2.3 60Gb/s NRZ and PAM4 Transmitters for 400GbE in 65nm CMOS

Ping-Chuan Chiang[1,2], Hao-Wei Hung[1], Hsiang-Yun Chu[1], Guan-Sing Chen[1], Jri Lee[1,2]

[1]National Taiwan University, Taipei, Taiwan,
[2]Atilia Technology, Taipei, Taiwan

Recent research indicates that data-link transceivers running at or below 40Gb/s are practical to implement in CMOS technology [1]. However, next-generation datacom and telecom systems require transceivers to operate at even higher data rates. For example, a 400Gb/s Ethernet system may need 8×50Gb/s PAM2 (NRZ) or PAM4 channels [2]. This paper introduces fully integrated solutions for NRZ and PAM4 transmitters. The 60Gb/s operating speed demonstrates sufficient bandwidth even for standards with coding overhead.

Figure 2.3.1 illustrates the NRZ transmitter architecture. It consists of a 4:1 multiplexer in a tree structure, a 60GHz PLL with adaptive phase aligner to optimize the clock phase in the last stage, and a built-in quarter-rate PRBS generator to facilitate testing. There are two independent modes to select the incoming data and clock by means of switches Sel_1 and Sel_2. In normal operation, the 4 input data ports are fed by 4 external independent data sequences (15Gb/s each). In self-testing mode, on the other hand, the inputs come from the built-in quarter-rate 2^7-1 PRBS generator. Similarly, the synchronizing clock can be selected from either from the internal PLL or from the external clock source. At 60Gb/s, the phase relationship between clock and data is critical to a functioning system. In the first multiplexing stage, delays ΔT_1 and ΔT_2 are inserted to balance the sample timing. These delays are designed to match the internal skews over a wide temperature range. At 60Gb/s, the phase-alignment issue becomes so severe that a static delay does not work. For instance, the acceptable sampling window in the last stage (60Gb/s output) is about 8 to 10ps, but the phase drifting caused by PVT variations can be as large as 15 to 20ps. To accommodate the random phase relationship, we put a phase aligner in front of the second multiplexing stage to dynamically track the optimal clock and data phases. The phase tracking operates as follows. First, the synchronization clock (wherever it comes from) is divided by 2 to generate quadrature clocks at 30GHz. The data transition is examined by using a roughly 16.5ps delay ΔT_3 with a mixer (M1) to detect the arrival of the internal 30Gb/s data. With the help of the 30GHz phase interpolator (PI) and the second mixer (M2), we arrive at a feedback loop that forces the PI to produce the clock phase that aligns with the data transition. As a result, the 60Gb/s multiplexer can properly sample and serialize the data, even under extreme PVT variations.

The final 2:1 selector stage needs to provide wide bandwidth and reasonable gain. As shown in Fig. 2.3.2, it is made of a 3-stage distributed amplifier with CML switching pair in the bottom. Two data inputs and one switching clock are applied into transmission lines, travelling along the 3 stages until the end terminations. The output ports are also connected through transmission lines. Here, one end is terminated while the other is open as an output port. The final D_{out60} can be ac coupled to external loading with 50Ω termination. Simulation shows that the output matching S_{22} is kept below –10dB from dc up to 70GHz. The data (L_G) and clock (L_C) paths are designed to have characteristic impedance of 50Ω as well, and the two paths have identical group velocity. Gate capacitances are absorbed into the transmission lines. It is essential to properly choose the number of stages (N) to achieve good performance. Taking into consideration the transmission-line loss and active device g_m, we determine that a 3-stage structure provides the best performance. As illustrated in Fig. 2.3.3(a), the total gain starts to roll off as N becomes larger than 3, whereas the overall power dissipation continues to increase. To avoid long routing, gate transmission lines are realized as lumped inductors in congested areas.

