A High-Speed PAM2 and PAM4 Voltage-Mode Transmitter With Auxiliary Push-Pull Current-Mode Swing Enhancment

Hanson Liu, Eamon Lee, Ethan Tran

I. Introduction

GIVEN the rapid explosion of data traffic in modern data centers, the emergence of large-scale artificial intelligence workloads, and the increasing power demands of large computing infrastructure, there is strong motivation to find ways to push the transmit speed of digital data. In particular, chip-to-chip wireline communication has become increasingly critical, as the majority of data movement within and between data centers occurs over high-speed electrical and optical I/O's rather than wireless links.

However, in addition to the pursuit of pushing speed, the concern of power efficiency remains a critical design constraint in high-speed wire-line design. While reducing transmit power is critical to improve overall energy efficiency, sufficient signal integrity must be maintained to ensure the receiver can properly function. Particularily, the signal provided by the transmitter must provide an adequate signal-to-noise ratio (SNR) at the input of the receiver despite attenuation from the channel, intersymbol interference (ISI), additive random noise, and additive random jitter. This presents the trade-off between power consumption and link performance, and is the motivation to explore different transmitter architectures that maximize output swing and signal quality per bit per unit power.

This report presents a hybrid voltage-mode transmitter complemented with an **auxiliary push–pull current-mode (CM)** swing enhancement stage, which improves transmit eye height degraded by channel-induced high-frequency loss. Due to the hybrid nature of the design, equalization can now also be implemented in voltage and current-mode. However, the available current in the CM stage is limited by the headroom provided by supply. To address this constraint, a **high-speed level shifter** is introduced to relax the headroom requirement, enabling increased current injection and, consequently, larger output swing. This enhancement, if done correctly so that it is not the main bottle-neck of the link, will therefore support even higher achievable data rates. Furthermore, the proposed design targets high tap resolution, enabling finer control over equalization as the data rate increases.

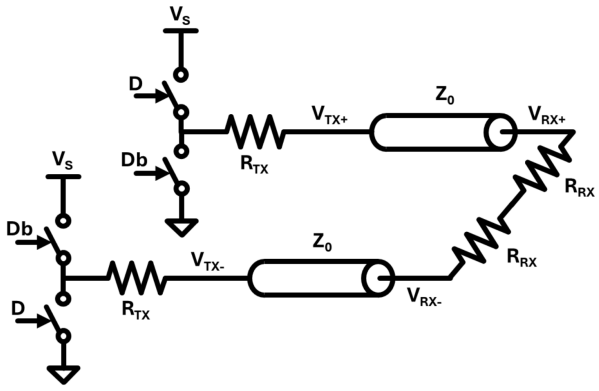

Fig. 1. SST Transmitter With Differential Termination

II. Design of the PAM2 SST Driver

The most basic implementation of a voltage-mode (VM)/source-series terminated (SST) transmitter consists of switches driven by the input data stream that pull the output either to a logic high or a logic low. For low-swing applications, both the switching function and the resistive termination can be realized using a single n-channel device. However, in high-swing VM architectures, the pull-up and pull-down paths can be implemented using pull-up and pull-down switches in combination with series resistors, as shown in Fig. 1. In this case, the effective output impedance, consisting of the sum of active and passive terminations, is designed to match the characteristic impedance of the channel, Z_0.

There are several reasons why one would opt to include the extrinsic resistor. First and foremost, it's hard to achieve a large enough resistance on the switch when designing for high feed-forward equalization (FFE) tap resolution. Furthermore, there is a non-linearity that comes with the resistance of an active device switch, as any changes in the drain to source voltage could affect the resulting resistance seen looking into the output of the device. Therefore, the optimal design philosophy when it comes to designing a fast SST transmitter is to design the switches to be as fast, and therefore small as possible, while still maintaining a series resistance that is large enough to counteract the linearity issue.

