A Scalable 0.128–1 Tb/s, 0.8–2.6 pJ/bit, 64-Lane Parallel I/O in 32-nm CMOS

Mozhgan Mansuri, *Member, IEEE*, James E. Jaussi, *Member, IEEE*, Joseph T. Kennedy, *Member, IEEE*, Tzu-Chien Hsueh, *Member, IEEE*, Sudip Shekhar, *Member, IEEE*, Ganesh Balamurugan, *Member, IEEE*, Frank O'Mahony, *Member, IEEE*, Clark Roberts, Randy Mooney, *Member, IEEE*, and Bryan Casper, *Member, IEEE*

***Abstract*—A scalable 64-lane chip-to-chip I/O, with per-lane data rate of 2–16 Gb/s is demonstrated in 32-nm low-power CMOS technology. At maximum aggregate bandwidth of 1.024 Tb/s across 50-cm channel length, the link consumes 2.7 W from a 1.08-V supply, corresponding to 2.6 pJ/bit. As bandwidth demand decreases, scaling the per-lane data rate to 4 Gb/s and power supply to 0.65 V provides 1/4 of the maximum bandwidth while consuming 0.2 W. Across a 1-m channel, the link operates at a maximum per-lane data rate of 16 Gb/s; thus, providing up to 1.024 Tb/s of aggregate bandwidth with 3.2 pJ/bit power efficiency from a 1.15-V supply. A length-matched dense interconnect topology allows clocking to be shared across multiple lanes to reduce area and power. Reconfigurable current/voltage mode transmitter driver and CMOS clocking enable a highly scalable power-efficient link. Optional low-dropout regulators provide >22-dB supply noise rejection at the package resonance frequency of 200 MHz. System-level optimization of duty-cycle and quadrature error correctors across the clock hierarchy provides optimized clock phase placement and, thus, enhances link performance and power. A lane failover mechanism provides design robustness to mitigate channel or circuit defects. The active circuitry occupies 1.3 mm^2.**

***Index Terms*—CMOS clocking, high-speed I/O, link, low-area, low-power, scalable circuits.**

I. Introduction

HIGH-PERFORMANCE graphics and computing systems continue to demand aggressive I/O scaling in which terabits/s of bandwidth are required for chip-to-chip data communication. While circuit bandwidth continues to scale well with process technology, interconnect remains the bandwidth bottleneck. As the channel loss increases, the link power efficiency degrades due to the need for complex equalization, larger transmit swing, and low-jitter clock requirements. Additionally, I/O power as a proportion of total system power will grow if bandwidth demand is not accompanied by proportional I/O power efficiency scaling. For example, using Intel Xeon as a baseline and assuming I/O bandwidth doubles every five years,

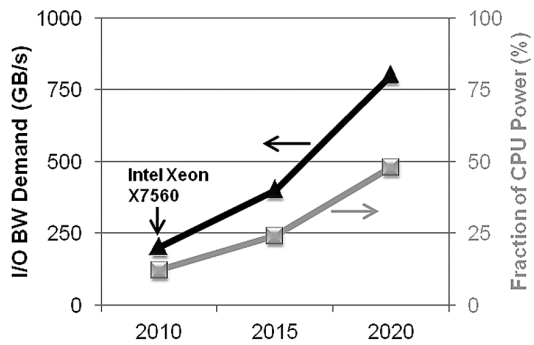

Fig. 1. High-performance computing I/O bandwidth and power trends.

Fig. 1 predicts I/O bandwidth demand of 750 GB/s for supercomputers by the year 2020. Assuming one-sided I/O power efficiency of 10 pJ/bit, I/O power could be up to 50% of CPU power at 135-W thermal design power by the year 2020. Practical realization of an I/O system with such large bandwidth requires cooptimization of circuit and interconnect design where link power, area, and cost are minimized.

