A 40-Gb/s 14-mW CMOS Wireline Receiver

Abishek Manian and Behzad Razavi, *Fellow, IEEE*

***Abstract*— A 40-Gb/s receiver includes a continuous-time linear equalizer, a discrete-time linear equalizer, a two-tap decision-feedback equalizer, a clock and data recovery circuit, and a one-to-four deserializer. Hardware minimization and charge steering techniques are extensively used to reduce the power consumption by a factor of ten. Fabricated in 45-nm CMOS technology, the receiver exhibits a bathtub curve opening of 0.28 UI with a recovered clock jitter of 0.5 ps_{rms}.**

***Index Terms*— Charge steering, clock and data recovery (CDR), decision feedback, discrete time, equalizer, linear equalizer, phase detector (PD).**

I. Introduction

TRANSMISSION of high-speed data over copper media has for some years targeted a power efficiency (PE) of 1 mW/Gb/s, striving for a reasonable chip power consumption as the number of input/output (I/O) pins continues to climb. Indeed, transceivers operating at a few gigabits per second have provided such a performance [1]–[4], but, for rates exceeding about 25 Gb/s, the PE is about 8 to 50 times higher [5]–[9]. For example, 40-Gb/s receivers draw 150 mW [10] to 1 W [9].

This paper introduces a 40-Gb/s receiver that achieves a tenfold reduction in power by means of new circuit and architecture techniques [11]. A byproduct of this effort is the small footprint of the receiver, about 110 μm $\times$ 175 μm, which makes it possible to realize a multi-lane system in a small area with short interconnects. Fabricated in 45-nm CMOS technology, the receiver compensates for a channel loss of 18.6 dB at Nyquist with a recovered clock jitter of 0.515 ps_{rms} and a bit error rate (BER) less than 10^{-12}.

II. Receiver Architecture

Fig. 1 conceptually shows the receiver architecture, with the overlapping boundaries denoting hardware sharing between the functions. The input signal D_{in} travels through a continuous-time linear equalizer (CTLE) and subsequently drives a clock and data recovery (CDR) circuit, a decision-feedback equalizer (DFE), and a "discrete-time linear equalizer" (DTLE) [12]. The retimed data are deserialized to four 10-Gb/s channels.

Manuscript received November 7, 2016; revised February 18, 2017 and April 11, 2017; accepted May 8, 2017. Date of publication June 6, 2017; date of current version August 22, 2017. This work was supported by Texas Instruments and Realtek Semiconductor. This paper was approved by Associate Editor Jack Kenney. *(Corresponding author: Abishek Manian.)*

A. Manian was with the Electrical Engineering Department, University of California at Los Angeles, Los Angeles, CA 90095 USA. He is currently with Texas Instruments Inc., Santa Clara, CA 95051 USA (e-mail: a-manian@ti.com).

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

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2017.2705913

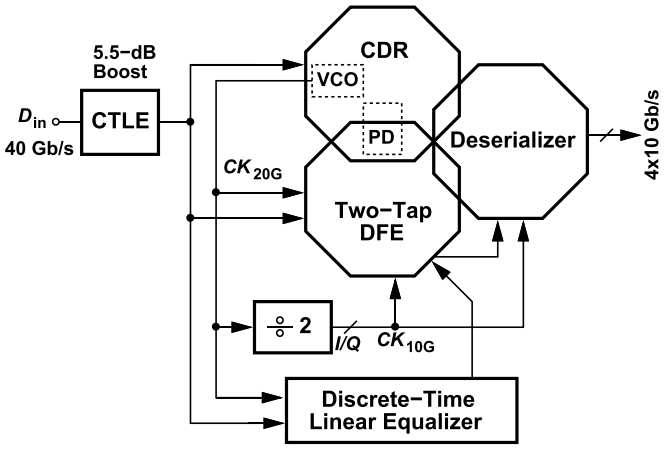

Fig. 1. Conceptual receiver architecture.

Our receiver development draws upon three principles to improve the performance. First, we uphold a "minimalist" approach, recognizing that, at these speeds, hardware means power and that every additional stage in the data path also reduces the bandwidth. We therefore limit the number of stages in the CTLE to one, and avoid quadrature oscillators in the CDR circuit, buffers in the data and clock paths, and phase interpolation circuits. The minimalist mentality also leads to sharing building blocks among different functions, thereby reducing both the power and the capacitance seen in the data and clock paths.

The receiver of Fig. 1 exemplifies how minimalist approach produces "growing" returns: the compact architecture contains no high-speed data interconnects longer than 25 μm, which in turn avoids the need for buffers, inductive peaking (except in the CTLE), and so on.

The second principle is to employ charge steering as a power-efficient logic design paradigm [13], [14]. We introduce several new latch topologies that extend the speed of charge steering, obviating the need for current-steering circuits in the entire receiver.

The third principle is the use of discrete-time linear equalization as a means of creating high-frequency boost with a PE much higher than that of CTLEs. The next section focuses on this principle.

Though attractive for compact efficient design, the foregoing principles pose other challenges; specifically, full-rate operation without clock buffers becomes very difficult, while half-rate operation doubles the load capacitance that the CDR and the DFE present to the CTLE.

III. Problem of CTLE

The minimum acceptable CTLE boost factor is determined by two requirements: 1) to open the eye to some extent for the CDR loop; an excessively small eye height translates to a

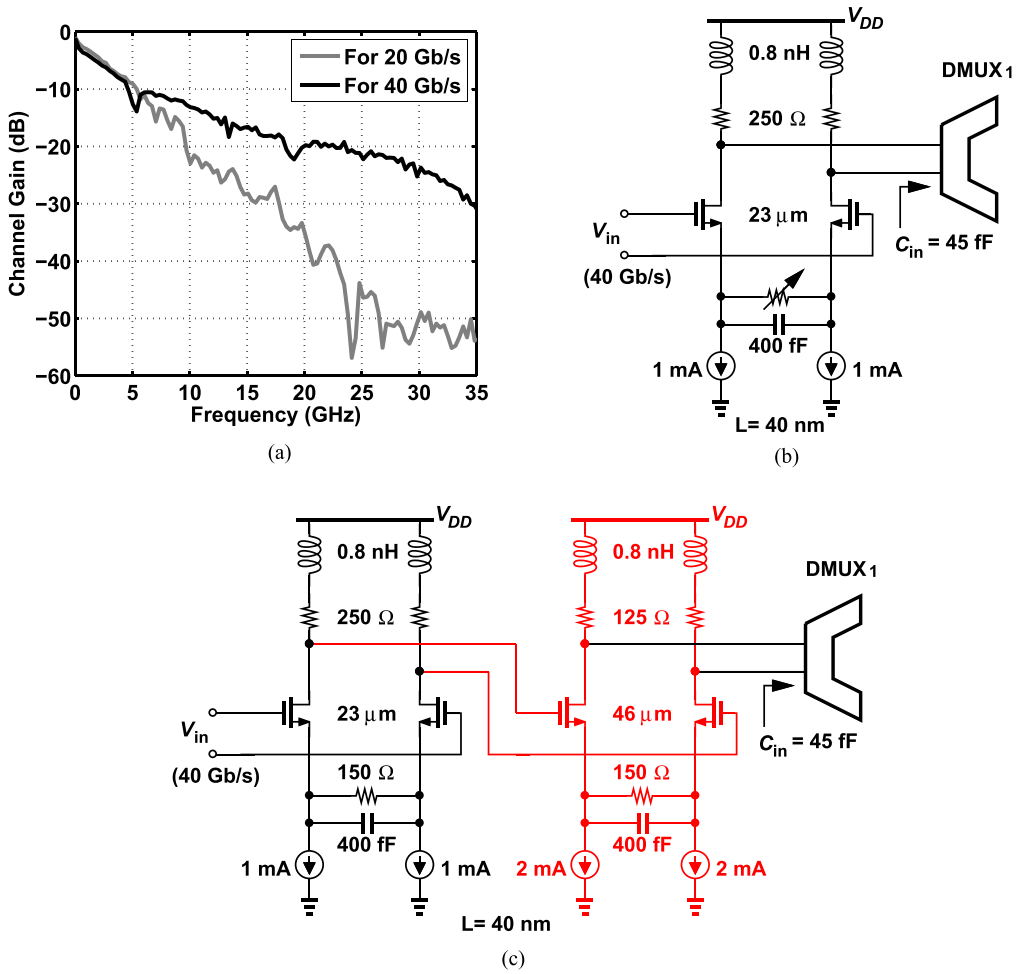

Fig. 2. (a) Measured channel frequency responses (excluding 1-dB insertion loss of probes). (b) One-stage CTLE. (c) Addition of a second stage to the CTLE.

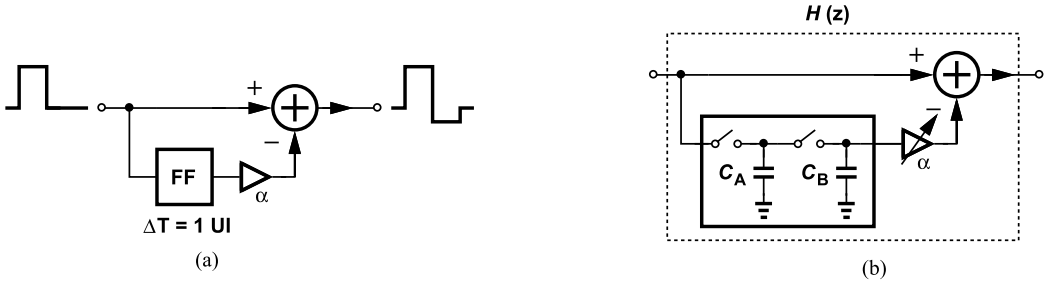

Fig. 3. (a) Discrete-time linear equalization in transmitters. (b) Discrete-time linear equalization in receivers using a master–slave sampling circuit.

low phase detector (PD) gain, limiting the CDR loop capture range and degrading its stability and 2) to ensure sufficient eye opening at the DFE summing junction to overwhelm noise and offsets referred to this port. While the latter condition can instead be met through the use of more DFE taps, the large number of high-speed flip-flops (FFs) necessary for these paths would draw significant power and present a heavy load to the recovered clock and the DFE summing junction.

To quantify the tradeoffs in CTLE design, consider the measured channel profile in Fig. 2(a) and the single boost stage in Fig. 2(b). This stage provides a maximum boost of 5.5 dB at 20 GHz and consumes 2 mW.[1] From SS 80 °C to FF 0 °C corners, the maximum CTLE boost varies from 5.4 to 6 dB, respectively, with a 1-dB variation in the dc gain. This stage is designed to handle single-ended input swings ranging from 200 to 400 mV_{pp} with the input common-mode (CM) level ranging from 0.6 to 0.8 V. If as shown in Fig. 2(c), we add a second stage to double the boost, the power penalty is another 4 mW, a significant fraction with respect to our

[1]In this design, the CTLE boost can be programmed from 0 to 5.5 dB in six steps of approximately 1.1 dB each.

receiver's overall power consumption. The key point here is that the CTLE power grows almost exponentially with the number of stages if the overall bandwidth must remain relatively constant.

