

A 0.3 pJ/bit 20 Gb/s/Wire Parallel Interface for Die-to-Die Communication

Behzad Dehlaghi, *Student Member, IEEE*, and Anthony Chan Carusone, *Senior Member, IEEE*

Abstract—A high-density low-power parallel I/O for die-to-die communication is presented. The proposed interface includes a low-power transceiver and a high-density low-cost silicon interposer. The link architecture exploits single-sided and capacitive termination, passive equalization in the transmitter, and CMOS logic-style circuits to reduce the power consumption. To achieve a high bump/wire efficiency, single-ended signaling is used. A 4-layer Aluminum silicon interposer is fabricated providing 2.5 mm and 3.5 mm links between prototype transceivers. The transceiver prototype includes 3 transmitters and 3 receivers fabricated in 28 nm STM FD-SOI CMOS technology. The parallel interface operates at 20 Gb/s/wire and 18 Gb/s/wire data rates over the 2.5 mm and 3.5 mm channels with 5.9 and 7.7 dB of loss relative to DC (10.7 and 13.5 dB total loss) at $f_{bit}/2$ while consuming 0.30 and 0.32 pJ/bit excluding clocking circuits, respectively.

Index Terms—Chip-to-chip communication, die-to-die communication, high-speed I/O, low-power transceivers, passive equalizer, silicon Interposer, single-ended signaling, termination.

I. INTRODUCTION

TO MEET demands for 10+ Gb/s/link signaling between chips, the complexity of I/O circuits has grown for decades with an attendant increase in I/O power consumption. As shown in Fig. 1, the power consumption of a wireline transceiver including all the clocking, equalization, data serializing and deserializing circuits, increases ten-fold with 30 dB increase in the channel loss [1]. The hybrid memory cube (HMC) specification [2] and high bandwidth memory (HBM) standard [3] call for an aggregate bandwidth of 8 Tb/s in next generation memory-to-processor links. On the other hand, thermal constraints typically limit the total power consumption of a CMOS chip to roughly 120 W [4]. Allocating 15% of this power for I/O circuits, only 18 W can be consumed in the I/O circuits. This means that the energy efficiency of the 8 Tb/s interface must be better than 2.25 pJ/bit. To achieve this energy efficiency, the channel loss must be below 10–15 dB as shown in Fig. 1.

To simultaneously keep the channel loss below 15 dB and keep the overall system size small, dies may be packaged and interconnected on shared packaging substrates such as silicon interposers and organic substrates. In such systems,

Manuscript received April 13, 2016; revised July 7, 2016; accepted July 24, 2016. Date of publication August 25, 2016; date of current version October 29, 2016. This paper was approved by Associate Editor Jack Kenney.

The authors are with the Electrical and Computer Engineering Department, University of Toronto, Toronto, ON M5S3G4, Canada.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2016.2596773

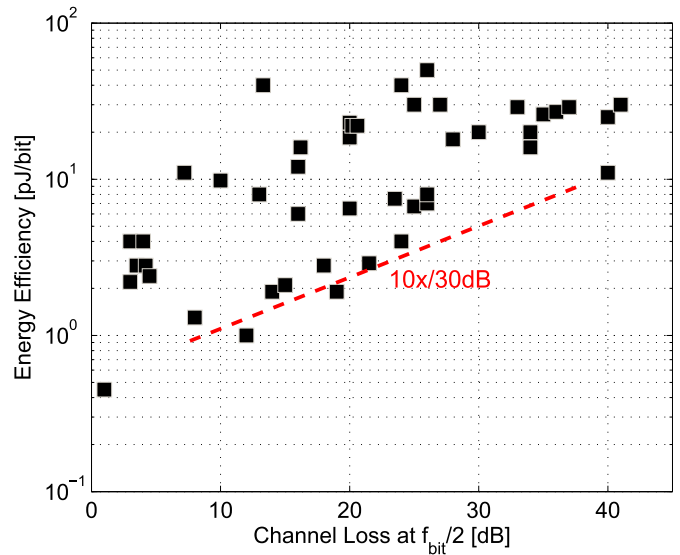


Fig. 1. Wireline transceiver energy efficiency vs. channel loss at one-half the bit rate among papers published at ISSCC from the past 10 years [1].

all the chips are flip chip mounted onto the shared substrate and interconnected via the redistribution layer as shown in Fig. 2. In recent years, these packaging technologies have been employed in short reach links [5]–[7]. There are two primary applications for these high-density die-to-die links. First, interfaces between two chips realized in different technologies which therefore cannot be integrated onto a single chip; for example a processor and memory. Second, multiple cores of an FPGA or CPU which cannot be integrated onto a single chip due to yield and cooling problems.

This paper presents a parallel interface for die-to-die communication. The proposed interface employs a silicon interposer providing the interconnects between chips. The interconnect spacing on the silicon interposer is chosen to maximize the bandwidth density without degrading performance due to crosstalk between adjacent lanes. Moreover, single-ended signaling is used to increase the data rate per bump/wire. The transceiver uses CMOS logic style circuits to minimize power consumption. A passive equalizer is used in the transmitter to compensate for channel losses up to 15 dB. Single-sided and capacitive termination techniques are employed to reduce the power consumption in the link front-end.

This paper is organized as follows. Section II discusses the proposed link architecture and the techniques used to decrease the power consumption. In Section III, the transmitter and

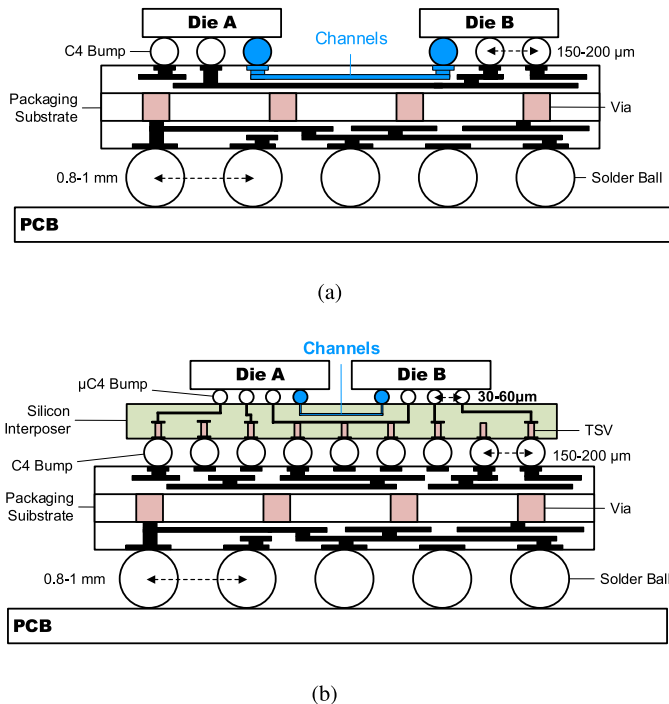


Fig. 2. Cross section of emerging packaging technologies in die-to-die communication: (a) organic substrates; (b) silicon interposers.




Fig. 3. Proposed link architecture.

receiver design is detailed. Section IV covers the design of the silicon interposer. Measurement results for the parallel interface are reported in Section V, and conclusions are drawn in Section VI.

II. PROPOSED LINK ARCHITECTURE

The proposed link architecture is shown in Fig. 3. A CMOS inverter is used as the transmitter output driver. Since the channel loss is not to be more than 15 dB, a passive equalizer is sufficient to compensate for the frequency dependent loss. In the receiver, a large capacitor (C_T) is used in series with the termination impedance to reduce power consumption and also generate the reference voltage, V_{ref} . At high frequencies, C_T acts as a short circuit and the receiver termination impedance becomes approximately equal to $R_T || (R_F/A)$, where A is the gain of pre-amplifier. Since the pre-amplifier gain varies significantly with the input signal swing, process and temperature variations, R_F must be large to reduce the

impact of these variations on receiver termination impedance. On the other hand, a large R_F can result in more duty-cycle distortion at the output of the pre-amplifier. In this work, R_F is 1 k Ω while the inverter gain is 4 for an input signal swing of 100 mV. This results in 250 Ω for the pre-amplifier's input impedance while the difference between the DC value of V_{sig} and V_{RX} is limited to 50 mV based on the simulation results with process variations. The DC offset between V_{sig} and V_{RX} is canceled using the following latch offset cancellation.

