

# An 80-Gb/s 44-mW Wireline PAM4 Transmitter

Yikun Chang<sup>1</sup>, *Student Member, IEEE*, Abishek Manian<sup>1</sup>, *Member, IEEE*,  
Long Kong<sup>1</sup>, *Member, IEEE*, and Behzad Razavi<sup>1</sup>, *Fellow, IEEE*

**Abstract**—A transmitter implemented in 45-nm CMOS technology serializes data from 312.5 Mb/s to an 80-Gb/s pulse-amplitude modulation 4 output with no need for latches. Utilizing the charge-steering techniques and a frequency divider that directly generates quadrature outputs with a 25% duty cycle, the design consumes 21.7 mW in the data path and 22.3 mW in the phase-locked loop and clock distribution while delivering a swing of 630 mV<sub>pp</sub> with a 1-V supply.

**Index Terms**—Charge-steering multiplexer (MUX), direct 4-to-1 MUX, divider, master-slave sampling filter (MSSF), multiplexer, serializer.

## I. INTRODUCTION

WITH the recent surge in the demand for high data rates, communication over copper media faces new challenges. The limited bandwidth removes so much of the signal's high-frequency energy that equalization and detection become very difficult. It is in this spirit that, after an initial appearance in 2000s [1]–[3], pulse-amplitude modulation 4 (PAM4) signaling has been resurrected. With a two-fold reduction in bandwidth occupancy compared to non-return-to-zero (NRZ) data, the PAM4 format allows significant speed improvement while introducing other issues.

This paper presents the design of an 80-Gb/s PAM4 transmitter (TX) in 45-nm CMOS technology that achieves significant improvement in power efficiency with respect to the state of the art. The prototype delivers a differential voltage swing of 630 mV<sub>pp</sub> and occupies an active area of 330 μm × 320 μm.

Section II provides the background for this paper and Section III introduces the TX architecture. Sections IV and V deal with the design of the serializer and the output driver, respectively. Sections VI and VII are concerned with the clock generation and phase-locked loop (PLL), respectively. Section VIII presents the experimental results.

## II. BACKGROUND

A number of PAM4 TXs operating at tens of gigabits per second has been reported [5]–[13]. Among these, the 56-Gb/s designs in [6] and [8] achieve a power

of 200 and 101 mW, respectively. The 64-Gb/s TX in [7] draws 145 mW. These values exclude the PLL. It is, therefore, prudent to identify the power-hungry functions in TXs before deciding on the architecture and its building blocks.

The fundamentally power-hungry circuit in a TX is the output driver. For a given voltage swing, this stage must deliver a certain current to the load (e.g., a 100-Ω differential resistance). In addition, at gigahertz speeds, the circuit must also include back termination resistors on the chip, which are approximately equal to the load impedance. This doubles the necessary supply current for a current-mode logic (CML) driver or the necessary supply voltage for a source-series-termination (SST) driver. Moreover, for a CML driver with PAM4 signaling, certain voltage headroom requirements must be met to ensure sufficient linearity. Thus, the supply voltage well exceeds the single-ended output swing, leading to a low efficiency.

To formulate the driver power consumption,  $P_{dr}$ , for a CML PAM4 topology, we consider the structure shown in Fig. 1, where half of the most significant bit (MSB) and the least significant bit (LSB) stages is shown for simplicity. We can view the circuit as a 2-bit digital-to-analog converter (DAC). Assuming  $R_T = R_L$ , and noting that the drain voltage has a common-mode (CM) level equal to  $V_{DD} - 3IR_T/2 = V_{DD} - 3IR_L/2$  and a single-ended peak-to-peak swing of  $V_{max} = 3I(R_T || R_L) = 3IR_L/2$ , we observe that the minimum supply voltage is given by  $V_{DD, min} = 3IR_L/4 + V_{max} + V_{DS} + V_{tail}$ , where  $V_{DS}$  and  $V_{tail}$  denote the minimum allowable drain-source voltage for the output transistors and the tail currents, respectively. It follows that  $V_{DD, min} = 1.5V_{max} + V_{DS} + V_{tail}$ , yielding a power consumption of:

$$\begin{aligned} P_{dr} &= V_{DD, min}(3I) \\ &= (1.5V_{max} + V_{DS} + V_{tail}) \frac{2V_{max}}{R_L} \\ &= \frac{3V_{max}^2}{R_L} + \frac{2V_{max}(V_{DS} + V_{tail})}{R_L}. \end{aligned} \quad (1)$$

Since  $V_{DS} + V_{tail}$  is comparable to  $V_{max}$ , the second term is nearly equal to the first. For example, if  $V_{max} = 350$  mV and  $V_{DS} + V_{tail} \approx 500$  mV, and  $R_L = 50$  Ω, we have  $P_{dr} \approx 14.35$  mW. The key point, here, is that the driver power consumption is given by a few fundamental parameters and cannot be reduced significantly. Note that these results also apply to NRZ output stages to some extent, with only  $V_{DS}$  being slightly more flexible due to the relaxed linearity requirement in that case.

The foregoing analysis can be repeated for voltage-mode drivers, specifically, those using SST [8], [10], [12]. Depicted

Manuscript received December 7, 2017; revised March 9, 2018 and April 19, 2018; accepted April 19, 2018. Date of publication June 1, 2018; date of current version July 20, 2018. This paper was approved by Associate Editor Jack Kenney. (Corresponding author: Yikun Chang.)

The authors are with the Department of Electrical and Computer Engineering, University of California at Los Angeles, Los Angeles, CA 90095-1594 USA (e-mail: changyk@ucla.edu; 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.2018.2831226




Fig. 1. CML output driver and its drain waveform.

in Fig. 2, such a topology incorporates the two scaled inverters and series termination resistors  $R_{T1}$  and  $R_{T2}$ . The choice of  $R_{T1} = 1.5R_L$  and  $R_{T2} = 3R_L$  yields proper PAM4 levels with a maximum single-ended swing of  $V_{\max} = V_{DD}/2$ , and  $R_{T1} || R_{T2} = R_L$  ensures proper back termination [10]. In this case, the inverter transistors must be so wide as to contribute an output resistance well below their respective series resistors. This circuit's power consumption is a function of the output voltage, exhibiting an average value given by  $13V_{DD}^2/36R_L = 13V_{\max}^2/9R_L$  if the PAM4 levels occur with equal probabilities. For a single-ended swing of 350 mV, we can choose  $V_{DD} = 700$  mV and obtain a total power of  $13V_{DD}^2/36R_L = 3.54$  mW for two such drivers operating differentially. While draining less power than its CML counterpart, the SST stage faces difficulties at the speeds of interest here.

To put matters in perspective, we ask, if the driver power can be maintained around roughly 10 mW, where does the remainder of the 100–200 mW go in actual designs, e.g., in [6]–[8]? We expect that the overall serializer that multiplexes the data from low speeds to the final data rate also draws considerable power. The issue is exacerbated in a PAM4 TX owing to the need for two separate multiplexer (MUX) chains for the MSB and LSB paths (Section III). For example, serialization from 312.5 Mb/s to 40 Gb/s (up to the inputs of the output driver/DAC) would require  $3 \times 254$  latches if three-latch MUX cells are used in a binary tree. Even though the number of latches drops by a factor of 2 from one rank to the next, the increase in speed at least doubles the power per latch. Consequently, the serializer can consume tens of milliwatts in 45-nm technology (Section IV).

The generation and distribution of the clock and its divided versions can also draw a high power. Among the prior PAM4 TXs, [7] includes the distribution in the overall power numbers but not the PLL and phase generation. The design in [6] reports a PLL power of 20 mW at 14 GHz, excluding phase generation and distribution. Thus, the PLL also merits investigation if the overall TX power must be minimized.

### III. TRANSMITTER ARCHITECTURE

Fig. 3 shows the proposed TX architecture, which consists of MSB and LSB data paths, an output driver/DAC, and a clock generation module. Each serializer consists of a CMOS MUX, a charge-steering MUX, and a direct 4-to-1 MUX. The co-design of the data paths and the PLL allows the former to employ new circuit topologies that substantially reduce the power. Specifically, the feedback dividers provide quadrature phases,  $\phi_1$ – $\phi_4$ ,  $45^\circ$  phases, select  $SEL_1$ – $SEL_4$ , etc., making it possible to avoid latches in the entire serializer (Section IV).

