A 112-Gb/s 58-mW PAM4 Transmitter in 28-nm CMOS Technology

Mahdi Forghani, Yu Zhao, Pawan K. Khanna, and Behzad Razavi
Electrical and Computer Engineering Department, University of California, Los Angeles, CA 90095, USA
forghani@ucla.edu

Abstract

A voltage-mode transmitter employs a resistorless output DAC, a 3-tap latchless FFE, a passive output skew compensation network, and a 56-GHz integer-N PLL. The prototype delivers an output swing of 0.8 $V_{pp,d}$ with an rms clock jitter of 160 fs and RLM = 96%.

Keywords: Wireline, PAM4, serializer.

PAM4 transmission at 112 Gb/s has proved popular in recent years [1]-[7], requiring either high power consumption [4, 6] and/or sub-10-nm processes [3, 5]. The motivation for targeting this speed in 28-nm technology is twofold. First, the cost can be lowered considerably. Second, the techniques affording this performance can be applied to more advanced nodes so as to achieve higher speeds and lower power consumption. This paper describes such a transmitter (TX).

Fig. 1 shows the proposed TX architecture, depicted in single-ended form for simplicity. The data path serializes the MSB and LSB NRZ data from 875 Mb/s to 56 Gb/s and applies the results to the output driver/DAC. The latter delivers PAM4 data through series-peaking inductors and a T-coil to the channel. The clock path consists of an on-chip 56-GHz PLL followed by a chain of ÷2 stages and buffers. The novelties in this work include the DAC, the FFEs, the skew compensation network, the MUX chain, and the ÷2 stage sensing the 56-GHz clock. The reasons for the choice of 56 GHz (rather than 28 GHz) for the PLL are twofold: a more compact inductor footprint is afforded, and the leakage of this clock to the output data is greatly attenuated by the channel.

Aiming for a low-power solution, we begin with the assumption that the output driver/DAC saves substantial power by incorporating voltage-mode (SST) operation rather than a current-mode (CML) approach. Thus, all data and clock waveforms preceding the TX output nodes must have rail-to-rail swings. The theoretical 4x power advantage of SST, however, does not readily materialize because both the predriver power and the crowbar (V_{DD}-to-GND) current of SST stages manifest themselves at these speeds. Moreover, SST output drivers (DACs) typically incorporate extensive strength programmability so as to calibrate both the output swing and the TX S_{22} [Fig. 2(a)], and series resistors, R_S, so as to partially rely on their value for swing and S_{22} definition. As a result, the overall input and output capacitances of the DAC are about twice or three times those of a single inverter delivering the same voltage swing to R_L. The predriver thus becomes nearly as power-hungry as the driver itself.

We propose an output drive adjustment scheme that avoids the "overhead" input capacitance described above. Depicted in Fig. 2(b), the idea is to design the SST DAC for the nominal output swing under nominal conditions and delegate the programmability to an LDO regulator that provides V_{DD}. Since typical systems do employ LDOs for different TX sections, this method requires negligible power or area penalty. According to simulations, V_{DD} need change by no more than 100 mV between SS and FF corners so as to yield a relatively constant output swing. The absence of linear resistors in series with the output degrades the PAM4 eye linearity to some extent because the on-resistance of MOSFETs depends on their drain-source voltage. We then present the DAC shown in Fig. 2(c), where the unit used in the LSB sections is scaled up by 14%. It can be proved that this value is obtained as follows. If the LSB and MSB resistances are respectively equal to $2R_u$ and R_u for an output level of +2, and the former rises to $2R_u(1+\alpha)$ for a level of +1, then the LSB path must be scaled up by a theoretical value of 16%. The DAC is designed to deliver a 0.8-V differential swing to a 100-Ω channel. The dramatic reduction of the DAC input capacitance propagates "backwards", allowing smaller transistors in the predriver, the MUX stages, and in the clock paths.

Another premise held in this TX development is that the last multiplexing stage – which produces NRZ data at 56 Gb/s – should be implemented as a 2-to-1 selector [Fig. 3(a)] rather than as a 4-to-1 tree [Fig. 3(b)] because the latter suffers from a higher output capacitance and demands I/Q calibration [5]. This choice leads to a host of other issues. Specifically, the stages driving the 2-to-1 MUX in Fig. 3(a) must deliver data at a rate of 28 Gb/s, a challenging task in 28-nm technology but feasible here because MUX_1 sees a smaller C_{pre} and, therefore, incorporates minimum-size transistors. Also, this MUX requires a 28-GHz clock. More importantly, the maximum tolerable divider delay in Fig. 3(a) is one-half of that in Fig. 3(b). This issue is resolved by inserting a delay replica in the clock path.

