A 0.47–0.66 pJ/bit, 4.8–8 Gb/s I/O Transceiver in 65 nm CMOS

Young-Hoon Song, *Student Member, IEEE*, Rui Bai, *Student Member, IEEE*, Kangmin Hu, *Member, IEEE*, Hae-Woong Yang, *Student Member, IEEE*, Patrick Yin Chiang, *Member, IEEE*, and Samuel Palermo, *Member, IEEE*

***Abstract*—A low-power forwarded-clock I/O transceiver architecture is presented that employs a high degree of output/input multiplexing, supply-voltage scaling with data rate, and low-voltage circuit techniques to enable low-power operation. The transmitter utilizes a 4:1 output multiplexing voltage-mode driver along with 4-phase clocking that is efficiently generated from a passive poly-phase filter. The output driver voltage swing is accurately controlled from 100–200 $\mathrm{mV_{ppd}}$ using a low-voltage pseudo-differential regulator that employs a partial negative-resistance load for improved low frequency gain. 1:8 input de-multiplexing is performed at the receiver equalizer output with 8 parallel input samplers clocked from an 8-phase injection-locked oscillator that provides more than 1UI de-skew range. In the transmitter clocking circuitry, per-phase duty-cycle and phase-spacing adjustment is implemented to allow adequate timing margins at low operating voltages. Fabricated in a general purpose 65 nm CMOS process, the transceiver achieves 4.8–8 Gb/s at 0.47–0.66 pJ/b energy efficiency for $\mathrm{V_{DD}} = 0.6$–0.8 V.**

***Index Terms*—High-speed I/O, injection-locked oscillator, low-power, low-voltage regulator, poly-phase filter, transceiver, voltage-mode driver.**

I. Introduction

TOTAL I/O bandwidth demand is growing in high-performance systems due to the emergence of many-core microprocessors and in mobile devices in order to support the next generation of multi-media features. High-speed serial I/O energy efficiency must improve in order to enable continued scaling of these parallel computing platforms in applications ranging from data centers to smart mobile devices.

Significant I/O energy efficiency improvements necessitate both advances in electrical channel technologies and circuit techniques in order to reduce complexity and power consumption. Examples of advanced inter-chip physical interfaces include high-density interconnect and Flex cable bridges, which allow operation at data rates near 10 Gb/s while only requiring modest equalization [1]. An I/O architecture that reduces clocking circuit complexity, while also allowing for wide-bandwidth jitter tracking, is a forwarded-clock system where a clock signal is transmitted in parallel with multiple data channels Fig. 1 [2], [3]. Furthermore, low-power transceivers often incorporate voltage-mode transmit drivers, as these output stages have the potential to consume one-quarter of the power compared to current-mode drivers [4].

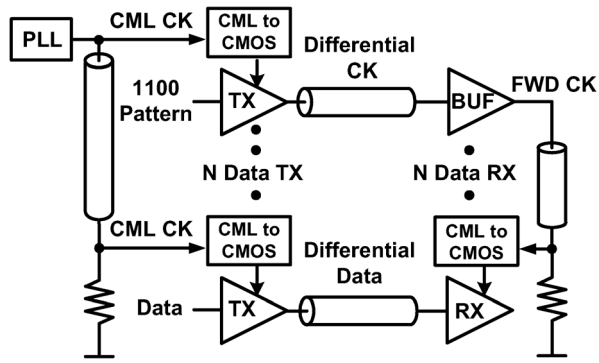

Fig. 1. A multi-data-channel forwarded-clock I/O architecture.

Further improvements in energy-efficiency are possible through reduction of the supply voltage $\mathrm{V_{DD}}$. Previously, this has enabled excellent energy/computation for digital systems [5] due to the exponential dependence of power on $\mathrm{V_{DD}}$. Leveraging supply scaling to improve energy efficiency motivates I/O architectures that employ a high level of output/input multiplexing, as this allows for the parallel transmit and receive segments to operate at lower voltages [6]. However, challenges exist in the design of an efficient output-multiplexed voltage-mode driver due to the relatively large driver transistor sizes required for output impedance control, as well as the reduced supply headroom for the output stage regulator. Furthermore, widespread adoption of low-$\mathrm{V_{DD}}$ transceivers has been limited due to questions regarding robust operation and severe sensitivity to process variations. In particular, the generation of precise multi-phase clocks and the ability to compensate for circuit mismatch is an issue both at the transmitter and receiver.

This paper describes a low-voltage forwarded-clock I/O architecture developed in 65 nm CMOS that is capable of 4.8–8 Gb/s operation while achieving an energy efficiency of 0.47 pJ/bit–0.66 pJ/bit. Section II discusses key circuit trade-offs associated with supply-scaling and multiplexing factor choice at both the transmitter and receiver. The proposed transmitter, which to the authors' knowledge, is the first to implement a

Manuscript received February 06, 2012; revised November 03, 2012; accepted January 27, 2013. Date of publication March 12, 2013; date of current version April 19, 2013. This paper was approved by Associate Editor Kenneth Shepard. This work was supported by the Semiconductor Research Corporation (SRC) under grant 1836.060, the Department of Energy Early Career program, and a gift from the Intel Labs Academic Research Office Wireline Signaling Center.

Y.-H. Song, H.-W. Yang, and S. Palermo are with Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA.

R. Bai, K. Hu, and P. Y. Chiang are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331 USA.

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.2013.2249812

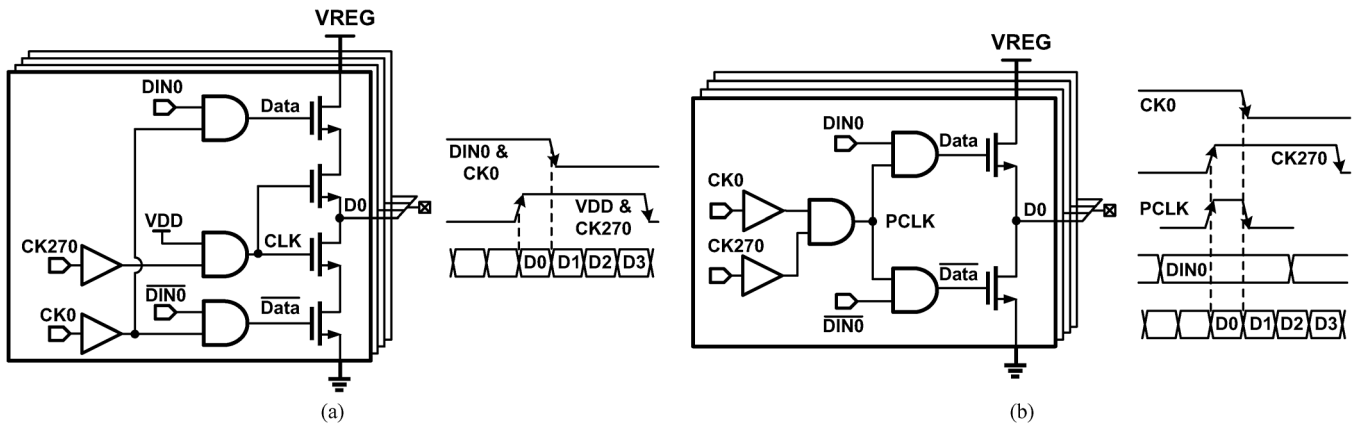

Fig. 2. Output multiplexing approaches for voltage-mode drivers: (a) producing an output data pulse with two-transistor output segments, (b) producing an output data pulse with a pulse-clock and a single-transistor output segment.

level-shifting pulse-clock pre-driver to reduce the transistor size and stack count in a voltage-mode output-multiplexing driver is detailed in Section III. Also discussed in this section is the use of a passive poly-phase filter for transmitter quadrate clock generation, which has been shown in previous work [7] as an efficient technique to generate quadrature receiver-side clocks. Section IV presents the 1:8 input de-multiplexing receiver which employs eight parallel input samplers clocked from an 8-phase injection-locked oscillator that provides more than 1UI de-skew range and utilizes AC-coupling injection for improved phase uniformity relative to transconductance injection [2]. The single-data-channel transceiver experimental results are summarized in Section V and a discussion on scaling this architecture to higher per-pin data rates is included in Section VI. Finally, Section VII concludes the paper.

II. Transceiver Architecture Considerations

Utilizing circuit parallelism in I/O transceivers allows for potential power savings, as the parallel transmit and receive segments operate at lower frequencies and potentially lower voltages [6]. Unfortunately, challenges exist in generating power-efficient multiple-phase clocks and maintaining critical circuit transmitter/receiver circuit bandwidths while operating under low voltage. This section analyzes the trade-offs associated with supply-scaling and multiplexing factor choices at both the transmitter and receiver.

A. Transmitter

Voltage-mode output stages are desired in low-power transmitter architectures due to the potential for significant current savings for a given output voltage swing. It is possible to implement output multiplexing in current-mode drivers through multiple two-transistor current-switch segments controlled by two overlapping clock signals and the data, thus avoiding any full data-rate signals until the final pad outputs Fig. 2(a) [6], [8]. Unfortunately, utilizing this approach in voltage-mode driver results in large output transistors in order to maintain proper channel impedance termination to minimize reflection-induced intersymbol interference and allow predictable transmit output swing levels. Driving these large output transistors increases dynamic power consumption and the series transistor combination degrades the output signal edge rates. Another output multiplexing approach suitable for a voltage-mode driver involves combining a one unit interval (UI) pulse-clock with the data before the output switch transistor, allowing for only one single-transistor output segment to be activated at a time Fig. 2(b). Hence, impedance control is achieved using smaller output transistors, resulting in reduced pre-driver power consumption and improved output signal edge rates. This pulse-clock output multiplexing scheme is utilized in the voltage-mode driver presented in this work.

