A True Full-Duplex IO (TFD-IO) with Background SI cancellation for High-Density Interfaces

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY-NC-SA 4.0

SUBMISSION DATE / POSTED DATE

25-08-2021 / 27-08-2021

CITATION

Goyal, Sandeep; Gupta, Shalabh; Parulekar, Ganpat Anant (2021): A True Full-Duplex IO (TFD-IO) with Background SI cancellation for High-Density Interfaces. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.16438071.v1>

DOI

[10.36227/techrxiv.16438071.v1](https://doi.org/10.36227/techrxiv.16438071.v1)

A True Full-Duplex IO (TFD-IO) with Background SI Cancellation for High-Density Interfaces

Sandeep Goyal, Ganpat Parulekar, and Shalabh Gupta

Abstract—In this work, we have proposed and experimentally demonstrated a true full-duplex IO (TFD-IO) for high-speed high-density interfaces. The proposed TFD-IO can be used as an independent module that converts a unidirectional IO/interconnect to a fully bidirectional IO/interconnect, to ideally double the throughput of the high-speed interface. The TFD-IO uses a correlation-based technique to cancel the self-interference (SI) and echoes adaptively in the background. The signals transmitted from the near-end and the far-end can use independent baud-rates and signaling schemes in a TFD-IO. A proof-of-concept design of the TFD-IO module has been fabricated in a 65 nm CMOS technology, and demonstrated with bidirectional throughputs of up to 12.8 Gb/s.

Index Terms—Adaptive echo cancellation (EC), chiplets, full-Duplex, high-speed serial links, heterogeneous integration, source-synchronous transceiver, simultaneous bidirectional (SBD).

I. INTRODUCTION

Recent growth in the requirements and capabilities of algorithms for machine learning and artificial intelligence based user applications has led to a huge demand for improved performance of computing systems. These systems need to store, process, and exchange large amounts of data to meet the growing demands. In these systems, data interfaces, such as processor-to-processor, processor-to-memories, and processor-to-network interfaces, may have to handle data throughputs of several terabits/second and interface densities of greater than terabits/second/mm [1], [2]. Throughputs of these high-speed interfaces are limited by interconnect bandwidths and chip packaging technologies [3], [4]. Multi-chip modules (MCM), system-in-package (SiP), and multi-chiplet based platforms with die-to-die interfaces have emerged as solutions to overcome some of the limitations [4]–[7]. The die-to-die interfaces need to support high data/pin efficiencies, low channel losses, and small form-factors [4], [5], [8].

The interface with a large number of lanes, although has low channel losses but suffers from the issue of crosstalk between them, which further deteriorates the performance with an increase in the lane density or baud-rate [9]–[11]. Reduction in the number of lanes is possible by using simultaneous bidirectional (SBD) or full-duplex (FD) communication between the high-performance modules. FD links use transceivers at the two ends, which are named as far-end and near-end. The strength of the near-end transmitted signal at the input-output (IO) port in the full-duplex link is much higher than that of the signal received from the far-end. Thus, the near-end

Fig. 1: Conventional bidirectional links: (a) without interference cancellation and (b) with the hybrid used for suppressing the SI. Various constituents of near-end interference, i.e., SI and echoes, are shown using dashed lines. Here, Tx - transmitter; Rx : receiver; D_T : data to be transmitted; SI : self-interference; S_{TR} : composite signal at transceiver IO port; S_{RI} : signal received with SI and echoes; and S_{FR} : recovered far-end signal.

transmitter causes a severe self-interference (SI) at the near-end receiver [11], [12]. In addition, there can be echoes of the transmitted signal at the receiver due to a mismatch between the source and load impedances, and back reflections from the channel. Therefore, transceivers in these links receive a composite signal S_{RI} , which consists of interference from the signal transmitted by the near-end transceiver along with the desired far-end signal, as shown in Fig. 1(a).

The interference hinders recovery of the far-end information at the near-end receiver. As a result, along with a data transmission mechanism, the transceivers also need a method to suppress a significant part of the SI at the receiver front-end (RFE) itself, before the received signal can be further processed to recover the data transmitted from the far-end. For example, as shown in Fig. 1(b), an analog front-end (AFE) uses a hybrid to suppress the SI to residual-SI, which may be cancelled in the digital domain along with the echoes.

Full-duplex links for Ethernet standards, such as IEEE 802.3ak and 802.3ap, typically use passive hybrids and auxiliary digital-to-analog-converters (DACs) to suppress the interference in the analog domain [13]–[15], before the received signal is further processed for data recovery. The passive hybrids are magnetic, bulky, and band-limiting in nature.

Fig. 2: Bidirectional chip-to-chip data transfer interface with the proposed true full-duplex IOs (TFD-IO)s.

