A 0.055pJ/bit/dB 42Gb/s PAM-4 Wireline Transceiver with Consecutive Symbol to Center (CSC) Encoding and Classification for 26dB Loss in 16nm FinFET

Ramin Javadi, Tejasvi Anand

Oregon State University, Corvallis, OR, USA

The growth of data-intensive applications such as large language AI models has increased the demand for higher data throughput in wireline links. Due to the bandwidth-limited nature of the wireline channel, increasing the data rates across the same physical distance of the communication channel results in more inter-symbol interference (ISI). Consequently, more channel equalization is required to compensate for ISI, which increases the energy/bit of the communication link. Researchers have discovered that machine learning (ML) inspired approaches [1-2] including feature extraction and classification provide a more efficient solution for compensating the channel loss compared to the conventional equalization techniques like FFE, DFE, and CTLEs [3-7]. However, the prior works on ML inspired links are limited to NRZ modulation only. In this work, we introduce an energy-efficient ML inspired transceiver that leverages feature extraction and classification to transmit encoded PAM-4 data across a wide range of channel loss (13dB to 26dB) while maintaining BER < 10^{-11} without using any conventional equalizers. Additionally, we propose a data encoding scheme, consecutive symbol to center encoding (CSC) to encode PAM-4 and provide identifiable attributes to the transmitted signal, which helps to increase the channel loss compensation range and reduce the complexity of the classifier. Since ISI is a deterministic non-ideality, the proposed decision-tree based classifier (on-chip) is designed to learn both the channel characteristics and the CSC data encoding, enabling it to accurately detect the original transmitted data in the presence of ISI with a latency of only 10 unit interval (UI). The decision tree classifier operates with a low-power feed-forward architecture without any feedback timing constraints, allowing the proposed transceiver to achieve 0.055pJ/bit/dB, which is ~2x lower than prior work [3-5] while compensating for nearly the same channel loss.

Fig. 1 shows the proposed transceiver scheme. In the proposed architecture, the multi-tap FFE, DFE, and high-peaking CTLE are replaced with data encoding at the transmitter (Tx) and feature extraction, classification, and decoding blocks at the receiver (Rx). The encoder adds recognizable characteristics to the standard PAM-4 data before sending the signal through the channel. At the Rx front-end, relevant features are extracted from the closed eye, and the classifier maps the received information to the voltage levels corresponding to the transmitted data. The decoder then converts the signal into the original PAM-4 data. In the proposed CSC encoding scheme, four levels of PAM-4 (+3, +1, -1, -3) are mapped to five distinct symbols (+3, +1.5, 0, -1.5, -3) with an additional voltage level at 0V, which ensures that any consecutive identical symbols (CIS) are shifted to a central common-mode voltage (0V). This encoding approach increases transition density during CIS, reduces the DC content of the data, and introduces redundancy within the symbol space (voltage domain) without increasing coding overhead. A conventional linear partial response signaling scheme such as Dicode encoding ($1-z^{-1}$) on PAM-4 could have also removed the CIS. However, Dicode on PAM-4 would have created 7 voltage levels, which would reduce the SNR by 6dB. The proposed non-linear CSC encoding limits the voltage levels to 5 and the SNR penalty is only 2.5dB.

Fig. 2 shows the off-chip training process of a decision tree classifier in MATLAB. The CSC encoded data is transmitted through an approximate channel model and received by the Rx front-end. Simultaneously, CSC encoded data is fed directly to the learning algorithm, bypassing the lossy channel. At the receiver 5 first-order features (F1-F5) are extracted at every UI. Higher order features such as slope are avoided to reduce the Rx front-end power consumption [1]. Each of these 5 features generates a 1-bit output based on the voltage thresholds Vth1 to Vth5 forming a feature vector. Feature vectors consisting of the current UI, past 6UIs, and future 3UIs (total 10UI data) are then given to the classifier. The classifier is trained with 5 features to compensate for 17dB to 26dB channel losses, while 3 features were sufficient to train it for 13dB to 17dB. The programmability of features is added in the transceiver to demonstrate this trade-off. By utilizing 5x10 bits of information in every UI (5 Features x (6pre+1+3post) UI) and the fact that there is no CIS in the data pattern, the decision tree classifier maps the received signal to one of the five encoded voltage levels. Using supervised learning, the decision tree learns from the labeled data during the training process, where the correct voltage level (label) corresponding to each feature vector is known. Since a large volume of data is essential for effectively training ML models, PRBS-17 pattern was selected for training. Patterns beyond PRBS-17 added complexity without improving classifier performance. Additionally, employing CSC mapping significantly increases redundancy, making PRBS-17 adequate to optimize training and compensate for a wide range of channel losses, which ensures robust performance in various conditions. Finally, the MATLAB classifier is converted to Verilog code for on-chip synthesis using a custom-developed algorithm.