The optimal output multiplexing ratio, with respect to power efficiency, is a function of both the minimum swing required to maintain the output eye margins and the complexity associated with the generation of precise multiple-phase clocks. Fig. 3 compares three 8 Gb/s transmitters that utilize output-multiplexing factors of 1:1 (multiplexing before the output driver), 4:1, and 8:1, respectively. The transmitters leverage supply scaling in the clock generation and serialization while the output stage is powered from a low-voltage regulator, discussed in Section III, which is capable of operating from a fixed 0.65 V supply. In order to avoid the challenges associated with global multiple-phase clock distribution in a multi-channel I/O system, all these topologies utilize a low-swing global differential clock distribution, with multiple-clock phases generated locally. The 1:1 multiplexing transmitter is a half-rate architecture [9]–[12] that utilizes a 2:1 CMOS mux before the output stage which is switched by two-phases of a 4 GHz clock generated by the local CML-to-CMOS clock buffer circuitry. For the 4:1 multiplexing transmitter, a 2 GHz low-swing global clock passes through a passive poly-phase filter to produce four clock phases, which are then converted to CMOS levels to actuate the pulse-clock predriver. The eight clock phases required for the 8:1 multiplexing transmitter are produced with a local injection-locked oscillator (ILO) locked to a 1 GHz low-swing global clock input.

Schematic simulation results are presented in Fig. 4(a), which compares the 8 Gb/s deterministic jitter (DJ) of the three transmitters driving an ideal channel as a function of the supply voltage. The 1:1 input multiplexing transmitter's DJ increases rapidly as the supply is reduced near 0.6 V due to degraded timing margin in the 2:1 CMOS multiplexer that switches at 4 GHz, while both the 4:1 and 8:1 output multiplexing designs

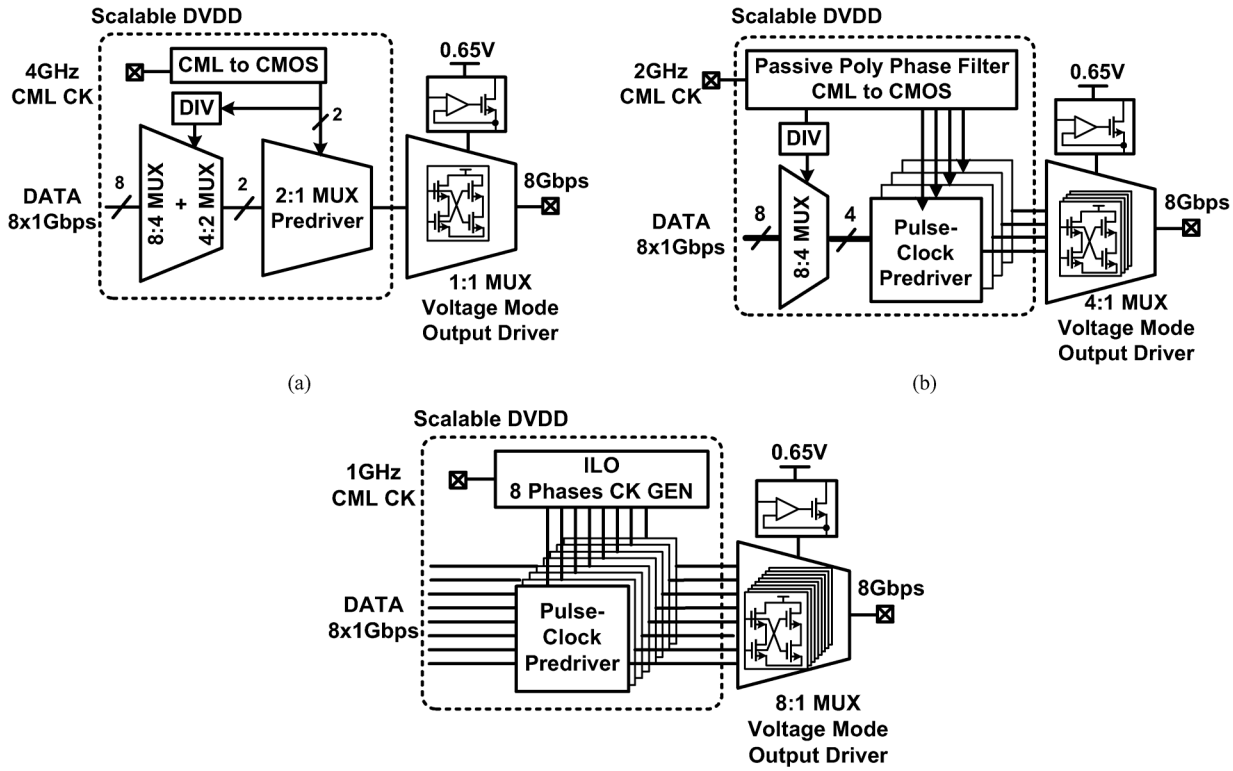

Fig. 3. Transmitter architectures with different output multiplexing factors: (a) 1:1, (b) 4:1, (c) 8:1.

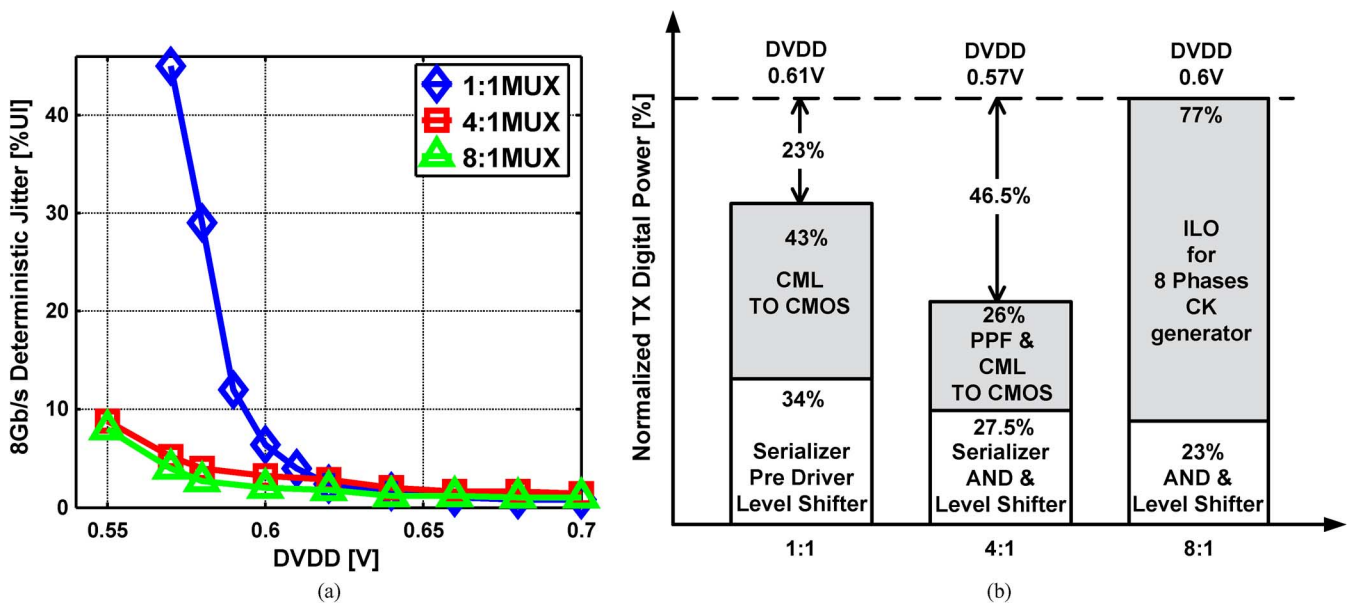

Fig. 4. Simulated 8 Gb/s transmitter performance with varying output multiplexing factors: (a) deterministic jitter versus supply voltage, (b) dynamic power consumption.

display similar performance and operate with reasonable DJ at lower voltages.

Fig. 4(b) compares the dynamic power consumption of the three transmitters normalized to the highest-power 8:1 architecture. Here the transmitter supply is set based on two constraints of 5% output DJ and acceptable output phase mismatch across Monte Carlo simulations. While the 8:1 transmitter is capable of less than 5% DJ at a supply lower than 0.6 V, the ILO displays excessive phase variation at these low voltages. Overall, the 4:1 output multiplexing architecture displays the best power consumption due to the superior timing margins relative to the 1:1 transmitter and reduced sensitivity to multi-phase clock generation enabled through the two-stage passive poly-phase filter. Hence, the 4:1 architecture is chosen and is discussed in detail in Section III.

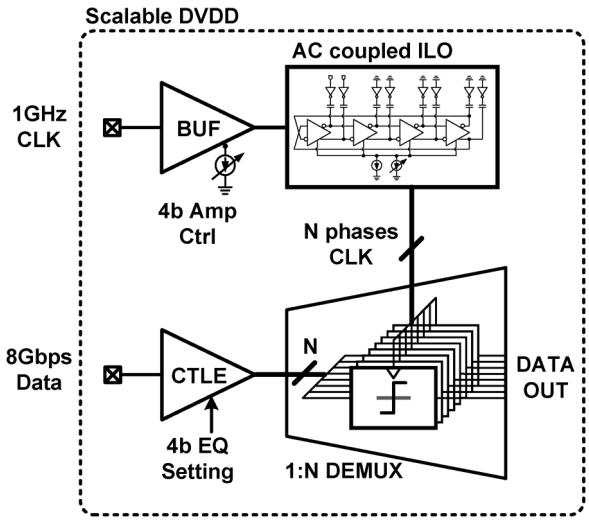

Fig. 5. A forwarded-clock 1:N receiver architecture.