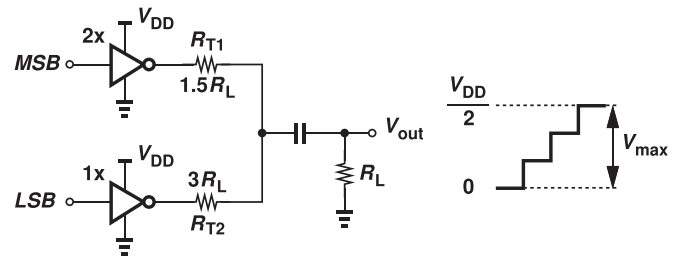


Fig. 2. SST driver and its output waveform.

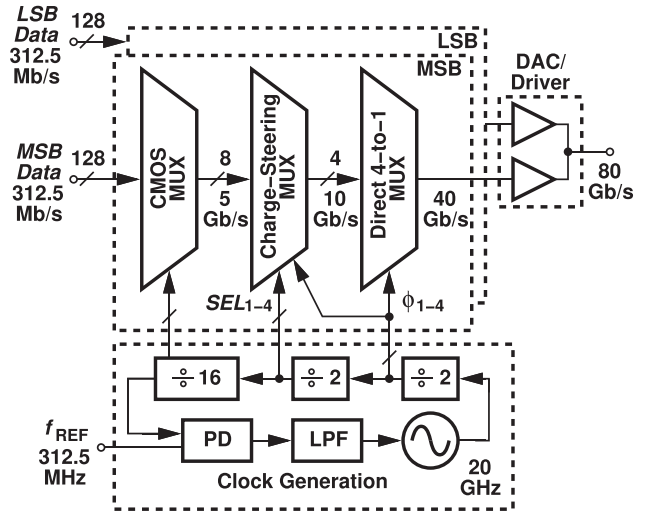


Fig. 3. Proposed TX architecture.

The interface between the MSB and LSB serializers and the driver/DAC in Fig. 3 entails a critical issue. Since the DAC MSB cell presents twice as much input capacitance as the LSB cell does, the two serializers preceding the DAC must have proportionally scaled drive strengths so to avoid a systematic skew between the MSB and the LSB waveforms. Such a skew manifests itself as jitter at the final output. Thus, the drive strength of the direct 4-to-1 MUX stage in the MSB serializer is scaled up by a factor of 2, but the stages before this MUX remain mostly unscaled.

### IV. SERIALIZER DESIGN

As mentioned in Section III, the TX must employ two serializer paths for the MSB and the LSB, potentially consuming a high power. In this paper, we propose a number of techniques to ameliorate this issue: 1) the use of three logic styles in Fig. 3 allows the optimum speed–power tradeoff; 2) a new “latchless” MUX design; 3) charge steering [15] as a paradigm that affords a higher speed than CMOS logic and a lower power consumption than CML; and 4) a direct latchless 4-to-1 MUX that considerably reduces the number of high-speed stages. We describe these concepts in the following.

#### A. CMOS MUX

Rail-to-rail CMOS logic provides robust operation with a power of the form  $fCV_{DD}^2$ , where  $f$  denotes the frequency at which  $C$  charges from 0 to  $V_{DD}$ . In the context of TX design,

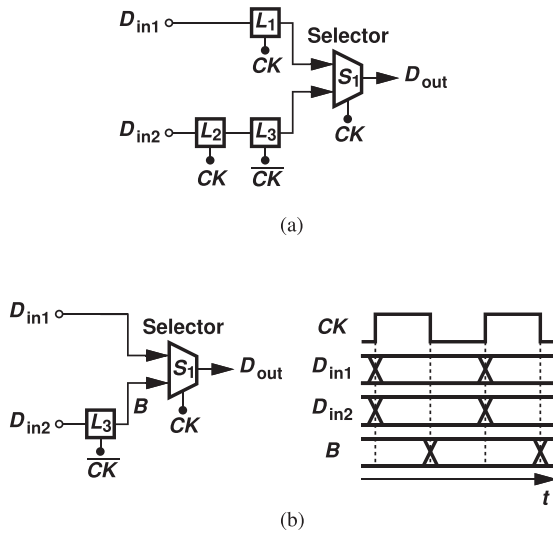


Fig. 4. (a) Conventional three-latch MUX cell. (b) Simplified MUX cell.

we must decide on the maximum reliable speed that this style can support. The architecture in Fig. 3 comfortably utilizes the rail-to-rail stages to serialize the data from 312.5 Mb/s to 5 Gb/s.

The 128-to-8 binary-tree CMOS MUX requires 120 2-to-1 MUX cells. As shown in Fig. 4(a), a typical cell comprises three latches and one selector, with  $L_1$  and  $L_2$  holding the inputs so as to block glitches from preceding stages, and  $L_3$  serving to avoid input change when the clock has selected that input. However, if the timing of  $D_{in1}$  and  $D_{in2}$  is known and well-controlled,  $L_1$  and  $L_2$  can be omitted [Fig. 4(b)] [14]. In this case, the assumption is that  $D_{in1}$  and  $D_{in2}$  change on one edge of the clock and settle before the next edge of the clock. Also,  $L_3$  ensures that the selector inputs do not make simultaneous transitions.

If the multiplexing clock is available in the quadrature phases, clock  $CK_I$  and  $CK_Q$ , the serializer design can be improved. For example, [6] utilizes such phases to establish a longer hold time for the MUX input. We introduce a new serialization approach that exploits  $CK_I$  and  $CK_Q$  to eliminate all latches in the data path.<sup>1</sup> Illustrated in Fig. 5, the idea is to create the necessary delay between each selector's inputs by proper choice of the clock edges in consecutive stages. Let us consider how  $D_{even}$  and  $D_{odd}$  avoid simultaneous transitions, noting that selectors  $S_2$  and  $S_3$  are driven by  $CK_{2,I}$  and  $CK_{2,Q}$ , respectively. We make two observations: 1) the edges of these two clocks have an offset equal to  $T_{CK2}/4$ , and hence  $D_{odd}$  changes  $T_{CK2}/4$  seconds after  $D_{even}$  does and 2) since the edge separation between  $CK_I$  and  $CK_2$  ( $\approx 200$  ps) is long enough for  $D_{even}$  or  $D_{odd}$  to settle, no glitch appears at the input of  $S_1$ . Thus, the three-cell structure consisting of  $S_1$ ,  $S_2$ , and  $S_3$  can be repeated in the preceding ranks so long as the clock phases are chosen accordingly.

The 120 selectors necessary for multiplexing 312.5 Mb/s to 5 Gb/s can incur a high power consumption in their clock

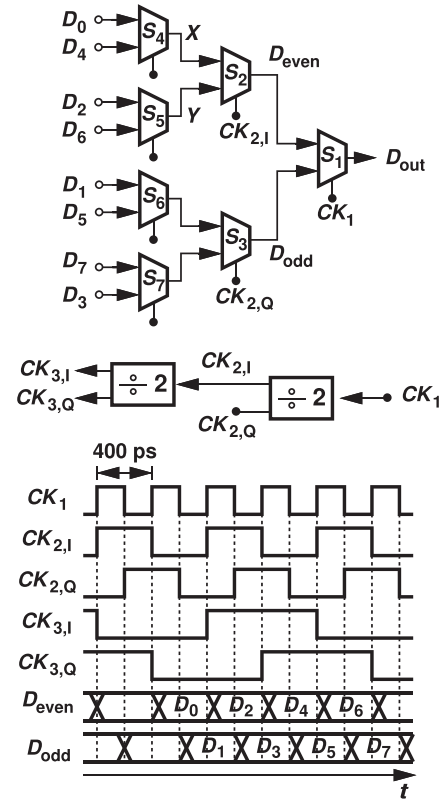


Fig. 5. Proposed timing scheme to remove latches by applying  $I$  and  $Q$  clocks.

path. We therefore wish to minimize the dimensions of the clocked transistors and the length of the clock wires. On the other hand, the drive strength of the last selector must suffice for the operation of the subsequent (charge-steering) MUX, calling for wide transistors.

