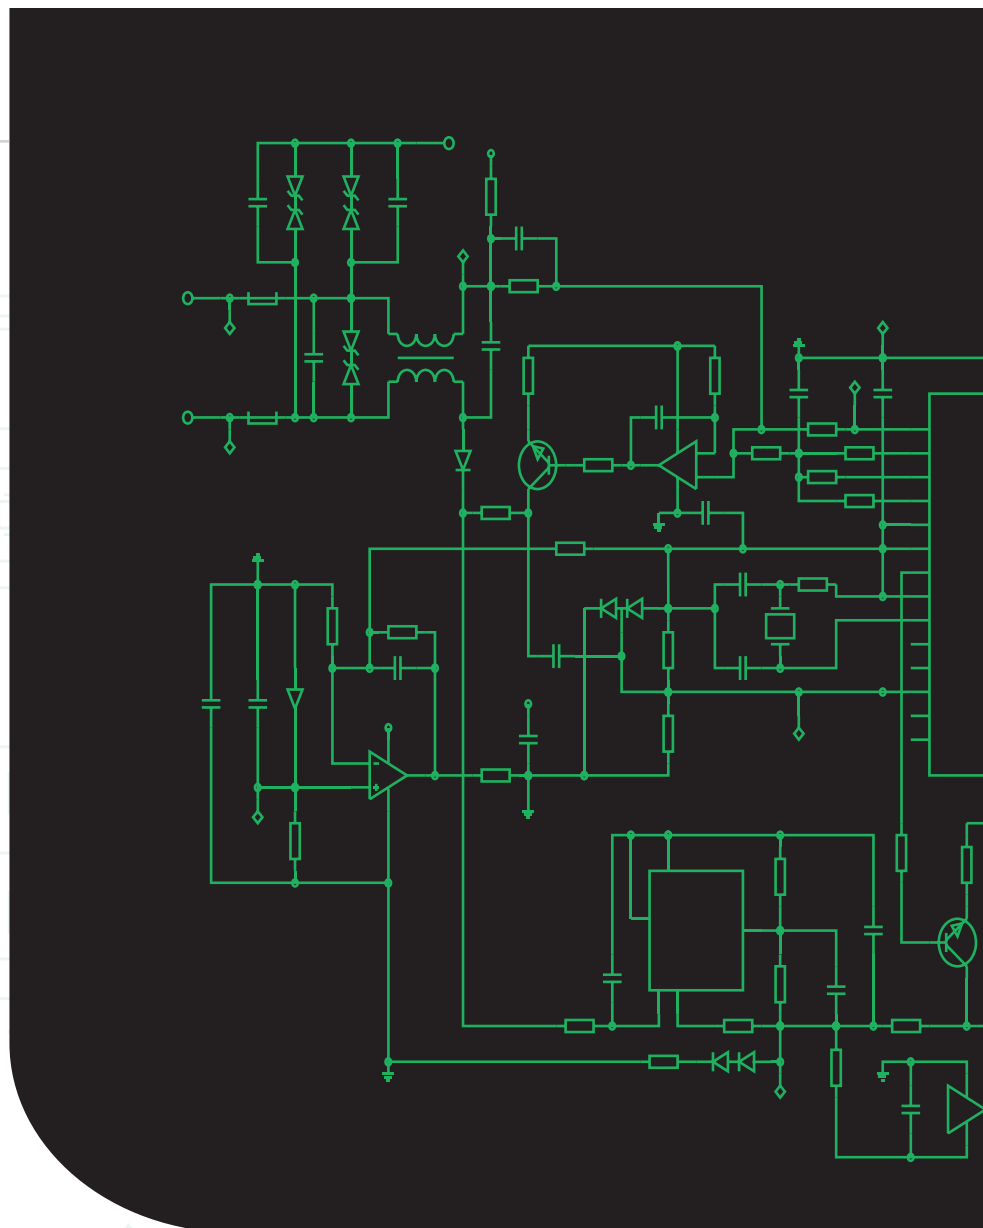


THE POWER CONSUMPTION OF broadband circuits used in wireline systems becomes increasingly more critical as higher speeds are sought. This article presents a number of design techniques that greatly relax the tradeoffs between the speed and power consumption of functions such as multiplexers (MUXs), frequency dividers, and equalizers. Examples include quadrature multiplexing, charge steering, and feedforward techniques. The concepts have been demonstrated in CMOS transceivers up to 56 GHz.

INTRODUCTION

The problem of power consumption assumes new dimensions as communication circuits target greater data rates. The reasons are many, ranging from package cost and heat removal issues to environmental impact. For example, data centers—each drawing some megawatts—continue to pose daunting challenges as they accommodate higher speeds and as their number also keeps rising. Considerable efforts are therefore expended on alleviating the speed-power tradeoffs in broadband (wireline)



©SHUTTERSTOCK.COM/ANEK KHUMPHANG

Breaking the Speed-Power Tradeoffs in Broadband Circuits

Reviewing design techniques for transceivers up to 56 GHz.

BEHZAD RAZAVI

Digital Object Identifier 10.1109/MNANO.2022.3160770
Date of current version: 11 April 2022

The power consumption of broadband circuits used in wireline systems becomes increasingly more critical as higher speeds are sought.

sary for these functions are provided by a phase-locked loop (PLL). The Tx receives multiple low-speed data streams at D_{in} , converts them to a single data sequence by the serializer, preconditions the result by the FFE, and delivers it to the channel.

Figure 1(b) depicts a basic Rx comprising a continuous-time linear equalizer (CTLE), a clock and data recovery (CDR) circuit, a decision-feedback equalizer (DFE), and a demultiplexer (DMUX). The CTLE and the DFE compensate for the channel imperfections experienced by the data, the CDR recovers the clock from D_{in} , and the DMUX produces the original data streams. Typical links target a bit error rate of 10^{-12} to 10^{-14} .

We expect that the Tx and Rx building blocks operating at high speeds face the most severe power tradeoffs and merit the greatest attention. Moreover,

we surmise that techniques that improve the *speed* of a given circuit can also relax these tradeoffs.

The foregoing architectures are suited to random binary information, also known as *nonreturn-to-zero (NRZ)* data. At higher speeds, it is preferable to employ four-level pulse-amplitude (PAM4) modulation. Most of the techniques described here are applicable to PAM4 transceivers as well. Our quantitative results are reported for 40- and 28-nm CMOS technologies.

HIGH-SPEED LOGIC

Many functions in broadband transceivers incorporate digital circuits, such as latches, flip-flops (FFs), and frequency dividers. The environment in Figure 1, for example, employs these structures in the serializer, the FFE, the PLL, the CDR, the DFE, and the DMUX. The

circuits employed in such an environment [1]–[24].

This article describes a number of techniques that significantly ease speed-power tradeoffs in communication circuits. The concepts have been experimentally demonstrated in CMOS technology in transmitters and receivers (Rxs) up to 56 GHz [5], [13]–[15], [18], saving power by as much as a factor of 10.