There are two primary benefits of a voltage-mode driver topology compared to a current-mode driver. One key advantage is that the transceiver exhibits an average power consumption of $\frac{1}{4}V_{DD}^2 Z_0$ with a differential termination, compared to $\frac{3}{4}V_{DD}^2 Z_0$ for a current-mode driver with differential termination. Another benefit of the voltage-mode driver is that its output impedance is inherently on a first order matched to the characteristic impedance of the channel.

Upon inspection of Fig. 1, we can see that for a differentially terminated SST transmitter, the maximum swing we can achieve is $V_{DD,ppd}$.

$$\max\left(\mathrm{V_{tx}^{+/-}}\right) = \mathrm{V_{DD}} \cdot \frac{\mathrm{Z_0 + 2Z_0}}{\mathrm{2Z_0 + 2Z_0}} = 0.75\,\mathrm{V_{DD}}$$

$$\min\left(\mathrm{V_{tx}^{+/-}}\right) = \mathrm{V_{DD}} \cdot \frac{\mathrm{Z_0}}{\mathrm{2Z_0 + 2Z_0}} = 0.25\,\mathrm{V_{DD}}$$

$$\text{When } D = \text{logic “1”}, \quad \mathrm{V_{rxd}} = 0.75\,\mathrm{V_{DD}} - 0.25\,\mathrm{V_{DD}} = 0.5\,\mathrm{V_{DD}}$$

$$\text{When } D = \text{logic “0”}, \quad \mathrm{V_{rxd}} = 0.25\,\mathrm{V_{DD}} - 0.75\,\mathrm{V_{DD}} = -0.5\,\mathrm{V_{DD}}$$

$$\mathrm{V_{rx,ppd}} = \mathrm{V_{DD,ppd}} \tag{1}$$

In order to prevent eye closure and meet stringent eye-height specifications based on a required SNR, this paper's take on swing-enhancement is introduced in the next section. The SST transmitter in this report is implemented with a simple inverter switch with resistance R_{on} tied in series to an extrinsic poly-resistor R_e, designed to meet a ratio of $\frac{R_e}{R_{on}} = 3$ for improved linearity and speed. On its own, the SST driver can reach speeds up to 12Gbps before compression from the channel prevents the optical eye seen at the receiver from meeting a minimum eye-height and aperture to meet a bit error rate (BER) of 10^{-12}.

III. Auxiliary Current-Mode Driver

The previously presented SST driver is now complemented with a current-mode logic (CML) push-pull driver as shown in Fig. 2, which serves to source and sink current simultaneously from the differential output nets depending on the data pattern.

The reason why this driver works is due to the concept of superposition, as the output voltage consists of the combined contributions from the SST driver, which establishes a baseline through resistive voltage divider of the supply, and the CML driver, which injects a differential current into the termination network. The superposition of these effects results in an increased differential output swing at the output of the transmitter.

One important feature of this hybrid VM–CM driver is that the current-mode (CM) stage does not require explicit common-mode feedback (CMFB) to operate correctly. This is because the differential outputs of the CM driver see a

