

Design Techniques for CMOS Wireline NRZ Receivers Up To 56 Gb/s

BEHZAD RAZAVI^{id}

Department of Electrical Engineering, University of California at Los Angeles, Los Angeles, CA 90095, USA

CORRESPONDING AUTHOR: B. RAZAVI (e-mail: razavi@ee.ucla.edu)

This work was supported in part by Realtek Semiconductor and in part by Oracle.

ABSTRACT Wireline receivers continue to target higher data rates, posing great challenges at circuit and architecture levels. Governed by tradeoffs among speed, power consumption, and channel loss (CL), receiver designs can benefit from new methods that push the performance envelope. This paper presents a number of techniques that allow non-return-to-zero data rates as high as 40 and 56 Gb/s in 45-nm and 28-nm CMOS technologies, respectively. The prototypes operate with a CL of 19-25 dB and a bit error rate of less than 10^{-12} .

INDEX TERMS Continuous-time linear equalizer (CTLE), demultiplexers (DMUX), equalization, feed-forward, SERDES, serial links.

I. INTRODUCTION

THE GROWING demand for greater throughput rates in data centers and edge computing presents significant challenges to physical layer designers. Wireline transceivers have been under intense development [1], [2], [3], [4], [5], [6], [7], [8], [9], targeting speeds as high as 224 Gb/s. This trend is also accompanied by issues regarding the power consumption—both in absolute value (which dictates packaging and heat removal costs) and as the amount of energy per bit (which determines the efficiency of serialization and hence the number of lanes).

This paper serves as a companion to [10] and describes receiver (RX) design techniques that can improve the achievable data rate while saving power. The ideas are presented in the context of 40-Gb/s [11] and 50-Gb/s [12] receivers operating with non-return-to-zero (NRZ) data. Realized in 28-nm and 45-nm technologies, respectively, the designs demonstrate concepts that can lead to higher speeds in more advanced process nodes.

II. GENERAL CONSIDERATIONS

A. CHANNEL CHARACTERIZATION

The design of a wireline receiver is dictated by the properties of the channel that precedes it. Imperfections, such as loss and impedance discontinuities, “distort” the data as it travels through the channel, requiring that the RX provides

sufficient compensation for successful data recovery.¹ We must therefore employ a reasonably realistic channel model in our RX design efforts.

A given channel can be modeled by an electromagnetic field simulator or a network analyzer, with the results typically expressed as S-parameters. In transceiver design, however, we prefer a *scalable* model so that the link behavior can be assessed for different amounts of loss. The scalability proves especially critical to the design of RX building blocks as it reveals the limits of their performance.

Copper media, such as printed-circuit-board traces, suffer from three nonidealities: 1) loss due to skin effect; 2) loss due to the dielectric underneath or surrounding the signal line; and 3) impedance discontinuities arising from connectors and line cards. The former two require a frequency-dependent model, as exemplified by the section shown in Fig. 1(a) [13]. Obtained empirically from simulations of 50- Ω traces on FR4 boards, this scalable representation accounts for skin effect by R_2 and L_2 (at high frequencies, the series resistance rises from $R_1 || R_2$ to R_1) and dielectric loss by R_3 and R_4 . As an example, [13] reports the following values for a section corresponding to a 1-in trace: $L_1 = 77.25$ nH, $C_1 = 30.9$ pF (such

1. The transmitter also offers a modest amount of compensation for the channel.

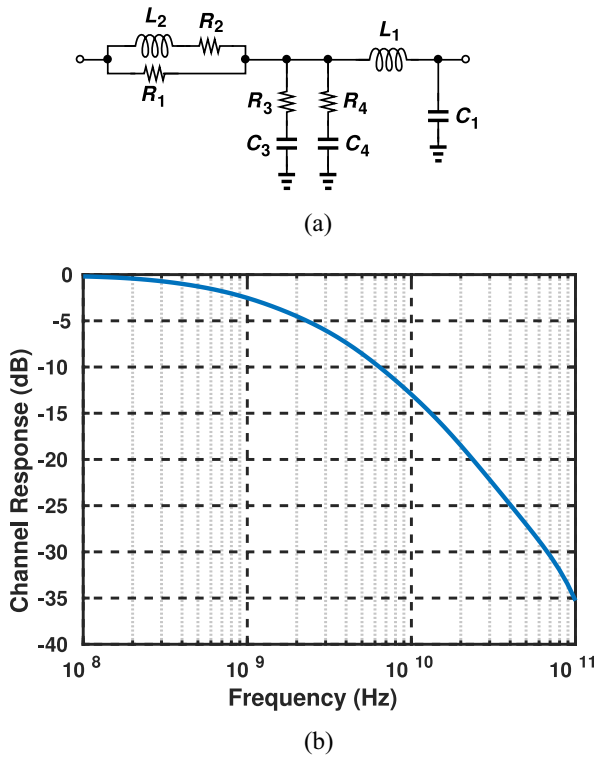


FIGURE 1. (a) One section of a scalable channel model. (b) Loss profile for 12 cascaded stages.

that the characteristic impedance, $Z_0 = \sqrt{L_1/C_1} = 50 \Omega$, $R_1 = 5.55 \Omega$, $R_2 = 150 \text{ m}\Omega$, $L_2 = 468.9 \text{ pH}$, $R_3 = 2 \text{ k}\Omega$, $C_3 = 200 \text{ fF}$, $R_4 = 100 \Omega$, and $C_4 = 80 \text{ fF}$. The trace simulations in [13] suggest a reasonable agreement with this model. Additional RL and RC branches can be included so as to refine the model. Fig. 1(b) plots the magnitude response of a channel consisting of 12 such sections, displaying a loss of 21 dB at 28 GHz. In this paper, the term “loss” will refer to that at the Nyquist rate.

We also wish to study the effect of impedance discontinuities on the link performance. We observe that such a nonideality can lead to deep notches in the channel frequency response. As an example, consider the scenario depicted in Fig. 2(a), where Z_p denotes a parasitic impedance at some point along the channel, e.g., at a connector, but the link is otherwise ideal. Since the impedance seen to the right of node X is equal to Z_0 , we note that $Z_p || Z_0$ is transformed by the transmission line on the left to create Z_{in} . The impedance rotation by a length of L_1 can move $Z_p || Z_0$ to a high Z_{in} , thus lowering the power delivered by V_{in} to the line and causing a notch in the frequency response [Fig. 2(b)]. In the time domain, the data experiences reflection at node X, a benign effect if $R_S = Z_0$. In other words, even though the reflection is absorbed on the TX side, the removal of the signal energy by the discontinuity still demands compensation.

The frequency-domain view of the channel proves useful for the design of circuits such as continuous-time linear equalizers (CTLEs). For discrete-time structures, on the other hand, a time-domain perspective becomes necessary. For

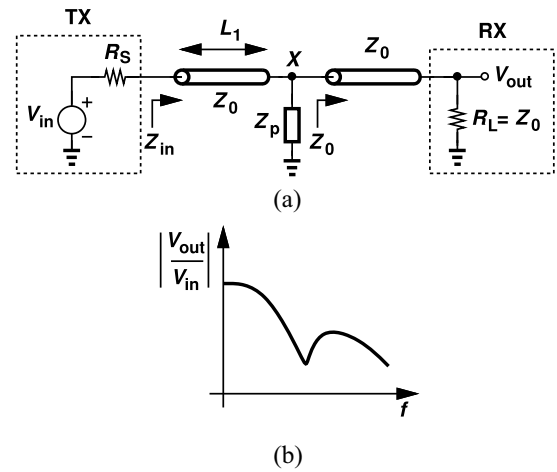


FIGURE 2. (a) Impedance discontinuity along a channel. (b) Resulting notch in the frequency response.




FIGURE 3. Impulse response of a lossy channel.

example, decision-feedback equalizers (DFEs) are designed according to the impulse response of the channel. Plotted in Fig. 3 is such a response, where T_B denotes the unit interval (UI), i.e., the bit or symbol period. The precursor at $-T_B$ and the post-cursors at $T_B, 2T_B$, etc., introduce intersymbol interference (ISI).

B. RECEIVER ARCHITECTURES

In the past decade, two general RX architectures have become common [1], [2], [3], [4], [5], [6], [7], [8], [9]. In “analog” receivers, equalization and clock and data recovery (CDR) occur in the analog domain. Fig. 4(a) illustrates this approach, which is better suited to NRZ data. A CTLE provides some high-frequency boost so as to partially compensate for the channel, and the result is applied to a DFE for further equalization. In addition, a CDR circuit senses the data and generates a clock with proper frequency and phase values for driving the DFE and the data demultiplexer (DMUX). Even though this architecture incorporates latches in the DFE, the CDR, and the DMUX, it is still considered an analog solution as most of its building blocks are crafted by analog designers.