IV. Discrete-Time Linear Equalization

A. Basic Idea

Linear equalization need not be a continuous-time function. In fact, wireline transmitters typically perform discrete-time discrete-amplitude equalization (de-emphasis) [15]. Shown in Fig. 3(a) is an FIR example incorporating an FF as a delay element and realizing a transfer function given by $H(z) = 1 - \alpha z^{-1}$.

We wish to extend the above technique to receivers, projecting that it may offer substantial power savings. However, we must deal with the continuous-time continuous-amplitude nature of the received dispersed data. That is, the data must be sampled with sufficient linearity so as to retain the channel profile information. Since dispersed data subjected to nonlinearity equivalently experience greater intersymbol interference (ISI) [14], we cannot perform the sampling and delay operations by means of FFs.

These considerations point to the conceptual arrangement shown in Fig. 3(b) as a potential candidate. Assuming $C_B \ll C_A$ for now to minimize charge sharing in the master–slave sampling circuit, we obtain $H(z) = 1 - \alpha z^{-1}$, and hence a high-frequency boost of $(1+\alpha)/(1-\alpha)$ with a low-frequency gain of $1 - \alpha$. As with the CTLE stage in Fig. 2(a), this approach too entails a tradeoff between the boost factor and the dc gain.

Linear equalization for a low-pass channel generally amplifies the high-frequency noise of the preceding stages. If placed after the CTLE, the DTLE of Fig. 3(b) tends to amplify the CTLE's high-pass-shaped noise. Nonetheless, with a one-stage CTLE, the overall output noise of the DTLE remains negligible. In this paper, the total noise at the DFE summing junction, including the CTLE and DTLE contributions, is around 1.1 $\mathrm{mV_{rms}}$.

B. Half-Rate DTLE

As mentioned in Section VI-A, we prefer half-rate operation for the CDR. In order to accommodate half-rate operation in the DTLE, we modify the receiver front end as conceptually illustrated in Fig. 4, where both the demultiplexer, $\mathrm{DMUX_1}$, and the DTLE are clocked at 20 GHz. The odd and even data bits produced by $\mathrm{DMUX_1}$ are delayed by 1 UI, 25 ps, scaled by a factor of α, and subtracted from the output of the other branch. Of course, $\mathrm{DMUX_1}$ must be sufficiently linear. We observe that, as D_{in} reaches S_{odd} and S_{even}, it experiences the high-frequency boost of both the CTLE and the DTLE, about 5.5 dB + 5.4 dB in our design.[2] As shown in Section IV-C, the additional boost provided by our DTLE costs only 0.3 mW.

The passive samplers in Fig. 4 entail two imperfections. First, in practice, some charge sharing occurs between

[2]In this design, the DTLE coefficient α can be programmed from 0 to 0.3 in 32 uniform steps.

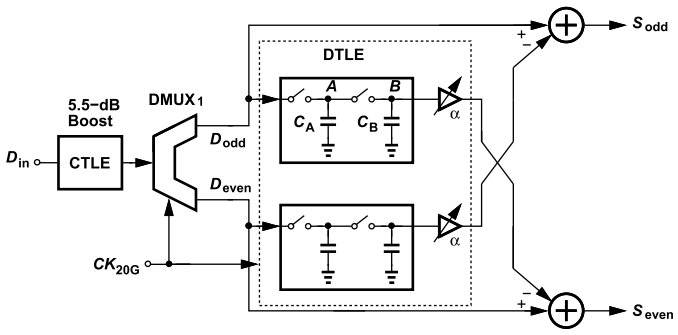

Fig. 4. Half-rate DTLE.

C_A and C_B. To determine the resulting transfer function, we denote the sampled voltage on C_A and C_B in the nth clock cycle by $V_{A,n}$ and $V_{B,n}$, respectively. For the first bit, we have $V_{B,1} = C_A V_{A,1}/(C_A + C_B)$ in the top branch. Due to half-rate operation, the next update corresponds to the third bit and generates $V_{B,3} = [C_A/(C_A+C_B)][V_{A,3}+V_{A,1}C_B/(C_A+C_B)]$. Continuing these calculations, one can show that the transfer function is given by

$$H(z) = 1 - \alpha \frac{C_A}{C_A + C_B}\left[z^{-1} + \frac{C_B}{C_A + C_B}z^{-3} + \frac{C_B^2}{(C_A + C_B)^2}z^{-5} + \ldots\right]. \quad (1)$$

Since this is a binomial infinite series, we can rewrite this transfer function as

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

Interestingly, charge sharing between C_A and C_B yields additional odd-order taps for the equivalent FIR filter, resulting in an IIR response. For example, if $C_B = 0.2C_A$, then

$$H(z) = 1 - \alpha \frac{0.83z^{-1}}{1 - 0.17z^{-2}}. \quad (3)$$

The second imperfection relates to kT/C noise in Fig. 4, especially if C_B is chosen small to minimize charge sharing. We formulate this effect in Section IV-C.

C. DTLE Implementation

The receiver front end shown in Fig. 4 in fact samples the data three times, once by $\mathrm{DMUX_1}$ and twice by C_A and C_B. We can therefore remove the DTLE's master samplers (the first set of switches and capacitors) and utilize $\mathrm{DMUX_1}$ in their stead. Fig. 5(a) illustrates the result, portraying $\mathrm{DMUX_1}$ as a passive sampler. At maximum boost, however, the low-frequency losses of the CTLE and the DTLE, 0 and -3 dB, respectively, make it desirable to have some gain in $\mathrm{DMUX_1}$. As shown in Fig. 5(b) for one branch, this is accomplished by means of a charge-steering amplifier [13]. Here, M_1 and M_2 are OFF while the CTLE output is sampled on X and Y, and regeneratively amplify $V_X - V_Y$ when the circuit enters the

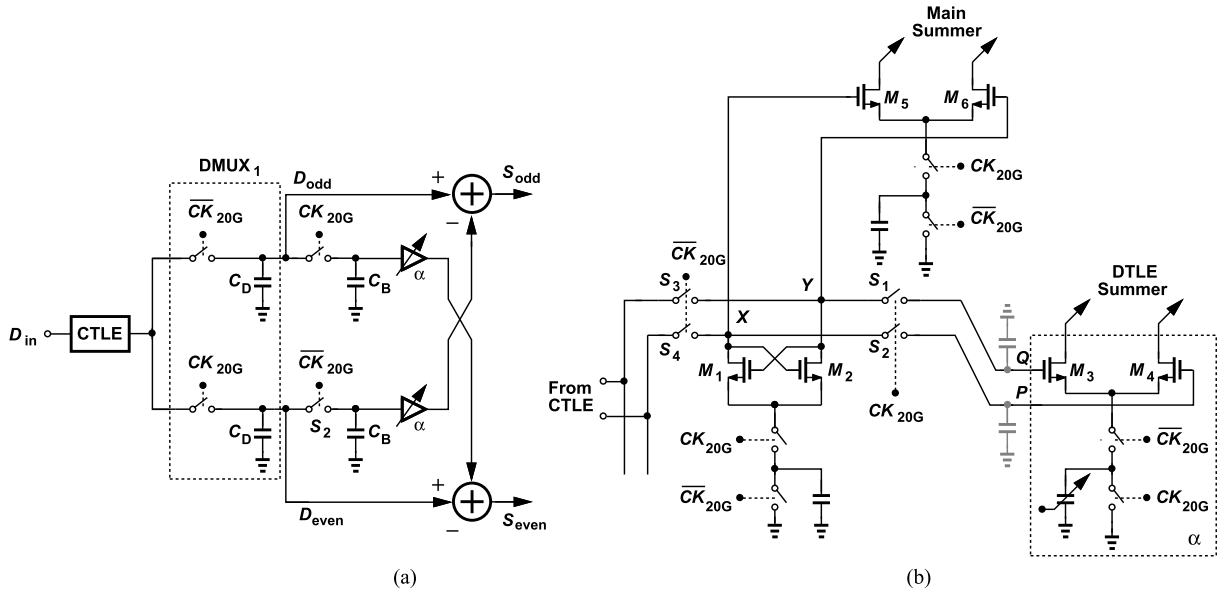

Fig. 5. Implementation of DTLE. (a) Conceptual diagram. (b) Transistor-level circuit. (Except for pMOS switches S_1–S_4, all switches are nMOS devices.)