In the remainder of this section, the details of the termination and equalization techniques are discussed.

A. Capacitive and Single-Sided Termination

Fig. 4 shows different possible receiver termination schemes for single-ended signaling with their average power consumption assuming the same signal swing in all cases. In Fig. 4(a) Thevenin termination, the receiver DC voltage is $V_{DD}/2$. Hence, the reference voltage can be generated using another resistive divider, with larger value resistors to reduce power. This scheme is not power efficient because of the DC current in the receiver termination. The second method, Fig. 4(b), that is commonly used in memory interfaces due to its better power efficiency is parallel termination [7], [8]. The DC voltage of the received signal depends on the signal swing and is near ground. Hence, the reference voltage is generated by a digital-to-analog converter (DAC) that is calibrated by sending a training sequence from the transmitter [7], [8]. In this work, capacitive or AC termination is used due to its excellent power efficiency. As shown in Fig. 4(c), the power consumption of capacitive termination is half that of parallel termination, once the capacitor C_T is charged up to approximately $V_{DD}/2$. When transmitting a "1", the current flows from V_{DD} to C_T , and when transmitting a "0", C_T discharges to ground through the transmitter pull-down network. It has to be noted that the current equations are valid as long as C_T holds its voltage or in other words the signal is DC balanced without long sequences of consecutive identical digits (CIDs).

Another advantage of capacitive termination is that the reference signal is extracted from the incoming signal. Fig. 5 shows how the DC voltage is set using the transmitter and the receiver pre-amplifier, where C_T is considered an open circuit. Assuming the same supply voltage of V_{DD} for both transmitter and receiver, the received DC voltage V_{RX} when transmitting a "1" or "0" can be found as follows:

$$V_{RX1} = \left(\frac{R_{NRX} + R_F}{R_{NRX} + R_F + R_{ch} + R_{eq} + R_{PTX}} \right) \times V_{DD}, \quad (1)$$

$$V_{RX0} = \left(\frac{R_{ch} + R_{eq} + R_{NTX}}{R_{NTX} + R_F + R_{ch} + R_{eq} + R_{PRX}} \right) \times V_{DD}. \quad (2)$$

where $R_{NRX}, R_{NTX}, R_{PRX}, R_{PTX}$ are the on-resistance of transistors N_{RX}, N_{TX} and P_{RX}, P_{TX} respectively, and R_{ch} is the DC resistance of the channel. Assuming a DC balanced signal at the receiver, the DC voltage at V_{RX} can be found by taking the average of (2) and (1). Note these voltages can be made relatively insensitive to variations in the on-resistance of the NMOS and PMOS transistors by making them smaller than R_F and R_{eq} . The termination R_T and C_T form a low-pass

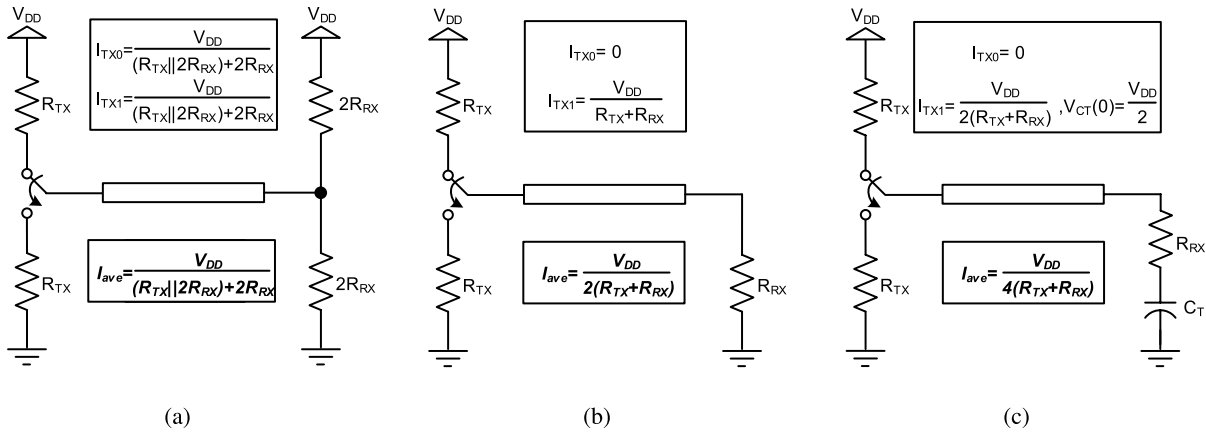


Fig. 4. Different receiver termination schemes in single-ended signaling: (a) Thevenin; (b) parallel; (c) capacitive (assuming C_T is relatively large and hence the voltage across it is unchanging).




Fig. 5. DC current path of signals when transmitting a long sequence of “0” or “1” (C_T is considered an open circuit).

filter of V_{RX} whose DC voltage is stored on capacitor C_T . This voltage, V_{ref} , is used as the reference signal in the receiver sampler which automatically adjusts to variations in input/output impedance, supply voltage and swing without a specific training sequence. However, capacitor C_T must be large enough to avoid large drift on the reference voltage during CIDs. In this work, C_T is 30 pF which translates into a 7.5 ns time constant for the capacitor, assuming $R_{TX} = 200 \Omega$ and $R_{RX} = 50 \Omega$. This means that the reference voltage drifts 10% after 16 CIDs. Moreover, dc balance is required to ensure the voltage at V_{ref} is an appropriate reference. Hence, this method is only effective when the input signal is line coded or adequately scrambled.

To reduce the link power consumption even further, single-sided termination is used. For example, in a voltage-mode driver with a 1 V supply voltage and capacitive termination, using 50 Ω termination in both the transmitter and receiver results in 2.5 mW power consumption. This power can be reduced to only 1 mW if a 50 Ω termination impedance is used on one side and a 200 Ω termination on the other side. Although, singly-terminated links reduce power consumption, they can degrade signal integrity due to reflections. Fig. 6(a)–(c) and (d)–(f) show the simulated received pulse responses and eye diagrams for doubly- and singly-terminated links respectively over a 4 cm organic substrate. In the

presence of channel discontinuities (two PCB vias in these simulations), single-sided termination degrades the received eye diagram significantly as shown in Fig. 6(f), but the eye diagram in Fig. 6(e) still has enough horizontal and vertical eye opening. The main reason is that the matched termination impedance on one side adequately absorbs reflections as long as there are no significant discontinuities in the channel. This is the case in silicon interposers and organic substrates as shown in Fig. 2 since there are no solder balls or PCB vias in the channel.

