

8.9 A 72Gb/s/pin Single-Ended Simultaneous Bi-Directional Transceiver with C-Peaking Leakage Cancellation and Dual-Loop Hybrid Impedance Calibration for Chiplet Interfaces

Xuxu Cheng, Hongzhi Wu, Zhenghao Li, Weitao Wu, Xiongshi Luo, Yangyi Zhang, Quan Pan

Southern University of Science and Technology, Shenzhen, China

Abstract

This paper presents a 72Gb/s/pin single-ended simultaneously bi-directional (SBD) TRX in 28nm CMOS. Capacitive peaking leakage cancellation (CPLC) suppresses the high-frequency leakage due to the main and hybrid driver mismatch in SBD links by 63%. Dual-loop hybrid

impedance calibration reduces the coefficient errors of the hybrid circuits under PVT variations by 92%. The TRX achieves a 72Gb/s data rate with an eye-opening of 0.45UI and 243mV and an energy efficiency of 1.5pJ/b.

The explosive growth of artificial intelligence has accelerated the development of high-performance computing and high-bandwidth interconnects. Single-ended parallel links, like chiplet interfaces, are favored for their high bandwidth density and low bit-error rate (BER) [1-5]. Simultaneous bi-directional (SBD) links have shown potential to double the system throughput [6-7]. However, the design of high-speed bi-directional links is challenging. The top-left of Fig. 8.9.1 shows the single-ended SBD interconnect. Signals are simultaneously coupled to both sides of the channel and then decoupled by the hybrid circuits. As shown in the top-right of Fig. 8.9.1, the signal integrity of the SBD link is compromised by two main challenges: (1) high-frequency leakage due to the mismatch between main and hybrid drivers [8]; and (2) the residual outbound signal caused by imprecise coefficients of the hybrid circuits [6]. Consequently, these effects shrink the eye opening and limit the overall link performance. This work presents a 72Gb/s/pin single-ended SBD transceiver (TRX) with capacitive-peaking leakage cancellation (CPLC) to suppress the high-frequency leakage by 63%. A dual-loop hybrid impedance calibration scheme adaptively adjusts the hybrid coefficient, achieving accurate cancellation of the outbound signal with residual error less than 5mV.

The overall TRX architecture is shown in the bottom of Fig. 8.9.1. In the transmitter (TX), the pattern generator outputs 1/8-rate low-speed data, which is fed to an 8:4 multiplexer (MUX) to generate the outbound signal D4 and inversed outbound signal D4B. These two signals are serialized by 4:1 MUXs and distributed to the output stage by the slew rate (SR)-controlled pre-drivers. The output stage employs voltage mode (VM) drivers to maximize signal swing to improve the signal-to-noise ratio (SNR) of the SBD link. 24 main-driver slices and 6 hybrid-driver slices are connected through hybrid resistor R_H , with a CPLC driver to suppress the high-frequency leakage. The dual-loop hybrid impedance calibration technique calibrates the output impedance and hybrid impedance of drivers. The receiver (RX) receives the output signal of the hybrid circuit and amplifies it by the inverter-based analog front-end (AFE). To mitigate the channel inter-symbol interference (ISI), the RX adopts a single-ended track and hold feed-forward equalizer (TAH FFE) with feedthrough cancellation that eliminates the feedthrough error during the holding phase. The output of the TAH FFE is sampled by the quarter-rate slicer, deserialized by the 1:2 DEMUX, and finally fed into the pattern checker for BER measurement.