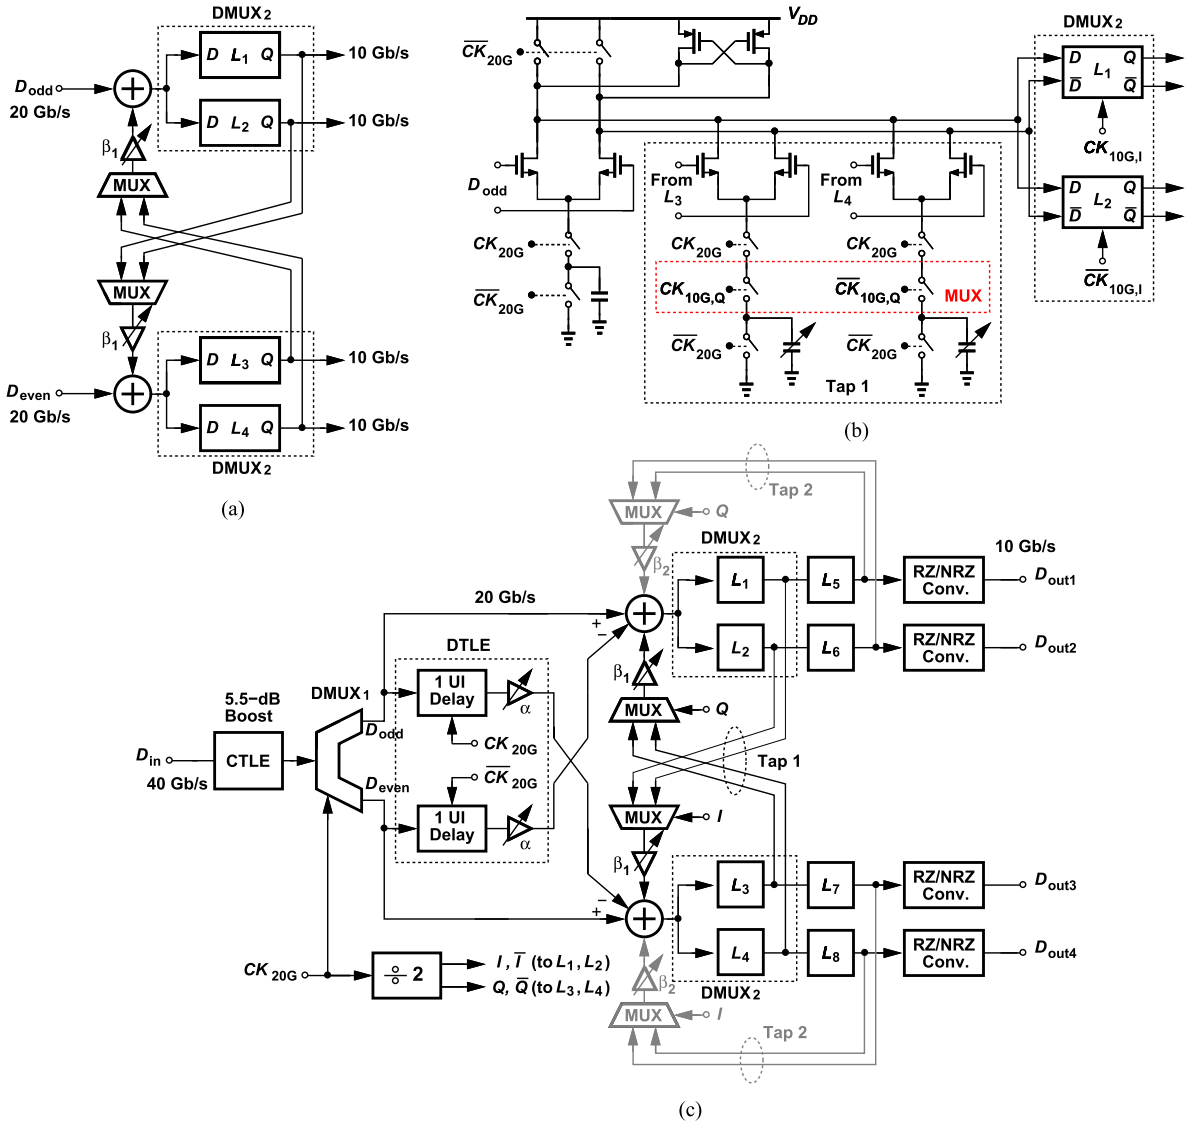

Fig. 6. (a) Half-rate/quarter-rate DFE. (b) Odd data path. (c) Proposed equalizer architecture.

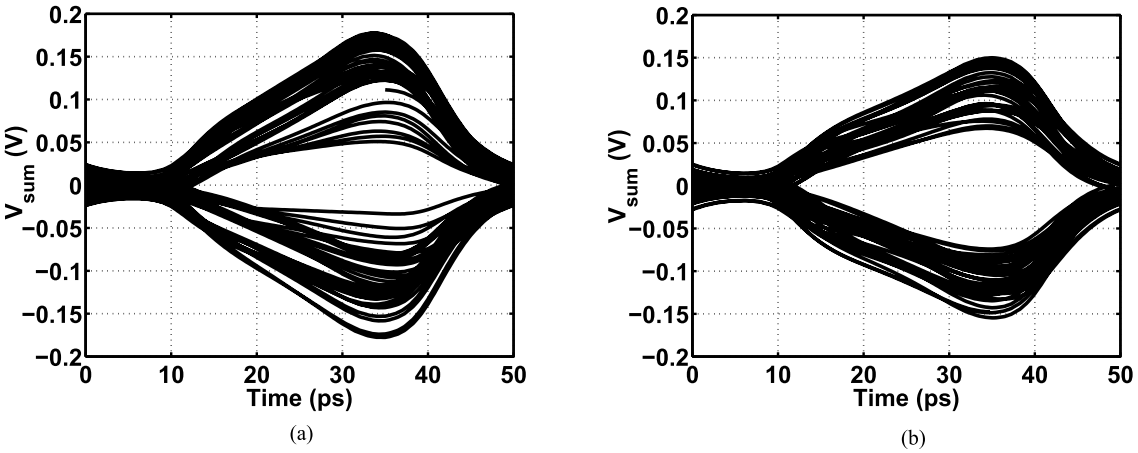

Fig. 7. Eye diagrams at the summing junction with DFE ON. (a) Without DTLE. (b) With DTLE.

hold mode, thereby providing a gain of 6 dB.[3] Charge-steering differential pairs M_3–M_4 and M_5–M_6 act as discrete-time G_m stages and drive the summing junctions according to Fig. 5(a). To minimize charge sharing, the parasitic capacitances at P and Q are kept less than one-fifth of those at X and Y.

The nonlinearity of the regenerative pair in Fig. 5(b) merits attention. As described in [14], the requirement on the linearity is that the main cursor at the input be smaller than 1.5 times the 1-dB compression point, $A_{1\ \mathrm{dB}}$. For $A_{1\ \mathrm{dB}} = 120\ \mathrm{mV_p}$, this stage contributes negligible ISI. With a 19-dB loss in the channel, the main cursor at the CTLE output is much less than the peak swing and hence, the output swing of the CTLE can be much larger than 120 $\mathrm{mV_p}$. Note that the regenerative pair dominates the linearity, making that of the switches less important.

Since the sampling operations in Fig. 5(b) rely on small parasitic capacitances, kT/C noise must be considered. The kT/C noise in Fig. 5(a) is obtained by first writing the sampled and amplified noise on each C_D as $4(kT/C_D)$, where the factor of four accounts for the 6-dB gain. The noise on C_B thus consists of kT/C_B and the noise stored on C_D after charge sharing occurs, that is, $kT/C_B+(4kT/C_D)[C_D/(C_D+C_B)]^2$. Multiplying this result by α^2 and adding to the noise on C_D, we obtain the total kT/C noise in S_{odd}. In this design, $C_D \approx 5\ C_B \approx 45$ fF and $\alpha_{\max} = 0.3$, yielding a total noise of approximately $(2 \times\ (4.7kT)/C_D)^{1/2} \approx 0.93\ \mathrm{mV_{rms}}$, where the factor of two accounts for the differential signal path.

V. DTLE/DFE Combination

The half-rate DTLE developed in Section IV-C can naturally merge with a half-rate DFE. We begin with a charge-steering implementation of such a DFE [14], whose first tap is shown in Fig. 6(a) for reference. Here, the half-rate data streams, D_{odd} and D_{even}, are applied to the summers, whose outputs are demultiplexed by latches L_1–L_4 by another factor of two and multiplexed before driving the summer in the other branch. The loop timing constraint of this DFE is given by [14]

$$t_{\mathrm{cq}} < 1\ \mathrm{UI} \tag{4}$$

[3]The gain of the regenerative pair varies from 5.66 dB at SS 80 °C to 6.789 dB at FF 0 °C.

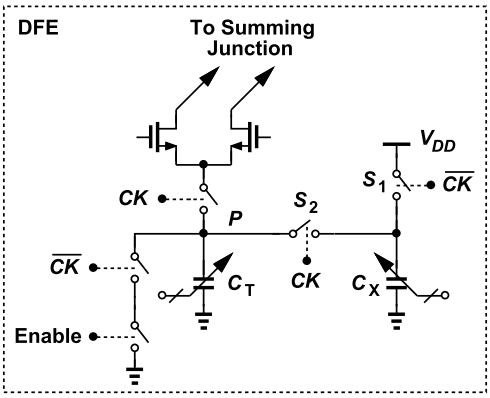

Fig. 8. Vernier charge delivery.

where t_{cq} denotes the delay from the odd (or even) summer's clock to the output of one of the latches. Avoiding the setup time and the feedback delay, this DFE obviates the need for loop unrolling. Fig. 6(b) depicts the implementation of the odd path for the first DFE tap. The multiplexer in the feedback path is merged with the tap, thus avoiding additional delay. We now combine the DTLE and the DFE as illustrated in Fig. 6(c), where the second DFE tap is also realized.

In order to study the improvement afforded by the DTLE, we perform transistor-level simulations on the CTLE–DTLE–DFE cascade in the presence of a channel loss of 20 dB at Nyquist. Unless otherwise mentioned, the system simulations are performed with a single-ended input swing equal to 200 $\mathrm{mV_{pp}}$. Fig. 7 plots the eye diagrams at the output of one summer before and after the DTLE is added.[4] We observe that the vertical opening at $t = 35$ ps increases from ± 30 to ± 70 mV.[5] This means that, for example, the DFE latches can now have 40 mV more input-referred offset voltage, benefiting from less input capacitance and lower power consumption.

As with other DFE topologies, additional taps can be added to Fig. 6(c), with alternate latch stages clocked by the quadrature phases. Operating at quarter rate and sensing large

[4]The return-to-zero behavior stems from the charge-steering operation of the summers [14].

[5]With a single-ended input swing of 400 $\mathrm{mV_{pp}}$, the vertical opening at the summer output increases to ± 105 mV.