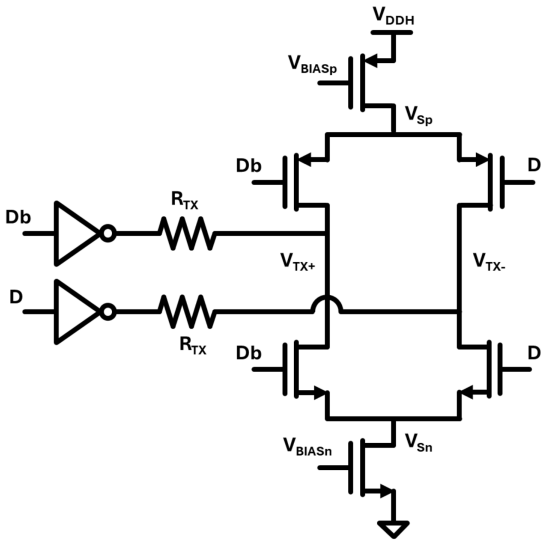

Fig. 2. Hybrid Transmitter featuring CM driver for swing-enhancement

low impedance in common mode set by the small termination resistance of the SST stage, which keeps the output common-mode voltage well defined. As a result, any mismatch in the branch currents produces only a small perturbation in the output voltage, and thus the common-mode level remains well defined without the need for dedicated CMFB circuitry. However, the presented design uses controlled biasing using scaled replica circuits of the transmitter, in which the tail current source bias voltages are regulated through feedback loops implemented with **two-stage, low-side Miller-compensated OTAs** designed for 55 dB midband gain and a 60 MHz unity-gain bandwidth. The replica circuits generate the appropriate bias voltages by enforcing the desired common-mode operating conditions, and these bias voltages are subsequently copied to the main driver. This feedback mechanism ensures that the output common-mode voltage of the transmitter is consistent across different process corners.

Another nice feature of the CML's is that their impedance does not affect the overall driver impedance, since the input of the channel will see a large output impedance of the current source in parallel with the small impedance of the SST driver.

Upon inspection, it can be inferred that for every 1mA of current supplied and sinked from the differential output, there is an extra $100\mathrm{mV_{ppd}}$ that can be added to the swing, assuming a 50Ω termination. This naturally raises the question: Why can't the source currents be further increased to maximize the achievable swing?

First and foremost, the CM is implemented with current-steering switches that steer the tail current sources to the differential output depending on the data pattern, which means that it'll be a constant source of static power consumption. Secondly, there are diminishing returns in increasing the transmitter's swing, since

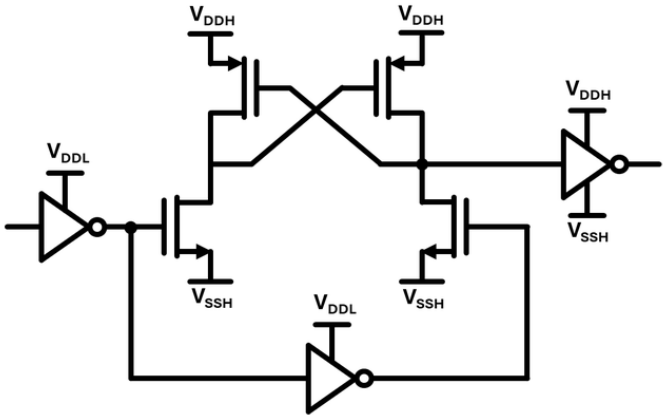

Fig. 3. Type-1 High-Speed Level Shifter

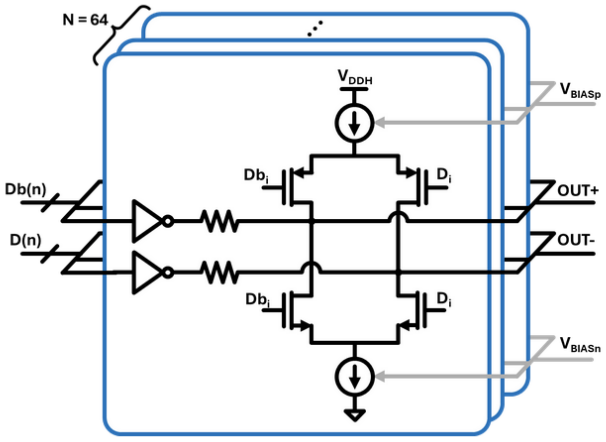

Fig. 4. Multi-Slice Implementation of Hybrid Transmitter for FFE

the channel loss scales in dB with frequency. Therefore, any additional voltage swing provided at the transmitter is exponentially attenuated, so incremental increases in swing do not translate to meaningful improvements at the receiver and will eventually degrade energy efficiency.

Thirdly, and most important of all, the main limitation on how much current you can use is actually the headroom provided to the tail current sources by the supply. To highlight the issue of headroom, consider the case of PAM2 with differential termination, 64 slices, and a $V_{DD} = 1\,V$ supply.

