

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

Ramin Javadi, Xiaohui Lin, and Tejasvi Anand

Oregon State University, Corvallis, OR, USA

Abstract

An energy efficient (0.068pJ/b/dB) NRZ wireline transceiver is presented that leverages 3-tap FFE in the transmitter and feature extraction with classification in the receiver for compensating 30dB to 47dB channel loss with $BER < 10^{-11}$. The proposed random forest classifier learns long-reach channel characteristics enabling the transceiver to achieve 3.2pJ/b at 25Gb/s for compensating 47dB loss.

Keywords: random forest, classification, wireline.

Introduction

The rapid growth of large language models (LLMs), like GPT-4 etc., is driving higher data rate demands, as massive amounts of information are communicated through wireline links between servers in data centers [1-3]. A majority of these links use copper wires and suffer from high insertion loss at high frequencies. The loss becomes significantly worse in long-reach applications, causing substantial inter-symbol interference (ISI). Therefore, numerous equalization taps must be implemented into the transmitter (Tx) and receiver (Rx) to compensate for the channel loss, which increases the power consumption [4-11]. Machine learning (ML) based transceivers are an effective solution that can reduce the number of equalization taps needed in communication links, thereby lowering power consumption. However, ML-inspired transceivers introduced by [12-14] are incompatible with present wireline standards such as non-return-to-zero (NRZ) as they require data encodings like Dicode and Hybrid-ternary schemes. These encoding schemes add an extra voltage level to the transmitted signal, which reduces the SNR by 3dB and requires a decoding block at the receiver. Furthermore, [12] utilizes second-order features and a decision-tree classifier, which results in higher Rx front-end power and limited channel loss compensation range, respectively. To address these limitations, an energy-efficient NRZ transceiver is proposed without data encoding that leverages first-order features plus classification in Rx and only 3-tap FFE in Tx to compensate for 47dB loss with an energy efficiency of 3.2pJ/b while maintaining $BER < 10^{-11}$. Additionally, a random forest classifier is introduced to learn the channel characteristics and recover dispersed NRZ signal across a wide range of channel loss (30dB to 47dB) with 10 unit interval (UI) latency. This on-chip classifier is a low-power feed-forward architecture without any feedback timing constraints that enables the proposed transceiver (16nm FinFET) to achieve a low energy/bit per channel loss of 0.068pJ/b/dB which is $\sim 1.4\times$ lower than prior-art [2] (5nm FinFET) for similar channel loss.

ML-Inspired Transceiver Concept & Training Process

A. Concept: Fig. 1 illustrates the concept of the proposed transceiver. The PRBS-generated data is de-emphasized by a 3-tap FFE and transmitted through a long-reach channel. The far-end eye is closed at the Rx due to channel ISI. Since the ISI is deterministic, a classifier that can compensate for a wide range of channel loss is trained and inserted at the Rx backend. The feature extraction block extracts channel attributes from the received signal and feeds them to the trained classifier. The classifier then maps the received information to the transmitted logic levels with low BER.

B. Training: Fig. 2 shows the off-chip training process of the random forest classifier in MATLAB. A PRBS generated data pattern passes through a practical channel model and is received by the feature extraction block in the Rx. This block includes three comparators, each generating 1-bit information based on threshold voltages (V_{th1} - V_{th3}) for each UI. Feature vectors, incorporating pre, main, and post-cursor information,

are formed by storing multiple UIs in a register block. Simultaneously, the transmitted UI bypasses the channel and serves as the label for training. The classifier builds a hierarchical decision structure to map received data to the corresponding label, learning the channel's non-idealities. The PRBS checker logic evaluates BER. Then, the number of pre/post UIs and thresholds are adjusted to minimize BER. Three decision-tree classifiers are trained with different channel models, and they are combined with a majority vote function to form one random forest classifier. This approach of classification helps to prevent overfitting and increases the channel loss compensation range. Fig. 3 depicts an example of implementing a simple tree using logic gates. Each node represents one of the comparator's outputs (feature) in a specific UI. Any branch from the root to a final leaf defines a particular condition that results in a decision bit and consists of inputs to an AND gate. The classifier is synthesized on-chip.