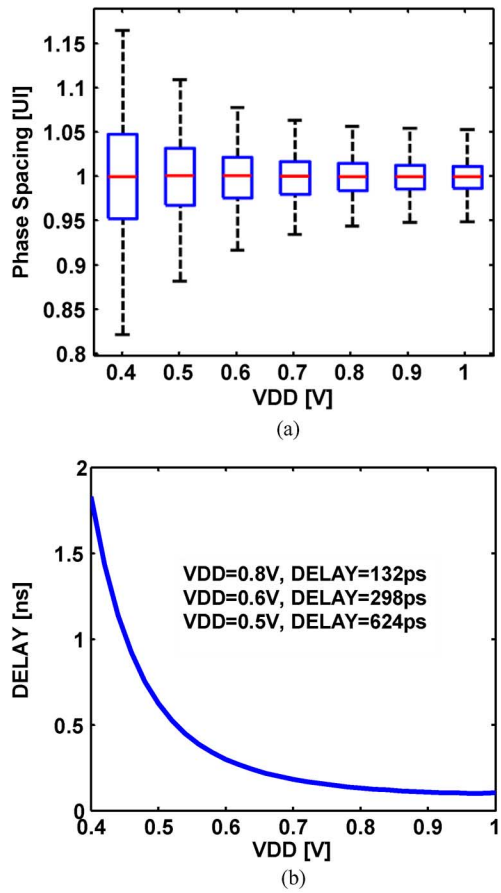

Fig. 6. Key receiver circuitry simulated performance versus supply voltage: (a) ring oscillator phase variation, (b) quantizer delay.

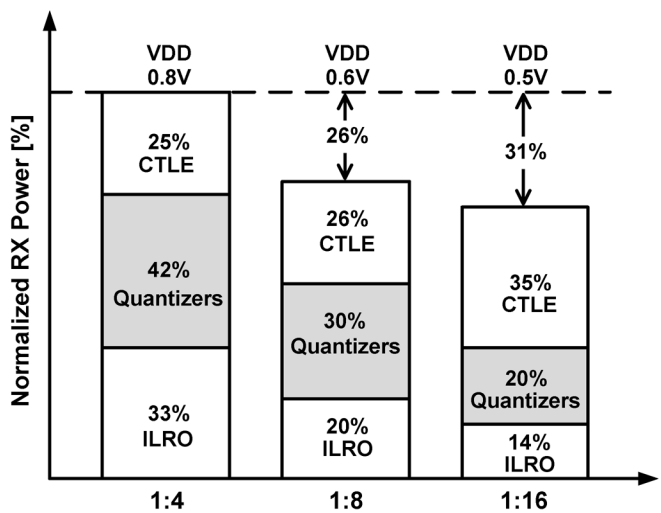

Fig. 7. Receiver power consumption versus de-multiplexing factor.

B. Receiver

At the receiver, the optimal input de-multiplexing ratio, in terms of power efficiency, is a function of the minimum voltage required to produce precise multi-phase clocks while maintaining adequate circuit speed. An input continuous-time linear equalizer (CTLE), consisting of a RC-degenerated differential amplifier, is used to compensate for the channel loss. Fig. 5. shows a high-level diagram of the receiver architecture in which it drives the N quantizers clocked by multi-phase clocks from an ILO locked to the forwarded clock. The ILRO also provides the ability to adjust for the skew between data and the sampling clock by adjusting its own free-running frequency, as demonstrated in [2].

CTLE equalization is chosen versus transmit feed-forward equalization (FFE) in this transceiver architecture, as link modeling studies [13] have found that including a CTLE can achieve less power than a design without TX equalization or designs which include 2-tap TX equalization without a CTLE. This is because the CTLE allows for a peak gain above 0 dB near the Nyquist frequency, which improves the sensitivity of the RX and allows scaling down the transmit output swing significantly. TX FFE, on the other hand, reduces the effective transmitted signal swing, placing more stringent requirements on the RX and also increases the TX circuit complexity. This is especially true for voltage-mode drivers, where significant output-stage segmentation and pre-drive logic is often necessary to achieve a given equalization range and resolution, both in designs which control the output impedance [14] and those that don't [15].

All of the receiver circuits share the same scalable power supply. A higher de-multiplexing ratio relaxes the quantization delay requirement for each quantizer, allowing quantization speed to be traded off for lower supply voltage. For the chosen quantizer structure, which is similar to [16], near-quadratic power reduction is observed associated with supply voltage scaling.

It is important to note that while a highly parallel architecture sees improved power efficiency by operating at lower voltage, several limitations prevents carrying out this methodology indefinitely. The first limitation is that lower overdrive and headroom reduce the performance of analog components in the critical high-speed path. In the case of the CTLE, larger current is needed to maintain its bandwidth at a lower supply voltage, contradicting the effort to reduce power consumption. In turn, larger current and lower headroom also limit the size of the load resistor, making it difficult to achieve the required gain. The

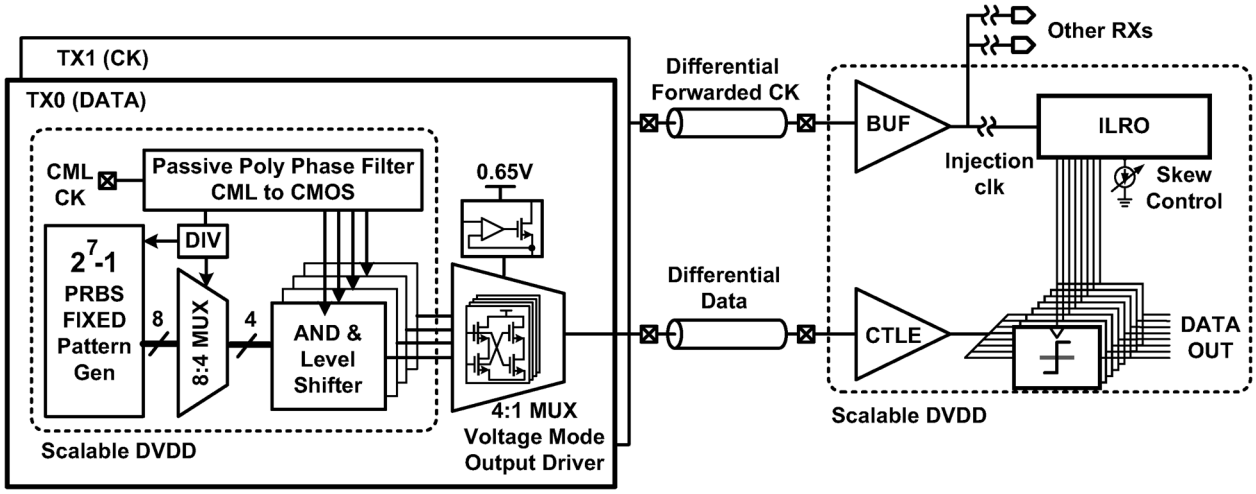

Fig. 8. The implemented single-data-channel low-power forwarded-clock transceiver block diagram.

second limitation is that the use of more quantizers in parallel increases the loading of CTLE, thus decreasing the bandwidth. This loading includes the input capacitance of the quantizer itself, as well as the wiring parasitic, which becomes more significant as longer wires are needed for higher parallelism. The third limitation is that the variation of certain blocks is more sensitive to supply voltage than others. For example, Fig. 6(a) shows the simulated phase mismatch from 100 Monte-Carlo runs of an 8-phase ring oscillator across different supply voltages. Here the phase mismatch is normalized to the UI value corresponding to the frequency achievable at a given supply voltage. It can be observed that σ grows faster as it approaches the near-threshold region. In a receiver, large phase mismatch makes it difficult to align every clock edges for all the parallel quantizers to the proper position in the data eye simultaneously. As a result, the combined BER becomes worse as phase mismatch increases. While individual skew adjustment could be added to each clock phase, this comes at the expense of additional mismatch detection and correction circuitry.

To evaluate the effectiveness of different de-multiplexing ratio and supply voltage combinations in the presence of these limitations, three receivers with different de-multiplexing ratios and supply voltages are simulated. The de-multiplexing ratios are chosen according to the different quantizer delays shown in Fig. 6(b) to meet the same 8 Gb/s throughput target, with constant CTLE output bandwidth maintained for all three designs. Fig. 7 summarizes the power consumption obtained from schematic simulations. Although the power consumption of quantizers and oscillator generally scales down with increased de-multiplexing factor and reduced supply voltage, the CTLE consumes the most power at 0.5 V for the reasons discussed above. This increase in CTLE power consumption nearly cancels all the power savings from scaling V_{DD} from 0.6 V to 0.5 V. Moreover, comparator offset increases significantly at extremely low voltages [17], necessitating excessive offset cancellation circuitry range. Considering the limited total power savings, corresponding CTLE bandwidth degradation, and the increased susceptibility to variation, reducing supply voltage beyond 0.6 V exhibits diminishing returns.

C. Transceiver Architecture

Fig. 8 shows the block diagram of the entire implemented transceiver. In order to optimize power efficiency, the transceiver is implemented with a 4:1 output multiplexing transmitter and an 8:1 de-multiplexing receiver. Except for the transmitter output stage, which is powered by a fixed 0.65 V regulator, all circuitry utilizes a supply which is scaled to the minimum voltage that satisfies the target BER specification for a given data rate.

III. Transmitter

Fig. 9 shows the I/O transmitter block diagram configured for 8 Gb/s operation. Eight bits of parallel input data are serialized in two stages, an initial 8:4 multiplexer and a final 4:1 output multiplexing voltage-mode driver. The clocks which synchronize the serialization are generated by passing a differential quarter-rate clock through a poly-phase filter to generate four quadrature-spaced phases. Two of these phases are divided by two to perform the initial 8:4 multiplexer operation, generating 4 parallel input data streams for the output multiplexing driver. A 4:1 output multiplexing voltage-mode driver is utilized in order to allow low-V_{DD} operation of the serialization stages.

A. Local Multi-Phase Clock Generation

A passive poly-phase filter is utilized to generate the four quadrature clock phases from a globally distributed low-swing quarter-rate clock. In order to enable operation over a wide range of data rates, a two-stage design with staggered time constants is implemented [18], [19]. As shown in Fig. 10, this two-stage design provides quadrature outputs over a range of 1 to 2 GHz with a phase error less than 6°, which is far superior to a single-stage design. In addition, this passive quadrature clock generation structure is well suited for scalable-supply designs, as the clock phase spacing is decoupled from the supply voltage.