The second architecture employs an analog-to-digital converter (ADC) and delegates some of the functions to the digital domain [Fig. 4(b)]. Called “ADC-based” receivers, such systems are suited to PAM4 data—especially for channel losses (CLs) greater than 20 dB. They do incorporate a CTLE in the front end so as to provide a boost of 10–20 dB, thus relaxing the ADC resolution to some extent. The ADC output drives a digital processor performing equalization and

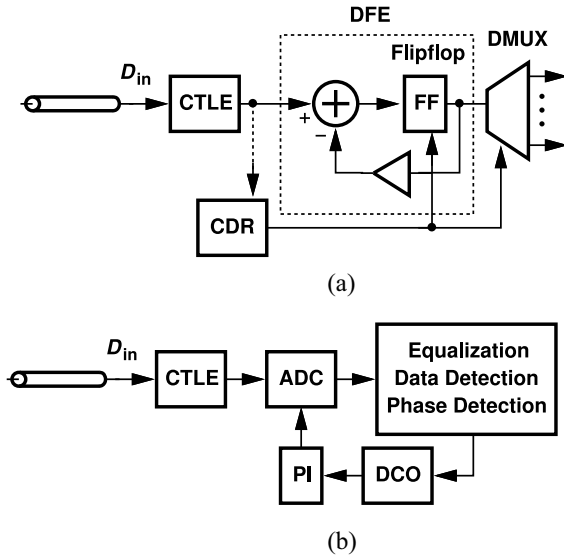


FIGURE 4. (a) Analog and (b) ADC-based receiver architectures.

data detection. The result also drives a CDR loop containing a phase detector (PD), a digitally controlled oscillator (DCO), and a phase interpolator (PI), which delivers the ADC's sampling clock(s). This RX architecture consumes substantial power in the ADC, the digital processor, and the clock generation and distribution network.

This paper focuses on analog NRZ receivers. For extra-short-reach or medium-reach links (with a CL of less than 20 dB), this architecture draws markedly less power, an important advantage because a given system contains many more such links than long-reach channels.

C. CHOICE OF CIRCUIT TOPOLOGIES

The analog and mixed-signal processing required in high-speed receivers can be realized by means of current-mode differential and regenerative pairs, but at the cost of significant static power consumption. For most of the operations beyond the CTLE, it is possible to employ “charge steering” [14].

Depicted in Fig. 5(a), a basic charge-steering differential stage replaces the tail current source with a “charge source” consisting of C_T , S_1 , and S_2 , and also the load resistors with precharge switches S_3 and S_4 . The output nodes are first tied to V_{DD} while C_T is discharged. Next, X and Y are released, and C_T switches into the tail node. The charge then flows from M_1 , M_2 , and their drain capacitances, amplifying the input and ceasing when V_P reaches about one threshold below the input common-mode (CM) level. The circuit can serve as an amplifier and/or a latch. A key difference between charge-steering and integrating stages, e.g., that shown in Fig. 5(b), is that, by design, the former does not allow V_X and V_Y , and hence $V_X - V_Y$, to collapse to zero whereas the latter does. Thus, the timing margins are more relaxed for charge steering. Moreover, this style can operate across a much wider speed range with no adjustment.

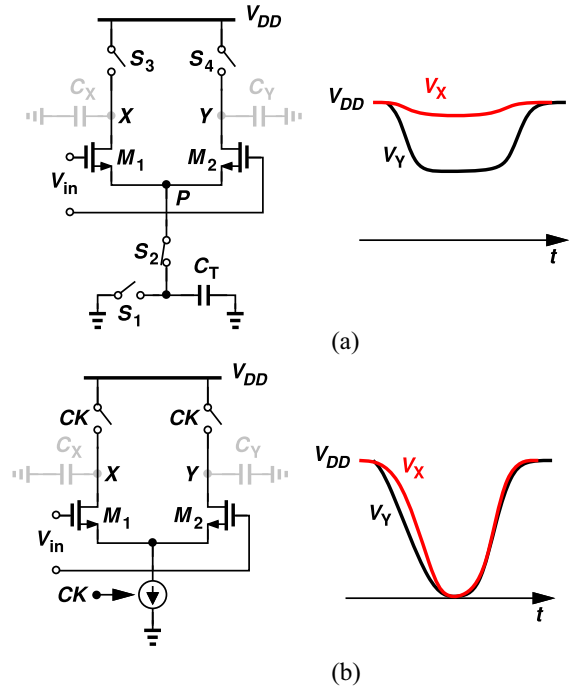


FIGURE 5. (a) Basic charge-steering stage and (b) its integrating counterpart.

Charge steering has been used in a multitude of RX and TX designs to save power [11], [12], [13], [14], [15].

D. LINEARITY REQUIREMENTS

The generation of NRZ data in transmitters does not dictate any linearity for their front end unless feedforward equalization is used. In NRZ receivers, on the other hand, some linearity is necessary before the data is sliced by the DFE because channel properties manifest themselves in the received signal amplitude. This issue proves important because we wish to amplify the input so as to maximize the eye height but must also be mindful of nonlinearity.

We investigate this point by considering the simple model shown in Fig. 6(a), where the RX front end is represented by a constant gain, k , and a static nonlinear stage [16]. Let us examine the impulse response of the entire chain, noting that $h_{in}(t)$ is that of the channel, which is then amplified by a factor of k . The result is subjected to compressive nonlinearity and exhibits a main cursor equal to h_0 and a first postcursor equal to h'_1 . In other words, nonlinearity equivalently raises the normalized postcursor level.

The nonlinearity is modeled by $y = \alpha_1 x + \alpha_3 x^3$ and thus an input 1-dB compression point $A_{1dB} = \sqrt{0.145|\alpha_1/\alpha_3|}$. Suppose $h_{in}(t)$ in Fig. 6(b) contains a main cursor equal to $\beta_m A_{1dB}$ and a first postcursor given by $\beta_1 A_{1dB}$. It can be shown that [16]

$$\frac{h'_1}{h_0} = \left(\frac{\beta_1}{\beta_m} \right)^3 \frac{(\beta_m/\beta_1)^2 - 0.145\beta_m^2}{1 - 0.145\beta_m^2}. \quad (1)$$

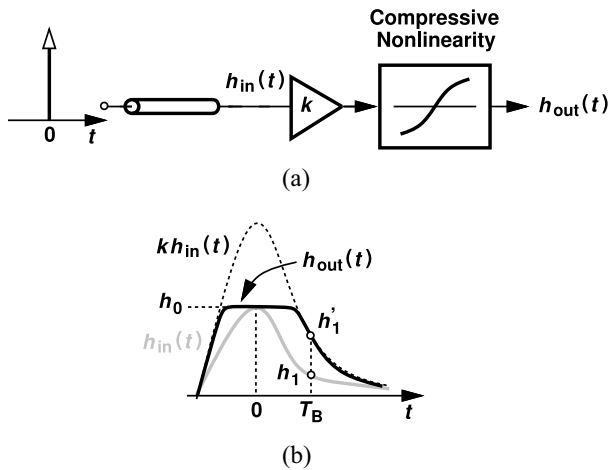


FIGURE 6. (a) Receiver containing nonlinearity and (b) effect of nonlinearity on postcursors.

As the front-end gain and hence β_m increases, h'_1/h_0 exceeds the input ratio, β_1/β_m . According to the findings in [16], this effect manifests itself if β_m reaches $1.5A_{1\text{dB}}$.

E. CHOICE OF CLOCK RATE

The simplest, most compact receivers operate with a full-rate clock, i.e., one whose frequency is equal to the input data rate. However, the generation and distribution of clocks at high speeds present formidable challenges. For this reason, we opt for half-rate or quarter-rate architectures—at the cost of doubling or quadrupling the hardware, respectively. An immediate consequence is that the CTLE in Fig. 4(a) now sees a greater load capacitance. As a compromise, we select half-rate clocking in the front end.

Half-rate clocking also becomes a natural choice in transceivers where the TX employs such a clock for its last multiplexer stage and the RX utilizes this clock along with phase interpolation to implement the CDR loop.

III. CTLE DESIGN

The CTLE in Fig. 4(a) must provide a high boost factor so as to 1) increase the eye opening at the DFE summing junction and 2) deliver a sufficient swing to the CDR, thus ensuring an adequate PD gain, lock range, and loop bandwidth (BW). We begin with the basic stage shown in Fig. 7(a), and note that the output pole, ω_0 , should preferably lie above $\omega_p = 1/(R_S C_S)$ [Fig. 7(b)], allowing the circuit to provide its maximum boost factor, $A_2/A_1 = 1 + g_m R_S/2$. In fact, ω_0 must exceed approximately $2.5\omega_p$ [17], a daunting challenge at high speeds that dictate the use of inductive peaking.

The design of the basic CTLE stage entails a tradeoff between the low-frequency gain, $g_m R_D/(1 + g_m R_S/2)$ (also called the “dc” gain), and the boost factor. For the output eye depicted in Fig. 7(c), a greater R_S reduces the outer height, H_1 , while raising the inner height, H_2 . An optimum can therefore be achieved for the latter as dictated by the

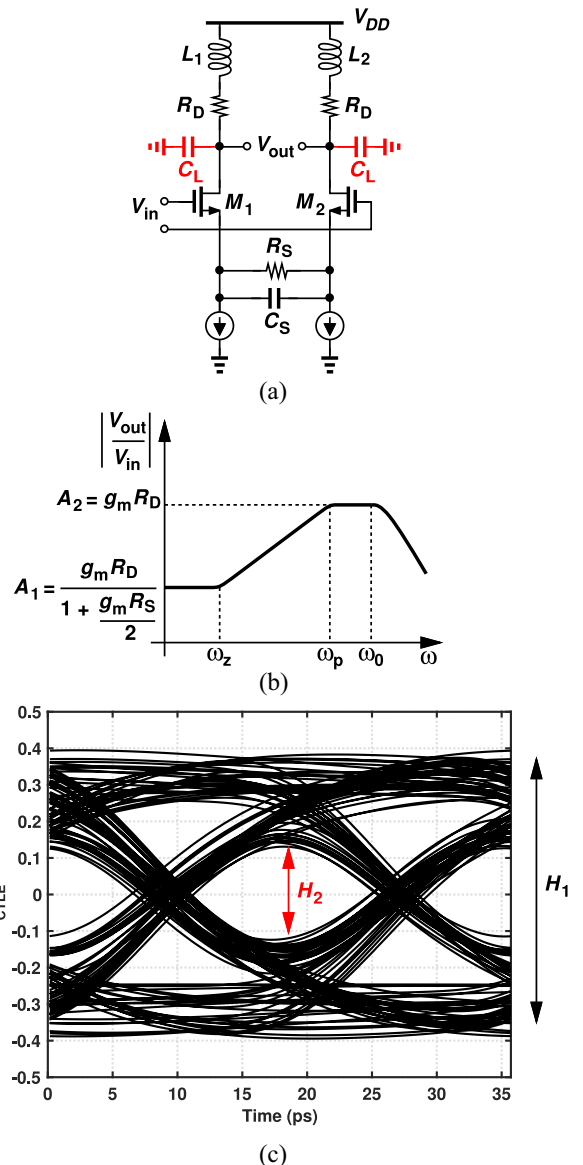


FIGURE 7. (a) Basic CTLE stage, (b) its frequency response, and (c) eye diagram showing inner and outer heights.

channel. We typically target a low-frequency gain of around 0 dB, thereby facing a boost factor bound of about 6 dB per stage due to the limited voltage headroom.