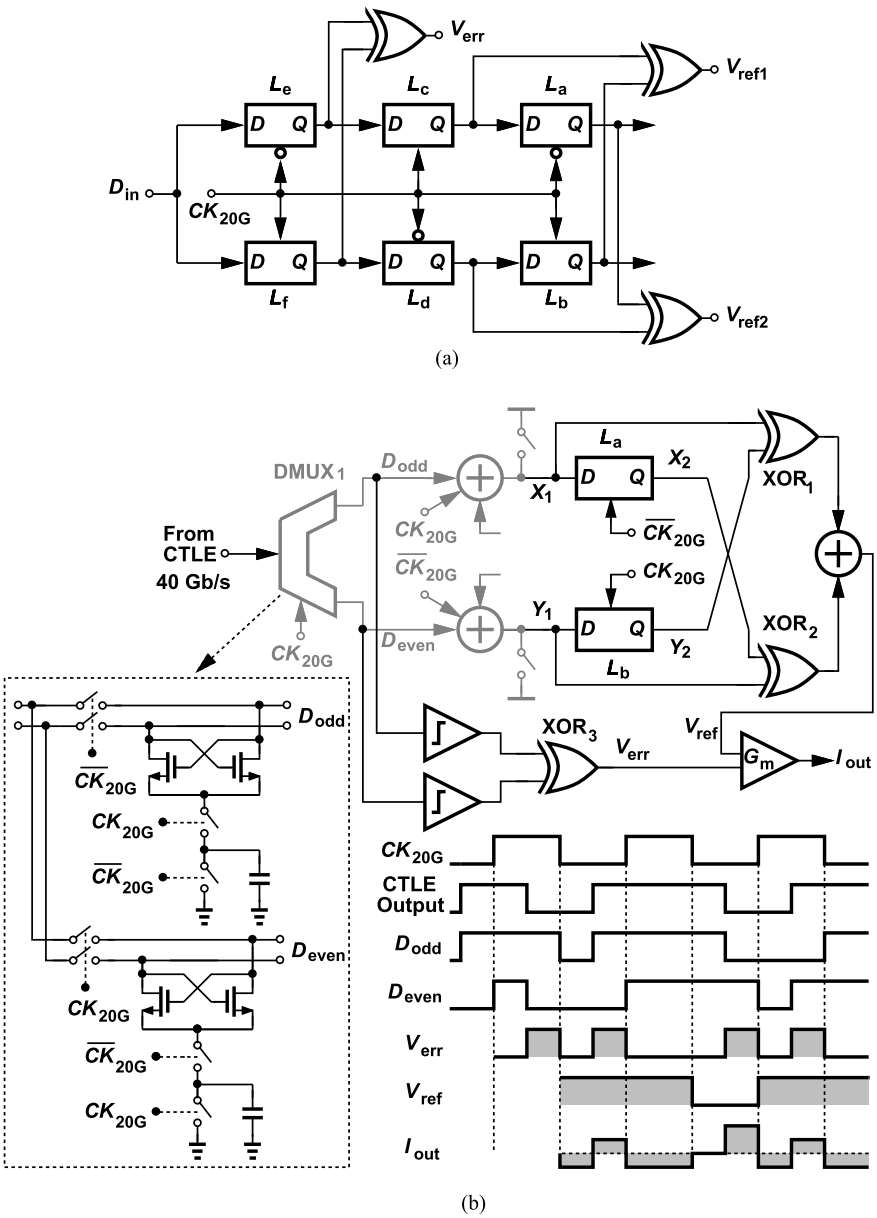

Fig. 9. (a) Half-rate charge-steering PD. (b) Proposed PD.

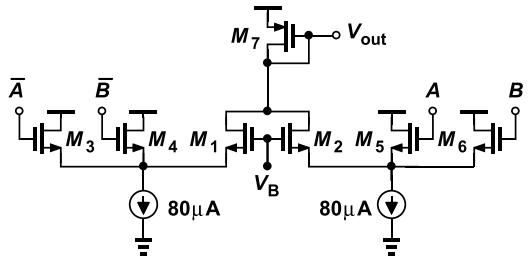

Fig. 10. Symmetric XOR gate.

data swings, such latches would have a simple design and consume less power. The higher capacitance at the summing junction would require a larger capacitor in the summer's tail.

A. Vernier Charge Delivery

The summers in Fig. 6(c) incorporate four charge-steering differential pairs at their inputs. As exemplified by M_3–M_4 and M_5–M_6 in Fig. 5(b), the strength of these pairs can be controlled by adjusting the value of the tail capacitors, but, even with these set to zero, the parasitic capacitances in the tail path draw some current when CK_{20G} is asserted. This effect becomes critical if one of the taps in the DFE must be set to zero for a particular channel.[6]

We introduce a Vernier charge delivery technique that can force a tap value to zero. As illustrated in Fig. 8, the idea is to charge C_X to V_{DD} while C_T is discharged to zero. When charge steering begins, C_X can provide the charge required by C_T and other parasitics, thereby yielding a nominally zero charge available for the differential pair. With a resolution of 1 fF for C_X, the tap value can be reduced to nearly zero.

The Vernier technique alleviates another issue as well. Since the CDR's VCO drives the switches in Fig. 8 without a

[6]Feedthrough from input to output of charge steering stages is negligible as the input transistors reside deep in saturation when the circuit is clocked, thus yielding low C_{GD} for these transistors.

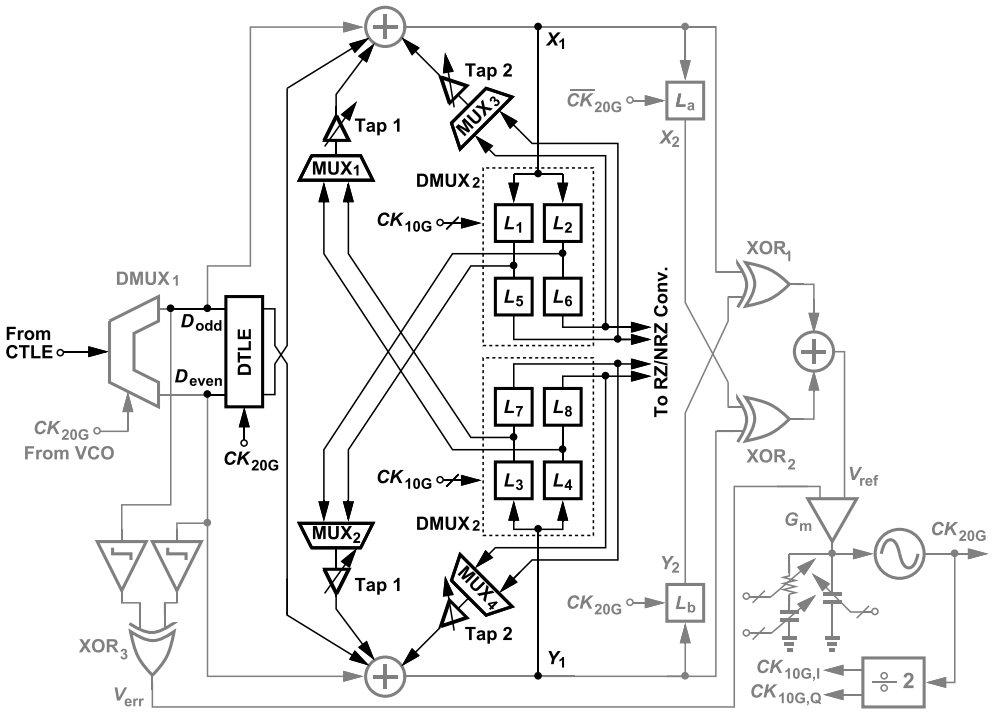

Fig. 11. Proposed receiver architecture with half-rate/quarter-rate CDR and DFE.

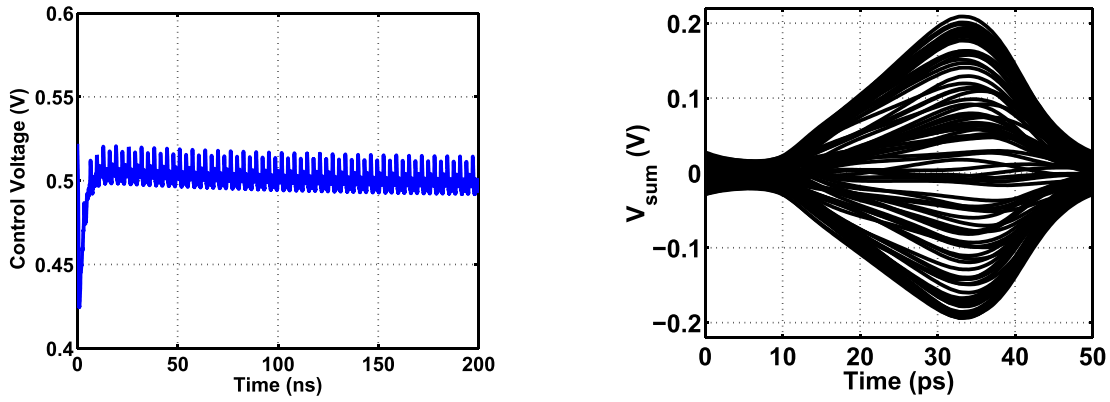

(a) (b)

Fig. 12. CDR behavior with DFE/DTLE tap coefficients set to zero. (a) Control voltage transient. (b) Eye diagram at the summing junction after the CDR has locked.

buffer (Section VII-B), its frequency shifts if the tail capacitors are changed by a large ratio so as to provide the necessary coefficients (α in the DTLE and β_1 and β_2 in the DFE). With the Vernier method, on the other hand, the change is reduced by a factor of three, from 200 to 70 MHz.

VI. Hardware Sharing

A. Merging Techniques

We wish to share the high-speed blocks between the receiver functions. We return to the front end shown in Fig. 6(c) and seek to lower the power consumption and capacitances in the data and clock paths. In particular, since both the front end of Fig. 6(c) and the CDR's PD incorporate samplers in the data path, we surmise that some sharing is possible. We should remark that: 1) the full-rate merged CDR/DFE described in [16] does not lend itself to half-rate operation and 2) the half-rate CDR/DFE topology in [10] cannot operate with charge steering or without quadrature clock phases. The other challenge is that the charge-steering summer outputs in Fig. 6(c) are precharged for half of the clock cycle, thereby producing return-to-zero data. For this reason, the CDR cannot be placed after these summers.

We begin with the PD and note that a half-rate charge-steering design [Fig. 9(a)] must take three samples of the full-rate data by means of six latches [13]. We then return to the front end shown in Fig. 6(c) and seek interfaces that provide such samples: DMUX_1 itself samples the data and the (charge-steering) summing circuits also act as samplers. Thus, only one more stage of sampling is necessary for unambiguous phase detection.

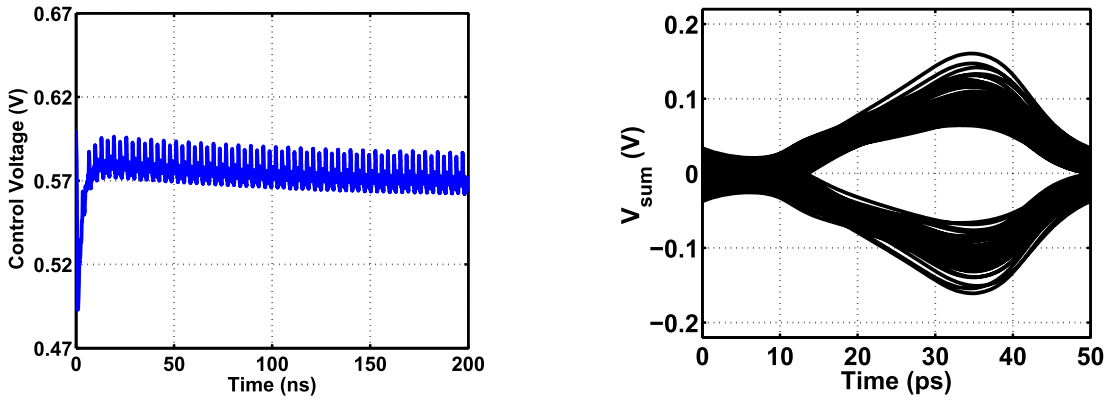

Fig. 13. CDR behavior with DFE/DTLE tap coefficients set correctly initially. (a) Control voltage transient. (b) Eye diagram at the summing junction after the CDR has locked.