Traditional off-chip interconnects utilizing sockets, connectors and PCB material suffer from both density and bandwidth bottlenecks. Channel technology with top-side connectors reduces loss by bypassing socket and core package layers as well as alleviating the package breakout constraints [1]. Additionally, top-side connectors provide higher density than a conventional CPU-socket interface. For a given aggregate bandwidth target, this wide interface enables lower per-lane data rate. Operating lanes at lower data rates reduces channel loss and relaxes required circuit bandwidth; thus, providing better power efficiency. However, the area density of I/O circuitry and interconnect is crucial to minimize area and cost for such a highly parallel interface.

While supercomputer demand for I/O peak bandwidth is high, the average bandwidth usage is typically much lower than the peak bandwidth. As the aggregate bandwidth demand is reduced, the link bandwidth can be decreased by powering down as many data lanes are required and serializing the data across the remaining lanes. This results in power scaling which is approximately linearly proportional to the I/O bandwidth scaling. Alternatively, the per-lane data rate can be scaled according to the required aggregate bandwidth. At lower data rates, where the channel loss is reduced and circuit bandwidth is more relaxed, aggressive supply scaling significantly improves power efficiency; thus, the link power scales at a rate faster than the link bandwidth reduction. In other words, power

Manuscript received April 10, 2013; revised June 18, 2013 and July 15, 2013; accepted July 15, 2013. Date of publication September 06, 2013; date of current version November 20, 2013. This paper was approved by Associate Editor Azita Emami. This work was supported in part by the U.S. Government under Contract HR0011-10-3-0007.

The authors are with the Intel Corporation, Hillsboro, OR 97124 USA (e-mail: mozhgan.mansuri@intel.com).

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.2013.2279052

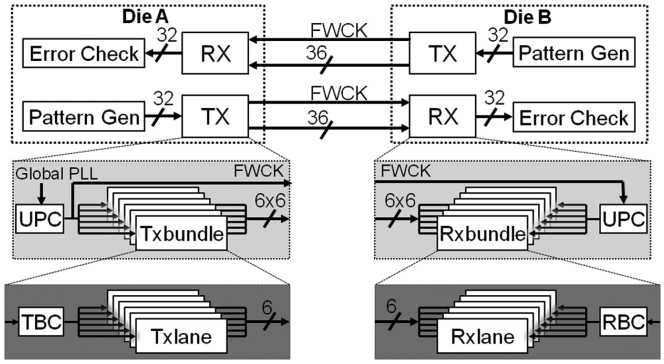

Fig. 2. A 64-lane forwarded clock I/O port: unified port clock (UPC), transmitter bundle clock (TBC), receiver bundle clock (RBC) and data lanes (Txlane or Rxlane).

scales nonlinearly with bandwidth. Additionally, a scalable link should operate over a wide range of data rates and supply voltages while meeting performance and power efficiency requirements as well as being tolerant to process, voltage and temperature (PVT) variation.

This paper describes a scalable low-power, dense 64-lane per-port forwarded clock I/O system, targeting high-performance computing applications. This prototype provides an aggregate bandwidth of up to 1.028 Tb/s per port and 2.6 pJ/bit power efficiency in 32-nm low-power CMOS technology. Aggressive power supply scaling from 0.6 to 1.08 V enables nonlinear power efficiency scaling of 0.8–2.6 pJ/bit as the per-lane data rate scales from 2 to 16 Gb/s, providing scalable aggregate bandwidth of 0.128 to 1.028 Tb/s. Highly power-scalable circuits such as CMOS clocking and reconfigurable current-mode (CM) or voltage-mode (VM) transmitter drivers enable both bandwidth and power-efficiency scaling. A high-density connector and cable, attached to the top-side of the package, enables this high interconnect density. System-level clock calibration and lane redundancy provide design robustness and improve performance.

This paper is organized as follows. Section II discusses the overall link architecture and interconnect topology. In Section III, the details of circuit implementation for this link are described. The system-level duty-cycle and quadrature error correction are discussed in Section IV. The link measurement results are presented in Section V.

