terms respectively represent the flip-flop (FF) clock-to-output delay, the FF setup time, and the feedback delay. The last delay term arises from the feedback tap(s) and the time constant at the DFE summing junction. For unrolled loops [see Fig. 12(b)], this term is replaced with a multiplexer (MUX) delay, which is not necessarily shorter if the MUX demands a rail-to-rail swing from the flip-flop.

It is important to note that these constraints are not relaxed in half-rate architectures. In fact, half-rate operation can exacerbate the timing if the clock waveform is roughly a sinusoid: since the clock transition time is doubled at half rate, the actual time for the loop operation is less. The use of resonant clocks

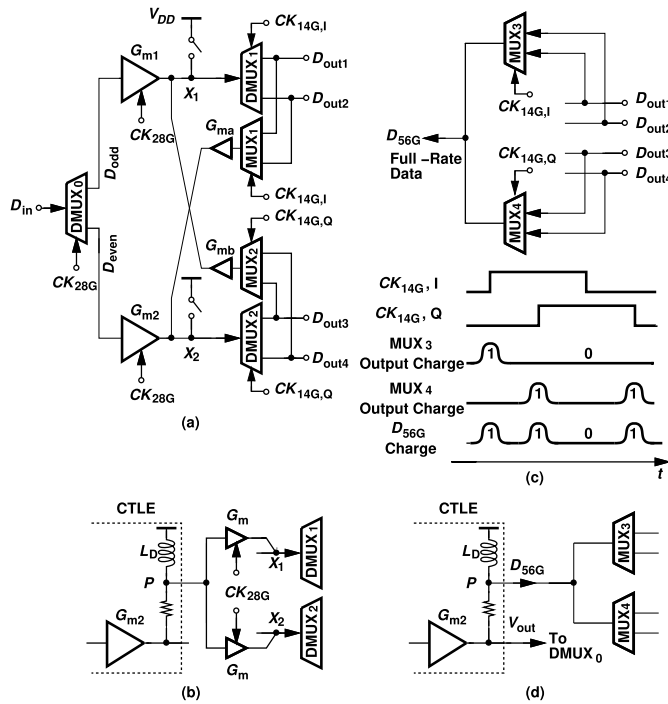


Fig. 15. (a) DFE core (D_{in} is the same node as V_{out} in Fig. 6), (b) high-pass action in feedforward path, (c) reconstruction of full-rate data and charge packets, and (d) high-pass action in feedback path.

to the voltage at X_2 . This bound is typically more relaxed than the first. In our work, for example, the G_{m1} -DMUX₁ delay and the MUX₁ delay are about 16.5 and 12.5 ps, respectively.

We can now add the high-pass branches shown in Fig. 13 to the DFE of Fig. 15(a). But we must ponder how the high-pass functions αs and βs can be realized in our half-rate architecture without additional complexity. The former is implemented by first recognizing that nodes X and Y in Fig. 3(b) carry the high-pass content of V_{out} . This point is also seen in Fig. 4 if we neglect C_p and C_{out} and write $V_{out} = (I_3 + I_1)L_D s + I_1 R_D$ and $V_X = (I_3 + I_1)L_D s$. We simply sense these nodes (called node P in the conceptual diagram of Fig. 6) by clocked G_m stages, as depicted in Fig. 15(b), and inject the results to the summing junctions.⁷

For $H(s) = \beta s$, one possibility is to subject to high-pass action the half-rate data streams at the inputs of G_{m1} and G_{m2} in Fig. 15(a) and at the outputs of MUX₁ and MUX₂, and somehow add the results to X_1 and X_2 . However, a simpler solution is to perform the high-pass filtering at the full rate. For this purpose, we have two options: 1) we can multiplex the outputs of MUX₁ and MUX₂ to obtain full-rate data, and apply the result to a high-pass filter; this is not possible here because our implementation merges the feedback G_m stages with their corresponding MUXes [17]; or 2) we can replicate MUX₁ and MUX₂ and multiplex the replica outputs at the

⁷The delay mismatch between the high-pass feedforward path [node P to node $X_{1,2}$ in Fig. 15(b)] and the main path [node D_{in} to node $X_{1,2}$ in Fig. 15(a)] varies from 1 ps in FF 0 °C to 3.5 ps in SS 80 °C. In view of the 36 ps bit period at this interface, the mismatch negligibly impacts the performance improvement provided by the feedforward path.

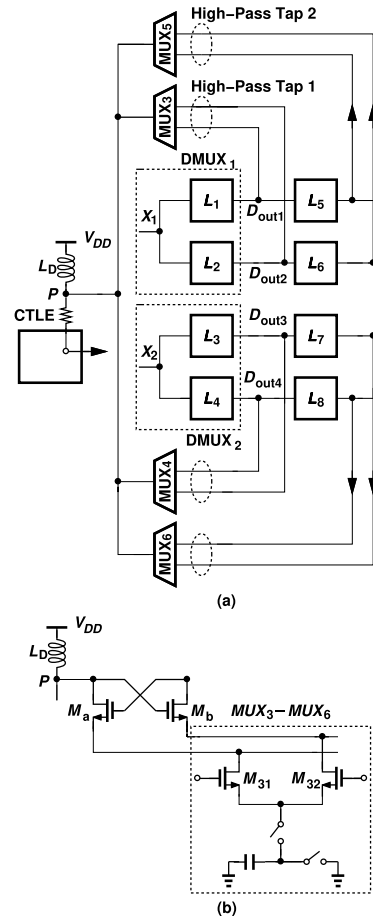


Fig. 16. (a) Two-tap DFE, and (b) reduction of capacitance at P .

cost of additional loading presented to D_{out1} - D_{out4} . We choose the second option, but multiplexing half-rate data lines to obtain 56 Gb/s proves extremely power-hungry.

Fortunately, we recognize that driving MUX₁ and MUX₂ by the quadrature phases of the 14-GHz clock creates a certain phase relationship between the half-rate data streams that simply allows shorting the two multiplexers' outputs. As illustrated in Fig. 15(c), charge-steering replica multiplexers, MUX₃ and MUX₄, produce charge packets corresponding to their inputs at each transition of their respective clocks; that is, each MUX presents a "tri-state" output after delivering its charge, allowing the other MUX to impress its data at the output as well. From another perspective, even though $CK_{14G,I}$ and $CK_{14G,Q}$ are not nonoverlapping clocks, charge steering creates nonoverlap action and hence a 4-to-1 MUX. The strength of the injection can be controlled by adjusting the amount of charge drawn by the MUXes.

