

A 5Gb/s Transmitter with Reflection Cancellation for Backplane Transceivers

Ricky Yuen, Marcus van Ierssel, Ali Sheikholeslami, William W. Walker¹, and Hirotaka Tamura²

Department of Electrical and Computer Engineering, University of Toronto, Canada

1. Fujitsu Laboratories of America, 2. Fujitsu Laboratories Limited, Kawasaki, Japan

Abstract — We present a 5Gb/s transmitter that cancels the reflected signals from any impedance discontinuity located at up to 64UI away from the transmitter and spread over 8UI interval. Measured results from our 0.11 μ m CMOS design reveal a 150mV eye-opening, from a nearly closed eye, when reflection cancellation is activated. The design consumes 510 μ A for the PLL operation, 60mA for data generation, and 50mA for data transmission, all from a 1.2V supply.

I. INTRODUCTION

One of the main challenges of high-speed chip-to-chip signaling is signal integrity in the face of channel loss, intersymbol interference (ISI), cross-talk, and signal reflections. Much attention has been paid to channel loss and ISI cancellation [1], but it is only recently that the issues of crosstalk [2] and signal reflections have begun to be addressed [3]. Crosstalk and reflections consume a bigger portion of the signal margin as the signaling rate increases well into the Gb/s regimes, resulting in worse bit-error rates (BER). One direct way to reduce reflections is to reduce channel discontinuities, but this is often a difficult task to achieve, especially at the interfaces of the channel with the transmitter and the receiver, and at the points where the signal passes through a connector (e.g. the connector between a daughter card and the backplane). Zerbe et al proposed a scheme to cancel the signal reflections at the *receiver* [3] as shown in Fig. 1(a). Zerbe’s scheme can cancel the *first* reflection signal caused by a channel discontinuity and the transmitter, but not *secondary* reflections. In contrast, our proposed scheme, shown in Fig. 1(b), cancels the reflections at the *transmitter*, effectively terminating the first reflection signal and not generating the secondary reflections. In addition, since the reflected pulse spreads in duration as it travels the length of a band-limited channel, it is easier to cancel it at the transmitter, before it could actually spread; i.e., the transmitter requires fewer taps than the receiver to cancel the same reflection.

To illustrate the effect of reflections, Fig. 2 shows an example where the receiver is terminated at 200 Ω (instead of at ideally 50 Ω) and where the transmitter is left unterminated. This causes several reflected pulses (two are shown) to arrive at the receiver in addition to the main pulse. By carefully mimicking the reflection signal path, we recreate the reflected signal in the

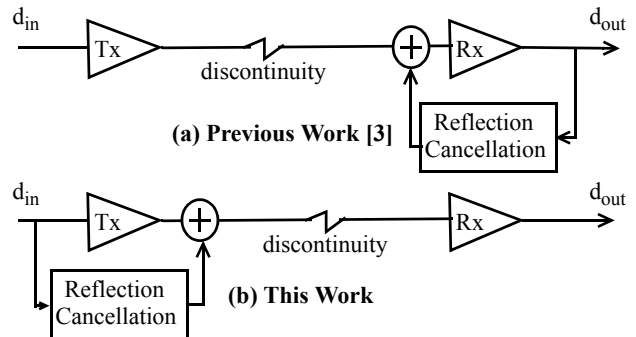


Fig. 1: Transceiver with reflection cancellation block (a) at the receiver (previous work), (b) at the transmitter (this work).




Fig. 2: Received signal (a) without reflection cancellation, (b) with reflection cancellation

transmitter side in order to subtract it from the outbound signal, eliminating both reflections. We confirm the validity of this approach by measurement results of a testchip implemented in 0.11 μ m CMOS.

II. TRANSMITTER ARCHITECTURE

The transmitter consists of four main blocks as shown in Fig. 3: A PLL, two data blocks, a MUX, and the drivers. The PLL multiplies an externally-provided 78MHz reference clock by 32, and produces two phases of a 2.5GHz clock, $Clk0$ and $Clk180$ in the figure. These two clock phases trigger two corresponding data blocks, producing two sequences of 2.5Gb/s data streams. Each data block also produces 11 delayed versions of its data streams (three 1UI-delayed streams for a 3-tap pre-emphasis, and eight 1UI-delayed streams for an 8-tap reflection cancellation as explained later). The 2x12 streams of




Fig. 3: Transmitter architecture




Fig. 4: VCO Implementation

data are then merged (using a total of 12 2:1 MUXes) to produce 12 streams of 5Gb/s, to be used in the front-end drivers. All the signals are fully differential, although for ease of readability we show them as single-ended in most of this paper. Next, we will describe the implementation detail of each block.