TRANSCIEVER ENVIRONMENT

To appreciate the role of the circuits presented here, we briefly study a generic broadband transceiver architecture. Shown in Figure 1(a) is a transmitter (Tx) whose signal path consists of a serializer, a feedforward equalizer (FFE), and an output driver. The various clock frequencies and phases neces-

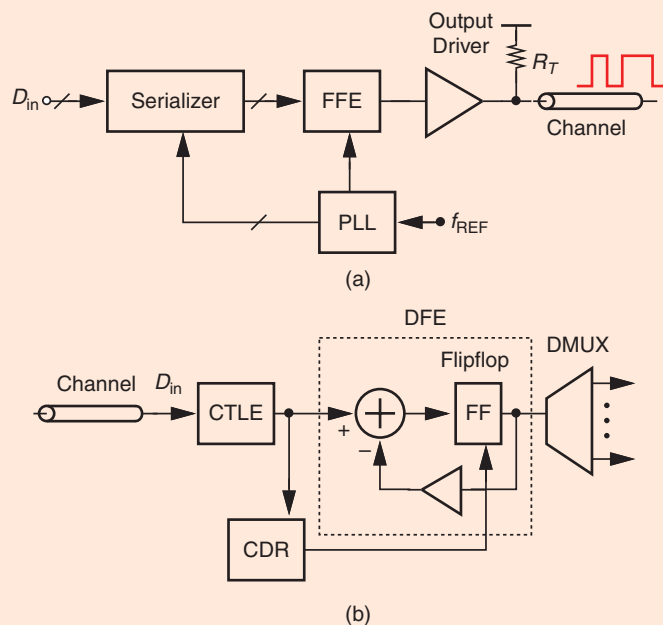


FIGURE 1 (a) and (b) a generic wireline Tx (a) and Rx (b).

power consumed by these logical functions thus becomes critical.

For speeds up to about 20 GHz in 28-nm CMOS technology, we can

consider two low-power logic styles, namely, clocked CMOS (C^2 MOS) circuits [25] and true single-phase clock (TSPC) topologies [26]. Shown in

Figure 2(a), the former can be viewed as an inverter that is enabled when the clock, CK, is high and as a state storage circuit (a latch) when CK is low. To construct an FF, we cascade two such latches and swap the clock phases in the second stage [Figure 2(b)]. The circuit consumes power in both the signal path and the clock path as the node capacitances charge and discharge. Each component of the power can be expressed in the form of $f_1 C_1 V_{DD}^2$, where f_1 is the rate at which a given node toggles between zero and the supply voltage, V_{DD} , and C_1 is the total capacitance at that node. Note that this topology requires four clocked transistors. While drawing low power, the C^2 MOS FF must deal with the transparency that occurs from D_{in} to Q when CK and \overline{CK} change in a finite time [25].

Shown in Figure 3(a) is a TSPC FF [26]. When CK is low, node B tracks D_{in} , while A remains high and X maintains the previous state. When CK goes high, we have $A = \overline{B}$ and $X = \overline{A}$. This FF, too, requires four clocked transistors.

The TSPC FF of Figure 3(a) entails a race condition that can potentially create a large glitch at the output [27]. As depicted in Figure 3(b), if B and X are high and CK rises, both M_4 and M_3 are activated. Since M_3 turns on while A is still high (around $t = t_1$), X begins to fall until A has dropped enough to turn M_1 on. During this time, X experiences a glitch that may be misinterpreted by subsequent stages.

Another TSPC FF variant is shown in Figure 4 [26]. Using only two clocked transistors, this topology is suited to applications that drive a large number of FFs at a high speed and must therefore deal with a large capacitance in the clock path. The drawback is that the logical levels at A_1 and B_2 are degraded, reducing the switching speeds of M_1 and M_5 , respectively.

For higher speeds, we turn to current-mode logic (CML). Illustrated in Figure 5 is a CML latch consisting of an input differential pair and a positive-feedback pair. The former senses the input when CK is high, and the latter regeneratively amplifies $V_X - V_Y$ when CK goes low. The CML speed advantage arises from moderate voltage swings

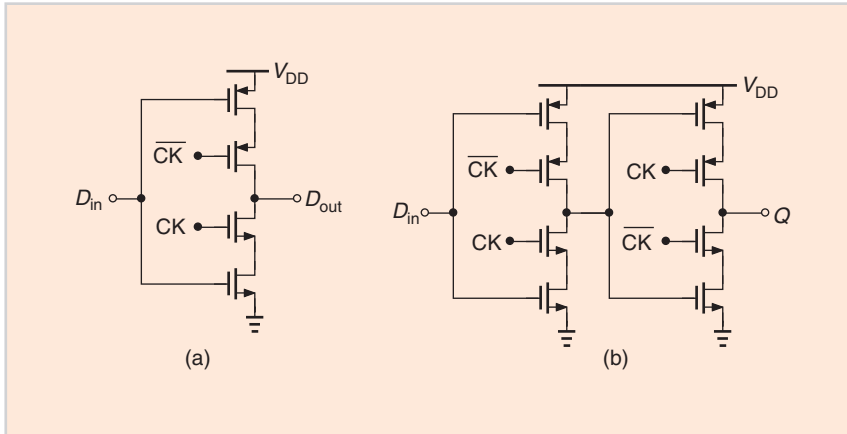


FIGURE 2 (a) A C^2 MOS latch (a). (b) An FF.

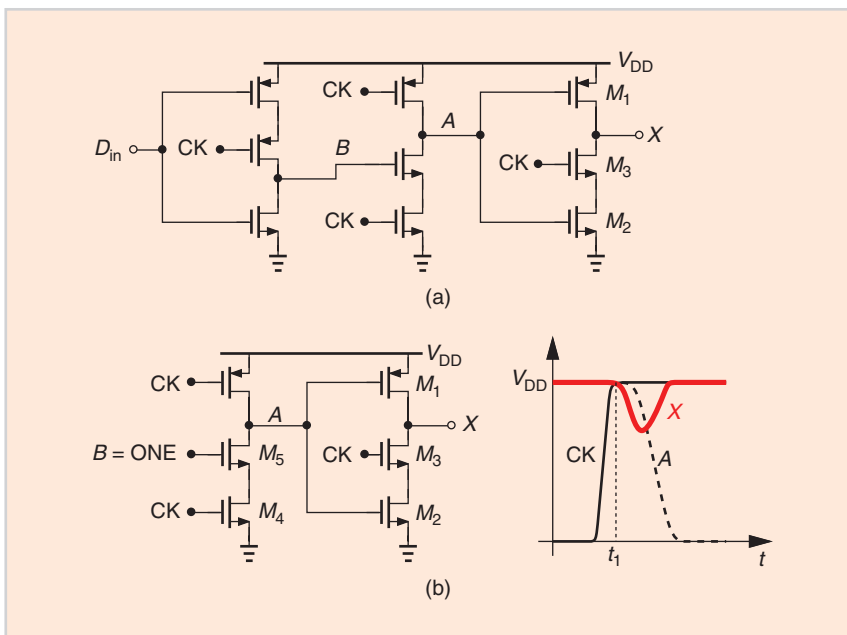


FIGURE 3 (a) A TSPC FF and (b) a glitch issue in TSPC logic.

