37.4 A 100Gb/s, 1.92pJ/b Aggregate 4-Lane Single-Ended NRZ Transceiver with 15dB Far-End Crosstalk Cancellation via On-Chip Feature Extraction and Classification in 16nm FinFET

Xiaohui Lin, Ramin Javadi, Bella Bose, Tejasvi Anand

Oregon State University, Corvallis, OR

Abstract

A low-power (1.92pJ/bit), high-speed (100Gb/s), equalizer-free, 4-lane single-ended transceiver that leverages machine learning principles for crosstalk compensation is proposed. The transceiver uses feature extraction and on-chip classification. It achieves more than 0.2UI horizontal opening for BER < 10^{-12}. The classification compensates the crosstalk from 3 aggressors and -15dB far-end crosstalk power sum. The classification only accounts for 7% of the total transceiver power.

Data-driven AI applications demand increasing bandwidth for processor-to-HBM or processor-to-module interconnects [1]. To meet this required higher throughput, multiple single-ended links are routed in close proximity to each other: resulting in electromagnetic coupling (crosstalk) between these densely packed interconnects. Crosstalk degrades signal integrity, posing a major challenge to achieving <10^{-12} BER. Prior crosstalk cancellation or reduction techniques, see Fig. 37.4.1, can be categorized into three categories: (1) wire-engineering methods, (2) bus-encoding schemes, and (3) active crosstalk cancellation using cross-coupled equalizers. Wire-engineering methods such as ground shielding [2] and wire swizzling [3] incur routing overhead at the PCB or package level; thus, limiting data density. Bus encoding, such as Fibonacci encoding [4], enhances signal integrity at the expense of coding complexity and additional pin overhead arising from redundant encoding. The active crosstalk cancellation (XTC) approach, using cross-coupled equalizers, adds a high-bandwidth cancellation circuit, operating at the data rate, into the transmitter or receiver front-end, resulting in increased loading and reduced energy efficiency. While prior work has shown cross-coupled equalizer-based crosstalk cancellation for just 2-lane high-speed (>20Gb/s) off-chip links [5], the energy inefficiency, due to the front-end loading of the XTC circuit and the long-distance routes of equalizer outputs across lanes, becomes more pronounced when crosstalk cancellation is required for four or more lanes; as is required for memory interfaces. This work proposes an equalizer-free approach that employs machine learning (ML) principles to compensate for crosstalk in a 4-lane off-chip wireline link. Although prior work adopted ML for wireline links, compensation was limited to single-channel imperfections: that is, inter symbol interference (ISI) [6-9]. This work investigates an ML approach to first learn the cross-channel coupling responses, revealing the inherently deterministic nature of the coupling effects that exist between channels, and then compensates for the crosstalk. Figure 37.4.1 presents the proposed approach of employing ML principles to compensate for crosstalk, where the receiver extracts the features of the received data from the victim and aggressor lanes and provides it to the classifier. The decision-tree classifier, first trained in supervised learning mode using PRBS data patterns, is then synthesized on-chip for classification. Classification is performed in the digital domain at the receiver back-end (low frequency: 781.25MHz), which results in energy-efficient crosstalk compensation; the classifier has a latency of only 6UI. The proposed 4-lane 25Gb/s single-ended transceiver, with an aggregate data rate of 100Gb/s, compensates for far-end crosstalk (FEXT) with a power sum up to 15dB; achieving a <10^{-12} BER with an energy efficiency of 1.92pJ/b, including clock distribution power, in a 16nm FinFET process. While the proposed transceiver is a 4-lane link, when compared to prior high-speed 2-lane links operating at a similar or higher data rate and when compensating similar FEXT, the proposed transceiver achieves a 31% lower energy/bit compared to [10] and 34% compared to [11].

Figure 37.4.2 shows the proposed 4-lane transceiver architecture, where the conventional receiver front-end crosstalk cancellation is replaced with feature extraction and classification in the receiver. All 4 lanes of the transmitter are identical, each with its own programmable PRBS generator. The PRBS generator can be initialized with different seeds to generate un-correlated data across 4-lanes, which provides a more representative distribution of coupled data transition events. The PRBS generator is followed by a 32:1 serializer and then a source-series-terminated (SST) output driver. Each transmitter lane has its own duty cycle corrector (DCC) and a programmable-delay line to de-skew phase mismatch among channels. The receiver is a half-rate architecture: its front-end consists of a T-coil that compensates for ESD and termination parasitics, followed by a feature extraction block, which consists of 2×3 slicers (even and odd, with each capturing three features) per lane. Feature extraction is followed by six 2:32 de-serializers per lane to reduce the processing data rate for the combined decision-tree classifier for all 4 lanes. Each lane supports 5 decision trees, which are trained on different channels; a preprogrammed classifier is selected, depending on channel conditions, via the SELECT signal. There are no equalizers such as FFE or CTLE present in the proposed transceiver as the ISI of channel will also be learned and compensated by the decision-tree classifier along with crosstalk.