The quadrature poly-phase filter outputs are converted to CMOS levels by a CML-to-CMOS converter, as shown in Fig. 11. AC-coupling from the poly-phase filter outputs directly to the input inverter with resistive feedback improves the level converter duty cycle performance [20]. A combination

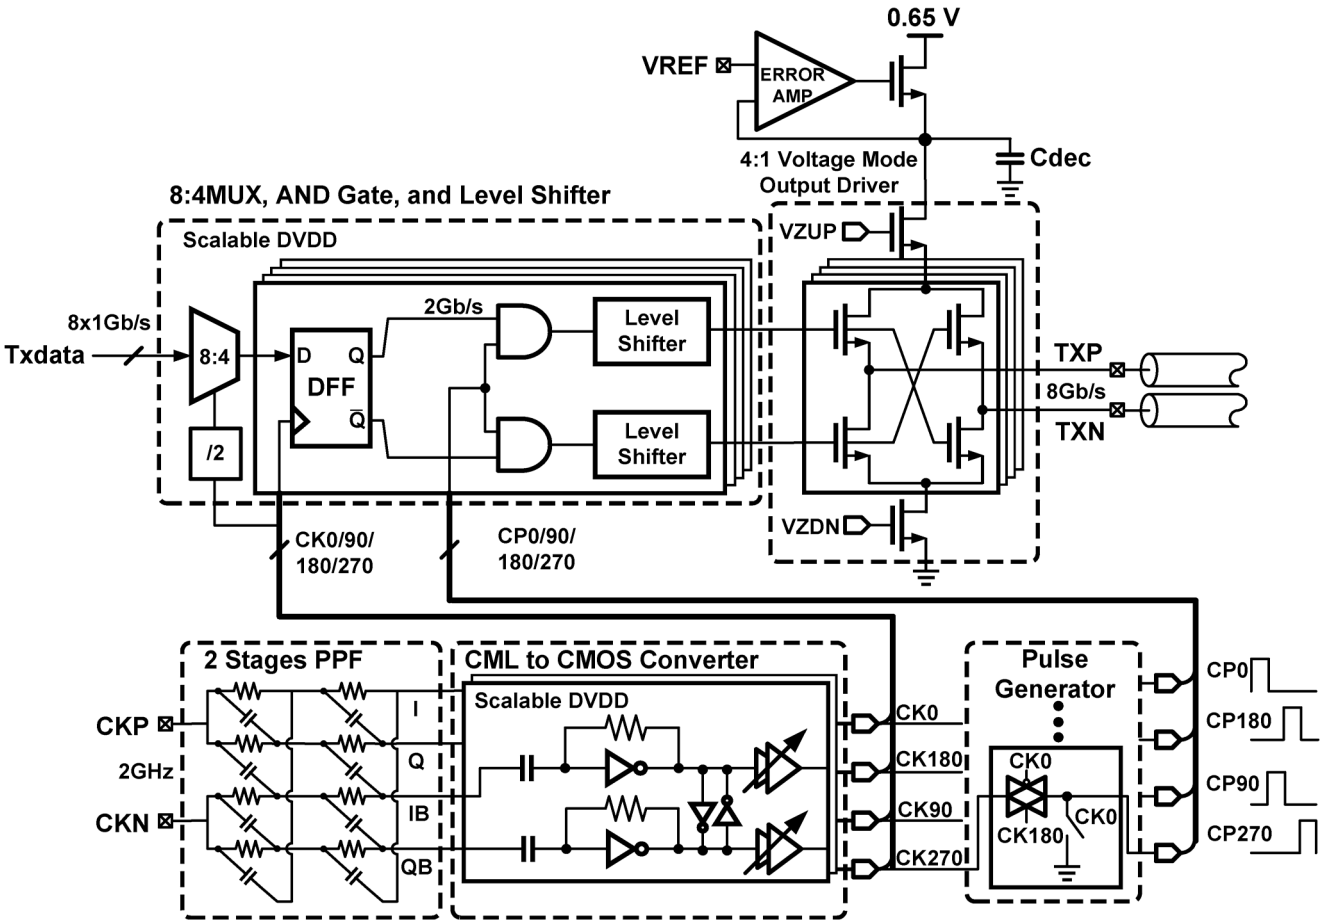

Fig. 9. 4:1 output multiplexing transmitter block diagram.

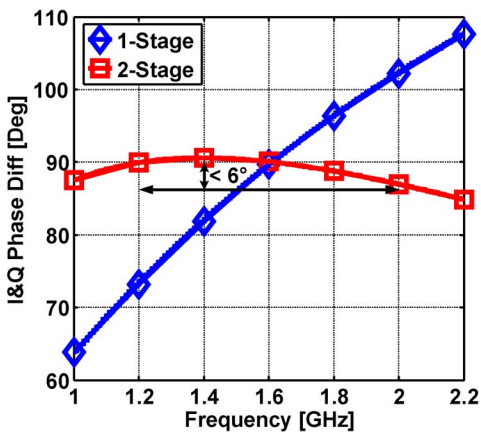

Fig. 10. Passive poly-phase filter I and Q phase spacing versus frequency.

of programmable p-n ratio inverter buffers and two stages of capacitive DACs compensate for both errors in duty cycle and quadrature phase spacing. As shown in Fig. 9, the final pulse-clocks for the output-multiplexing driver are produced by passing the CMOS level quadrature clocks through a transmission-gate AND logic block.

B. Level-Shifting Pre-Driver

One of the challenges associated with scalable-supply designs with voltage-mode output drivers involves maintaining proper channel termination at low-supply voltages without dramatic increases in the output stage transistors. In order to alleviate this problem, a level-shifting pre-driver block (Fig. 12) is utilized to drive the final switch transistors of the voltage-mode output stage with a full DVDD swing above the nominal nMOS threshold voltage, V_{thn}. This level shifting stage, consisting of a feed-forward capacitor that biases the output switches near V_{thn} when off and pulses up to $V_{thn} + DVDD$ when on, allows for a full-DVDD gate overdrive on the output switch transistors, as shown in the simulation results of Fig. 13. This minimizes the size of the output switch transistors required to match the channel impedance, allowing for low supply operation and reduced dynamic power consumption.

C. Output Driver

The low-swing voltage-mode driver is comprised of nMOS transistors, with four parallel switch segments implementing the 4:1 output multiplexing. Driver output impedance is formed by the series combination of the switch transistors driven by the level-shifting pre-drivers and the impedance control transistors shared by the four output segments. A global impedance control loop produces VZUP and VZDN voltages to independently set the pull-up and pull-down impedance, respectively. A voltage regulator sets the power supply of the voltage-mode driver to a value V_{REF}, which due to impedance control is equal to the peak-to-peak differential output swing, allowing for an

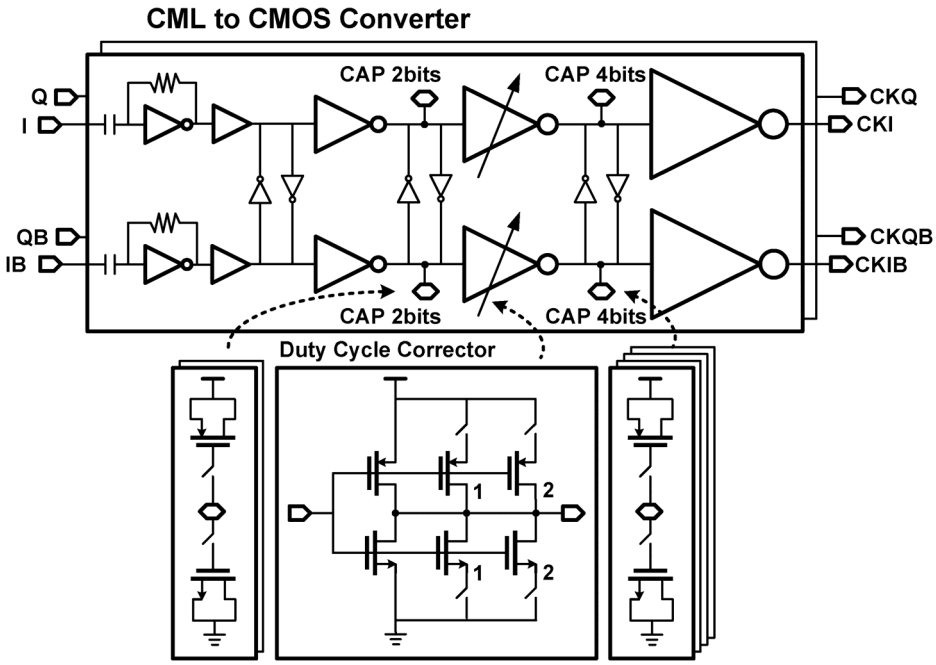

Fig. 11. CML-to-CMOS converter with duty-cycle and phase spacing compensation.

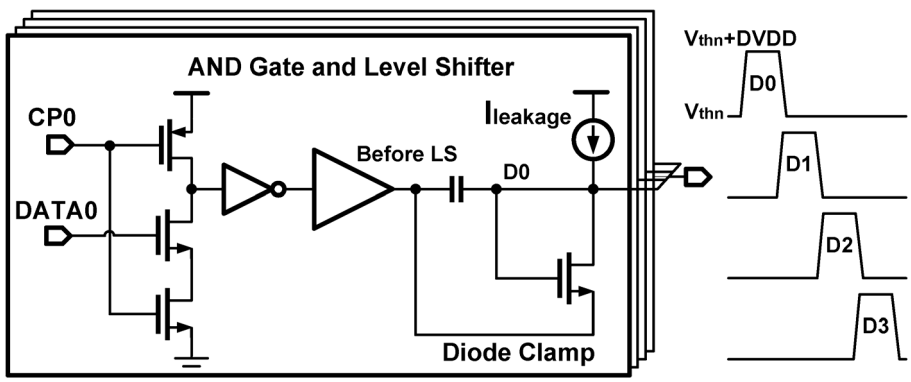

Fig. 12. Level-shifting pre-driver.