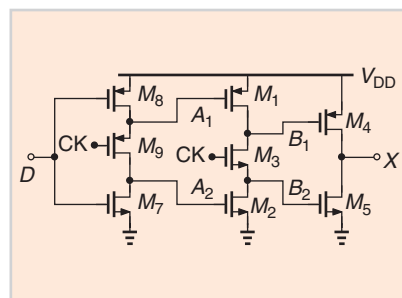


FIGURE 4 A TSPC FF with split paths.

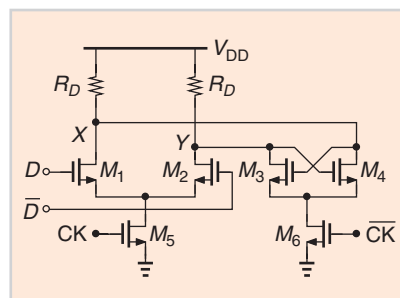


FIGURE 5 A CML latch.

(typically, 400 mV_{pp} at X or at Y) and the regenerative action of M_3 and M_4 . The cost is static power consumption.

A logic style that achieves a higher speed than C²MOS and TSPC circuits and consumes less than CML topologies is based on “charge steering” [14]. Shown in Figure 6 is the basic structure, which resembles a differential pair but with the tail current source replaced by a “charge source” and the loads replaced with capacitances and switches. In the precharge mode, nodes X and Y are reset to V_{DD} , while C_T is discharged. In the evaluation mode, C_T is switched to P , drawing current from M_1 and M_2 . According to the value of V_{in} , the difference between I_{D1} and I_{D2} creates a differential output component until C_T charges up and M_1 and M_2 turn off. We observe that the circuit can act as an amplifier and/or a latch. The speed and power advantages of this approach stem from the moderate output swings, about 400 mV_{pp} at X and at Y . It can be shown that, for a given data rate, the power consumed by charge-steering topologies is roughly a factor of 1.4π lower than that of their CML counterparts [14].

One drawback of charge steering is that $V_X - V_Y$ in Figure 6 exhibits an RZ behavior in the precharge mode, requiring that the circuits be properly architected. Nonetheless, this design paradigm has proven useful in various Tx and Rx functions in 65-, 45-, and 28-nm technologies [5], [13]–[15], [18].

When optimizing a complex transceiver for power, we must bear in mind that, for a given function and a given logic style, the speed-power tradeoff is linear up to some frequency and non-linear beyond. Conceptually illustrated in Figure 7, this behavior manifests itself in CMOS (i.e., TSPC and C²MOS) stages as the data rate and/or the clock frequency reach f_1 , above which the transistors must be chosen excessively wide so as to improve the speed by a small amount. CML circuits, on the other hand, experience such diminishing returns at a higher frequency, f_2 . (We assume that CML stages are custom designed for each frequency.) Charge-steering logic offers a solution between the two.

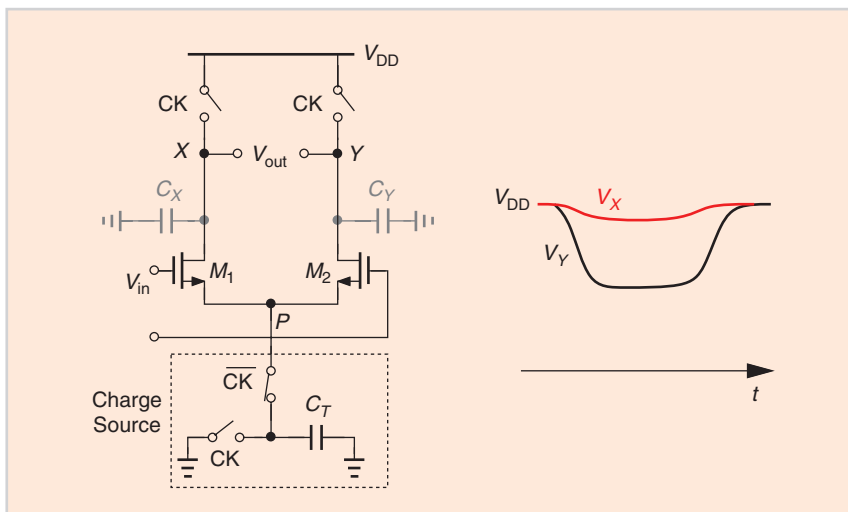


FIGURE 6 A charge-steering latch and its waveforms.

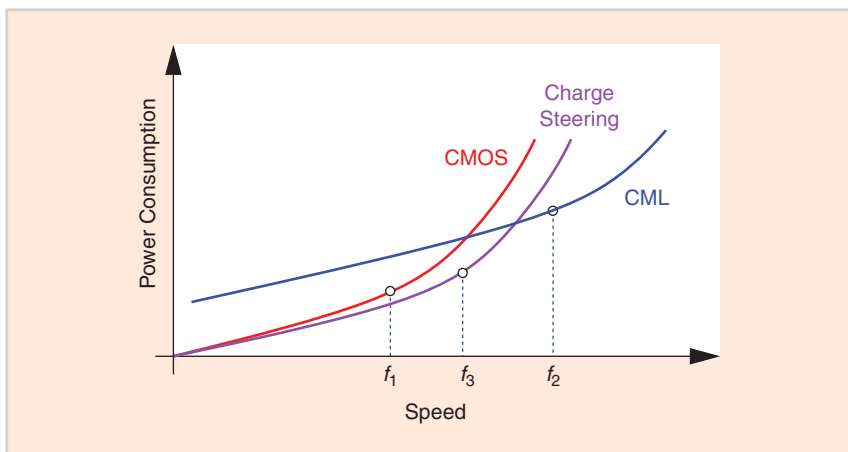


FIGURE 7 A conceptual illustration of power-speed tradeoffs in different logic families.

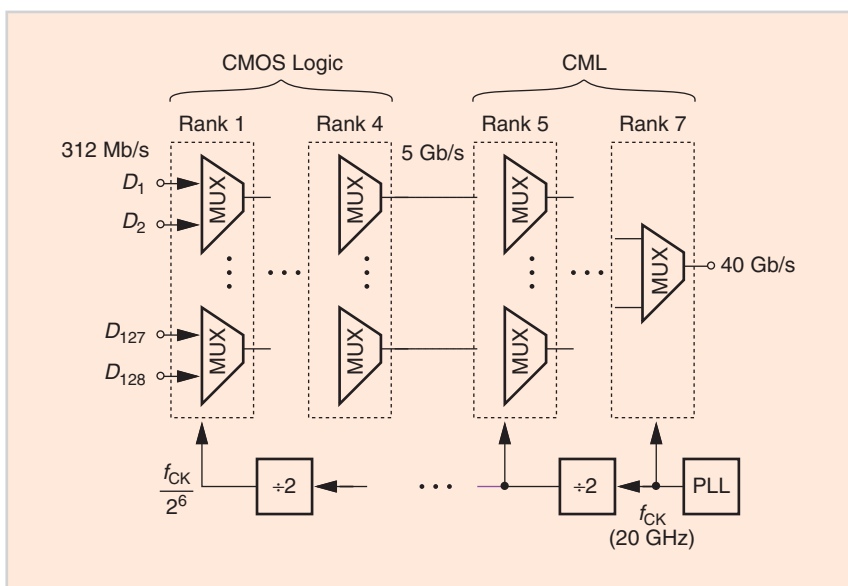


FIGURE 8 A binary MUX chain.