The PAM4 transmitter design is depicted in Fig. 2.3.4. It includes a built-in PLL for clock generation, and a two-path half-rate FFE with 3 taps and ×2 weighting factor. The original data is split into two sequences (D_{inA} and D_{inB}) of 28Gb/s, which are pre-emphasized (with the same coefficients) before combination. To ensure signal integrity, all high-speed paths are realized as transmission lines or equivalent peaking circuits. A key component that significantly affects the performance is the combiner (i.e., output driver). At tens of GHz, large-area elements such as inductors can no longer be considered lumped components, but rather distributed devices. In that sense, the peaking and signal-travelling circuits must be combined as a distributed network so as to minimize skews, reflection, and other non-idealities. Figure 2.3.3(b) reveals the combiner design. Here, peaking inductors L_D and L_G are inserted between taps to (1) absorb the gate and drain capacitance, (2) balance the travelling time. Back to the transmitter architecture in Fig. 2.3.4, we have the master clock designed in a way that it can be either provided externally or generated by the PLL (by switch Sel). The 28GHz nominal frequency is for some standard requirements defined in [3]. Again with the help of a SSB mixer-based PFD [4] and a sub-harmonic injection locking technique [5], the PLL provides a pure clock from 26.9 to 28.5GHz with jitter as low as 508fs_{rms}. Based on the design of the inside matching network, the transmitter supports a wider operating range via the external clock. Simulation indicates the internal peaking and transmission lines behave well from dc to 105GHz, and the transmitter is verified by measurement to provide a flat data response from 1Gb/s to 60Gb/s. The tail currents in combiner II are twice as much as those in combiner I to realize PAM4 waveforms.

Both TX circuits are designed and fabricated in 65nm CMOS technology. The NRZ transmitter consumes 450mW of power and the PAM4 TX 290mW, both from a 1.2V supply. Figure 2.3.5 shows the measurement results for the NRZ transmitter. The 30Gb/s and 60Gb/s outputs are shown in Fig. 2.3.5(a) and (b), respectively, presenting output magnitude of 100mV with open eyes. The rms data jitter of 30Gb/s output measured from oscilloscope is equal to 1.08ps, and its peak-to-peak data jitter is measured as 5.33ps. The 30GHz clock output from the built-in PLL is also recorded as shown in Fig. 2.3.5(c). It presents rms jitter of 461fs (integrated from 1kHz to 20MHz offset), and phase noise of –100dBc/Hz at 1MHz offset. The PAM4 transmitter is also tested thoroughly. Figure 2.3.6(a) depicts the output waveform and phase noise plot of the 28GHz built-in PLL. It shows an integrated rms jitter (from 100Hz to 1GHz offset) of the divided-by-2 clock (i.e., 14GHz) of 508fs, and –98.5dBc/Hz phase noise at 1MHz offset. Using an external clock, we confirm the output waveform at different data rates. The PAM4 TX is verified to operate from less than 1Gb/s to 62Gb/s. Figure 2.3.6(b) and (c) reveal the PAM4 waveform at 32 and 60Gb/s, implying rising/falling time (20-to-80%) of 12.8ps and minimum eye opening of 50mV. The sharp transition and clean eyes ensure proper data delivery. Figure 2.3.7 shows the die micrograph of the two transmitter chips, which occupy 2.1×1.0mm^2 and 1.2×0.95mm^2, respectively. A table summarizing the performance of this work and that of other state-of-the-art transmitters is shown in Fig. 2.3.7 as well.

Acknowledgment:
The authors thank TSMC university shuttle program and National Chip Implementation Center (CIC) for chip fabrication.