From the previously derived results in Eq. 1, one differential output on the transmitter side is $750\,mV$, while the other is $250\,mV$. If a total current of $I_s = 3\,mA$ is selected, these values become $825\,mV$ and $175\,mV$, respectively.

Assuming a switch resistance of $R_{switch} = 1.5\,k\Omega$, the resulting node voltages V_{sp} and V_{sn} on Fig. 2 are:

$$V_{sp} = 825\,mV + \frac{3\,mA}{64} \cdot 1.5\,k \approx 900\,mV$$
$$V_{sn} = 175\,mV - \frac{3\,mA}{64} \cdot 1.5\,k \approx 100\,mV \quad (2)$$

From this analysis, it can be concluded that the tail current sources have only **100mV of headroom** to remain in saturation. One potential approach to address this limitation is to increase the supply voltage. However, doing so necessitates the use of a high-speed level shifter, as the logic levels generated by the PRBS source are constrained to 0V and 1V. Without level shifting, the increased supply voltage may prevent certain p-channel devices from fully turning off, and may also introduce reliability concerns for devices not rated for V_{GS} or V_{DS} exceeding 1V.

IV. High-Speed Level Shifter

The level shifter implemented in this work shown in Fig. 3 is a Type-1 cross-coupled latch inverter. The operation begins with an inverter supplied by the original supply rails, $V_{DD,low}$ and $V_{SS,low}$, which drives a cross-coupled differential latch pair. This latch utilizes positive feedback to shift the output levels to the higher supply domain, $V_{DD,high}$ and $V_{SS,high}$.

The inverter stages and cross-coupled latch are designed with minimum channel length while maintaining equal drive strength. Since the delay of the cross-coupled latch pair is proportional to $\frac{C_x}{g_m}$, increasing the transistors transconductance reduces the effective time constant associated with node transitions, thereby improving the speed of level conversion.

Simulations show a rise time of $15.6\,ps$ and a fall time of $13.13\,ps$, resulting in a total delay of $28.73\,ps$. This indicates that for data rates up to $35\,Gb/s$, the level shifter is **not expected to be the bottleneck of the link**.

Since the level shifter is required to drive 64 slices of the hybrid driver, its intrinsic drive strength is insufficient if directly connected. To address this, a pre-driver chain of inverters is introduced, designed with a fan-out of 4 to ensure adequate drive strength while maintaining high-speed operation.

V. Equalization

As the data rate increases, equalization becomes necessary to counteract the frequency-dependent loss of the channel, which increasingly attenuates the high-frequency components of the signal and leads to intersymbol interference (ISI). A common approach to mitigate this effect is feedforward equalization (FFE), which can be modeled as a discrete-time finite-impulse-response (FIR) filter operating at the symbol rate. Qualitatively the FFE acts as a **form of predistortion**, shaping the transmitted waveform such that channel-induced distortion is compensated at the receiver.

In CM implementations, equalization is achieved by steering different amounts of current across multiple slices