B. Passive Equalization in Transmitter

Capacitive transmitters for on-die interconnects were introduced in [9], [10] to save power and compensate for interconnect loss at the same time. In [11], a capacitive transmitter in parallel with a voltage-mode transmitter was used to implement pre-emphasis in a single-ended memory transceiver. Similar to these works, as shown in Fig. 5, passive equalization is used in the transmitter here. This contrasts with [12] which used a low-impedance transmitter and a high-impedance receiver for links on silicon interposers. Having a low-impedance termination in the transmitter results in more drive strength and faster slew rate; however it increases the signal swing on the channel, the pad and ESD capacitance. For example, $V_{DD} = 1$ V, $R_{TX} = 50 \Omega$, $R_{RX} = 200 \Omega$, and $C = 500$ fF at 20 Gb/s results in 0.8 V signal swing and 1.6 mW of dynamic power consumption, while switching to $R_{TX} = 200 \Omega$, $R_{RX} = 50 \Omega$ results in 0.2 V signal swing and only 0.1 mW of dynamic power consumption. Therefore, the passive equalizer is used in the transmitter to improve the link energy efficiency. The main drawback of putting the equalizer in the transmitter is the inability to adapt the equalization; however, if the channels are known *a priori*, which is the case in most of the applications in die-to-die communication, the equalizer setting can be set manually. The impact of high-impedance DC termination in the transmitter on the drive strength and slew rate is not significant because the capacitor in the passive equalizer decreases the transmitter impedance at high frequencies.

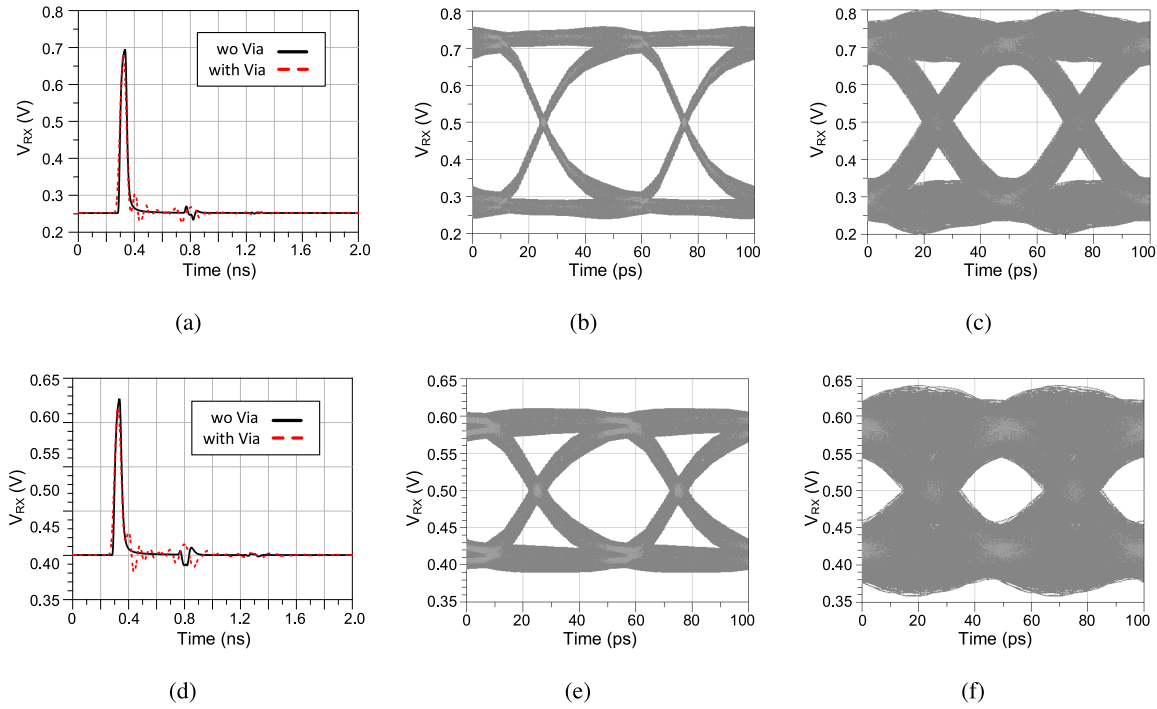


Fig. 6. (a) Simulated pulse responses with double termination without and with vias; simulated received eye diagrams with double termination (b) without vias, (c) with vias; (d) simulated pulse responses with single-sided termination without and with vias; simulated received eye diagrams with single termination (e) without vias, and (f) with vias.

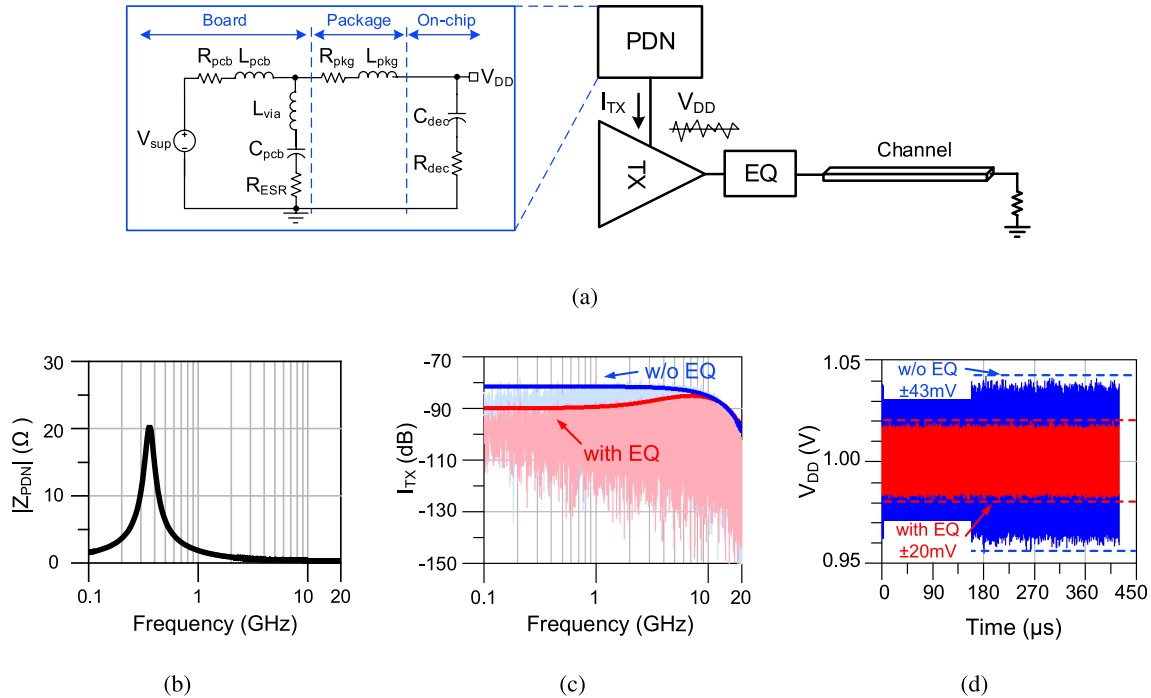


Fig. 7. (a) Power supply noise simulation; (b) PDN impedance magnitude; (c) current spectrum drawn from the power supply with and without passive equalizer assuming PRBS23 data at 20 Gb/s; (d) ripples on the supply voltage with and without passive equalizer.

The passive equalizer also mitigates power supply noise which is one of the main challenges in single-ended signaling. As shown in Fig. 7(b), the inductive and capacitive components of a typical power distribution network (PDN) cause its impedance to resonate at hundreds of MHz. Hence, currents drawn at those frequencies may result in large ripples on the

supply voltage. Using a passive equalizer in the transmitter reduces the supply current spectrum in that critical frequency range as shown in Fig. 7(c). It should be noted that the transmitter strength at high frequencies is similar with and without the passive equalizer because the impedance of the passive equalizer decreases with the increase in frequency.




Fig. 8. Block diagram of transceiver prototype.

As can be seen in Fig. 7(d), using the passive equalizer results in significant reduction in power supply noise for the PDN model in Fig. 7(a).

