

Design Techniques for High-Speed Wireline Transmitters

BEHZAD RAZAVI^{ID}

(Invited Paper)

Department of Electrical Engineering, University of California, 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 transmitters operating at tens of gigabits per second pose challenging design issues ranging from limited bandwidths to severe sensitivity to jitter. This paper presents a number of analog and digital circuit techniques that allow data rates as high as 80 Gb/s in 45-nm CMOS technology. A PAM4 prototype delivers an output swing of 630 mV_{pp} with a clock jitter of 205 fs_{rms} while drawing 44 mW.

INDEX TERMS SERDES, serial links, multiplexers, oscillators, phase noise, crystal oscillators, integrated jitter.

I. INTRODUCTION

WITH the dramatic rise of data transport over the Internet, wireline systems are pressed for increasingly higher speeds. Recent projections indicate that the data traffic climbs by 25% per year, possibly reaching 20 zettabytes (20×10^{21} bytes) in 2025 [1]. Also challenging the engineers is the matter of power consumption and how it impacts package and module design and heat removal.

Wireline transceivers have been under intense development for two decades [1]–[17], inheriting broadband concepts from optical communication circuits as well as dealing with other issues that are specifically related to copper media. This paper proposes circuit and architecture techniques that prove useful in the design of transmitters (TXs) operating at tens of gigabits per second. The methods are introduced in the context of a 40-Gb/s TX [11] and an 80-Gb/s TX [10], which have been developed in 45-nm CMOS technology.

Sections II and III provide a tutorial background on TX design. Section IV describes the transmitter architecture, and Section V the design of its building blocks. Section VI presents the experimental results.

II. BASIC PRINCIPLES

A wireline TX generally performs three functions (Fig. 1): it converts a large number of parallel, low-speed data streams to a single high-speed output (“serialization”), it subjects the data to equalization so as to partially compensate for the loss




FIGURE 1. Generic wireline transmitter.

of the channel through which the information is transmitted, and it delivers sufficient output swings to the channel. These operations require various clock frequencies and phases that are generated by a phase-locked loop (PLL). The output transistors are protected by electrostatic discharge (ESD) devices.

As we seek greater data rates, a number of trends emerge. First, the signal path in Fig. 1 requires a larger number of broadband stages, posing more difficult circuit design and signal distribution issues. For example, cascaded functions that use inductors dictate long interconnects. Second, as commonly practiced today, higher speeds are carried by PAM4 data, presenting more severe challenges than non-return-to-zero (NRZ) data does (Section III-B). Third, the equalizer inevitably poses a trade-off between the amount of channel loss that it can compensate and the attenuation that

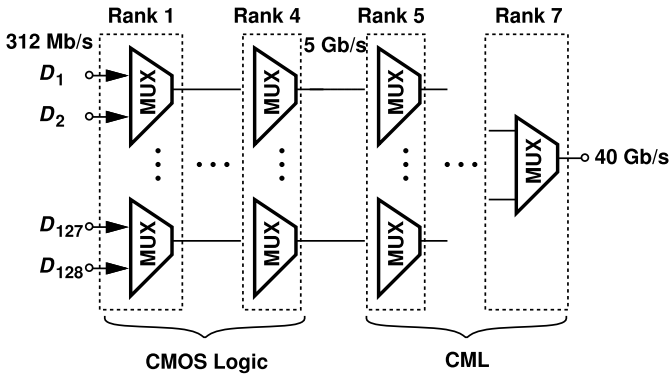


FIGURE 2. A binary-tree MUX.

it introduces in the output voltage swing (Section III-C). Fourth, the line driver becomes a serious bandwidth bottleneck because its output voltage and current swings are fairly unscalable and so are its transistor widths and capacitances (in a given process node) (Section III-D). Fifth, for a given level of protection, the ESD devices exhibit a certain capacitance that ultimately limits the output bandwidth even if the driver itself does not. Sixth, the numerous clocks reaching the high-speed stages lead to complex routing problems. Seventh, the PLL jitter must be commensurate with the TX output bit (or symbol) period (Section III-E). These trends imply that high-speed TX design must draw upon both advances in CMOS technology and new circuit techniques.

III. GENERAL CONSIDERATIONS

A. CHOICE OF LOGIC STYLES

The data-path functions illustrated in Fig. 1 require a great deal of logic, primarily in the form of multiplexers (MUXs) and flipflops (FFs). We expect that digital CMOS (rail-to-rail) realizations suffice for multiplexing up to a certain speed beyond which the delays and rise and fall times cause failure. For multiplexing to higher data rates, we can resort to current-mode logic (CML). For example, suppose 128 inputs at 312 Mb/s must be serialized to obtain a single 40-Gb/s data stream. Considering a binary-tree MUX structure (Fig. 2), we envision that the first four ranks can comfortably operate with rail-to-rail swings, generating a rate of 5 Gb/s, and the next three should employ CML. This boundary shifts to higher frequencies in more advanced process nodes.

A logic design style that operates faster than CMOS logic and draws less power than CML is based on “charge steering” [18]. Illustrated in Fig. 3, the basic structure resembles a differential pair but with the load resistors replaced with capacitors and the tail current source with a charge source. In the reset mode, nodes X and Y are precharged to V_{DD} while C_T is discharged to the ground. In the evaluation mode, X and Y are released and C_T switches to node P , drawing a current from M_1 and M_2 and hence from C_X and C_Y . This continues until C_T charges to approximately $V_{CM} - V_{TH}$, where V_{CM} denotes the input common-mode (CM) level. The output difference thus developed is proportional to V_{in} . We note that the circuit can act as an amplifier




FIGURE 3. Basic charge-steering stage.




FIGURE 4. Problem of divider and MUX delay.

and/or a latch. The capacitances are so chosen as to provide a moderate output swing, e.g., around 400 mV.

The signal swings used in charge steering allow it to support higher speeds than does CMOS logic. In addition, such swings reduce the power consumption by a theoretical factor of 1.4π with respect to CML stages [18]. One drawback of this design style is that $V_X - V_Y$ in Fig. 3 follows a return-to-zero (RZ) waveform, requiring that the circuits be properly architected. For TX design, charge steering proves particularly useful in realizing MUX stages that interface CMOS ranks to CML ranks (Section V-B).

B. NRZ AND PAM4 ISSUES

The binary nature of NRZ data requires a single signal path with, in principle, no need for linearity. Thus, NRZ signal swings are dictated by primarily speed and power considerations.

Certain delays become problematic near the TX front end. Consider the serializer shown in Fig. 4(a), where Rank n is driven by f_{CK} and Rank $n - 1$ by $f_{CK}/2$. We predict that the divider and MUX delays, ΔT_1 and ΔT_2 , respectively,




FIGURE 5. Distortion of PAM4 signal due to driver nonlinearity.




FIGURE 6. A basic PAM4 TX.

introduce timing issues. The waveforms in Fig. 4(b) reveal that D_a arrives at Rank n with a total skew of $\Delta T_1 + \Delta T_2$ with respect to f_{CK} . If this skew is a significant fraction of $T_{CK}/2$, Rank n may fail.

PAM4 data generation must deal with two additional issues. First, the stages processing PAM4 signals must be sufficiently linear so that they do not compress the “top” and “bottom” eyes (Fig. 5). The principal concern here is the loss of eye height and hence the higher error rate that the receiver may experience. This is quantified by the “ratio of level mismatch” (RLM), defined as the smallest eye height divided by one-third of the total eye height [20]. We typically target an RLM of greater than 95%.

The second PAM4 issue relates to the actual generation of the four-level waveform. We surmise that a 2-bit digital signal, D_2D_1 , applied to a digital-to-analog converter (DAC) yields a PAM4 output. The bits D_2 and D_1 must be generated by independent most-significant bit (MSB) and least-significant bit (LSB) paths (Fig. 6). The need for two serializers leads to greater complexity and power consumption than those of NRZ transmitters.

The MSB and LSB serializers in Fig. 6 cannot be identical as they drive different DAC input capacitances. Since $C_2 \approx 2C_1$, we expect that, at least, the last stage in the MSB path must have twice the strength of its LSB counterpart so that the D_2 and D_1 transitions occur at approximately the same time. Without such a precaution, a skew arises between D_2 and D_1 , creating jitter at the DAC output.