For higher boost factors, we cascade multiple CTLE stages, bearing in mind the proportional rise in the power consumption and the *reduction* in the bandwidth. For n identical stages, we have [18]

$$\text{BW}_{\text{tot}} = \text{BW}_0 \sqrt[m]{2^{1/n} - 1} \quad (2)$$

where BW_0 denotes the bandwidth of one stage and $m = 4$ for second-order stages. A cascade of two thus suffers from a 20% bandwidth shrinkage, i.e., ω_0 in Fig. 7(b) falls by this amount. For these reasons, typical front-end designs, comprising a CTLE and possibly a variable-gain amplifier, contain no more than three stages.

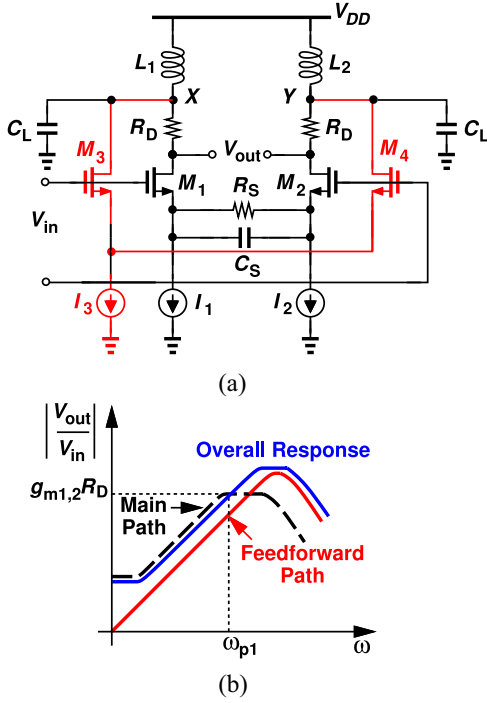


FIGURE 8. (a) CTLE using feedforward and (b) its frequency response.

The boost factor limitations outlined above call for additional high-frequency equalization techniques. We propose the concept of “feedforward” in this regard [12]. Illustrated in Fig. 8(a), the idea is to create a high-pass branch that contributes boost with negligible voltage headroom consumption. Transistors M_3 and M_4 and inductors L_1 and L_2 play such a role. The overall response is quantified as

$$\frac{V_{out}}{V_{in}} = -\frac{g_{m1,2}(R_D + L_D s)}{1 + g_{m1,2}\left(\frac{R_S}{2} \parallel \frac{1}{2C_S}\right)} - g_{m3,4}L_D s \quad (3)$$

where $L_1 = L_2 = L_D$ and the capacitances at the drains are neglected for now. The second term on the right-hand side represents the zero created by feedforward.

At high frequencies, source degeneration in Fig. 8(a) vanishes and the fraction on the right-hand side of (3) approaches $-g_{m1,2}(R_D + L_D s)$, yielding $V_{out}/V_{in} \approx -g_{m1,2}R_D - (g_{m1,2} + g_{m3,4})L_D s$. This implies that feedforward raises the apparent value of L_D and could be simply avoided by making L_D larger. The key point, however, is that C_L constrains the value of L_D if the output pole must lie above the Nyquist frequency. Thus, feedforward provides greater flexibility in shaping the frequency response.

We now consider the capacitances at the drains in Fig. 8(a) and sketch the responses created by the two paths. As shown in Fig. 8(b), the feedforward path is designed such that it dominates as the main path’s response reaches a plateau at ω_{p1} . The feedforward path should take over for $\omega > \omega_{p1} = (1 + g_{m1,2}R_S/2)/(R_S C_S)$; i.e., we must have $g_{m3,4}L_D \omega_{p1} <$




FIGURE 9. Extensive use of feedforward in a CTLE.

$g_{m1,2}R_D$ and hence

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

The advantages of feedforward become more pronounced if it is applied to both stages of a CTLE. As illustrated in Fig. 9, we exploit all three possible feedforward paths. The stage consisting of G_{m1} , G_{mf1} , and its RL load is identical to the circuit shown in Fig. 8(a), and so is the stage formed by G_{m2} , G_{mf2} , and its RL load. The values of G_{mf1} and G_{mf2} follow (4). The path consisting of G_{mf3} and L_{D2} manifests itself as the rest of the circuit approaches a flat response.

The performance of CTLEs must be studied in both frequency and time domains. Owing to the significant effect of layout parasitics, we report simulation results for only extracted circuits. The inductors are modeled by RLC networks obtained from Cadence’s EMX tool. We also include the input capacitances of the stages fed by the CTLE, namely, the CDR and the DFE. In the frequency domain, we perform two tests and study 1) the stand-alone CTLE and 2) the channel-CTLE cascade. Fig. 10(a) plots the proposed CTLE response as feedforward paths are added to the circuit. We observe that feedforward increases the boost factor by about 7 dB but it also lowers the corresponding frequency. Whether or not this result is acceptable is determined by additional tests. As depicted in Fig. 10(b), we cascade the channel profile of Fig. 1(b) with the CTLE. Notably, the overall response becomes flatter as feedforward branches are inserted, but the 3-dB bandwidth decreases to some extent. The ultimate test examines the eye diagram at the summing junction of the DFE with and without these branches. As explained in Section V-C, the three paths increase the eye height from 55 to 160 mV and the eye width from 18.5 to 20.5 ps.

IV. DISCRETE-TIME LINEAR EQUALIZATION

The notion of boosting high-frequency components can be pursued in the time domain as well. As illustrated in Fig. 11(a), a pulse experiencing the channel’s loss is broadened and introduces ISI at $t = T_B = 1$ UI. If this pulse is shifted by 1 UI, scaled by a factor of α , and negated,

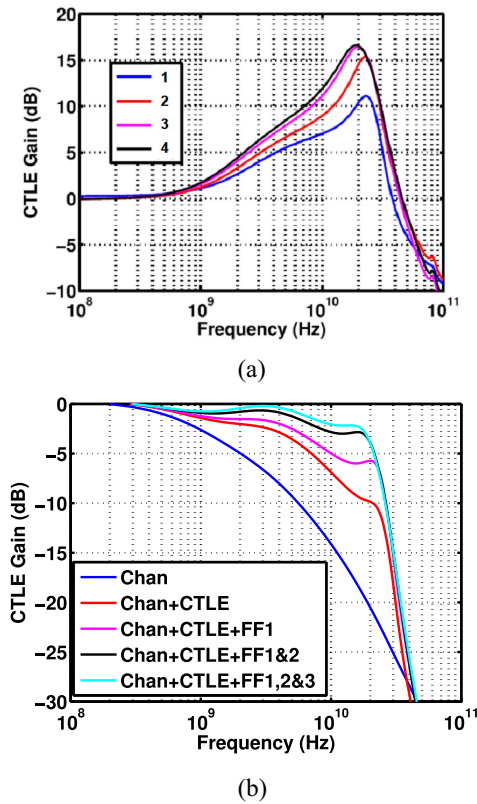


FIGURE 10. (a) CTLE frequency response for different configurations: (1) no feedforward, (2) with G_{mr1} , (3) with G_{mr1} and G_{mr2} , and (4) with all feedforward paths, and (b) corresponding responses for the channel-CTLE cascade.

it leads to its broadened counterpart at the output. Thus, $p(t) - \alpha p(t - T_B)$ produces less ISI. Implementing the operation as shown in Fig. 11(b), we write $Y = (1 - \alpha z^{-1})X$ and recognize that this “feedforward equalizer” (FFE) yields

$$Y = (1 - \alpha)X + \alpha(1 - z^{-1})X \quad (5)$$

where $0 < \alpha < 1$. The input is therefore subjected to two effects.

- 1) It is scaled by a factor of $1 - \alpha$, suffering from attenuation and displaying *smaller* low-frequency swings. This can be seen by applying a long sequence of ONEs and noting that they settle to a smaller amplitude [Fig. 11(c)].
- 2) The input is differentiated and scaled by a factor of α , thereby benefiting from high-frequency amplification. The boost factor is equal to $(1 + \alpha)/(1 - \alpha)$. A greater α translates to both a higher “dc” loss and a larger boost factor.

In TX design, the unit delays necessary for FFE are readily realized by flipflops as the NRZ data can be processed nonlinearly before the final summation point. FFE can also be formed in the analog domain in receivers. We call such a circuit a “discrete-time linear equalizer” (DTLE) [11]. Unlike TX FFEs, however, RX DTLEs process dispersed data and must provide some linearity so as to preserve the channel profile information. That is, they cannot rely on




FIGURE 11. (a) Illustration of FFE, (b) its implementation, and (c) its effect on long runs.

flipflops. Depicted in Fig. 12(a) is a DTLE example where the 1-UI delay is formed by a two-stage passive sampler. If $C_A \gg C_B$, the circuit delays $x(t)$ and scales it by a factor of α , but α itself can be realized by ratioing C_B with respect to C_A .