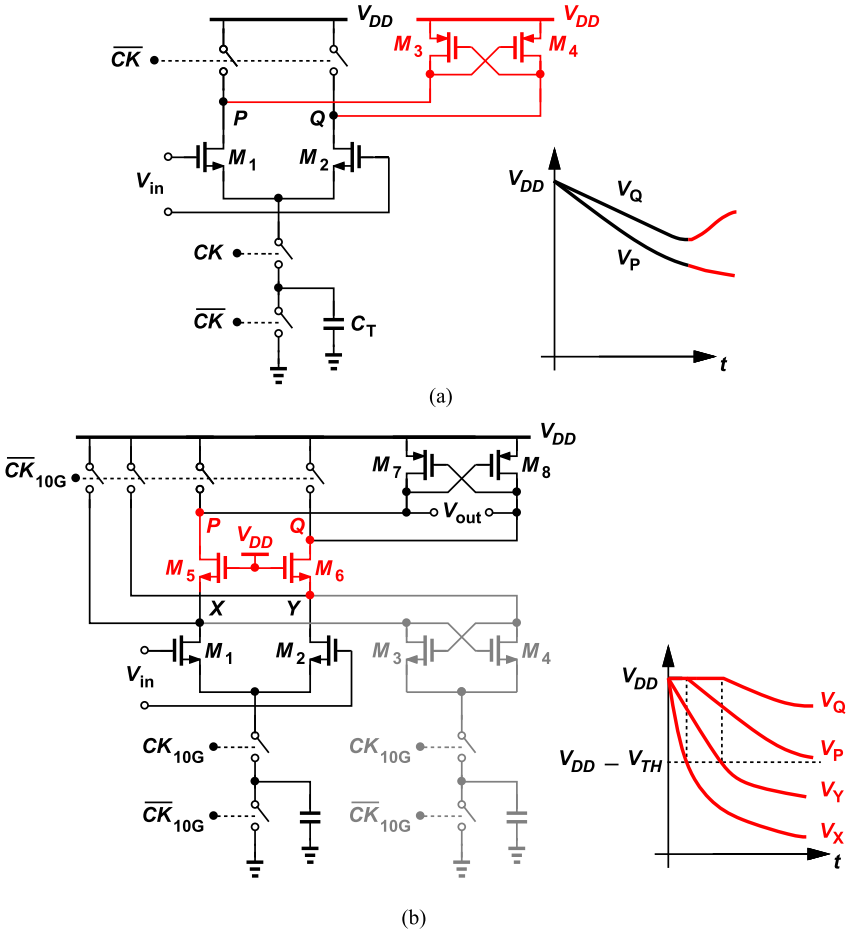

Fig. 14. (a) Operation of RZ charge-steering latch with a cross-coupled pMOS pair. (b) Adding a cascode pair to the charge-steering latch.

The foregoing thoughts lead to the PD topology depicted in Fig. 9(b), where the gray blocks are shared with the receiver front and the DFE. Here, DMUX_1 acts as L_e and L_f in Fig. 9(a), and the summers as L_c and L_d. As illustrated by the waveforms, XOR_3 measures the phase difference between D_{odd} and D_{even}, generating V_{err}. In addition, XOR_1 and XOR_2 produce a pulse of constant width on V_{ref} for each data transition. The difference, $V_{\text{err}} - V_{\text{ref}}$, uniquely represents the phase error regardless of the data pattern. The symmetric XOR gates [13] are implemented as shown in Fig. 10. Note that the G_m stage measures the area under $V_{\text{err}} - V_{\text{ref}}$ and need not operate at high speeds.

The limiting stages preceding XOR_3 in Fig. 9(b) are included for the following reason. We note that DMUX_1 is a linear sampler whereas L_e and L_f are not. Thus, a 1010 sequence traveling through the channel and heavily attenuated by its loss appears in D_{odd} and D_{even} with no slicing but in X_2 and Y_2 with slicing. As a result, if directly sensing D_{odd} and D_{even},

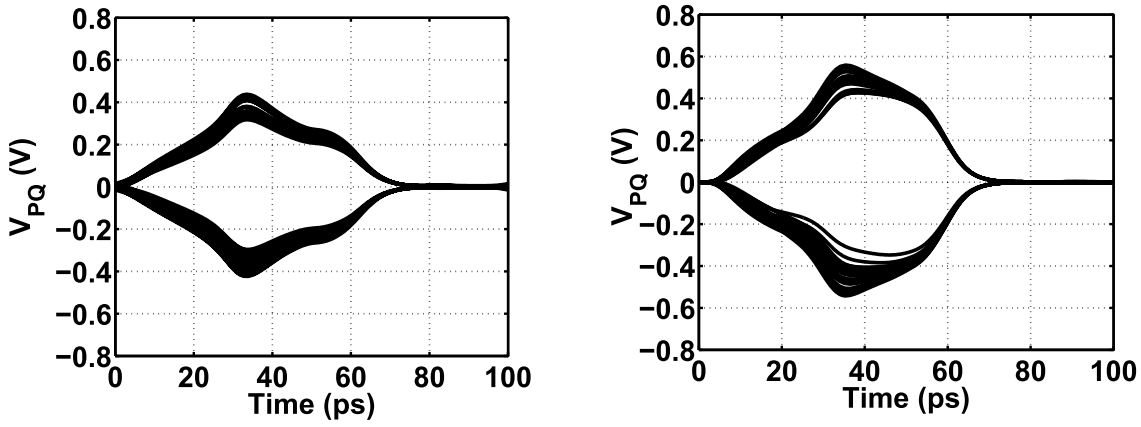

(a) (b)

Fig. 15. Eye diagram at the output of latch (a) without and (b) with cascode devices.

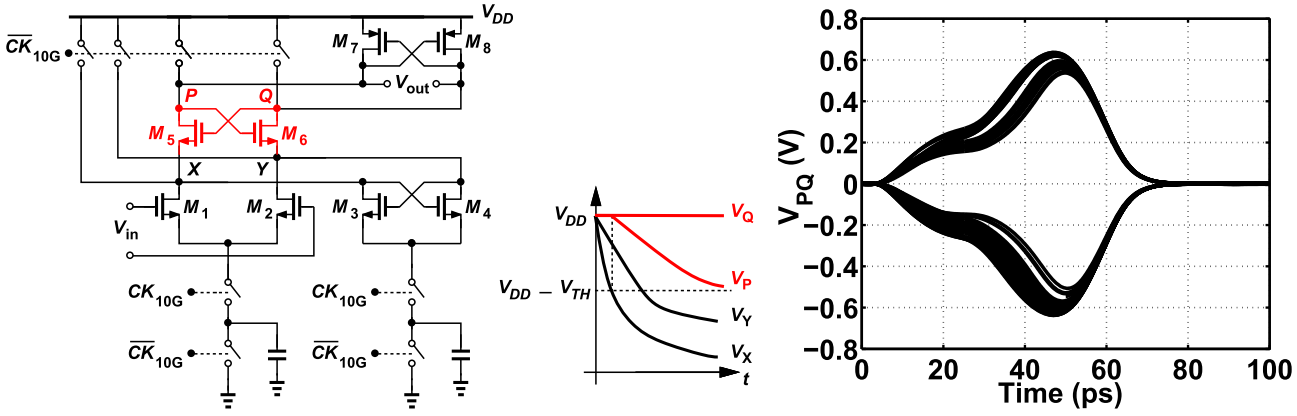

(a) (b)

Fig. 16. (a) Improved charge-steering latch with two cross-coupled nMOS pairs. (b) Its simulated output eye diagram.

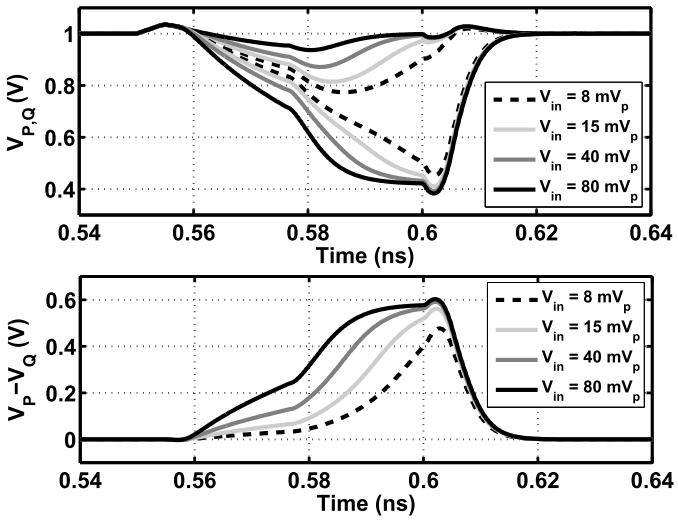

Fig. 17. Outputs of charge-steering latch with two cross-coupled nMOS pairs with various input swings.

XOR_3 would sometimes operate in the small-signal regime, generating an output inconsistent with those of XOR_1 and XOR_2. Realized as simple differential pairs with a tail current of 0.2 mA, the slicers perform limiting on D_{odd} and D_{even} by virtue of the 6-dB gain provided by $DMUX_1$, and hence its moderate output swings. (360-mV_{pp} differential swing).

B. Complete Receiver

Fig. 11 shows the overall receiver design. A single-stage CTLE [Fig. 2(b)] drives $DMUX_1$, which, along with the summers, latches L_a and L_b, and $XOR_{1\text{-}3}$, acts as the CDR PD. The PD outputs are combined and converted to current by a G_m stage. This half-rate CDR loop is completed by the second-order filter and the VCO. The half-rate/quarter-rate DFE consists of the summers, latches $L_{1\text{-}4}$, and $MUX_{1,2}$ for the first tap. The second tap employs $L_{5\text{-}8}$ and $MUX_{3,4}$. The retimed demultiplexed RZ data produced by $L_{5\text{-}8}$ are converted to NRZ form as described in [13].

In this design, the CTLE draws 2 mW, the PD, DTLE, and all of the latches 5.3 mW, the VCO 2.8 mW, the ÷2 circuit 3.4 mW, and the RZ/NRZ converter 0.53 mW.

From Monte Carlo simulations, the overall rms offset referred to the summer output is 15.1 mV, out of which 8.3 mV is contributed by the CTLE, 8.1 mV by the DFE latches, and about 9.7 mV by the remaining charge-steering circuits.[7] This number encompasses the offsets due to mismatch in transistors and also the mismatch between even and odd paths. (In this prototype, no offset correction scheme is used.)

The intertwined CDR and DFE loops in Fig. 11 can fight and fail to converge. The operation can begin by setting the

[7]This offset can be removed by a programmable threshold for the latch.