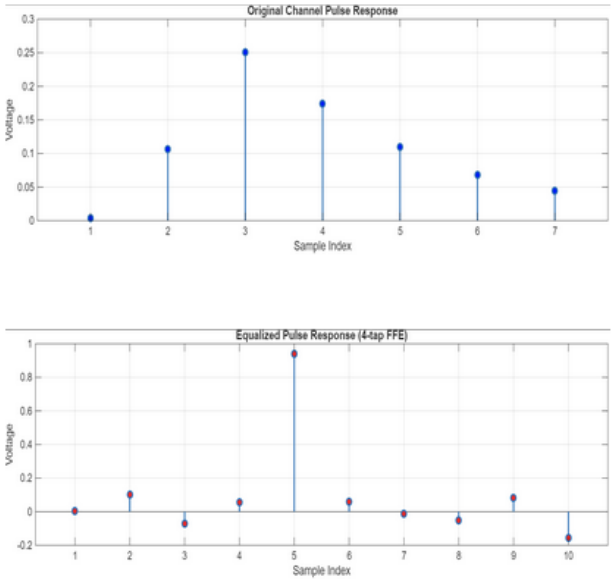

Fig. 5. Simulated MATLAB Pulse Response After 4-Tap FFE

to realize the pre-cursor, main-cursor, and post-cursor tap weights, with contributions summed in parallel at the output. In contrast, VM FFE implemented weights taps by segmenting the driver into multiple slices with appropriately scaled output impedances, thus adding voltage contributions "in series". If done correctly, the total number of slices remains constant, ensuring that the effective output impedance seen by the channel is still maintained at Z_0. Because the presented design in this report is a combination of the two, each slice associated with a given tap coefficient simultaneously contributes through both series (voltage) and parallel (current) summation mechanisms, as illustrated in Fig. 4. This hybrid operation preserves full flexibility in implementing equalization without compromising control over the tap weighting.

This work employs a **zero-forcing (ZF) equalization** approach, where the tap coefficients are selected such that the pre-cursor and post-cursor ISI components are ideally driven to zero at the sampling instants, as shown in Fig. 5. These coefficients can be obtained through multiple methods. One approach is to simulate the pulse response of the combined transmitter and channel, sample it at the symbol rate, and solve for the coefficients under the ZF condition. Alternatively, the channel impulse response can be extracted from the provided S-parameters using external tools such as the MATLAB SerDes Toolbox, and the coefficients can then be computed algorithmically for a target data rate. Both approaches were utilized in this work to derive the FFE tap coefficients.

VI. PAM4 Design

This report also features a PAM4 design of the same hybrid architecture that was designed for PAM2, as shown in Fig. 6 . There are several clear new differences in order

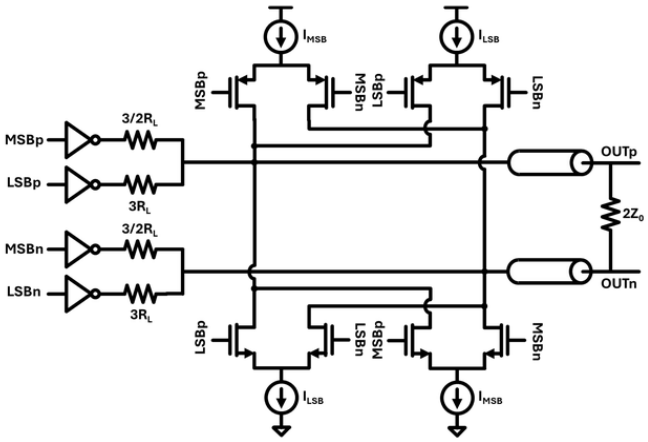

Fig. 6. PAM4 Hybrid Design

to achieve PAM4. First and foremost, the four different voltage levels for the PAM4 must now be created with a specific voltage divider network that equally adds up the contribution from the MSB and LSB datastream respectively.

Compared to PAM2, PAM4 receives much more benefit from swing enhancement due to the multilevel nature of the modulation scheme. To implement swing enhancement, a similar case of creating the voltage levels is shown in the following figure, where now different levels of current is sourced/sinked depending on what values the MSB and LSB datastream values are provided to the current steering switches.

Although Linearity is less critical in PAM2, where only two output voltage levels are required, issues of nonlinearity becomes significantly more critical in PAM4 due to the need to generate the two intermediate levels. Therefore, to properly design a PAM4 transmitter, several sources of nonlinearity must be considered.

One such source arises from the SST driver: if the external termination resistance is not sufficiently larger than the on-resistance of the switching devices, different data patterns will produce varying output voltage levels. This occurs because the effective switch resistance depends on the drain to source voltage of the transistors, which varies with the output voltage, leading to signal-dependent nonlinearity.