TX FFEs are another power-hungry function as they employ a large number of latches. We propose a "latchless FFE" topology that avoids this difficulty. Recognizing that the FFE should primarily provide a high-frequency boost, we note that its taps need not be precisely defined with respect to the symbol period. We thus present the arrangement shown in Fig. 4(a), where the delay, T_1, is also realized by inverters. Along with the DAC, the circuit synthesizes $1 - 0.17\,z^{-0.6} - 0.12\,z^{-1.2}$, offering 4.3 dB of boost at Nyquist even with PVT variations. The MSB and LSB FFEs draw a total of only 4 mW.

A serious concern in serial links relates to the skew experienced by differential signals as they incur random length mismatches on their way to the receiver. Such a mismatch translates to large common-mode noise at the RX input. We introduce a skew compensation method: as shown in Fig. 4(b), the 56-Gb/s differential NRZ data produced by the last MUX travels through the predriver inverters and feedforward resistors R_1 and R_2. Realized by complementary MOS devices, R_1 and R_2 offer ±3.2 ps of skew control.

Fig. 4(c) shows the TX die photo (for flexibility, the LDO is not integrated). The cable and probe losses are compensated in the oscilloscope, and the TX FFE is enabled. Fig. 5(a) depicts the measured NRZ output eye at 56-Gb/s. The jitter is about 160 fs. Fig. 5(b) shows the 112-Gb/s PAM4 output with

RLM = 96%, with a total power of 58 mW. Raising V_{DD} by 150 mV yields the eye in Fig. 5(c) with RLM = 99% while requiring another 14 mW. Disabling the TX FFE produces the eye shown in Fig. 5(d), proving the efficacy of the FFE. Fig. 6 depicts the ±3.2 ps skew control. The measured S_{22} is around -10 dB up to 67 GHz for all output levels. Table 1 compares the performance to prior work. We remark that [3] only reports the on-chip PAM4 eye (by statistical measurements).

Acknowledgements Research supported by Realtek Semiconductor. The authors thank the TSMC University Shuttle Program for chip fabrication.

References

[1] Z. Toprak-Deniz et al., *IEEE JSSC*, pp. 19-26, Jan. 2020.
[2] J. Kim et al., *IEEE JSSC*, pp. 29-42, Jan. 2019.
[3] C. F. Poon et al., *IEEE JSSC*, pp. 1199-1210, Apr. 2022.
[4] P. -J. Peng et al., *IEEE JSSC*, pp. 2123-2131, July 2021.
[5] E. Groen et al., *IEEE JSSC*, pp. 30-42, Jan. 2021.
[6] X. Zheng et al., *IEEE JSSC*, pp. 1864-1876, July 2020.
[7] B. Ye et al., *IEEE JSSC*, pp. 19-29, Jan. 2023.

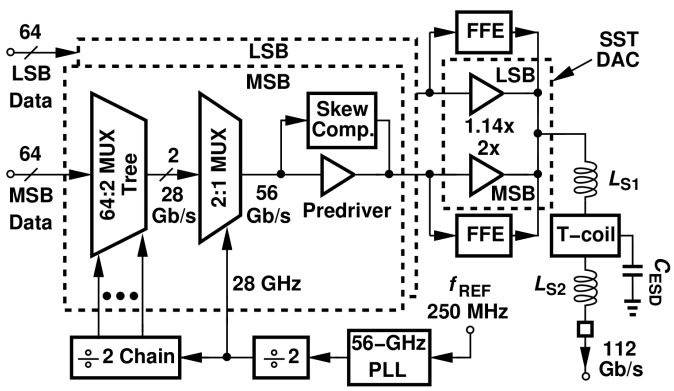

Fig. 1. Proposed TX architecture.

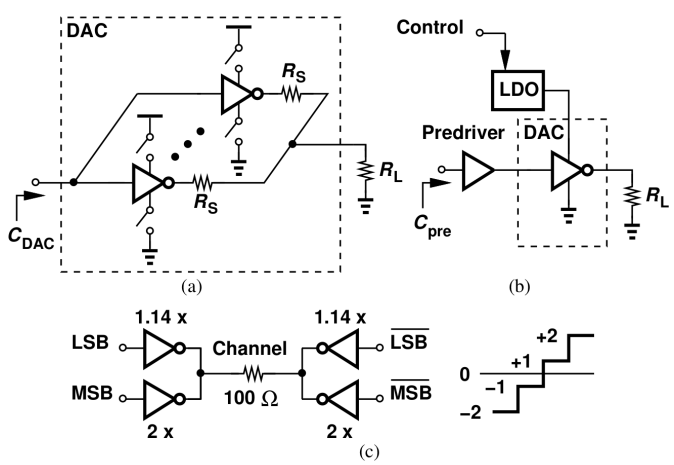

Fig. 2. DAC output swing and S_{22} calibration: (a) conventional method, (b) proposed method, and (c) proposed linearization technique.