II. Link Architecture

Fig. 2 shows the link block diagram for one port. This forwarded-clock link is symmetric full-duplex with 36 data lanes (including four redundant lanes) and one forwarded clock in each direction. As a result, there are a total of 72 data lanes and two forwarded clock lanes per port. The data lanes are organized into "bundles" of six, with a total of six transmitter and receiver bundles per port. The six data signals in each bundle are length-matched on the die as well as on the package, connector, and cable to within 1 mm. This allows the bundle clock to be shared among lanes to conserve power and area. The unified port clock (UPC) is a common design shared by the transmitter and receiver to reduce validation, circuit design, and layout design overhead. A quarter-rate forwarded-clock signaling architecture is chosen for this scalable parallel link. Operating the UPC at 1/4 data rate minimizes channel loss of the forwarded clock (FWCK) and thus improves the received clock quality. By limiting the frequency of clock generation and distribution, a CMOS-based voltage-controlled oscillator (VCO) clock generator is enabled. The UPC injection-locks to the global phase-locked loop (PLL) clock and FWCK for the transmitter and receiver, respectively. The injection-locked port clock rejects the high-frequency jitter from the global PLL/FWCK and provides low-jitter clocks to transmitter/receiver bundle clocking (TBC/RBC) [2]. The CMOS transmitter bundle clock receives 1/4-rate quadrature clock phases from the port clock, doubles the frequency and passively distributes the differential clock signal to the double data rate clocking transmitter lanes within a bundle. The transmitter driver utilizes 3-tap pre-emphasis (TXLE). The CMOS receiver bundle clock receives a 1/4-rate clock from the port clock and rotates the phase of quadrature clock signals for optimum data sampling of four-way interleaved receivers. The 3-tap TXLE along with the receiver continuous-time linear equalizer (CTLE) provides adequate equalization for channel loss of $\leq$18 dB at the Nyquist rate.

The transmitter and receiver floorplans for a pair of bundles are shown in Fig. 3(a) and (b), respectively. In each bundle, there are six differential C4 bumps for data, located in four rows. A pair of I/O power (Vccio) and ground (Vss) bumps are shared between the right and left bundles which supply transmitter and receiver active circuitry. The floorplan of the active circuitry for one transmitter (or receiver) bundle is shown in Fig. 3(c). The bundle circuitry consists of six lanes which share the bundle clock with an optional low-dropout (LDO) regulator. The six lanes within each bundle share test and visibility circuits. Receiver lanes also share a common data acquisition (DAQ) unit for bit error ratio (BER) measurement. On-die differential transmission lines (TLs) route the bundle data signals between C4 bumps and the centrally located transmitter and receiver circuitry. The passive TL routing decouples scaling of C4 bump pitch from process; thus, it reduces package cost by scaling C4 pitch less aggressively than circuit area as process scaling progresses. The data TLs are length-matched within each bundle and are designed as differential coplanar without an explicit microstrip shield [3]. Removing the microstrip shield reduces capacitance and loss. Fig. 4(a) and (b) show differential coplanar TLs with and without a microstrip shield, respectively. Compared to TLs with a microstrip shield, the simulated loss for nonshielded TLs is 0.7 dB lower at 16 Gb/s and also is less frequency dependent [Fig. 4(d)]. To reduce series resistance, the data TLs are constructed of metal 8 in parallel with metal 7 [Fig. 4(c)]. Compared to a nonshielded TL constructed of metal 8 only, the simulated loss is 0.4 dB less at 16 Gb/s. To reduce loss, Metal 6 usage underneath of TLs is limited to power routes. Routing below metal 6 is unrestricted.