In the last step of our development, we recognize that the high-pass branch, $H(s)$, envisioned in Fig. 13 can be implemented by injecting D_{56G} in Fig. 15(c) into node P in Fig. 6 [see Fig. 15(d)].

The foregoing ideas are also applicable to other DFE taps. Fig. 16(a) shows the details of our implementation. Here, L_j denotes a latch. As D_{out1} - D_{out2} travel through L_5 - L_8 , respectively, they are delayed by 1 UI. We apply the results to MUX₅ and MUX₆ for the second tap. The second tap also

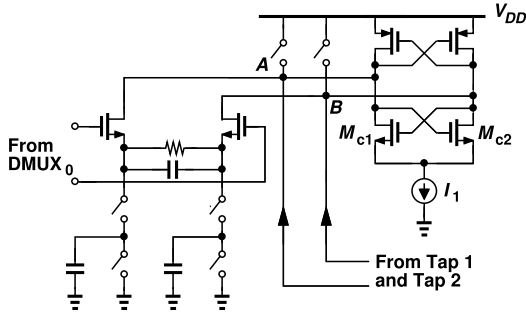


Fig. 17. Proposed summer topology.

includes the conventional G_m stages, not shown here for the sake of clarity. In this work, we decided not to reuse L_1 – L_4 in the DFE as part of the CDR's PD. The extra loading of the edge samplers at the DFE summing junction would exacerbate the DFE timing constraint.

The four MUXes attached to node P in Fig. 15(d) carry the output capacitances of eight differential pairs, thereby reducing the self-resonance frequency of L_D and degrading the CTLE performance. To alleviate this issue, we interpose (in the fully-differential implementation) two cross-coupled cascode transistors between the MUXes and the inductors [see Fig. 16(b)]. For power savings, MUX₃–MUX₆ in Fig. 16(b) are implemented as charge-steering logic [16]. The tail of MUXes employs the structure with two switches and a capacitor that acts as a charge source. The near-VDD CM level at P and the large swings arriving at the gates of M_{31} and M_{32} allow complete charge steering even though these transistors enter the triode region. That is, the voltage headroom consumed by M_a and M_b negligibly affects the speed.

Shown in Fig. 17, each charge-steering DFE summer incorporates a continuous-time cross-coupled pair, M_{c1} and M_{c2} , so as to increase the voltage swings by 50%. The value of I_1 is chosen such that the common-mode droop at nodes A and B remains less than 20 mV in the 18-ps evaluation time.

The efficacy of the proposed CTLE and DFE techniques can be studied by examining the eye diagram at the latter's summing junction in the presence of a channel loss of 25 dB at 28 GHz. Plotted in Fig. 18(a) is the simulated eye before any of the techniques is applied and in Fig. 18(b) after all are. We observe that the width improves from 18.5 to 25 ps and the height from 55 to 200 mV. The incremental improvements arising from each technique are illustrated in Fig. 19.

The minimum acceptable differential eye height in Fig. 15(a) is computed from the BER equation in [18]

$$\text{BER} \approx \frac{1}{2} Q \left(\frac{V_{\text{sj,pp}}/2 - V_{\text{os}}}{\sqrt{V_n^2}} \right) \quad (6)$$

where Q is the error function, $V_{\text{sj,pp}}$ the peak-to-peak swing at the summing junction, V_{os} the total offset, and V_n the total rms noise; the last two are referred to the summing junction. We first subtract from the height the 3σ dc offsets contributed by the CTLE and the summer itself and the input offset of




Fig. 18. Eye diagrams at the DFE summing junction, (a) without, and (b) with new techniques.




Fig. 19. Eye height and width improvement due to proposed techniques (A: original design; B: CTLE feedforward 1; C: CTLE feedforward 1 and 2; D: CTLE feedforward 1, 2, and 3; E: DFE high-pass feedback branch; F: DFE high-pass input branch; and G: cross-coupled pair at the summing junction).

the latches within DMUX₁ and DMUX₂ in Fig. 15(a). The remainder must exceed approximately $14V_n$ for $\text{BER} < 10^{-12}$. According to Monte Carlo simulations, the 3σ dc offset is approximately equal to 31 mV.⁸ Also, pss and pnoise simulations in Cadence yield the noise spectrum plotted in Fig. 20 at the summer output (excluding the input-referred noise of the subsequent latches). Integrated from 1 MHz to 28 GHz, this noise amounts to 4.2 mV_{rms}.⁹ The input noise of the latches is estimated from [19] to be about 0.9 mV_{rms}. Thus, the differential eye height must exceed $31 \text{ mV} + 14 \times 4.26 \text{ mV} = 90 \text{ mV}$.

The dual-loop two-tap DFE and the DTLE draw 4 mW, making the CTLE still the dominant power-hungry block in the data path.

VI. CDR DESIGN

A. CDR Input Interface

The receiver data path offers a multitude of ports that can feed the CDR circuit. Fig. 21(a)–(d) summarize some possible candidates. The final choice of the CDR input interface is given by the tradeoff between: 1) how much the data path degrades as a result of the CDR's input capacitance; and 2) how much the CDR's bandwidth and capture range degrade

⁸In this work, the offset of these latches is not canceled because the vertical eye opening is sufficiently large to ensure robust operation.

⁹Analysis of the DFE noise should take into account the correlation between the CTLE and DFE noise sources and perhaps the CDR bandwidth. If, for example, the integration starts from the CDR bandwidth (50 MHz), the noise amounts to 4 mV_{rms} instead of 4.2 mV_{rms}. Nevertheless, we have selected the DFE swings to be far higher than the limitations due to the DFE noise.




Fig. 20. Noise spectrum at summing junction.