Figure 8.9.2 shows the conventional leakage cancellation scheme and the proposed CPLC scheme. The SBD signal fed to the hybrid is superimposed by the inbound and outbound signals. The relative timing of the inbound and outbound signal transition edges depend on the channel delay. In the worst case, the high-frequency leakage adds directly at the center of the inbound signal eye. The conventional leakage cancellation technique (top-left of Fig. 8.9.2) using the tunable LC delay line [8] aligns edges of the inbound and outbound signals, thereby limiting the residual leakage to the edge regions of the inbound signal to improved eye-opening. However, it introduces significant jitter to the inbound signal, and the implementation of the passive LCs occupy large area. FIR-based leakage cancellation in the middle-left of Fig. 8.9.2 utilizes digital signal processing (DSP) to detect the leakage and adjust cancellation coefficients to minimize residual leakage. However, since the leakage is primarily high-frequency noise, the FIR-based solution only eliminates the leakage partially. To suppress high-frequency leakage, a VM hybrid circuit with CPLC is proposed as shown in the top-right of Fig. 8.9.2. The CPLC driver extracts high-frequency components from the input signal to generate reversed leakage signal $V_{Leakage,B}$, which is then fed forward to the hybrid's output for the cancellation of $V_{Leakage}$. To enable precise elimination of leakage, the strength of pull-up and pull-down of CPLC driver is controlled by the 6b switchable header and footer. Additionally, a series inductor is adopted to broaden the bandwidth and reduce the leakage. The simulated waveforms of leakage are shown in the middle-right of Fig. 8.9.2. By using the series inductor and CLPC, the amplitude of the leakage signal is suppressed by 63% (from 148mV to 55mV). The bottom of Fig. 8.9.2 shows the simulated eye diagrams of the output signal of the hybrid circuit. The eye is completely closed without the series inductor and CPLC. The eye of the hybrid output signal is slightly opened with an eye height of 38mV and eye width of 0.23UI by introducing the series inductor. When the CPLC is further enabled, the eye-opening is improved to an eye height of 134mV and an eye width of 0.65UI. The eye-opening is improved by 6 times, ensuring sufficient SNR for SBD links.

The accuracy of the hybrid coefficients across PVT variations is critical for effective cancellation of the outbound signal. Thus, hybrid coefficient adaptation is required to improve the robustness of SBD links. In the VM hybrid circuit, the hybrid coefficients are determined by the output impedance ratio between the main driver and the hybrid driver, along with an appropriate hybrid resistor R_H value [6]. As shown in the top-left of Fig. 8.9.3, conventional impedance calibration uses an impedance control loop shared by the main and hybrid drivers to match their impedance to a precise off-chip resistor, thereby ensuring an accurate ratio between them. However, the fixed output impedance of drivers fails to compensate for the R_H variation, resulting in incomplete cancellation of the outbound signal. To address this issue, this work introduces a dual-loop hybrid impedance calibration scheme that calibrates the impedance of hybrid driver according to the impedance of R_H and main driver. As depicted in the top-right of Fig. 8.9.3, Loop#1 calibrates the pull-up (PU) driver and pull-down (PD) driver branches to a high-precision off-chip resistor, generating the bias voltage V_{BP} and V_{BN} for the replica main driver. In Loop#2, two operational amplifiers adjust the bias voltage of the hybrid driver based on the impedance of the replica main driver and R_H . This process forces the node V_{RX} to remain at VDD/2 for both 'high' and 'low' states of the outbound signal, thereby achieving cancellation of the outbound signal. The bottom-left of Fig. 8.9.3 shows the simulation results of the conventional impedance calibration scheme and the proposed dual-loop hybrid impedance calibration scheme when varying R_H . With the conventional scheme, the residual hybrid error is up to 60mV as R_H varies by 20%. In contrast, the proposed scheme ensures that the impedance of the hybrid driver tracks R_H variation, minimizing the error to below 5mV, reducing by 92% compared to the conventional scheme. Transient simulation waveforms shown in the bottom-middle of Fig. 8.9.3 further confirm that the proposed calibration scheme significantly reduces residual error compared to the conventional approach. The simulated output impedance is shown in the bottom-right of Fig. 8.9.3. The output impedance can be adjusted within a range from 38 to 63 Ω by configuring the number of active driver slices, thereby enabling impedance matching to the channel.