Another significant source of nonlinearity arises from the CM driver. Similar to the SST driver, different bit combinations produce different differential output voltage levels, which in turn modulate the drain to source voltage of the tail current source devices. As a result, the current sources may deviate from ideal behavior if sufficient voltage headroom is not maintained. This creates a design tradeoff: the current-steering switches must be kept small to support high data rates, but reducing their size limits the voltage headroom available to the tail current sources. Furthermore, since the current sources are then naturally

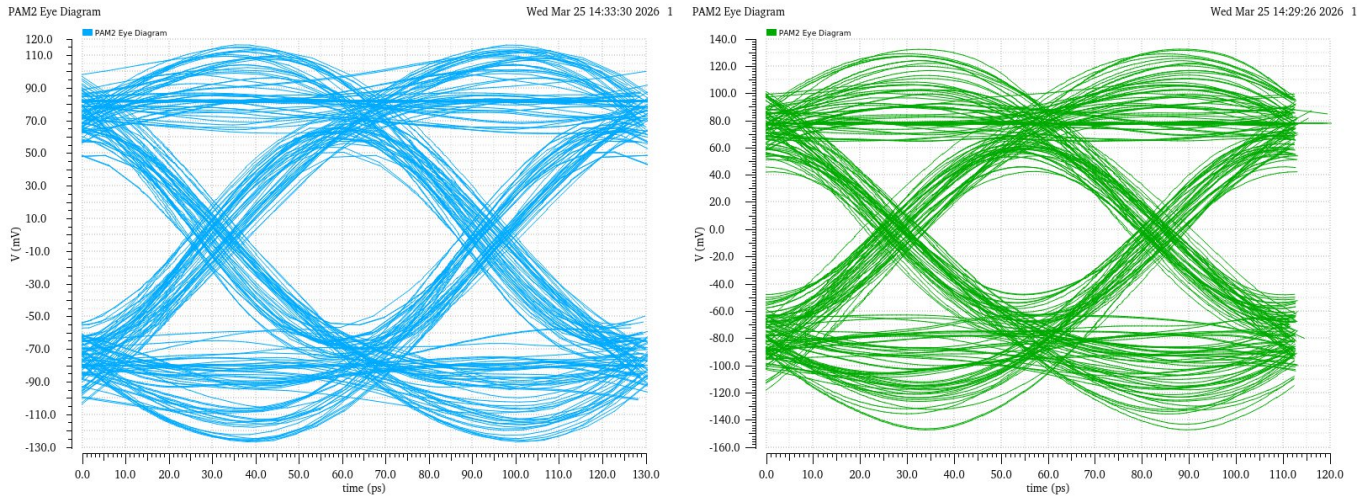

(a) PAM2 16Gbps RX Input Eye

(b) PAM2 18Gbps RX Input Eye

Fig. 7. PAM2 RX input eye diagrams meeting additive random noise and jitter specifications for: (a) $V_{DD} = 1\,V$ design without level shifters, (b) $V_{DD} = 1.3\,V$ design with level shifters.