As explained in Section II-E, we prefer half-rate operation so as to ease the generation and distribution of clocks. This points to the topology shown in Fig. 12(b) [11], where both DMUX₁ and the DTLE are driven by a half-rate clock, $CK_{1/2}$. The odd and even data produced by DMUX₁ are delayed by 1 UI, scaled, and injected into the DFE’s summing junctions.

We make two remarks. First, DMUX₁ in Fig. 12(b) must perform sampling and can thus be merged with the first stage of the DTLE. This leads to the implementation shown in Fig. 12(c), where two-stage sampling is performed in the odd path by S_3 , S_4 , and the charge-steering stage, M_1 - M_2 , which injects the result into the DFE summing junction. In addition, the charge-steering regenerative pair consisting of M_3 and M_4 provides a gain of 6 dB. The nonlinearity introduced by this pair is studied in [11].

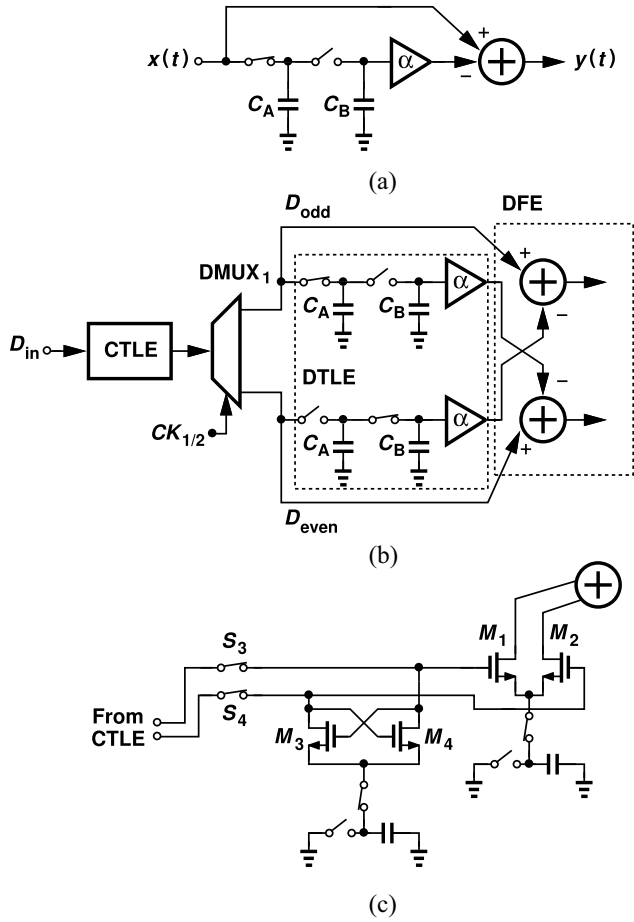


FIGURE 12. (a) Discrete-time linear equalization, (b) half-rate RX using DTLEs (C_A in odd branch tracks while DMUX₁ produces D_{odd}), and (c) charge-steering realization.

Second, the DTLE transfer function emerges as

$$H(z) = 1 - \alpha \frac{\frac{C_A}{C_A + C_B} z^{-1}}{1 - \frac{C_B}{C_A + C_B} z^{-2}}. \quad (6)$$

That is, if C_B is not much less than C_A , then the circuit also displays an infinite impulse response (IIR) tap equal to $C_B/(C_A + C_B)$.

V. DFE DESIGN

DFE architectures have been studied extensively. For most, the loop around the first tap must “close” in 1 UI regardless of the clock rate/data rate ratio. (In “unrolled” or “speculative” topologies, a loop consisting of a multiplexer still dictates a 1-UI timing budget [19].)

A. EYE OPENING CONSIDERATIONS

The eye height observed at the DFE summing junction must be large enough to satisfy the target bit error rate (BER), e.g., 10^{-12} . As shown in Fig. 13, five imperfections must be discounted from this height. These include 1) V_{OS1} : the CTLE and summer dc offsets; 2) V_{OS2} : the flipflop (FF) input-referred offset; 3) V_{n1} : the CTLE and summer noise;




FIGURE 13. Five sources of error in a DFE.

4) V_{n2} : the FF input-referred noise; and 5) V_{sen} : the FF sensitivity. We define V_{sen} as the input difference that allows the FF output to reach roughly 90% of its full swing in 1 UI [19] so that the first tap, h_1 completely switches.² If an eye monitor is available in the system, then V_{OS1} and V_{OS2} can be canceled. The BER is expressed as

$$\text{BER} \approx \frac{1}{2} Q \left(\frac{V_{pp}/2 - V_{OS} - V_{sen}}{\sqrt{V_n^2}} \right) \quad (7)$$

where Q denotes the error function, V_{pp} denotes the differential eye opening, V_{OS} denotes the total offset (with or without cancellation), and $\sqrt{V_n^2}$ denotes the total rms noise referred to the summing junction. An error rate of 10^{-12} demands that the argument of the Q function exceed 7.

In the absence of an eye monitor, V_{OS} in (7) must remain sufficiently small by proper design. For example, with an eye opening of 200 mV_{pp} and a total noise of $5 \text{ mV}_{\text{rms}}$, the offset must not exceed 65 mV (if V_{sen} is neglected). In practice, we would confine the 3σ offset to about 30 mV to leave a margin for the sensitivity and other imperfections.

The horizontal eye opening determines how much clock jitter and phase offset the equalizer can tolerate. The acceptable eye width depends, to some extent, upon the height: the greater the latter, the more the clock phase can depart from the center. This relationship is formulated in [13].

B. PROPOSED DFE TOPOLOGIES

A number of circuit and architecture techniques can improve the performance of high-speed DFEs. We begin by applying the concept of charge steering to summation and latching in a half-rate/quarter-rate environment. Consider the topology shown in Fig. 14(a), where half-rate data streams D_{odd} and D_{even} drive the summers and 2-to-1 DMUXs. The quarter-rate outputs of each DMUX are then multiplexed, scaled, and subtracted from the input data in the other path. The DMUX and MUX stages utilize the quadrature phases of the quarter-rate clock, generated by a $\div 2$ circuit that receives the half-rate clock. Illustrated in Fig. 14(b), the circuit implementation employs charge-steering differential pairs for the summer, the latch, and the MUX/tap 1 combination [16]. Moreover, the summer exploits RC degeneration so as to provide a few dB of boost.

2. Additionally, the FF kickback noise and hysteresis become problematic in some implementations.

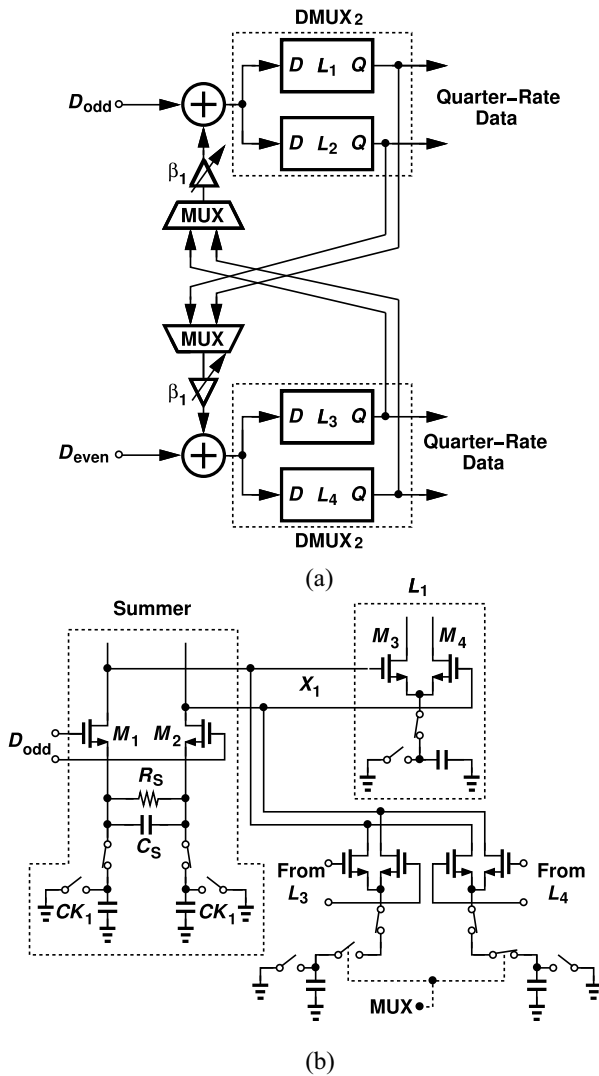


FIGURE 14. (a) Half-rate/quarter-rate DFE and (b) its charge-steering implementation.

A remarkable attribute of this architecture is its relaxed first-tap timing budget. In a conventional loop, we must have $t_{CK-Q} + t_{MUX} + t_{sum} + t_{setup} < 1$ UI, where the four terms, respectively, denote the flipflop clock-to- Q delay, the MUX delay, the summing node delay, and the FF setup time. In the charge-steering realization, on the other hand, we have $t_{CK-Q} < 1$ UI, where t_{CK-Q} is the delay from CK_1 to the output of the latch [16]. This constraint does not include a setup time because, in contrast to continuous-time current-mode latches, here the input data need not propagate to the precharged drain nodes of the MUX before this stage is clocked.

It is possible to reach a similar timing budget by injecting the feedback signal into the output of the first latch in the FF [22]. But this is not possible in the half-rate architecture of Fig. 14(a).