Transceiver Architecture

Fig. 4 shows the proposed transceiver where two clock phases are distributed to the Tx and Rx in half-rate architecture. The Tx consists of a 32-bit PRBS Gen., a 32:1 multiplexer, and a 3-tap FFE with SST output drivers. The Rx includes front-end amplifiers, 6 half-rate slicers to extract 3 features, three 2:16 de-multiplexers, registers, and 16 classifiers. The classifiers are positioned at the back-end, where they operate at a lower frequency, which results in lower power consumption. Fig. 5 shows the AFE including 4 amplifier stages with an inductive BW extension at the last stage. Sweeping the bias current source allows a variable gain from 6dB to 19dB at 12.5GHz.

Measurement Results

The chip is fabricated in 16nm FinFET and Fig. 6 shows Tx eye measurements. PRBS-7 data is transmitted via an SST driver at 0.9V. At 20Gb/s, the Tx far-end eye is open in ch1, showing FFE is effective for up to 20dB loss. However, at 25Gb/s the far-end eye is completely closed in ch2, indicating severe ISI impact. The Tx near-end output shows an open eye at 25Gb/s where the losses exhibited by the chip package and PCB traces are compensated by the FFE. Fig. 7 presents the transceiver measurement results. The channel losses measured at 8GHz and 12.5GHz are 30dB and 47dB, respectively (ch2). At data rates of 25Gb/s and 16Gb/s, the transceiver's bathtub plot shows horizontal openings of 0.05UI and 0.1UI for a $BER < 10^{-11}$. Overall, the transceiver compensates for 30dB to 47dB channel losses while maintaining a $BER < 10^{-11}$, demonstrating the effectiveness of random forest training in preventing overfitting. FFE coefficients were kept constant during the channel loss range measurements. Vertical margins for the three features are measured by sweeping the threshold voltages (V_{th1} - V_{th3}) (Fig. 8). A comparison with state-of-the-art (SOTA) designs is shown in Table I. When operating over a 47dB channel loss, the proposed transceiver consumes 80mW at 25Gb/s, achieving an energy efficiency of 3.2pJ/b. This power includes CLK distribution buffers, delay lines, and dividers. The classifier consumes only 9% of the total power, enabling this work to achieve a low energy/bit per channel loss of 0.068 pJ/b/dB, outperforming other SOTA designs (Fig. 9). The die photo is shown in Fig. 10.

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] B. Zhang *et al.*, *ISSCC*, 2023, pp. 5-7.

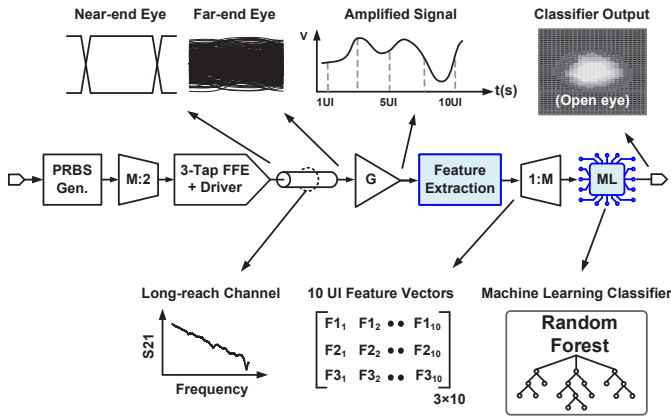


Fig. 1. Machine learning inspired wireline transceiver concept with feature extraction and random forest classification.