C. EQUALIZER

The equalizer in Fig. 1 provides partial compensation for the loss of the channel. As conceptually illustrated in Fig. 7(a), this circuit ideally provides a frequency response, $|H_{eq}|$, that is the inverse of the channel’s, $|H_{ch}|$, so that the cascade exhibits a flat passband. In practice, however, only losses less than about 6 dB can be accommodated. The implementation

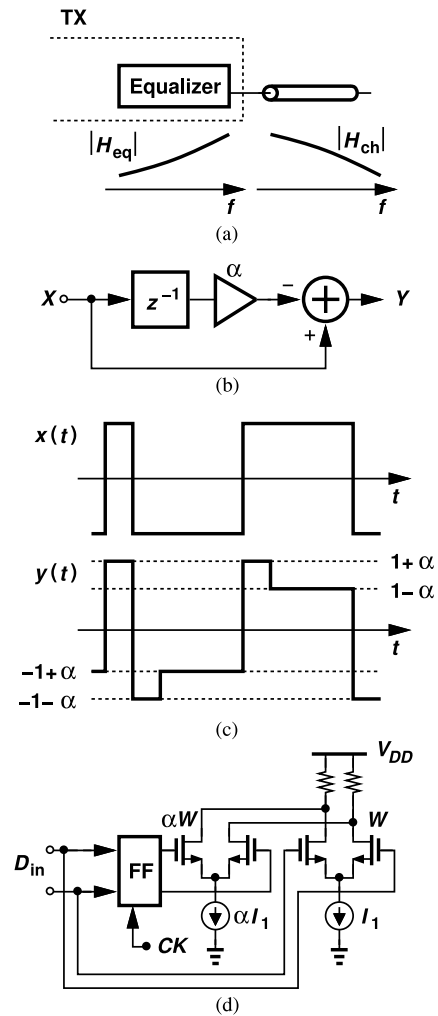


FIGURE 7. (a) Equalizer and channel frequency response, (b) an FFE stage, (c) effect of equalization on data, and (d) FFE circuit implementation example.

is shown in Fig. 7(b) and called the “feedforward equalizer” (FFE). The circuit delays the input by one clock cycle (one bit period), scales it by a factor of $\alpha < 1$, and subtracts the result from X . Characterized by $Y/X = 1 - \alpha z^{-1}$, this topology approximates a differentiator and hence a high-pass filter. Since $z = \exp(j2\pi f T_{CK})$, it can be shown that $|Y/X|$ reaches a peak value of $1 + \alpha$, i.e., the high-frequency content is amplified by a factor of $1 + \alpha$. But we observe that, if X varies slowly in the time domain, then its delayed copy is approximately equal to itself and $y(t) \approx (1 - \alpha)x(t)$. That is, the low-frequency content (representing the “dc swings”) is attenuated.

This effect can also be seen in the time-domain waveforms shown in Fig. 7(c): the output amplitude jumps to $1 + \alpha$ immediately after a transition but drops to $1 - \alpha$ for a consecutive sequence of ONES or ZEROS. The swing reduction translates to additional challenges in receiver design. Figure 7(d) shows a basic FFE realization where the differential pair driven by the delayed data is scaled by a factor of α with respect to the other one.

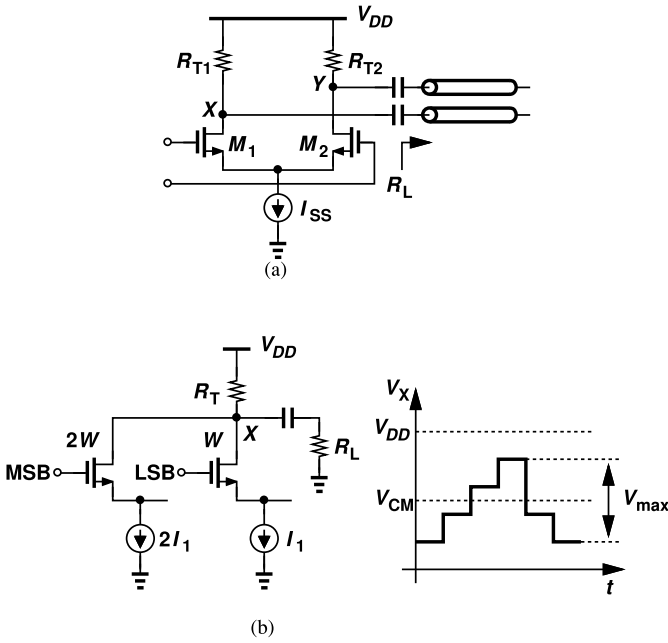


FIGURE 8. (a) NRZ CML driver, and (b) PAM4 CML driver.

D. LINE DRIVER

The most challenging building block in TX design is typically the line driver as it faces the most difficult demands of the standard: it must deliver high current levels to the channel while meeting the bandwidth requirements.

Consider the CML driver shown in Fig. 8(a), where back-termination resistors R_{T1} and R_{T2} , both having a value of R_T , minimize the effect of reflections from the channel. Suppose we wish to create single-ended voltage swings at X and Y equal to 400 mV_{pp} . Since R_{T1} and R_{T2} are chosen approximately equal to the channel's single-ended characteristic impedance, $R_L \approx 50 \Omega$, we must select a value of at least 16 mA for I_{SS} . In addition to burning high power, such a current dictates large widths for M_1 and M_2 , thereby introducing substantial capacitance at the input and output of the driver. This issue in turn requires the use of various inductive and T-coil peaking techniques in both the stage preceding the driver and in the driver output nodes [21]. We express the power consumption as $V_{DD} \cdot I_{SS} = V_{DD}(2V_{max}/R_L)$, where V_{max} denotes the single-ended peak-to-peak output swing.

One can alleviate the foregoing issues by selecting R_{T1} and R_{T2} in Fig. 8 to be somewhat greater than their ideal values. For example, a value of 75Ω still reduces the reflections but allows lesser tail currents for a given output swing [23]. However, the lower output CM level may degrade the circuit's speed.

For PAM4 signaling, the situation becomes more severe. The CML driver depicted in Fig. 8(b) incorporates an MSB and an LSB branch to generate a single-ended peak-to-peak swing of $V_{max} = 3I_1(R_T || R_L) = 3I_1R_L/2$ around a common-mode level, $V_{CM} = V_{DD} - 3I_1R_T/2 = V_{DD} - 3I_1R_L/2$. One difficulty here is that the CM level is given by the

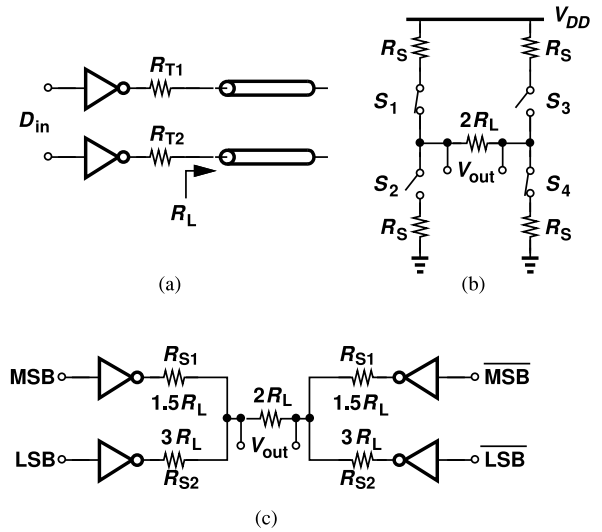


FIGURE 9. (a) NRZ SST driver, (b) its equivalent circuit, and (c) PAM4 SST driver.

tail currents and R_T whereas the output voltage swing is defined by $R_T || R_L$. The low output CM level tends to push the transistors into the triode region and degrade the linearity. The minimum supply voltage is given by $V_{DD,min} = 3I_1R_L/4 + V_{max} + V_{DS} + V_{tail}$, where V_{DS} and V_{tail} denote the minimum drain-source voltages necessary for the output transistors and the tail currents, respectively. We then have $V_{DD,min} = 1.5V_{max} + V_{DS} + V_{tail}$, obtaining a minimum of $I_{SS}(1.5V_{max} + V_{DS} + V_{tail})$ for the driver's power consumption.