Based on the above-mentioned considerations, the selector unit is realized as shown in Fig. 6(a). This topology occupies a small area, allowing short interconnects for the entire CMOS serializer, and achieves sufficient speed. For  $S_1$  in Fig. 5, the transistor dimensions are chosen as  $W_N = 1 \mu\text{m}$ ,  $W_P = 1.5 \mu\text{m}$ , and  $L = 40$  nm, leading to a power consumption of  $22 \mu\text{W}$  for this unit (in both the data and the clock paths). The eye diagram shown in Fig. 6(b) represents this output. Since the stages preceding this selector operate at progressively lower frequencies, the unit design is scaled down by a factor of 2 from one rank to the rank preceding it, until a minimum allowable width of 120 nm is reached. Note that the latchless topology does not exhibit glitches because it benefits from ample timing margin between the  $I$  and  $Q$  edges. Also, the clocking action applied to the selector does not allow device or timing mismatches to accumulate through the serializers. The entire 128-to-8 serializer draws  $365 \mu\text{W}$  in the data path.<sup>2</sup> The single-ended output is converted to complementary form by means of an inverter after the final CMOS MUX stage.

<sup>1</sup>The use of quadrature clocks does not translate to a power penalty because every selector would need a clock in any other architecture as well.

<sup>2</sup>A three-latch approach would require a power consumption of about 11 mW for the 128-to-8 serializer including the clock path.

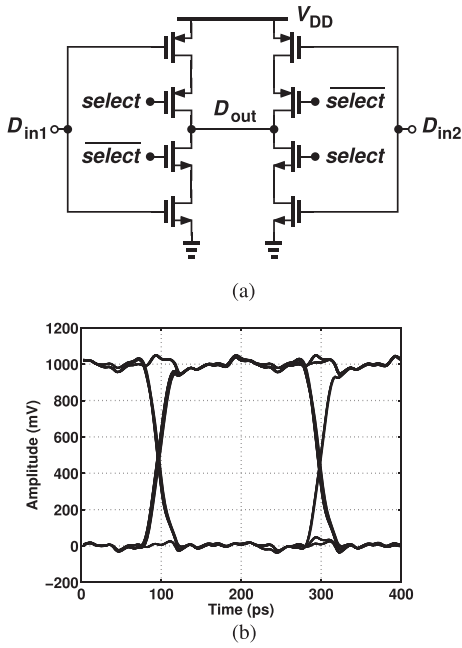


Fig. 6. (a) CMOS selector used in this paper. (b) Simulated output eye diagram of the last stage of CMOS MUX.

**B. Charge-Steering MUX**

For operation above 5 Gb/s, charge steering proves more viable than CMOS logic. By virtue of their moderate voltage swings ( $\approx 300$  mV<sub>pp</sub> single-ended), charge-steering circuits achieve a higher speed [15]. We propose a number of techniques that improve the performance of charge-steering stages in the context of the 8-to-4 MUX in Fig. 3.

We begin with the simple charge-steering selector shown in Fig. 7(a). When  $CK$  is low,  $S_1$ – $S_3$  are ON,  $C_T$  is discharged to ground and  $X$  and  $Y$  are precharged to  $V_{DD}$ . When  $CK$  rises, the output begins to track  $V_{in1}$  or  $V_{in2}$  depending on the logical value of  $SEL$ . Capacitor  $C_T$  continues to draw charge from  $X$  or  $Y$  until its voltage reaches approximately one threshold voltage below the input CM level, at which point  $V_X$  or  $V_Y$  approaches its minimum value.<sup>3</sup>

The reset action at the output nodes removes ISI but occupies about half of the clock cycle, during which the next stage must not sense  $X$  and  $Y$ . Note that none of the transistors need operate in saturation because the rail-to-rail input and clock swings guarantee complete steering of the charge. In this topology,  $CK$  runs at twice the  $SEL$  frequency, which itself is equal to the input data rate.

If used in the TX architecture of Fig. 3, the above charge-steering selector faces a critical issue: the levels produced at  $X$  and  $Y$  deteriorate due to the kickback noise of the next stage, namely, the direct 4-to-1 MUX (Section IV-C).<sup>4</sup> Fortunately,  $V_{in1}$  and  $V_{in2}$  in Fig. 7(a) are produced by the CMOS serializer and have rail-to-rail swings. Exploiting these swings, we add a small helper of an PMOS selector to the circuit as depicted

<sup>3</sup>The charge-steering MUX does not allow the output low level to reach zero regardless of the clock period, a point of contrast to current-integrating circuits.

<sup>4</sup>This issue is also present if a current-integrating MUX is used.

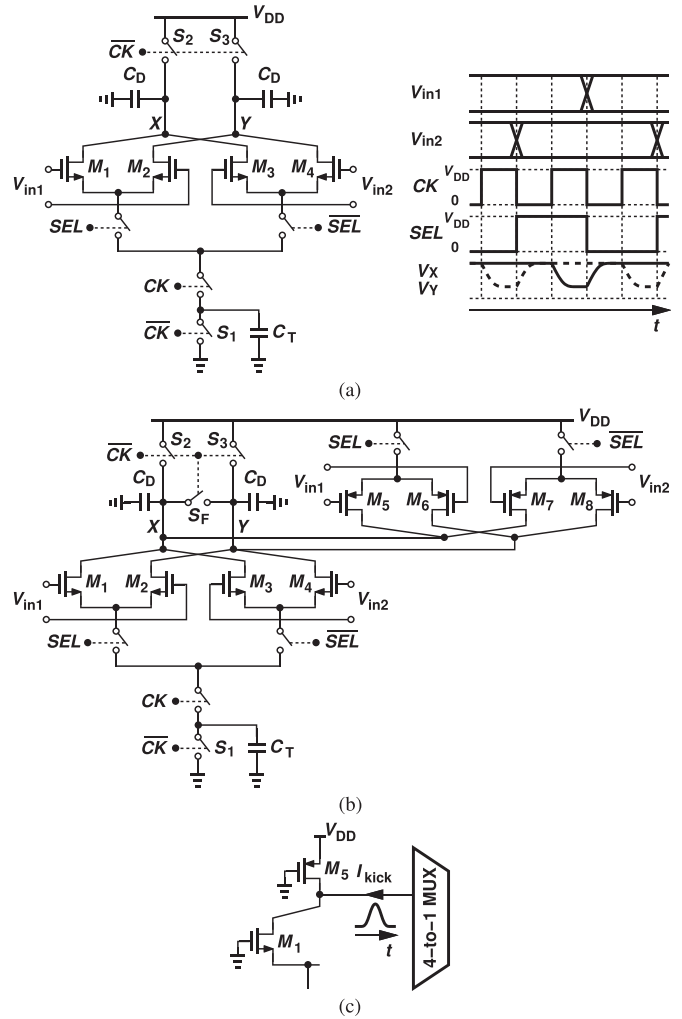


Fig. 7. (a) Simple charge-steering 2-to-1 MUX. (b) Proposed charge-steering MUX. (c) Role of PMOS pull-up device in suppressing the effect of kickback noise.

in Fig. 7(b). Here, for a given input state, one of  $M_5$ – $M_8$  conducts, providing a resistive path from  $X$  or  $Y$  to  $V_{DD}$  [Fig. 7(c)], and hence, restoring the high level even in the presence of kickback noise from the next stage. Fig. 8 plots the selector’s simulated output waveforms with and without the PMOS differential pairs, indicating an improvement of about 100 mV in the high level.<sup>5</sup>

Another difficulty in the charge-steering selector design is that, at 10 Gb/s, nodes  $X$  and  $Y$  in Fig. 7(a) do not precharge to  $V_{DD}$  completely, thereby suffering from ISI and degraded levels. This is alleviated by introducing switch  $S_F$  in Fig. 7(b), which ensures  $V_X \approx V_Y$  during precharge.

Since the above selector’s output is unavailable in the precharge mode, the 8-to-4 charge-steering MUX and the direct 4-to-1 MUX in Fig. 3 must be co-designed to ensure compatibility between their timings. We propose the use of quadrature clock phases with 25% duty cycle for both.

<sup>5</sup>The PMOS devices primarily restore the output CM level, providing a greater voltage headroom for the direct 4-to-1 MUX tail devices.






Fig. 8. Simulated output waveforms of charge-steering MUX with and without PMOS differential pairs.