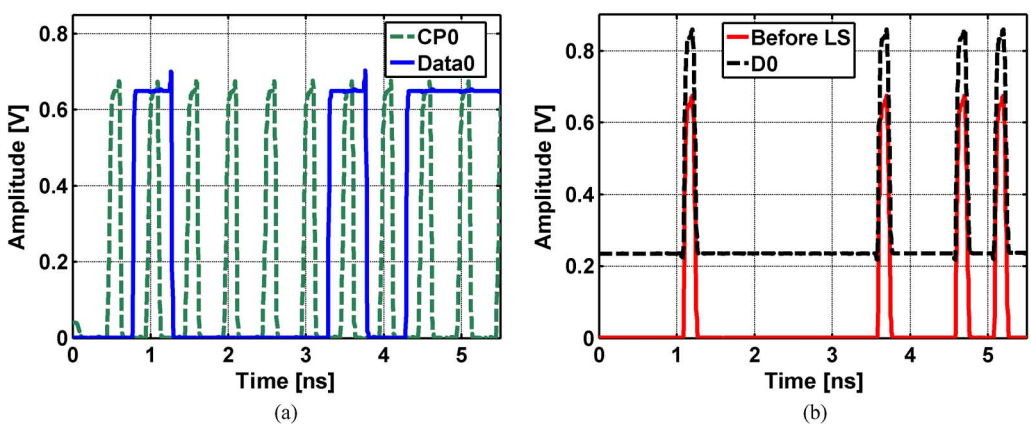

Fig. 13. Level-shifting pre-driver simulated operation: (a) input pulse-clock and data signals, (b) output data pulse before and after level shifting.

adjustable output swing from 100–200 $\mathrm{mV_{ppd}}$. The driver's low common-mode output voltage allows for the regulator to have a source-follower output stage, which offers improved supply-noise rejection relative to common-source output stages.

Utilizing a low supply voltage to power the output stage regulator dramatically improves the transmitter power efficiency. In a multi-channel I/O system, this common regulator supply could be generated by a global I/O regulator with high efficiency, such

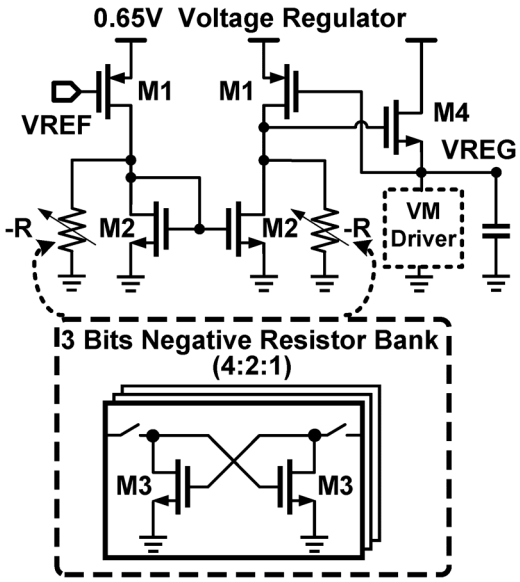

Fig. 14. Low-voltage regulator utilizing a pseudo-differential error amplifier with partial negative-resistance load.

as a switching regulator topology, where the per-channel voltage regulators would allow for improved isolation and output swing optimization. For the per-channel voltage regulator, it is important to achieve a high gain-bandwidth within the error amplifier to minimize the output swing error and provide noise rejection. However, this can be difficult to achieve as the voltage headroom is reduced in low-voltage operation.

In order to achieve a high gain-bandwidth error amplifier at a low 0.65 V supply voltage, a pseudo-differential topology with negative resistance gain boosting is utilized in this design, rather than a conventional simple OTA stage [21] in Fig. 14. Low voltage operation is enabled by the transmit output impedance control, which allows for a tight range of V_{REF} values for a given output swing, and eliminating the typical tail current source while still maintaining a simulated 22 dB power-supply rejection ratio. A programmable negative resistance load increases the DC gain of the error amplifier to

$$|A_{err}| = \frac{g_{m1}}{g_{ds1} + g_{ds2} - g_{m3}}. \quad (1)$$

Fig. 15 shows that this negative resistive load boosts the low frequency error amplifier gain by approximately 12 dB, while still maintaining adequate stability. The low frequency error amplifier gain can be further increased to near 30 dB by increasing the negative resistance strength; however stability is compromised, as shown in the supply step response simulations. In order to guarantee regulator stability over process variations, a three-bit digital control is utilized to tune the negative impedance value.

D. Global Impedance Controller

Fig. 16 shows the global output driver impedance controller that produces the output voltages, VZUP and VZDN, which controls multiple output drivers' pull-up and pull-down impedance, respectively, allowing for impedance control loop power amortization among the number of transmitter channels [21]. A replica transmitter stage with a precision off-chip 100 Ω resistor is placed in two feedback loops, one which sets the top-most transistor gate voltage, VZUP, to force a value of (3/4)*VREF at the replica transmitter positive output, and the other which sets the bottom-most transistor gate voltage, VZDN, to force a value of (1/4)*VREF at the replica transmitter negative output. While other voltage-mode impedance control schemes primarily utilize the pre-driver supply voltage [4], [10], utilizing dedicated transistors for impedance control allows the pre-drive swing value to be decoupled from the impedance control, providing a degree of freedom to allow for potential pre-drive voltage scaling for improved power efficiency [21]. A replica bias circuit consisting of a diode-connected nMOS whose source is connected to the scalable DVDD biases the replica switch transistors to a voltage level, $\mathrm{VLS} = \mathrm{V_{thn}} + \mathrm{DVDD}$, consistent with the level shifting pre-driver output. The driver output resistance is partitioned with nominally 30 Ω switch transistors and 20 Ω impedance control transistors in order to reduce the switch transistor size and obtain lower dynamic power consumption.

IV. Receiver

The receiver consists of an input CTLE that drives eight parallel data quantizers [17] and provides up to 8 dB of peaking to support low-loss channels. While a multi-stage CTLE could potentially provide higher gain and peaking, it would lower bandwidth due to additional poles in the signal path. The quantizers are each clocked from eight phases generated by an ILRO locked to an eighth-rate forwarded clock from the transmitter chip.

Injection locking has been demonstrated as an energy-efficient scheme for both clock generation and de-skewing due to its reduced complexity relative to other approaches such as PLL- or DLL-based timing recovery [2], [22]. In addition, when ILRO-based de-skew is combined with aggressive supply voltage scaling, excellent receiver energy-efficiency of <0.2 pJ/b at 8 Gb/s has been demonstrated in a previous work [17].

Fig. 17 shows the ILRO used in this design, which consists of a 4-stage differential current-starved ring oscillator. The oscillation frequency is controlled by a tail current source that is split into two parts, one controlled by an external frequency-locked loop to nominally oscillate at the forwarded eighth-rate frequency, and the other portion controlled by a 6-bit binary code for de-skew. In order to enable ILRO operation over a wide frequency range, the relative strength between the frequency-tuning current source and de-skewing current sources is adjustable, effectively decoupling the frequency tuning range from the de-skew step resolution. The frequency locking process, which is performed at start-up or during periodic link re-training, insures that the ring oscillator free-running frequency is at the desired forwarded eighth-rate clock frequency. This also ensures that the ring oscillator operates near the center of the locking range before injection, and has enough tuning range to provide either positive or negative skew.

The forwarded differential clock is first buffered and converted to full scale before being distributed to the ILRO. In order to support different data rates and channel conditions,

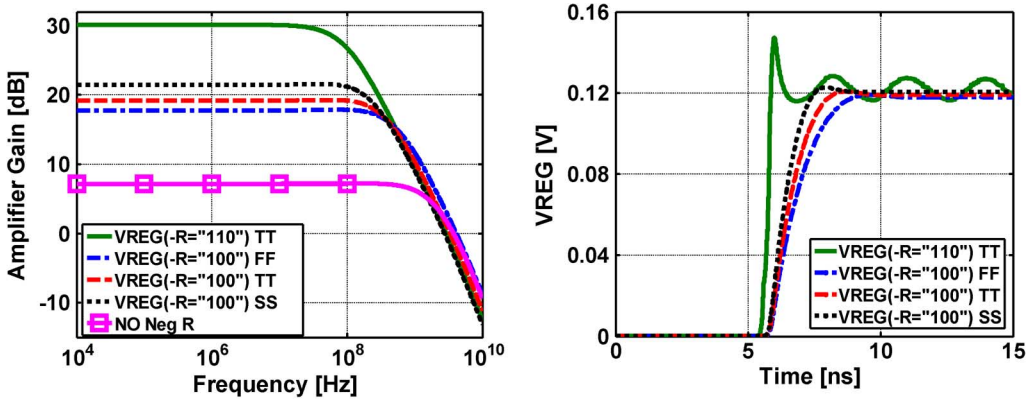

Fig. 15. Low-voltage regulator simulated performance with various negative resistance settings: (a) error amplifier gain versus frequency, (b) supply step response from 0 to 0.65 V with $\mathrm{VREF} = 120\ \mathrm{mV}$.

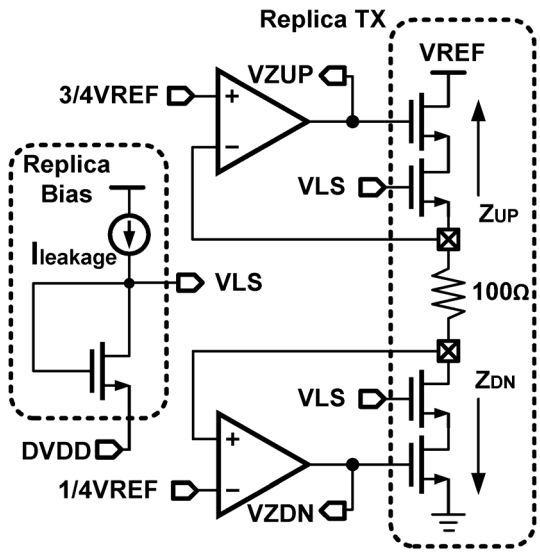

Fig. 16. Global output driver impedance controller.