In comparison to CML topologies, voltage-mode structures, also called "source-series termination" (SST) circuits [24], draw substantially less power. Depicted in Fig. 9(a) is an example for differential NRZ data, where the on-resistance, R_{on} , of the transistors within the inverters plus R_{T1} or R_{T2} is equal to R_L . Denoting $R_{on} + R_{T1}$ and $R_{on} + R_{T2}$ by R_S , we recognize from the equivalent circuit shown in Fig. 9(b) that the driver provides a peak-to-peak differential output voltage swing equal to V_{DD} . Moreover, the class-D action reduces the power to $V_{DD}[V_{DD}/(4R_L)] = V_{DD}[V_{max}/(2R_L)]$, a factor of 4 lower than that found for the CML NRZ driver studied above.

The SST topology of Fig. 9(a) faces two drawbacks. First, unlike the differential pair in Fig. 8(a), the inverters draw a large transient current from V_{DD} during data transitions, demanding a heavy bypass capacitance to minimize supply bounce. Second, at very high speeds, it is difficult to generate the rail-to-rail input swings necessary for the inverters.

SST operation can be extended to PAM4 signaling as well. Shown in Fig. 9(c) is a realization where the transistor on-resistances are included in R_{S1} and R_{S2} . The net back-termination is equal to R_L and the peak-to-peak differential output swing is equal to V_{DD} . The drawbacks mentioned above apply here as well. Furthermore, if R_{on} is a significant fraction of the back-termination resistance, then its voltage dependence translates to nonlinearity.




FIGURE 10. Spur levels corresponding to deterministic jitter.

E. PLL JITTER

The random and deterministic jitters generated by the PLL in Fig. 1 directly corrupt the transmitted data. As a rule of thumb, we wish to keep the rms value of each below roughly one-hundredth of the bit or symbol period. The reason for this bound is explained below.

For a bit period of, say, 25 ps, we target a random jitter of less than 250 fs_{rms}. The PLL's voltage-controlled oscillator (VCO) phase noise budget is determined by both this constraint and the reference phase noise, S_{REF} . It can be shown that the optimum loop bandwidth, f_{BW} , makes the reference and VCO jitter contributions approximately equal [25]. That is, no more than $250 \text{ fs}/\sqrt{2}$ must arise from the reference. Assuming a one-pole transfer function for the PLL,¹ we write the integrated phase noise as $2(\pi/2)S_{REF}f_{BW} = \pi S_{REF}f_{BW}$, where the factor of 2 is included if S_{REF} represents the phase noise on only one side of the carrier, and the factor of $\pi/2$ originates from the one-pole model. This contribution is bounded according to

$$\sqrt{\pi S_{REF}f_{BW}} \frac{T_{REF}}{2\pi} \leq \frac{250 \text{ fs}}{\sqrt{2}}, \quad (1)$$

where T_{REF} is the reference period. If, for example, $S_{REF} = -150 \text{ dBc/Hz}$ for a 312-MHz crystal oscillator, we have $f_{BW} \approx 38 \text{ MHz}$, slightly greater than $f_{REF}/10$. This result does not quite agree with the one-pole response, indicating that the actual jitter is higher. In reality, a loop bandwidth of about 20 MHz is necessary in this example [10].

The deterministic jitter in PLLs results from periodic modulation of the output frequency—primarily by the reference—and is quantified as follows. If the VCO output spectrum contains sidebands at $\pm f_{REF}$ around the carrier and their normalized amplitude is denoted by β (Fig. 10), we express the output in the time domain as follows:

$$V_{out}(t) = A_0 \cos(\omega_0 t + 2\beta \cos \omega_{REF} t), \quad (2)$$

where $\omega_0 = 2\pi f_0$ and $\omega_{REF} = 2\pi f_{REF}$. The phase modulation term signifies a sinusoidal jitter having a peak value of 2β radians and hence an rms value of $\sqrt{2}\beta/\omega_0$ seconds. For this jitter to be about one-hundredth of the period, $T_0 = 2\pi/\omega_0$, we require that $\beta < \sqrt{2}\pi/100 \equiv -27 \text{ dB}$, a relaxed constraint.

In order to combine the effects of random and deterministic jitter, we can add the squares of their rms values as the two phenomena are uncorrelated. But if we wish to estimate the peak-to-peak jitter, J_{pp} , and hence the horizontal

1. A reasonable model if $f_{BW} < f_{REF}/20$.




FIGURE 11. Illustration of TX phase noise removed by tracking action of RX CDR.

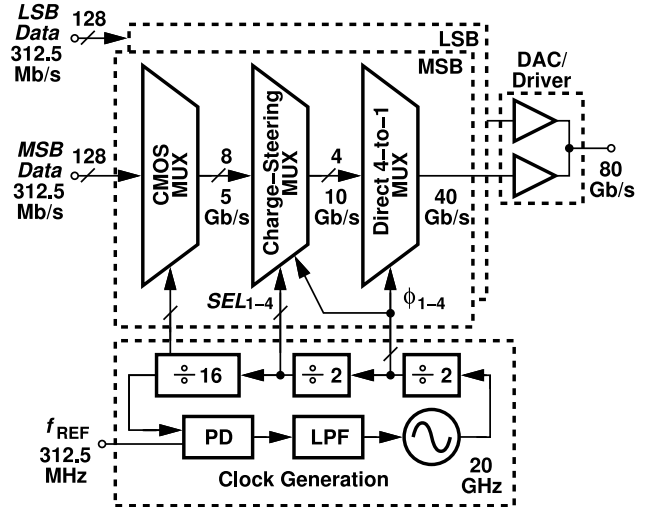


FIGURE 12. Proposed TX architecture.

eye closure at the TX output, we write

$$J_{pp} \approx 6\sigma_r + \frac{4\beta}{\omega_0}, \quad (3)$$

where σ_r denotes the rms random jitter. If σ_r and $\sqrt{2}\beta/\omega_0$ are around $T_0/100$, we have $J_{pp} \approx (6 + 2\sqrt{2})T_0/100 \approx 8.5\%T_0$, a reasonable amount of eye closure due to only the PLL jitter.

We should make two additional remarks. First, since typical PLLs exhibit sidebands well below -40 dBc , the effect of deterministic jitter is negligible, leaving a greater budget for the random component. Second, the actual tolerable random jitter at a TX output may be higher than what we have assumed. This is because wireline standards recognize that the low-frequency phase noise components produced by the TX PLL are “tracked out” by the clock and data recovery (CDR) circuit in the receiver (Fig. 11), requiring that only the high-frequency content of the PLL phase noise be taken into account.

IV. 80-Gb/s PAM4 TX ARCHITECTURE

The high-speed circuit techniques to be presented here are employed in an 80-Gb/s PAM4 TX [10]. Figure 12 shows the proposed architecture. The MSB and LSB data paths consist of a 128-to-8 CMOS MUX, an 8-to-4 charge-steering MUX, a 4-to-1 “direct” CML MUX, and a 2-bit DAC acting as the line driver. A PLL generates the clocks necessary for multiplexers from a 312.5-MHz reference. As explained below,

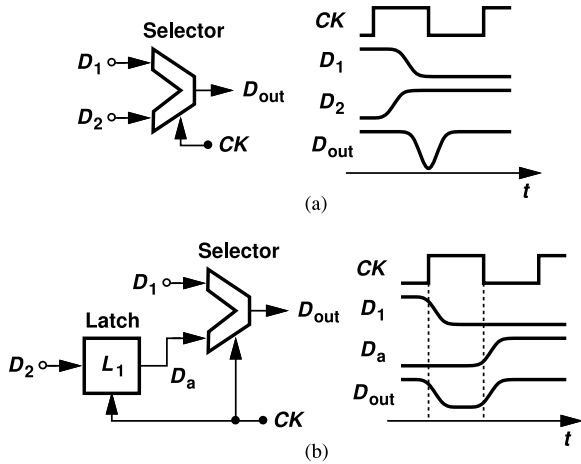


FIGURE 13. (a) A 2-to-1 selector, and (b) use of a latch to avoid simultaneous data transitions.

the use of various clock phases can dramatically improve the serializers' performance, but it is feasible if the PLL can deliver such phases. Specifically, the PLL feedback dividers provide quadrature phases, ϕ_1 - ϕ_4 , with a duty cycle of 25%, 45° phases, select commands SEL_1 - SEL_4 , etc., making it possible to avoid latches in serializer design (Section V-A).

As explained in Section III-B, the MSB and LSB paths in Fig. 12 must provide 2:1 drive strengths, respectively, so as to avoid a systematic skew between the MSB and LSB waveforms arriving at the DAC inputs. For this reason, the direct 4-to-1 MUX in the MSB path is scaled up by a factor of 2 with respect to its counterpart in the LSB path.