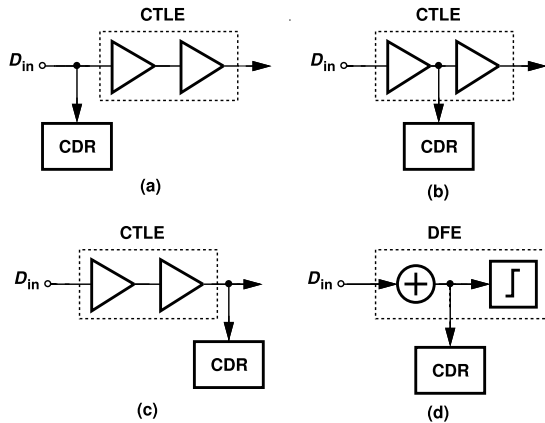


Fig. 21. Possible data path ports for driving the CDR circuit. (a) Main input. (b) Output of first CTLE stage. (c) Output of second CTLE stage. (d) DFE summing junction.

as a result of applying attenuated data to the CDR, a point that we further discuss below. The former effect generally intensifies as we go from Fig. 21(a) to (d) whereas the latter improves because of the lesser eye closure as we select a port deeper in the receive chain.

With a channel loss of 25 dB at 28 GHz, the arrangement in Fig. 21(a) does not enable the CDR to lock. This is because the heavily attenuated data dramatically reduces the gain of the CDR's PD. According to transistor-level simulations, our PD gain drops by a factor of 10 if the channel loss rises from 6 to 25 dB. In Fig. 21(b)–(d), the CDR input capacitance nearly doubles the capacitance at the sensing port, severely degrading the equalization function. While it is possible to drive the CDR by a replica CTLE, such a solution would draw substantial power and would require additional inductors, thereby complicating the floor plan.

The CTLE topology described in Section IV does provide two ports that are more tolerant of load capacitance, namely, the bottom terminals of the inductors in the first and second stages in Fig. 6. We then ask whether it is possible to drive a CDR from a high-pass node. Specifically, we must answer two questions: 1) can a CDR operate properly if its input data is subjected to a high-pass filter (HPF)? and 2) if so, which one of the nodes P and Q in Fig. 6 is preferable?

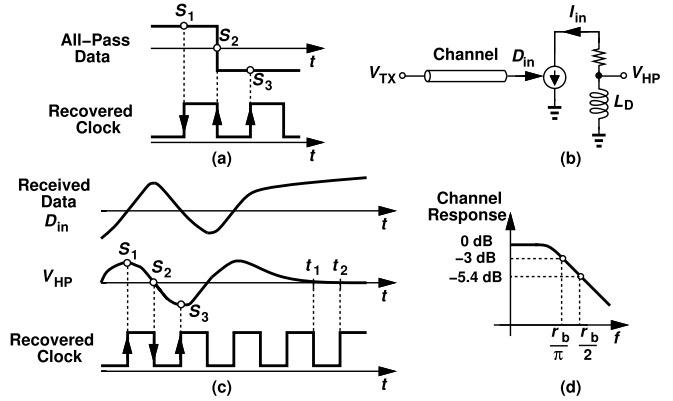


Fig. 22. (a) Early-late PD operation, (b) lossy channel followed by HPF, (c) circuit's waveforms, and (d) channel response necessary for lock.

To answer the first question, let us first consider a full-rate Alexander PD [20], whose input data and clock waveforms are shown in Fig. 22(a). The PD takes samples S_1, S_2 , etc., by means of flipflops and forms $S_1 \oplus S_2, S_2 \oplus S_3$, etc., to decide whether the clock is early or late. In the absence of data transitions, these XORed quantities are zero, and the PD does not update the voltage stored on the CDR loop filter. We now study the behavior of the circuit if the data travels through the channel and a high-pass stage [see Fig. 22(b)]. Suppose the limited bandwidth of the channel yields a transition time of about 1 UI for D_{in} . As a result, the high-pass voltage, $V_{HP} = L_D dI_{in}/dt$, remains high or low for roughly 1 UI [see Fig. 22(c)]. That is, for data sequences such as 101010, V_{HP} is similar to V_{TX} and the Alexander PD can still receive proper samples from this waveform.

To obtain a simple rule of thumb for the 1-UI transition time, we approximate the channel by a one-pole LPF and assume the transition time of I_{in} in Fig. 22(b) is about two time constants. That is, we wish to have $2\tau \approx 1$ UI, arriving at a -3 -dB bandwidth of $1/(2\pi\tau) \approx 1/(2\pi \times 0.5 \text{ UI})$, i.e., about $1/\pi$ times the bit rate, r_b [see Fig. 22(d)]. This implies that the loss at half of the bit rate (i.e., at the Nyquist frequency) is approximately equal to

$$\text{Loss at Nyquist} = 10 \log \left[1 + \left(\frac{r_b/2}{r_b/\pi} \right)^2 \right] \quad (7)$$

$$\approx 5.4 \text{ dB}. \quad (8)$$

We conclude that a minimum loss of roughly 5.4 dB is necessary so that the high-pass response still produces a proper V_{HP} waveform in Fig. 22(c), allowing acceptable PD gain and hence lock.¹⁰

Let us now consider the case of longer runs in the input data. As mentioned above, an Alexander PD operating with all-pass data yields no update on the loop filter during long runs. For a high-pass CDR, on the other hand, we observe in Fig. 22(c) that V_{HP} approaches zero. The PD flipflops thus become metastable at sampling points such as t_1 and

¹⁰From another perspective, the input data edge rate cannot be arbitrarily high for the CDR to lock; this point is relatively independent of the data rate.

t_2 . Fortunately, the XOR gates sensing these flipflops' outputs produce a zero differential output if one or both of their inputs are metastable,¹¹ negligibly disturbing the loop filter.

The foregoing calculations also answer the second question, namely, which port in Fig. 6 is preferable. The additional bandpass gain provided by the second CTLE stage raises the 5.4-dB lower bound in (8) by about 6 dB. That is, only channels having losses greater than 12 dB allow lock. We therefore select node Q . The performance of the high-pass PD is quantified in Section VI-C.

B. Problem of Phase Alignment

Before describing the CDR architecture, we should make a remark about the problem of phase alignment between the data and the clock. The CDR circuit must recover the clock with proper phase with respect to the waveforms appearing at the DMUX inputs in Fig. 15(b) such that latches L_1 and L_3 sample the data in the middle of the eye. If the CDR senses these signals, e.g., if L_1 and L_3 also act as part of the CDR's PD, then the phase alignment occurs naturally. But since we have chosen a port within the CTLE for driving the CDR, it is unlikely that these latches sample at optimal points.