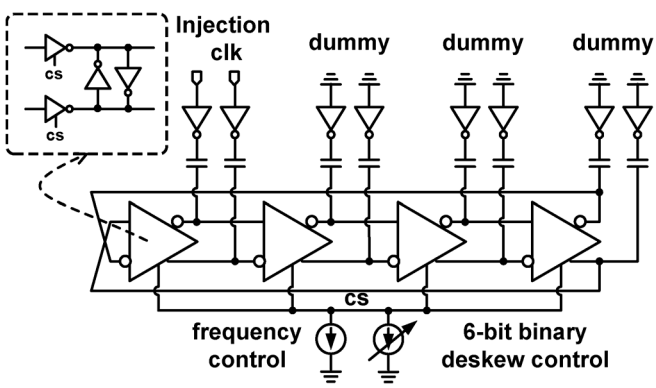

Fig. 17. ILRO schematic.

4-bit amplitude control is included in the clock input buffer. The buffered clocks are then injected into two complementary oscillator stages through coupling capacitors, with dummy capacitors placed at the other stages to equalize the load capacitances. Fixed injection strength is used for this design in order to minimize excessive phase spacing errors. As shown in the simulation results of Fig. 18, this fixed-strength AC-coupled injection approach results in a more uniform phase spacing compared to DC-coupled injection schemes that use V/I converters, such as the technique incorporated in [2], while exhibiting a similar locking range. Similar to the transmitter multi-phase clocking paths, capacitive DACs in the clock buffer stages following the ILO compensate for phase spacing errors.

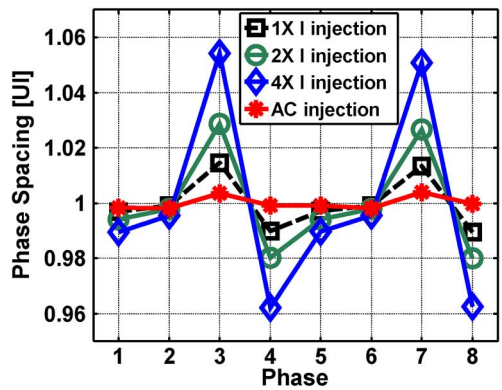

Fig. 18. Simulated impact of clock injection approach on phase spacing uniformity.

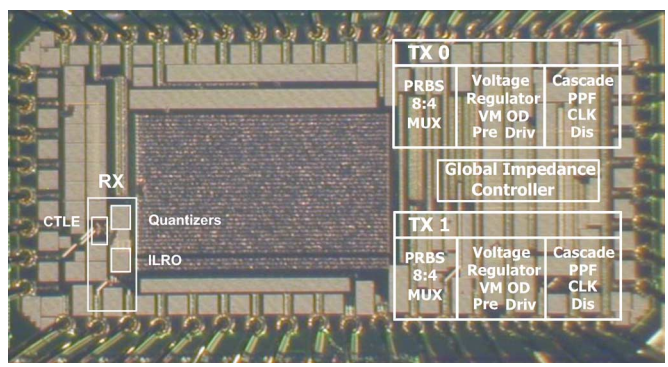

Fig. 19. I/O transceiver chip micrograph.

V. Experimental Results

The transceiver was fabricated in a 65 nm CMOS general purpose process. As shown in the die micrograph of Fig. 19,

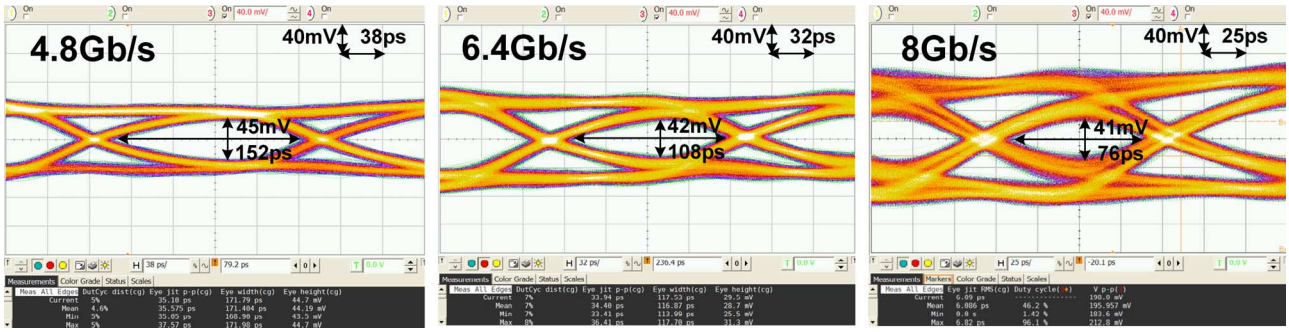

Fig. 20. 4.8 Gb/s, 6.4 Gb/s, and 8 Gb/s transmitter output eye diagrams.

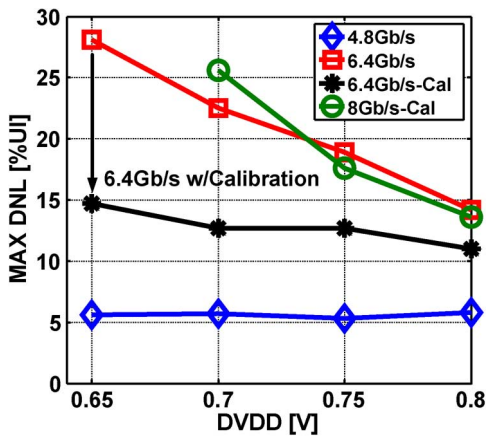

Fig. 21. 4:1 output-multiplexing transmitter phase spacing maximum DNL versus supply voltage.

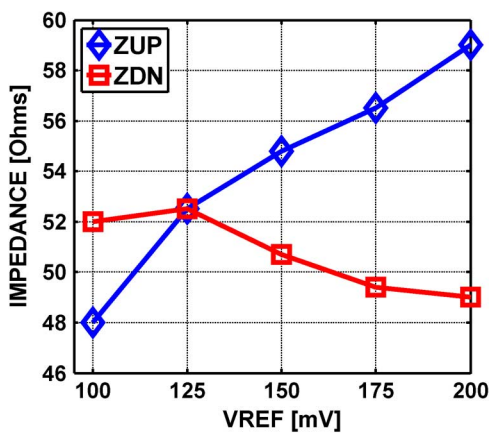

Fig. 22. Transmitter output impedance versus VREF.

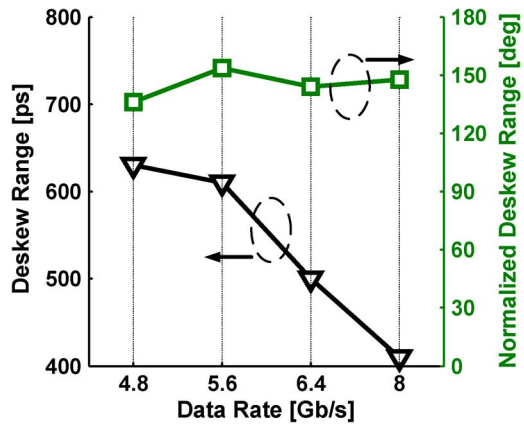

Fig. 23. Receiver de-skew range.

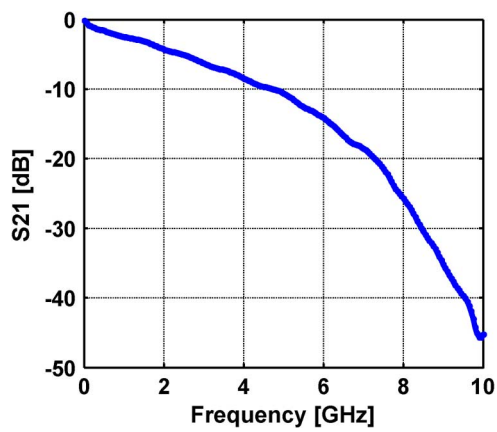

Fig. 24. Frequency response of 3.5″ FR4 trace and interconnect cables.

the total active area for the transmitter is $214 \times 104\ \mu\mathrm{m}^2$, the global impedance controller is $140 \times 31\ \mu\mathrm{m}^2$, and the receiver is $139 \times 230\ \mu\mathrm{m}^2$, for a total transceiver area of $0.057\ \mathrm{mm}^2$ and a bandwidth density of $0.007\ \mathrm{mm}^2/\mathrm{Gb/s}$. Conservatively considering a minimum of 4 wire-bond pads at a 100 μm pitch for the differential TX and RX data signals, the design has a circuit/pad area ratio of 2.9, and could be considered active-area limited. While if the design was implemented with coarser-pitch C4 bumps [1], the circuit/bump area ratio falls to 0.46 for 4 C4 bumps, and could be considered bump-limited. Given the slower pitch scaling of both bondpads and C4 bumps, this architecture is projected to be both pad and bump limited in a 22 nm CMOS node.