They support suppression only at low baud-rates and have larger form factors, making them unattractive for high-density and high-speed interconnects. To achieve higher throughputs, recent high-speed FD interconnects use hybrids with active circuits for interference suppression [11], [12], [16]–[20]. These hybrids effectively consist of one of the following: (i) a scaled replica generator with a subtractor [17]; (ii) a comparator with dynamic referencing [18], [21], [22]; (iii) a resistive or a capacitive bridge [11], [19]; (iv) a resistor-transconductor (R-gm) cell [12], [16]; and (v) a directional inverter/buffer (DIB) with weighted cancellation paths [23]. The main limitation in these schemes is that the AFEs do not account for the delay spread of the transmitted pulses that causes SI over multiple bit periods. These hybrids may also require manual tuning or foreground calibration to support data transfer over a wide range of interface conditions. Furthermore, most of these techniques require timing synchronization between the two transceivers, to ensure that the interference is suppressed/cancelled sufficiently, as the amount of cancellation depends on the receiver sampling phase [12], [18], [24], [25].

To overcome the above mentioned limitations, digital techniques have to be implemented, which can support interference cancellation adaptively [17]. However, implementing interference cancellation using digital approaches need a high resolution analog-to-digital converter, thus is often power and area inefficient [13], [17]. Furthermore, a significant training time may be required for adaption of the canceller’s weight coefficients using the digital techniques.

In this paper, we propose a true full-duplex IO (TFD-IO) module that supports simultaneous bidirectional data transmission over an electrical interconnect. The TFD-IO uses a correlation-based SI-canceller, whose weight coefficients are adjusted in the background, adaptively, using a modulation format and data-rate independent technique [26], [27]. Figure 2 shows a conceptual diagram of a high-density chip-to-chip interface, wherein, the TFD-IOs can be used as plug-n-play modules to double the data efficiency per lane and per IO port of the chip. The use of the bidirectional ports, i.e., the TFD-IOs, can also simplify board level design while interfacing two chips on a PCB with a large number of high-bandwidth interconnects between them.

In the following section, we review the correlation-based interference cancellation technique. In Section III, we present the proposed TFD-IO and its implementation details. Section IV presents the measurement setup and results obtained for the TFD-IO cells, and Section V concludes this paper.

Fig. 3: System model for a correlation-based interference canceller. Here, RFE: receiver front-end; D_T : Data to be transmitted; $D[k]$: k^{th} delayed tap D_T ; $T[k]$: k^{th} weighted tap; c_k : k^{th} weight coefficient; WA_k : k^{th} weighting amplifier (WA); S_{TR} : composite signal at transceiver IO port; S_F : far-end signal; and S_{NI} : near-end interference.

II. CORRELATION-BASED SI CANCELLATION

The near-end transceiver in a FD link receives a composite signal S_{RI} , which contains: (i) the far-end signal S_F ; and (ii) S_{NI} - includes the SI and echoes. The major challenge in transferring data over FD links is to extract the desired far-end information from S_{RI} . This extraction becomes further challenging when the near-end transceiver should suppress the interference unconscious of the far-end transceiver. The used correlation-based cancellation approach allows for such an interference suppression, independent of the operating state of the far-end transceiver. This approach uses a practical assumption that the data signals transmitted by two independent data sources are typically uncorrelated. For such uncorrelated signals, this cancellation approach allows the SI-canceller to suppress the interference by adapting its coefficients in the background. This approach thus eliminates the need of training the SI-canceller in foreground, which is required to track the variations in data transmission path.

Figure 3 illustrates the model of a transceiver with a correlation-based cancellation loop used for suppressing the near-end interference. The cancellation approach uses a simple fact that the two signals, D_T and S_{NI} , have the same origin, and hence are correlated. In this approach, to suppress S_{NI} , the cancellation loop minimises only the correlation present between output of the canceller S_{FR} (the recovered far-end signal) and $D[1:k]$ – a set of delayed taps (copies) of D_T . These delayed taps are required to mimic the delay spreads and magnitude responses of the interference path. In addition, the use of multiple delayed taps also enables cancellation for the echoes. The cancellation loop, thus extracts the uncorrelated far-end information from the S_{RI} . For interference cancellation, the delayed taps $D[1:k]$ are first weighted with the corresponding weight coefficients $c_{[1:k]}$, and are then summed together to generate a replica cancellation signal S_C . Next, S_C is subtracted from S_{RI} to minimize S_{NI} , as given by Eq. (1).

$$\begin{aligned} S_{FR} &= S_{RI} - S_C, \\ &= S_F + S_{NI} - \sum_{p=1}^k c_p \times D[p], \end{aligned} \quad (1)$$

Fig. 4: Proposed true full-duplex IO with a correlation-based SI-canceller. Here, DCOC: DC offset compensator and Clk_{in} : input clock to D-latches.

where, $D[p]$ denotes the p^{th} delayed tap of D_T , and c_p is the weight coefficient (gain) of weighting amplifier WA_p . To extract the far-end information by suppressing the interference, i.e., to make S_{FR} free of S_{NI} or of any interfering component due to the delayed taps $D[1:k]$, Eq. (2) must be satisfied.