The prototype described here does not include FFE action. As explained in [10], the FFE method in [11] can also be applied to this TX.

V. DESIGN OF BUILDING BLOCKS

In this section, we study the transistor-level design of the TX building blocks, including the CMOS, charge-steering, and 4-to-1 multiplexers, the line driver, and the PLL.

A. CMOS MUX DESIGN

The serializer in Fig. 1 must typically employ a large number of latches and selectors so as to aggregate more than 100 input data streams. The chain in Fig. 2, for example, requires about 2^7 MUX cells. The principal issue here is the power consumption associated with the serializer's clock path. The challenge becomes more severe in the dual-path PAM4 TX shown in Fig. 12.

In order to arrive at a standard MUX cell design, we begin with the 2-to-1 selector depicted in Fig. 13(a). For simplicity, suppose the structure consists of two differential pairs that sense D_1 and D_2 and are controlled by CK . The difficulty here is that the output can exhibit excessively narrow pulses or glitches if the input transitions occur at arbitrary times. This effect is avoided if the selector inputs are guaranteed to change at different times, which is possible if D_1 or D_2 is delayed by a flipflop or a latch. Shown in Fig. 13(b) is an

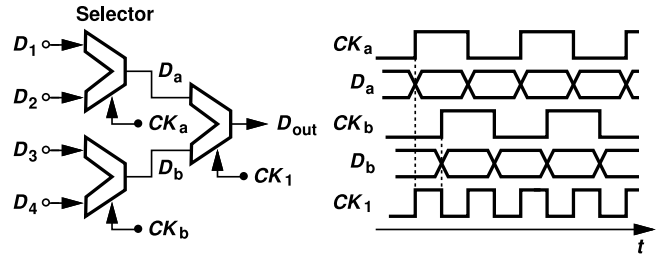


FIGURE 14. Proposed MUX cell.

example [26] where the latch and the selector are controlled by CK such that, when the former is in the sense mode, the latter selects D_1 . When CK goes low, the latch enters the store mode and the selector reads D_a . The output can still suffer from narrow pulses if the transitions in D_1 occur close to the falling edges of CK —unless D_1 has a well-defined timing relationship with respect to CK .

Even with a single-latch MUX cell, a PAM4 serializer contains hundreds of latches and selectors. In the architecture of Fig. 2, the number of 2-to-1 MUXs drops by a factor of 2 from one rank to the next, but the increase in speed at least doubles the power consumed by the MUX cells.

It is possible to architect the serializer so that it utilizes *no latches*, thereby reducing the complexity and power considerably. We first recognize that, as shown in Fig. 4, the clock for each lower rank is generated by dividing the clock frequency of the higher rank by 2. We can thus utilize the quadrature clock phases provided by the $\div 2$ stages [10]. Illustrated in Fig. 14, the idea is to drive two selectors in the same rank by quadrature phases CK_a and CK_b so that D_a changes only on the edges of CK_a , and D_b on the edges of CK_b . This means that the inputs to the next selector are properly offset in time, avoiding glitches in D_{out} . This three-cell topology acts as a 4-to-1 MUX and can be repeated to form a complete serializer.

The 2-to-1 selector cell in Fig. 14 can be realized by CMOS logic for speeds up to about 5 GHz. To minimize the power consumption in its clock path, we prefer to employ small transistors. Figure 15(a) depicts a simple, efficient topology based on complementary CMOS (C^2 MOS) logic and Fig. 15(b) its simulated output eye diagram at 5 Gb/s [10]. This structure occupies a small area, allowing short interconnects for the entire CMOS serializer.

The CMOS serializer design begins with the last 2-to-1 selector, which must provide enough strength to drive the charge-steering MUX in Fig. 12. This C^2 MOS selector employs PMOS and NMOS widths equal to 2 μm and 1 μm , respectively, with a channel length of 40 nm, and hence draws 22 μW . Since the stages preceding this selector operate at progressively lower frequencies, the 2-to-1 selector is scaled down by a factor of 2 from one MUX rank to the rank preceding it, until a minimum allowable transistor width of 120 nm is reached (Fig. 16). The entire 128-to-8 serializer draws 365 μW in the data path.

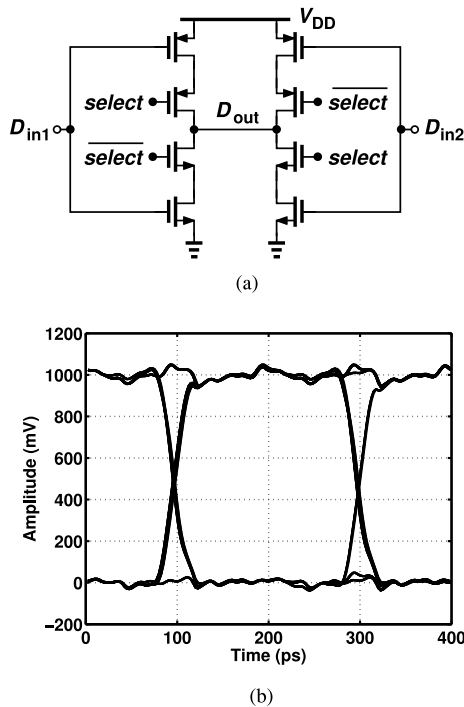


FIGURE 15. (a) Selector cell, and (b) its simulated output eye diagram at 5 Gb/s.




FIGURE 16. Tapering in CMOS MUX chain.

B. CHARGE-STEERING MUX DESIGN

For operation above 5 Gb/s in 40-nm CMOS technology, charge steering proves more viable than CMOS logic. In this spirit, we wish to apply this concept to the 8-to-4 MUX in Fig. 12.

The charge-steering stage of Fig. 3 can be readily extended to form a selector. Illustrated in Fig. 17, the result senses the inputs by means of two differential pairs and performs the selection by enabling the tail path in one. As the waveforms demonstrate, V_X and V_Y are precharged to V_{DD} when CK is low and C_T is discharged. After CK goes high, depending on the logical value of SEL , the output responds to V_{in1} or V_{in2} , allowing V_X or V_Y to fall. Note that the rail-to-rail swings arriving from the preceding C^2MOS MUX ensure that the selected differential pair steers the tail charge completely. In this topology, CK runs at twice the SEL frequency, which itself is equal to the input data rate (5 Gb/s).

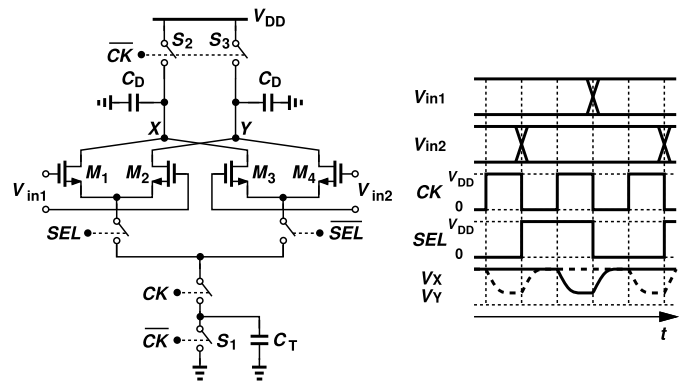


FIGURE 17. Basic charge-steering MUX.

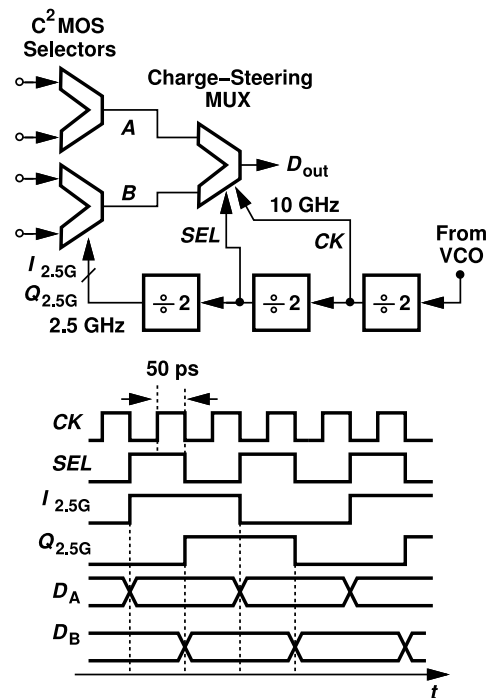


FIGURE 18. Use of 10-GHz, 5-GHz, and 2.5-GHz clock phases to drive the CMOS and charge-steering multiplexers.