Fig. 5 shows the system interconnect topology. Seventy–four pairs of 32AWG micro-twinax cables, configured into six ribbons, are connected to the top-side of two packaged dies using dense Ardent Concepts connectors. This type of connector achieves 40% better density than a conventional CPU socket interface. Compared with [1], the pitch of these connectors is designed less aggressively to reduce cost. The top-side package connector minimizes the breakout area and facilitates lane

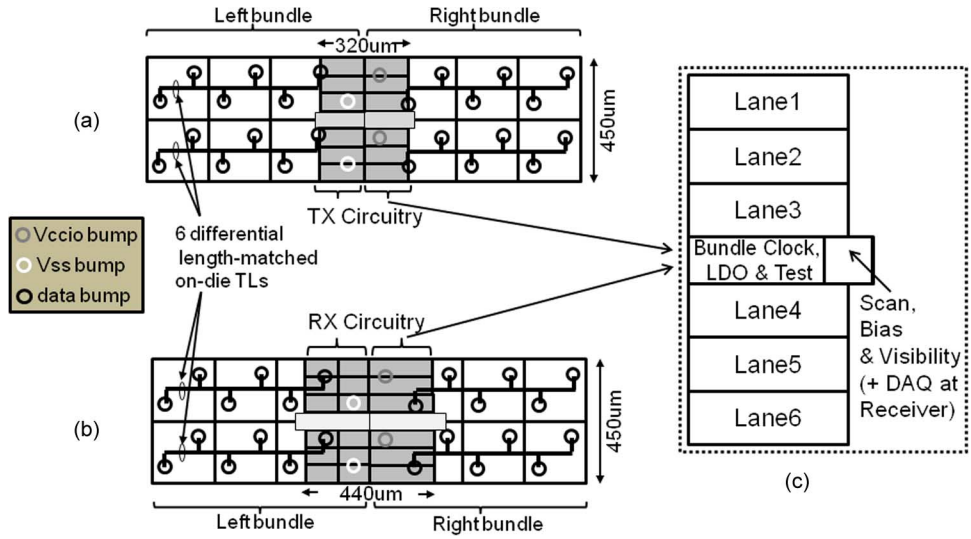

Fig. 3. Floorplan of (a) a pair of transmitter bundles, (b) a pair of receiver bundles, and (c) active circuitry of a transmitter (or receiver) bundle.

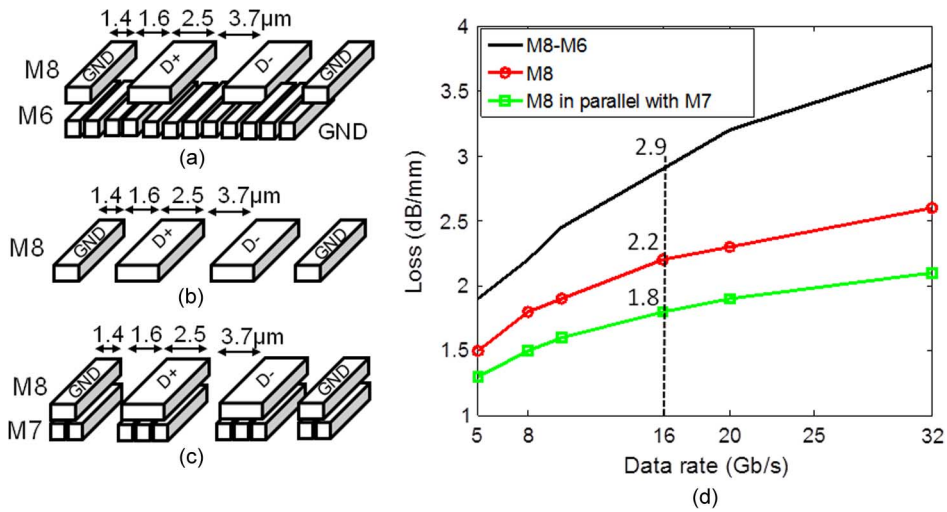

Fig. 4. Data transmission lines. (a) M8 differential coplanar with M6 microstrip shield. (b) M8 differential coplanar. (c) M8 in parallel with M7 differential coplanar. (d) Loss comparison at different data rates.