III. TRANSMITTER/RECEIVER DESIGN

The block diagram of the transceiver prototype that includes three pairs of transmitters and receivers is shown in Fig. 8. In each transmitter, a parallel PRBS generator produces a quarter-rate length- $(2^7 - 1)$ pattern. These signals are then serialized in a 4-to-1 multiplexer. The full rate signal goes through two stages of pre-amplification and then to the channel through the output driver discussed in Section II. The passive equalizer in the transmitter is programmable to provide different amounts of boost or can be bypassed using EQ_BP switch in Fig. 8. R_{eq} is programmable from 155 Ω to 1 k Ω using four switchable resistors, while the on-resistance of the transistors in the output driver is 40 Ω . C_{eq} is programmable from 200 fF to 575 fF in steps of 25 fF. The area of the passive equalizer is 1500 μm^2 . In the receiver, capacitive termination is used to save power and generate the reference voltage. The received signal goes into a 1-to-4 demultiplexer after the pre-amplifier, consisting simply of four time-interleaved flip-flops. Each quarter-rate demultiplexer output can be selected using the corresponding Output Select switch. The selected output is then sent off-chip through a 50 Ω output driver for bit error rate (BER) measurement.

In high-density I/Os, multiple data lanes are usually bundled and one clock signal per bundle is forwarded from the transmitters to the receivers. This simplifies clock alignment

in the receiver by ensuring zero frequency offset between the transmitter and receiver. In [6], the delay through all the data lanes and the clock lane is matched, so there is no need for per lane clock de-skew or continuous phase alignment. In [5], [18], the phase alignment is done in a periodic “round-robin” fashion for all the data lanes and 1/4-UI clock phase alignment per lane is provided to overcome mismatch between data lanes within the bundle. Typically, clock recovery and distribution circuits consume around 50% of the total link power based on the published results in [5], [6], [18]. These circuits are not the focus of this work. Therefore, a half-rate differential clock signal is provided from off-chip and the phase alignment is done manually using off-chip clock synthesizers and on-chip delay lines. After termination, the half-rate clock is distributed to each transmitter and receiver using on-chip transmission lines. In each transmitter and receiver, the half-rate clock is divided into quarter-rate quadrature clock signals.

All latches in the PRBS7 generators and 1-to-4 demultiplexers are double-tail latches [13]. The 4-to-1 multiplexer and pre-driver stages in the transmitter use CMOS logic circuitry to save power compared with CML logic. In the rest of this section, these building blocks are discussed.

A. 4-to-1 Multiplexer

Using higher-order multiplexers (i.e. greater than 2:1) as the final stage of the serializer can lead to lower power consumption when the data rate is high relative to the transistor’s intrinsic speed as it does not require any half-rate latches; however, this has two drawbacks. First, higher-order multiplexers usually exhibit lower bandwidth than 2:1 multiplexers. Second, the phase mismatch between the multiphase clock signals degrades the output signal quality [14]. In [5], [15], CMOS 2-to-1 multiplexers have been used due to these problems. In this work, a 4-to-1 multiplexer is used to reduce the power consumption of clocking circuits. The limited bandwidth of the multiplexer is compensated in the pre-driver stage. A fully symmetric layout for the clock signals and the multiplexer is used. Moreover, a half-rate clock signal is distributed between multiple lanes and is divided using an injection-locked frequency divider in each lane to minimize the phase mismatch between the quadrature clock signals. It should be noted that distributing the clock signal half-rate or quarter-rate consumes a similar amount of power if CMOS inverters can be used in the clock distribution buffers.

Fig. 9(a) shows the 4-to-1 CMOS multiplexer schematic. The pull-down network which serves as both logical-AND and latch is a modified CMOS version of the CML 4-to-1 multiplexer presented in [14]. The multiplexer only requires single-ended signals for its operation and is comprised of four pull-up/pull-down cells. The timing diagram of the first cell is shown in Fig. 9(b). As long as Clk0 is high (Clk270 is low), the “Np” and “Nn” nodes are charged to VDD and discharged to ground respectively. When Clk90 goes low, the “Np” and “Nn” nodes will track the “D0” and its complement respectively. At the rising edge of Clk0, the input data is sampled on the parasitic capacitance on the “Np” and “Nn” nodes. These nodes hold their values until Clk90 goes high. During the time




Fig. 9. 4-to-1 CMOS multiplexer: (a) schematic; timing diagram of (b) one pull-up/pull-down cell; (c) 4-to-1 multiplexer.

window when Clk0 is high and Clk90 is low, the cell 1 output $M1_{out1}$ is equal to the input data $D0$. When Clk0 is low or Clk90 is high, $M1_{out1}$ is disconnected from M_{out} . Therefore, each pull-up/pull-down cell controls the output signal during the time when its clock (Clk0 in cell 1) is high and its delayed clock (Clk90 in cell 1) is low. By connecting the appropriate clock phases to each of the pull-up and pull-down cells, the output is driven by only one cell at a time and each clock phase sees the same capacitive load. Fig. 9(c) shows the timing diagram of the 4-to-1 multiplexer.

To increase the timing margin for the sampling clock, there is a 180 phase difference between D0-D1 and D2-D3. In [14], this phase difference is introduced by using different clock phases for the multiplexers in the preceding 8:4 serializer. In the prototype, since there is no 8:4 serializer this is done by retiming two of the four inputs prior to the multiplexer. The proposed 4-to-1 CMOS multiplexer consumes less power than its CML counterpart but has significant parasitic capacitance on its output from both the pull-up and pull-down networks preventing the multiplexer from operating at 20 Gb/s. In the next section, a bandwidth enhancement technique is presented to alleviate this problem.

B. Transmitter Pre-driver

Fig. 10 illustrates the proposed CMOS bandwidth enhancement technique. As shown in Fig. 10(a), the pre-driver circuit is comprised of a resistive feedback inverter with source degeneration followed by a conventional CMOS inverter. In the absence of this technique, the bandwidth of the multiplexer is $1/(R_{mux} \cdot C_{mux})$ resulting in inter-symbol interference at 20 Gb/s as shown in Fig. 10(b). This increases the jitter in the transmitter and leads to performance degradation. Using a feedback resistor (R_{pdf}) in the next stage improves the bandwidth of the multiplexer by pushing the pole location to $1/((R_{mux} || R_{pdf}) \cdot C_{mux})$, assuming R_{pdf} is larger than on-resistance of transistors N1 and P1, and the capacitance

at the output of the inverter with resistive feedback (PD_M) is negligible. The simulated eye diagram at PD_M after adding the feedback resistor is shown in Fig. 10(c).

Adding source degeneration capacitors (C_{pds}) and resistors (R_{pds}) to the sources of both P1 and N1 creates a zero at $1/((R_{on} || R_{pds}) \cdot C_{pds})$ which can be used to further improve the bandwidth. As shown in Fig. 10(d), data-dependent jitter is eliminated at PD_M after applying both source degeneration and resistive feedback in the pre-driver. In the prototype, the switches En are included to disable both the feedback resistor and source degeneration and show the effectiveness of this technique.

As can be seen in 10(c)–(d), the signals at PD_M do not swing rail-to-rail after applying the proposed technique so a conventional CMOS inverter is needed to bring the signals back to rail-to-rail. If used to achieve large high-frequency boost, the circuit's reduced output swing makes it more sensitive to power supply noise. Therefore, this technique is only used to introduce between approximately 0 to 5 dB of high-frequency boost.

IV. SILICON INTERPOSER

To implement high density interconnects, silicon interposers have been used in electrical and optical transceivers [5], [16]. In this work, a very low-cost (0.35 μm) standard CMOS technology is used to build an interposer due to its ready availability. In addition to providing the high-speed interconnect, the interposer can also be used to realize metal-metal decoupling capacitors for the power supplies.