In addition to charge steering, we investigate greater interactions between the CTLE and the DFE to open the eye further. In contrast to conventional cascades, wherein the

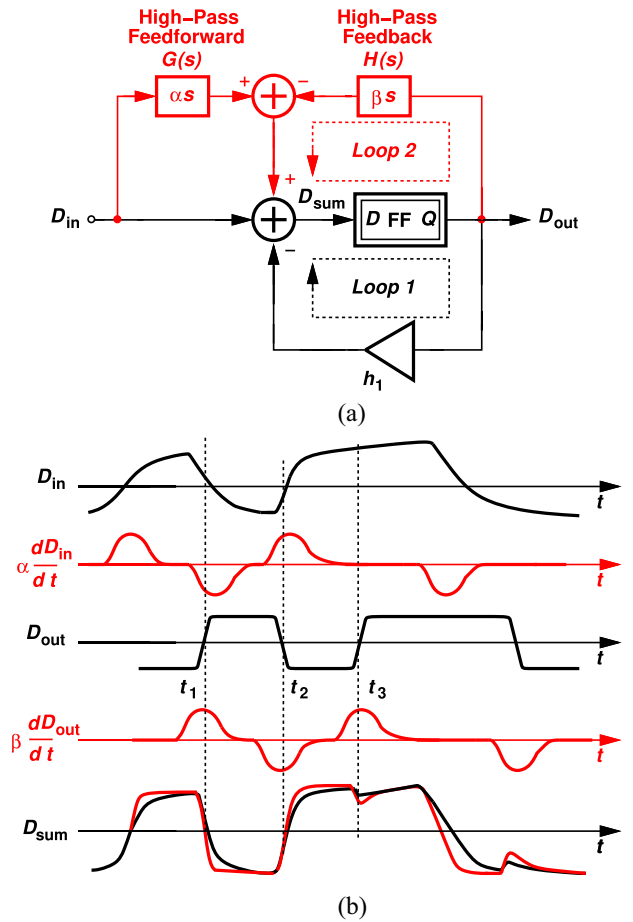


FIGURE 15. (a) Use of high-pass branches in a DFE and (b) associated waveforms.

CTLE drives the DFE unilaterally and only at one port, we can envision some feedforward and feedback paths between the two [12]. Depicted in Fig. 15(a) is a full-rate example: we allow a high-pass feedforward branch, $G(s)$, to inject the CTLE output into the summing junction. Furthermore, we create a high-pass feedback branch, $H(s)$, that returns the slicer output to D_{sum} (Loop 2). If $G(s) = \alpha s$ and $H(s) = \beta s$, we have

$$D_{sum}(n) = (1 + \alpha s)D_{in} - (h_1 + \beta_0 s)D_{out}(n - 1). \quad (8)$$

The high-frequency boost thus imparted to D_{in} and D_{out} improves the performance, a point that can be verified in the time domain as well. From the waveforms shown in Fig. 15(b), we observe that $\alpha dD_{in}/dt$ and $\beta dD_{out}/dt$ pulsate only on the data edges. Upon adding these derivatives to the summer output, we note that the rise and fall times are shortened. If two consecutive bits are the same, D_{sum} exhibits a kink due to $\beta dD_{out}/dt$ (e.g., at $t = t_3$), a benign effect as the kink occurs at bit boundaries.

The proposed feedforward and feedback techniques readily lend themselves to circuit implementation. As shown in Fig. 16(a), dD_{in}/dt is available at node P within the CTLE and travels through G_m stages to reach the summing junctions. For dD_{out}/dt , we first multiplex the quarter-rate outputs

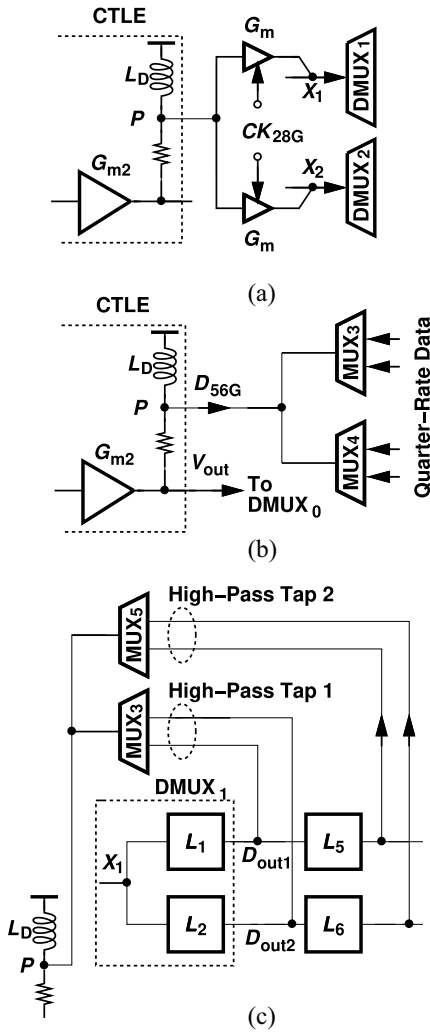


FIGURE 16. (a) Use of a high-pass signal within the CTLE. (b) Use of the same node for high-pass DFE feedback. (c) Addition of second tap.

of the latches so as to obtain full-rate data [Fig. 16(b)]. This topology can be viewed as a direct 4-to-1 MUX, except that it is driven by *overlapping* quadrature phases. It is shown that charge steering still delivers nonoverlapping charge packets to this output. We then inject the result into node P , granting L_D the task of differentiation. The strength of the injection, i.e., β , is defined by the amount of charge that each MUX branch draws.

The second DFE tap is accommodated by adding secondary latches to each quarter-rate arm, multiplexing their outputs, and injecting the results into each summing node and node P [Fig. 16(c)].

One may wonder how precisely one must control the timing alignment of the data that returns to node P in Fig. 16(c). In this work, no adjustment has been included as simulations reveal that this timing is no more critical than that of the main tap. If an eye monitor is present, one can adjust this path's delay for optimum performance.

In contrast to IIR DFEs [20], [21], the proposed method returns the shaped signal to the DFE input rather than to its

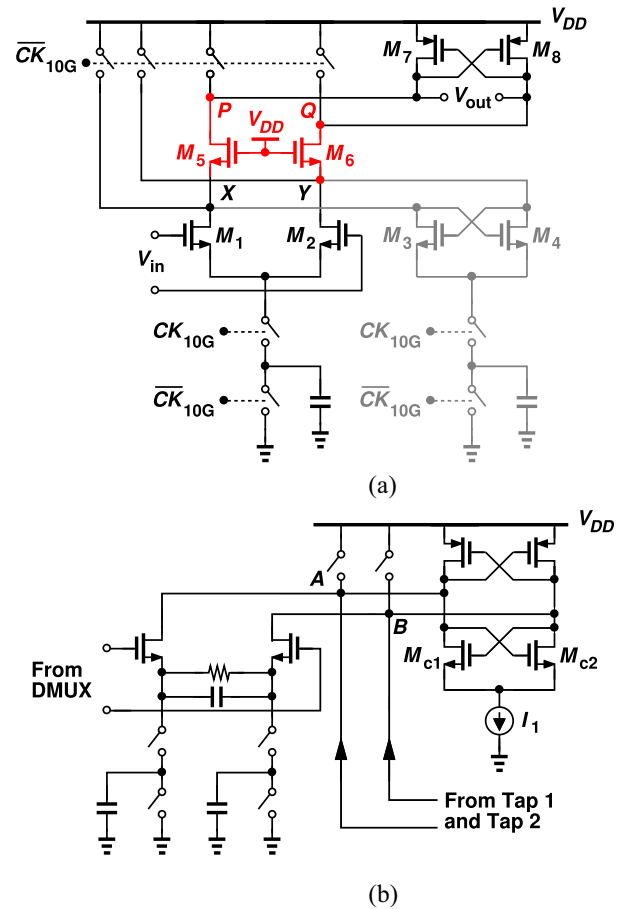


FIGURE 17. (a) Modified charge-steering latch and (b) improved summer.




FIGURE 18. 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).

summing junction. According to the foregoing analysis and simulations, this approach yields a greater eye opening.

C. REFINEMENTS

We incorporate additional circuit techniques to further improve the DFE's performance, striving to maximize the NRZ eye opening at its summing junctions. First, we modify the basic charge-steering latch of Fig. 5 as shown