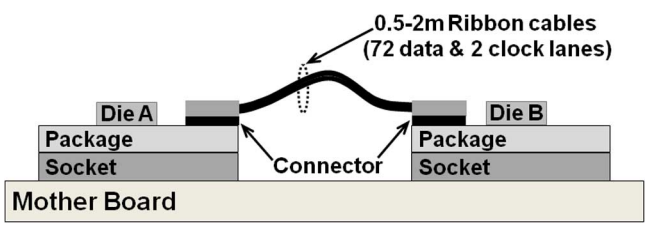

Fig. 5. System interconnect topology.

delay matching. The top-side launch improves signal integrity by bypassing the socket and core package layers, thus reducing loss and crosstalk. To compensate for manufacturing variation and flaws in the channel and circuits, we provide a lane failover mechanism to choose the best 32 out of 36 lanes in each direction of a port based on eye margin measurements.

III. Circuit Implementation

A. Unified Port Clock

The port clock is composed of clock amplifier, injection-locked VCO, and passive clock distribution as shown in Fig. 6. CMOS inverters with resistive feedback amplify low-swing input clocks of the global PLL for the transmitter (FWCK for the receiver) to full rail. For this testchip, a global PLL was not implemented and instead, the clock for the transmitter port clock is provided externally through C4 bumps. The feedback resistor (Rtrim) is implemented with complementary NMOS and PMOS devices in linear region as shown in Fig. 7. The input clock is ac coupled to a dc common-mode voltage that is set to half the I/O supply. The ac-coupling cutoff frequency is set with the capacitor and feedback resistor where Rtrim is smaller than Rbias. At 16-Gb/s operation, where the input clock frequency is 4 GHz, the cutoff frequency is 400 MHz. As I/O data rate and supply voltage scale, the gate-to-source voltage of feedback resistor scales; thus, the cutoff frequency

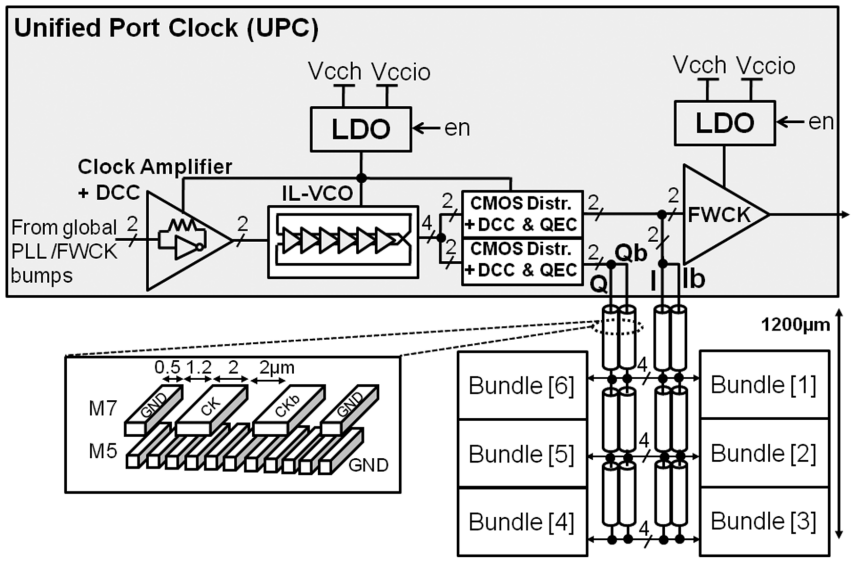

Fig. 6. Transmitter (or receiver) unified port clock.