The foregoing issue appears in receivers that do not attach the CDR directly to the DFE summing junctions, demanding that the CDR output phase be adjustable [10]. In this work, we provide an adjustment range from -0.25 UI to $+0.25$ UI by simple tunable buffers. The output buffer of the CDR employs delay control by means of programmable capacitors. Alternatively, a phase interpolator can be used [2].

In this work, the delay is adjusted through a serial bus, but a number of automatic methods can be envisaged. For example, if an eye monitor is employed to determine how the DFE coefficients must be set, it can also command this delay adjustment so as to maximize the eye height and width. But a simpler approach is to benefit from loop-back operation performed in transceivers [24], wherein a known data pattern is delivered by the transmitter to the receiver input, the BER is measured, and the various coefficients are so adjusted as to optimize the performance [25]. Fig. 23(a) depicts the idea conceptually. Here, the select command, SEL, can disable G_{m1} , G_{mf1} , and G_{mf3} , and enable G_{m0} so that the loop-back data is applied to the CTLE. According to simulations, if G_{m0} is scaled down by a factor of 2 with respect to G_{m1} it has a negligible effect on the bandwidth and, by virtue of the CTLE boost, yields large swings at the DFE summing node [see Fig. 23(b)].

C. CDR Architecture

The use of a half-rate CDR greatly simplifies the generation and distribution of the recovered clock. However, half-rate PDs present other challenges. The linear PD in [21] does not require quadrature clock phases but its gain drops considerably with high-loss channels. This is because this PD produces

¹¹In this work, each XOR gate is implemented as in [15]. However, even Gilbert-cell XORs produce a zero differential output if one or both of their inputs are metastable.

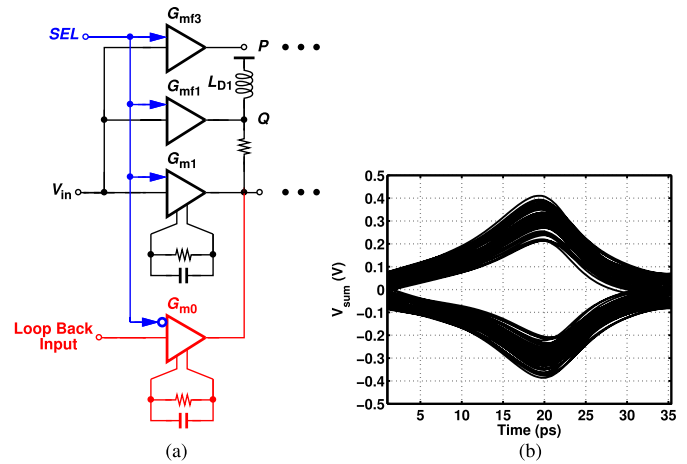


Fig. 23. (a) Conceptual receiver front-end depicting the loop back path, and (b) eye diagram at the DFE summing junction in the loopback mode.

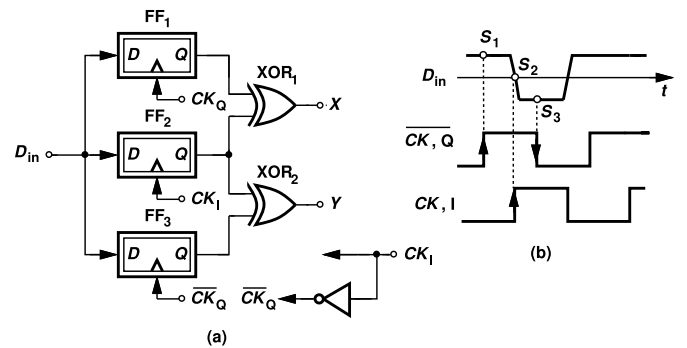


Fig. 24. (a) Half-rate Alexander PD using inverter for quadrature clock generation, and (b) its waveforms.

the samples by latches rather than by flipflops, experiencing metastability frequently.

It is possible to construct a half-rate Alexander PD as shown in Fig. 24(a). Here, the sampled outputs, S_1 , S_2 , and S_3 , are provided by flipflops and hence the problem of metastability is much less severe. Nevertheless, the need for quadrature clock phases, CK_I and CK_Q , leads to substantial complexity and power consumption.

We can then ask, is it possible to employ only a differential VCO and generate these phases using simple, inevitably poorly-matched circuits, as depicted by the inverter in Fig. 24(a)? In other words, how does the mismatch between CK_I and CK_Q affect the performance of the CDR circuit?

Assuming that the CDR is locked and the rising edges of CK_I sample the data transitions, we recognize that an inverter delay of ΔT displaces the edges of CK_Q from the center of the data eye [see Fig. 24(b)]. Thus, if $T/2 - \Delta T$ is moderately small,¹² it has no effect on the PD operation. Indeed, simulations reveal that the PD gain changes negligibly as ΔT varies from 6 to 11.5 ps (see Fig. 25). We should remark that, even if S_1 or S_3 in Fig. 24(b) shift so much as to incur errors, the data integrity is preserved because the DFE latches—and not the PD flipflops—produce the output data.

¹²Here, T denotes UI ≈ 17.5 ps.

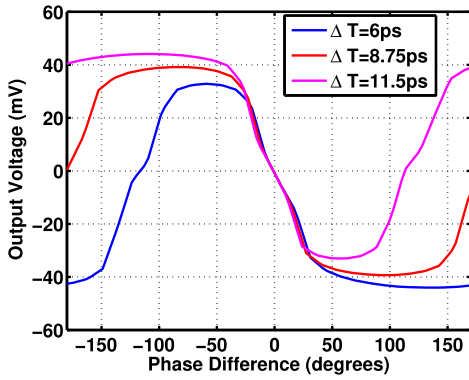


Fig. 25. PD characteristics with imprecise quadrature clock phases.

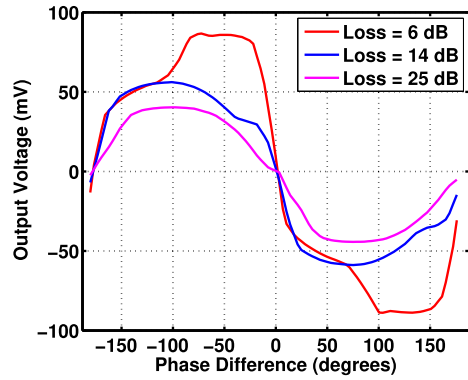


Fig. 27. PD characteristics of high-pass CDR for channel loss from 6 to 25 dB.

