

A 32-mW 40-Gb/s CMOS NRZ Transmitter

Yikun Chang, Abishek Manian, Long Kong, and Behzad Razavi
Electrical Engineering Department, University of California, Los Angeles, CA

Abstract— A wireline transmitter with 2-tap feedforward equalization achieves more than a two-fold improvement in the power efficiency through the use of a new integrating multiplexer, as well as quadrature clock phases with 25% duty cycle. Serializing data by a factor of 128 and including an on-chip phase-locked loop, the transmitter is realized in 45-nm CMOS technology and delivers a differential output swing of 460 mV_{pp}.

Keywords— Serialization, output driver/multiplexer, integrating multiplexer, multi-phase multiplexing, phase mismatch.

I. INTRODUCTION

The demand for greater data rates in serial links continues unabated, making the problem of power consumption increasingly more difficult. Transmitters operating in the range of tens of gigabits per second have been reported [1] - [4], but the lowest power is in the range of 80 to 88 mW, yielding a power efficiency of 1.82 to 2 mW/Gb/s [2], [4].

This paper describes a 40-Gb/s transmitter (TX) that improves the efficiency by a factor of 2.28 with a clock jitter of 332 fs_{rms}. This is accomplished through the use of new architecture and circuit techniques.

II. DESIGN CONSIDERATIONS

In transmitter design, the output driver tends to consume a high power as it must deliver large currents with relatively large voltage swings. This issue is exacerbated at tens of gigabits per second for two reasons: (1) the need for on-chip back-termination resistors doubles the power, and (2) the use of feedforward equalization (FFE) requires additional strength in the driver.

Two other difficulties arise in conventional transmitters due to the use of a full-rate retimer and hence the need for a full-rate frequency divider [Fig. 1(a)]. First, the divider must operate at 40 GHz while driving at least two multiplexers, potentially drawing a high power. Second, the divider delay is subtracted from the timing margin available to the retimer, severely limiting the speed [2], [5].

Another challenge in the TX front end of Fig. 1(a) relates to the driver's large input capacitance, C_{dr} . To achieve sufficient bandwidth ($\approx 0.7 \times 40 \text{ Gb/s} = 28 \text{ GHz}$) at this interface, we can either introduce a predriver or design the retimer with high currents and low impedances, both power-hungry solutions. For example, the two-stage driver in [2] draws 26.4 mW. It is possible to remove the full-rate retimer and divider [Fig. 1(b)] so as to avoid the speed and delay constraints imposed by the latter. In this case, however, the clock duty cycle error and the data path mismatches within each MUX directly translate to jitter at the output.

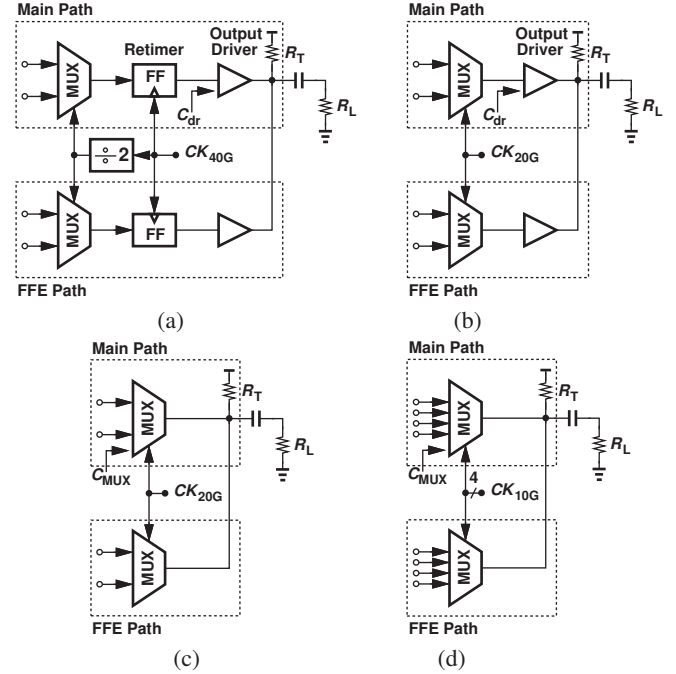


Fig. 1. (a) Full-rate front end with retimer and divider, (b) half-rate front end, (c) half-rate front end with combined 2-to-1 MUX and driver, and (d) quarter-rate front end with combined 4-to-1 MUX and driver.