At higher speeds, it is preferable to employ four-level pulse-amplitude modulation.

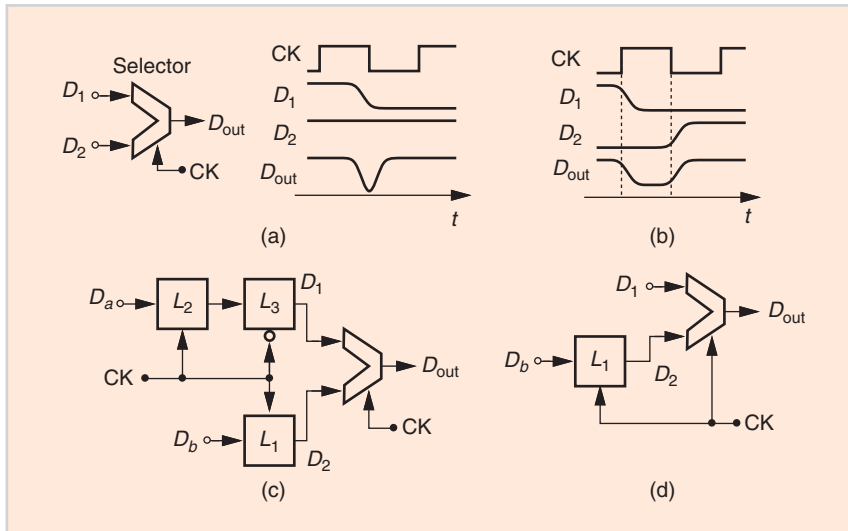


FIGURE 9 (a) A basic binary selector and its waveforms, (b) the proper input data alignment, (c) three-latch MUX, and (d) one-latch MUX.

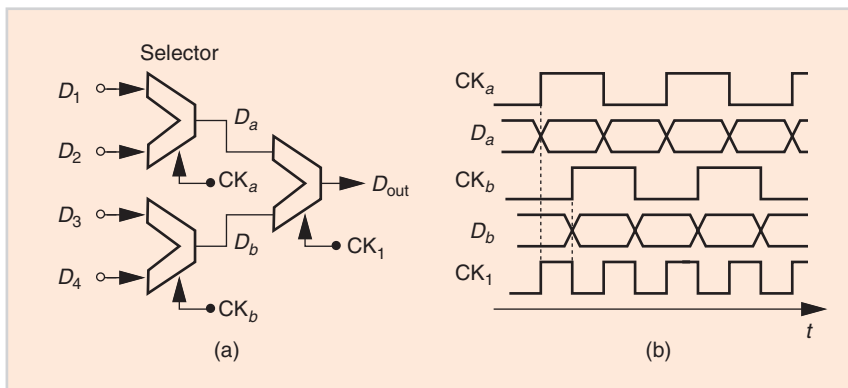


FIGURE 10 A binary MUX tree with no latches and its waveforms.

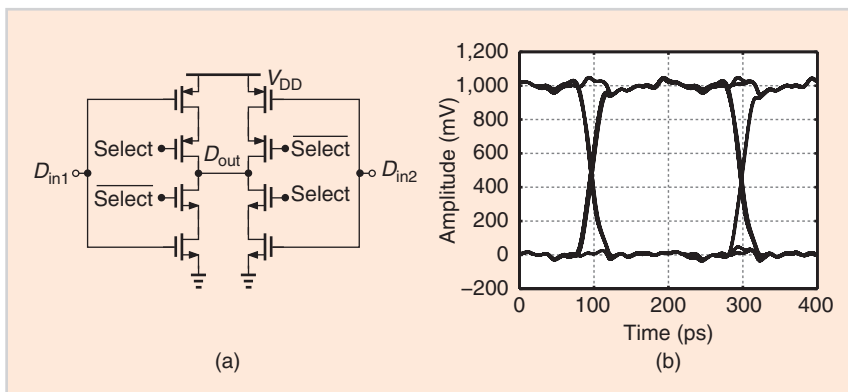


FIGURE 11 (a) A selector cell and (b) its simulated output eye diagram at 5 Gb/s.

LOW-POWER MUX DESIGN

The serializer in Figure 1(a) employs a large number of latches and selectors so as to aggregate more than 100 input data streams at D_{in} . Based on the speed barriers quantified in the previous section, we envision a binary-tree MUX chain such as that in Figure 8, where CMOS (e.g., C²MOS) logic performs serialization up to about 5 Gb/s and CML circuits deal with higher speeds. Each rank is driven by a clock frequency equal to its input data rate. The clocks are thus obtained by a cascade of $\div 2$ stages. This MUX chain requires 2^7 MUX cells, drawing significant power in its clock path. Moreover, the CML section also consumes a great deal in its data path. As explained later, charge steering proves beneficial here as well.

In its simplest form, a MUX cell selects one of its two inputs according to a command (a clock) [Figure 9(a)]. If, however, D_1 , D_2 , and CK bear an arbitrary timing relationship, the output can suffer from excessively narrow pulses or glitches. For example, suppose CK is high and $D_{out} = D_1$, while D_2 is high. If D_1 begins to fall before the clock does, then D_{out} tracks this change and then becomes equal to D_2 , experiencing a glitch.

This issue is resolved by guaranteeing that D_1 and D_2 arrive at proper times, e.g., D_1 changes only on the rising edge of CK , and D_2 changes only on the falling edge [Figure 9(b)]. This is accomplished as shown in Figure 9(c), where latches L_2 and L_3 retime D_a and latch L_1 retimes D_b . Thus, so long as the latches are not metastable, D_1 and D_2 change as prescribed in Figure 9(b). In the interest of power consumption, we can reduce the circuit to the one-latch topology in




FIGURE 12 Tapering in the CMOS MUX chain.

Figure 9(d) if proper timing of D_1 is guaranteed by the preceding stage, as is the case in a cascade of such MUXs.

CMOS MUX DESIGN

The one-latch MUX topology still proves power hungry for the long chain of Figure 8, where the number of MUXs drops by a factor of two from one rank to the next, but the increase in speed at least doubles the power consumed by the cell. We then ask: is it possible to omit this latch as well? Indeed, one can design the serializer so that it requires no latches, hence reducing both the complexity and the power dramatically. We first recognize that the $\div 2$ stages in Figure 8 can generate quadrature phases, which can be utilized to ensure proper input timing for the MUX cells [18]. As depicted in Figure 10, two selectors in the same rank are driven by quadrature phases CK_a and CK_b so that D_a changes only on the edges of CK_a and D_b changes on the edges of CK_b . Consequently, the inputs to the next selector are properly offset in time, causing no glitches. This three-cell configuration serves as a four-to-one MUX and can be repeated to form a complete serializer [18].

The two-to-one selector cell in Figure 10 can be realized by C²MOS logic for speeds up to about 5 GHz. To minimize the power consumption in its clock path, we prefer to employ small transistors. Figure 11(a) depicts a simple, efficient topology, and Figure 11(b) depicts its simulated output eye diagram at 5 Gb/s [18]. This structure occupies a

small area, allowing short interconnects for the entire CMOS serializer.