Fig. 11 shows the stripline configuration of the channel in a standard CMOS interposer (with 0.64 μm -thick Al interconnect) compared to the dedicated interposer technology in [5] (with 3 μm copper interconnect). The dedicated (copper back-end) interposer technology exhibits lower channel insertion loss at DC and high frequencies. As shown in Fig. 12(a), channel insertion loss at 10 GHz for a 3 cm channel with

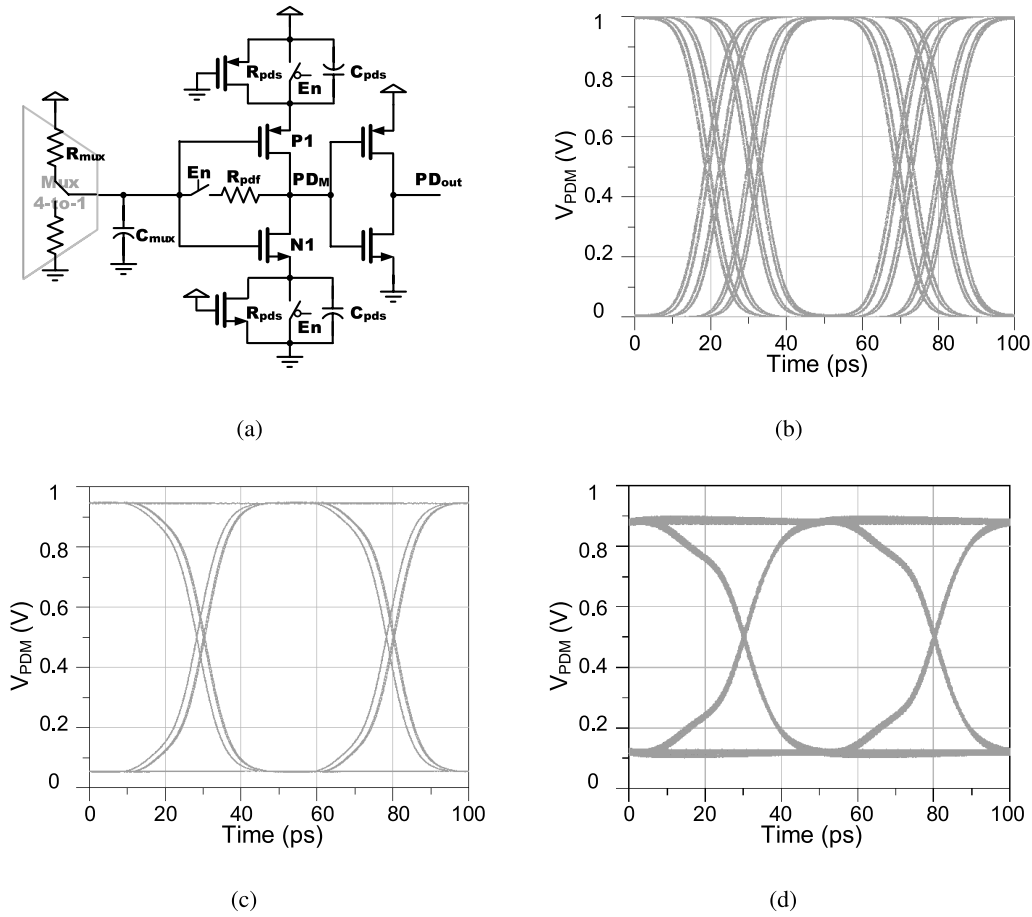


Fig. 10. (a) Pre-driver bandwidth enhancement circuit, simulation results at 20 Gb/s; (b) without bandwidth enhancement; (c) with resistive feedback; (d) with resistive feedback and source degeneration.




Fig. 11. The stripline configuration in (a) standard CMOS silicon interposer; (b) dedicated silicon interposer in [5].

the thick copper interconnect is similar to a 3.5 mm channel on the low-cost standard interposer. As shown in Fig. 12(b), the return loss of the channels on the copper interconnect is constant over a wider range of frequencies compared to the standard interposer. This is mainly because the interconnects on the standard interposer are RC dominant at low frequencies.

To achieve 50 Ω characteristic impedance using the stack-up in Fig. 11(a), the trace width must be 1.4 μm resulting in 7.5 dB of DC loss as shown in Fig. 12(a). For the same DC loss as in [5], the trace width may be increased to 4.2 μm but this leads to poor return loss. Consequently, a tapered transmission line is used where the trace width is increased

from 1.4 μm to 4.2 μm and then back to 1.4 μm in steps of 0.4 μm. The tapered transmission line exhibits less low frequency insertion loss compared to a 1.4 μm wide transmission line while still maintaining return loss below 10 dB above 5 GHz. In other words, the complex characteristic impedance of the interconnects becomes mostly real and the interconnects perform as low-loss transmission lines. So the link needs to be properly terminated to avoid reflections.

One of the main challenges in high-density I/Os is crosstalk. In [17], the measured far-end crosstalk (FEXT) and near-end crosstalk (NEXT) of differential channels on the dedicated silicon interposer are reported to be less than 30 dB

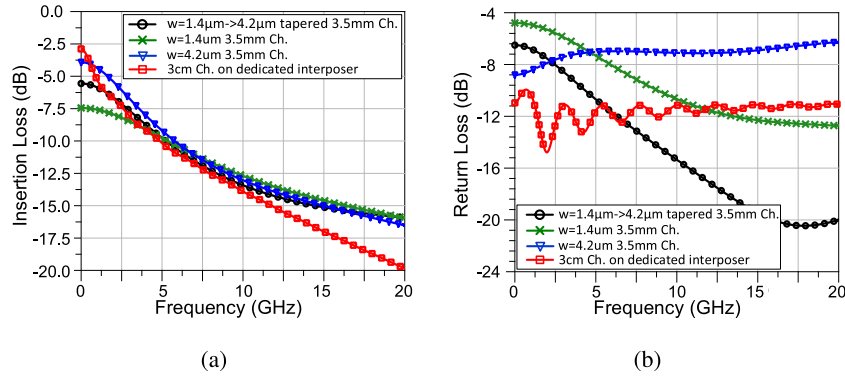


Fig. 12. Comparison between the 2D electromagnetic field solver simulation results of 3.5 mm interconnects on standard silicon interposer and 3 cm interconnects on dedicated silicon interposer: (a) insertion loss; (b) return loss.




Fig. 13. 2D electromagnetic field solver crosstalk simulation results for different 3.5 mm interconnect spacings: (a) near-end crosstalk (NEXT); (b) far-end crosstalk (FEXT).

and 37 dB without and with ground shielding respectively. Fig. 13 shows the FEXT and NEXT simulation results for a 3.5 mm interconnect on our standard interposer with different channel spacings compared to a 3 cm interconnect on the dedicated interposer with 6 μm spacing. As can be seen, the crosstalk in the dedicated interposer is worse than the numbers reported in [17] since the channels are single-ended in the simulations shown in Fig. 13. The crosstalk in the standard interposer is negligible once the channel spacing is above 4 μm . Both interposers use stripline configuration which is more immune to crosstalk compared to microstrip configuration, but the standard interposer exhibits lower crosstalk. This is mainly because of the thin traces and their RC dominated behavior. The channel spacing is chosen to be 6 μm . This leaves coupling between neighboring bumps/pads as the main source of crosstalk. The channel pitch available in the low-cost standard silicon interposer is 10.2 μm which is comparable to the 8 μm and 22 μm channel pitches used in [17].