$$S_{NI} = \sum_{p=1}^k G_{\text{Loop}} \times R_p \times D[p], \quad (2)$$

where, R_p denotes the correlation present between S_{FR} and $D[p]$, and G_{Loop} denotes the loop gain of the cancellation loop. Let us consider for a delayed tap $D[p]$, ε_p is the desired tap weight to suppress the amount of interference present in the S_{RI} . Therefore, for the cancellation, ε_p is the required value for c_p , and it can be determined as per the following Equation,

$$\varepsilon_p \times D[p] = G_{\text{Loop}} \times R_p \times D[p], \quad (3)$$

where, R_p for a system in its linear range of operation can be further divided into three parts as given by,

$$R_p = R_{PI} + R_{PF} + R_{PC}, \quad (4)$$

where, R_{PI} denotes the correlation between $D[p]$ and S_{NI} , R_{PF} denotes the correlation between $D[p]$ and S_F , and R_{PC} denotes the correlation between $D[p]$ and S_C . For interference cancellation, the used correlation-based approach supports minimisation of these correlation components as follows:

- R_{PF} : should be equal to zero as the two transmitted signals, i.e., the near-end signal and the far-end signal, should be practically uncorrelated.
- R_{PC} : should be low as the negative feedback of the loop minimises the correlation between S_{FR} and S_C .
- R_{PI} : the required correlation for minimising the SI present in the S_{FR} corresponding to $D[p]$ is equal to $\varepsilon_p / G_{\text{Loop}}$.

The R_{PI} acts as a control input for the negative feedback SI-cancellation loop. Additionally, R_{PI} is also the residual-SI present in S_{FR} , and its value depends upon the G_{Loop} . Therefore, G_{Loop} should be large for reducing the residual-SI, and to faithfully recover the far-end information from S_{RI} .

Using this correlation-based cancellation approach, this work proposes a TFD-IO cell that supports bidirectional data transfer over an electrical interconnect. The use of the TFD-IO cell supports background suppression of the SI and near-end echoes using a correlation-based SI-canceller.

III. PROPOSED TRUE-FULL DUPLEX IOS

The proposed TFD-IO, as shown in Fig.4, consists of a line-driver (LD) for signal transmission, a buffer as the RFE for signal reception, and the correlation-based SI-canceller. The canceller supports interference cancellation spanning over multiple bits of the near-end transmitted data. The correlative approach allows the canceller to adapt in the background for suppressing the time-variations of interference.

Figure 4 shows the architecture of correlation-based SI-canceller. The SI-canceller is used to cancel S_{NI} present in S_{RI} , which includes the SI received from LD, and the immediate echoes/reflections from the channel-IO interface. The canceller uses the correlation-based cancellation algorithm to decide its weight coefficients $c_{[1:k]}$ for the delayed taps $D[1:k]$. The SI-canceller generates the replica tap S_C , as a sum of the weighted taps $T[1:k]$ to suppress the interference, as given by Eq. (5),

$$S_C = \sum_{p=1}^k T[p] = \sum_{p=1}^k c_p \times D[p]. \quad (5)$$

The cancellation loop uses the coefficients $c_{[1:k]}$ to minimize the correlation between S_{FR} and $D[1:k]$, to reduce the interference. As the canceller minimises only the correlation, by subtracting S_C from S_{RI} , it preserves the uncorrelated far-end information at the output of summer, and helps in extracting the far-end information from S_{FR} . Hence, the value of c_k depends only on the correlation between $D[k]$ and constituents of S_{NI} , and is independent of the far-end signal. Deviations in the value of c_k from the actual correlation may occur due to: (i) the presence of some finite correlation between the far-end and near-end transmitted signals; (ii) DC-offsets in the cancellation loops; and (iii) the finite loop gain of the correlator, leaving a residual-SI in S_{FR} at canceller's output.

Fig. 5: Correlator logic circuit with DCOC. Here, B_k : k^{th} control bit of DCOC.

Fig. 6: Die-micrograph of TFD-IO.

The canceller is implemented using current mode logic. Each canceller tap includes a correlator and a weighting amplifier WA – implemented as a Gilbert-cell. As shown in Fig. 5, the correlator uses a multiplier for determining the instantaneous correlation between a delayed tap $D[k]$ and the summer output S_{FR} . The multiplier's output is then amplified and integrated to determine c_k using a gain block and a low pass filter, respectively. The weight coefficient c_k controls the gain of WA_k for generating the weighted tap $T[k]$. Depending upon the data rates and interfacing scenarios, the c_k can have both the polarities and can take a wide range of values.

Figure 5 shows the circuit implementation of the correlator. The multiplier is implemented using a differential double balanced configuration with programmable load resistances, which serves as a DC offset compensator (DCOC) to correct offset in the loop. These offsets can occur due to various system anomalies, such as process, voltage, and temperature variations. The multiplier output is amplified using the gain block. The gain block is a cascade of two differential op-amps with an inbuilt common-mode feedback circuit and two differential amplifiers with variable gains. Variable gain provides additional flexibility to adjust the loop gain and maintain the stability of the loop. The output of the cascaded gain block is further integrated using the low pass filter to determine the weight coefficient c_k . The negative feedback action of the cancellation loop reduces the interference component at

Fig. 7: (a) Experimental setup and (b) the channel response. Here, I_T : control for the output swing of near-end transmitter; I_R : control for the linearity of near-end RFE; I_L : control for loop gain of SI-canceller; A1: source meters; A2: digital storage oscilloscope; A3: frequency synthesizer; and A4: arbitrary waveform generator.

the output of summer. The summer is implemented using resistances to add currents from various stages of the loop. For testing purpose, the TFD-IO cell has a built-in $2^7 - 1$ pseudo-random-bit-sequence (PRBS) generator.