The charge-steering MUX of Fig. 17 entails a number of issues. First, its data inputs must make transitions only on clock edges, a condition fulfilled by the last C^2MOS selector's clocking. As seen in Fig. 18, the in-phase (I) and quadrature (Q) components of the 2.5-GHz clock produce the 5-Gb/s data streams at A and B (V_{in1} and V_{in2} in Fig. 17, respectively) with properly positioned edges. The 5-GHz select command enables one differential pair around the transition times of A or B . Also, the 10-GHz clock, CK , precharges the charge-steering MUX for 50 ps before the select command changes. The MUX therefore has 50 ps for evaluation.

The second issue is that the MUX in Fig. 17 generates high levels at both X and Y in its precharge mode and its output must not be sensed by the next MUX during this

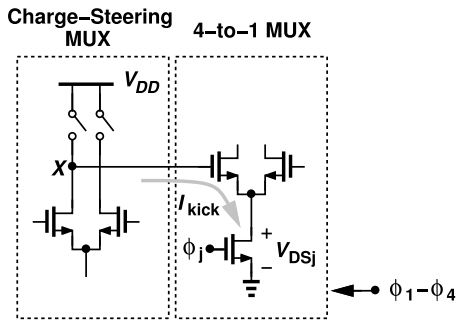


FIGURE 19. Kickback noise injected onto charge-steering MUX by 4-to-1 MUX.

time. This is guaranteed in the 4-to-1 MUX by means of clocks having a 25% duty cycle.

The third issue relates to the kickback noise of the 4-to-1 MUX. Depicted in Fig. 19, this effect occurs on the edges of this MUX's 10-GHz select commands, ϕ_1 - ϕ_4 , and drops the CM level at X and Y by more than 100 mV. This fall in turn causes the tail current sources in the 4-to-1 MUX to collapse. We resolve the difficulty by changing the charge-steering MUX's differential pairs to complementary input stages [Fig. 20(a)]. With rail-to-rail inputs, the PMOS devices also switch completely, pinning either X or Y to V_{DD} . Plotted in Fig. 20(b) are V_X and V_Y before and after the PMOS transistors are added, displaying less variation in their CM level in the presence of the pull-up devices.

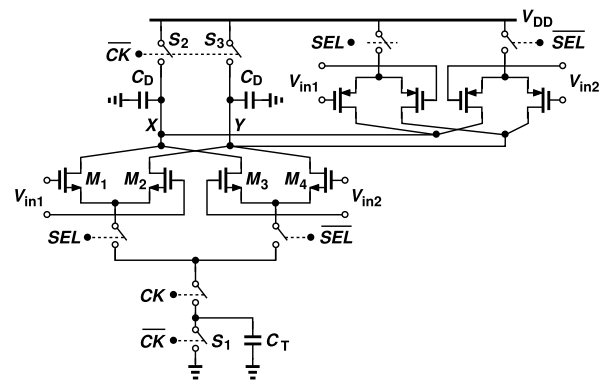
The fourth issue concerns the skew between the SEL and CK commands in Fig. 17. Owing to the divider delay in Fig. 20, SEL arrives slightly later than CK does. This effect is benign as it does not interfere with charge steering. On the other hand, this delay also means that the circuit enters the precharge mode before SEL changes, again a benign situation as the tail is disabled by CK .

C. CML MUX DESIGN


In the TX presented here, the charge-steering MUX delivers four 10-Gb/s data streams, which must next be multiplexed to reach a single 40-Gb/s output. A binary-tree CML topology would then necessitate at least three latches and three selectors. These amount to 12 tail current sources for the MSB and LSB paths in Fig. 12, drawing high power.

Another challenge in a high-speed binary-tree structure is that it can fail due to the skew illustrated in Fig. 4. Recall that we wish to maintain $\Delta T_1 + \Delta T_2$ well below $T_{CK}/2 = 25$ ps, an unrealistic goal in 45-nm technology in view of the layout parasitics and the finite clock transition times.

The two foregoing issues are ameliorated by means of a direct 4-to-1 MUX, depicted in Fig. 21. The four differential pairs are enabled in succession by clocks having a 25% duty cycle. The output therefore tracks each of the inputs for 25 ps. Inductive peaking extends the bandwidth in the presence of the large input capacitance of the next stage (the DAC) and the drain capacitance of the four differential pairs. We observe that this topology contains only one active tail



(a)



(b)

FIGURE 20. (a) Modified charge-steering MUX, and (b) its output waveforms.

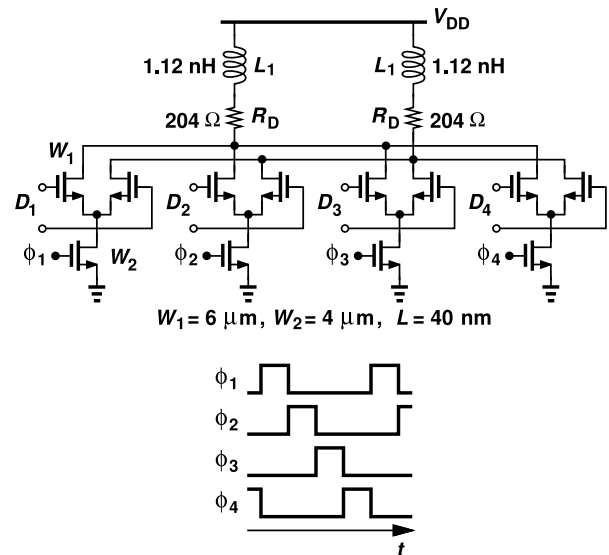


FIGURE 21. Direct 4-to-1 MUX.

current. Moreover, the skew now must be well below 50 ps rather than 25 ps.

Direct 4-to-1 MUX topologies have been reported [26], but our approach merits two remarks. First, the prior art employs stacked tail transistors that are driven by overlapping phases

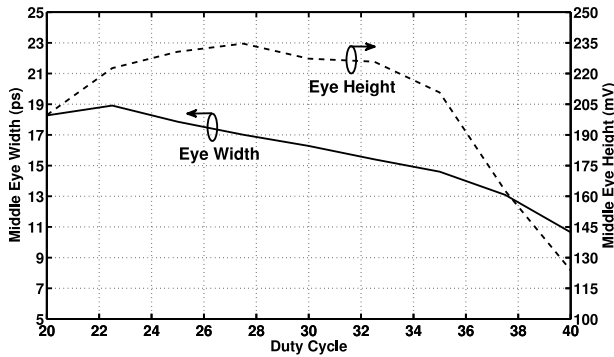


FIGURE 22. PAM4 vertical and horizontal eye openings as a function of clock duty cycle.

having a 50% duty cycle [27]. The stacking degrades the tail current waveforms at high speeds. We instead implement the select path by a single tail transistor and rely on a new $\div 2$ circuit that directly provides a 25% duty cycle. Second, in Fig. 21, W_2 need be only $4 \mu\text{m}$ whereas in a stacked structure it must be twice as wide. The clock path power consumption would then rise by a factor of 4.

For power efficiency, our 4-to-1 MUX does not incorporate current sources; rather, it drives the tail transistor gates by rail-to-rail swings. This means that the MUX output voltage swing depends, to some extent, on the process, supply voltage, and temperature. Nonetheless, this variation can be tolerated so long as the worst-case output swing is still sufficient to ensure complete switching in the following stage, namely, the DAC.

The nonoverlapping clocks, ϕ_1 - ϕ_4 , in Fig. 21 are directly generated by a divide-by-2 stage. The divider design is described in Section V-F.

D. DIRECT 4-TO-1 MUX ISSUES

As noted in the previous section, direct multiplexing offers certain advantages over binary-tree realizations. However, this multiplexer's data is not retimed as it travels to the TX output. In other words, any edge misalignment in the 4-to-1 MUX output accompanies the transmitted data. We address two issues related to this absence of retiming.