Fig. 3 shows the proposed transceiver with half-rate architecture. The transmitter consists of two 32-bit PRBS generators (MSB and LSB), a CSC encoder, four 32:1 multiplexers, and a source-series terminated output driver. The receiver includes five half-rate samplers to extract 5 features, five 2:16 de-multiplexers, shift registers, 16 classifiers, and 16 CSC decoders. The encoder, decoders, and classifiers are positioned at the back-end, where they operate at a lower frequency, which results in lower power consumption. A Feature Enable signal can activate or deactivate features/slicers based on the channel loss requirement. An external clock source supplies the clock tree where two clock phases are distributed to the Tx and Rx. The chip is fabricated in 16nm FinFET.

Fig. 4 depicts the measured power spectrum density (PSD) of both PAM-4 and proposed PAM-4+CSC which indicates PSD reduction at low frequencies in the CSC scheme. This PSD reshaping along with added pattern redundancy in CSC, archives a 10dB gain in channel loss compensation compared to standard PAM-4 while using fewer features with lower latency. Three near-end transmitter eye measurements show the voltage levels of the CSC encoded PAM-4 and the effect of ISI on CSC signaling across different channel losses and data rates. PRBS-7 data is transmitted through an SST driver with a supply voltage of 0.9V. At 8Gb/s, the Tx near-end output clearly shows five distinct voltage levels corresponding to 5 symbols of the CSC encoding scheme. At 32Gb/s, the signal becomes affected by ISI due to the losses exhibited by the chip package and PCB, leading to a partial closure of the eye. At 42Gb/s, the eye is completely closed at the Tx near-end, indicating severe ISI impact (~12dB loss at 10.5GHz).

Fig. 5 shows the transceiver measurement results. The measured channel losses at 8GHz and 10.5GHz are 17dB and 26dB, respectively. At 42Gb/s and 32Gb/s with 5 features, the measured transceiver's bathtub after the CSC decoder shows 0.06UI and 0.1UI horizontal opening for BER < 10^{-11}, respectively. At 32Gb/s and 17dB channel loss, the transceiver can operate with both 3 and 5 features. Reducing the number of features from 5 to 3 improves the energy efficiency by 20% (1.53pJ/b to 1.23pJ/b). With 3 features, the transceiver can compensate 13dB to 17dB, and with 5 features, it can compensate 17dB to 26dB with a BER < 10^{-11}. The transceiver's ability to compensate for a wide channel loss range (13dB to 26dB) demonstrates the effectiveness of the decision tree training in preventing overfitting or underfitting. Vertical and horizontal margins for the 5 features are measured by sweeping the sampling time and individual threshold voltages (Vth1-Vth5). At 42Gb/s and a BER < 10^{-11}, the average horizontal and vertical margins of 5 features are 0.06UI and 10mV. The power breakdown and comparison with the state-of-the-art are shown in Fig. 6. When operating with 5 features over 26dB channel loss, the proposed transceiver consumes 60.3mW at 42Gb/s and energy efficiency of 1.43pJ/b, with the encoder, decoders, and classifiers together accounting for 4.3mW, which represents only 7% of the total power consumption. This power includes the clock tree power: clock distribution buffers, delay lines, and dividers. Compared to other state-of-the-art designs, this work achieves a low energy efficiency per channel loss of 0.055pJ/b/dB. The die photo is shown in Fig. 7.