Figure 37.4.3 illustrates the supervised learning process for all four-lane decision-tree classifiers. The training setup is modeled in MATLAB, where four PRBS-23 sources with various starting seeds drive four parallel lanes with practical channel models that include both insertion loss and inter-lane crosstalk. At the receiver front-end, each lane contains a feature extraction block with three slicers. These slicers compare the sampled incoming waveform against three threshold voltages (V_{TH1}, V_{TH2} and V_{TH3}) to extract first-order features (F1, F2 and F3). Received signal temporal information from both the victim and aggressor lanes is captured by recording several past UIs, the present UI, and a few future UIs forming a feature vector. The feature vectors from all four lanes are merged into a composite vector that represents both the channel impairments on each signal and the crosstalk coupling between them. This combined feature vector is then provided to the classifier for training. The original transmitted 4-lane PRBS data serves as labels for supervised learning. The classifiers build a hierarchical decision structure that maps the received information to the transmitted bits (labels), learning the per-lane channel impairments and crosstalk across all lanes. The trained decision-tree classifiers are then evaluated by analyzing the accuracy using a PRBS checker. System parameters, including the number of past/future UIs considered and the threshold voltages used for feature extraction, are iteratively adjusted through exhaustive search based on the measured accuracy. The classifiers are subsequently retrained with the updated parameters, and this loop continues until the accuracy converges to its highest achieved value among all the training conditions; resulting in an optimized decision-tree configuration for the given channel and crosstalk conditions. Finally, the classifiers generated in MATLAB are converted into Verilog code for on-chip synthesis. The test channel used shows that the training loop converges to a solution using 3 past and 2 future UIs from both the victim and aggressor lanes. To ensure a loose fit, the trained decision-tree model was evaluated across varying data rates, data patterns and channel models before being finalized, and therefore enabling it to compensate for a range of crosstalk.

Figure 37.4.4 depicts the classifier operation when compensating for the crosstalk. For simplicity, the decision-tree classifier is trained on a two-lane transceiver scenario. Both lanes experience mutual crosstalk, but in this example lane 1 is designated as the victim and lane 2 as the aggressor. Figure 37.4.4 shows the trained lane-1 decision tree. For each UI, three features, based on the three threshold voltages, are extracted from the victim at the current (F_{V1}, F_{V2} and F_{V3}) and past UI (F_{V4}, F_{V5} and F_{V6}), as well as from the aggressor at current and past UIs (F_{A1} - F_{A6}); these features are combined and provided to the decision tree. Consider the time instance (highlighted in yellow) when the victim's 0 is corrupted by crosstalk from the aggressor's transition to 1. The decision tree begins classification at the root node using feature F_{V1}, representing the lowest threshold voltage of the victim lane. With F_{V1}=1, it checks the middle-threshold feature F_{V2}. Since F_{V2}=1, the classifier examines the aggressor using F_{A1} and F_{A6}, corresponding to the current and past UI. With F_{A1}=1 and F_{A6}=0, a 0→1 transition is detected on the aggressor. Then, the classifier checks F_{V6} and F_{V5} to extract the victim's previous bit. With both being 0 and with F_{A3}=1, the classifier determines that the aggressor transitioned and no change occurred on the victim. The decision tree outputs 0, which represents its final decision based on the decisions taken, as previously described. This example shows that the transmitted bit, 0, on the victim lane would have been detected as a 1, due to crosstalk, if only the midpoint (V_{TH2}) sample was used. By incorporating current- and past-UI features from both lanes, the decision tree accurately detects the aggressor transition and outputs the correct decision. In this example, the decision tree uses 12 features and has 9 leaves, but for a more realistic channel like the one used to validate this work, the decision tree leverages 72 features (3×4×6) and has more than a hundred leaves.