To this end, we modify the selector's clocks as shown in Fig. 9(a). Here, the clock phase  $RST$  performs precharge and reset for 25 ps and the  $EVL$  phase evaluates the input also for 25 ps. The command  $SEL$  selects one input after each precharge interval. Thus, the output is available from  $t_3$  to  $t_4$ .

The 8-to-4 MUX requires four two-input selectors whose timings must agree with those of the direct 4-to-1 MUX. This is accomplished as illustrated in Fig. 9(b), where  $\phi_1$ – $\phi_4$  denote the four phases of the 10-GHz clock with 25% duty cycle and  $SEL_1$ – $SEL_4$  are the 45° phases of the 5-GHz clock with 50% duty cycle. The first selector on the left operates with  $\phi_1$  and  $\phi_2$  in the same manner as in Fig. 9(a), i.e.,  $RST = \phi_1$ ,  $EVL = \phi_2$ . For the next selector,  $\phi_2$  and  $\phi_3$  act as  $RST$  and  $EVL$ , respectively, and  $SEL_2$ , which is 25 ps behind  $SEL_1$ , drives the  $SEL$  input. The remaining two selectors run on other rotated phases, and the four outputs  $D_a$ – $D_d$  appear in succession.

The idealized situation depicted in Fig. 9(b) assumes a zero delay between the rising edge of  $\phi_2$  and the rising edge of  $SEL_1$  and similarly for other phases. In reality, however,  $SEL_1$  is obtained by frequency division and incurs a delay of about 20 ps. Thus, the charge-steering action is delayed by this amount, shortening the time available for evaluation to about zero. To resolve this issue, we recognize that the select command in Fig. 9(a) can be asserted even before  $EVL$  arrives. We therefore apply  $SEL_4$ , rather than  $SEL_1$ , to the first selector and rotate the rest accordingly. Fig. 9(c) shows the resulting assignment of  $SEL_1$ – $SEL_4$ .

While the present prototype does not include feedforward equalization (FFE), our scheme makes it possible to add FFE with minimal power penalty. We briefly explain the idea, here, based on the charge-steering MUX of Fig. 7(b) and refer the reader to a similar FFE implementation based on an integrating MUX in [18]. To create a post cursor tap, we first decompose the following direct 4-to-1 MUX and the output driver into, for example, four slices, three of which are driven by the main cursor and the fourth by the post cursor. Since the MUX of Fig. 7(b) holds the output for 2 unit interval (UI) [Fig. 9(c)], the second UI (from  $t_3$  to  $t_4$ ) can be used to drive the post cursor without adding any latches. This overall strategy can be applied to both the MSB and the LSB paths.

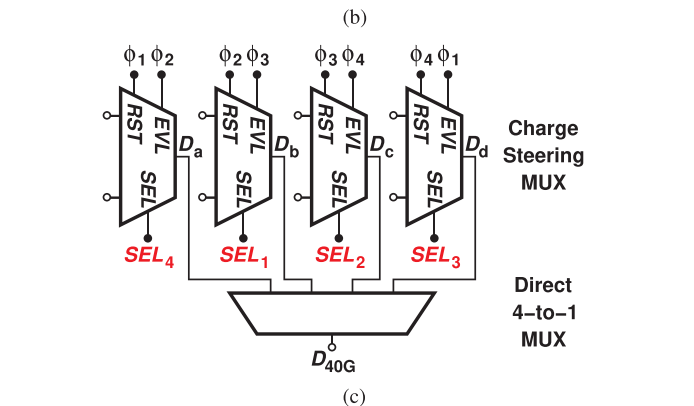

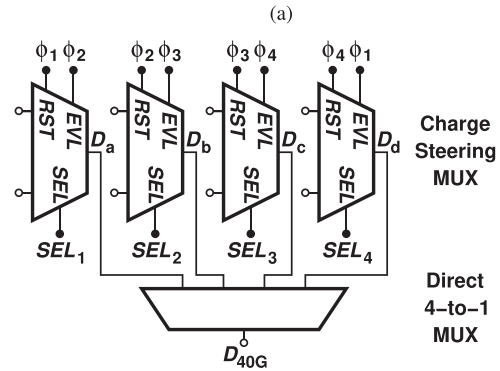
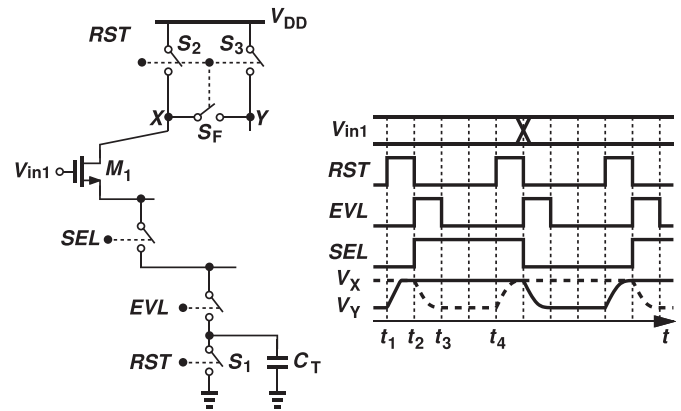


Fig. 9. (a) Timing diagram of charge-steering MUX with 25% duty-cycle clocks. (b) Four charge-steering MUXes with idealized waveforms. (c) Rotation of  $SEL_1$ – $SEL_4$  in four charge-steering MUXes to accommodate the clock delay.

### C. Direct 4-to-1 MUX

Serialization of data from 10 Gb/s to 40 Gb/s in 45-nm CMOS technology inevitably calls for CML implementations.

However, a 4-to-1 binary-tree topology employing one latch per selector would require 12 tail currents (for MSB and LSB paths), at least eight transistors clocked at 20 GHz, and at least 16 at 10 GHz. The latchless topology described in Section IV-A could potentially save a total of six high-speed latches but would necessitate quadrature phases of the 10-GHz clock with a 50% duty cycle. The charge-steering MUX, on the other hand, requires 25%-duty-cycle phases at 10 GHz. We must therefore develop a CML MUX that can operate with the latter.

We opt for a direct 4-to-1 CML structure that can utilize these phases. Fig. 10 depicts the result. The four differential pairs are enabled in succession such that each senses an input that is evaluated and held by the preceding charge-steering selector. Inductive peaking deals with the heavy capacitive load ( $\approx 82$  fF for the MSB path and  $\approx 40$  fF for the LSB path) presented by the large input transistors of the next stage (the output driver/DAC) and the self-load from the four differential pairs.

Direct 4-to-1 MUX topologies have been reported [14], but our approach merits some remarks. First, at a clock frequency of 10 GHz, the use of single clocked transistors driven by  $\phi_1$ – $\phi_4$  proves more efficient than generating overlapping quadrature phases and using stacked transistors to perform a NAND gate [17]. Second, with the rail-to-rail swings for  $\phi_1$ – $\phi_4$ , the clocked transistors need only be  $8 \mu\text{m}$  wide for the MSB path and  $4 \mu\text{m}$  wide for the LSB path to draw a sufficient current, but the MUX output swing exhibits some dependence upon the PVT. Nevertheless, so long as the output swing is large enough to ensure complete switching in the following driver, this dependence is benign. The values shown in Fig. 10 correspond to the LSB path; the design is linearly scaled up by a factor of 2 for the MSB path.

In Section VI, we address the task of generating the clock phases and observe that their duty cycle can be slightly less or greater than 25% depending on the circuit topology. We must, therefore, quantify the effect of this systematic departure upon the MUX performance. Plots in Fig. 11(a) are the width and height of the TXs output eye as a function of the duty cycle. Here, the middle eye of PAM4 is examined. We note that: 1) the width in fact prefers a duty cycle of about 23%<sup>6</sup> and 2) the height is less sensitive, prefers about 28%, and can tolerate from about 22% to 33%. Fig. 11(b) and (c) depicts the simulated examples, indicating that erring toward smaller values is more tolerable because the eye in the former exhibits a greater opening. The simulations leading to Fig. 11 include the direct 4-to-1 MUX and the output driver (with inductive peaking) with a clock transition time of 15 ps. These simulations can be repeated with a channel model and other imperfections to determine the optimum duty cycle.