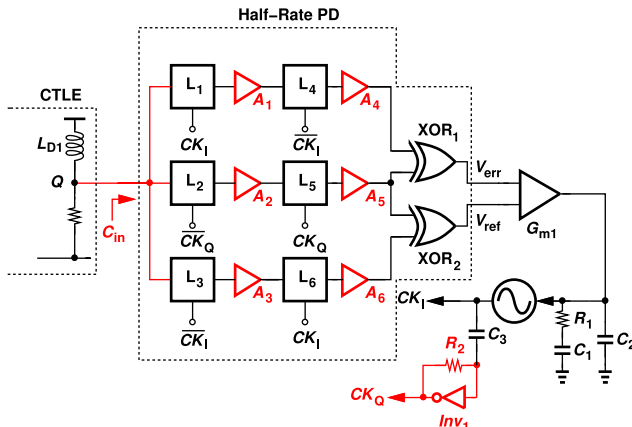


Fig. 26. CDR architecture depicting proposed PD.

The PD topology of Fig. 24(a) faces another challenge with respect to sensing 56-Gb/s data: for a channel loss of 25 dB, the vertical eye opening at the PD input is around 20 mV. Thus, the PD flipflops must achieve both a high speed and a high sensitivity, demanding inductive peaking. However, the use of six latches in the three flipflops leads to an exceedingly complex floor plan for the receiver. Without inductors, the flipflops deliver small swings (~ 25 mV_{pp}) to the XOR gates, and the PD gain is insufficient to guarantee lock. We, therefore, follow each latch with a gain stage consisting of a simple differential pair (see Fig. 26). According to simulations, this approach raises the eye opening at the XOR inputs to 65 mV_{pp}. As a result, the PD gain rises by a factor of 2.6, affording a robust lock for the CDR loop.

Fig. 27 plots the simulated characteristics of the high-pass PD for different channel losses. We observe an adequate gain for losses as low as 6 dB. As with conventional designs, the PD gain falls at high loss values [22].

Fig. 28 demonstrates that the simulated characteristic of the high-pass PD does not change much with the data pattern, again because the PD “tri-states” for long runs.

The CDR design shown in Fig. 26 merits two more remarks. First, the 28-GHz LC VCO¹³ is designed as a conventional NMOS cross-coupled pair with a bias current of 5.2 mA.

¹³For multi-lane transceivers, LC VCOs can utilize “candy-shaped” inductors to minimize pulling [23].

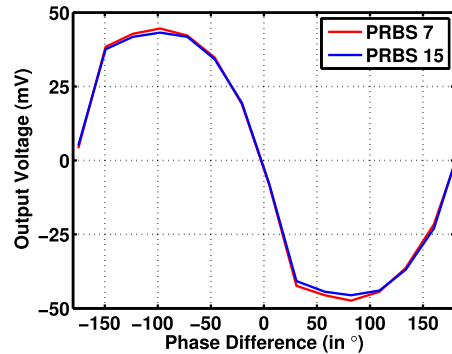


Fig. 28. PD characteristics of high-pass CDR for PRBS length 7 and 15.

Second, a three-stage buffer drawing 7.5 mA follows the VCO to drive the DFE, the DTLE, and a $\div 2$ stage. The relatively high power values are dictated by the interconnect parasitics.

VII. RECEIVER IMPLEMENTATION

Fig. 29 illustrates how the building blocks described in the previous sections form the overall receiver. For simplicity, the second DFE tap is not shown. Highlighted here are the CTLE feedforward paths, the DFE high-pass branches, the gain stages in the PD, and the simple quadrature clock generator.

The prototype reported here does not contain automatic gain control (AGC). According to simulations, the CTLE output, the DFE summing node eye diagram, and the CDR performance remain acceptable for differential input swings as large as 600 mV_{pp}. To accommodate larger swings, AGC techniques can be employed [26], [27].

VIII. EXPERIMENTAL RESULTS

The 56-Gb/s NRZ receiver has been fabricated in TSMC’s 28-nm CMOS technology. Fig. 30 shows the die photograph with an active area $250 \mu\text{m} \times 275 \mu\text{m}$. The die has been directly mounted on a printed-circuit board and high-speed signals are carried through probes. The prototype consumes 48.7 mW: 9.1 mW in the CTLE, 6.3 mW in the DFE, DTLE, and DMUXes, 19.3 mW in the CDR (11 mW of which are

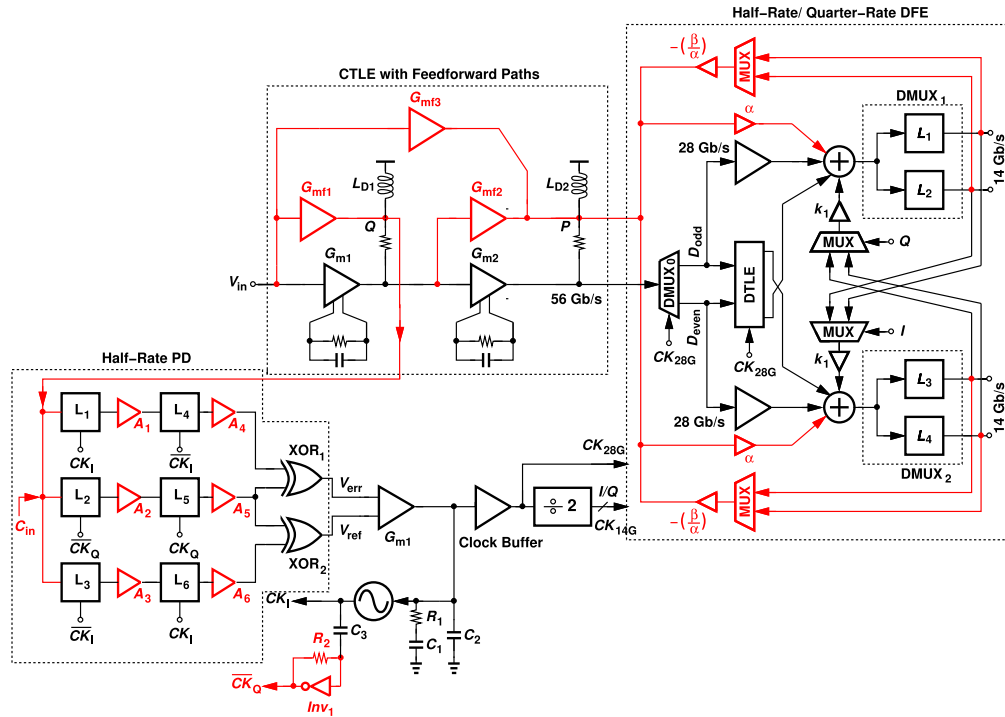


Fig. 29. Overall RX architecture.