References:
[1] J. Jiang et al., "100Gb/s Ethernet Chipsets in 65nm CMOS Technology," *ISSCC Dig. Tech. Papers*, pp. 120-121, Feb. 2013.
[2] S. Zhai et al., "The Requirement Analysis of 400GE FEC for Gen1 PMDs," *IEEE 400Gb/s Ethernet Study Group*, July 2013. [Online]. Available: http://www.ieee802.org/3/400GSG/public/13_07/zhai_400_01_0713.pdf
[3] *40 Gb/s and 100 Gb/s Ethernet Task Force*. [Online]. Available: http://www.ieee802.org/3/ba/index.html
[4] Jri Lee et al., "A 75-GHz Phase-Locked Loop in 90-nm CMOS Technique," *IEEE J. Solid-State Circuits*, vol. 43, pp. 1414-1426, Jun. 2008.
[5] Jri Lee et al., "Study of Subharmonically Injection-Locked PLLs," *IEEE J. Solid-State Circuits*, vol. 44, pp. 1539-1553, May 2009.
[6] D. Yamazaki et al., "A 25GHz Clock Buffer and a 50Gb/s 2:1 Selector in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 240-241, Feb. 2004.
[7] K. Kanda et al., "A Single-40Gb/s Dual-20Gb/s Serializer IC with SFI-5.2 Interface in 65nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 360-361, Feb. 2009.
[8] C. Menolfi et al., "A 25Gb/s PAM4 Transmitter in 90nm CMOS SOI," *ISSCC Dig. Tech. Papers*, pp. 72-73, Feb. 2005.

 978-1-4799-0920-9/14/$31.00

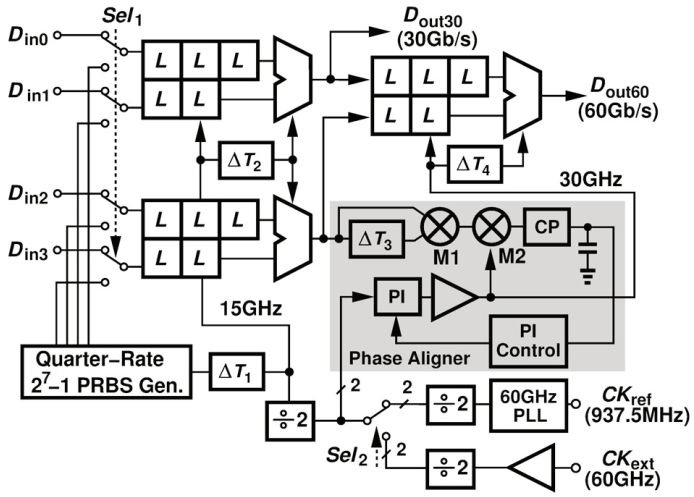

Figure 2.3.1: NRZ TX architecture.

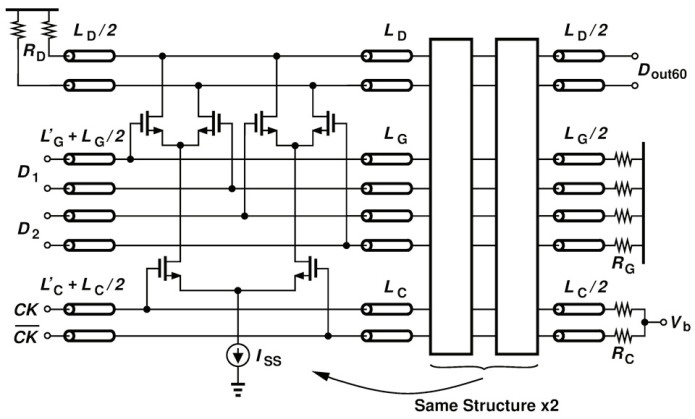

Figure 2.3.2: 60Gb/s 2:1 MUX.