To mitigate ISI in the channel, a TAH FFE is employed in the RX side, as illustrated in the top of Fig. 8.9.4. The TAH FFE is implemented using an inverter-based design to reduce power consumption while achieving high data rate. A single-ended transmission gate is adopted for the track and hold switch due to its wide input and output common-mode range and strong immunity to clock feedthrough. However, during the holding phase, data feedthrough occurs if the input signal D_{in} changes, thereby affecting the holding signal. The feedthrough can be alleviated by using the differential signals, but with high power consumption. To address this issue, a feedthrough cancellation (FTC) tap is introduced in the RX TAH FFE. As shown in the timing diagram in the top-right of Fig. 8.9.4, taking the THIB path as an example, during the holding phase, the transition edge between data symbols D2 and D3 feeds through to THIB, which introduces an error on the holding value. When the post tap (THQ) is enabled, the timing margin for the FFE operation is 1.5UI; during this period, the error is further amplified. The FTC tap is introduced to eliminate the feedthrough by using the THQB signal. Simulation results are shown in the bottom of Fig. 8.9.4. The optimal clock timing is set for TAH operation. The eye-opening of the TAH signal is degraded by the ISI and feedthrough. With the post tap of RX FFE enabled only, the ISI is eliminated but the feedthrough remains. With the FTC tap further enabled, the feedthrough is eliminated, and the eye-height is improved from 233 to 285mV. The overshooting is outside the TAH FFE timing margin, and thus it has little impact on the link performance.

The proposed SBD TRX is fabricated in 28nm CMOS and evaluated using a PCB. Figure 8.9.5 shows the measurement results. In the uni-directional (UD) mode, the TX output eye exhibits a height of 189mV and a width of 20.7ps at 36Gb/s, which is 214mV and 25.7ps at 32Gb/s. In the bi-directional mode, the TRX is measured over an off-chip channel with an insertion loss of 6dB at 18GHz. The measured single-channel RX bathtub curve at data rate of 72Gb/s (36Gb/s in each direction) is shown in the top-middle of Fig. 8.9.5. The TRX reaches error-free margin of 0.2UI at 1e-12 BER when the CPLC and RX TAH FFE are turned off. The error free margin is extended to 0.39UI with the CPLC enabled; it is further extended to 0.45UI with the TAH FFE enabled. As shown in the top-right of Fig. 8.9.5, with both the CPLC and TAH FFE enabled, the RX internal eye height is improved from 131mV to 244mV. Measurement results of the SBD TRX at data rate of 64Gb/s are shown in the bottom-right of Fig. 8.9.5. With the CPLC and TAH FFE enabled, the error free margin at 1e-12 BER is

extended from 0.22UI to 0.52UI, and the eye height is increased from 158mV to 263mV. Figure 8.9.7 shows the die photo and power breakdown.

The TRX performance summary and comparison with the prior works are shown in Fig. 8.9.6. With the proposed CPLC, dual-loop hybrid impedance calibration and TAH FFE, the SBD TRX achieves the data rate of 72Gb/s/pin and the vertical eye-opening of 244mV, which compares favorably among prior works.

Acknowledgement:

The work is supported by the Science and Technology Plan of Shenzhen under Grant KJZD20231023100159002 and KJZD20240903100208012, and in part by the SUSTech High-Level Special Funds under Grant G03034K007. Corresponding author: Quan Pan.

References:

[1] K. Seong et al., "A 4nm 32Gb/s 8Tb/s/mm Die-to-Die Chiplet Using NRZ Single-Ended Transceiver With Equalization Schemes And Training Techniques," *ISSCC*, pp. 114–116, Feb. 2023. <http://doi.org/10.1109/ISSCC49657.2024.10454481>
 [2] K. Seong et al., "A 4nm 48Gb/s/wire Single-Ended NRZ Parallel Transceiver with Offset-Calibration and Equalization Schemes for Next-Generation Memory Interfaces and Chiplets," *ISSCC*, pp. 250–252, Feb. 2024. <http://doi.org/10.1109/ISSCC49657.2024.10454481>
 [3] Y. Nishi et al., "A 0.190-pJ/bit 25.2-Gb/s/wire Inverter-Based AC-Coupled Transceiver for Short-Reach Die-to-Die Interfaces in 5-nm CMOS," *IEEE JSSC*, vol. 59, no. 4, pp. 1146–1157, April 2024. <http://doi.org/10.1109/JSSC.2023.3338478>
 [4] D. T. Melek et al., "A 0.29pJ/b 5.27Tb/s/mm UCIe Advanced Package Link in 3nm FinFET with 2.5D CoWoS Packaging," *ISSCC*, pp. 590–592, Feb. 2025. <http://doi.org/10.1109/ISSCC49661.2025.10904754>
 [5] M. -S. Lin et al., "A 32Gb/s 10.5Tb/s/mm 0.6pJ/b UCIe-Compliant Low-Latency Interface in 3nm Featuring Matched-Delay for Dynamic Clock Gating," *ISSCC*, pp. 586–588, Feb. 2025. <http://doi.org/10.1109/ISSCC49661.2025.10904767>
 [6] 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, Apr. 2023. <http://doi.org/10.1109/JSSC.2022.3232024>
 [7] Z. Wang et al., "A 64Gb/s/wire 10.5Tb/s/mm/Layer Single-Ended Simultaneous Bi-Directional Transceiver with Echo and Crosstalk Cancellation for a Die-to-Die Interface in 28nm CMOS," *ISSCC*, pp. 588–590, Feb. 2025. <http://doi.org/10.1109/ISSCC49661.2025.10904631>
 [8] R. Farjadrad et al., "An Echo-Cancelling Front-End for 112Gb/s PAM-4 Simultaneous Bidirectional Signaling in 14nm CMOS," *ISSCC*, pp. 194–196, Feb. 2021. <http://doi.org/10.1109/ISSCC42613.2021.9365852>
 [9] Y.-H. Fan et al., "A 32-Gb/s Simultaneous Bidirectional Source-Synchronous Transceiver With Adaptive Echo Cancellation Techniques," *IEEE JSSC*, vol. 55, no. 2, pp. 439–451, Feb. 2020. <http://doi.org/10.1109/JSSC.2019.2956369>

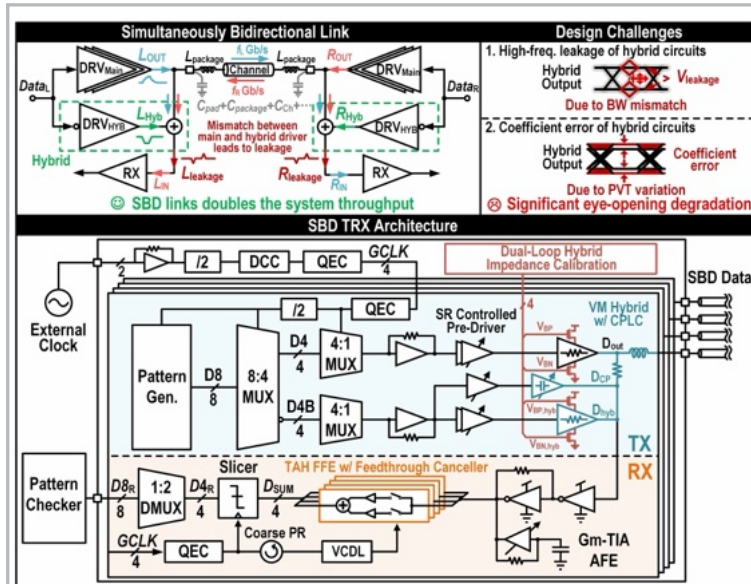


Figure 8.9.1: SBD link and its design challenges (top); SBD transceiver architecture (bottom).