A prototype 4-lane single-ended transceiver, using the proposed ML approach, is fabricated in a 16nm FinFET process. Figure 37.4.5 shows the measured Tx near-end eye diagrams for the four lanes at 25Gb/s/pin with PRBS7 data. Note that to measure the Tx outputs on the oscilloscope, a different and longer PCB channel was used due to the spacing required for the SMA connectors; the ISI for this PCB channel is characterized using a single-bit pulse response, shown in Fig. 37.4.5(top left). To validate that the decision-tree has learned the channel ISI and is now able to compensate for it, bathtub plots for all 4 transceiver (Tx & Rx) lanes were measured by turning on one lane at a time with and without the classifier. Measurements show that a >0.2UI improvement in the horizontal eye opening is achieved at <10^{-12} BER when the classifier is enabled. Each lane's measured channel response has a ~8dB insertion loss. The results presented show that the classifier compensates for the channel's ISI non-idealities as well. Figure 37.4.6 demonstrates the resilience of the proposed ML approach from crosstalk by enabling 1, 2 and all aggressor lanes: the crosstalk

 979-8-3315-8936-3/26/$31.00

interference translates to a power sum FEXT between -23.5 to -15dB for 1 to 3 aggressors. Notably, with classification OFF, the victim lane is unable to achieve a 10^{-12} BER in the presence of one aggressor and the received data was full of errors in the presence of two and more aggressors. With the classifier enabled a ~0.1UI eye opening is achieved at 10^{-12} BER for 1 aggressor, and a ~0.05UI opening for 3 aggressors. Figure 37.4.7 presents the performance comparison with prior work. The proposed equalizer-free transceiver achieves an energy efficiency of 1.92pJ/b operating at a 100Gb/s aggregate data rate. The decision-tree based 4-lane classifiers consume 13.6mW: accounting for 7% of the total transceiver power. The die micrographs for the Tx and Rx, and the power breakdown chart is shown in Fig. 37.4.7.

Acknowledgement:
This work was supported in part by the Center for Ubiquitous Connectivity (CUbiC), by the Semiconductor Research Corporation (SRC), and the Defense Advanced Research Projects Agency (DARPA) under the JUMP 2.0 program. Additional support in part by CDADIC. We would like to thank Intel for chip fabrication, and Tektronix for the measurement equipment.

References:
[1] A. Gholami et al., “AI and Memory Wall,” *IEEE Micro*, vol. 44, no. 3, pp. 33-39, May-June 2024. https://doi.org/10.1109/MM.2024.3373763
[2] J. Zhang and E. G. Friedman, “Effect of shield insertion on reducing crosstalk noise between coupled interconnects,” *IEEE ISCAS*, pp. 529-532, 2004. https://doi.org/10.1109/ISCAS.2004.1329325
[3] P. Gupta and A. B. Kahng, “Wire swizzling to reduce delay uncertainty due to capacitive coupling,” *Intl. Conf. on VLSI Design*, pp. 431-436, 2004. https://doi.org/10.1109/ICVD.2004.1260960
[4] Q. Liu et al., “A 0.90-Tb/s/in 1.29-pJ/b Wireline Transceiver With Single-Ended Crosstalk Cancellation Coding Scheme for High-Density Interconnects,” *IEEE JSSC*, vol. 58, no. 8, pp. 2326-2336, Aug. 2023. https://doi.org/10.1109/JSSC.2023.3261125
[5] L. Zhong et al., “A 2×56 Gb/s 0.78-pJ/b PAM-4 Crosstalk Cancellation Receiver with Active Crosstalk Extraction Technique in 28-nm CMOS,” *IEEE JSSC*, vol. 59, no. 9, pp. 3008-3020, Sept. 2024. https://doi.org/10.1109/JSSC.2024.3387355
[6] Y. Chun and T. Anand, “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 JSSC, vol. 55, no. 3, pp. 567-579, Mar. 2020. https://doi.org/10.1109/JSSC.2019.2959487
[7] 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,” *IEEE Symp. VLSI Circuits*, 2021. https://doi.org/10.23919/VLSICircuits52068.2021.9492510
[8] R. Javadi and T. Anand, “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,” *IEEE CICC*, 2025. https://doi.org/10.1109/CICC63670.2025.10982737
[9] R. Javadi et al., “A 3.2pJ/b 0.068pJ/b/dB 25Gb/s NRZ Wireline Transceiver with 3-Tap FFE and Random Forest Classification for Compensating 47dB Loss in 16nm FinFET,” *IEEE Symp. VLSI Tech. & Circuits*, 2025. https://doi.org/10.23919/VLSITechnologyandCir65189.2025.11075116
[10] L. Zhong et al., “A 112Gb/s/pin Single-Ended Crosstalk-Cancellation Transceiver with 31dB Loss Compensation in 28nm CMOS,” *ISSCC*, pp. 134-136, 2024. https://doi.org/10.1109/ISSCC49657.2024.10454508
[11] H. Wu et al., “A 2 × 24 Gb/s Single-Ended Transceiver With Channel-Independent Encoder-Based Crosstalk Cancellation in 28-nm CMOS,” *IEEE JSSC*, vol. 60, no. 8, pp. 2959-2970, Aug. 2025. https://doi.org/10.1109/JSSC.2024.3507855
[12] H.-G. Ko et al. “An 8Gb/s/μm FFE-Combined Crosstalk-Cancellation Scheme for HBM on Silicon Interposer with 3D-Staggered Channels,” *ISSCC*, pp. 128-130, 2020. https://doi.org/10.1109/ISSCC19947.2020.9063162
[13] Y. Nishi et al., “A 0.297-pJ/Bit 50.4-Gb/s/Wire Inverter-Based Short-Reach Simultaneous Bi-Directional Transceiver for Die-to-Die Interface in 5-nm CMOS,” *IEEE JSSC*, vol. 58, no. 4, pp. 1062-1073, April 2023. https://doi.org/10.1109/JSSC.2022.3232024