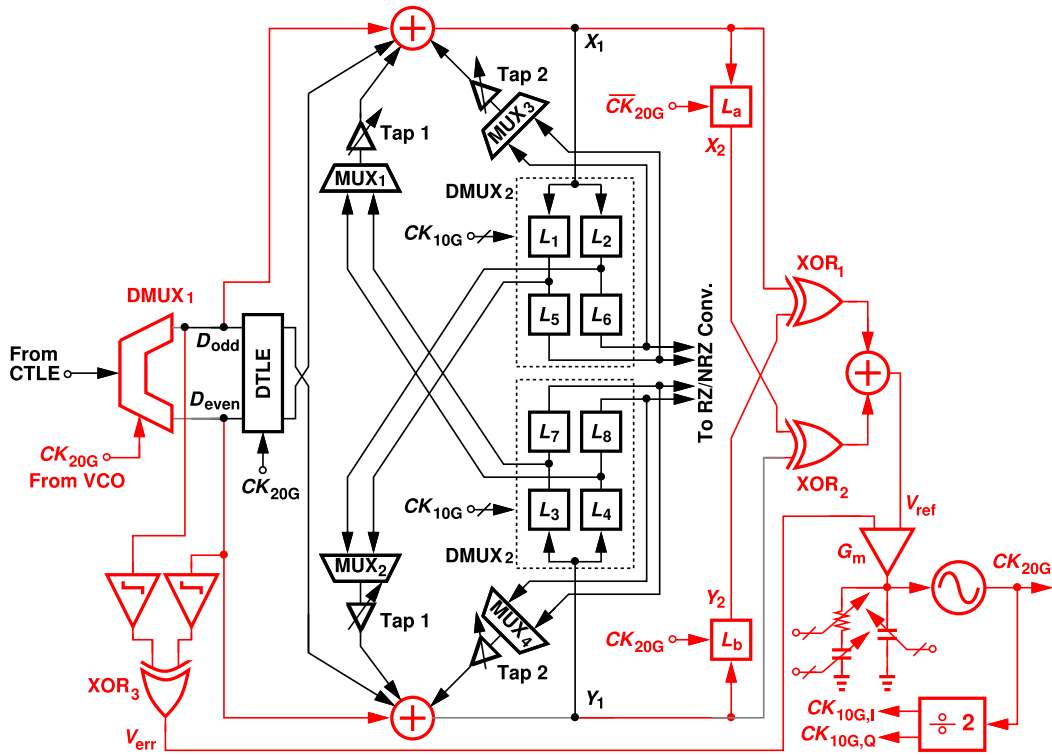


FIGURE 19. Overall architecture of 40-Gb/s RX.

in Fig. 17(a), where a cascode pair, M_5 – M_6 , and two cross-coupled pairs, M_3 – M_4 and M_7 – M_8 , boost the output voltage swings [11]. These transistors play the following roles: the first pair isolates X and Y from the large capacitance at P and Q , raising the voltage gain from V_{in} to these nodes; the second pair also increases this gain by means of regeneration; the third pair restores the high level at P or Q to V_{DD} , avoiding the CM drop observed in Fig. 5(a).

The second method relates to the DFE summing node itself. As illustrated in Fig. 17(b), we attach two cross-coupled pairs to this interface, thus increasing the eye height by 50% [12]. The continuous-time CM drop caused by I_1 at A and B is less than 20 mV in the 18-ps evaluation mode of the 56-Gb/s RX.

We quantify the improvements afforded by some of our proposed techniques for the 56-Gb/s RX in the presence of a CL of 25 dB. Fig. 18 illustrates the incremental improvements due to each concept. The eye width increases from 18.5 to 25 ps, and the eye height from 55 to 200 mV.

VI. 40-Gb/s AND 56-Gb/s RECEIVERS

The 40-Gb/s and 56-Gb/s NRZ RX examples reported here operate with a CL of 19–25 dB at the Nyquist frequency. The former’s architecture is shown in Fig. 19 [11]. A single CTLE stage drives DMUX₁, the DTLE, and the DFE, which consists of two summers, latches L_1 – L_8 , and MUX₁–MUX₄. The retimed and demultiplexed return-to-zero (RZ) data is converted to NRZ as described in [14].

The CDR utilizes the signals processed by DMUX₁ and the DFE to reduce the number of latches that it requires [11]. Specifically, XOR₃ measures the phase difference between D_{odd} and D_{even} , while XOR₁ and XOR₂ generate a constant-width pulse on V_{ref} for each data transition. The resulting difference, $V_{err} - V_{ref}$, uniquely represents the phase error regardless of the data pattern.

The 56-Gb/s RX is depicted in Fig. 20 [12]. (For simplicity, the second DFE tap is not shown.) In this case, the higher speed is accommodated by driving the CDR from node Q in the CTLE so that C_{CDR} negligibly affects the signal path’s bandwidth. The data presented to the CDR thus displays a high-pass spectrum, but it still allows locking [12].

The half-rate PD requires quadrature clocks at 28 GHz, a condition fulfilled by simply delaying the output of a differential LC oscillator by a self-biased inverter. It is shown that this stage’s delay variability does not affect the PD gain significantly [12].

VII. EXPERIMENTAL RESULTS

This section presents the measured results for the 40-Gb/s and 56-Gb/s NRZ receivers. The prototypes have been mounted directly on printed-circuit boards and tested on a high-speed probe station. Unless otherwise stated, all measurements are carried out with a 1-V supply at the full data rate and with a pseudo-random bit sequence (PRBS) pattern of $2^7 - 1$. Fig. 21 depicts a test setup example for characterizing receivers. A BER tester (BERT) generates NRZ data, which is then subjected to a lossy channel such as M8049A.

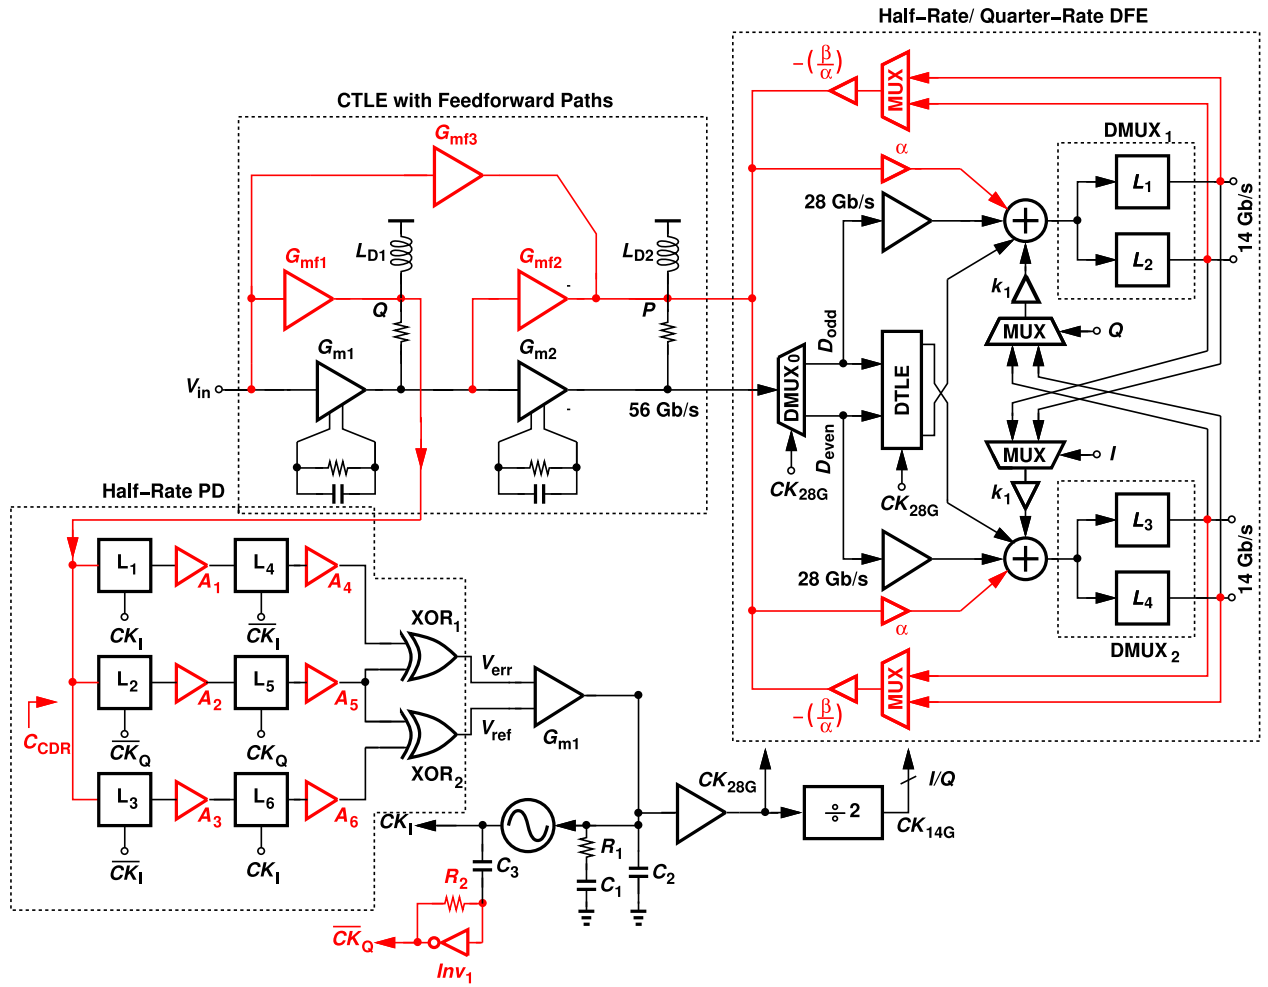


FIGURE 20. Overall architecture of 56-Gb/s RX.

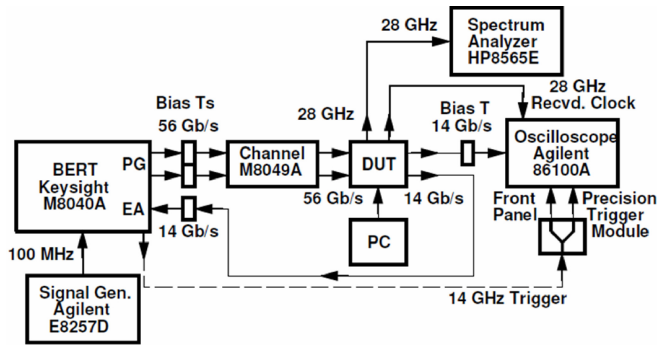


FIGURE 21. Test setup example.

The result drives the device under test (DUT) and the output is captured by an oscilloscope. The recovered clock too is monitored on a spectrum analyzer.