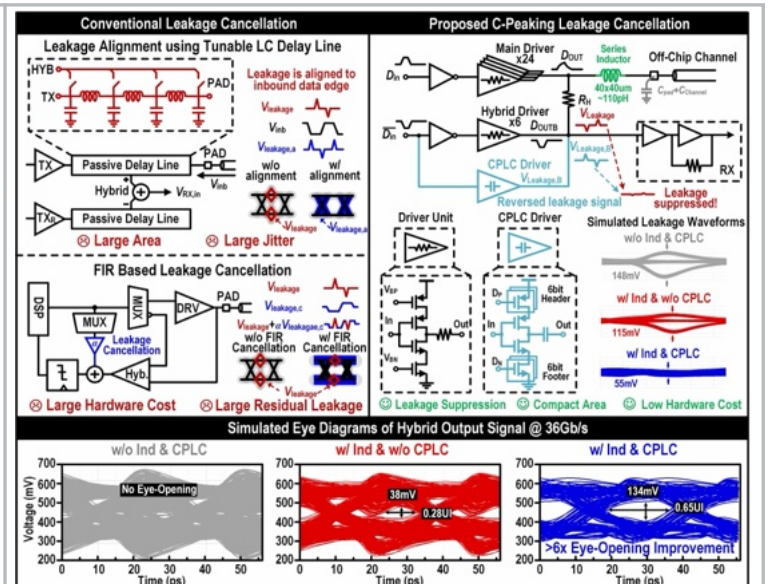


Figure 8.9.2: Conventional leakage cancellation (top-left) and proposed CPLC (top-right); simulated eye diagrams of the SBD links (bottom).

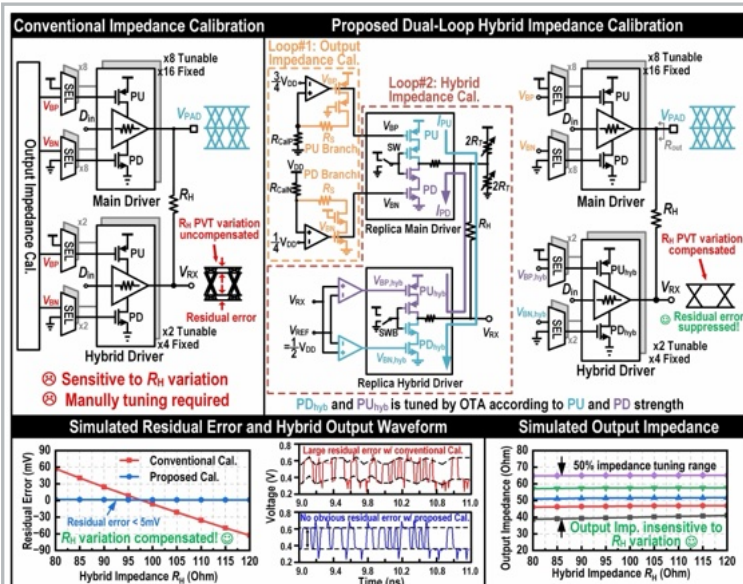


Figure 8.9.3: Conventional impedance calibration (top-left); proposed dual-loop hybrid impedance calibration scheme (top-right); simulation results (bottom).

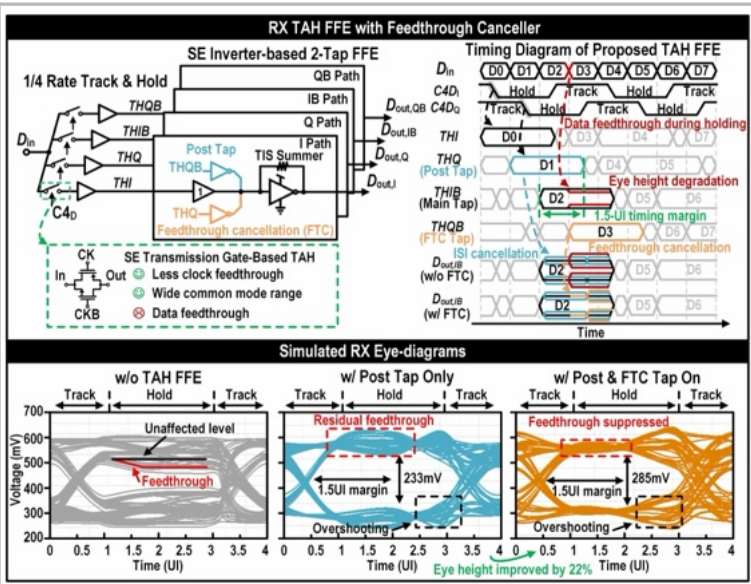


Figure 8.9.4: Proposed RX TAH FFE with feedthrough cancellation and corresponding timing diagram (top); simulation results of TAH FFE (bottom).