CMOS serializer design begins with the last two-to-one selector rank. This C²MOS selector employs PMOS and NMOS widths equal to 2 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 two-to-one selector is scaled down by a factor of two from one MUX rank to

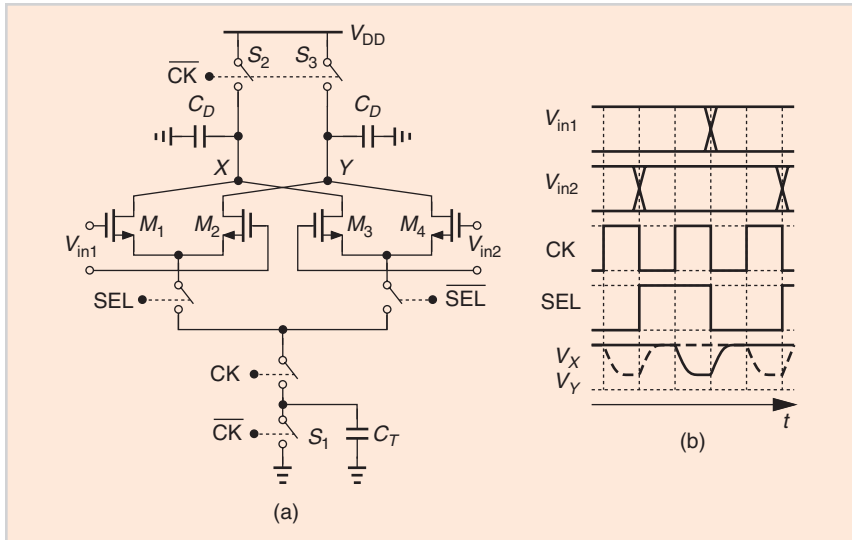


FIGURE 13 Basic charge-steering MUX.

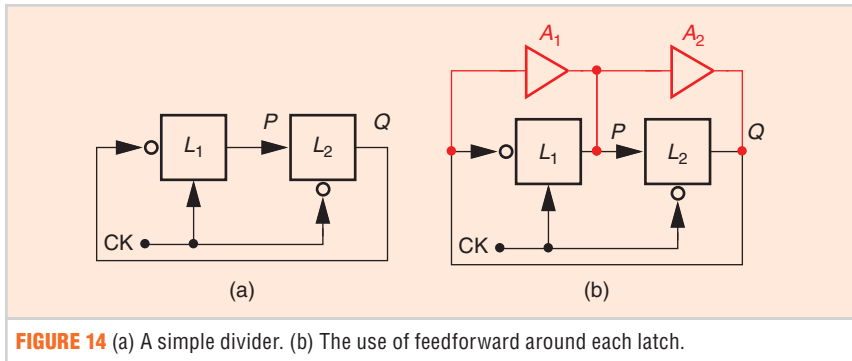


FIGURE 14 (a) A simple divider. (b) The use of feedforward around each latch.

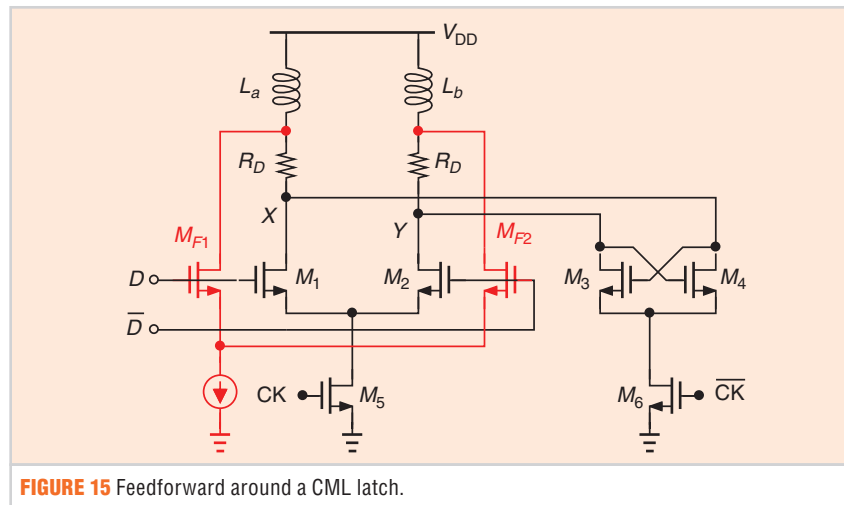


FIGURE 15 Feedforward around a CML latch.

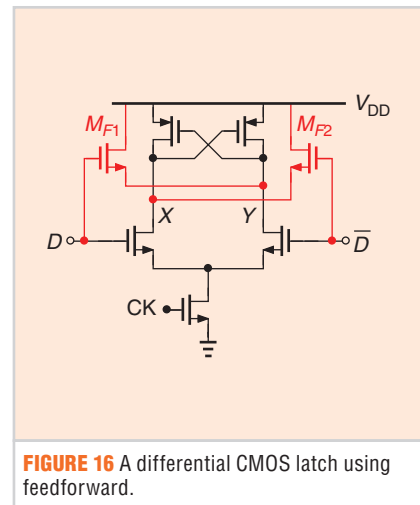


FIGURE 16 A differential CMOS latch using feedforward.

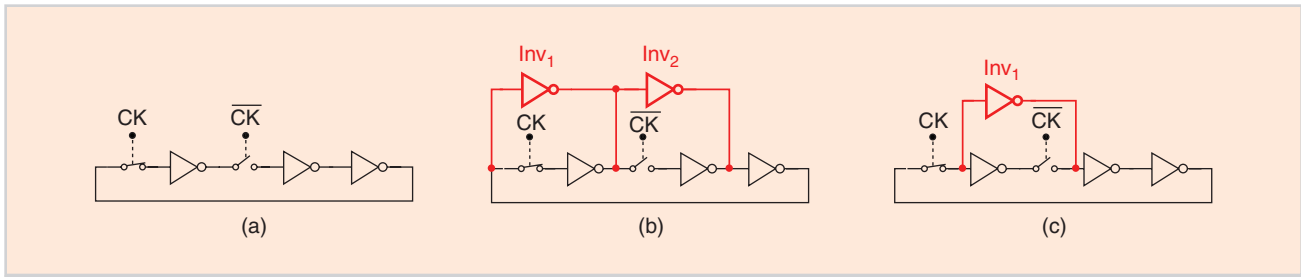


FIGURE 17 (a) A divider using dynamic latches. (b) Feedforward type 1i. (c) Feedforward type 2.

the rank preceding it, until a minimum allowable transistor width of 120 nm is reached (Figure 12). The entire 128-to-8 serializer draws 365 μ W in the data path.

CHARGE-STEERING MUX DESIGN

The binary-tree CML MUX chain in Figure 6 can potentially draw a high power. It is possible to implement rank 5 by means of charge steering. This eight-to-four MUX receives inputs at 5 Gb/s and delivers outputs at 10 Gb/s.

Let us extend the charge-steering stage of Figure 6 to form a two-to-one selector. Illustrated in Figure 13 [18], 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^2 MOS 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).