Fig. 30. Receiver die photograph.

consumed by the PD), 8.9 mW in the CDR buffer, and 5.1 mW in a $\div 2$ stage that generates the quarter-rate clocks.

About a dozen different measurements have been made to characterize the receiver performance. All are carried out with a channel loss of at least 25 dB at the Nyquist rate using the Keysight’s boards, M8049A-002 and M8049A-003. Fig. 31 plots the measured responses of traces 2 and 4 on the M8049A-003 board, respectively, while driving a 30-in cable. To the loss at 28 GHz, we add 1.7 dB to account for the probes and interconnects. All measurements use a 56-Gb/s pseudorandom bit sequence (PRBS) of length $2^7 - 1$.

We first describe the experimental results for the receiver with the CDR circuit disabled. An external 28-GHz clock allows us to quantify the performance of the CTLE, the DTLE, and the DFE and construct the bathtub curve. In this measurement, the bit error rate tester (BERT) (Keysight’s M8040A)




Fig. 31. Measured channel frequency response (excluding 1.45 dB @ 14 GHz and 1.7 dB @ 28 GHz insertion loss of probes and interconnects).


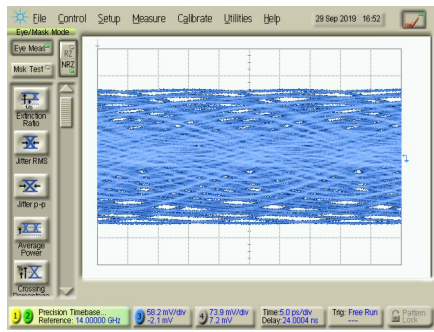
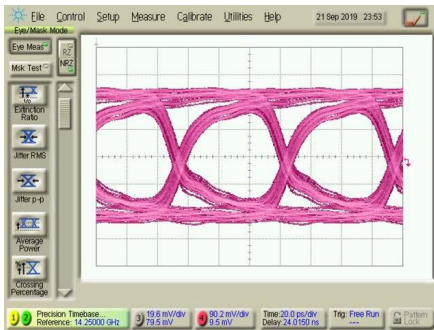


Fig. 32. Measured bathtub curve for 56-Gb/s data.

provides the capability to emulate a 2-tap transmit feedforward equalizer (FFE) in the data applied to the channel. Fig. 32 plots the measured bathtub curves for two cases: 1) for channel A,

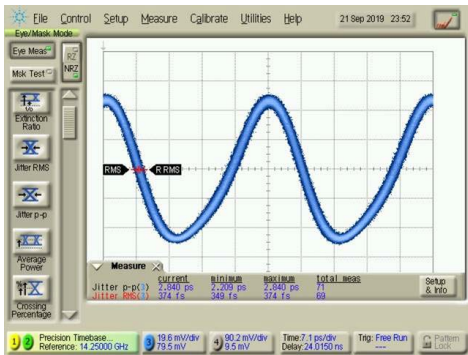


(a)




(b)

Fig. 33. Measured eye diagram of (a) channel output at 56 Gb/s (58.2 mV/div, 5 ps/div), and (b) receiver demultiplexed output data at 14.25 Gb/s (90.2 mV/div, 20 ps/div).



(a)



(b)

Fig. 34. Measured (a) receiver recovered clock eye diagram (19.6 mV/div, 7.1 ps/div), and (b) Spectrum of recovered clock.

which has a loss of 25 dB at 28 GHz, (with no FFE); and 2) for channel B, which has a loss of 30 dB at 28 GHz, while the BERT is set to an FFE function of the form $-0.2 +$




Fig. 35. Phase noise of the recovered clock at frequency offsets 100 Hz–100 MHz.

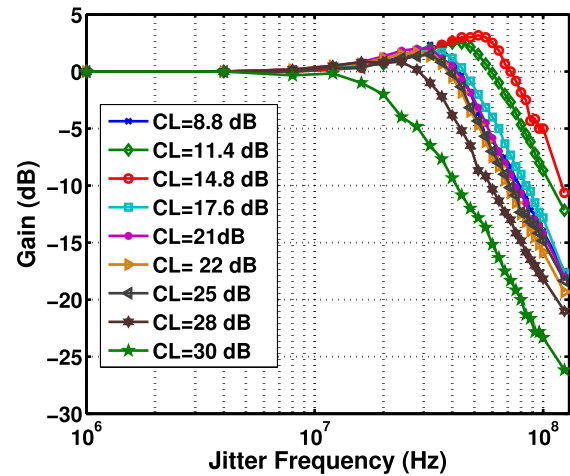


Fig. 36. Measured jitter transfer versus channel loss.

$0.8 z^{-1}$. The horizontal eye openings are 0.4 UI and 0.33 UI, respectively.

The remaining measurements include the CDR circuit as well. Fig. 33(a) and (b) respectively plot the measured eye diagrams at the channel output and at the receiver output. The BER is less than 10^{-12} . Fig. 34 shows the recovered clock in the time and frequency domains. The spectrum reveals a VCO noise-shaping bandwidth of about 50 MHz. The recovered clock phase noise is plotted in Fig. 35 for frequency offsets up to 100 MHz (due to equipment limitation), at which it is equal to -124.4 dBc/Hz. For greater offsets, we have measured the phase noise directly from the spectrum, which falls to -128 dBc/Hz at 14-GHz offset. We thus obtain an rms jitter of 500 fs integrated from 100 Hz to 14 GHz. The measured VCO tuning range is 26.9 to 30 GHz.