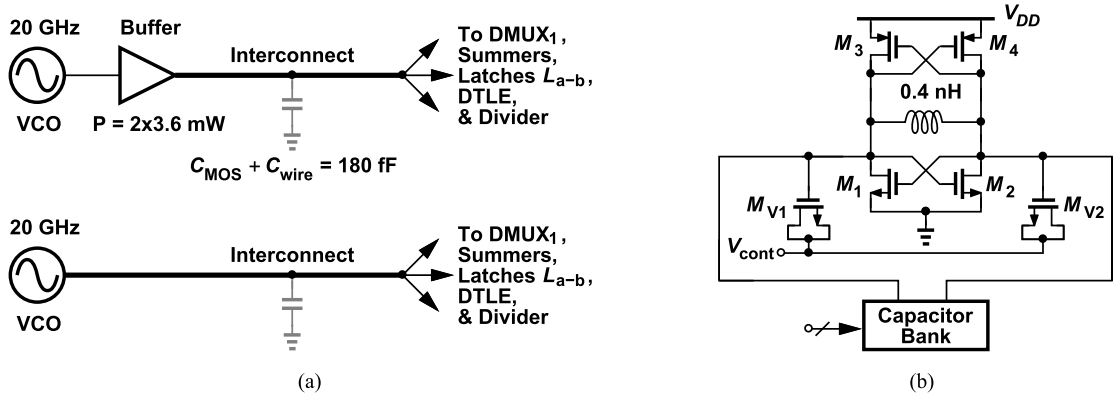

Fig. 18. (a) VCO: to buffer or not to buffer? (b) VCO implementation.

DFE and DTLE tap coefficients to zero and the CTLE boost factor to its maximum value. The gray path detects the phase error between the data and the VCO output, delivering a proportional current to the loop filter and driving the oscillator toward 20 GHz. Despite the heavy ISI at the CTLE output, the CDR locks, as shown in Fig. 12, because data patterns having several consecutive ONEs or ZEROs make full transitions and provide sufficient phase information. Moreover, with about 11 dB of voltage gain through DMUX_1 and the limiters,[8] the transitions presented to XOR_3 become sharp enough to produce proper phase error. The CDR takes approximately 100 ns to lock, after which the CTLE boost and the DFE and DTLE tap coefficients are adapted to complete the equalization.[9] In this design, a peak-to-peak ripple amplitude of 1 mV translates to 17 fs of output jitter.

If we begin with the correct DFE and DTLE tap coefficient, the CDR still locks in about 100 ns with an open eye at the summing junction (Fig. 13). One can also apply automatic coefficient adaptation provided that the update steps are long enough for the CDR to lock each time.

It is interesting to note that the settling behavior of the CDR is similar in the two cases of Figs. 12 and 13 with approximately the same ripple on the control voltage. This is because the input to the PD in Fig. 9 comes from the CTLE, and hence is independent of the DFE/DTLE coefficients.

Also, owing to the 70-MHz shift in the VCO frequency when the DFE and DTLE taps are changed (Section V-A), the control voltage settles to different values in Figs. 12(a) and 13(a). For $K_{\text{VCO}} = 1$ GHz/V, this change is equal to $\Delta f / K_{\text{VCO}} = 70$ mV.

In steady-state operation, the CDR is locked in Fig. 11 and the DFE produces properly equalized data at the summer outputs, X_1 and Y_1. In the presence of a lossy channel, therefore, the BER of X_1 and Y_1 (or X_2 and Y_2) is much lower than that of D_{odd} and D_{even}. This points to some inconsistency between the output of XOR_3 and those of XOR_1 and XOR_2, which ultimately manifests itself as a static phase error within the CDR loop. Nevertheless, so long as the error in the number of transitions at the CTLE output is less than about 1 in 50, the resulting phase error is negligible. In addition, we choose the gain associated with the V_{err} path to be higher than that of the V_{ref} path so as to remove the phase offset. This gain difference is adaptively created along with the other receiver coefficients.

[8]The voltage gain through DMUX_1 and the limiters changes from 11 dB at SS 80 °C to 10.46 dB at FF 0 °C.

[9]In this prototype, since we do not have an error/observation path to drive coefficient adaptation, the adaptation is done "manually" through a serial bus to find a setting that works based on BER.

VII. Building Blocks

In this section, we deal with the transistor-level design of the critical building blocks. We should remark that proper choice of various CM levels allows us to avoid bootstrapped switches. Also, the switches are chosen wide enough to provide proper speed while their capacitance is absorbed by the VCO tank (Section VII-B).

A. Latch Design

As illustrated in Fig. 11, the receiver incorporates eight latches, $L_{1\text{-}8}$, at a clock frequency of 10 GHz. In order to improve the speed of charge-steering implementations, we introduce several new topologies here.

Consider the charge-steering latch shown in Fig. 14(a) [14]. We begin with a differential pair whose outputs are precharged to V_{DD} and has a charge source rather than a current source in its tail which is completely discharged. When CK goes high, there is charge sharing between capacitances at P and Q, and C_T, which results in a differential swing at the output. Without the pMOS pair, the circuit exhibits a limited voltage gain because the output CM level drops significantly as the tail charge flows from the capacitances at P and Q, driving M_1 and M_2 into the triode region. With the pMOS pair present, the high level can be restored, but only if the circuit has enough time. In our design, however, the clock period does not allow much restoration.

In order to improve the performance, we add two cascode devices as shown in Fig. 14(b), with nodes X and Y also precharged to V_{DD} during reset. When CK goes high, the CM level at X and Y falls and $|V_X - V_Y|$ increases, but P and Q remain at V_{DD} until the voltage at either X and Y drops to about $V_{\text{DD}} - |V_{\text{TH}}|$. The key point here is that the cascode devices isolate X and Y from the large capacitance at the output nodes, raising the gain from V_{in} to V_{XY}. As a result, either M_5 or M_6 remains OFF for some time while the other

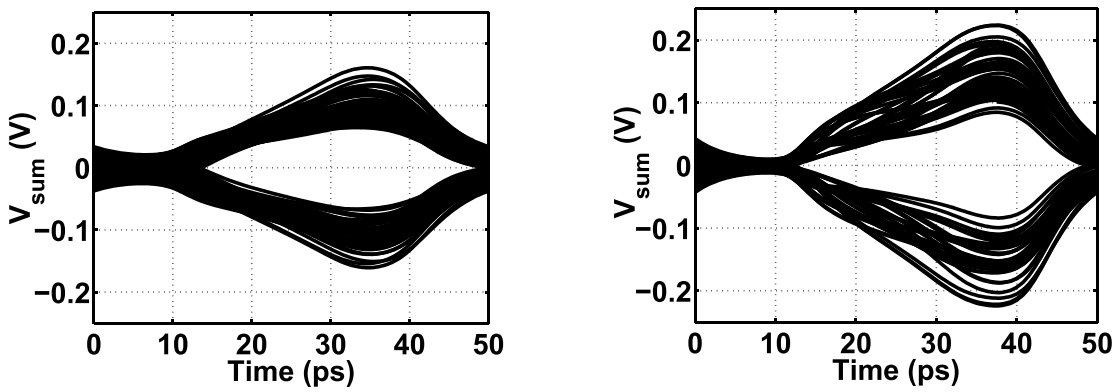

Fig. 19. Eye diagrams at the summing junction with (a) $V_{DD} = 1$ V and (b) $V_{DD} = 1.1$ V.

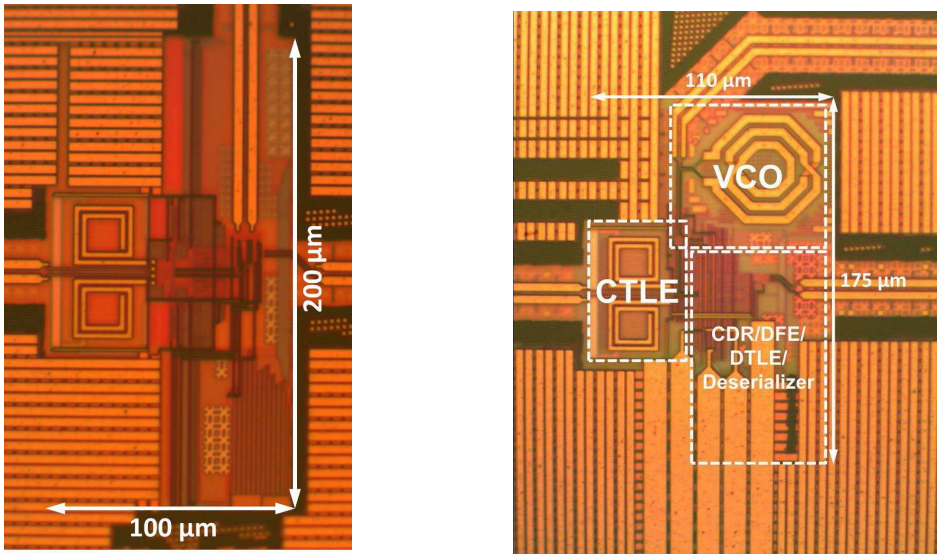

Fig. 20. Die photograph of (a) stand-alone equalizer and (b) complete receiver.

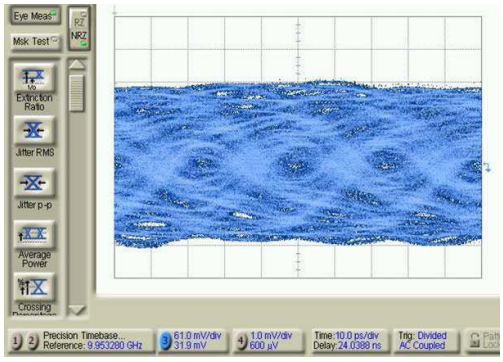

Fig. 21. Measured eye diagram at the equalizer input at 40 Gb/s (10 ps/div., 61 mV/div.).

transfers the amplification to the output. The performance is further improved by tying another cross-coupled pair, M_3–M_4, to X and Y so as to increase the voltage gain as V_X and V_Y fall from V_{DD}. Fig. 15 plots the simulated outputs of the latches in Fig. 14(a) and (b), demonstrating about 30% greater eye opening for the latter.

In the circuit of Fig. 14(b), neither of the cascode devices remains OFF for the entire amplification period because both V_X and V_Y eventually fall below $V_{DD} - V_{TH}$. Consequently, some CM drop occurs at P and Q, wasting the tail charge. This issue can be avoided if M_5 and M_6 are reconfigured to act as a cross-coupled pair [Fig. 16(a)]. Now, if, for example, V_X reaches $V_{DD} - V_{TH}$ before V_Y does, then M_5 conducts and lowers V_P, thereby postponing the turn-on of M_6. As plotted in Fig. 16(b), the simulated output exhibits a greater swing than the previous two latches. In the receiver of Fig. 11, $L_{5–8}$ are realized as the former [Fig. 14(b)] and, $L_{a,b}$ and $L_{1\text{-}4}$ as the latter [Fig. 16(a)]. This is because the cascode latches have a lower parasitic capacitance at their output node and are somewhat faster. Fig. 17 plots the single-ended and differential outputs of latch L_1 for various swings at its input. It can be seen that even with a summer output swing of ±8 mV, the latch output still reaches about 80% of its final value.[10]

While reminiscent of the StrongArm latch, the topology of Fig. 16(a) bears two differences.