V. EXPERIMENTAL RESULTS

Fig. 14 shows the packaged prototype, where two identical transceiver dies are flip-chip mounted onto a standard CMOS silicon interposer using gold stud bumping. The bump pitch is 100 μm . The interposer is then wire-bonded to a high-speed PCB. As described in Section III, there are three identical transmitter/receiver pairs on each transceiver die. All three transmitters on Die B (TX1B-3B) are connected to the

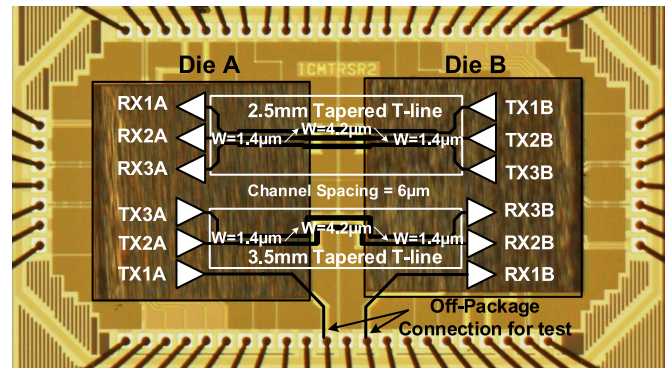


Fig. 14. The system prototype package where two transceiver dies are flip-chip mounted onto a silicon interposer with 10.2 μm channel pitch.

receivers on Die A (RX1A-3A) through a 2.5 mm interconnect. Also TX2A and TX3A are connected to RX2B and RX3B through a 3.5 mm interconnect. TX1A and RX1B are brought off-package to measure the performance of the transmitter and receiver independently.

The interposer prototype was fabricated in AMS standard 0.35 μm CMOS technology. Fig. 15 shows the measured insertion/return loss of the interposer along with their loss at one-half the maximum bit rate tested. The measurements were done by wafer probing and using a vector network analyzer (VNA). Moreover, a sampling oscilloscope was used

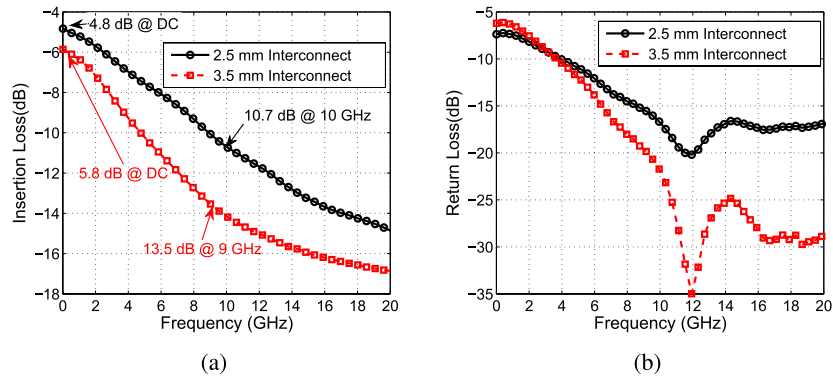


Fig. 15. Measurement results of 2.5 mm and 3.5 mm interconnects: (a) insertion loss; (b) return loss.

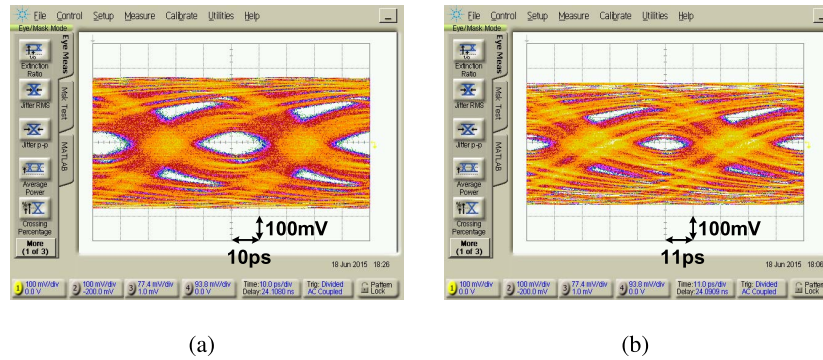


Fig. 16. Measured interposer eye diagrams with 1.7 V launch swing (a) at 20 Gb/s over 2.5 mm interconnect, and (b) at 18 Gb/s over 3.5 mm interconnect.

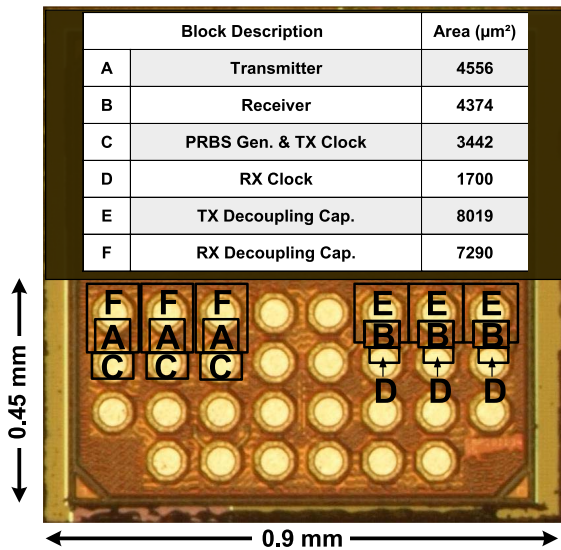


Fig. 17. Die photo and area breakdown of the 28 nm FD-SOI CMOS transceiver prototype.

to measure the eye diagrams after the channels on the interposer. Fig. 16 shows the eye diagrams when driving the 2.5 mm and 3.5 mm interconnects through wafer probes with test equipment at 20 Gb/s and 18 Gb/s, respectively.

The transceiver prototype was fabricated in 28 nm STM FD-SOI CMOS. The die photo and area breakdown in the transceiver is shown in Fig. 17. Each transmitter and receiver occupies $67.5 \times 67.5 \mu\text{m}^2$ and $81 \times 54 \mu\text{m}^2$ of silicon

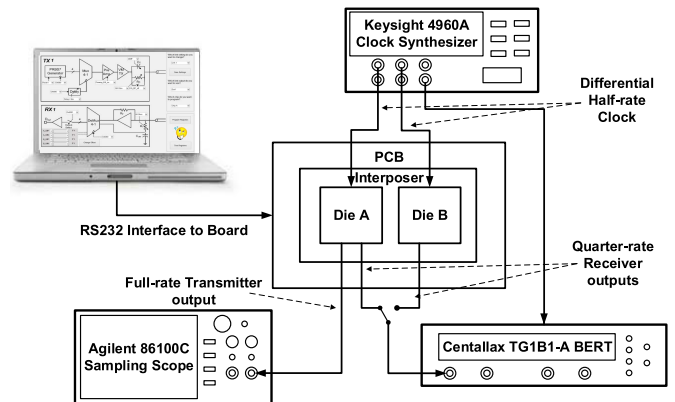


Fig. 18. Measurement setup.

area, respectively. Therefore, both transmitter and receiver including all their decoupling capacitors and clocking circuits easily fit under 2 pads even with our reduced pitch, so that for example a complete transceiver could fit under 4 pads for V_{DD} , Gnd , TX_{out} , RX_{in} , an important feature for high-density IOs.

Fig. 18 shows the measurement setup used to characterize the link performance. The interposer is wire-bonded on a high speed board where all the high-speed connectors are mounted. A DC board is used to provide the power supplies and low-speed control signals for both transceivers through a laptop GUI using a RS232 interface. The half-rate differential clock signals for both transmitter and receiver are provided from a Keysight 4960A clock synthesizer. The full-rate

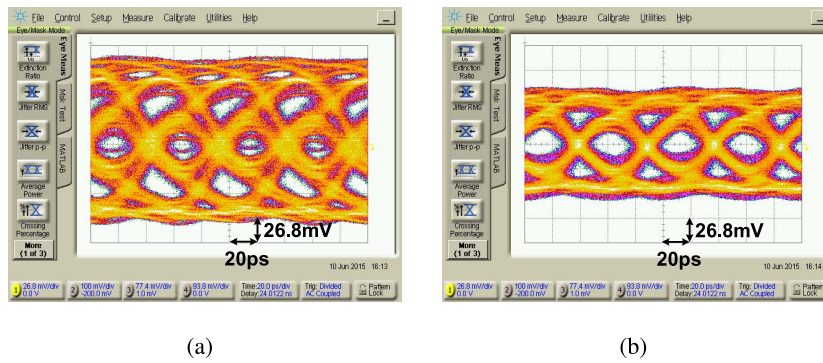


Fig. 19. The transmitter output eye diagram taken off-package with 20 Gb/s PRBS7: (a) bypassing the transmitter passive equalizer; (b) with passive equalizer.