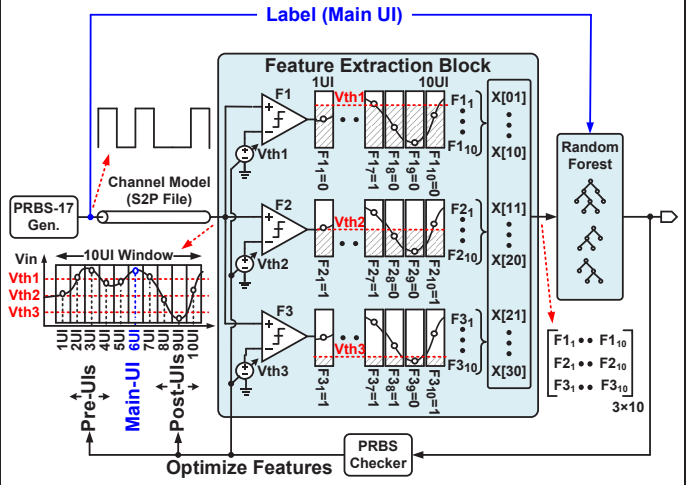


Fig. 2. Random Forest training process using supervised learning in MATLAB.

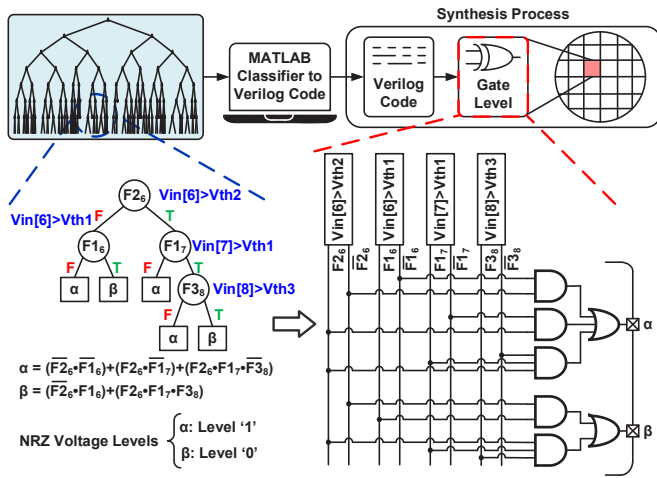


Fig. 3. A simplified example showing how a classifier is synthesized on silicon.

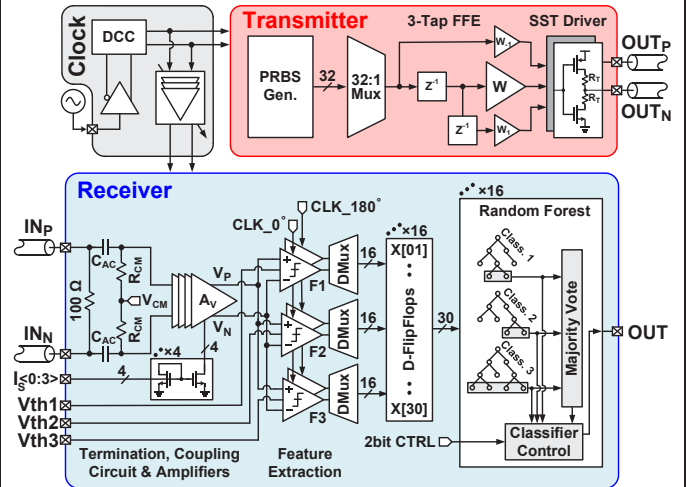


Fig. 4. Block diagram of the proposed transceiver with 3-tap FFE in Tx and feature extraction with random forest classification in Rx.

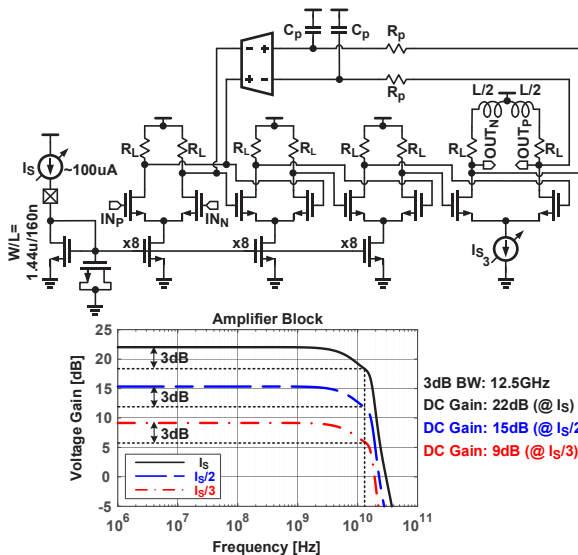


Fig. 5. Analog front-end with amplifier stages and simulated AC gain response with different bias current values.