First, given that the output of the MUX in Fig. 21 tracks one input when a tail transistor is turned on, we ask how departures in the clock duty cycle from 25% affect the output. To quantify the ultimate impact of such errors, we examine the PAM4 data eye generated by the DAC. Shown in Fig. 22 are the simulated width and the height of the middle PAM4 eye as a function of the duty cycle of ϕ_1 - ϕ_4 . It is interesting to note that the width prefers about 23% and height about 28%, but some variability is tolerable. This point requires that the circuit delivering ϕ_1 - ϕ_4 and its layout parasitics be carefully simulated.

Second, the 4-to-1 MUX produces jitter at its output due to both duty cycle mismatches and delay mismatches among ϕ_1 - ϕ_4 [11]. Illustrated in Fig. 23(a), the former random mismatches can be represented by ΔT_{H1} - ΔT_{H4} , where

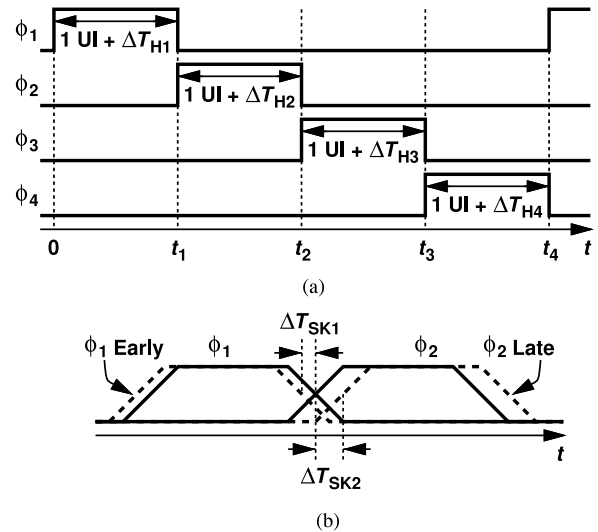


FIGURE 23. Effect of (a) duty cycle mismatch, and (b) delay mismatch on MUX output waveform.

$\Delta T_{H1} + \dots + \Delta T_{H4} = 0$. We observe that the rising edge of ϕ_2 at $t = t_1$ is displaced by ΔT_{H1} , that of ϕ_3 at $t = t_2$ by $\Delta T_{H1} + \Delta T_{H2}$, etc. Thus, the peak-to-peak jitter at the MUX output can be expressed as

$$J_{pp} = \max(\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4) - \min(\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4), \quad (4)$$

where $\epsilon_1 = \Delta T_{H1}$, $\epsilon_2 = \Delta T_{H1} + \Delta T_{H2}$, etc.

The effect of delay mismatches is depicted in Fig. 23(b), where we assume the falling edge of ϕ_1 incurs an error of ΔT_{sk1} and the rising edge of ϕ_2 , and error of ΔT_{sk2} . In this case, the MUX's differential output suffers from a zero-crossing displacement equal to $(\Delta T_{sk1} + \Delta T_{sk2})/2$. Extending this result to all four phases yields

$$J_{pp} = \max(\delta_1, \delta_2, \delta_3, \delta_4) - \min(\delta_1, \delta_2, \delta_3, \delta_4), \quad (5)$$

where $\delta_1 = (\Delta T_{sk1} + \Delta T_{sk2})/2$, $\delta_2 = (\Delta T_{sk2} + \Delta T_{sk3})/2$, etc.

The foregoing random mismatches can be quantified by running Monte Carlo simulations on the extracted layout of the 4-to-1 MUX, the frequency divider providing ϕ_1 - ϕ_4 , and the charge-steering MUX. But we also wish to measure the effect of these mismatches in the laboratory, a task more easily carried out in the frequency domain. Since duty cycle and delay mismatches repeat every clock cycle (with $T_{CK} = 100$ ps) [Fig. 24(a)], we must devise a test that displays this periodicity. This is accomplished by assigning a static pattern to the TX low-speed inputs so that the final output is nominally periodic. In our TX example, we can generate a 0101 NRZ sequence at the output with a frequency of 20 GHz. The mismatches modulate the phase of this waveform at a rate of 10 GHz, thereby yielding spurs at ± 10 GHz around the carrier [Fig. 24(b)]. From the results in Section III-E, we conclude that the peak jitter is equal to 2β radians.




FIGURE 24. (a) Generation of a periodic waveform at TX output, and (b) effect of mismatches in the frequency domain.




FIGURE 25. DAC/line driver circuit.

E. LINE DRIVER

The output DAC in Fig. 12 combines the 40-Gb/s MSB and LSB data streams while also acting as a line driver. As such, it bears the greatest speed burden in the entire TX. Shown in Fig. 25, the circuit incorporates three identical differential pairs having a unit tail current of 4.3 mA.

The DAC design merits two remarks. First, the 300-pH inductors provide series peaking in the presence of the DAC output capacitance (≈ 73 fF) and the pad and ESD capacitance (≈ 50 fF). The series inductors simplify the layout because they serve as part of the routing to the pads. In practice, larger ESD devices embedded in a T-coil can be used [22].

Second, the finite output resistance of the differential pairs in Fig. 25 generally translates to nonlinearity. This effect is particularly pronounced at the extremes of PAM4 swings because the transistors that are turned on reside in the triode region. Nevertheless, it can be shown that the nonlinearity is still sufficiently small for PAM4 signal generation [10].




FIGURE 26. PLL architecture.


F. CLOCK GENERATION

As explained in Section IV, the transmitter relies on quadrature and 45° clock phases with 25% or 50% duty cycles. The generation and distribution of these phases play a critical role in the overall TX performance.


The PLL architecture is shown in Fig. 26. Unlike conventional topologies, this work realizes the phase detector (PD) as an exclusive OR (XOR) gate and a master-slave sampling filter (MSSF) [28]. Eliminating the phase/frequency detector and the charge pump, the PLL potentially achieves lower phase noise. The master-slave action also offers a wide capture range, obviating the need for frequency acquisition [28]. With $f_{REF} = 312.5$ MHz and a loop bandwidth of 20 MHz, the LC VCO phase noise requirement is greatly relaxed, allowing an oscillator power consumption of only 3.5 mW for a free-running phase noise of -119 dBc/Hz at 10-MHz offset.

The most critical clocks in the TX of Fig. 12 are ϕ_1 - ϕ_4 as their mismatches directly introduce jitter. To generate such waveforms, we have three options: (1) employ a 10-GHz quadrature LC VCO to create overlapping clocks and use AND gates to change the duty cycle to 25%, (2) apply the output of a 20-GHz differential VCO to a conventional $\div 2$ circuit and AND gates, or (3) apply the output of a 20-GHz differential VCO to a $\div 2$ circuit that inherently delivers clocks with a 25% duty cycle. We pursue the last method here for it is potentially more efficient.


The LC VCO in Fig. 26 is followed by a $\div 2$ circuit to generate the nonoverlapping clocks necessary for the direct 4-to-1 MUX of Fig. 21. Before introducing the new divider topology, we consider the circuit shown in Fig. 27(a) [19]. From the waveforms in Fig. 27(b), we note that each output voltage is high for about one-half cycle of the input clock, providing a duty cycle close to 25%. In reality, the duty cycle is 25% plus one gate delay. Moreover, the logical low level is slightly degraded for part of the cycle because one PMOS pull-up transistor and one input coupling transistor conduct simultaneously before the latter turns off. These issues are resolved in the latch shown in Fig. 27(c), where transistors M_c and M_d are driven by CK to reduce the transition delay at the output, and transistors M_a and M_b cut the path from V_{DD} to ground. The series two PMOS devices, however, degrade the speed. We thus change all of the transistors to their opposite type, arriving at the latch shown in Fig. 28(a). The simulated waveforms in Fig. 28(b) reveal nonoverlapping



(a)



(b)



(c)

FIGURE 27. (a) Latch with 25% duty cycle output, (b) its output waveforms, and (c) modified version.


phases with a duty cycle of 75%. After inversion by buffers, the duty cycle changes to 25%.

The second $\div 2$ stage in Fig. 12 runs at 5 GHz but it is driven by a duty cycle of 25%. For this divider to generate quadrature phases, we introduce the ring counter shown in Fig. 29(a), which exploits ϕ_1 - ϕ_4 . Each latch is implemented as depicted in Fig. 29(b), where the cross-coupled inverters guarantee differential operation.²

VI. EXPERIMENTAL RESULTS


The PAM4 TX has been fabricated in TSMC's 45-nm CMOS technology. Shown in Fig. 30 is the die photograph; the active area is about $330 \mu\text{m} \times 320 \mu\text{m}$. The die has been directly mounted on a printed-circuit board and tested on a high-speed probe station. All of the measurements are carried out with a 1-V supply.