As mentioned in Section IV-B, the MUX of Fig. 10 draws the transient kickback currents from its inputs. The kickback arises when one tail device turns on and its current must initially flow from the  $C_{GS}$  of the corresponding differential




Fig. 10. Direct 4-to-1 CML MUX.

pair transistors. For the MSB path, the resulting gate current has a peak of  $260 \mu\text{A}$  and lasts about 20 ps. The PMOS differential pairs in the charge-steering selector alleviate the issue. For the MSB path, the tail capacitance in Fig. 7(b) is doubled to ensure sufficient voltage swings at X and Y, and the precharge switches are widened by a factor of 2 to guarantee proper reset.

## V. OUTPUT DRIVER/DAC

The 40-Gb/s MSB and LSB data streams are combined in the output driver to produce the final 80-Gb/s PAM4 signal.

Fig. 12 shows the realization, where three nominally identical differential pairs act as a 2-bit DAC. The 300-pH inductors broaden the bandwidth in the presence of the driver output capacitance ( $\approx 73$  fF) and the pad and ESD capacitance ( $\approx 50$  fF).<sup>7</sup> The overall circuit consumes 13 mW from a 1-V supply.

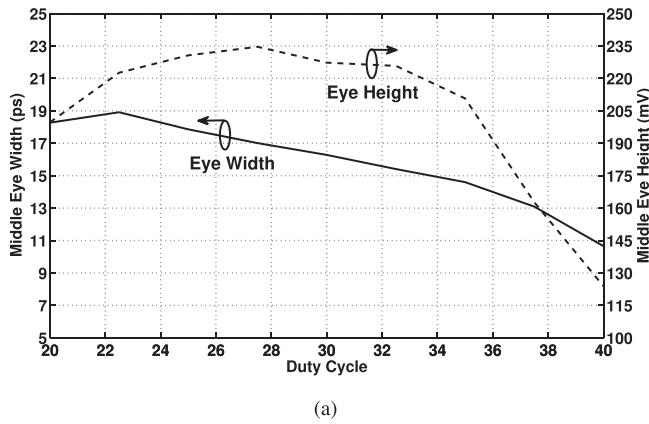
The use of short-channel devices raises concern regarding the nonlinearity of the DAC: since the output resistance varies with the digital input, the output eye can be distorted. The effect is exacerbated by the fact that the input high level is close to  $V_{DD}$ , forcing the transistors into the triode region for some output PAM4 levels. However, it can be shown that, unlike general current-steering DAC, the two-bit topology does not exhibit much nonlinearity arising from the finite output resistance of the units (Appendix).

## VI. CLOCK GENERATION


As explained in Section IV, the TX in Fig. 3 extensively exploits quadrature and  $45^\circ$  clock phases with 25% or 50% duty cycles to perform serialization without the use of latches. The generation and distribution of these phases, thus,

<sup>7</sup>Series peaking in this case simplifies the layout as the inductors become part of the routing to the pads. In practice, larger ESD devices embedded in a T-coil can be used [19]. An octagonal pad structure consisting of metal 8 and metal 9 and with a diameter of  $50 \mu\text{m}$  helps to reduce the pad capacitance to 30 fF.


<sup>6</sup>Since the turn-off and turn-on delays of the direct 4-to-1 MUX tails are not equal, the neighboring branches briefly overlap in time for a duty cycle of 25%.



(a)



(b)



(c)

Fig. 11. (a) Dependence of height and width of middle eye in PAM4 upon duty cycle. (b) Output eye for 20% duty cycle. (c) Output eye for 37.5% duty cycle.

play a central role in the overall performance and power consumption.

The most critical clock phases are those running at 10 GHz with a duty cycle of 25% because their mismatches directly translate to jitter at the output of the 4-to-1 MUX. To create these phases, we can: (1) directly generate 10-GHz overlapping quadrature clocks by means of two coupled  $LC$  oscillators and use AND gates to convert the duty cycle to 25%; (2) generate a 20-GHz differential clock, apply it to a standard




Fig. 12. Topology of the PAM4 CML output driver/DAC.

$\div 2$  circuit and AND the results; or (3) generate a 20-GHz differential clock and apply it to a  $\div 2$  circuit that inherently produces outputs with a 25% duty cycle. From Fig. 11(a), we target an optimal duty cycle of around  $25\% \pm 3\%$ . The first approach is less attractive as quadrature  $LC$  VCOs suffer from a high phase noise and require at least two symmetric inductors, complicating the floor plan. The second method demands that CMOS static AND gates operate at 10 GHz, a difficult and power-hungry task. The third solution is potentially the most efficient since it avoids the logic altogether.

We begin with the divider topology illustrated in Fig. 13(a) [20], whose outputs have a duty cycle of approximately 25%. While achieving a high speed, this structure faces two drawbacks: 1) the logical low levels at the output are degraded for about one quarter of the time and 2) the duty cycle is in fact greater than 25% by one gate delay. To understand the cause of these issues, we examine the circuit's operation with the aid of the waveforms shown in Fig. 13(b). Suppose  $CK$  is low,  $V_{X1}$  is high, and the other three outputs are low. At  $t = t_1$ ,  $CK$  rises and  $\overline{CK}$  falls, turning on  $M_{10}$  and pulling  $V_{Y2}$  to  $V_{DD}$  at  $t = t_2$  (while  $M_{12}$  is OFF). Since  $V_{X1}$  is still high,  $M_{11}$  is ON, but  $M_9$  has also turned on. Thus, the low level in  $V_{X2}$  degrades and a static current flows. Now, the rising edge at  $Y_2$  drives  $M_5$  and brings  $V_{X1}$  down at  $t = t_3$ . That is, the high-to-low transition at  $V_{X1}$  occurs two gate delays after the rising edge on  $CK$ . The operation proceeds in a similar manner until  $t = t_4$ , when  $CK$  falls, causing  $V_{X1}$  to rise at  $t = t_5$ . In summary,  $V_{X1}$  incurs one gate delay on its falling edge and two on its rising edge, exhibiting a duty cycle of 25% plus one gate delay, a significant error at 10 GHz.

In order to eliminate the static current, a cross-coupled pair can be inserted in series with the drains of the clocked transistors [21], but, owing to the greater gate delay, the duty cycle error increases further. As an alternative approach, let us consider the static latch topology shown in Fig. 13(c), where  $M_a$  and  $M_b$  are controlled by the inputs. If, for example,  $CK$  falls when  $D_{in}$  is high,  $M_5$  does not fight  $M_3$  anymore. Nevertheless, the duty cycle still remains well above the desired value. To address this issue, we recognize in Fig. 13(b) that any rising edge on  $CK$  can be allowed to pull  $V_{X1}$  to zero.