III. BUILDING BLOCKS AND CIRCUIT IMPLEMENTATION

Phase-Locked Loop

The PLL is designed to produce a low-jitter 2.5GHz clock output from an off-chip 78MHz reference clock. The PLL uses a standard NOR-based phase-frequency detector (PFD), a charge pump, followed by a second-order passive loop filter [4] that provides the control voltage to the VCO. The VCO's frequency is divided by 32, and fed to the PFD, closing the loop. With these settings, the VCO output frequency is measured to be 2.5Gb/s.

The VCO is implemented using a negative Gm and an LC tank as shown in Fig. 4. The L is a spiral inductor implemented with top-layer metal, with a simulated inductance of 1nH. Tuning is performed by two back-to-back MOS varactors with their gates connected to the output nodes of the VCO. Our simulation results indicate that the VCO has a tuning range of 2.5-3.8GHz with a gain of 1.66GHz/V.

Data Block

For testing purposes, the data block (Fig. 5) allows the selection of either an on-chip $2^{31}-1$ Pseudo-Random Bit Sequence (PRBS) or an on-chip programmable repetitive 128-bit sequence, both aligned with either Clk180 (for data block

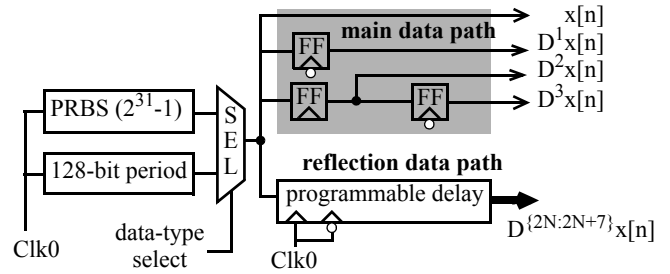


Fig. 5: Data block (refer to Fig. 3). D^1 refers to 1UI delay




Fig. 6: Programmable Delay Implementation

#1) or Clk0 (for data block #2). The selected data then proceeds along the *main data path* and the *reflection data path* simultaneously. In the main data path, the data stream is simply delayed successively by 1UI (=200ps) per stage, producing three delayed versions of the original data stream. Since the data rate at this point is 2.5Gb/s, the flip-flops are triggered with opposite phases of the 2.5GHz clock in order to produce 1UI delays. In the reflection data path, the data stream undergoes a programmable delay of $2N$ UI, where N can be programmed to be any integer between 0 and 63. This is to match the signal delay as it travels from the transmitter to the point of discontinuity and back to the transmitter. At 5Gb/s signaling, the 126UI delay corresponds to about 6 meters of cable/trace, allowing the maximum channel length to be 3 meters.

The transmit signal is both delayed and *spread* in time. To mimic the spread of the reflected signal, the delay block also provides eight 1UI-delayed versions of the $2N$ -UI-delayed sequence. A weighted sum of these streams will reproduce the shape of the reflected signal as described later in this section.

Programmable Delay

By combining six 2:1 MUXes and a total of 63 flip-flops clocked with Clk0, the architecture in Fig. 6 creates a delay that can be programmed between 0UI and 126UI, with a granularity of 2UI (since the flip-flops are clocked at 2UI intervals). Seven additional flip-flops at the end create 1UI-delay granularity at $2N$ -UI-delay point. This is done simply by triggering alternating flip-flops with the opposite phase of the 2.5GHz clock, as shown in the figure.

By introducing only a 2:1 MUX input and a single flip-flop input, this architecture minimizes the parasitic capacitance on

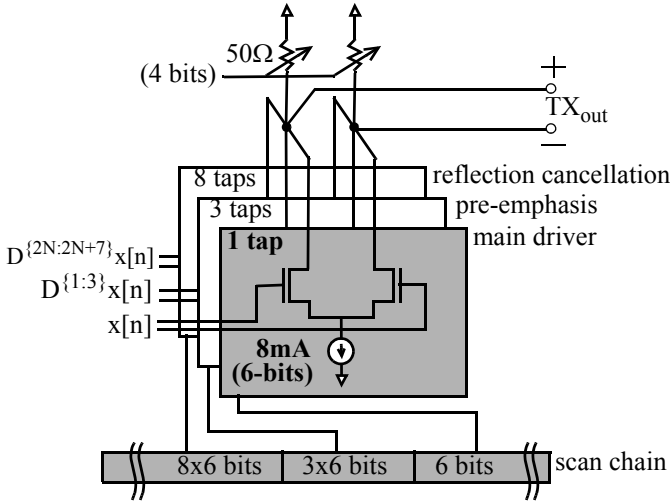


Fig. 7: Front-end drivers: main, pre-emphasis, reflection cancellation

the node shared between the two data paths (see Fig. 5). Note that this block has an inherent 6UI delay due to the use of three MUX-FF's (each contributing 2UI delay) [5]. The same delay has been added to the data path (not shown) in order to match this inherent delay.