Fig. 36 plots the measured CDR jitter transfer for different channel losses, which are obtained by cascading different sections of Keysight’s boards, M8049A-002 and M8049A-003, and different cable lengths. For the 25-dB loss case, the -3 -dB loop bandwidth is around 55 MHz, consistent with the VCO noise-shaping bandwidth reported above. The high-pass nature of the CDR gives rise to some peaking for low loss values,

TABLE I
PERFORMANCE SUMMARY

Reference	[5]	[6]	[19]	[7]	[9]	[8]	This Work
Modulation	NRZ	PAM4	NRZ	PAM4	NRZ	PAM4	NRZ
Data Rate (Gb/s)	56	56	60	64	56	56	56
Architecture	CTLE 1-tap DFE	CTLE 3-tap DFE	2-tap RX FFE CTLE 3-tap DFE	CTLE	3-tap DFE	CTLE 1-tap FIR DFE 1-tap IIR DFE	CTLE with High-Pass FF Path DTLE Dual-loop DFE 2 conv. & 2 HP taps
Channel Loss	18.4 dB* @ 28 GHz	24 dB** @ 14 GHz	21 dB** @ 30 GHz	16.8 dB*** @16 GHz	37.8 dB* @ 28 GHz	20.8 dB* @ 28 GHz	30 dB* @ 28 GHz 16.5 dB* @ 14 GHz 25 dB @ 28 GHz 13.5 dB @ 14 GHz
Horizontal Eye (UI)	0.28 @ 10 ⁻⁹	0.25 @ 10 ⁻¹²	0.3 @ 10 ⁻¹²	0.19 @ 10 ⁻⁶	0.44 @ 10 ⁻¹²	0.19 @ 10 ⁻¹²	0.4 @ 10 ⁻¹²
Clock Jitter (fs, rms)	–	688 (100 Hz–1 GHz)	–	–	–	–	500 (100 Hz–14 GHz)
PRBS	15	7	7	Q 13	15	15	7
Power (mW)	Incl. ^s	141.7	382	136	–	–	49.56
	Excl. ^{ss}	–	–	–	180	112	43.6
Power Eff. (pJ/bit)	Incl. ^s	2.53	6.82	2.26	–	–	0.88
	Excl. ^{ss}	–	–	–	2.81	2	0.77
Area (mm ²)	1.4 [#]	1.26	2.03	0.32	0.053	0.51	0.102
Technology	28-nm CMOS	40-nm CMOS	65-nm CMOS	28-nm FDSOI	14-nm FINFET	65-nm CMOS	28-nm CMOS

*Includes 2-tap TX FFE **Includes 3-tap TX FFE ***Includes 4-tap TX FFE # Includes TX area
^s Includes Clock Gen. ^{ss} Excludes Clock Gen.




Fig. 37. Measured jitter tolerance of the receiver.

but it helps the CDR maintain a reasonable bandwidth (about 25 MHz) for losses as high as 30 dB.

Fig. 37 plots the measured CDR jitter tolerance for a channel loss of 25 dB, yielding a value of 1.1 UI_{pp} at 5 MHz and exceeding the CEI-56G-VSR mask.

Table I summarizes our receiver’s performance and compares it with that of the prior art. We should make several remarks. First, some of the reported channels are preceded with FFEs having 2 to 4 taps. Second, some receivers do not include clock generation. Third, in comparison to the 14-nm NRZ design in [12], we have achieved a 2.2× reduction in power. But since [12] requires an external 28-GHz clock, excluding the 6 mW that our VCO draws

gives an improvement factor of greater than 2.5× in 28-nm technology, albeit for a loss of 30 dB.

IX. CONCLUSION

This paper presents a number of circuit and architecture techniques that alleviate the tradeoffs among speed, power consumption, and channel loss compensation. A 56-Gb/s NRZ receiver employing these concepts has been demonstrated, achieving more than twofold reduction in power in 28-nm CMOS technology.

REFERENCES

- [1] P. Upadhyaya *et al.*, “A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with a configurable ADC in 16nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 108–110.
- [2] T. Ali *et al.*, “6.4 A 180 mW 56Gb/s DSP-based transceiver for high density I/Os in data center switches in 7nm FinFET technology,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 118–120.
- [3] J. Im *et al.*, “6.1 A 112Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR-ADC and inverter-based RX analog front-end in 7nm FinFET,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 116–118.
- [4] T. Ali *et al.*, “6.2 A 460 mW 112Gb/s DSP-based transceiver with 38dB loss compensation for next-generation data centers in 7nm FinFET technology,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2020, pp. 118–120.
- [5] A. Atharav and B. Razavi, “11.7 A 56Gb/s 50 mW NRZ receiver in 28nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2021, pp. 192–194.
- [6] T. Shibusaki *et al.*, “3.5 A 56Gb/s NRZ-electrical 247 mW/lane serial-link transceiver in 28nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Jan. 2016, pp. 64–66.
- [7] P. J. Peng, J. F. Li, L. Y. Chen, and J. Lee, “A 56Gb/s PAM-4/NRZ transceiver in 40nm CMOS,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 110–111.