2025 IEEE Custom Integrated Circuits Conference (CICC) | 979-8-3315-1745-8/25/$31.00 ©2025 IEEE | DOI: 10.1109/CICC63670.2025.10982737

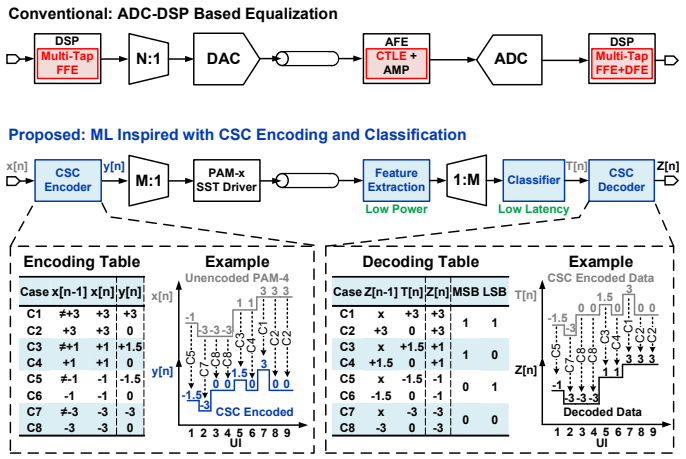

Fig. 1. Conventional ADC-DSP based wireline transceiver and proposed ML inspired wireline transceiver with the consecutive symbol to center (CSC) encoding.

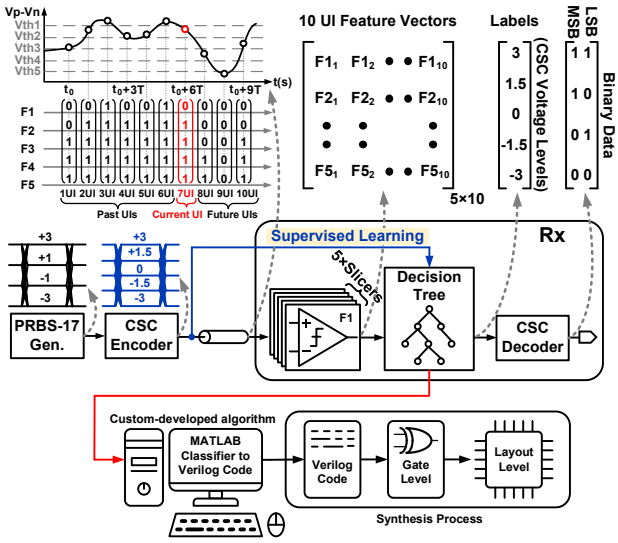

Fig. 2. Decision tree training process using supervised learning.

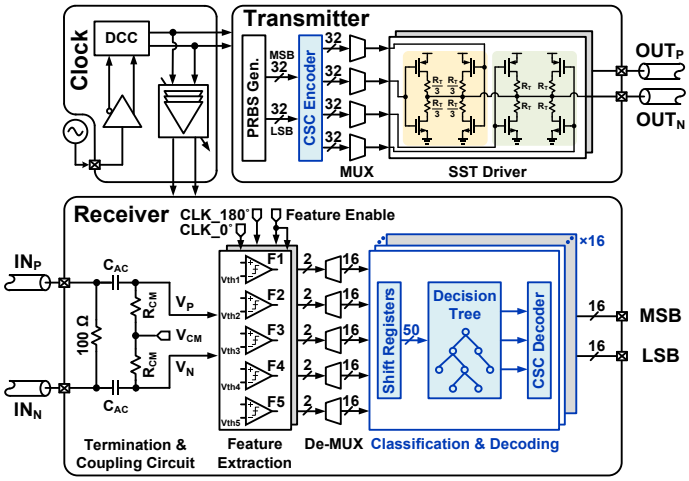

Fig. 3. Block diagram of the proposed PAM-4 transceiver with consecutive symbol to center encoding, feature extraction and classification.

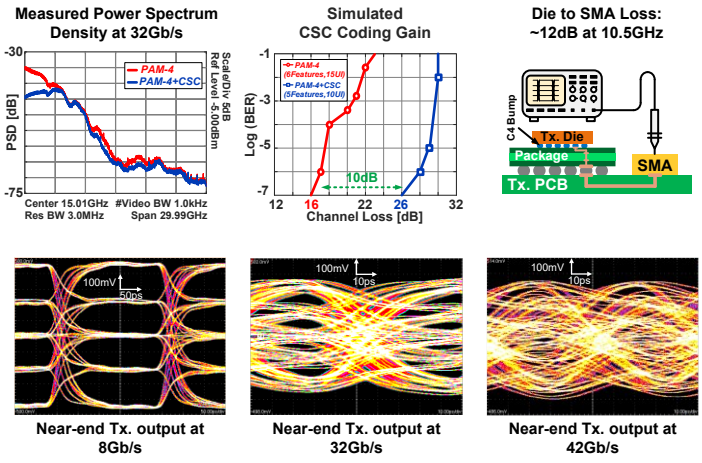

Fig. 4. Measured PSD of PAM-4 and PAM-4+CSC, transmitter's near-end output at 8Gb/s, 32Gb/s, and 42Gb/s with PRBS-7 data along with simulated coding gain of the CSC encoding.