A chip-on-board test setup is utilized, with the die directly wirebonded to the FR4 board. In order to demonstrate the transmitter functionality, the eye diagrams of Fig. 20 are produced with a short 1.5″ channel. In order to demonstrate transmitter operation, both the transmitter scalable power supply and output swing are optimized at a given data rate to achieve a minimum 40 $\mathrm{mV_{ppd}}$ eye height and 0.6UI eye width at the channel output, with 0.65 V and a 150 $\mathrm{mV_{REF}}$ DC output swing at 6.4 Gb/s. Fig. 21 shows the results of the 4:1 output-multiplexing transmitter for its phase-spacing mismatch versus the scalable power

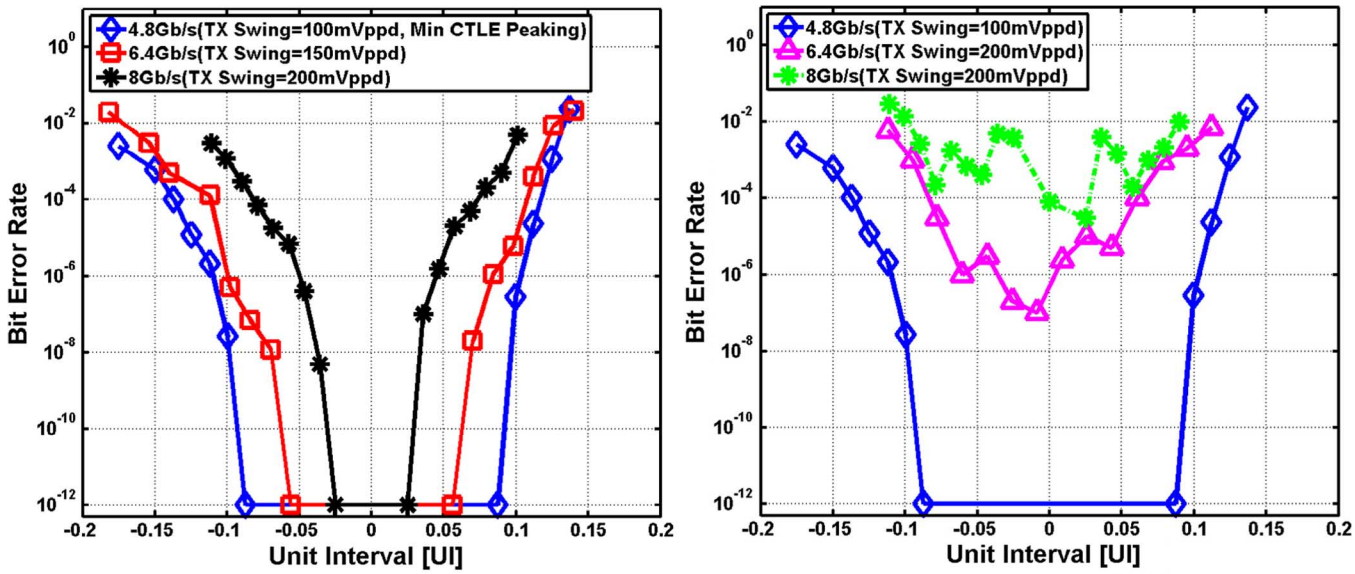

(a) (b)

Fig. 25. (a) Transceiver BER performance with optimal TX/RX supply voltages and CTLE settings, (b) transceiver BER with minimum CTLE peaking settings.

supply. Phase spacing mismatches increase with higher data rate, resulting in a minimum supply voltage for an acceptable phase DNL at a given data rate. Duty-cycle control circuitry and tunable-delay quadrature clock buffers allow for calibration that improves phase DNL. For example, calibration at 6.4 Gb/s and 0.65 V improves from the max phase DNL from 28% UI to 15% UI, with further improvement limited by an oversight in the chip layout that resulted in asymmetrical clock routing. Fig. 22 shows the effectiveness of the impedance loop, where both Z_{UP} and Z_{DN} are between 48 to 59 Ω as the output swing, V_{REF}, varies from 100–200 mV_{ppd}. While tighter impedance control is not essential [15], this could be achieved by sizing the output drivers' impedance control transistors to achieve a wider tuning range, at the cost of larger switch transistors and increased dynamic power.

Fig. 23 shows the measured de-skew range of the receiver ILRO versus data rate. When normalized to the clock period, the achievable de-skew range is more than 120° across the entire operating range. Since in the 1:8 de-multiplexing receiver 1UI is 45°, this translates into a de-skew range that exceeds 2UI.

Transceiver performance is verified with BER measurements of $2^7 - 1$ PRBS data over the channel shown in Fig. 24, which consists of a 1.5 inch FR4 TX-side trace, a 0.5 m SMA cable, and a 2 inch FR4 RX-side trace, and displays −8.4 dB loss at 4 GHz. BER results with optimized TX/RX supply voltages, TX output swing, and CTLE settings are shown in Fig. 25(a), and CTLE performance impact is shown in Fig. 25(b). A fixed 130 fF capacitor and a programmable 100–650 Ω resistor makes up the CTLE degeneration network. At 4.8 Gb/s, a 16% UI timing margin is achieved with a 100 mV_{ppd} TX swing and the minimum 100 Ω CTLE degeneration resistor setting. While the CTLE could perhaps be eliminated at 4.8 Gb/s, operation at 6.4 Gb/s requires 350 Ω degeneration and 8 Gb/s requires the maximum 650 Ω setting. Due to the channel loss and increased sensitivity to phase mismatches, the required transmit swing is increased to 150 mV_{ppd} and 200 mV_{ppd} at 6.4 Gb/s and 8 Gb/s, respectively.

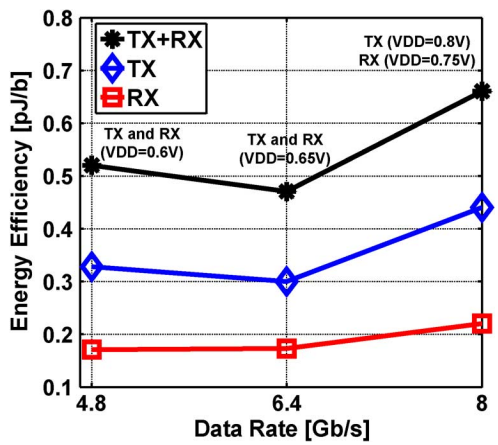

Fig. 26. Transceiver energy efficiency versus data rate.

TABLE I
TRANSCEIVER POWER BREAKDOWN AT 6.4 Gb/s

TX Power Breakdown (6.4Gb/s at 0.65V)	
LDO & Output Driver (150mV_{ppd})	793uW
Serializer, Pre-drivers, Clocking	933uW
Global Impedance Control (amortized across 9 TX)	193uW
TX Energy Efficiency	0.3pJ/b
RX Power Breakdown (6.4Gb/s at 0.65V)	
CTLE, Quantizers, ILRO	1.07mW
Clock Distribution	38uW
RX Energy Efficiency	0.17pJ/b
Total Energy Efficiency	**0.47pJ/b**

Fig. 26 shows transceiver energy efficiency measurement results at various data rates and supply voltages. The transmitter and receiver supply is equal at 0.6 V and 0.65 V for 4.8 Gb/s and 6.4 Gb/s, respectively. However in order to achieve 8 Gb/s operation, the transmitter requires a slightly higher 0.8 V supply to maintain sufficient margin in the 4:1 output multiplexing phase

TABLE II
LOW-POWER I/O TRANSCEIVER COMPARISONS

	[1]	[4]	This Work
Technology	45nm CMOS	90nm CMOS	65nm CMOS
Supply Voltage	0.8V/1.5V	1.2V	0.6-0.8V
Data Rate	10Gb/s	0.5-4Gb/s	4.8-8Gb/s
Clocking	Source-Synchronous	Plesiochronous	Source-Synchronous
Energy Efficiency	1.4pJ/b @ 10Gb/s	1.9pJ/b @ 3.2Gb/s	0.47pJ/b @ 6.4Gb/s
Transmitter			
Driver	CML, 2:1 Input Mux	VM, 2:1 Input Mux	VM, 4:1 Output Mux
Swing	150mVppd	100mVppd	100-200mVppd
Equalization	2-Tap FFE	None	None
Energy Efficiency	0.65pJ/b	0.6pJ/b	0.3pJ/b
Channel			
	2" HDI	Not reported	3.5" FR4 + 0.5m SMA cable
Loss at Nyqu Freq	8 dB		8.4
Receiver			
Equalization	None	CTLE	CTLE
Energy Efficiency	0.75pJ/b	1.3pJ/b	0.17pJ/b

spacing, which has a greater impact on the output transmitter eye at high data rates due to the low-pass filtering of the high-speed off-chip data. While the receiver CTLE and quantizers would work fine at this 0.8 V supply at 8 Gb/s, unfortunately this voltage is somewhat high for the ILRO and pushes the injection lock range above 1 GHz. Thus, 0.75 V is required at the receiver to allow the ILRO to operate at the 1 GHz frequency required for 8 Gb/s operation. In the event the I/O system demands that the transmitter and receiver operate with equal supply voltages, this could be achieved by adding switchable capacitor loads to the ILRO. While the transceiver operates at the lowest voltage at 4.8 Gb/s, optimal energy efficiency is achieved at 6.4 Gb/s due to the amortization of the static power consumed in the final output line driver.

Table I shows the measured transceiver power breakdown at 6.4 Gb/s. The total transceiver energy-efficiency is 0.47 pJ/b, with 0.3 pJ/b and 0.17 pJ/b efficiency achieved in the transmitter and receiver, respectively. Table II compares this design with recent energy-efficient serial links that either employ source-synchronous clocking [1] or utilize a voltage-mode driver [4]. On the transmitter side, compared to the current-mode output driver in [1] and conventional 2:1 input multiplexing voltage-mode output driver in [4], the 4:1 output multiplexing voltage-mode driver in this design improves energy efficiency by more than 50%. On the receiver side, supply scaling and the use of ILRO have also resulted in significant power efficiency improvements over similar designs with linear equalization to compensate for moderate-loss channels.

VI. SCALING TO HIGHER PER-PIN DATA RATES

As future systems will demand per-pin data rates in excess of 10 Gb/s in the near future, it is interesting to consider how this architecture can support this under various scenarios. The first scenario considered is with a fixed channel-loss ($\sim$10 dB) at the Nyquist frequency and scaling the technology node with increased data rate. This allows the architecture to remain somewhat unchanged, except for adjusting the transmit poly-phase filter passives to support the higher clock frequencies. Thus, the data rate which can be supported at a given energy efficiency scales with the improvement in technology speed, with 24 Gb/s at 0.7 pJ/bit projected in a 22 nm CMOS node. In another scenario where the channel bandwidth does not scale with data rate, the architecture would need to be modified to include extra equalization. At the transmit side this can be accomplished in an efficient manner by modifying the existing voltage-mode driver to a hybrid driver which includes parallel current-mode equalization with very low pre-drive complexity [21]. Including a multi-stage CTLE topology [23] at the receiver would also allow support of higher-loss channels with relatively low overhead. The inclusion of these two efficient equalization blocks would allow operation with an additional 10–15 dB loss.