Fig. 6. Transmitter's near-end and far-end at 20Gb/s and 25Gb/s with PRBS-7 data.

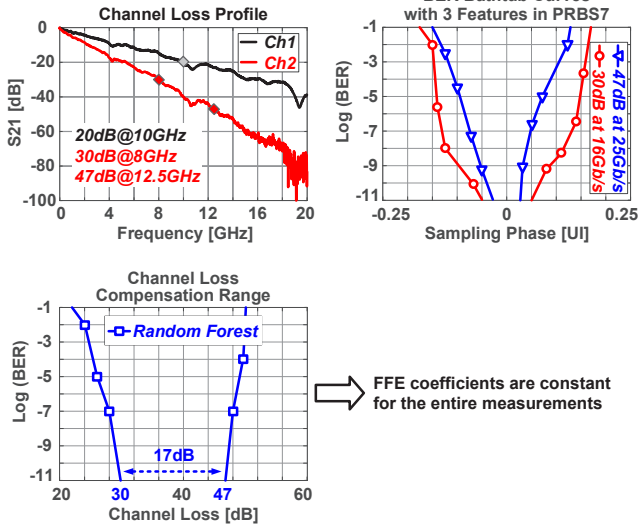


Fig. 7. Measured bathtub plot, channel loss profile, and channel compensation range with PRBS-7.

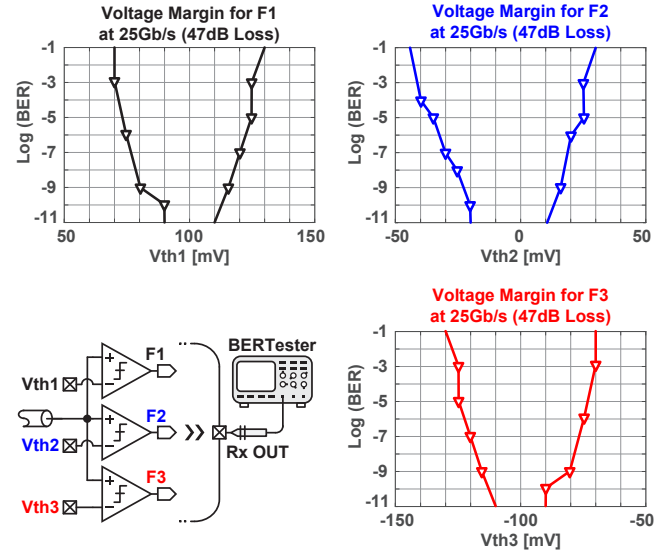


Fig. 8. Effects of feature's threshold voltage on BER for the proposed transceiver at 25Gb/s with PRBS-7.

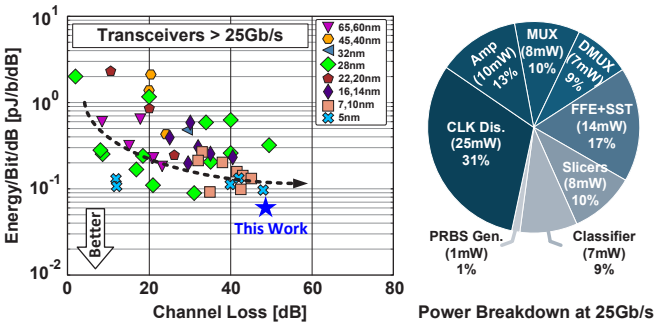


Fig. 9. Performance benchmark with prior arts and power breakdown. (The FoM chart includes links that include digital + analog power)