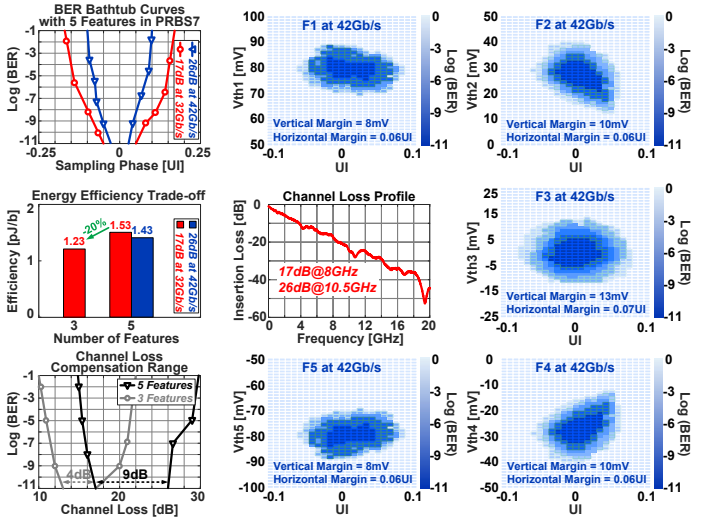

Fig. 5. Measured bathtub plot (CSC decoder output), channel loss profile, impact of number of features in channel loss compensation range, energy efficiency, and effects of feature's threshold voltage on BER for the proposed transceiver at 42Gb/s with PRBS-7.

	This Work			[1] VLSI'21	[2] JSSCC'20	[3] ISSCC'18	[4] ISSCC'18	[5] ISSCC'22	[6] ISSCC'22	[7] ISSCC'24
Technology	16-nm FinFET			65-nm CMOS	65-nm CMOS	16-nm FinFET	16-nm FinFET	28-nm CMOS	5-nm FinFET	28-nm CMOS
Data Rate [Gb/s]	32	32	42	13.8	13.6	56	64	112	113	112
Encoding /Modulation	PAM-4 + CSC			HT	Dicode	PAM-4	PAM-4	PAM4	PAM4	PAM4
Bits per Symbol	2			1	1	2	2	2	2	2
Number of Features	3	5	5	3	X	X	X	X	X	X
Latency [UI]	10			19	5	-	-	-	-	-
Equalization	X			CTLE 2-TapFFE	X	CTLE FFE, DFE	CTLE FFE	CTLE 4-TapFFE	CTLE FFE	CTLE 3-TapFFE
Log_{10}(BER)	<-11			<-6	<-12	<-12	<-4	-11	<-11	<-12
Power [mW]	39.4	49.1	60.3	100.9	34.8	545	373.6	256	175	241.9
Loss@Nyquist [dB]	17	17	26	44.7	24.2	32	29.5	20.8	11.5	8.5
Power Efficiency [pJ/b]	1.23	1.53	1.43	7.3	2.56	9.7	5.8	2.29	1.55	2.16
FoM [pJ/b/dB]	0.072	0.09	0.055	0.16	0.106	0.3	0.2	0.11	0.13	0.25
Area [mm²]	0.218			0.236	0.167	x	0.252	0.123	0.264	0.236

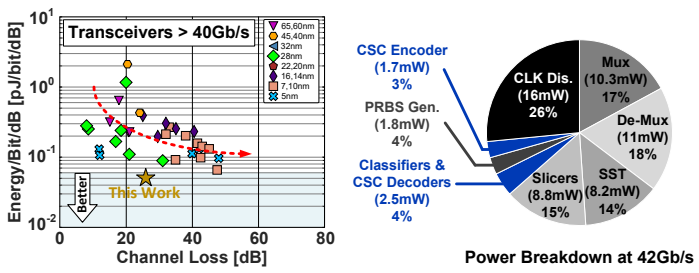

Fig. 6. Performance summary, power breakdown and comparison table with state-of-the-art designs.