LOW-POWER FREQUENCY DIVIDERS

As exemplified by the serializer of Figure 8, wireline Tx's incorporate $\div 2$ stages to generate various clock frequencies (and phases). Similarly, PLLs providing the local oscillator waveforms in millimeter-wave radio frequency transceivers often follow the oscillator with one or more $\div 2$ circuits. It is therefore beneficial to examine methods of reducing their power consumption.

We introduce the concept of “feedforward” as a means of relaxing the speed-power tradeoffs of dividers. We describe three different topologies targeting the range of 25–60 GHz. We begin with the basic structure shown in Figure 14(a), where two latches form a loop with a net inversion. Output P changes on, for example, the rising edge of CK and Q on the falling edge. The speed of this circuit is limited by the delay through the latches. We now add feedforward paths A_1 and A_2 [Figure 14(b)], which are *not* clocked, i.e., they continuously inject P to Q and Q to P . As a result, a change in Q propagates to P *before* CK places L_1 in the sense mode, thereby reducing the apparent delay. The speed improvement can then be traded for power.

One issue facing feedforward is that the unlocked paths continue to operate even at low clock frequencies, thus overwhelming the latches and causing failure. That is, the upper end of the divider frequency range is improved at the cost of

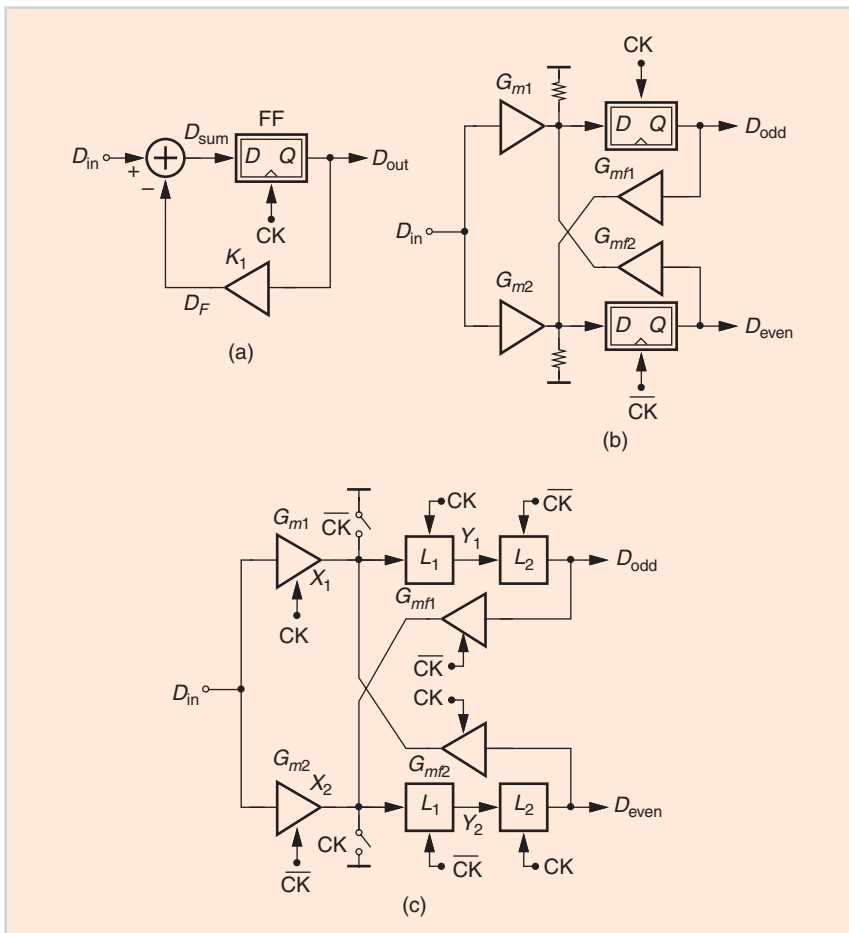


FIGURE 18 (a) The basic full-rate DFE, (b) half-rate DFE, and (c) charge-steering implementation of half-rate DFE.

the lower end. We must therefore ensure that the circuit meets the desired range across process, supply voltage, and temperature corners.

The first realization of the feedforward technique is illustrated in Figure 15 for a CML latch [28]. We observe that M_{F1} and M_{F2} bypass the clocked input differential pair and inject their outputs into the load inductors. This scheme is slightly different from that of Figure 14(b); the feedforward path does not lead to X and Y , allowing its current to flow only through L_a and L_b and hence saving voltage headroom. This divider has been employed in a 60-GHz wireless transceiver [29].

The second embodiment begins with a latch having complementary inputs and outputs and bypasses the circuit by source followers (Figure 16) [14]. In this case, M_{F1} and M_{F2} also provide additional pull-up strength at nodes X and Y . As explained in [14], this feedforward method raises the maximum speed by about 25% and, in fact, lowers the power consumption by 50% as well.

The third feedforward structure can be appreciated in the context of the $\div 2$ topology depicted in Figure 17(a), where each switch and the inverter following it form a latch. The third inverter is necessary for negative feedback, and hence, the toggling of the states. We can envision a number of feedforward paths here. For example, Figure 17(b) shows that Inv_1 and Inv_2 bypass their respective latches, and Figure 17(c) prescribes that an inverter can sense and inject signals at the *outputs* of the switches. Simulations of these circuits with layout parasitics indicate that the latter offers the optimum frequency range. (Two feedforward inverters raise

the *lower* end excessively in the fast-fast corner of the process.) As a result, the maximum input frequency rises from 55 to 68 GHz; the lower end is about 45 GHz. Note that the circuit requires no inductors. Incorporated in a 56-GHz fractional- N PLL [30], this stage draws about 2 mW. Without feedforward, the necessary power would be about three times as much.

LOW-POWER DFE DESIGN

High-speed DFEs have been under extensive development [6], [9], [31], [32]. For low-power operation, we explore the use of charge steering in such an environment. We begin with the simple, full-rate topology shown in Figure 18(a). Here, the previous bit, stored in the FF, is scaled by a factor of K_1 and subtracted from the present input, thereby removing some of the




FIGURE 20 A Tx die. DAC: digital–analog converter. (Source: [18].)

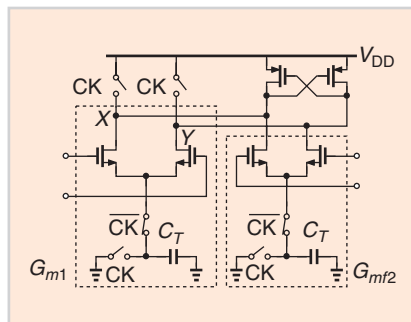


FIGURE 19 A charge-steering stage with a cross-coupled PMOS load.

TABLE 1 A 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	1,200	900	630
RLM	N/A	0.94	N/A	0.99
RMS jitter (fs)	688	290	318	205
Integ. range (MHz)	0.0001–1,000	0.5–8,000	N/A	10–1,000
Power (mW)	Exc.* 200 Inc.** 220	145***	101	25.8 44.1
Power Eff. (pJ/bit)	Exc.** 3.57 Inc.** 3.93	2.26***	1.8	0.32 0.55
Active Area (mm ²)	0.8*	N/A	0.035*	0.1

Max.: maximum; N/A: not applicable; Integ.: integer; Eff.: efficiency; Exc.: excluding; Inc.: including; SST: series source termination; RLM: ratio of level mismatch.
 *Excluding PLL power but including clock distribution.
 **Including PLL power and clock distribution.
 ***Without in-phase and quadrature clock generation.