IV. MEASUREMENT SETUP AND EXPERIMENTAL RESULTS

A prototype of the proposed TFD-IO is designed, and fabricated in a standard 65-nm CMOS low-leakage process with

Fig. 8: Configurations of TFD-IO to verify the performance of its: (a) transmitter, (b) SI-canceller, and (c) receiver. Inactive components and connections are shown using gray color. Dashed lines show path followed by the signals.

Fig. 9: Measured eye diagrams with the near-end transmitter active and the far-end source inactive: (a) Near-end transmitted signal S_T at port-1, (b) Near-end interference S_{NI} at port-2, and (c) residual interference at port-2 with interference cancellation and offset correction. These measurements were performed with the TFD-IO operating at 6.4 Gb/s.

a supply voltage of 1 V. Figure 6 shows the die micrograph of TFD-IO. The active core of TFD-IO has an area of $\approx 0.1 \text{ mm}^2$. For the measurements, a chip-on-board (CoB) arrangement is used with both the DC and the RF pads wire-bonded to a four-layer FR-4 PCB. PCB traces are designed to have a differential impedance of 100Ω for high-speed data and clock signals. Figure 7 shows the experimental setup along with the channel response. The channel has a transmission loss (S_{21}) of -8 dB at 3.2 GHz and return losses (S_{11}) better than -10 dB upto 3.7 GHz . TFD-IO is AC-coupled at port-1 using 10 nF capacitors to isolate DCs of the two ends.

DCOC is manually programmed to perform a coarse correction for DC-offsets present in the SI-canceller. DCOC is initialized using an on-chip serial-to-parallel interface (SPI) with the correction word required to correct for the offsets, which may arise due to variations in the process, voltage, and temperature (PVT). The near-end and far-end data are generated using an on-chip 2^7-1 PRBS generator and an external arbitrary waveform generator (A4), respectively. Our proposed TFD-IO is configured to transmit and receive signals with differential swings of 200 mV_{pp} . The performances of TFD-IO's transmitter, receiver, and correlation-based SI-canceller are verified as follows:

Transmitter: The system configuration utilised for the measurements is shown in Fig. 8(a). TFD-IO is used to generate the near-end signal and the far-end source is configured as a 100Ω differential load. The near-end signal, as a result of TFD-IO's bidirectional architecture, appears as S_T and S_{NI} , at port-1 and port-2, respectively. Eye-diagram for S_T

with a differential swing of 200 mV_{pp} is plotted in Fig. 9(a). Similarly, eye-diagram for S_{NI} with a differential swing of approximately 230 mV_{pp} is shown in Fig. 9(b), respectively. Here, S_T is the desired near-end transmitted signal, whereas S_{NI} is the undesired near-end interference.

Receiver: To verify the performance of TFD-IO as a receiver, the PRBS generator as well SI-canceller are deactivated and the far-end data is transmitted from the arbitrary waveform generator, as shown in Fig. 8(c). Figure 10(a) shows the far-end signal received at port-2. The received far-end signal has a differential swing of approximately 240 mV_{pp} . TFD-IO shows similar swings for the S_{NI} and the received far-end signal, as plotted in Fig. 9(b) and Fig. 10(a), respectively.

In the above two configurations, TFD-IO can reliably transmit and receive data without any processing over simplex links but need interference suppression for full-duplex operations.

SI-canceller: TFD-IO uses the correlation-based SI-canceller to suppress the interference, in the analog domain, at the output of the summer. To validate canceller's performance, near-end transmitter along with the SI-canceller is active, and the far-end source is configured as a 100Ω differential load, as shown in Fig. 8(b). Figure 9(c) shows the residual-SI present as S_{NI} at port-2, post interference cancellation plus DC offset correction. The cancellation along with offset correction reduces the residual-SI from 74 mV_{rms} to 12.8 mV_{rms} , and thus provides an interference suppression of 15.2 dB .

Interference suppression performed using the proposed SI-canceller is required to recover far-end information in TFD links. As discussed in Section II, the correlation-based ap-

Fig. 10: Measured eye diagrams for S_{FR} in TFD mode: (a) received far-end signal without near-end interference, (b) With near-end interference, and (c) With interference cancellation and DC offset correction. These measurements were performed with the two ends operating at 6.4 Gbps, the link thus supports an aggregate bidirectional throughput of 12.8 Gb/s.

Fig. 11: Measured eye diagrams for S_{FR} in TFD mode: (a) received far-end signal along with near-end interference, (b) recovered far-end signal (6.25 Gb/s) with interference cancellation plus offset correction, and (c) recovered far-end signal (6.4 Gb/s) with interference cancellation plus offset correction. These measurements were performed with the two ends operating at different baud rates.

proach for interference cancellation allows the SI-canceller to adapt its tap values in the background, independent of the far-end signal. The TFD-IO is operated in the following system configurations to validate the concept of true full-duplexing.

A. Same baud-rates and signalling schemes

The measurements are performed for a full-duplex link with the two ends, simultaneously transmitting, at data rates of 6.4 Gb/s using NRZ signalling scheme. Figure 10(b) shows eye-diagrams of the composite signal received at port-2, which comprises the far-end signal and the near-end interference. The presence of interference closes the eye of far-end signal as shown in Fig. 10(b) as compared to Fig. 10(a), where near-end interference is absent.