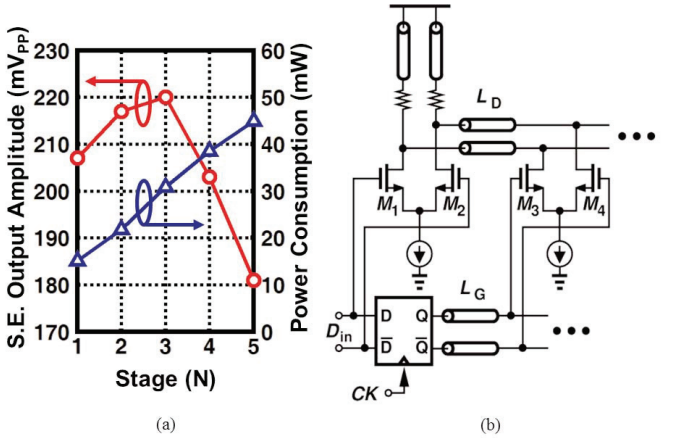

Figure 2.3.3: (a) Output amplitude and power consumption of distributed amplifier, (b) combiner.

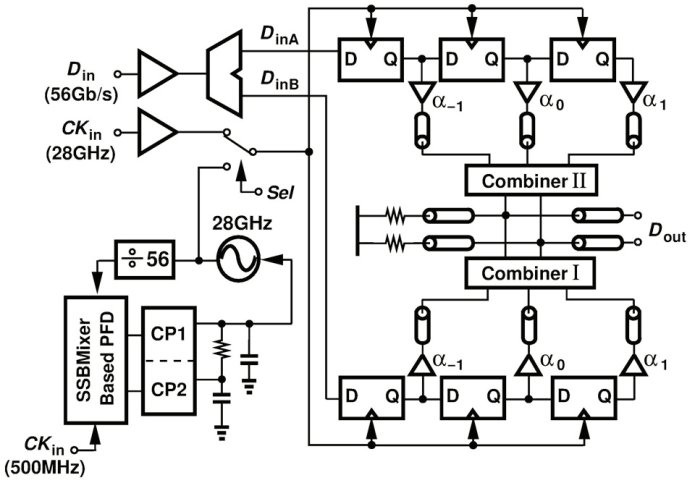

Figure 2.3.4: PAM4 TX architecture.

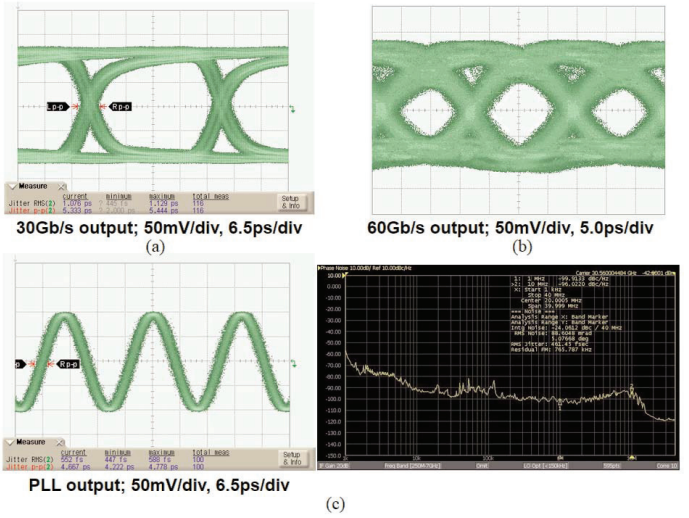

Figure 2.3.5: NRZ TX measurement results: (a) 30Gb/s output, (b) 60Gb/s output, (c) 30GHz clock waveform and its phase noise.