This article presents a number of circuit techniques that alleviate the tradeoffs between speed and power in broadband transceivers.

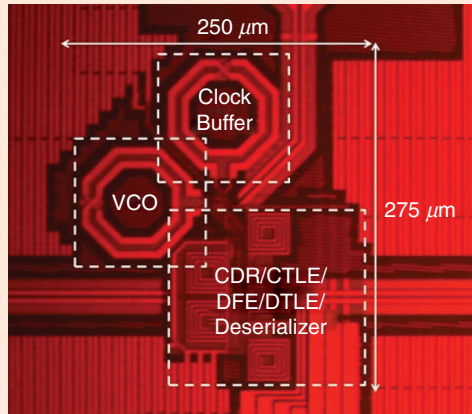


FIGURE 21 An Rx die. VCO: voltage-controlled oscillator. (Source: [5].)

intersymbol interference that the channel in Figure 1(b) introduces. This circuit's speed is limited by the following loop timing constraint:

$$t_{CK-Q} + t_{FB} + t_{setup} < T_b, \quad (1)$$

where the three terms on the left denote the FF clock-to-output delay, the feedback delay, and the FF setup time; T_b is the input bit period.

At high speeds, we prefer the half-rate scheme depicted in Figure 18(b) [32], where the FFs also act as a DMUX and generate half-rate data streams D_{odd} and D_{even} . This approach can employ charge steering in the FFs as well as in the G_m stages. Illustrated in Figure 18(c) [15], such an arrangement precharges summing junction X_1 , while G_{m1} and G_{m2} are reset. Next, CK changes, these two G_m stages begin to evaluate, and so does latch L_1 . The DFE thus draws no static power. It can be proved that the timing constraint in this case relaxes to

$$t_{CK-Q} + t_{setup} < T_b, \quad (2)$$

TABLE 2 A summary of measured Rx performance and comparison with the prior art.

REFERENCE	[5]	[6]	[19]	[7]	[9]	[8]	THIS WORK
Modulation	NRZ	PAM4	NRZ	PAM4	NRZ	PAM4	NRZ
Data rate (Gb/s)	56	56	60	64	56	56	56
Channel loss	18.4 dB* @ 28 GHz	24 dB** @ 14 GHz	21 dB** @ 30 GHz	16.8 dB*** @ 16 GHz	37.8 dB* @ 28 GHz	20.8 dB* @ 28 GHz	30 dB* @ 28 GHz 16.5 dB* @ 14 GHz 25 dB @ 28 GHz 13.5 dB @ 14 GHz
Horizontal eye (UI)	0.28 @ 10 ⁻⁹	0.25 @ 10 ⁻¹²	0.3 @ 10 ⁻¹²	0.19 @ 10 ⁻⁶	0.44 @ 10 ⁻¹²	0.19 @ 10 ⁻¹²	0.4 @ 10 ⁻¹²
Clock jitter (fs, RMS)	—	688 (100 Hz–1 GHz)	—	—	—	—	500 (100 Hz–14 GHz)
PRBS	15	7	7	Q 13	15	15	7
Power (mW)	Incl. \$ 141.7 Excl. \$\$ —	382	136	— 180	— 112	259	49.56 43.6
Power Eff. (pJ/bit)	Incl. \$ 2.53 Excl. \$\$ —	6.82	2.26	— 2.81	— 2	4.63	0.88 0.77
Area (mm ²)	1.4 [#]	1.26	2.03	0.32	0.053	0.51	0.102
Technology	28-nm CMOS	40-nm CMOS	65-nm CMOS	28-nm FDSOI	14-nm FIN-FET	65-nm CMOS	28-nm CMOS

FIR: finite-impulse response; IIR: infinite-impulse response; FINFET: fin field-effect transistor; RMS: root mean square, UI: unit interval.
 *Includes two-tap Tx FFE.
 **Includes three-tap Tx FFE.
 ***Includes four-tap Tx FFE.
 #Includes a Tx area.
 \$Includes clock gen.
 \$\$Excludes clock gen.

where t_{CK-Q} represents the delay from the clock edge to when G_{m1} and L_1 create a reasonable swing at Υ_1 [15].

The charge-steering DFE is further simplified by noting that L_2 and G_{mf1} (and L_4 and G_{mf2}) can be merged because they evaluate concurrently. For this reason, latches L_3 and L_4 can be omitted, a unique property of charge steering.

The G_m cells in Figure 18(c) can be realized in the basic charge-steering form depicted in Figure 6. Nonetheless, the performance is improved if a cross-coupled PMOS pair is added to the output nodes (Figure 19) [13]. The positive feedback thus afforded helps regenerate the output voltages if the common-mode level drops due to the charge drawn by the G_m stages. Charge-steering DFEs have been employed from 25 Gb/s [15] to 56 Gb/s [5].

DESIGN EXAMPLES

The concepts described in the previous sections have been employed in a number of designs [5], [13]–[15], [18]. We present two here.

The first is an 80-Gb/s PAM4 Tx realized in TSMC's 45-nm CMOS technology. Shown in Figure 20 is the die photograph. The measured performance is summarized in Table 1 along with that of the prior art.

The second is a 56-Gb/s NRZ Rx fabricated in TSMC's 28-nm CMOS technology. Figure 21 shows the die photograph, and Table 2 summarizes the performance.

CONCLUSION

This article presents a number of circuit techniques that alleviate the tradeoffs between speed and power in broadband transceivers. Described are charge steering, quadrature multiplexing, and feedforward-based frequency division. The methods have been realized in CMOS Tx and Rx functions ranging from 25 to 56 Gb/s.

ABOUT THE AUTHOR

Behzad Razavi (razavi@ee.ucla.edu) is with the Electrical and Computer Engineering Department, University of California, Los Angeles, Los Angeles, 90095, California.

REFERENCES

[1] P. Upadhyaya *et al.*, "A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with

a configurable ADC in 16nm FinFET," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2018, pp. 108–110, doi: 10.1109/ISSCC.2018.8310207.

[2] T. Ali *et al.*, "A 180mW 56Gb/s DSP-based transceiver for high-density IOs in data center switches in 7nm FinFET technology," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2019, pp. 118–120, doi: 10.1109/ISSCC.2019.8662523.

[3] J. Im *et al.*, "A 112Gb/s PAM-4 long-reach wireline transceiver using a 36-way time-interleaved SAR-ADC and inverter-based RX analog front-end in 7nm FinFET," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2020, pp. 116–118.

[4] T. Ali, E. Chen, H. Park, and R. Yousry, "6.2 A 460mW 112Gbps DSP-based transceiver with 38dB loss compensation for next generation data centers in 7nm FinFET technology," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2020, pp. 118–120, doi: 10.1109/ISSCC19947.2020.9062925.