Conventionally, to extract the far-end information, adaptive interference cancellers in such links may require a foreground training and a phase locked arrangement [12], [18], [24], [25]. Whereas, the proposed TFD-IO because of the correlation-based SI-canceller can adapt its weights in presence of the far-end signal. Figure 10(c) shows eye-diagram of the far-end signal recovered post interference cancellation. The eye-diagrams show a reduction in the eye-height and eye-width because of the presence of residual-SI.

B. Different baud-rates and same signalling schemes

To validate the ability of SI-canceller to support interference cancellation independent of the phase of far-end

TABLE I: Error performance of the far-end signal in the TFD mode.

Near-end	Gbaud	6.4	6.4	6.4	6.25
	PRBS	2^7-1	2^7-1	2^7-1	2^7-1
Far-end	Gbaud	6.4	6.4	6.25	6.4
	PRBS	2^7	2^{15}	2^7	2^{15}
BER (Raw)		1.3E-10	7.9E-10	4.1E-11	7.6E-10
BER (FFE+DFE)		<1E-15	<1E-15	<1E-15	<1E-15

signal, the measurements are performed with the two ends operating at different baud-rates using same NRZ signalling scheme. Figure 11(a) shows eye-diagram of the composite signal received at port-2, which comprises the far-end signal and the interference due to near-end, operating at 6.25 Gb/s and 6.4 Gb/s, respectively. In such an operating condition, the relative phase of the two ends continuously varies. The interference, thus completely closes the eye-opening of far-end signal. SI-canceller is used to suppress the interference for recovering the far-end signal. Figure 11(b) shows eye-diagram of the recovered far-end signal. In addition, Fig. 11(c) shows eye-diagram of the recovered signal, for the measurement, where the far-end and the near-end are operated at 6.4 Gb/s and 6.25 Gb/s, respectively. Both eye-diagrams show a reduction in the eye-width and eye-height similar to Fig. 10(c).

Bit-error-rates (BER)s are estimated to evaluate the error

Fig. 12: Measured eye diagrams for S_{FR} in TFD mode: (a) received far-end PAM4 signal without near-end interference, (b) received far-end PAM4 signal without interference cancellation, and (c) recovered far-end PAM4 signal with interference cancellation plus offset correction. These measurements were performed with NRZ and PAM4 used as the near-end and far-end signalling schemes, respectively.

performance of the link, and are presented in Table I for the far-end signal recovered at port-2. BERs are estimated using Eq. (6) as given in [28], by determining the signal-to-noise ratio (SNR) at the sample point.

$$BER = \frac{1}{2} \times Q(\sqrt{SNR}). \quad (6)$$

Table I presents the error performance for the recovered signal corresponding to two PRBS sequences (used for transmitting the far-end data): 2^7 and 2^{15} . These cases show a similar BERs, and thus advise a nearby signal-to-interference ratio for the recovered signals. Here, correlation-based SI-canceller adapts its weights in the background independent of the data sequence (or information) used for transmitting the far-end signal. In addition, the Table also shows the estimated BERs with the two ends operating at the same and different baud-rates. The BERs are for the following three combinations of near-end and far-end baud-rates: (i) 6.4 Gb/s (phase of near-end and far-end data change at the same rate), (ii) 6.4 Gb/s / 6.25 Gb/s (phase of near-end data changes faster than the far-end data), and (iii) 6.25 Gb/s / 6.4 Gb/s (phase of far-end data changes faster than the near-end data). The three combinations show a similar error rate for the recovered far-end signal, and thus verify the ability of SI-canceller to support interference cancellation independent of the phase or baud-rate of far-end data. In addition, BERs has also been presented for the recovered far-end with equalization performed using a 2-tap feed-forward equalizer (FFE) and a 1-tap decision feedback equalizer (DFE) to compensate for the channel loss. Eye-diagram for the equalized far-end signal corresponding to Fig. 10(c) is shown in Fig. 13.

C. Same baud-rates and different signalling schemes

To validate SI-canceller's ability to support interference cancellation at the near-end, independent of the signalling scheme used by far-end transceiver, a proof of concept demonstration is done with different signalling schemes used at the two ends. The near-end and far-end signals are transmitted at 500 MBd using NRZ and PAM4 signalling schemes, respectively. To operate in the linear region, the gain of the Rx buffers is lowered using I_R , and the output swing of the near-end transmitted signal is reduced to 160 mVpp using I_T . Figure 12(a)

Fig. 13: Far-end signal received post interference cancellation in TFD-IO mode, and equalized using a 2-tap FFE and a 1-tap DFE.

and Fig. 12(b) show eye-diagrams of the far-end PAM4 signals received without and with near-end interference, respectively. In contrast to Fig. 12(a), the eyes for S_{FR} in Fig. 12(b) are entirely closed. Figure 12(c) shows eye-diagram of the far-end PAM4 signal received post near-end interference suppression performed using SI-canceller. Figure 12(c) shows a significant eye opening for the received signal as compared to Fig. 12(b). Thus, the SI-canceller can faithfully recover the far-end PAM4 signal while using NRZ signalling scheme for near-end data transmission.