- [8] B. Razavi, "Jitter-power trade-offs in PLLs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 4, pp. 1381–1387, Apr. 2021.
- [9] B. Dehlaghi *et al.*, "A 1.41-pJ/b 56-Gb/s PAM-4 receiver using enhanced transition utilization CDR and genetic adaptation algorithms in 7-nm CMOS," *IEEE Solid-State Circuits Lett.*, vol. 2, no. 11, pp. 248–251, Nov. 2019.
- [10] 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," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3474–3485, Dec. 2017.
- [11] E. Depaoli *et al.*, "A 4.9pJ/b 16-to-64Gb/s PAM-4 VSR transceiver in 28nm FDSOI CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 112–114.
- [12] A. Cevrero *et al.*, "A 100Gb/s 1.1pJ/b PAM-4 RX with dual-mode 1-tap PAM-4 / 3-tap NRZ speculative DFE in 14nm CMOS FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2019, pp. 112–114.
- [13] A. Roshan-Zamir *et al.*, "A 56-Gb/s PAM4 receiver with low-overhead techniques for threshold and edge-based DFE FIR- and IIR-tap adaptation in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 672–684, Mar. 2019.
- [14] S. Gondi and B. Razavi, "Equalization and clock and data recovery techniques for 10-Gb/s CMOS serial-link receivers," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1999–2011, Sep. 2007.
- [15] A. Manian and B. Razavi, "A 40-Gb/s 14-mW CMOS wireline receiver," *IEEE J. Solid-State Circuits*, vol. 52, no. 9, pp. 2407–2421, Sep. 2017.
- [16] J. W. Jung and B. Razavi, "A 25-Gb/s 5-mW CMOS CDR/deserializer," *IEEE J. Solid-State Circuits*, vol. 48, no. 3, pp. 684–697, Mar. 2013.
- [17] J. W. Jung and B. Razavi, "A 25 Gb/s 5.8 mW CMOS equalizer," *IEEE J. Solid-State Circuits*, vol. 50, no. 2, pp. 515–526, Feb. 2015.
- [18] S. Ibrahim and B. Razavi, "Low-power CMOS equalizer design for 20-Gb/s systems," *IEEE J. Solid-State Circuits*, vol. 46, no. 6, pp. 1321–1336, Jun. 2011.
- [19] P. Nuzzo, F. De Bernardinis, P. Terreni, and G. Van der Plas, "Noise analysis of regenerative comparators for reconfigurable ADC architectures," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 6, pp. 1441–1454, Jul. 2008.
- [20] J. D. H. Alexander, "Clock recovery from random binary data," *Elect. Lett.*, vol. 11, no. 22, pp. 541–542, Oct. 1975.
- [21] J. Savoj and B. Razavi, "A 10-Gb/s CMOS clock and data recovery circuit with a half rate linear phase detector," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 761–768, May 2001.
- [22] M. Hossain, Aurangozeb, and N. Nguyen, "DDJ-adaptive SAR TDC-based timing recovery for multilevel signaling," *IEEE J. Solid-State Circuits*, vol. 54, no. 10, pp. 2833–2844, Oct. 2019.
- [23] M. Pisati *et al.*, "A 243-mW 1.25–56-Gb/s continuous range PAM-4 42.5-dB IL ADC/DAC-based transceiver in 7-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 55, no. 1, pp. 6–18, Jan. 2020.
- [24] M. Erett *et al.*, "A 126 mW 56Gb/s NRZ wireline transceiver for synchronous short-reach applications in 16nm FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 274–276.
- [25] E. H. Chen *et al.*, "Near-optimal equalizer and timing adaptation for I/O links using a BER-based metric," *IEEE J. Solid-State Circuits*, vol. 43, no. 9, pp. 2144–2156, Sep. 2008.
- [26] T. Ali, M. Abdullatif, H. Park, E. Chen, R. Awad, and M. Gandara, "56/112Gbps wireline transceivers for next generation data centers on 7nm FinFET CMOS technology," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Apr. 2021, pp. 1–6.
- [27] J. Im *et al.*, "A 40-to-56 Gb/s PAM-4 receiver with ten-tap direct decision-feedback equalization in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3486–3502, Dec. 2017.



Atharav Atharav (Member, IEEE) received the B.E. degree in electronics and communication engineering from Birla Institute of Technology and Science at Pilani, Pilani, India, in 2012, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles (UCLA), Los Angeles, CA, USA, in 2015 and 2020, respectively.

He worked as an Analog Design Intern with Texas Instruments, Bangalore, India, Silicon Labs, Sunnyvale, CA, USA, and Broadcom Inc., Irvine, CA, USA, and an Analog Design Engineer with Redpine Signals, Hyderabad, India. He is currently a Staff-Engineer with High-Speed Wireline Group, Mediatek, Irvine. His current research interest includes ultrahigh-speed wireline transceivers.

Dr. Atharav was a recipient of the Analog Devices Outstanding Student Designer Award in 2017.



Behzad Razavi (Fellow, IEEE) received the B.S. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1985, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1988 and 1992, respectively.

He was an Adjunct Professor with Princeton University, Princeton, NJ, USA, from 1992 to 1994, and Stanford University in 1995. He was with AT&T Bell Laboratories, Holmdel, NJ, USA, and Hewlett-Packard Laboratories, Palo Alto, CA, USA, until 1996. Since 1996, he has been an Associate Professor and a subsequently Professor of electrical engineering with the University of California at Los Angeles, Los Angeles, CA, USA. He is the author of *Principles of Data Conversion System Design* (IEEE Press, 1995), *RF Microelectronics* (Prentice Hall, 1998, 2012) (translated to Chinese, Japanese, and Korean), *Design of Analog CMOS Integrated Circuits* (McGraw-Hill, 2001, 2016) (translated to Chinese, Japanese, and Korean), *Design of Integrated Circuits for Optical Communications* (McGraw-Hill, 2003 and Wiley, 2012), *Design of CMOS Phase-Locked Loops* (Cambridge University Press, 2020), and *Fundamentals of Microelectronics* (Wiley, 2006) (translated to Korean, Portuguese, and Turkish). His current research interests include wireless and wireline transceivers and data converters.

Prof. Razavi is currently a member of the U.S. National Academy of Engineering. He has served as an IEEE Distinguished Lecturer. He received the Beatrice Winner Award for Editorial Excellence from the 1994 International Solid State Circuits Conference (ISSCC), the Best Paper Award from the 1994 European Solid-State Circuits Conference, the Best Panel Award from 1995 ISSCC and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, the Best Paper Award from the IEEE Custom Integrated Circuits Conference in 1998, and the McGraw-Hill First Edition of the Year Award in 2001. He was a co-recipient of both the Jack Kilby Outstanding Student Paper Award and the Beatrice Winner Award for Editorial Excellence at 2001 ISSCC. He also received the Lockheed Martin Excellence in Teaching Award in 2006, the University of California at Los Angeles (UCLA) Faculty Senate Teaching Award in 2007, and the Custom Integrated Circuits Conference (CICC) Best Invited Paper Award in 2009 and 2012. He was the co-recipient of the 2012 and 2015 VLSI Circuits Symposium Best Student Paper Awards and the 2013 CICC Best Paper Award. He was also recognized as one of the top ten authors in the 50-year history of ISSCC. He also received the 2012 Donald Pederson Award in Solid-State Circuits. He was a recipient of the American Society for Engineering Education PSW Teaching Award in 2014. He also received the 2017 IEEE CAS John Choma Education Award. He has served on the Technical Program Committee for ISSCC from 1993 to 2002 and the VLSI Circuits Symposium from 1998 to 2002. He has also served as the Guest Editor and an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and *International Journal of High Speed Electronics*. He is also an Editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996) and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003). He also serves as the Editor-in-Chief for the IEEE SOLID-STATE CIRCUITS LETTERS.