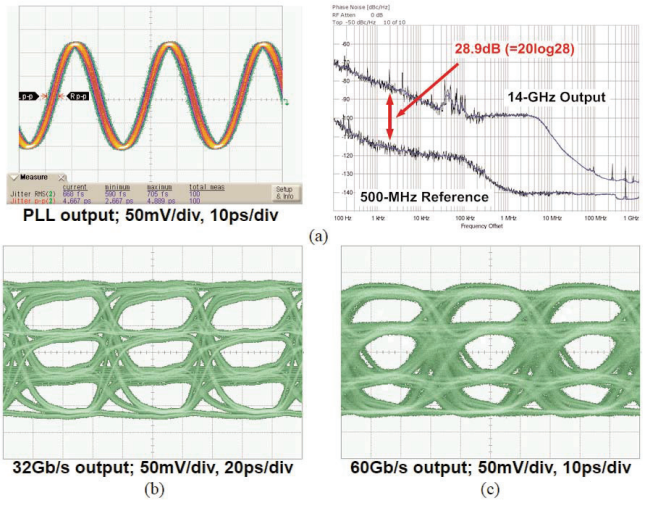

Figure 2.3.6: PAM4 TX measurement results: (a) 28GHz clock waveform and its phase noise, (b) 32Gb/s output, (c) 60Gb/s output.

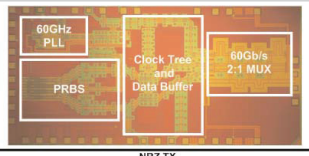

NRZ TX

	[6]	[7]	This Work
Data Rate	50Gb/s	40Gb/s	60Gb/s
Function	2:1 MUX Only	SFI 5.2 Receiver + 4:1 MUX (40Gb/s) + 4:2 MUX (20Gb/s) + 20GHz PLL	4:1 MUX + 60GHz PLL + Built-in PRBS
TX Clock PNoise @ 1MHz	N/A	N/A	-100dBc/Hz
TX Clock RMS Jitter	N/A	N/A	461fs (1kHz~20MHz)
Data Input Range (S.E.)	$1V_{PP}$	N/A	$50\sim300mV_{PP}$
Clock Input Range (S.E.)	$1V_{PP}$	N/A	$100\sim300mV_{PP}$
Data Swing (S.E.) Full-Rate Half-Rate	$70mV_{PP}$ (50Gb/s) N/A	$325mV_{PP}$ (40Gb/s) $400mV_{PP}$ (20Gb/s)	$250mV_{PP}$ (60Gb/s) $250mV_{PP}$ (30Gb/s)
Data Jitter	N/A	783fs,rms (40Gb/s)	1.08ps,rms (30Gb/s)*
20-80% Rise/Fall Time Full Rate Half-Rate	10.4ps (50Gb/s) N/A	10.22ps (40Gb/s) 14.67ps (20Gb/s)	8.0ps (60Gb/s) 7.5ps (30Gb/s)
Power Consumption	43mW	1.8W (40Gb/s mode)	450mW (MUX: 31mW)
Chip Area	1.8 x 1mm²	4.2 x 4.2mm²	2.1 x 1mm²
Technology	90nm CMOS (Shrunk to 48nm)	65nm Digital CMOS	65nm Digital CMOS

* For half-rate D_{OUT}, Full-rate data jitter is not measurable due to the limited bandwidth of oscilloscope.

PAM4 TX

	[8]	This Work
Data Rate	25Gb/s	60Gb/s
Function	Combiner + FFE	Combiner + FFE + Built-in PLL
TX Clock PNoise @ 1MHz	N/A	-98.5dBc/Hz
TX Clock RMS Jitter	N/A	508fs (100Hz~1GHz)
Data Input Range (S.E.)	N/A	$50\sim300mV_{PP}$
Clock Input Range (S.E.)	N/A	$100\sim300mV_{PP}$
D_{OUT} Full Mag. (4 Levels)	$430mV_{PP}$	$250mV_{PP}$
Min. Eye Opening	$90mV_{PP}$	$50mV_{PP}$
Horizontal Eye Opening	>0.4UI @ 25Gb/s	>0.6UI @ 60Gb/s
20-80% Rise/Fall Time	29ps @ 25Gb/s	12.8ps @ 60Gb/s
Power Consumption	101.8mW	290mW
Chip Area	1 x 0.5mm²	1.2 x 0.95mm²
Technology	90nm SOI CMOS	65nm Digital CMOS

Figure 2.3.7: Die micrograph and performance summary.

978-1-4799-0920-9/14/$31.00