In the above configurations, the far-end information was recovered with the SI-canceller operated in the background. These measurements, thus validate SI-canceller's ability to adapt in the presence of far-end signal, independent of the signalling schemes and baud rates used by the transceivers at the two ends. TFD-IO with the correlation-based SI-canceller, thus eliminates the need to have a phase/frequency synchronization or clock forwarding between the two transceivers. Additionally, TFD-IO can also support interference cancellation in re-configurable transceivers [29], [30] that allow data transfers in multiple signalling schemes.

A summary of the TFD-IO performance, and its comparison with prior-art transceivers is given in Table II. Power breakdown for the design with one SI-cancellation tap and three echo cancellation taps is provided in Fig. 14. The proposed TFD-IO achieves an energy efficiency of 1.85 pJ/bit. This proof-of-concept implementation focuses on verifying the

TABLE II: PERFORMANCE SUMMARY AND COMPARISON TO PRIOR ART

References	[22]	[16]	[17]	[12]	[18]	[11]	[19]	This work
Technology	180-nm CMOS	110-nm CMOS	14-nm FinFet	28-nm CMOS	14-nm FinFet	65-nm CMOS	65-nm CMOS	65-nm CMOS
SBD Throughput	8 Gb/s/port	20 Gb/s/port	112 Gb/s/port	32 Gb/s/port	32 Gb/s/port	15 Gb/s/port	6 Gb/s/port	12.8 Gb/s/port
Architecture	TX:CM, Hybrid: SC based comparator having variable offsets and a replica driver	TX:CM, Hybrid: R-gm	TX:DAC, Hybrid: passive + digital echo canceller	TX:VM, Hybrid: Tunable R-gm hybrid plus a SSLMS based echo canceller	TX:VM, Hybrid: comparator with a replica driver	TX:VM, Hybrid: resistive-bridge	TX:VM, Hybrid: passive and a replica driver	TX:CM, Hybrid: correlation-based SI-canceller
Canceller	Fixed	Tunable	Adaptive	Adaptive	Fixed	Tunable	Tunable	Adaptive
Canceller training	NA	NA	Foreground	Foreground	NA	NA	NA	Background
Clock synchronization	Required	Not Required	Required	Required	Required	Not Required	Not Required	Not Required
Signalling schemes	NRZ	NRZ	PAM4	NRZ	NRZ	NRZ	PAM4	NRZ, PAM4*
Equalization	2-tap FFE @ TX	2-tap FFE @ TX, CTLE @ RX	16-tap FFE @ TX, CTLE + 32-tap FFE + 1-tap DFE @ RX (offline)	CTLE @ RX	No	No	De-emphasis @ TX	2-tap FFE + 1-tap DFE @ RX (offline)
Echo cancellation	No	No	1 NE taps	2 NE taps 2 FE taps	No	No	No	3 NE taps
Channel loss @ Nyquist	0.6 dB	5 dB	35 dB	10.2 dB	<4 dB	2.5 dB	-	8 dB
Area	0.13 mm ²	1.02 mm ²	≈1.5 mm ²	0.182 mm ²	0.02 mm ²	0.012 mm ²	0.3 mm ²	0.1 mm ²
Supply	-	1.2 V	1.2 V/0.9 V	0.9 V	0.9 V	1.2 V	1 V	1 V
Power consumption (single-side)	Transceiver: 158 mW	TX:126 mW RX:140 mW	Transceiver: 994 mW	Transceiver: 29.2 mW	Transceiver: 11.2 mW	Transceiver: 10.1 mW	Transceiver: 37.3 mW	Transceiver: 21 mW
Power Efficiency	39.5 pJ/bit	26.6 pJ/bit	17.75 pJ/bit	1.83 pJ/bit	0.7 pJ/bit	1.35 pJ/bit	12.4 pJ/bit	1.85 ^a pJ/bit 3.28 ^b pJ/bit

*-Supported by the SI-canceller and Rx; a and b without and with echo cancellation, respectively. Here, NE: near-end; FE: far-end; VM: voltage mode; CM: current mode; FFE: feed-forward equalizer; DFE: decision feedback equalizer; CTLE: continuous time linear equalizer.

Fig. 14: TFD-IO power dissipation.

concept of true full-duplexing, and the TFD-IO uses 44% of its power for realising three echo cancellation taps, which however lowers the efficiency to 3.28 pJ/bit but also allows to minimise the echoes adaptively in the background. The tap count and the power consumption can be lowered for links with better impedance matching and smaller reach. In addition, the energy efficiency can be further improved by using LDs with a higher power efficiency [31]. The demonstrated TFD link supports an aggregate bidirectional data throughput of

12.8 Gb/s, with a raw BER better than 10^{-10} , and with a BER better than 10^{-15} post equalization performed offline. The use of proposed TFD-IO helps in lowering the effect of cross-talk and thus improving the performance of the high-density chip-to-chip interconnects.