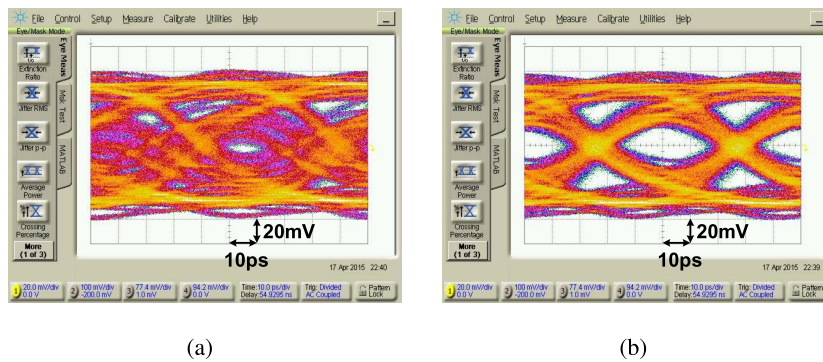


Fig. 20. The transmitter output eye diagrams taken off-package with 20 Gb/s PRBS7: (a) without pre-driver bandwidth enhancement; (b) with pre-driver bandwidth enhancement.

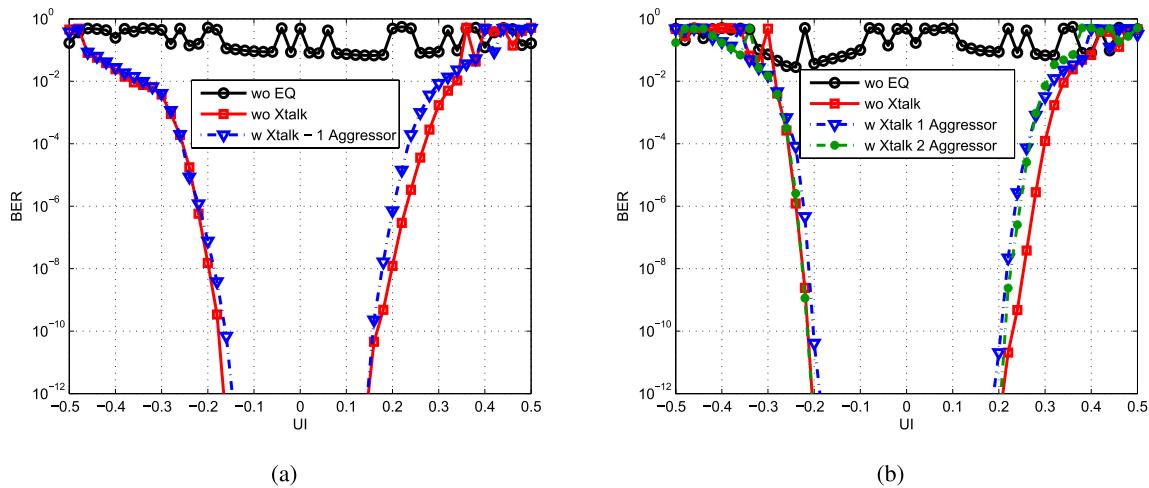


Fig. 21. Bathtub curves with and without crosstalk and passive equalizer with PRBS7: (a) 3.5 mm at 18 Gb/s; (b) 2.5 mm at 20 Gb/s.

transmitter output (TX1A) is sent to a Keysight 86100C sampling oscilloscope to look at the eye diagrams. The quarter-rate received signal is sent to a Centellax TG1B1-A BERT for BER measurements.

The transmitter output (TX1A) eye diagrams are first illustrated. It must be noted that the transmitter signal in these measurements goes through a shorter interconnect on the interposer but package discontinuities such as wirebonds and a PCB trace to an SMA connector degrade the eye performance. Therefore, the eye quality does not represent the link performance; the only purpose of these eye diagrams is

to illustrate the effectiveness of the passive equalizer and pre-driver bandwidth enhancement techniques.

Fig. 19 shows the TX1A output eye diagrams at 20 Gb/s without and with the transmitter passive equalizer. Also, the TX1A output eye diagrams without and with the pre-driver bandwidth enhancement technique are shown in Fig. 20, illustrating its effectiveness to compensate for the 4:1 multiplexer’s limited bandwidth.

Fig. 21 shows the measured bathtub curves of the received signal on both 2.5 mm and 3.5 mm interconnects with and without passive equalizer and crosstalk. As can be seen in this

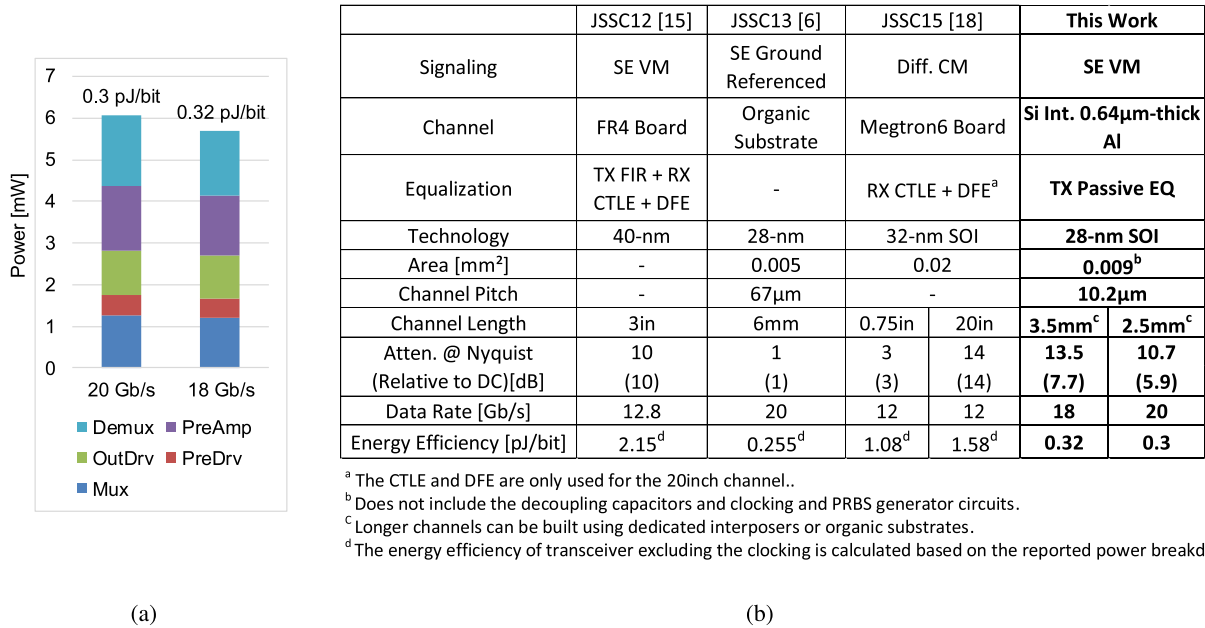


Fig. 22. (a) Power breakdown of the transceiver at 20 Gb/s and 18 Gb/s. (b) Performance summary comparison.

figure, the performance degradation due to crosstalk is negligible as expected. As shown in Fig. 21(a), when one aggressor is present, the eye opening is 0.29 UI with 10^{-12} BER at 18 Gb/s over the 3.5 mm interconnect with 5.8 dB and 13.5 dB loss at DC and Nyquist frequency respectively. Also, when two aggressors are present, the eye opening is 0.38 UI with 10^{-12} BER at 20 Gb/s over the 2.5 mm interconnect with 4.8 dB and 10.7 dB loss at DC and Nyquist frequency.