The TX front end can be further simplified if the MUX and driver stages are merged [Fig. 1(c)] [4]. Here, the power consumed by the MUXes is not “wasted.” The large input capacitance of the MUX, C_{MUX} , is now driven at 20 Gb/s, and the latches within the multiplexers operate at 20 GHz, both still challenging problems. In the next step, we contemplate the use of multi-phase multiplexing [6] and combine the idea with Fig. 1(c), arriving at the front end shown in Fig. 1(d). In this case, C_{MUX} is as large as C_{dr} in Fig. 1(a), and is driven at 10 Gb/s, but the main path contains four of these input capacitances. Moreover, multi-phase clocking still requires four latches for each 4-to-1 MUX [2]. Our proposed TX architecture addresses both of these issues.

III. TRANSMITTER ARCHITECTURE

Fig. 2 shows the TX architecture. It consists of a main serializer path, an FFE path with a programmable strength, and a phase-locked loop (PLL) for clock generation. The main path comprises a 128-to-8 CMOS MUX, which produces data at a rate of 5 Gb/s, an 8-to-4 “current-integrating” MUX (IMUX), and a 4-to-1 current-mode-logic (CML) MUX. The FFE employs four programmable 4-to-1 CML MUX slices. The PLL receives a reference frequency of 312.5 MHz and

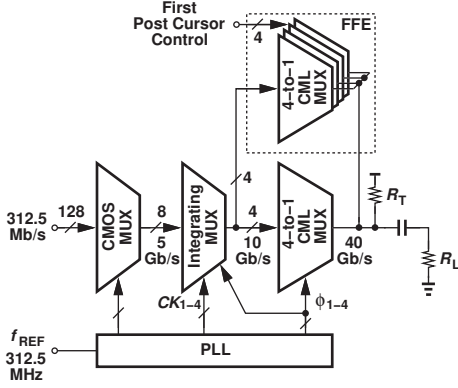


Fig. 2. Proposed NRZ transmitter architecture.

delivers 25%-duty-cycle clock phases, ϕ_1 - ϕ_4 , at 10 GHz, 50%-duty-cycle phases, CK_1 - CK_4 , at 5 GHz, etc.

The proposed transmitter achieves more than a two-fold improvement in the power efficiency as a result of three new concepts: (1) the integrating MUX drives the large input capacitance of the CML MUX with very low power consumption ($410 \mu\text{W} \times 4$ for the 8-to-4 selector), (2) the use of quadrature clock phases with 25% and 50% duty cycles completely eliminates high-speed latches in the data path, and (3) the integrating selector incorporates a timing scheme that readily accommodates the first FFE post cursor.

IV. INTEGRATING MUX

The direct 4-to-1 MUX/driver in Fig. 2 presents two issues, namely, a large input capacitance, $C_{MUX} \approx 96 \text{ fF}$, and proper timing in the preceding stage to guarantee that each input is available when one of ϕ_1 - ϕ_4 is asserted. The integrating MUX efficiently deals with both issues.

In order to drive C_{MUX} at a bit rate of $r_b = 10 \text{ Gb/s}$, we can opt for a CML stage [Fig. 3(a)]. Here, we must choose $1/(2\pi R_L C_{MUX}) \approx 0.7r_b$ for minimal ISI, and $R_L = V_0/I_{SS}$ to obtain a single-ended peak-to-peak swing of V_0 . That is, the CML stage consumes $1.4\pi r_b C_{MUX} V_0 V_{DD}$. For example, if $C_{MUX} \approx 100 \text{ fF}$, $V_0 \approx 400 \text{ mV}$, and $V_{DD} = 1 \text{ V}$, the four CML stages driving the 4-to-1 MUX consume a total of 7 mW.

Alternatively, the MUX can be driven by an integrating stage [Fig. 3(b)], where first the output is reset to V_{DD} and then the tail current turns on to impress the data level on C_{MUX} . In this case, the power consumption is given by $r_b C_{MUX} V_0 V_{DD}$, a factor of 4.4 lower than that of the CML topology. Additionally, the differential pair transistors in the integrating stage present less input capacitance.

It is desirable to incorporate multiplexing within the integrating stage of Fig. 3(b). As shown in Fig. 4, two differential pairs receive D_{in1} and D_{in2} , but only one is enabled according to the select command, CK_1 . Thus, when ϕ_2 goes high, D_{in1} or D_{in2} travels to the output. As in a standard selector, CK_1




Fig. 3. Driving the 4-to-1 MUX by (a) a CML stage, or (b) an integrating stage.