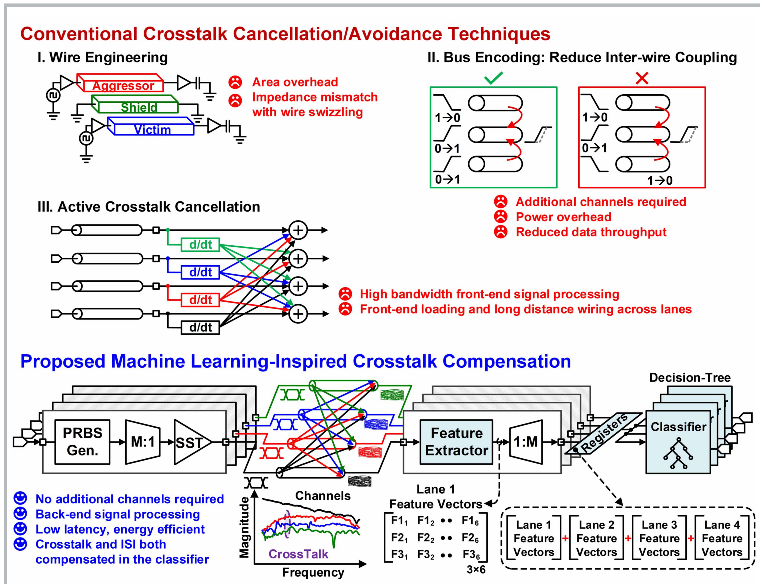

Figure 37.4.1: Conventional crosstalk compensation techniques and the proposed machine learning-inspired approach.

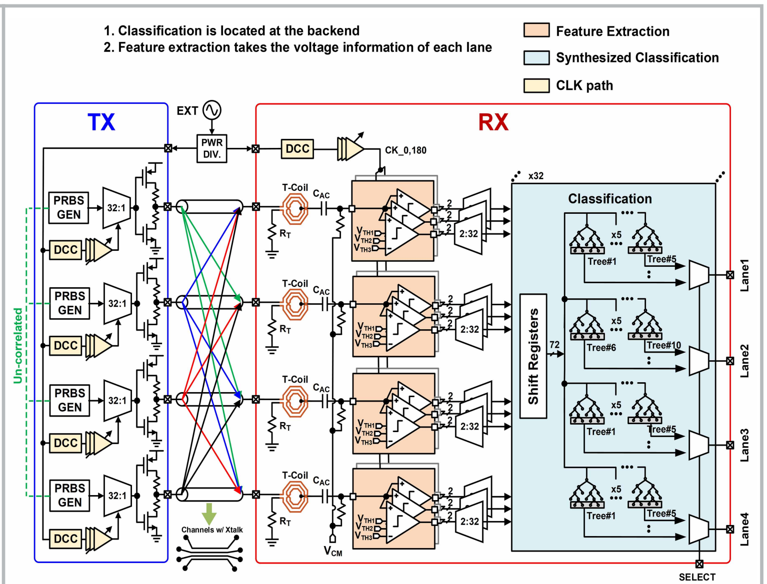

Figure 37.4.2: Proposed 4-lane ML-inspired equalizer-free transceiver architecture for crosstalk compensation.