A. 40-Gb/s RX

Realized in TSMC's 45-nm technology, the 40-Gb/s RX die is shown in Fig. 22 and occupies an active area of about $110 \mu\text{m} \times 175 \mu\text{m}$. Another version accepting an external


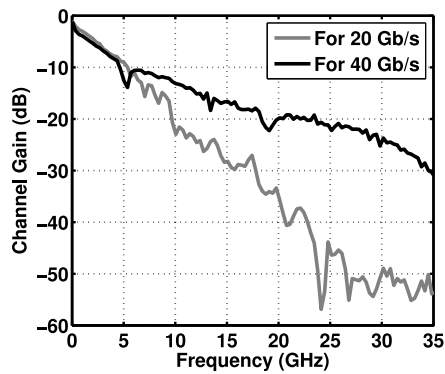
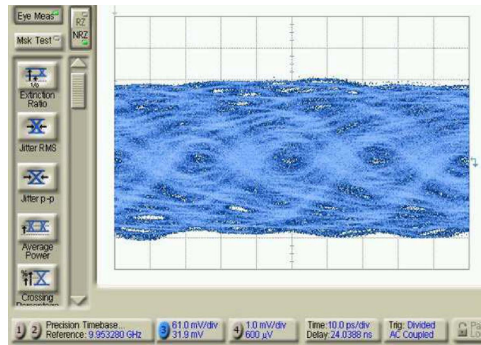


FIGURE 22. 40-Gb/s RX die photograph.

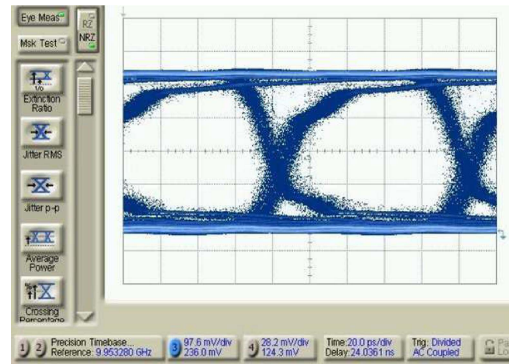
clock has also been fabricated to permit the characterization of the equalizer. We first employ a channel having the black profile shown in Fig. 23(a) and producing the eye in Fig. 23(b).



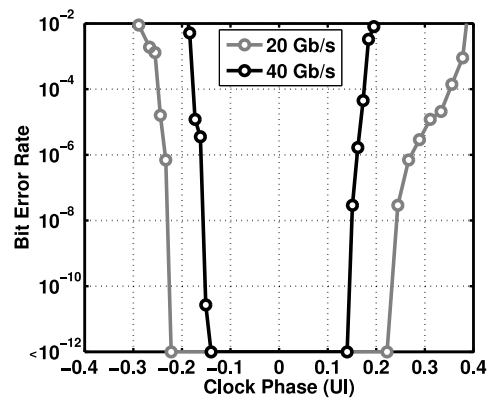
(a)



(b)



(a)



(b)

FIGURE 23. (a) Two channel profiles used in measurements and (b) received eye diagram (with the gray profile).

We begin with the RX path measurements using an external clock. The output data at 10 Gb/s is depicted in Fig. 24(a) and the equalizer bathtub curve in Fig. 24(b). The horizontal eye opening is 0.28 UI. Part of the eye closure arises from the PRBS generator's 8-ps_{rms} jitter. Also shown is the bathtub curve for an input data rate of 20 Gb/s and the gray loss profile in Fig. 23(a), demonstrating that charge-steering circuits can accommodate a wide range of frequencies.

The complete RX is characterized for jitter generation, transfer, and tolerance while it equalizes the dispersed data. The CDR bandwidth is set to 20 MHz unless otherwise stated. Fig. 25(a) and (b) plots the recovered clock spectrum and waveform, respectively. For phase noise measurements, the 20-GHz clock is divided by 2 off-chip, yielding the profile illustrated in Fig. 26. The integrated jitter amounts to 515 fs_{rms} from 100 Hz to 1 GHz.

Fig. 27 shows the measured jitter transfer and tolerance for different CDR bandwidths. The latter improves as the BW increases, reaching 0.45 UI_{pp} at 5 MHz with 19 dB of CL. (The maximum jitter amplitude of 20 UI is dictated by the equipment.)

Table 1 summarizes and compares the performance.

B. 56-Gb/s RX

This RX has been fabricated in TSMC's 28-nm technology. Fig. 28 shows the die with an active area of 250 $\mu\text{m} \times$

FIGURE 24. (a) RX output eye diagram at 10 Gb/s and (b) equalizer bathtub curves.

TABLE 1. Performance summary and comparison for 40-Gb/s RX.

Reference	Hsieh VLSI 2011	Chen JSSC Mar. 2012	Raghavan JSSC Dec. 2013	This Work
Data Rate (Gb/s)	40	40	40	40
Supply (V)	1.2 for DFE/CDR, 1.5 for CTLE	1.6	1	1
Channel Loss at Nyquist (dB)	23.5	19	>21	18.6
Bit Error Rate	10^{-12}	10^{-12}	10^{-12}	10^{-12}
Power (mW)	150	520	1050 [†]	14
Power Efficiency (pJ/bit)	3.75	13	26.25	0.35
Recovered Clock Jitter (ps)	6.8 pp	0.319 rms	-	0.515 rms
Jitter Tolerance	-	≈ 0.65 UI _{pp} at 10 MHz	0.95 UI _{pp} at 10 MHz [‡]	0.45 UI _{pp} at 5 MHz
Area (mm ²)	0.278	1.1475 [*]	3.9 [*]	0.019
Technology	65-nm CMOS	65-nm CMOS	40-nm CMOS	45-nm CMOS


^{*} Includes pads

[†] Includes SFI-5.2 TX; 350 mW for line-side RX


[‡] Measured for BER = 10^{-9}

275 μm . The tests are carried out with Keysight's boards, MS8049A-002 and MS8049A-003, which, along with a 30-in cable, provide the loss profiles plotted in Fig. 29. To these losses at 28 GHz, we add 1.7 dB to account for the probes and the interconnects.

We first report the RX performance while the CDR is disabled and an external 28-GHz clock is used. In this



(a)



(b)

FIGURE 25. (a) Measured recovered clock spectrum and (b) its waveform.

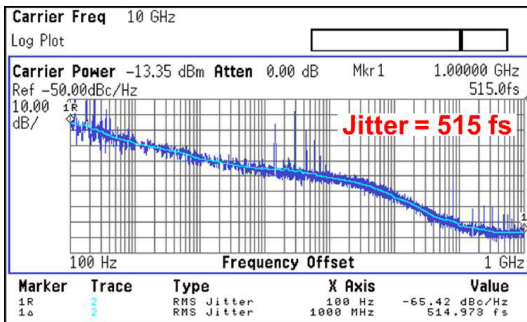


FIGURE 26. Measured phase noise of recovered clock.

measurement, Keysight’s M8040A BERT has the capability to emulate a 2-tap TX FFE in the data applied to the channel. Fig. 30 plots the bathtub curves for two cases: 1) for channel A, which has a loss of 25 dB, and no FFE and 2) for channel B, which has a loss of 30 dB, while the BERT implements an FFE function of the form $-0.2 + 0.8z^{-1}$. The horizontal eye openings are 0.4 and 0.33 UI, respectively.

We next present results with the CDR enabled. Shown in Fig. 31 are the outputs of channel A and the RX. The BER is less than 10^{-12} . Fig. 32 plots the recovered clock waveform and spectrum for a CDR noise-shaping bandwidth of 50 MHz. The phase noise profile of Fig. 33 reaches a 100-MHz offset, at which it is equal to -124.4 dBc/Hz. For greater offsets, we measure the phase noise directly from

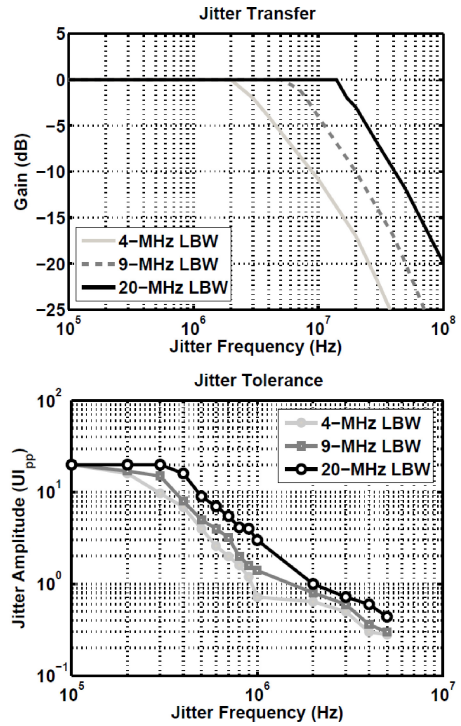


FIGURE 27. Measured jitter transfer and tolerance.

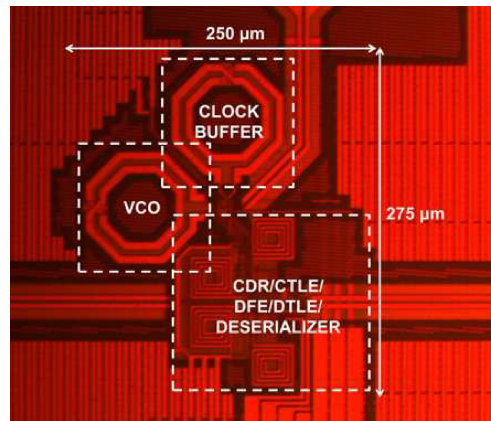


FIGURE 28. 56-Gb/s RX die photograph.