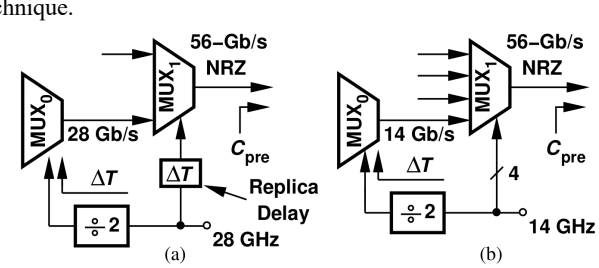

Fig. 3. (a) Half rate TX and (b) quarter rate TX.

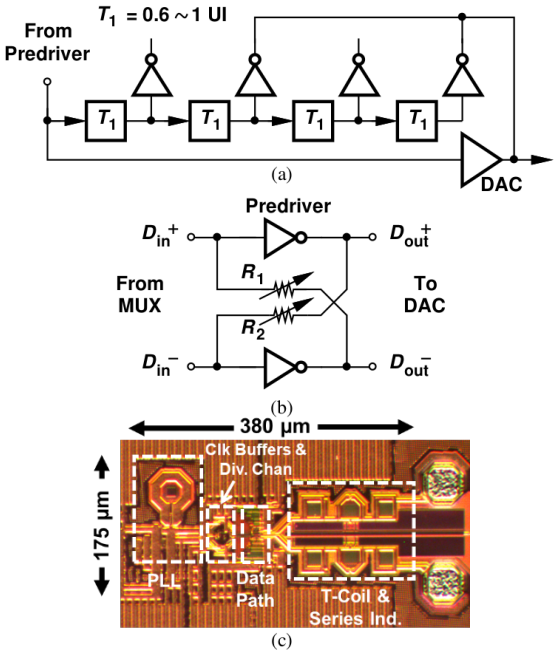

Fig. 4. (a) 3-tap latchless FFE, (b) skew compensation network, and (c) die photograph.

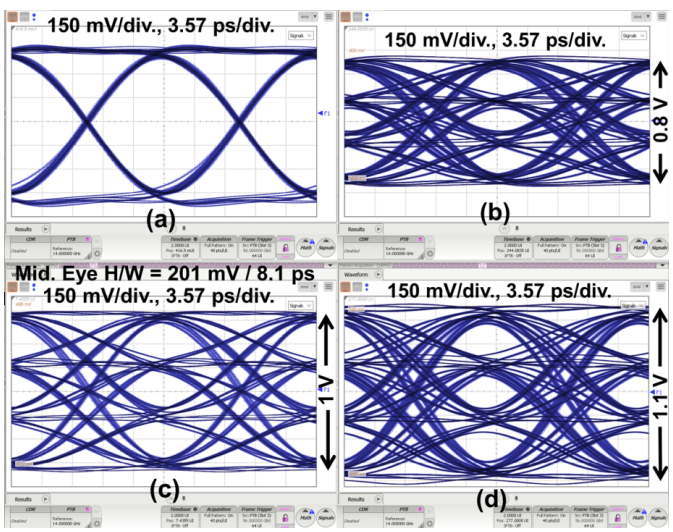

Fig. 5. Measured eye diagrams for (a) 56-Gb/s NRZ, (b) 112-Gb/s PAM4, (c) high-swing mode, and (d) FFE disabled.

Fig. 6. Single-ended eye diagrams with skew adjustment.

TABLE I. Performance Summary.

		[1]	[2]	[3]	[4]	[5]	[6]	[7]	This Work
Technology (nm)		14	10	7	40	7	65	28	28
Data Rate (Gb/s)		112/128	112	112	112	112	112	112	112
$V_{pp,d}$ (V)	Max.	0.6[a]/1	0.75	0.7	1	1.2	1.2	0.8	1.1
	112Gb/s w/ FFE	0.4/0.7	0.5	0.4	0.64	0.36	0.75	0.65	0.8
RLM (%)		99/98.6	98.5	–	97.6	94	99.7	–	96
RMS RJ (fs) Integ. Range (Hz)		–	154*	114** 10k–14G	210*** 100–40M	171**** 4M–28G	–	128*	160 1k–3.5G
Power (mW)	Exc. PLL	112[a]/170	193	54.1 (30[b])	–	130 (83[b])	243 (122[b])	59	49
	Total	N/A	232	62	436	175	–	109	58
Active Area (mm²)		0.048	0.03	0.055	0.56	0.193	0.694	0.05	0.066

*Measured by oscilloscope ** 28–GHz NRZ Clk pattern *** 14–GHz PLL **** 28–GHz PLL
[a] Low–power mode [b] Excluding all clocking