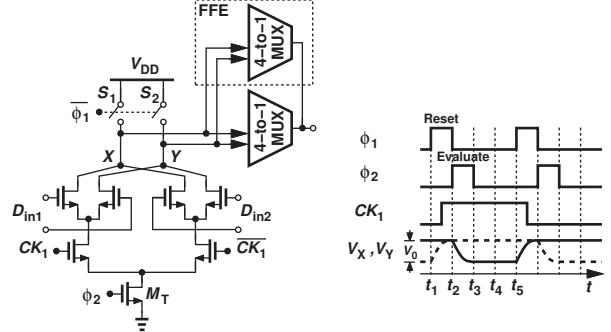


Fig. 4. Proposed 2-to-1 integrating MUX.

has a 50% duty cycle and the same rate as the inputs (5 GHz), but ϕ_1 and ϕ_2 have a 25% duty cycle and run at 10 GHz, creating much more flexibility in the overall architecture (explained below).

The integrating MUX operates as follows. First, X and Y are reset to V_{DD} while the tail current source, M_T , is off. Next, CK_1 arrives to select D_{in1} or D_{in2} , and then ϕ_2 rises to perform evaluation. In this mode, V_X or V_Y falls for about 25 ps, providing the desired swing, V_0 . When ϕ_2 goes low, the tail current ceases and the output is held for approximately 50 ps.

The use of both 25% and 50% duty cycles enhances two aspects of the design. First, the main 4-to-1 MUX senses V_X and V_Y from t_3 to t_4 whereas the FFE branch receives these values from t_4 to t_5 . Since this time offset is equal to 1 unit interval (UI) at 40 Gb/s, feedforward equalization is implemented with no latches, thus saving power. Second, the topology in Fig. 4 provides a hold period, during which V_X and V_Y are constant, so that the subsequent stages can sense the signals reliably. Without the 25%-duty-cycle phases, on the other hand, V_X or V_Y would continue to fall after t_3 , creating unequal swings for the main and FFE paths and hence substantial ISI.

The integrating MUX of Fig. 4 merits two more remarks. First, the stacking of transistors still lends itself to a 1-V supply because all of the inputs have rail-to-rail swings. Second, since the value of V_0 is PVT-dependent, the circuit is designed so as to produce a sufficient swing for the 4-to-1 MUXes if M_T is weak and also reset X and Y to V_{DD} in 25 ps if M_T is strong and S_1 and S_2 are weak.

The proposed transmitter employs four integrating multiplexers to serialize data from $8 \times 5 \text{ Gb/s}$ to $4 \times 10 \text{ Gb/s}$. These outputs directly drive the direct 4-to-1 multiplexers in the main and FFE paths.

V. MAIN AND FFE MULTIPLEXERS/DRIVERS

The waveforms in Fig. 4 indicate that the integrating MUX provides a stable output for two time slots, each 25 ps long. We allocate the first slot to the 4-to-1 MUX in the main path and the second to that in the FFE path.

Fig. 5 shows the main and FFE 4-to-1 MUX/driver circuits in simplified form. Four differential pairs controlled by ϕ_1 - ϕ_4 select one of the inputs for 25 ps, delivering the 40-Gb/s data to the 50- Ω on-chip back-termination resistors and the 50- Ω loads. The differential output voltage swing (without FFE action) is at least 440 mV across PVT corners.

The FFE path consists of four programmable slices that provide a relative tap coefficient ranging from 0 to 0.4. Each slice contains four differential pairs controlled by ϕ_1 - ϕ_4 and scaled down by a factor of 10 with respect to those in the main path.

The interface between the IMUX and the final drivers is illustrated in Fig. 6. Here, three operations occur in succession: first, ϕ_2 is high from t_2 to t_3 for the IMUX to generate proper levels at X and Y ; next, ϕ_3 is high, allowing the 4-to-1 MUX in the main path to sense V_X and V_Y ; last, ϕ_4 is high, enabling the FFE MUX.

VI. EFFECT OF MISMATCHES

Multiplexers generally produce some jitter at their output due to the mismatches in the data paths and in the clock paths.




Fig. 5. Main and FFE data paths.