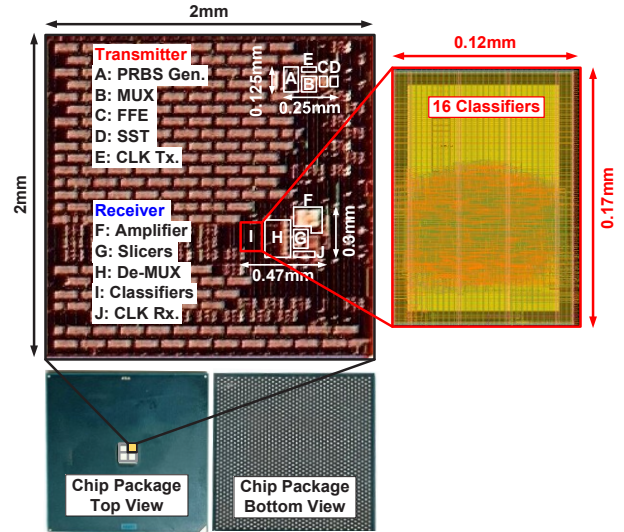


Fig. 10. Die photo, packaged chip photo and layout of 16 classifiers for 1:16 demultiplex output.

TABLE I. Comparison with state-of-the-art transceivers

	This Work	[12] VLSI'21	[13] JSSC'20	[1] ISSCC'23	[2] ISSCC'23	[4] ISSCC'21	[5] ISSCC'20
Technology	16-nm FinFET	65-nm CMOS	65-nm CMOS	7-nm FinFET	5-nm FinFET	16-nm FinFET	10-nm FinFET
Data Rate [Gb/s]	25	13.8	13.6	112	112.5	56	56
Encoding / Modulation	NRZ	HT	Dicode	PAM-4	PAM4	NRZ	PAM4
Bits per Symbol	1	1	1	2	2	1	2
Latency [UI]	10	19	5	-	-	-	-
Equalization	3-Tap FFE	CTLE 2-TapFFE	X	3-Tap FFE 18-TapDFE	CTLE, DFE 24-TapFFE	CTLE FFE	CTLE FFE, DFE
Log ₁₀ (BER)	<-11	<-6	<-12	<-6	-6	<-6	<-5
Tx+Rx Power [mW]	55	100.9	34.8	690	521	529	431.2
CLK Dis. Power [mW]	25	-	-	-	-	-	-
Loss@ Nyquist [dB]	47	44.7	24.2	43.9	48	40.4	38
Efficiency [pJ/b]	2.2†	3.2	7.3†	2.56†	6.16	4.63	9.44
FoM [pJ/b/dB]	0.046†	0.068	0.16†	0.106†	0.14	0.096	0.23
Area [mm ²]	0.172	0.236	0.167	0.63	0.461	0.996	0.72

† Excluding CLK distribution power to Tx and Rx in the chip.

References

- [1] H. Park *et al.*, *ISSCC*, 2023, pp. 5–7.
- [2] T. Ali *et al.*, *ISSCC*, 2020, pp. 118–120.
- [3] D. Xu *et al.*, *ISSCC*, 2021, pp. 134–136.
- [4] B. Yoo *et al.*, *ISSCC*, 2020, pp. 122–124.
- [5] A. Varzaghani *et al.*, *Symp. on VLSI*, 2022, pp. 26–27.
- [6] M. LaCroix *et al.*, *ISSCC*, 2021, pp. 132–134.
- [7] P. Mishra *et al.*, *ISSCC*, 2021, pp. 138–140.
- [8] R. Farjadrad *et al.*, *ISSCC*, 2021, pp. 194–196.
- [9] M. Jalali *et al.*, *ISSCC*, 2018, pp. 106–108.
- [10] Z. Guo *et al.*, *ISSCC*, 2022, pp. 116–118.
- [11] Z. Wang *et al.*, *Symp. on VLSI*, 2021, pp. 1–2.
- [12] Y. Chun *et al.*, *JSSC*, vol. 55 no. 3, pp. 567–579, 2020.
- [13] R. Javadi *et al.*, in press, *CICC*, 2025.