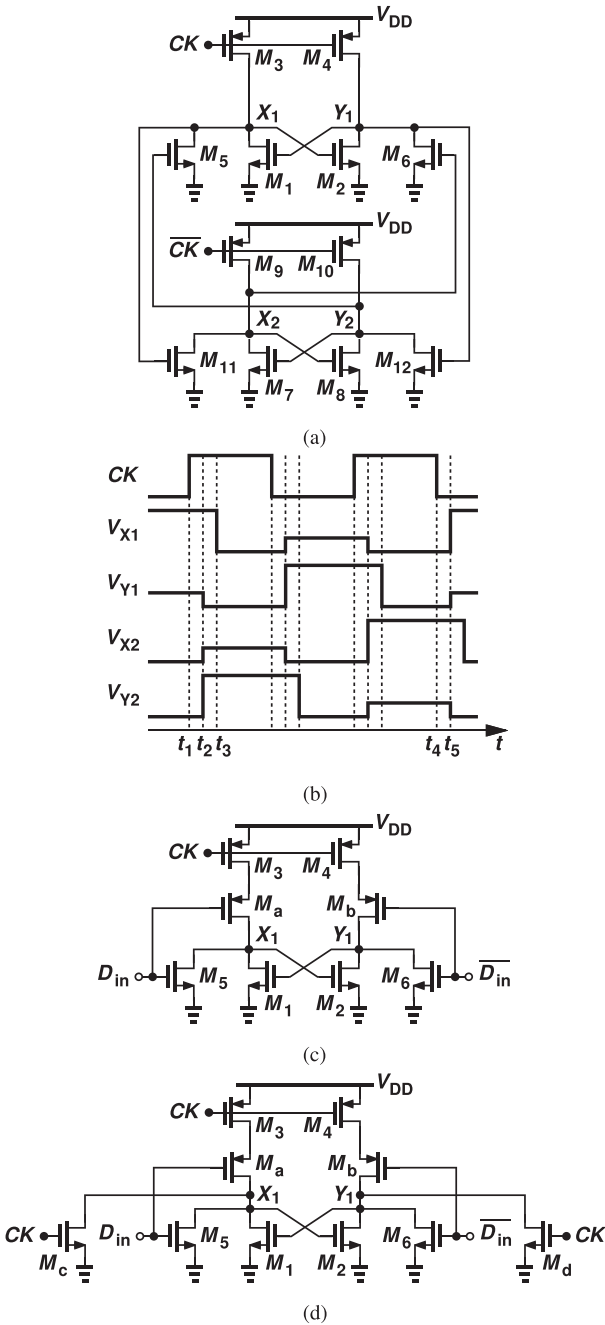


Fig. 13. (a) Divider topology to generate 25%-duty-cycle clocks directly [20]. (b) Divider’s waveforms. (c) Latch topology to remove static current of  $M_a$  and  $M_b$ . (d)  $M_c$  and  $M_d$  driven by  $CK$  to reduce transition delay of falling edge on  $V_{X1}$  and  $V_{Y1}$ .

In other words,  $CK$  can directly lower  $V_{X1}$  rather than through  $V_{Y2}$ . This observation leads us to add two clocked devices,  $M_c$  and  $M_d$ , as shown in Fig. 13(d) such that they can, respectively, force  $V_{X1}$  or  $V_{Y1}$  to zero when  $CK$  goes high. Proper rating of  $W_{5,6}$  and  $W_{c,d}$  yields the desired duty cycle.

The series combination of PMOS devices in Fig. 13(d) degrades the divider’s speed significantly. We then change all of the transistors to their opposite type, arriving at the proposed latch design depicted in Fig. 14(a)<sup>8</sup> and the simu-

<sup>8</sup>The ratios chosen here lead to a duty cycle range of 24%–32% across SS, SF, FS, FF, and TT corners.




Fig. 14. (a) Proposed latch topology with stacked NMOS devices, and simulated waveforms of (b) divider outputs, and (c) after three buffers.

lated waveforms in Fig. 14(b) and (c). According to simulations, the topology of Fig. 13(d) reaches a maximum speed of 23 GHz and that in Fig. 14(a), 29 GHz. The divider is followed by an inverter first to generate the complementary phases and by another two inverters to drive the next divider and deliver the four phases to the charge-steering MUX and the direct 4-to-1 MUX in Fig. 3. The divider core consumes 3.7 mW at an input frequency of 20 GHz, the first inverter, 1.8 mW, and the second set of inverters, 6.3 mW.

As mentioned in Section III, with no retimer after the 4-to-1 MUX, the mismatches between the clock phases produce jitter. Monte Carlo simulations of the divider, its buffers, the four charge-steering 2-to-1 MUXes, the direct 4-to-1 MUX, and the output driver/DAC indicate a one-sigma jitter of 75 fs<sub>rms</sub>



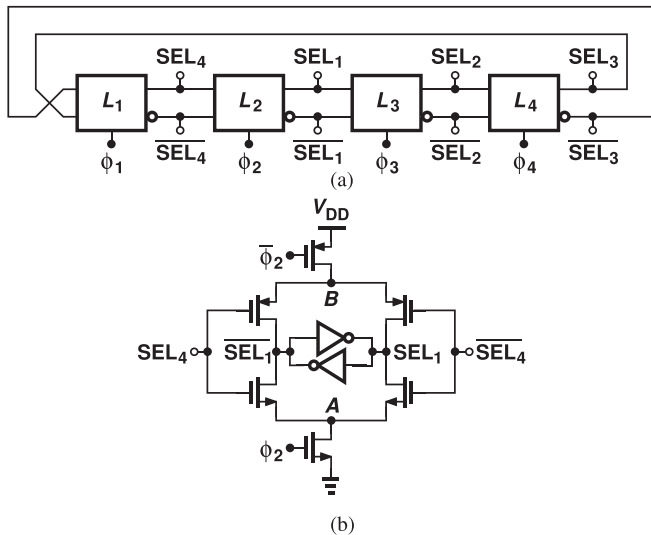


Fig. 15. (a) Divide-by-2 stage to generate eight-phase clocks. (b)  $C^2MOS$  latch used in the divider.

due to mismatches. We also observe in Section VIII that the measured TX output jitter in the 40-Gb/s NRZ mode is only 479  $f_{s_{rms}}$  and the measured duty cycle distortion (DCD) is 100  $f_{s_{rms}}$ , concluding that the matching is acceptable.

The second divide-by-2 stage in Fig. 3 runs at an input frequency of 10 GHz but, with only 25%-duty-cycle phases available from the preceding divider, it must operate with a clock high level that lasts less than 25 ps. Moreover, the circuit must provide eight output phases,  $SEL_j$  and  $\overline{SEL}_j$  for  $j = 1, \dots, 4$ . For this purpose, we introduce another new divider topology that exploits all four 10-GHz phases. Shown in Fig. 15(a), the circuit incorporates four latches that are consecutively driven by  $\phi_1$ – $\phi_4$ , thereby shifting two ONEs and two ZEROs by 25 ps every time  $\phi_j$  pulsates. Fig. 15(b) depicts the  $C^2MOS$  latch used here, with the cross-coupled inverters guaranteeing differential operation. The overall circuit draws 1.9 mW at an input frequency of 10 GHz.

## VII. PLL DESIGN

In most high-speed wireline TXs, the PLL and the clock distribution network draw considerable power. In this paper, the PLL generates a 20-GHz output that is subsequently divided to produce the phases and frequencies necessary for serialization. With  $UI = 25$  ps, we target an overall PLL jitter of 300  $f_{s_{rms}}$  for negligible degradation of the transmitted data.

The PLL jitter arises from the reference spurs, the VCO phase noise, and the multiplied reference phase noise. The closed-loop bandwidth,  $f_{BW}$ , must therefore be optimized in terms of these three imperfections.

To quantify the deterministic jitter due to the reference spurs, we write  $V_0 \cos(\omega_c t + \beta \sin \omega_m t) \approx V_0 \cos \omega_c t - \beta V_0 \sin \omega_c t \sin \omega_m t = V_0 \cos \omega_c t - 0.5\beta V_0 \cos(\omega_c - \omega_m)t + 0.5\beta V_0 \cos(\omega_c + \omega_m)t$  and note that the normalized spur level is  $\beta/2$ . Also, the peak jitter in radians is equal to  $\beta$ . Thus, if the normalized spur level in the spectrum is multiplied by  $\sqrt{2}$ , it yields the rms jitter. For example, if the spurs




Fig. 16. PLL with MSSF.

are at  $-50$  dBc, the jitter is around 36  $f_{s_{rms}}$ , and hence negligible. We also note that a crystal oscillator phase noise,  $S_{REF}$ , of about  $-150$  dBc/Hz at 312.5 MHz rises by  $20 \log 64 = 36$  dB within the loop bandwidth as it reaches the output. Thus,  $f_{BW}$  must be chosen so as to minimize the sum of  $64 S_{REF} f_{BW}$  and the shaped VCO phase noise. This PLL design chooses  $f_{BW} = 20$  MHz.

In order to achieve a wide bandwidth with acceptable spur levels, we modify the RF synthesizer architecture introduced in [23] for operation with  $f_{REF} = 312.5$  MHz and  $f_{VCO} = 20$  GHz. Shown in Fig. 16, the loop consists of an XOR phase detector (PD), a master-slave sampling filter (MSSF), a VCO, and a divider chain. As described in [23], the master-slave sampling action yields a small ripple on the control line and hence low spurs at the output. Owing to a closed-loop bandwidth of 20 MHz, the phase noise requirement for the  $LC$  VCO is greatly relaxed, allowing the oscillator power to be as low as 3.5 mW. Implemented as an  $LC$  oscillator with complementary cross-coupled transistors, the VCO exhibits a phase noise of  $-119$  dBc/Hz at 10-MHz offset, contributing roughly the same amount of jitter as the reference. Since PSS simulations in Cadence do not converge for the PLL, we have used transient noise simulations to obtain an rms jitter of 169 fs for the entire PLL circuit (excluding the reference noise).