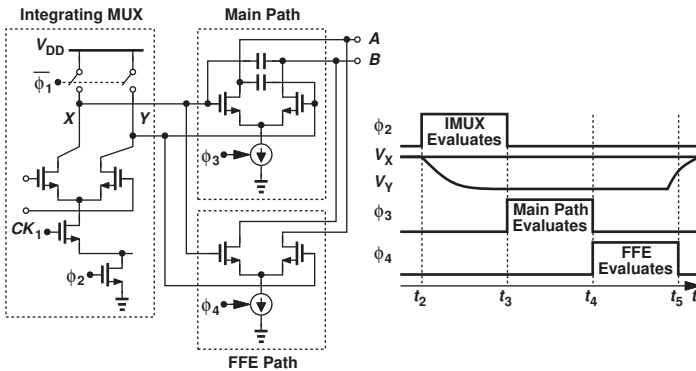


Fig. 6. Interface between the integrating MUX and the main and FFE drivers/MUXes.

Even a 2-to-1 MUX suffers from output jitter if its selection clock duty cycle deviates from 50%, a difficulty avoided by placing a retiming flip-flop after the MUX. Without such a retimer, however, the mismatches must be sufficiently small. Prior work evidently has not addressed this issue.

The MUX output jitter arises from both duty cycle mismatches and delay mismatches among the clock phases. But only some of the transition errors translate to jitter. We consider two cases. As shown in Fig. 7(a), the high times of ϕ_1 - ϕ_4 incur errors equal to ΔT_{H1} - ΔT_{H4} , respectively, where $\Delta T_{H1} + \Delta T_{H2} + \Delta T_{H3} + \Delta T_{H4} = 0$. 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

$$J_{pp} = \max(\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4) - \min(\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4), \quad (1)$$

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

The effect of delay mismatches is illustrated in Fig. 7(b), where we assume the falling edge of ϕ_1 incurs an error of ΔT_{sk1} , and the rising edge of ϕ_2 , an error of ΔT_{sk2} . In this case, the differential output of the MUX suffers from a zero-crossing displacement equal to $(\Delta T_{sk1} + \Delta T_{sk2})/2$. Extending this result to all four phases, we have

$$J_{pp} = \max(\delta_1, \delta_2, \delta_3, \delta_4) - \min(\delta_1, \delta_2, \delta_3, \delta_4), \quad (2)$$

where $\delta_1 = (\Delta T_{sk1} + \Delta T_{sk2})/2$, $\delta_2 = (\Delta T_{sk2} + \Delta T_{sk3})/2$, etc. These results are for differential outputs; the single-ended output jitter can be shown to be *larger*.

To quantify the jitter, Monte Carlo simulations have been performed on the chain consisting of the frequency dividers, buffers, the 8-to-4 IMUX and the 4-to-1 MUX, all extracted from the layout. The simulations indicate a total rms jitter of 284 fs in the single-ended output and 205 fs in the differential output. We verify this result experimentally in Section VII. The key point here is that the inevitably large transistors in the data and clock paths provide sufficiently small mismatches.

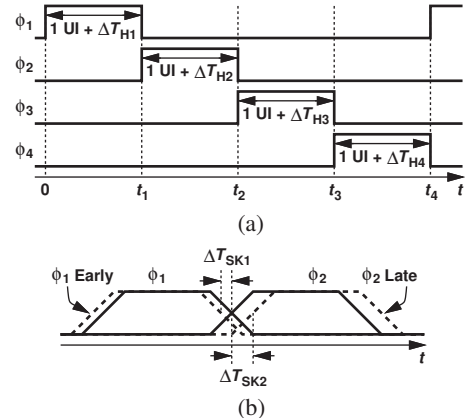


Fig. 7. (a) Duty cycle mismatch, and (b) delay mismatch.

VII. EXPERIMENTAL RESULTS

The 40-Gb/s NRZ transmitter has been fabricated in TSMC's 40-nm CMOS technology and tested with a 1-V supply. Fig. 8 shows a photograph of the die, whose active area measures $330 \mu\text{m} \times 175 \mu\text{m}$.

Fig. 9(a) plots the measured output spectrum of the PLL at 20 GHz and Fig. 9(b) shows the measured phase noise after this clock is divided by 2. The phase noise is -110 dBc/Hz at 10 GHz. Integrated from 10 kHz to 100 MHz, the jitter is equal to $332 \text{ fs}_{\text{rms}}$. The reference spurs are at -45 dBc .