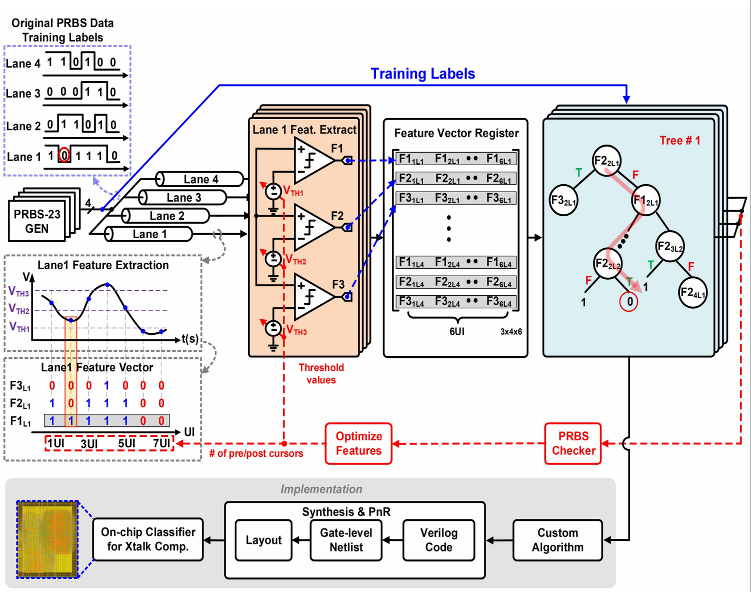

Figure 37.4.3: Decision-tree training process using supervised learning approach.

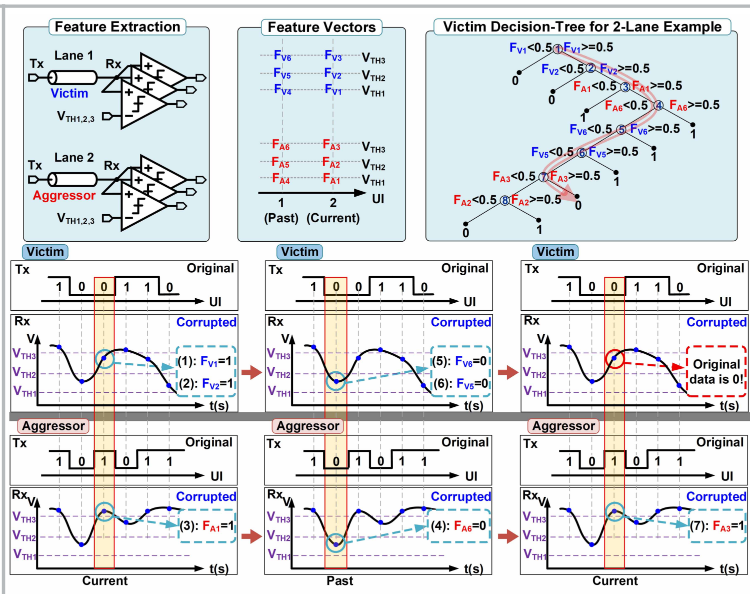

Figure 37.4.4: An example showing how crosstalk is compensated in a 2-lane coupled channel.