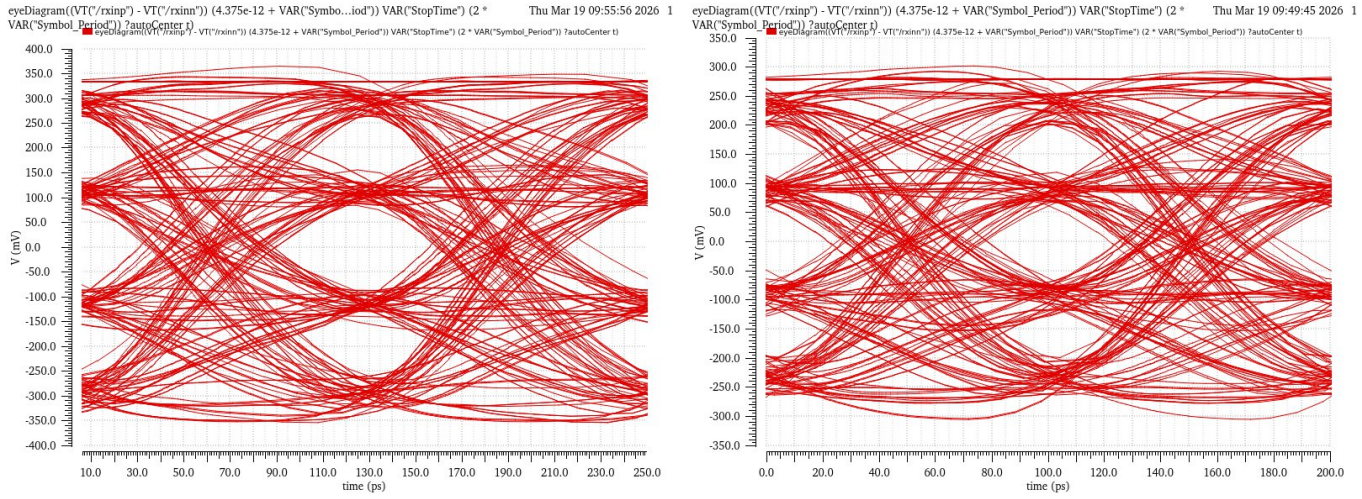

(a) PAM4 16Gbps RX Input Eye

(b) PAM4 20Gbps RX Input Eye

Fig. 8. PAM4 RX input eye diagrams meeting additive random noise specifications for: (a) $V_{DD} = 1\,V$ design, (b) $V_{DD} = 1.3\,V$ design

operating near the edge of saturation, their output current becomes dependent on the drain to source voltage, resulting in data-dependent nonlinearity. This shows up in performance as uneven spacing between PAM4 levels, where intermediate levels (e.g., (MSB, LSB) = (1,0) and (0,1)) may exhibit larger than ideal magnitudes compared to the intended uniform step size based on the full scale of the supply.

There are several ways the presented design aimed to counteract this degradation in Ratio Level Mismatch (RLM). First and foremost, the extrinsic resistor is designed to be much larger in scale compared to the switch resistance, thus reducing the nonlinearity of the VM transmitter. Furthermore, the supply was increased to a voltage of 1.3V, providing the CLM with more headroom in order to be more agnostic of variations in different data patterns.

Another critical impairment introduced by PAM4 is the significantly reduced deterministic jitter budget PAM4 allows for in comparison to PAM2. Unlike PAM2, where the data transitions from logic high and logic low are very close in transition time, PAM4 has several different possible transition types between symbols. Each transition will experience a different time due to varying voltage step sizes and driver conditions, leading to timing variations depending on the transmitted bit pattern. Furthermore, channel bandwidth limitations further exacerbate this issue, as the inner eye transitions are now also susceptible to different ISI. Because of these reasons, there will unavoidably be far more deterministic jitter and reduced timing margin at the receiver.

TABLE I
PERFORMANCE SUMMARY OF THE PROPOSED TRANSMITTERS.

Design	VDD	Tap Resolution / # of Taps	Max Rate	BER (w/ add. Noise)	BER (w/ add. Jitter)	Energy
PAM2	1 V	6 bits / 4 taps	16 Gb/s	10^{-12}	10^{-12}	0.71 pJ/bit
PAM2 w/ HS Level Shifters	1.3 V	6 bits / 4 taps	18 Gb/s	10^{-12}	10^{-12}	3.5 pJ/bit
PAM4	1.3 V	5 bits / 3 taps	20 Gb/s	10^{-12}	3.3×10^{-3}	0.65 pJ/bit