Fig. 10(a) shows the TX output eye diagram with no FFE action. The differential voltage swing is $460 \text{ mV}_{\text{pp}}$. Fig. 10(b) shows the output with the FFE tap strength of 0.4, yielding a 7.4-dB boost. The output bit stream has also been captured and checked to ensure correct serialization of the 128 312.5-Mb/s inputs to the 40-Gb/s output.

In order to examine the effect of mismatches (Section VI), we apply the input data so as to create a 20-GHz periodic 0101 sequence at the TX output. The duty cycle and delay mismatches produce spurs at 10-GHz offset. Fig. 11 shows the single-ended measured spectrum, indicating a spur level of -34 dBc . Translating this value to rms jitter in the single-ended output, we arrive at 225 fs, in reasonable agreement with Monte Carlo simulations.

The TX consumes 32 mW from a 1-V supply: 9.0 mW in the main and FFE 4-to-1 MUX/drivers, 1.6 mW in the four integrating MUXes, 3.4 mW in the VCO and 18.0 mW in the 128-to-8 serialization, PLL core, divider chain and clock distribution. Table I summarizes the measured performance of our TX and compares it to the prior art. We have achieved a factor of 2.28 improvement in the power efficiency.




Fig. 8. TX die photograph.




Fig. 9. (a) Measured spectrum of 20-GHz clock, and (b) phase noise profile of 10-GHz clock.

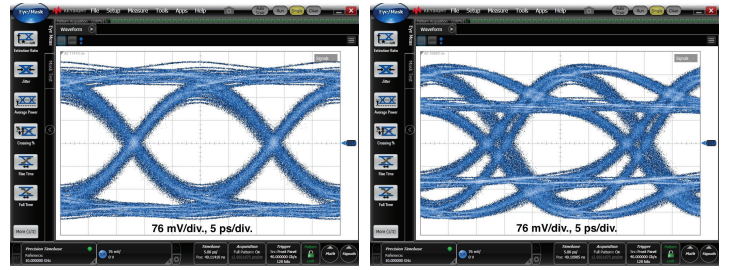


Fig. 10. Measured eye diagrams with (a) FFE off, and (b) four FFE slices on.




Fig. 11. Measured spectrum of single-ended output delivering 0101 sequence.

TABLE I. PERFORMANCE SUMMARY.

Reference	[1]	[2]	[3]	[4]	This Work
Technology (nm)	14	65	28	65	45
Data Rate (Gb/s)	16 - 40	31.68 - 48.4	40	40	40
FFE	4-tap	no	2-tap	2-tap	2-tap
PN (dBc/Hz) f_{offset} (MHz)	-	-127.5 10	-128 100	-	-104 10
RMS Jitter (fs) Integ. Range (MHz)	-	251 0.0001 - 10	162 10 - 10000	-	332 0.01 - 100
Power (mW)	Data Path	-	41.4	-	11
	Whole TX*	518**	88	-	80
Power Eff. (pJ/bit)	Data Path	-	0.86	-	0.28
	Whole TX*	12.95**	1.82	-	2

* Data path and clock path (PLL, phase generation and clock distribution).
** Excluding power of PLL.

ACKNOWLEDGMENTS

The authors thank the TSMC University Shuttle Program for chip fabrication. This research was supported by Oracle, Realtek Semiconductor, and Texas Instruments.

REFERENCES

- [1] J. Kim *et al.*, "A 16-to-40Gb/s Quarter-Rate NRZ/PAM4 Dual-Mode Transmitter in 14nm CMOS," *ISSCC Dig. Tech. papers*, pp. 60-61, Feb. 2015.
- [2] A. A. Hafez, M.-S. Chen and Ch.-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, pp. 763-775, Mar. 2015.
- [3] R. Navid *et al.*, "A 40 Gb/s Serial Link Transceiver in 28 nm CMOS Technology," *IEEE J. Solid-State Circuits*, vol. 50, pp. 814-827, Apr. 2015.
- [4] K. Huang *et al.*, "A 190mW 40Gbps SerDes Transmitter and Receiver Chipset in 65nm CMOS Technology," *Proc. CICC*, Sep. 2015.
- [5] B. Razavi, *Design of Integrated Circuits for Optical Communications*, McGraw-Hill, 2003.
- [6] Ch.-K. K. Yang, R. Farjad-Rad and M. A. Horowitz, "A 0.5- μm CMOS 4.0-Gbit/s Serial Link Transceiver with Data Recovery Using Oversampling," *IEEE J. Solid-State Circuits*, vol. 33, pp. 713-722, May 1998.