## VIII. EXPERIMENTAL RESULTS

The PAM4 TX has been fabricated in TSMC's 45-nm digital CMOS technology. Fig. 17 shows a photograph of the die, whose 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 have been performed with a 1-V supply.

The overall TX consumes 44 mW. Table I shows the measured breakdown of the power consumption at 80 Gb/s. To separate the power of the clock distribution from the PLL, we simulate the divider chain in two cases: 1) while it drives the data path and 2) while it does not. The difference between the power values, 4.1 mW, is that necessary for clock distribution.

Fig. 18 shows the measured TX output in the NRZ mode at 40 Gb/s. Fig. 19 shows the output in the PAM4 mode at 40 Gb/s and 80 Gb/s. The differential voltage swing is 630 mV<sub>pp</sub>. The use of a 1-V supply for the entire system limits the output swing to about 630 mV. If the output driver supply is raised to 1.2 V and the tail currents in Fig. 12


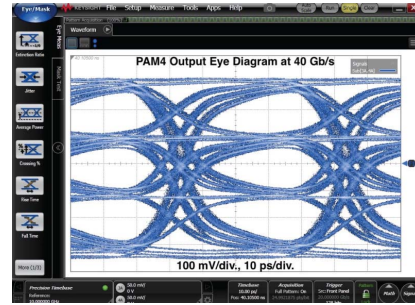


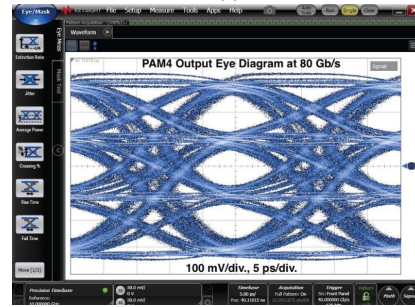
Fig. 17. Die photograph.

TABLE I  
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



(a)



(b)

Fig. 19. PAM4 output eye diagrams. (a) At 40 Gb/s. (b) At 80 Gb/s.

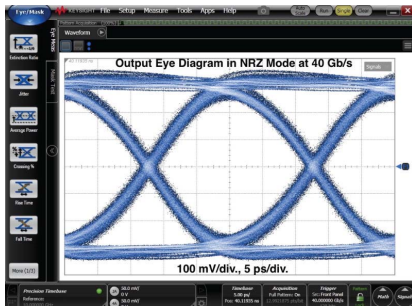


Fig. 18. Output eye diagram in NRZ mode at 40 Gb/s.

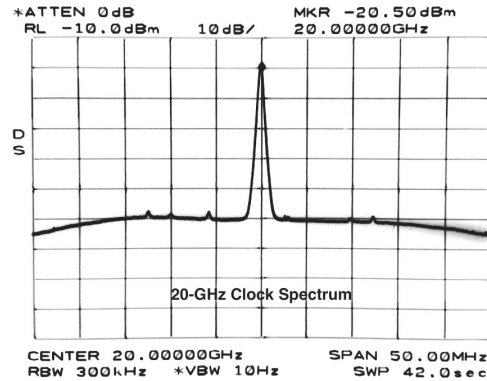


Fig. 20. Spectrum of 20-GHz clock.

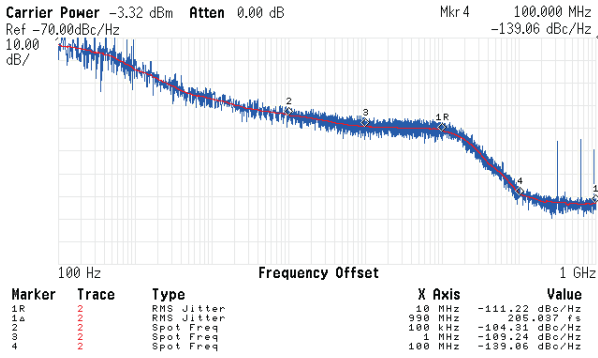
to 24 mA, the swing can reach 1.2 V. The data pattern is PRBS7. The vertical eye opening is 170 mV, the horizontal eye opening is 0.56 UI for the middle eye and 0.43 UI for the top and bottom eyes. The output bit pattern has been captured and checked against the input data to verify correct serialization.

The linearity of the PAM4 waveform is quantified by the “ratio of level mismatch” (RLM) [4], defined as the smallest eye height divided by one-third of the total eye height. To measure the RLM, the input data pattern is chosen so that the output PAM4 waveform contains 10 symbols with each lasting for 16 UI [4], [10]. Our measured RLM is around 99%, exceeding the 92% specification [4].

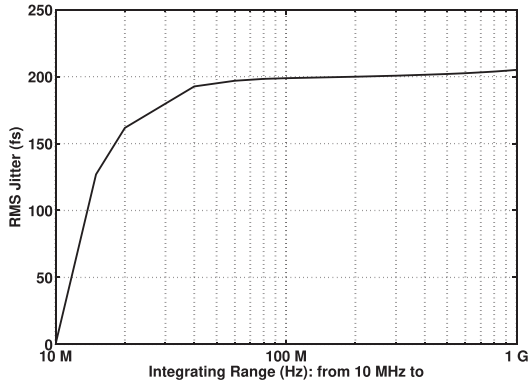
The 20-GHz clock generated by the PLL has also been characterized. The measured spectrum is shown in Fig. 20. The reference spurs are at  $-45$  dBc. Fig. 21(a) plots the measured phase noise of the 10-GHz clock. Due to our equipment limitation, the maximum offset is 1 GHz, but

we note from Fig. 21(b) that the integrated jitter reaches a plateau of 200 fs beyond approximately 200 MHz. In fact, noting that the phase noise is around  $-140$  dBc/Hz for offsets greater than 200 MHz, we observe that the range from 1 GHz to 5 GHz (the Nyquist rate) contributes  $[(4 \text{ GHz} \times 10^{-14})^{1/2}/2\pi] \times 100 \text{ ps} \approx 100 \text{ fs}$ , which, combined with the 205-fs value found in Fig. 21(a), amounts to 228 fs. That is, the phase noise beyond 1 GHz is negligible. This is also verified by simulation of the data path, including the output driver, and observing a flat phase noise up to 10 GHz.

To examine the effect of mismatches in  $\phi_1$ – $\phi_4$ , we apply the input data so as to create a 20-GHz periodic 0101 NRZ sequence at the TX output. Shown in Fig. 22, a spur level of  $-41$  dBc at 10-GHz offset in the single-ended output indicates a deterministic jitter of 100 fs<sub>rms</sub> jitter due to mismatches



(a)



(b)

Fig. 21. (a) Phase noise profile. (b) Relation of jitter and integrating range of 20-GHz clock divided by two externally.




Fig. 22. Measured spectrum of single-ended output delivering 20-GHz 0101 NRZ sequence.

among  $\phi_1$ – $\phi_4$  and within the 4-to-1 MUX. The relation between the spur and the jitter is obtained in Section VII.

Table II compares our measured performance with that of the prior art. We note that, if the PLL power consumption is excluded, our design achieves a nearly six-fold improvement in power efficiency. Even if we prorate the power consumption of our output DAC from 13.7 mW to about 32 mW to account for the larger output swing of 1.2  $V_{pp,d}$  in [7], our power efficiency is still higher by approximately a factor of 4 (excluding the PLL). Even though our prototype does not

TABLE II  
PERFORMANCE SUMMARY

	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 <sup>2</sup> )	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.




Fig. 23. Equivalent circuit of CML PAM4 output driver.

include FFE, the discussion in Section IV-B shows that adding FFE would entail negligible power penalty.

## IX. CONCLUSION

The power efficiency of ultrahigh-speed PAM4 TXs can be improved by means of techniques such as charge steering, latchless multiplexers, direct multi-phase multiplexers, frequency dividers with a 25% output duty cycle, and type-I PLL using MSSF. This paper has demonstrated an 80-Gb/s PAM4 TX achieving considerably higher efficiency than the prior art.