[5] A. Atharav and B. Razavi, "A 56Gb/s 50mW NRZ receiver in 28nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2021, pp. 192–194.

[6] T. Shibasaki *et al.*, "3.5 A 56Gb/s NRZ-electrical 247mW/lane serial-link transceiver in 28nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2016, pp. 64–66, doi: 10.1109/ISSCC.2016.7417908.

[7] P. J. Peng, J. Li, L. Chen, and J. Lee, "6.1 A 56Gb/s PAM-4/NRZ transceiver in 40nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2017, pp. 110–111, doi: 10.1109/ISSCC.2017.7870285.

[8] B. Dehlaghi *et al.*, "A 1.41-pJ/b 56-Gb/s PAM-4 receiver using enhanced transition utilization CDR and genetic adaptation algorithms in 7-nm CMOS," *IEEE Solid-State Circuits Lett.*, vol. 2, no. 11, pp. 248–251, Nov. 2019, doi: 10.1109/LSSC.2019.2938677.

[9] J. Han, N. Sutardja, Y. Lu, and E. Alon, "Design techniques for a 60-Gb/s 288-mW NRZ transceiver with adaptive equalization and baud-rate clock and data recovery in 65-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3474–3485, Dec. 2017, doi: 10.1109/JSSC.2017.2740268.

[10] E. Depaoli *et al.*, "A 4.9pJ/b 16-to-64Gb/s PAM-4 VSR transceiver in 28nm FDSOI CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2018, pp. 112–114, doi: 10.1109/ISSCC.2018.8310209.

[11] A. Cevrero *et al.*, "6.1 A 100Gb/s 1.1pJ/b PAM-4 RX with Dual-Mode I-Tap PAM-4/3-Tap NRZ Speculative DFE in 14nm CMOS FinFET," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2019, pp. 112–114, doi: 10.1109/ISSCC.2019.8662495.

[12] A. Roshan-Zamir *et al.*, "A 56-Gb/s PAM4 receiver with low-overhead techniques for threshold and edge-based DFE FIR- and IIR-Tap adaptation in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 672–684, Mar. 2019, doi: 10.1109/JSSC.2018.2881278.

[13] 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, doi: 10.1109/JSSC.2017.2705913.

[14] J. W. Jung and B. Razavi, "A 25-Gb/s 5-mW CMOS CDR/deserializer," *IEEE J. Solid-State Circuits*, vol. 48, no. 3, pp. 684–697, Mar. 2013, doi: 10.1109/JSSC.2013.2237692.

[15] J. W. Jung and B. Razavi, "A 25 Gb/s 5.8 mW CMOS equalizer," *IEEE J. Solid-State Circuits*, vol. 50, no. 2, pp. 515–526, Feb. 2015, doi: 10.1109/JSSC.2014.2364271.

[16] J. Kim *et al.*, "A 112-Gb/s PAM4 transmitter with 3-Tap FFE in 10-nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2018, pp. 102–104, doi: 10.1109/ISSCC.2018.8310204.

[17] M. Hossain, Aurangozeb, and N. Nguyen, "DDJ-adaptive SAR TDC-based timing recovery for multilevel signaling," *IEEE J. Solid-State Circuits*, vol. 54, no. 10, pp. 2833–2844, Oct. 2019, doi: 10.1109/JSSC.2019.2923586.

[18] Y. Chang, A. Manian, L. Kong, and B. Razavi, "An 80-Gb/s 40-mW wireline PAM4 transmitter," *IEEE J. Solid-State Circuits*, vol. 53, no. 8, pp. 2214–2226, Aug. 2018, doi: 10.1109/JSSC.2018.2831226.

[19] B.-J. Yoo *et al.*, "A 56Gb/s 7.7mW/Gb/s PAM-4 wireline transceiver in 10nm FinFET using MM-CDR-based ADC timing skew control and low-power DSP with approximate multiplier," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2020, pp. 122–124, doi: 10.1109/ISSCC19947.2020.9062964.

[20] Groen *et al.*, "6.3 A 10-to-112Gb/s DSP-DAC-Based Transmitter with 1.2Vpp Output Swing in 7nm FinFET," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2020, pp. 120–122, doi: 10.1109/ISSCC19947.2020.9063130.

[21] J. Kim *et al.*, "8.1 A 224Gb/s DAC-based PAM-4 transmitter with 8-Tap FFE in 10nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2021, pp. 126–128, doi: 10.1109/ISSCC42613.2021.9365840.

[22] M. Choi *et al.*, "4.8 An output-bandwidth-optimized 200Gb/s PAM-4 100Gb/s NRZ transmitter with 5-Tap FFE in 28nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2021, pp. 128–130, doi: 10.1109/ISSCC42613.2021.9366012.

[23] M. Kossel *et al.*, "8.3 An 8b DAC-based SST TX using metal gate resistors with 1.4pJ/b efficiency at 112Gb/s PAM-4 and 8-Tap FFE in 7nm CMOS," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2021, pp. 130–132, doi: 10.1109/ISSCC42613.2021.9365784.

[24] P. Mishra *et al.*, "8.7 A 112Gb/s ADC-DSP-based PAM-4 transceiver for long-reach applications with >40dB channel loss in 7nm FinFET," in *Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2021, pp. 138–140, doi: 10.1109/ISSCC42613.2021.9365929.

[25] Y. Suzuki, K. Odagawa, and T. Abe, "Clocking CMOS calculator circuitry," *IEEE J. Solid-State Circuits*, vol. 8, no. 6, pp. 462–469, Dec. 1973, doi: 10.1109/JSSC.1973.1050440.

[26] J. Yuan and C. Svensson, "High-speed CMOS circuit technique," *IEEE J. Solid-State Circuits*, vol. 24, no. 1, pp. 62–70, Feb. 1989, doi: 10.1109/4.16303.

[27] R. Rogenmoser, N. Felber, Q. Huang, and W. Fichtner, "1.16 GHz dual-modulus 1.2 μ m CMOS prescaler," in *Proc. IEEE Custom Integr. Circuits Conf.*, May 1993, pp. 27.6.1–27.6.4, doi: 10.1109/CICC.1993.590807.

[28] B. Razavi, "The role of PLLs in future wireline transmitters," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 8, pp. 1786–1793, Aug. 2009, doi: 10.1109/TCSI.2009.2027507.

[29] B. Razavi *et al.*, "A low-power 60-GHz CMOS transceiver for WiGig applications," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2013, pp. 300–301, pp. C300–C301.

[30] Y. Zhao and B. Razavi, "A 56-GHz fractional-N 22-mW PLL," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2022, pp. 288–289.

[31] A. Garg, A. C. Carusone, and S. P. Voinigescu, "A 1-tap 40-Gb/s look-ahead decision-feedback equalizer in 0.18 μ m SiGe BiCMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 10, pp. 2224–2232, Oct. 2006, doi: 10.1109/JSSC.2006.878109.

[32] Y.-S. Sohn, S.-J. Bae, H.-J. Park, and S.-I. Cho, "A 1.2-Gbps CMOS DFE receiver with extended sampling time window for application to SSTL channel," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2002, pp. 92–93, doi: 10.1109/VLSIC.2002.1015055.

N