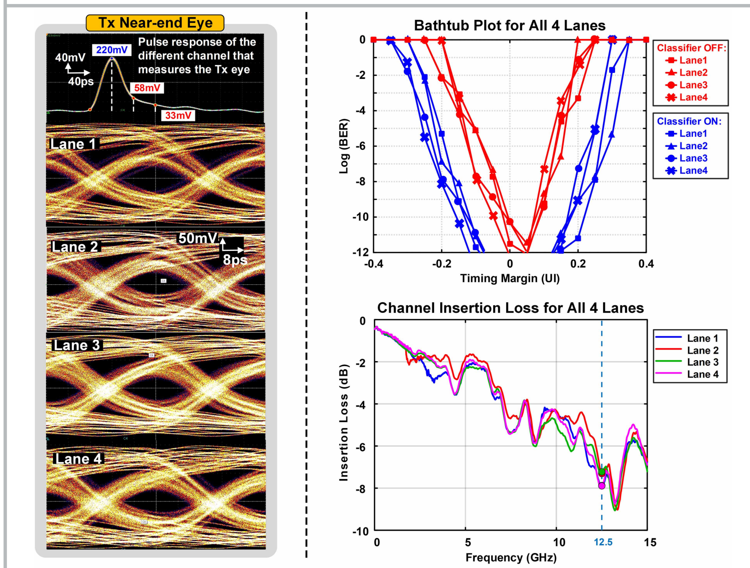

Figure 37.4.5: Measured Tx near-end eye diagram at 25Gb/s/pin, Tx + Rx bathtub plot for each lane without aggressors and with classifier ON/OFF at 25Gb/s/pin with PRBS7 data.

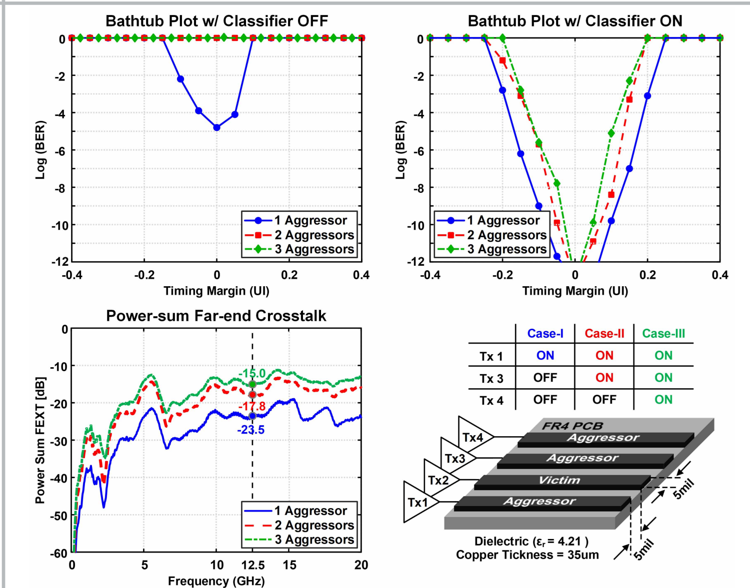

Figure 37.4.6: Measured bathtub plot with multiple aggressors with classifier ON/OFF, at 25Gb/s/pin with PRBS7 data, and 4-lane channel far-end crosstalk profile.

	Off-Chip Links				On-Chip Links	
	This work	Zhong ISSCC' 2024 [10]	Wu JSSC' 2024 [11]	Liu JSSC' 2023 [4]	Ko ISSCC' 2020 [12]	Nishi JSSC' 2023 [13]
Technology	16nm	28nm	28nm	28nm	65nm	5nm
Signaling	Single-ended	Single-ended	Single-ended	Single-ended	Single-ended	Single-ended
Modulation	NRZ	NRZ & PAM4	Orthogonal NRZ	NRZ	NRZ	NRZ SBD
# of Lanes	4	2	2	3+1	6	28
Crosstalk Compensation/ Cancellation	Machine Learning	CTLE 4-tap Tx FFE 4-tap Rx FFE	CTLE 3-tap FS-FFE	Fibonacci Coding	FFE-combined XTC	Ground Shielding
Data Rate [Gb/s/pin]	25	112	24	10	4	50.4
Loss at Nyquist [dB]	7.3	22	20	2	10.2	5.4
FEXT Power Sum at Nyquist [dB]	-15	-14	-10	-4.7	-28.9	-42.4
Power [mW]	70† 192	309.7	139	38.7	36.6	210
Energy Efficiency [pJ/b]	0.7† 1.92	2.77	2.9	1.29	1.5	0.297

†: Exclude clock distribution power from Tx and Rx

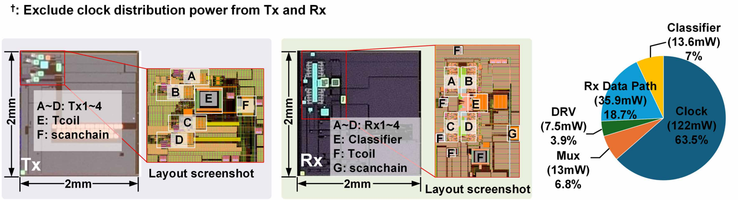

Figure 37.4.7: Comparison table, power breakdown, and die micrograph.

979-8-3315-8936-3/26/$31.00 ©2026 IEEE