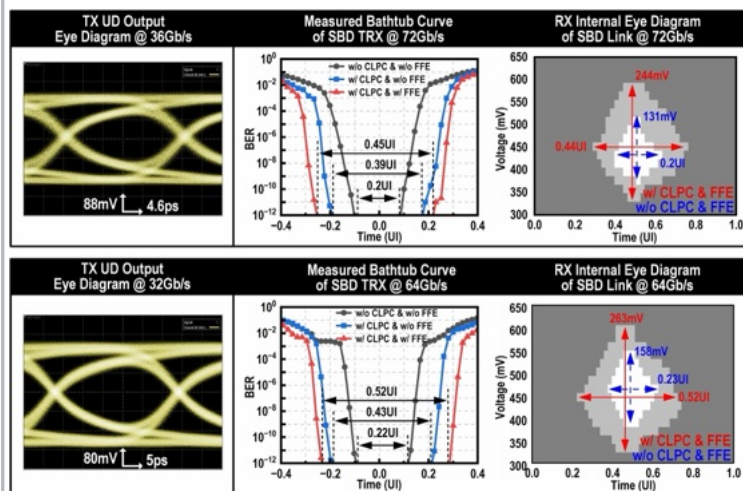


Figure 8.9.5: Measured TX output eye diagrams (left), RX bathtub curves (middle) and RX internal eye diagrams (right).

	JSSC'20 [9]	JSSC'23 [6]	ISSCC'23 [1]	ISSCC'24 [2]	JSSC'25 [3]	ISSCC'25 [7]	This Work
Process	28nm CMOS	5nm FinFET	4nm FinFET	4nm FinFET	3nm FinFET	28nm CMOS	28nm CMOS
Supply	0.9	0.75	0.9	0.9	0.75	0.9/0.45	0.9
Data Rate Per Pin (Gb/s/pin)	32	50.4	32	48	32	64	72
Signaling	SBD NRZ	SBD NRZ	NRZ	NRZ	NRZ	SBD NRZ	SBD NRZ
Energy Efficiency (pJ/bit)	1.83	0.30	0.44	0.67	0.36	1.21	1.50
Channel Loss @ f_{sys} (dB)	4.4	2	3.9	3	2.4	2.8	6
Channel Type	Off-Chip	On-Chip	Off-Chip	Off-Chip	Off-Chip	On-Chip	Off-Chip
Equalization	CTLE	No	DFE	C-Peaking	No	No	2-Tap RX FFE
Hybrid Impedance Cal.	No	No	-	-	-	-	Yes
Leakage Cancellation	No	No	No	No	No	No	C-Peaking
Eye Height (mV)	20	NA	170	88	NA	80	244 @ 72Gb/s 263 @ 64Gb/s
Eye Width (UI)	0.4	0.56	0.53	0.58	0.34	0.64	0.44 @ 72Gb/s 0.52 @ 64Gb/s
Core Area / Lane (mm ²)	0.182	0.004	0.004	0.067	0.34	0.019	0.012

Figure 8.9.6: Performance summary and comparison with prior works.

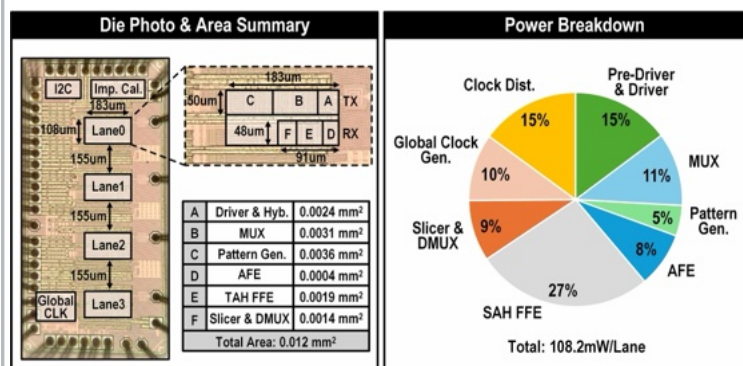


Figure 8.9.7: Die micrograph, area summary, and power breakdown.