The power breakdown of the proposed transceiver is illustrated in Fig. 22(a). At 20 Gb/s over the 2.5 mm interconnect, the transmitter and receiver power consumption are 2.8 mW and 3.3 mW, respectively. At 18 Gb/s over the 3.5 mm interconnect, the transmitter and receiver consume 2.7 mW and 3.0 mW of power, respectively.

Fig. 22(b) compares the performance of the proposed transceiver with recently published papers for short reach applications, where the channel loss is up to 10 dB [6], [15], [18]. The energy efficiency of the proposed transceiver at 20 Gb/s is 0.3 pJ/bit including the 4-to-1 multiplexer, transmitter pre-driver, output driver, receiver front-end, pre-amplifier, and 1-to-4 demultiplexer. The power for the clocking circuits are not included in the power numbers, but is also excluded from other works making for a fair comparison. Although [6] has better energy efficiency than this work, it includes no equalization. It is also worth mentioning that the length of the interconnects in this work are shorter because of the low-cost and higher-loss thin Al interconnect; hence, the interconnect length could be extended using a dedicated silicon interposer technology or an organic packaging substrate with lower loss.

VI. CONCLUSION

A 0.3 pJ/bit parallel interface has been introduced for die-to-die interconnect. The implemented prototype includes a standard 0.35 μ m CMOS silicon interposer and two 28 nm

CMOS transceivers. Different techniques in the transceiver design such as single-sided termination, capacitive termination in the receiver, and passive equalization in the transmitter have been employed to reduce the power consumption of the link front-end. Moreover, CMOS logic style circuits have been used in the transmitter to further decrease the power consumption. Single-ended signaling has been used to improve the link bandwidth per pin/wire. The proposed parallel interface consumes 6.1 mW power at 20 Gb/s over a 2.5 mm interconnect with 10.7 dB loss at Nyquist frequency and 4.8 dB loss at DC and it is immune to crosstalk noise. The excellent energy and pin efficiency of the proposed parallel interface and its ability to compensate for 10 dB of insertion loss make it a suitable candidate for die-to-die communication over silicon interposers or organic substrates.

ACKNOWLEDGMENT

The authors would like to acknowledge Huawei Canada for financial support and technical discussions, and Canadian Microelectronics Company (CMC) for fabricating and packaging the chips, provision of equipment, and CAD tools.

REFERENCES

- [1] D. Friedman, "Subcommittee wireline—2015 trends," in *IEEE ISSCC Dig. Tech. Papers*, 2015, pp. 88–90.
- [2] Hybrid Memory Cube Consortium. *Hybrid Memory Cube Specification 2.0*. [Online]. Available: <http://hybridmemorycube.org/>
- [3] High Bandwidth Memory (HBM) DRAM. [Online]. Available: <https://www.jedec.org/standards-documents/results/jesd235>
- [4] W. Huang, K. Rajamani, M. R. Stan, and K. Skadron, "Scaling with design constraints: Predicting the future of big chips," *IEEE Micro*, vol. 31, no. 4, pp. 16–29, Jul. 2011.
- [5] T. O. Dickson *et al.*, "An 8 \times 10-Gb/s source-synchronous I/O system based on high-density silicon carrier interconnects," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 884–896, Apr. 2012.

- [6] J. W. Poulton *et al.*, "A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short-reach serial link in 28 nm CMOS for advanced packaging applications," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3206–3218, Dec. 2013.
- [7] H. Lee *et al.*, "A 16.8 Gbps/channel single-ended transceiver in 65 nm CMOS for SiP based DRAM interface on Si-carrier channel," *IEEE J. Solid-State Circuits*, vol. 50, no. 11, pp. 2613–2624, Nov. 2015.
- [8] M. Bucher *et al.*, "A 6.4-Gb/s near-ground single-ended transceiver for dual-rank DIMM memory interface systems," *IEEE J. Solid-State Circuits*, vol. 49, no. 1, pp. 127–139, Jan. 2014.
- [9] E. Mensink, D. Schinkel, E. A. M. Klumperink, E. van Tuijl, and B. Nauta, "Power efficient gigabit communication over capacitively driven RC-limited on-chip interconnects," *IEEE J. Solid-State Circuits*, vol. 45, no. 2, pp. 447–457, Feb. 2010.
- [10] R. Ho *et al.*, "High speed and low energy capacitively driven on-chip wires," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 52–60, Jan. 2008.
- [11] H. Partovi *et al.*, "Single-ended transceiver design techniques for 5.33 Gb/s graphics applications," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 136–137.
- [12] B. Kim, Y. Liu, T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 10-Gb/s compact low-power serial I/O with DFE-IIR equalization in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3526–3538, Dec. 2009.
- [13] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18 ps setup+hold time," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 314–605.
- [14] A. A. Hafez, M. S. Chen, and C. K. K. Yang, "A 32–48 Gb/s serializing transmitter using multiphase serialization in 65 nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 3, pp. 763–775, Mar. 2015.
- [15] A. Amirkhany *et al.*, "A 12.8-Gb/s/link tri-modal single-ended memory interface," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 911–925, Apr. 2012.
- [16] H. Morita *et al.*, "A 12×5 two-dimensional optical I/O array for 600 Gb/s chip-to-chip interconnect in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 140–141.
- [17] X. Gu *et al.*, "High-density silicon carrier transmission line design for chip-to-chip interconnects," in *Proc. IEEE 20th Conf. Elect. Perform. Electron. Packag. Syst. (EPEPS)*, Oct. 2011, pp. 27–30.
- [18] T. O. Dickson *et al.*, "A 1.4 pJ/bit power-scalable 16 × 12 Gb/s source-synchronous I/O with DFE receiver in 32 nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 8, pp. 1917–1931, Aug. 2015.



Behzad Dehlaghi (S'08) received the B.Sc. degree from University of Tehran, Tehran, Iran, in 2009, and the M.Sc. degree from University of Calgary, AB, Canada, in 2012, both in electrical engineering. Since 2013, he has been working towards Ph.D. degree in electrical engineering at the University of Toronto, Toronto, ON, Canada.

During his M.Sc. studies, he worked on on-chip circuits for jitter measurement and signal capture with sub-picosecond resolution. From February 2012 to December 2012, he worked at Semtech Canada Corporation as an Analog Design Engineer where he was involved in the development and modeling of high-speed transceivers. His research interests include on-chip measurement circuits and high-speed chip-to-chip communication.



Anthony Chan Carusone (S'96–M'02–SM'08) received the Ph.D. degree from the University of Toronto, Toronto, ON, Canada, in 2002 and has since been with the Department of Electrical and Computer Engineering at the University of Toronto where he is currently a Professor. He is also an occasional consultant to industry in the areas of integrated circuit design, clocking, and digital communication.

Prof. Chan Carusone has co-authored over 90 conference and journal papers on integrated circuit design, including the Best Student Papers at the 2007, 2008 and 2011 Custom Integrated Circuits Conferences, the Best Invited Paper at the 2010 Custom Integrated Circuits Conference, the Best Paper at the 2005 Compound Semiconductor Integrated Circuits Symposium, and the Best Young Scientist Paper at the 2014 European Solid-State Circuits Conference. He authored, along with David Johns and Ken Martin, the 2nd edition of the classic textbook *Analog Integrated Circuit Design*. He was Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS in 2009, and has served on the technical program committee for the IEEE Custom Integrated Circuits Conference and the VLSI Circuits Symposium. He currently serves on the editorial board of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, as a member of the Technical Program Committee of the International Solid-State Circuits Conference, and as a Distinguished Lecturer for the IEEE Solid-State Circuits Society.