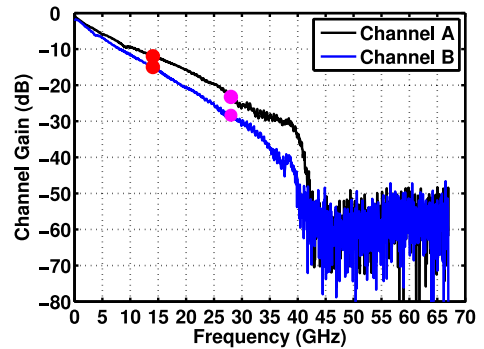

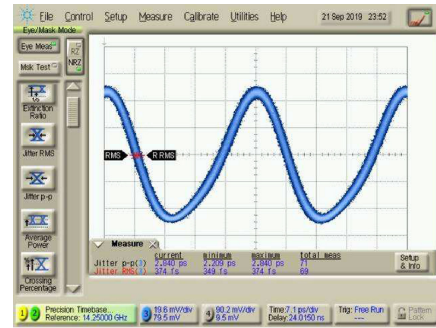



FIGURE 29. Measured CL profiles.

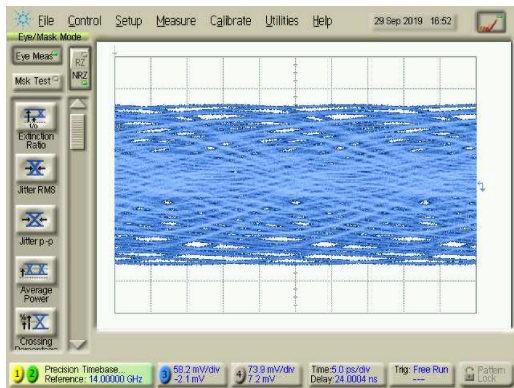
the spectrum, which falls to -128 dBc/Hz at 14-GHz offset. The integrated jitter from 100 Hz to 14 GHz amounts to 100 fs_{rms}.


FIGURE 30. Measured RX bathtub curves.


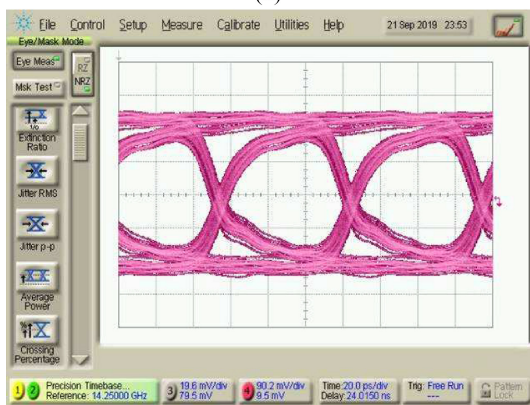
(a)



(b)

FIGURE 32. (a) Recovered clock waveform and (b) its spectrum.


(a)



(b)

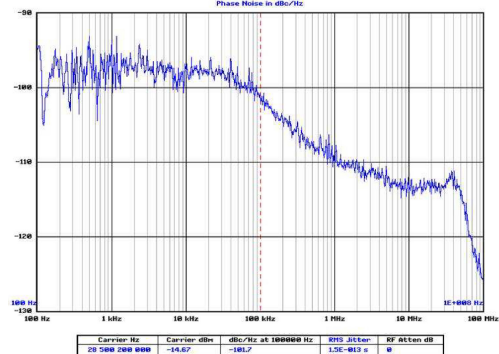
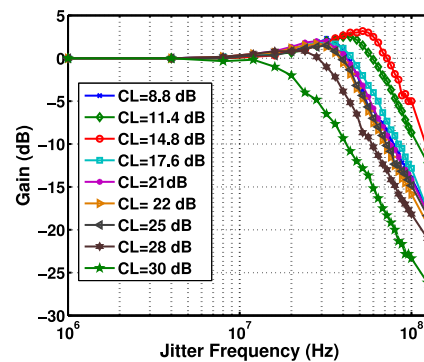
FIGURE 31. Eye diagrams (a) received from the channel and (b) at RX output.

FIGURE 33. Measured phase noise of recovered clock.

FIGURE 34. Measured jitter transfer.

Fig. 34 plots the CDR jitter transfer for different CLs, obtained by cascading different sections of Keysight's board and different cable lengths. For the 25-dB loss case, the 3-dB BW is around 55 MHz, consistent with the VCO noise-shaping BW observed in Fig. 32(b). The high-pass nature of the CDR input data leads to some peaking for low loss values, but it enables the CDR to achieve bandwidths as high as 25 MHz for a CL of 30 dB.

Fig. 35 plots the measured CDR jitter tolerance for a loss of 25 dB, yielding a value of 1.1 UI_{pp} at 5 MHz and exceeding the CEI-56G-VSR mask. Table 2 summarizes and compares the performance.

VIII. CONCLUSION

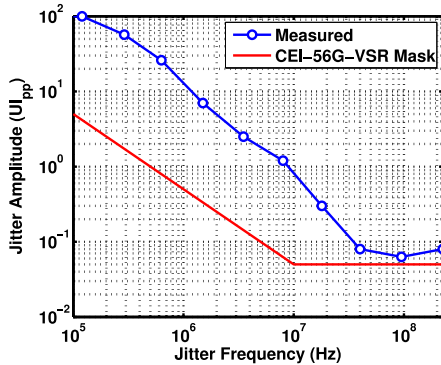
High-speed wireline receivers present a multitude of challenges, especially for greater CLs. This paper describes

TABLE 2. Performance summary and comparison for 56-Gb/s RX.

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	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. [§]	141.7	382	136	–	–	49.56
	Excl. ^{§§}	–	–	–	180	112	43.6
Power Eff. (pJ/bit)	Incl. [§]	2.53	6.82	2.26	–	–	0.88
	Excl. ^{§§}	–	–	–	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

§ Includes Clock Gen. §§ Excludes Clock Gen.


FIGURE 35. Measured jitter tolerance.

methods that improve the performance of CTLEs and DFEs and proposes concepts such as discrete-time linear equalization and charge steering. Collectively, these techniques lead to 40-Gb/s and 56-Gb/s receivers with low power consumption.

ACKNOWLEDGMENT

The author gratefully acknowledges the TSMC University Shuttle Program for chip fabrication.

REFERENCES

- [1] Y. Segal et al., “A 1.41pJ/b 224Gb/s PAM-4 SerDes receiver with 31dB loss compensation,” in *ISSCC Dig. Tech. Papers*, Feb. 2022, pp. 114–115.
- [2] T. Ali et al., “A 460mW 112Gbps DSP-based transceiver with 38dB loss compensation for next generation data centers in 7nm FinFET technology,” in *ISSCC Dig. Tech. Papers Slide Supplements*, Feb. 2020, pp. 118–120.
- [3] P. Upadhyaya et al., “A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with a configurable ADC in 16nm FinFET,” in *ISSCC Dig. Tech. Papers*, Feb. 2018, pp. 108–110.
- [4] T. Ali et al., “6.4 A 180mW 56Gb/s DSP-based transceiver for high-density I/Os in data center switches in 7nm FinFET technology,” in *ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 118–120.
- [5] 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 *ISSCC Dig. Tech. Papers*, Feb. 2020, pp. 116–118.
- [6] 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.
- [7] T. Shibusaki et al., “A 56-Gb/s receiver front-end with a CTLE and 1-tap DFE in 20-nm CMOS,” in *VLSI Circuits Symp. Dig.*, Jun. 2014, pp. 1–2.
- [8] 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.
- [9] A. Cevrero et al., “6.1 A 100Gb/s 1.1 pJ/b PAM-4 RX with dual-mode 1-tap PAM-4/3-tap NRZ speculative DFE in 14nm CMOS FinFET,” in *ISSCC Dig.*, Feb. 2019, pp. 112–113.
- [10] B. Razavi, “Design techniques for high-speed wireline transmitters,” *IEEE Open J. Solid-State Circuits Soc.*, vol. 1, pp. 53–66, 2021.
- [11] 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.
- [12] A. Atharav and B. Razavi, “A 56-Gb/s 50-mW NRZ receiver in 28-nm CMOS,” *IEEE J. Solid-State Circuits*, vol. 57, no. 1, pp. 54–67, Jan. 2022.
- [13] S. Gondi and B. Razavi, “Equalization and clock and data recovery techniques for 10-Gb/s CMOS serial links,” *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1999–2011, Sep. 2007.

- [14] 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.
- [15] Y. Chang, A. Manian, L. Kong, and B. Razavi, "An 80-Gb/s 40-mW wireline PAM4 transmitter," *IEEE J. Solid-State Circuits*, vol. 53, no. 8, pp. 2214–2226, Aug. 2018.
- [16] 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.
- [17] B. Razavi, "The design of an equalizer—Part I [the analog mind]," *IEEE Solid-State Circuits Mag.*, vol. 13, no. 4, pp. 7–160, 2021.
- [18] R. P. Jindal, "Gigahertz-band high-gain low-noise AGC amplifiers in fine-line NMOS," *IEEE J. Solid-State Circuits*, vol. 22, no. 4, pp. 512–521, Aug. 1987.
- [19] 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.
- [20] B. Kim, Y. Liu, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10-Gb/s compact low-power serial I/O with DFE-IIR equalization in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3526–3538, Dec. 2000.
- [21] O. Elhadidy and S. Palermo, "A 10-Gb/s 2-IIR-Tap DFE receiver with 35 dB loss compensation in 65-nm CMOS," in *Symp. VLSI Circuits Dig.*, Jun. 2013, pp. C272–C273.
- [22] Y. Lu and E. Alon, "Design techniques for a 66 Gb/s 46 mW 3-tap decision feedback equalizer," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3243–3257, Dec. 2013.