1) It operates with a finite tail charge, thus producing moderate swings at X, Y, P, and Q; since the voltages

[10]The latches L_{1-4} in Fig. 11 require an output differential swing of only ±150 mV to completely steer the feedback transconductance for the first tap.

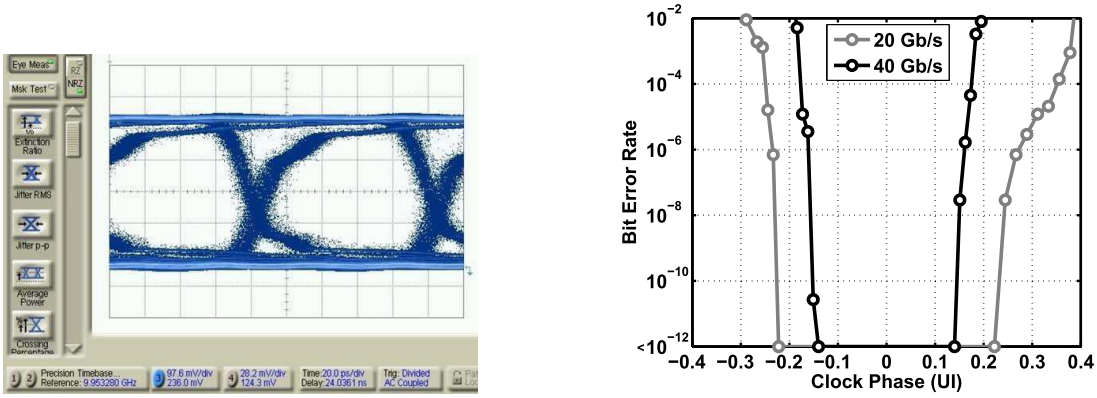

(a) (b)

Fig. 22. (a) Equalized and demultiplexed output data at 10 Gb/s (20 ps/div., 97.6 mV/div.). (b) Measured bathtub curves.

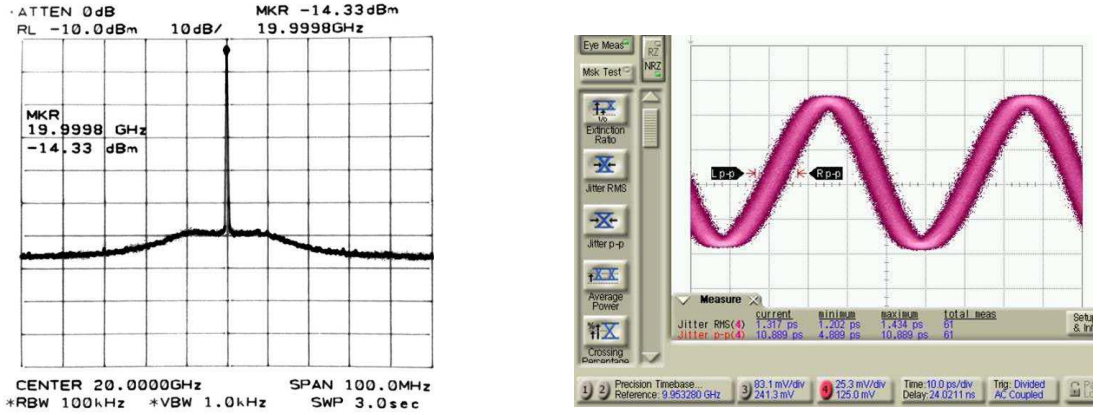

(a) (b)

Fig. 23. Recovered clock. (a) Spectrum. (b) Waveform (10 ps/div., 25.3 mV/div.).

at these nodes do not collapse to zero, both the speeds are improved and the power consumption (given by $fCV_{DD}V_{swing}$ for each node) is reduced.

2) The additional gain contributed by M_3–M_4 also enhances the speed.

B. VCO Design Considerations

As mentioned in Section II, the receiver eliminates all buffers in the data and clock paths. The VCO in Fig. 11 directly drives $DMUX_1$, the DTLE, latches L_a and L_b, and all of the interconnects. In a manner similar to that in [13], this approach saves considerable power while negligibly suffering from data coupling and de-Qing effects.

As shown in Fig. 18(a), the 20-GHz VCO drives charge-steering latches and, as such, must provide (nearly) rail-to-rail voltage swings. With a single-ended load capacitance of $C_L \approx 180$ fF, the differential inductor value is limited to about 0.4 nH, which yields a parallel equivalent resistance of $R_p = QL\omega = 402\ \Omega$ if $Q = 8$. For a 1-V single-ended swing, therefore, the oscillator must be biased at 2 mA. The actual design, shown in Fig. 18(b), draws 2.8 mA. (By comparison, two inverters driving these capacitances at 20 GHz would draw 7.2 mW.) In this design, $W_{1,2} = 9\ \mu$m and $W_{3,4} = 30\ \mu$m.

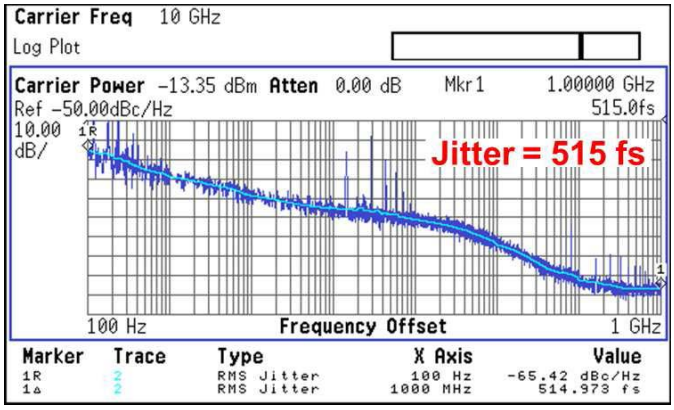

Fig. 24. Phase noise of recovered clock at 10 GHz.

The swing-dominated power consumption calculated for the VCO must be compared against that required for sufficiently low jitter. We express the integrated random jitter as $t_\sigma \approx T_{VCO}(4\ f_{BW}S_0)^{1/2}/(2\pi)$, where T_{VCO} denotes the VCO period, f_{BW} the CDR loop bandwidth, and S_0 the VCO phase noise at an offset equal to f_{BW}. For example, if $t_\sigma = 0.2$ ps_{rms} and $f_{BW} = 20$ MHz, then $S_0 = -111$ dBc/Hz. Using the following phase noise expression for an LC VCO with two

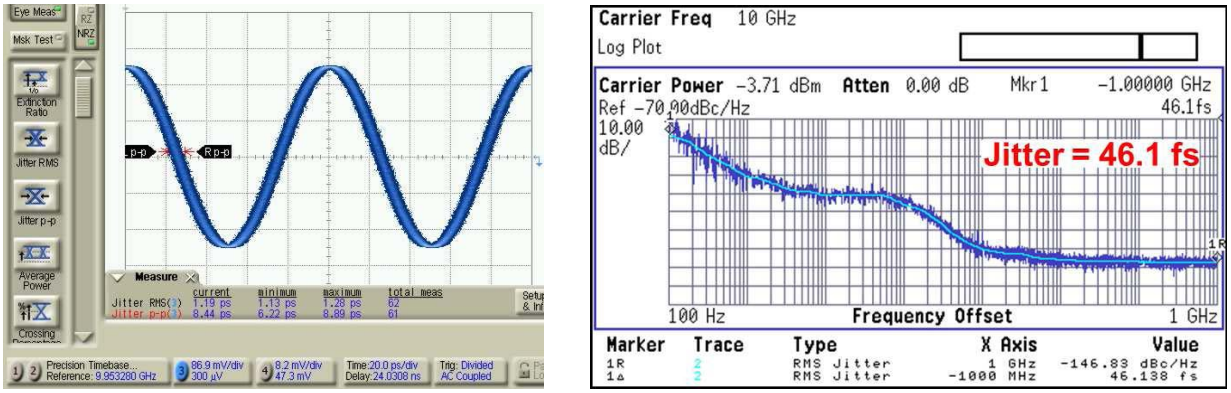

(a) (b)

Fig. 25. (a) Oscilloscope display of a low-noise RF tone (20 ps/div., 86.9 mV/div.). (b) Phase noise spectrum of the same tone.

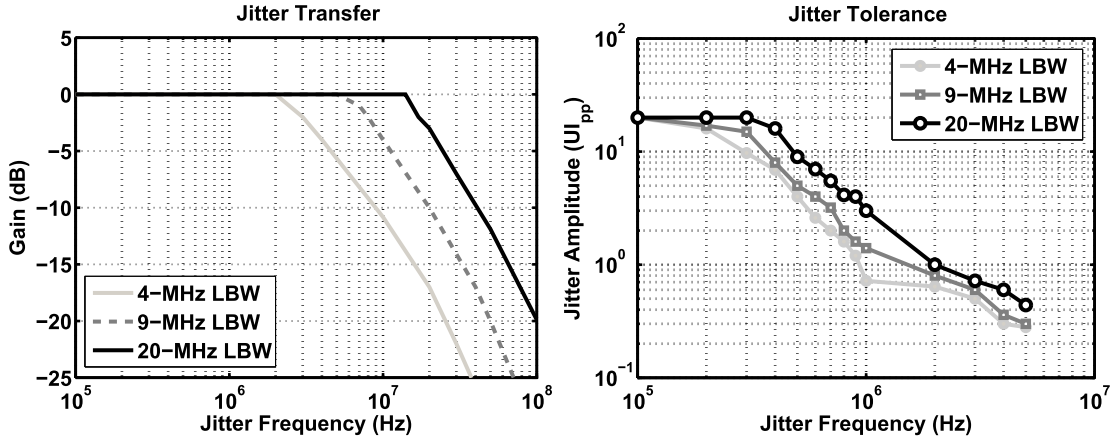

Fig. 26. Measured jitter transfer and tolerance curves.

cross-coupled pairs [17]:

$$S(\Delta\omega) = \frac{\pi^2}{R_P}\frac{kT}{I_{SS}^2}\left(\frac{\gamma_p+\gamma_n}{2}+1\right)\frac{\omega_0^2}{16Q^2\Delta\omega^2} \qquad (5)$$

for $\gamma_p = \gamma_n = 1$, we obtain $I_{SS} = 0.158$ mA, concluding that the power consumption is dictated by the required voltage swing rather than by the phase noise.