V. CONCLUSION

In this work, we presented a true full-duplex IO (TFD-IO) that can serve as a plug-n-play module for high-density interfaces to convert a unidirectional interconnect to a fully bidirectional interconnect. The TFD-IO performs interference suppression in the analog domain using a correlation-based SI-canceller to support bidirectional data transfer. The SI-canceller adaptively determines its coefficients even in the presence of the signal transmitted by the far-end transceiver. TFD-IO has been experimentally validated with a proof-of-concept demonstration for a bidirectional throughput of up to 12.8 Gb/s. The interference cancellation technique used in the TFD-IO allows for the use of different baud-rates and signalling schemes, such as NRZ/PAM4, by the two transceivers. The IO module presented in this work can be used in a practical design for doubling the throughput of the interface. This approach can be very useful in improving the data bandwidth of a high-density interface. The TFD-IOs should also help in simplifying the

floor planning of PCBs for chip-to-chip multi-lane high-speed interfaces, as well as for interfaces in chiplet based designs used in SiPs/MCMs.

VI. ACKNOWLEDGEMENT

The authors thank the SMDP-C2SD program of MeitY (Govt. of India), and Qualcomm Innovation Fellowship 2020-21, for supporting this activity in parts. We would also like to acknowledge Mr. Mukund Narasimhan and Mr. Shajith Sirajudeen from Qualcomm, for some useful discussions. Initial funding by SRC for activity in this area (task no. 2739.001) is also acknowledged.