VII. CONCLUSION

This paper presented an energy-efficient transceiver architecture that operates at low supply voltages. In order to reduce the transmitter dynamic power consumption, a passive poly-phase filter is utilized to produce the multi-phase clocks that switch a 4:1 output-multiplexing voltage-mode driver. A low power-supply linear regulator with negative-resistance gain-boosting allows further improvement in transmitter energy efficiency. In the forwarded-clock receiver, the use of injection-locked oscillator de-skew and a high 1:8 de-multiplexing ratio receiver architecture allows operation at low supply voltages. Overall, this I/O architecture provides scalable voltage and data rate operation at energy-efficiency levels demanded by future systems.

ACKNOWLEDGMENT

The authors would like to thank Dr. J. Liu of the University of Texas–Dallas for the use of measurement equipment.

REFERENCES

[1] F. O'Mahoney, J. E. Jaussi, and J. Kennedy *et al.*, "A 47 × 10 Gb/s 1.4 mW/Gb/s parallel interface in 45 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2828–2837, Dec. 2010.

[2] K. Hu, T. Jiang, J. Wang, F. O'Mahony, and P. Y. Chiang, "A 0.6 mW/Gb/s, 6.4–7.2 Gb/s serial link receiver using local injection-locked ring oscillators in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 899–908, Apr. 2010.

[3] B. Casper and F. O'Mahony, "Clocking analysis, implementation and measurement techniques for high-speed data links—A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 1, pp. 17–39, Jan. 2009.

[4] R. Inti *et al.*, "A highly digital 0.57–4 Gb/s 1.9 mW/Gb/s serial-link transceiver using currect-recycling in 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2011, pp. 152–153.

[5] A. P. Chandrakasan *et al.*, "Technologies for ultradynamic voltage scaling," *Proc. IEEE*, vol. 98, no. 2, pp. 191–214, Feb. 2010.

[6] J. Kim and M. Horowitz, "Adaptive supply serial links with sub-1 V operation and per-pin clock recovery," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1403–1413, Nov. 2002.

[7] R. Reutemann *et al.*, "A 4.5 mW/Gb/s 6.4 Gb/s 22+1-lane source synchronous receiver core with optional cleanup PLL in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2850–2860, Dec. 2010.

[8] C.-K. Yang and M. Horowitz, "A 0.8 μm CMOS 2.5 Gb/s oversampling receiver and transmitter for serial links," *IEEE J. Solid-State Circuits*, vol. 31, no. 12, pp. 2015–2023, Dec. 1996.

[9] H. Lee *et al.*, "A 16 Gb/s/link, 64 GB/s bidirectional asymmetric memory interface," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1235–1247, Apr. 2009.

[10] J. Poulton *et al.*, "A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2745–2757, Dec. 2007.

[11] G. Balamurugan, J. Kennedy, and G. Banerjee *et al.*, "A scalable 5–15 Gbps, 14–75 mW low power I/O transceiver in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 1010–1019, Apr. 2008.

[12] K. Fukuda *et al.*, "A 12.3 mW 12.5 Gb/s complete transceiver in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2010, pp. 368–369.

[13] A. Palaniappan and S. Palermo, "A design methodology for power efficiency optimization of high-speed equalized-electrical I/O architectures," *IEEE Trans. VLSI Systems*, 2012, early access.

[14] W. D. Dettloff, J. C. Eble, L. Luo, P. Kumar, F. Heaton, T. Stone, and B. Daly, "A 32 mW 7.4 Gb/s protocol-agile source-series terminated transmitter in 45 nm CMOS SOI," in *IEEE ISSCC Dig. Tech. Papers*, 2010, pp. 370–371.

[15] R. Sredojevic and V. Stojanović, "Digital link pre-emphasis with dynamic driver impedance modulation," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC'10)*, San Jose, CA, Sep. 2010, pp. 1–4.

[16] D. Schinkel *et al.*, "A double-tail latch-type voltage sense amplifier with 18 ps setup+hold time," in *IEEE ISSCC Dig. Tech. Papers*, 2007, pp. 314–605.

[17] K. Hu *et al.*, "0.16–0.25 pJ/bit, 8 Gb/s near-threshold serial link receiver with super-harmonic injection locking," *IEEE J. Solid-State Circuits*, vol. 47, no. 8, pp. 1842–1853, Aug. 2012.

[18] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[19] J. Kaukovuori, K. Stadius, J. Ryynänen, and K. A. I. Halonen, "Analysis and design of passive polyphase filters," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 10, pp. 3023–3037, Nov. 2008.

[20] C. Menolfi *et al.*, "A 16 Gb/s source-series terminated transmitter in 65 nm CMOS SOI," in *IEEE ISSCC Dig. Tech. Papers*, 2007, pp. 446–447.

[21] Y.-H. Song and S. Palermo, "A 6-Gbit/s hybrid voltage-mode transmitter with current-mode equalization in 90-nm CMOS," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 8, pp. 491–495, Aug. 2012.

[22] M. Hossain and A. C. Carusone, "7.4 Gb/s 6.8 mW source synchronous receiver in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 6, pp. 1337–1348, Jun 2011.

[23] J. Bulzacchelli *et al.*, "A 28 Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32 nm SOI CMOS technology," in *IEEE ISSCC Dig. Tech. Papers*, 2012, pp. 324–325.

Young-Hoon Song (S'12) received the B.S. and M.S. degrees in electrical engineering from University of Texas at Arlington, TX, USA, in 2002 and 2004, respectively. He is currently working toward the Ph.D. degree at the Analog and Mixed Signal Center (AMSC) at Texas A&M University, College Station, TX, USA.

In 2008 and 2012, he was a intern at Broadcom in Irvine, CA, USA, and at IBM T. J. Watson Research Center in Yorktown Heights, NY, USA. His research interests include mixed-signal integrated circuit design primarily as applied to power efficient serial link transceiver for high speed digital communications.

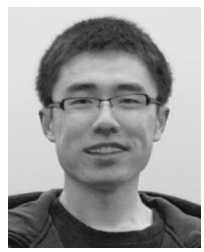

Rui Bai (S'11) received the B.S. degree in microelectronics from the University of Electronic Science and Technology of China (UESTC), Chengdu, China, in 2008. He has been working towards the Ph.D. in electrical engineering at Oregon State University, Corvallis, OR, USA, since 2009.

His current research interests include power-efficient high-speed circuits for electrical and optical I/O.

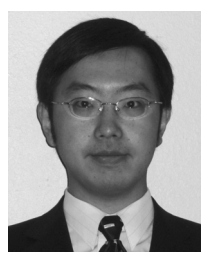

Kangmin Hu (S'08–M'12) received the B.S. degree in electrical engineering from Shanghai Jiao Tong University, Shanghai, China, in 2004, the M.S. degree in microelectronics from Fudan University, Shanghai, China, in 2007, and the Ph.D. degree in electrical engineering from Oregon State University, Corvallis, OR, USA, in 2011.

Since 2011, he has been with Broadcom Corporation, Irvine, CA, USA, as a Staff Scientist, where he is working on CMOS high-speed SerDes circuits. His technical interests include system analysis and circuit design on high-speed low-power I/O transceivers.

Dr. Hu was a recipient of the 2008 Best Poster Award from Center for Design of Analog-Digital Integrated Circuits (CDADIC), and the 2011 Intel/Helic Student Scholarship Award at the IEEE Custom Integrated Circuits Conference (CICC). He serves as a reviewer for several IEEE journals.

Hae-Woong Yang (S'10) received the B.S. degree and the M.E. degree in electrical and computer engineering from Texas A&M University, College Station, TX, USA, in 2007 and 2009, respectively. He is currently working toward the Ph.D. degree at the Analog and Mixed Signal Center (AMSC) in Texas A&M University, College Station, TX, USA.

His interests include low-power high-speed electrical link circuits, clock generation circuits, and signal integrity.

Patrick Yin Chiang (S'99–M'04) received the B.S. degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 1998, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 2001 and 2007.

He is an Associate Professor (on sabbatical) with Oregon State University, Corvallis, OR, USA. He currently is a 1000-Talents Young Professor at the ASIC & System State Key Laboratory at Fudan University in Shanghai, China. In 1998, he was a design engineer at Datapath Systems (now LSI), where he designed a low-power, standard cell library for xDSL. In 2003 he was a research intern at Velio Communications (now Rambus) investigating 10 GHz clock synthesis techniques. In 2004, he worked as a consultant at Telegent Systems (now Spreadtrum), where he analyzed low-phase VCOs for TV tuners. In 2006 he was a visiting NSF research fellow at Tsinghua University, China, investigating super-regenerative RF transceivers. His interests are in energy-efficient circuits and systems, such as low-power wireline and photonic interfaces, energy-constrained medical sensors, and reliable near-threshold computing.

Dr. Chiang was the recipient of a 2010 Department of Energy Early CAREER award and a 2012 NSF CAREER award. He is an associate editor of IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS, and is on the technical program committee for the IEEE Custom Integrated Circuits Conference.

Samuel Palermo (S'98–M'07) received the B.S. and M.S. degree in electrical engineering from Texas A&M University, College Station, TX, USA, in 1997 and 1999, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2007.

From 1999 to 2000, he was with Texas Instruments, Dallas, TX, USA, where he worked on the design of mixed-signal integrated circuits for high-speed serial data communication. From 2006 to 2008, he was with Intel Corporation, Hillsboro, OR, USA, where he worked on high-speed optical and electrical I/O architectures. In 2009, he joined the Electrical and Computer Engineering Department of Texas A&M University where he is currently an Assistant Professor. His research interests include high-speed electrical and optical links, clock recovery systems, and techniques for device variability compensation.

Dr. Palermo is a recipient of a 2013 NSF CAREER award. He is a member of Eta Kappa Nu. He currently serves as an associate editor for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II, and served on the IEEE CAS Board of Governors from 2011 to 2012. He was a coauthor of the Jack Raper Award for Outstanding Technology-Directions Paper at the 2009 International Solid-State Circuits Conference.