2. While the ring resembles an injection-locked divider, complete switching in the latches ensures a lock range extending to very low frequencies.




$$W_{1-2} = 12 \mu\text{m}, W_{5-6} = 9 \mu\text{m}, W_{c-d} = 3 \mu\text{m}, \\ W_{3-4} = 15 \mu\text{m}, W_{a-b} = 8 \mu\text{m}, L = 40 \text{ nm}.$$

(a)




(b)

FIGURE 28. (a) Latch used in 20-GHz divide-by-2 circuit, and (b) its output waveforms.



(a)



(b)

FIGURE 29. (a) Ring counter driven by clocks having a 25% duty cycle, and (b) C^2 MOS latch used in the divider.

The TX power breakdown is shown in Table 1. We observe that the line driver and the divider chain along with its clock distribution buffers constitute the most power-hungry functions.

Figure 31 depicts the measured TX output in the NRZ mode at 40 Gb/s and Fig. 32 plots the PAM4 waveforms at 40 Gb/s and 80 Gb/s. The differential voltage swing is




FIGURE 30. TX die photograph.

TABLE 1. TX power breakdown.

Blocks		Power (mW)
Data Path (MSB + LSB)	Output Driver/DAC	13.72
	CML MUX	5.66
	Charge-steering MUX	1.61
	CMOS MUX	0.73
Clock Path	Divider Chain and Buffers	18.25
	XOR + MSSF + Nonoverlap Gen.	0.62
	VCO	3.46
Total		44.05

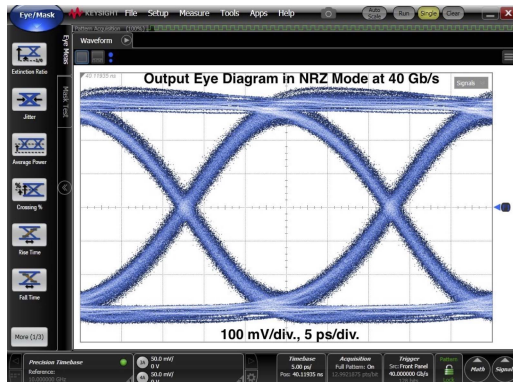
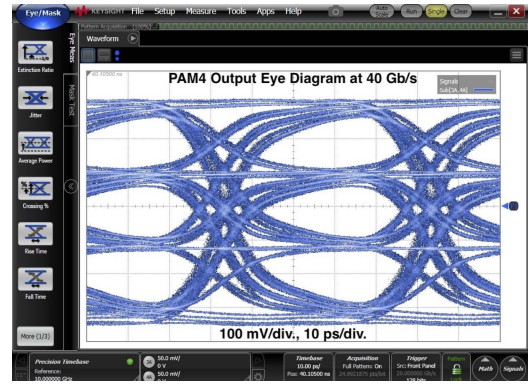


FIGURE 31. Measured output NRZ eye diagram at 40 Gb/s.

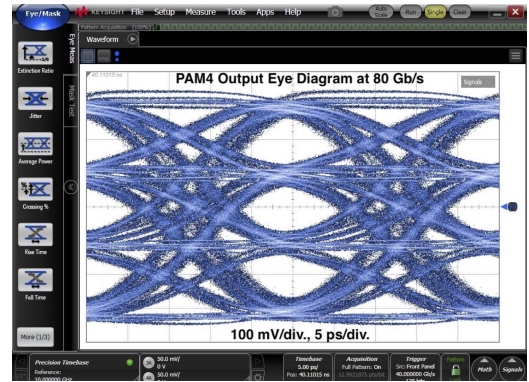
630 mV_{pp}, with a vertical eye opening of 170 mV. The horizontal openings are 0.56 unit intervals (UIs) for the middle eye and 0.43 UI for the top and bottom eyes. If the line driver's supply is raised to 1.2 V and the total tail current to 24 mA, the output swing reaches 1.2 V_{pp}.

As explained in Section III-B, the PAM4 waveform linearity is quantified by the RLM. In this measurement, the output contains 10 symbols, each lasting for 16 UI [20]. Our RLM is about 99%.

The 20-GHz clock generated by the PLL has also been characterized. Plotted in Fig. 33 is the measured spectrum, revealing a loop bandwidth of about 20 MHz. The reference spurs are at -45 dBc, higher than expected but still



(a)



(b)

FIGURE 32. Measured PAM4 output eye diagram at (a) 40 Gb/s, and (b) 80 Gb/s.


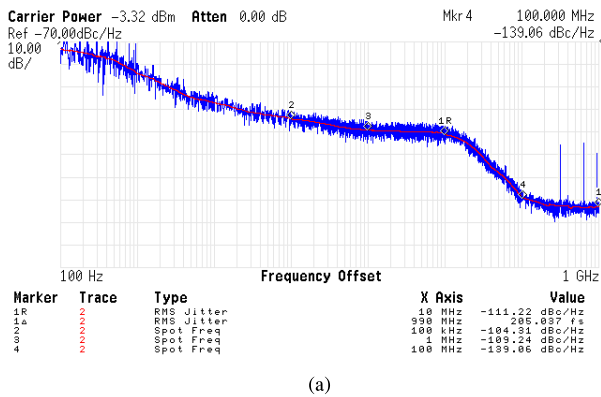
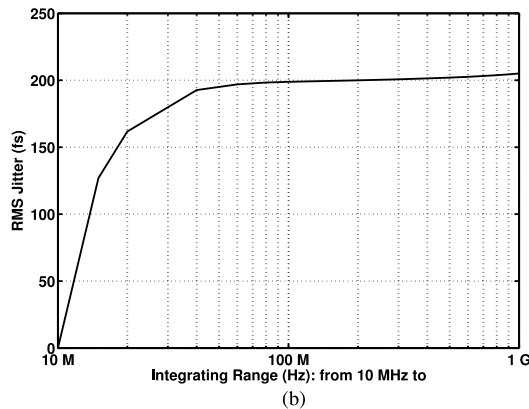


FIGURE 33. Measured 20-GHz clock spectrum.


yielding negligible deterministic jitter. Our phase noise measurement equipment faces two limitations, namely, the carrier frequency should be less than 13 GHz and the maximum offset frequency is 1 GHz. To address the former, we employ an external $\div 2$ circuit. Figure 34(a) plots the result. The phase noise within the loop bandwidth is around -110 dBc/Hz and falls to about -140 dBc/Hz at 200-MHz offset. These values are elevated by 6 dB for the 20-GHz clock. Figure 34(b) displays the jitter as a function of the integration bandwidth, suggesting that it reaches a relatively constant amount of



(a)



(b)

FIGURE 34. (a) Measured clock phase noise after an external divide-by-2 circuit, and (b) jitter as a function of integration bandwidth.

FIGURE 35. Spectrum of TX output for a nominally periodic 20-GHz waveform.

205 fs_{rms} beyond 200 MHz.³ As a worst-case estimate of jitter integrated up to 5-GHz offset (the Nyquist frequency), we integrate -140 dBc/Hz from 200 MHz to 5 GHz and obtain 100 fs_{rms}. Combined with 205 fs_{rms}, this result translates to a total jitter of 228 fs_{rms}.

As explained in Section V-D, the effect of mismatches upon the direct 4-to-1 MUX can be measured by generating a

3. Note that the integration begins from 100-Hz offset.

TABLE 2. Summary of measured TX performance and comparison with the prior art.

	Peng ISSCC'17	Steffan ISSCC'17	Dickson ISSCC'17	This Work
Technology (nm)	40	28	14	45
Data Rate (Gb/s)	56	64	56	80
Output Driver Type	CML	CML	SST	CML
Driver Supply (V)	1.5	1.2	0.95	1
Max. Output $V_{pp,d}$ (mV)	600	1200	900	630
RLM	N/A	0.94	N/A	0.99
RMS Jitter (fs)	688	290	318	205
Integ. Range (MHz)	0.0001 - 1000	0.5 - 8000	N/A	10 - 1000
Power (mW)	Exc.*	200	145***	101
	Inc.**	220	-	-
Power Eff. (pJ/bit)	Exc.**	3.57	2.26***	1.8
	Inc.**	3.93	-	-
Active Area (mm ²)	0.8*	N/A	0.035*	0.1

* Excluding PLL power but including clock distribution.