REFERENCES

- [1] F. O'Mahony, G. Balamurugan, J. E. Jaussi, J. Kennedy, M. Mansuri, S. Shekhar, and B. Casper, "The future of electrical I/O for microprocessors," in *IEEE Symp. on VLSI Design, Automation and Test*, 2009, pp. 31–34.
- [2] J. Macri, "AMD's next generation GPU and high bandwidth memory architecture: FURY," in *IEEE Hot Chips*, 2015, pp. 1–26.
- [3] Y. Wen, "The package bandwidth limitation of high speed broadband products," in *IEEE ECTC*, 2001, pp. 1457–1462.
- [4] J. W. Poulton, W. J. Dally, X. Chen, J. G. Eyles, T. H. Greer, S. G. Tell, and C. T. Gray, "A 0.54pJ/b 20Gb/s ground-referenced single-ended short-haul serial link in 28nm CMOS for advanced packaging applications," in *IEEE ISSCC Dig. Tech. Papers*, 2013, pp. 404–405.
- [5] A. C. Durgun, Z. Qian, K. Aygun, R. Mahajan, T. T. Hoang, and S. Y. Shumarayev, "Electrical Performance Limits of Fine Pitch Interconnects for Heterogeneous Integration," in *IEEE ECTC*, 2019, pp. 667–673.
- [6] A. Shokrollahi, D. Carnelli, J. Fox, K. Hofstra, B. Holden, A. Hormati, P. Hunt, M. Johnston, J. Keay, S. Pesenti, R. Simpson, D. Stauffer, A. Stewart, G. Surace, A. Tajalli, O. T. Amiri, A. Tschank, R. Ulrich, C. Walter, F. Licciardello, Y. Mogentale, and A. Singh, "A pin-efficient 20.83 Gb/s/wire 0.94 pJ/bit forwarded clock CNRZ-5-coded SerDes up to 12 mm for MCM packages in 28 nm CMOS," in *IEEE ISSCC*, 2016, pp. 182–183.
- [7] D. S. Green, "Common Heterogeneous Integration and Intellectual Property (IP) Reuse Strategies (CHIPS)," [Online]. Available: <https://www.darpa.mil/program/common-heterogeneous-integration-and-ip-reuse-strategies> [Accessed: 21-06-2021].
- [8] D. Kehlet, "Accelerating Innovation Through a Standard Chiplet Interface: The Advanced Interface Bus," [Online]. Available: <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/accelerating-innovation-through-aib-whitepaper.pdf> [Accessed: 18-07-2021].
- [9] J. Fourniols, M. Roca, F. Caignet, and E. Sicard, "Characterization of crosstalk noise in submicron CMOS integrated circuits: an experimental view," *IEEE Trans. Electromagn. Compat.*, vol. 40, no. 3, pp. 271–280, 1998.
- [10] F. D. Mbairi, W. P. Siebert, and H. Hesselbom, "High-Frequency Transmission Lines Crosstalk Reduction Using Spacing Rules," *IEEE Trans. Compon. Packag. Technol.*, vol. 31, no. 3, pp. 601–610, 2008.
- [11] C. Yuan, A. Naguib, and S. Shekhar, "On the Design of Low-Power Hybrids for Full Duplex Simultaneous Bidirectional Signaling Links," *IEEE TCASI: Regular Papers*, vol. 67, no. 4, pp. 1413–1422, 2020.
- [12] Y.-H. Fan, A. Kumar, T. Iwai, A. Roshan-Zamir, S. Cai, B. Sun, and S. Palermo, "A 32 Gb/s Simultaneous Bidirectional Source-Synchronous Transceiver With Adaptive Echo Cancellation Techniques," *IEEE J. Solid-State Circuits*, vol. 55, no. 2, pp. 439–451, 2020.
- [13] P. Roo, S. Sutardja, S. Wei, F. Aram, and Y. Cheng, "A CMOS transceiver analog front-end for gigabit ethernet over CAT-5 cables," in *IEEE ISSCC Dig. Tech. Papers.*, 2001, pp. 310–311.
- [14] "IEEE 802.3 Ethernet Working Group," [Online]. Available: <http://www.ieee802.org/3> [Accessed: 18-07-2021].
- [15] G. Chandra and M. Malkin, "A full-duplex 10GBase-T transmitter hybrid with SFDR > 65dBc Over 1 to 400 MHz in 40 nm CMOS," in *IEEE ISSCC*, 2011, pp. 144–146.
- [16] Y. Tomita, H. Tamura, M. Kibune, J. Ogawa, K. Gotoh, and T. Kuroda, "A 20 Gb/s Simultaneous Bidirectional Transceiver Using a Resistor-Transconductor Hybrid in 0.11- μ m CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 3, pp. 627–636, 2007.
- [17] R. Farjadrad, K. Kaviani, D. Nguyen, M. Brown, G. Geelen, C. Bastiaansen, N. Rao, V. Popuri, G. Shen, H. Khatibi, S. Dey, A. Chatterjee, D. Shen, P. Zijlstra, H. Gunnink, K. Zhang, V. Penumuchu, O. Weiss, E. Paulus, and J. Briaire, "An Echo-Cancelling Front-End for 112Gb/s PAM-4 Simultaneous Bidirectional Signaling in 14nm CMOS," in *IEEE ISSCC*, vol. 64, 2021, pp. 194–196.
- [18] R. Farjadrad, M. Kuemerle, and B. Vinnakota, "A Bunch-of-Wires (BoW) Interface for Interchiplet Communication," *IEEE Micro*, vol. 40, no. 1, pp. 15–24, 2020.
- [19] S. Mukherjee, A. Das, S. Seth, and S. Saxena, "An Energy-Efficient 3 Gb/s PAM4 Full-Duplex Transmitter With 2-Tap Feed Forward Equalizer," *IEEE TCASII: Express Briefs*, vol. 67, no. 5, pp. 916–920, 2020.
- [20] T. Kim, S. Jang, S. Kim, S. Chu, J. Park, and D. Jeong, "A Four-Channel 32-Gb/s Transceiver With Current-Recycling Output Driver and On-Chip AC Coupling in 65-nm CMOS Process," *IEEE TCASII: Express Briefs*, vol. 61, no. 5, pp. 304–308, 2014.
- [21] R. Drost and B. Wooley, "An 8 Gb/s/pin simultaneously bidirectional transceiver in 0.35- μ m CMOS," *IEEE J. Solid-State Circuits*, vol. 39, no. 11, pp. 1894–1908, 2004.
- [22] B. Casper, A. Martin, J. Jaussi, J. Kennedy, and R. Mooney, "An 8 Gb/s simultaneous bidirectional link with on-die waveform capture," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2111–2120, 2003.
- [23] N. Wary and P. Mandal, "Current-Mode Full-Duplex Transceiver for Lossy On-Chip Global Interconnects," *IEEE J. Solid-State Circuits*, vol. 52, no. 8, pp. 2026–2037, 2017.
- [24] H. Tamura, M. Kibune, Y. Takahashi, Y. Doi, T. Chiba, H. Higashi, H. Takauchi, H. Ishida, and K. Gotoh, "5 Gb/s bidirectional balanced-line link compliant with plesiochronous clocking," in *IEEE ISSCC Dig. Tech. Papers*, 2001, pp. 64–65.
- [25] T. Takahashi, T. Muto, Y. Shirai, F. Shirotori, Y. Takada, A. Yamagiwa, A. Nishida, A. Hotta, and T. Kiyuna, "110 Gb/s simultaneous bidirectional transceiver logic synchronized with a system clock," *IEEE J. Solid-State Circuits*, vol. 34, no. 11, pp. 1526–1533, 1999.
- [26] S. Goyal, P. Agarwal, and S. Gupta, "Demonstration of a Single-Lane 80 Gbps PAM-4 Full-Duplex Serial Link," in *IEEE Hot Interconnects*, 2019, pp. 32–35.
- [27] S. Goyal, G. A. Parulekar, and S. Gupta, "A True Full-Duplex IO for High-Density High-Speed Interconnects," in *IEEE ICECS*, 2020, pp. 1–4.
- [28] S. Haykin, *Communication Systems*, 5th ed. Wiley Publishing, 2009.
- [29] A. Roshan-Zamir, O. Elhadidy, H.-W. Yang, and S. Palermo, "A Reconfigurable 16/32 Gb/s Dual-Mode NRZ/PAM4 SerDes in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 52, no. 9, pp. 2430–2447, 2017.
- [30] V. Stojanovic, A. Ho, B. Garlepp, F. Chen, J. Wei, G. Tsang, E. Alon, R. Kollipara, C. Werner, J. Zerbe, and M. Horowitz, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, 2005.
- [31] K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, E. Suzuki, N. Masuda, T. Takemoto, F. Yuki, and T. Saito, "A 12.3-mW 12.5-Gb/s Complete Transceiver in 65-nm CMOS Process," *IEEE J. of Solid-State Circuits*, vol. 45, no. 12, pp. 2838–2849, 2010.