The VCO has a measured tuning range of 18.8–20.4 GHz, with five discrete tuning steps using a capacitor bank and $K_{VCO} \approx 1$ GHz/V.[11]

VIII. PVT Tolerance

The process–voltage–temperature tolerance of charge-steering circuits can be studied both analytically and through simulations. The process and temperature sensitivities have been provided in previous sections. In this section, we consider supply dependence. In general, since the input and output CM levels of the latches in Figs. 14 and 16 begin from V_{DD}, supply variation translates to a small change in the output voltage swing. According to simulations, a ΔV_{DD} of 10% translates to 11.5% change in the swing. Thus, supply noise primarily affects the output CM level—unless it contains frequencies comparable with the data rate.

The effect of supply can be quantified for the entire receiver. Fig. 19 plots the simulated eye at the DFE summing junction (with a channel loss of 20 dB) with $V_{DD} = 1$ or 1.1 V. We observe that the circuit operates satisfactorily in both cases.

[11]The simulated VCO tuning range at TT 27 °C is 19.21–20.8 GHz, at SS 80 °C, 18.77–20.29 GHz, and at FF 0 °C, 19.66–21.28 GHz.

IX. Experimental Results

This section summarizes our measured results for two prototypes: the stand-alone equalizer (including the CTLE, the DFE, the DTLE, and the RZ/NRZ converters) and the complete receiver. Both have been fabricated in TSMC's 45-nm CMOS technology and characterized with a 1-V supply at 40 Gb/s. The chips are mounted directly on printed-circuit boards and tested on a high-speed probe station. Fig. 20 shows their die photographs. The equalized and demultiplexed data at 10 Gb/s and the recovered clock at 20 GHz travel through open-drain pMOS buffers to drive off-chip 50-Ω cables and instrumentation.

Fig. 2(a) plots the frequency response of the channel used in the measurement for 40-Gb/s data rate in black indicating a channel loss of about 19 dB at Nyquist if the insertion loss of the probes is included. Fig. 21 shows the eye diagram received from the channel.

The equalizer's demultiplexed single-ended output at 10 Gb/s is displayed in Fig. 22(a) and the bathtub curve in Fig. 22(b). The horizontal eye opening at 40 Gb/s is 0.28 UI. It is important to note that the PRBS generator in our setup has a jitter of 8 ps_{pp}, contributing to the eye closure.

The measured frequency response of the channel used for 20-Gb/s data rate is shown in Fig. 2(a) in gray, exhibiting

TABLE I
PERFORMANCE SUMMARY AND COMPARISON WITH PRIOR ART

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	-	$\approx$ 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}

nearly 22 dB of loss at Nyquist. At 20 Gb/s, Fig. 22(b) indicates a horizontal eye opening of 0.44 UI, thus illustrating the scaling property of charge-steering circuits.

The complete receiver is tested with a PRBS of $2^7 - 1$ for jitter generation, transfer, and tolerance while it equalizes the dispersed data. Unless otherwise specified, the default CDR loop bandwidth of 20 MHz is used for all measurements. Fig. 23(a) and (b) shows the recovered clock spectrum and waveform, respectively.[12] To construct the phase noise profile using our instrumentation, we apply the 20-GHz clock to an external ÷2 circuit and monitor the result. As shown in Fig. 24, the profile should be raised by 6 dB to obtain that of the actual clock. The integrated rms jitter is 515 fs from 100 Hz to 1 GHz, which would also correspond to the jitter before the divider if the latter has negligible phase noise. Included in the jitter value, the tones in the phase noise profile correspond to the periodicity of the PRBS sequence and back-coupling of the half-rate data directly into the VCO. We should remark that direct jitter measurement in the time domain can be corrupted by oscilloscope jitter. As shown in Fig. 25(a), an oscilloscope measuring a low-noise tone provided by an RF generator exhibits 1.2-ps rms jitter whereas the phase noise spectrum of the same tone [Fig. 25(b)] yields an integrated jitter of 46 fs.

Fig. 26 plots the measured jitter transfer and tolerance, respectively, for three different CDR loop bandwidths. As expected, jitter tolerance improves as the bandwidth increases, reaching 0.45 UI_{pp} at 5-MHz offset for a loop bandwidth of 20 MHz in the presence of the 19-dB channel loss. Table I compares the performance of our receiver with that of the prior art at 40 Gb/s.

[12]The clock and data swings shown here are different due to different number of off-chip splitters and different speeds.

X. CONCLUSION

The ability to reduce high-speed hardware, whether by omitting stages or by merging functions, can lead to significant power savings in wireline receiver design. Along with discrete-time linear equalization and charge-steering techniques, this minimalist approach has afforded one order of magnitude reduction in power at a data rate of 40 Gb/s. The receiver also occupies a small area, 0.019 mm^2, and can serve in applications with a large number of I/O pins.

REFERENCES

[1] F. O'Mahony *et al.*, "A 47 × 10 Gb/s 1.4 mW/Gb/s parallel interface in 45 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2828–2837, Dec. 2010.

[2] K. Kaviani *et al.*, "A 0.4-mW/Gb/s near-ground receiver front-end with replica transconductance termination calibration for a 16-Gb/s source-series terminated transceiver," *IEEE J. Solid-State Circuits*, vol. 48, no. 3, pp. 636–648, Mar. 2013.

[3] J. Poulton *et al.*, "A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short-reach serial link in 28 nm CMOS for advanced packaging applications," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3206–3218, Dec. 2013.

[4] B. Dehlaghi and A. C. Carusone, "A 20 Gb/s 0.3 pJ/b single-ended die-to-die transceiver in 28 nm-SOI CMOS," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, Sep. 2015, pp. 1–4.

[5] J. F. Bulzacchelli *et al.*, "A 28-Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 3232–3248, Dec. 2012.

[6] M. Harwood *et al.*, "A 225 mW 28 Gb/s SerDes in 40 nm CMOS with 13dB of analog equalization for 100GBASE-LR4 and optical transport lane 4.4 applications," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 326–327.

[7] M.-S. Chen, Y.-N. Shih, C.-L. Lin, H.-W. Hung, and J. Lee, "A fully-integrated 40-Gb/s transceiver in 65-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, no. 3, pp. 627–640, Mar. 2012.

[8] H. Kimura *et al.*, "A 28 Gb/s 560 mW multi-standard SerDes with single-stage analog front-end and 14-tap decision feedback equalizer in 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 3091–3103, Dec. 2014.

[9] B. Raghavan *et al.*, "A sub-2 W 39.8–44.6 Gb/s transmitter and receiver chipset with SFI-5.2 interface in 40 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3219–3228, Dec. 2013.

[10] C.-L. Hsieh and S.-I. Liu, "A 40 Gb/s adaptive receiver with linear equalizer and merged DFE/CDR," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2011, pp. 208–209.
[11] A. Manian and B. Razavi, "A 40 Gb/s 14 mW CMOS wireline receiver," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2016, pp. 412–413.
[12] A. Manian and B. Razavi, "A 40-Gb/s 9.2-mW CMOS equalizer," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2015, pp. 226–227.
[13] 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.
[14] 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.
[15] M.-S. Chen and C.-K. K. Yang, "A 50-64 Gb/s serializing transmitter with a 4-tap, LC-ladder-filter-based FFE in 65 nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 8, pp. 1903–1916, Aug. 2015.
[16] L. Li and M. Green, "Power optimization of an 11.75-Gb/s combined decision feedback equalizer and clock data recovery circuit in 0.18-μm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 58, no. 3, pp. 441–450, Mar. 2011.
[17] P. Andreani and A. Fard, "More on the $1/f^2$ phase noise performance of CMOS differential-pair LC-tank oscillators," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2703–2712, Dec. 2006.

Abishek Manian received the B.E. degree in electronics and telecommunication from the University of Mumbai, India, in 2011, and the M.S. and Ph.D. degrees in electrical engineering from the University of California (UCLA), Los Angeles, CA, USA, in 2013 and 2016, respectively.

He is currently with the High-Speed Signal Conditioning Group, Texas Instruments Incorporated, Santa Clara, CA, USA.

Dr. Manian was a recipient of the Scholarship under Sir Ratan Tata Trust's Studies in India Program for 2008–2009, 2009–2010, and 2010–2011, the Sir Dorabji Tata Trust Travel Fellowship in 2011, the J.N. Tata Endowment Scholarship for Higher Studies in 2011, the Jamshedji Tata Trust Scholarship in 2012, the UCLA Henry Samueli Distinguished Fellowship Award 2013, the UCLA Graduate Division Fellowship for 2012–2013 and 2013–2014, the Best Student Paper Award for the 2015 Symposium on VLSI Circuits, the Dissertation Year Fellowship Award 2015–2016 at UCLA, and the Distinguished Ph.D. Dissertation Award in Circuits and Embedded Systems (2016) awarded by the EE Department, UCLA. He serves as a reviewer for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I AND II, and the IEEE International Symposium on Circuits and Systems.

Behzad Razavi (S'87–M'90–SM'00–F'03) is a member of the U.S. National Academy of Engineering. He received the B.S.E.E. degree from the Sharif University of Technology, Tehran, Iran, in 1985, and the M.S.E.E. and Ph.D.E.E. degrees 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 and Hewlett-Packard Laboratories until 1996. Since 1996, he has been an Associate Professor with the University of California, Los Angeles, CA, USA, where he is currently a Professor of Electrical Engineering. His current research includes wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters. He has authored *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; Wiley, 2012), and *Fundamentals of Microelectronics* (Wiley, 2006) (translated to Korean and Portuguese). He was the editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996), and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003).

Prof. Razavi was a recipient of the Beatrice Winner Award for Editorial Excellence at the 1994 International Solid-State Circuits Conference (ISSCC), the Best Paper Award at the 1994 European Solid-State Circuits Conference, the Best Panel Award at the 1995 and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, the Best Paper Award at the IEEE Custom Integrated Circuits Conference in 1998, the McGraw-Hill First Edition of the Year Award in 2001, the Lockheed Martin Excellence in Teaching Award in 2006, the UCLA Faculty Senate Teaching Award in 2007, the CICC Best Invited Paper Award in 2009 and 2012, the 2012 Donald Pederson Award in Solid-State Circuits, and the American Society for Engineering Education PSW Teaching Award in 2014. He was a co-recipient of the Jack Kilby Outstanding Student Paper Award, the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC, the 2012 and the 2015 VLSI Circuits Symposium Best Student Paper Awards, and the 2013 CICC Best Paper Award. He was also recognized as one of the top 10 authors in the 50-year history of ISSCC. He served on the Technical Program Committees of the ISSCC from 1993 to 2002 and the VLSI Circuits Symposium from 1998 to 2002. He has served as an IEEE Distinguished Lecturer. He has also served as Guest Editor and Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and the *International Journal of High Speed Electronics*.