** Including PLL power and clock distribution.

*** Without I&Q clock generation.

20-GHz periodic NRZ waveform at the TX output and examining the spurs at ± 10 GHz around the carrier. Figure 35 plots the resulting spectrum. A spur level of -41 dBc in the single-ended output represents a deterministic jitter of 100 fs_{rms} due to such mismatches.

Table 2 compares the proposed transmitter's measured performance to that of the prior art. We note that, if the PLL power consumption is excluded, our work achieves a nearly six-fold improvement in power efficiency. As mentioned in Section V-E, the DAC supply voltage and tail currents can be raised so as to deliver a $1.2 \cdot V_{pp}$ output. In this case, our power efficiency is higher by about a factor of 4 (excluding the PLL).

VII. CONCLUSION

The design of broadband wireline transmitters presents numerous challenges from the circuit level to the architecture level. This paper describes these challenges and proposes a number of techniques that lead to output data rates as high as 80 Gb/s in 45-nm technology.

ACKNOWLEDGMENT

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

REFERENCES

- [1] T. Ali *et al.*, "6.2 a 460 mW 112 Gbps DSP-based transceiver with 38 dB loss compensation for next generation data centers in 7 nm FinFET technology," in *ISSCC Dig. Tech. Papers Slide Supplements*, Feb. 2020, pp. 118–120.
- [2] 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.
- [3] T. Ali *et al.*, "6.4 a 180 mW 56 Gb/s DSP-based transceiver for high-density IOs in data center switches in 7 nm FinFET technology," in *ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 118–120.
- [4] J. Im *et al.*, "6.1 a 112 Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR-ADC and inverter-based RX analog front-end in 7 nm FinFET," in *ISSCC Dig. Tech. Papers*, Feb. 2020, pp. 116–118.
- [5] T. Shibasaki *et al.*, "3.5 a 56Gb/s NRZ-electrical 247mW/lane serial-link transceiver in 28nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2016, pp. 64–66.

- [6] P.-J. Peng, J.-F. Li, L.-Y. Chen, and J. Lee, "6.1 a 56 Gb/s PAM-4/NRZ transceiver in 40 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2017, pp. 110–111.
- [7] J. Han, N. Sutardja, Y. Lu, and E. Alon, "Design techniques for a 60-Gb/s 288-mW NRZ transceiver with adaptive equalization and baud-rate clock and data recovery in 65-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 52, pp. 3474–3485, Dec. 2017.
- [8] E. Depaoli *et al.*, "A 4.9pJ/b 16-to-64Gb/s PAM-4 VSR transceiver in 28nm FDSOI CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2018, pp. 112–114.
- [9] J. Kim *et al.*, "A 112-Gb/s PAM4 transmitter with 3-Tap FFE in 10-nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2018, pp. 102–103.
- [10] Y. Chang and L. Kong, "An 80-Gb/s 40-mW wireline PAM4 transmitter," *IEEE J. Solid-State Circuits*, vol. 53, no. 8, pp. 2214–2226, Aug. 2018.
- [11] Y. Chang, A. Manian, L. Kong, and B. Razavi, "A 32-mW 40-Gb/s CMOS NRZ transmitter," in *Proc. CICC*, Apr. 2018, pp. 1–4.
- [12] B.-J. Yoo *et al.*, "6.4 a 56Gb/s 7.7mW/Gb/s PAM-4 wireline transceiver in 10 nm FinFET using MM-CDR-based ADC timing skew control and low-power DSP with approximate multiplier," in *ISSCC Dig. Tech. Papers*, Feb. 2020, pp. 122–123.
- [13] E. Groen *et al.*, "6.3 a 10-to-112Gb/s DSP-DAC-based transmitter with 1.2Vppd output swing in 7 nm FinFET," in *ISSCC Dig. Tech. Papers*, Feb. 2020, pp. 120–121.
- [14] J. Kim *et al.*, "8.1 a 224Gb/s DAC-based PAM-4 transmitter with 8-tap FFE in 10nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 126–127.
- [15] M. Choi *et al.*, "8 an output-bandwidth-optimized 200Gb/s PAM-4 100Gb/s NRZ transmitter with 5-tap FFE in 28nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 128–129.
- [16] M. Kossel *et al.*, "An 8b DAC-based SST TX using metal gate resistors with 1.4pJ/b efficiency at 112 Gb/s PAM-4 and 8-tap FFE in 7 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 130–131.
- [17] P. Mishra *et al.*, "8.7 a 112 Gb/s ADC-DSP-based PAM-4 transceiver for long-reach applications with >40 dB channel loss in 7 nm FinFET," in *ISSCC Dig. Tech. Papers*, Feb. 2021, pp. 138–139.
- [18] 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.
- [19] B. Razavi, K. F. Lee, and R. H. Yan, "Design of high-speed low-power dividers and phase-locked loops in deep submicron CMOS," *IEEE J. Solid-State Circuits*, vol. 30, no. 2, pp. 101–109, Feb. 1995.
- [20] *IEEE P802.3bs 400 GbE Task Force*. Accessed: Sep. 17, 2021. [Online]. Available: <http://www.ieee802.org/3/bs/>
- [21] S. Galal and B. Razavi, "10-Gb/s limiting amplifier and laser/modulator driver in 0.18 μ m CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2138–2146, Dec. 2003.
- [22] S. Galal and B. Razavi, "Broadband ESD protection circuits in CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2334–2340, Dec. 2003.
- [23] H.-M. Rein and M. Moller, "Design considerations for very-high-speed Si-bipolar ICs operating up to 50 Gb/s," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1076–1090, Aug. 1996.
- [24] M. Kossel *et al.*, "A T-coil enhanced 8.5-Gb/s high-swing SST transmitter in 65-nm Bulk CMOS with \ll -16 dB return loss over 10-GHz bandwidth," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2905–2920, Dec. 2008.
- [25] B. Razavi, "Jitter-power trade-offs in PLLs," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 4, pp. 1381–1387, Apr. 2021.
- [26] A. A. Hafez, M.-S. Chen, and C.-K. K. Yang, "A 32-48 Gb/s serializing transmitter using multiphase serialization in 65-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 3, pp. 763–775, Mar. 2015.
- [27] C.-K. K. Yang, R. Farjad-Rad, and M. A. Horowitz, "A 0.5- μ m CMOS 4.0-Gb/s serial link transceiver with data recovery using oversampling," *IEEE J. Solid-State Circuits*, vol. 33, pp. 713–722, May 1998.
- [28] L. Kong and B. Razavi, "A 2.4-GHz 4-mW integer-N inductorless RF synthesizer," *IEEE J. Solid-State Circuits*, vol. 51, no. 3, pp. 626–635, Mar. 2016.



BEHZAD RAZAVI received the B.S.E.E. degree from the Sharif University of Technology in 1985, and the M.S.E.E. and Ph.D.E.E. degrees from Stanford University in 1988 and 1992, respectively.

He was with AT&T Bell Laboratories and Hewlett-Packard Laboratories until 1996. Since 1996, he has been an Associate Professor and subsequently a Professor of Electrical Engineering with the University of California at Los Angeles. He was an Adjunct Professor with Princeton

University from 1992 to 1994, and with Stanford University in 1995. 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), *Design of CMOS Phase-Locked Loops* (Cambridge University Press, 2020), and *Fundamentals of Microelectronics* (Wiley, 2006, 2014, and 2021) (translated to Korean, Portuguese, and Turkish), and the Editor of *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996), and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003). His current research includes wireless and wireline transceivers and data converters.

Prof. Razavi received 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, and the 2012 Donald Pederson Award in Solid-State Circuits. He was the 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 the recipient of the American Society for Engineering Education PSW Teaching Award in 2014 and the 2017 IEEE CAS John Choma Education Award. He was also recognized as one of the top ten authors in the 50-year history of ISSCC. He served on the Technical Program Committees of ISSCC from 1993 to 2002 and VLSI Circuits Symposium from 1998 to 2002. He has also served as a Guest Editor and an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART I: REGULAR PAPERS, and *International Journal of High Speed Electronics*. He served as the Founding Editor-in-Chief of the IEEE SOLID-STATE CIRCUITS LETTERS. He has served as an IEEE Distinguished Lecturer. He is a member of the U.S. National Academy of Engineering.