VII. RESULTS

The performance targets used to evaluate the transmitter are based on receiver (RX) input requirements for a specific bit-error rate. Specifically, the RX input eye sensitivity is 20mV_{ppd} at a BER $< 10^{-12}$ under additive voltage noise $\sigma_n = 5\text{mV}_{\text{rms}}$, and the RX input eye aperture time is $0.3\text{UI}_{\text{ppd}}$ at BER $< 10^{-12}$ under additive random jitter $\sigma_{RJ} = 0.035\text{UI}_{\text{rms}}$. Since additive noise and random jitter are not capturable in simulation, these specifications instead define the minimum required eye height and eye width, and the transmitted eye is evaluated against these targets to ensure sufficient margin. Simulated RX eye diagrams of the fastest possible data rate the presented PAM2 and PAM4 transmitters was able to meet is shown in Fig. 7 and Fig. 8, respectively.

The results for the PAM2 and PAM4 transmitters, shown in Table I, are obtained through simulation in Cadence Virtuoso using ideal PRBS10 signal generators with ideal delays for the different weighted FFE taps. In practice, the inclusion of non-ideal PRBS generators and serialization circuitry (e.g., 2:1 serializers, 4:1 serializers) would introduce additional deterministic jitter, thereby reducing the maximum achievable data rate. It is important to note that, due to the limited jitter margin in PAM4, achieving the target BER requires more careful control of equalization that the presented design could not achieve. As such, the reported PAM4 BER based on the specification's additive jitter is simulated at the data rate at which the simulated eye diagram meets the minimum RX input eye height specification.

As expected for a voltage-mode transmitter, the proposed design demonstrates strong energy efficiency compared to traditional current-mode (CM) transmitters, which typically exhibit energy consumption that is many times higher than the values reported in this work.

The primary contributor to power consumption in the proposed architecture is the static current drawn by the auxiliary CM driver. This is further increased in the configuration with high-speed level shifters as reflected in Table I, as CM driver were designed to source and sink up to 7 mA to maximize output swing. Despite this increase, the required current remains significantly lower than that of conventional CM transmitters, which may require up to 40 mA of total current to achieve a differential swing of $1\,\text{V}_{\text{ppd}}$ across the transmitter and channel impedance.

VIII. CONCLUSION

This work presents a hybrid voltage-mode/current-mode (VM–CM) transmitter architecture that leverages the efficiency of SST-based voltage-mode operation while incorporating an auxiliary current-mode stage for swing enhancement without sacrificing the inherent impedance matching and energy efficiency advantages of voltage-mode drivers. Headroom constraints on the tail current sources of the CM driver was addressed through the introduction of a high-speed level shifter, enabling increased current injection and increased headroom.

Simulation results demonstrate that the proposed PAM2 transmitter achieves data rates up to 18 Gb/s while meeting receiver eye requirements under noise and jitter constraints, with competitive energy efficiency compared to conventional current-mode transmitters. A PAM4 extension of the architecture was also explored, highlighting both the benefits of swing enhancement and the increased sensitivity to nonlinearity and deterministic jitter inherent to multilevel signaling.

Overall, the presented hybrid VM–CM approach provides an effective trade-off between power efficiency, output swing, and equalization flexibility, making it a promising solution for high-speed short-reach wireline communication links in data center environments and modern serial I/O interfaces such as memory, peripheral, and chip-to-chip interconnects.

ACKNOWLEDGMENT

The authors would like to thank Professor Tzu-Chien Hsueh and his teaching assistant Sumilak Chaudhury for all their help and guidance throughout this project!

REFERENCES

[1] ECE283 Lecture Slides Courtesy of Professor Hsueh
[2] 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,"* in IEEE Journal of Solid-State Circuits, vol. 51
[3] Y. -H. Song and S. Palermo, *"A 6-Gbit/s Hybrid Voltage-Mode Transmitter With Current-Mode Equalization in 90-nm CMOS,"* in IEEE Transactions on Circuits and Systems II: Express Briefs, vol. 59
[4] L. Fassio et al., *"A Robust, High-Speed and Energy-Efficient Ultralow-Voltage Level Shifter,"* in IEEE Transactions on Circuits and Systems II: Express Briefs, vol. 68