## APPENDIX

For the 2-bit DAC shown in Fig. 12, we construct the equivalent circuit in Fig. 23, where  $R_T = R_L = 50 \Omega$  represents the onchip termination and the load, respectively,  $r_o$  is the output impedance of each branch,  $k = 0, 1, 2, 3$ , and  $N = 3$ . We have

$$V_{out}(k) = \frac{(V_{DD} - I_0 r_o) r_o R_T (N - 2k)}{2r_o^2 + 1.5N r_o R_T + k(N - k) R_T^2}. \quad (2)$$

The levels corresponding to  $k = 1$  and  $k = 2$  exhibit integral nonlinearity (INL). Passing a straight line through the end points, finding its value at  $k = 2$ , subtracting it



from  $V_{\text{out}}(2)$ , and normalizing the result to the full scale,  $6(V_{\text{DD}} - I_0 r_o)R_T / (2r_o + 4.5R_T)$ , we obtain the INL as

$$\text{INL} = \frac{R_T^2}{6r_o^2 + 4.5r_oR_T + 6R_T^2}. \quad (3)$$

In this paper,  $r_o \approx 300 \Omega$ , yielding  $\text{INL} = 0.33\%$ , a negligible amount.

## REFERENCES

- [1] C. Menolfi *et al.*, "A 25 Gb/s PAM4 transmitter in 90 nm CMOS SOI," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 72–73.
- [2] V. Stojanovic *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.
- [3] B. Garlepp *et al.*, "A 1-10 Gbps PAM2, PAM4, PAM2 partial response receiver analog front end with dynamic sampler swapping capability for backplane serial communications," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2005, pp. 376–379.
- [4] *IEEE P802.3bs 400 GbE Task Force*. Accessed: Mar. 2015. [Online]. Available: <http://www.ieee802.org/3/bs/>
- [5] Y. Frans *et al.*, "A 56-Gb/s PAM4 wireline transceiver using a 32-way time-interleaved SAR ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 1101–1110, Apr. 2017.
- [6] P. J. Peng, J. F. Li, L. Y. Chen, and J. Lee, "A 56Gb/s PAM-4/NRZ transceiver in 40 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 110–111.
- [7] G. Steffan *et al.*, "A 64 Gb/s PAM-4 transmitter with 4-Tap FFE and 2.26 pJ/b energy efficiency in 28 nm CMOS FDSOI," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2017, pp. 116–117.
- [8] T. O. Dickson, H. A. Ainspan, and M. Meghelli, "A 1.8 pJ/b 56 Gb/s PAM-4 transmitter with fractionally spaced FFE in 14 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2017, pp. 118–119.
- [9] K. Gopalakrishnan *et al.*, "A 40/50/100 Gb/s PAM-4 ethernet transceiver in 28 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2016, pp. 62–63.
- [10] M. Bassi, F. Radice, M. Bruccoleri, S. Erba, and A. Mazzanti, "A high-swing 45 Gb/s hybrid voltage and current-mode PAM-4 transmitter in 28 nm CMOS FDSOI," *IEEE J. Solid-State Circuits*, vol. 51, no. 11, pp. 2702–2715, Nov. 2016.
- [11] A. Nazemi *et al.*, "A 36 Gb/s PAM4 transmitter using an 8 b 18 GS/S DAC in 28 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 58–59.
- [12] J. Kim *et al.*, "A 16-to-40 Gb/s quarter-rate NRZ/PAM4 dual-mode transmitter in 14 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 60–61.
- [13] J. Lee, P.-C. Chiang, and C.-C. Weng, "56 Gb/s PAM4 and NRZ SerDes transceivers in 40 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2015, pp. 118–119.
- [14] 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.
- [15] 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.
- [16] Y. Lu, K. Jung, Y. Hidaka, and E. Alon, "Design and analysis of energy-efficient reconfigurable pre-emphasis voltage-mode transmitters," *IEEE J. Solid-State Circuits*, vol. 48, no. 8, pp. 1898–1909, Aug. 2013.
- [17] C.-K. K. Yang, R. Farjad-Rad, and M. A. Horowitz, "A 0.5- $\mu\text{m}$  CMOS 4.0-Gbit/s serial link transceiver with data recovery using oversampling," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 713–722, May 1998.
- [18] Y. Chang, A. Manian, L. Kong, and B. Razavi, "A 32-Gb/s 40-mW CMOS NRZ transmitter," in *Proc. IEEE Custom Integr. Circuits Conf.*, Apr. 2018, pp. 1–4.
- [19] 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.
- [20] B. Razavi, K. F. Lee, and R. H. Yan, "Design of high-speed, low-power frequency dividers and phase-locked loops in deep submicron CMOS," *IEEE J. Solid-State Circuits*, vol. 30, no. 2, pp. 101–109, Feb. 1995.
- [21] I. Fabiano, M. Sosio, A. Liscidini, and R. Castello, "SAW-less analog front-end receivers for TDD and FDD," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 82–83.
- [22] A. Manian and B. Razavi, "A 40-Gb/s 14-mW CMOS wireline receiver," *IEEE J. Solid-State Circuits*, vol. 52, no. 9, pp. 2407–2421, Sep. 2017.
- [23] 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.
- [24] 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.



**Yikun Chang** (S'16) received the B.S. degree in microelectronics from Peking University, Beijing, China, in 2013, and the M.S. degree in electrical engineering from the University of California at Los Angeles, Los Angeles, CA, USA, in 2015, where she is currently pursuing the Ph.D. degree in circuits and embedded systems track.

Her current research interests include low-power techniques in wireline transceivers.

Ms. Chang was a recipient of the China National Scholarship in 2012 and the Analog Devices Outstanding Student Designer Award in 2016.



**Abishek Manian** (S'09–M'18) received the B.E. degree in electronics and telecommunication from the University of Mumbai, Mumbai, India, in 2011, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles (UCLA), Los Angeles, CA, USA, in 2013 and 2016, respectively.

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

Dr. Manian is a member of the Editorial Review Board for IEEE SOLID-STATE CIRCUITS LETTERS. He 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 in 2016 awarded by the Electrical Engineering Department, UCLA. He also serves as a reviewer for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I AND II, and the IEEE International Symposium on Circuits and Systems.



**Long Kong** (S'15–M'16) received the B.E. degree in microelectronics from Shanghai Jiao Tong University, Shanghai, China, in 2011, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles, Los Angeles CA, USA, in 2013 and 2016, respectively.

In 2016, he joined Oracle, Santa Clara, CA, USA, as a Senior Hardware Engineer working on high-speed SerDes transceivers. He is currently an RFIC Design Engineer with Apple, Cupertino, CA, USA.

His current research interests include frequency synthesizers, clock and data recovery for data communication systems, and wireless transceivers.

Dr. Kong was a recipient of the Qualcomm Innovation Fellowship in 2013–2014, the Analog Devices Outstanding Student Designer Award in 2015, and the Broadcom Fellowship in 2015–2016.





**Behzad Razavi** (S'87–M'90–SM'00–F'03) 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.

From 1992 to 1994, he was an Adjunct Professor with Princeton University, Princeton, NJ, USA, and Stanford University, Stanford, CA, USA, in 1995. He was with the AT&T Bell Laboratories, Murray Hill, NJ, USA, and Hewlett-Packard Laboratories, Palo Alto, CA, USA. Since 1996, he has been an

Associate Professor and subsequently a Professor in electrical engineering with the University of California at Los Angeles, Los Angeles, CA, USA. 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) and has edited *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996) and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003). His current research interests include wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters.

Dr. Razavi is a member of the U.S. Academy of Engineering. He served on the Technical Program Committees of the International Solid-State Circuits Conference (ISSCC) from 1993 to 2002 and very large-scale integration Circuits Symposium from 1998 to 2002. He has served as an IEEE Distinguished Lecturer. He was a recipient of the Beatrice Winner Award for Editorial Excellence at the 1994 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, and 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, and the CICC Best Invited Paper Award in 2009 and in 2012, the 2012 Donald Pederson Award in Solid-State Circuits, the American Society for Engineering Education PSW Teaching Award in 2014, and the 2017 IEEE CAS John Choma Education Award. He was a co-recipient of the 2012 and the 2015 VLSI Circuits Symposium Best Student Paper Awards and the 2013 CICC Best Paper Award. He was also a co-recipient of both the Jack Kilby Outstanding Student Paper Award and the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC. He was also recognized as one of the top 10 authors in the 50-year history of ISSCC. He has also served as a Guest Editor and an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and *International Journal of High Speed Electronics*.