Front-End Drivers

The front end consists of a total of 12 differential pairs, all tied together at their output nodes, as shown in Fig. 7. One of these differential pairs (main driver) is driven with the 5Gb/s bit sequence, three (pre-emphasis driver) are driven by immediate 1UI-delayed versions of the bit sequence, and eight (reflection cancellation drivers) with 1UI-delayed versions of the bit sequence with $2N$ UI base delay. The tail current in each of these pairs is controlled digitally with 6 bits, requiring a total of 72 bits. Together, these bits shape the ISI cancellation signal and the reflection cancellation signal that are added to the transmit signal. The 72 bits, in addition to 4 bits controlling the transmitter termination resistance are supplied through the scan chain in the testchip.

IV. MEASUREMENT RESULTS

The design is implemented in a $0.11\mu\text{m}$ CMOS process, and measures $700\mu\text{m} \times 375\mu\text{m}$, as shown in Fig. 8. The testchip is fully functional, as confirmed by the bathtub curves shown in Fig. 9 and the eye diagrams shown in Fig. 10 and 11, all using the on-chip $2^{31}-1$ PRBS generator.

To show the reflection cancellation scheme works under strong reflections, we have connected our unterminated transmitter (high-Z condition) to the receiver (Centellax's OTB1P1A PRBS/BERT board with an effective 22Ω termination) through a 1m SMA cable. The transmitter clock is also fed to the board with the same-type cable but with an extra length that is an integer multiple of 0.5cm (corresponding to 0.1UI of the 5Gb/s signal, that is 20ps). As a result, we can sample the received signal at 1UI intervals but with a granularity of 0.1UI, and




Fig. 8: chip photo (top layer to be removed for publication and presentation)




Fig. 9: Measured BER as a function of sample position, in units of 0.1UI

measure the BER accordingly. The results, as shown in Fig. 9, confirm that the BER improves significantly when the reflection cancellation circuitry is activated.

The eye diagrams of Fig. 10 correspond to bathtub curve measurements of Fig. 9 but using a TDX7704B oscilloscope instead of the BERT board. The eye diagrams confirm that the reflection cancellation recovers a healthy eye from an otherwise closed eye. We have repeated these measurements for a 200Ω termination (equivalent to a reflection coefficient of +60%). Again, with the reflection cancellation activated, the receiver observes an eye opening of 150mV.

The test chip consumes $510\mu\text{A}$ from the 1.2V analog PLL supply and 110mA from the 1.2V digital supply. Out of the 110mA, an estimated 60mA is consumed in the PRBS and the periodic sequence generator, and only 50mA in the actual transmitter.

V. CONCLUSIONS

We demonstrated a 5Gb/s transmitter implementing both an ISI cancellation and a reflection cancellation scheme at the transmitter. The reflection cancellation scheme is shown to successfully cancel both positive and negative reflections,




Fig. 10: Eye diagram of the 5Gb/s high-Z transmitter with a 200Ω receiver termination (a) with no pre-emphasis and no reflection cancellation, (b) with pre-emphasis only, (c) with both pre-emphasis and reflection cancellation




Fig. 11: Eye diagram of the 5Gb/s high-Z transmitter with a 200Ω receiver termination (a) with no pre-emphasis and no reflection cancellation, (b) with pre-emphasis only, (c) with both pre-emphasis and reflection cancellation

corresponding to receiver impedance being lower and higher than the characteristic impedance of the PCB trace.

ACKNOWLEDGMENT

The authors are grateful to N. Nedovic and N. Tzartzanis from Fujitsu Laboratories of America for reviewing the manuscript, NSERC and Fujitsu for funding, Fujitsu for test chip fabrication, and CMC for testing facilities.

REFERENCES

- [1] M. Horowitz et al., "High-speed electrical signalling: overview and limitations," IEEE Micro, Jan.-Feb. 1998, pp. 12-24
- [2] C. Pelard et al., "Realization of multigigabit channel equalization and crosstalk cancellation integrated circuits, JSSC, Oct. 2004, pp. 1659-69
- [3] J. Zerbe et al., "Equalization and clock recovery for a 2.5-10Gb/s 2-PAM/4-PAM backplane transceiver cell," JSSC, Dec. 2003, pp. 2121-30
- [4] F. M. Gardner, "Phaselock Techniques," J. Wiley & Sons Inc., Third Edition, 2005
- [5] B. Nikolic et al., "Sense amplifier-based flip-flop" ISSCC, 1999, pp. 282-3