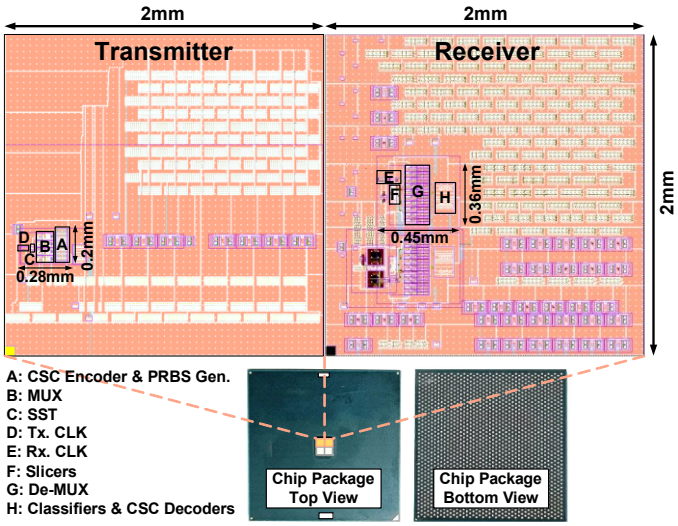

Fig. 7. Die photo.

Acknowledgements

This work was supported in part by the Center for Ubiquitous Connectivity (CUbiC), sponsored by the Semiconductor Research Corporation (SRC) and the Defense Advanced Research Projects Agency (DARPA) under the JUMP 2.0 program. We thank Intel for the chip fabrication.

References

[1] Z. Wang et al., "A machine learning inspired transceiver with ISI-resilient data encoding: hybrid-ternary coding + 2-Tap FFE + CTLE + feature extraction and classification for 44.7dB channel loss in 7.3pJ/bit," in Proc. IEEE Symp. VLSI Circuits, pp. 1-2, Jun. 2021. doi.org/10.23919/VLSICircuits52068.2021.9492510

[2] Y. Chun et al., "An ISI-resilient data encoding for equalizer-free wireline communication—Dicode encoding and error correction for 24.2-dB loss with 2.56 pJ/bit," IEEE Journal of Solid-State Circuits vol. 55, no. 3, pp. 567-579, Mar. 2020. doi.org/10.1109/JSSC.2019.2959487

[3] P. Upadhyaya *et al.*, "A fully adaptive 19-to-56Gb/s PAM-4 wireline transceiver with a configurable ADC in 16nm FinFET," in IEEE ISSCC Dig. Tech. Papers, pp. 108-110, Feb. 2018. doi.org/10.1109/ISSCC.2018.8310207

[4] L. Wang et al., "A 64Gb/s PAM-4 transceiver utilizing an adaptive threshold ADC in 16nm FinFET," in IEEE ISSCC Dig. Tech. Papers, pp. 110-112, Feb. 2018. doi.org/10.1109/ISSCC.2018.8310208

[5] B. Ye et al., "A 2.29pJ/b 112Gb/s wireline transceiver with RX 4-Tap FFE for medium-reach applications in 28nm CMOS," in IEEE ISSCC Dig. Tech. Papers, pp. 118-120, Feb. 2022. doi.org/10.1109/ISSCC42614.2022.9731591

[6] G. Gangasani et al., "A 1.6Tb/s chiplet over XSR-MCM channels using 113Gb/s PAM-4 transceiver with dynamic receiver-driven adaptation of TX-FFE and programmable roaming taps in 5nm CMOS," in IEEE ISSCC Dig. Tech. Papers, pp. 122-124, Feb. 2022. doi.org/10.1109/ISSCC42614.2022.9731636

[7] Y. -P. Lin et al., “A 2.16pJ/b 112Gb/s PAM-4 transceiver with time-interleaved 2b/3b ADCs and unbalanced baud-rate CDR for XSR applications in 28nm CMOS," in IEEE ISSCC Dig. Tech. Papers, pp. 136-138, Feb. 2024. doi.org/10.1109/ISSCC49657.2024.10454418