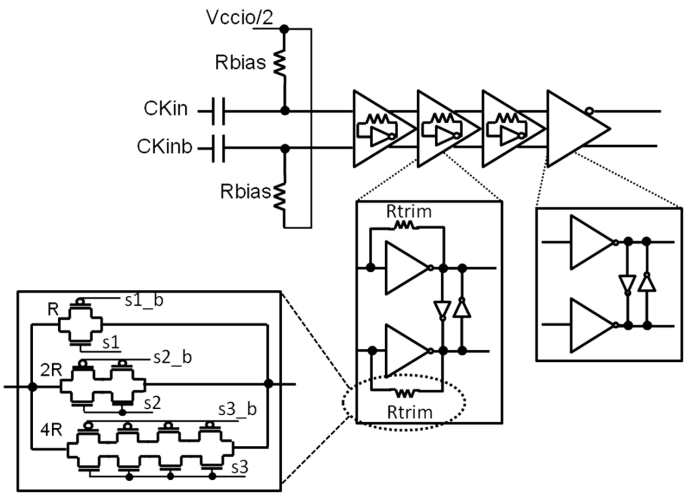

Fig. 7. AC-coupled input clock amplifier with resistive feedback.

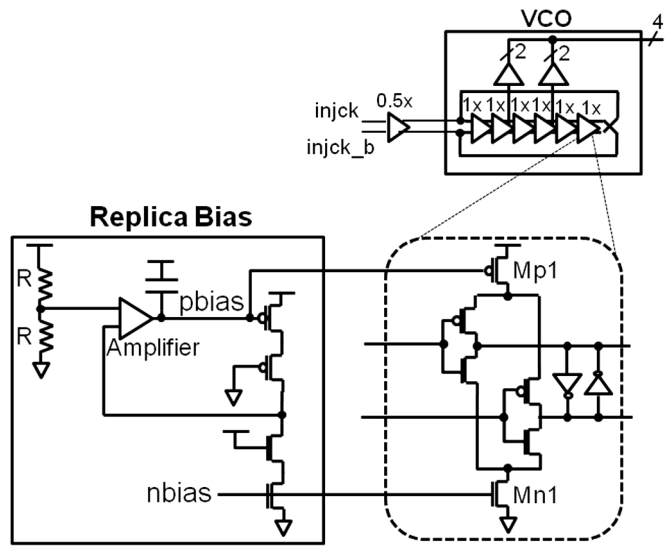

Fig. 8. Six-stage current-starved CMOS IL-VCO with replica bias.

tracks data rate. The cutoff frequency can also be adjusted by tuning feedback resistor switches (s1-s3). The differential clock signals at the output of clock amplifier injection-lock a six-stage current-starved CMOS VCO. The schematic of the VCO and its replica bias circuit is shown in Fig. 8. The VCO frequency is set by tuning the bias (nbias) of NMOS current source device, Mn1, which sets the bias (pbias) for the PMOS current source, Mp1, through the replica bias circuit. The VCO provides quadrature clock phases and operates from 0.5 to 4 GHz, at a quarter of the I/O data rate from 2 to 16 Gb/s. Injecting differential clocks into a VCO whose free-running frequency is not equal to the injection frequency causes systematic error at quadrature clock phases [4], [5]. To mitigate this systematic error, a periodically calibrated loop sets the VCO free-running frequency equal to the frequency of the injecting clock. The loop measures the phase error (or phase shift) between the VCO output clock and the injection clock and tunes the VCO free-running frequency by controlling "nbias" until the phase error goes to zero. This low-bandwidth loop also tracks the low-frequency phase drift due to PVT. Similar to a PLL, an injection-locked VCO (IL-VCO) is a low-pass filter which tracks the received clock jitter up to frequencies equal to the IL-VCO bandwidth and rejects high frequency jitter. The filtering bandwidth of IL-VCO, for zero phase error between the VCO output and injection clock signals, is calculated [5] as

$$f_{-3\ \mathrm{dB}} = \frac{f_{\mathrm{inj}}}{2Q} \times \frac{K}{K+1} \tag{1}$$

where f_{inj} is the clock frequency of injecting signal, Q is the quality factor of the oscillator, and K is the injection strength. The injection strength of this IL-VCO is designed to be 0.5. At 4-GHz input clock, the filtering bandwidth of IL-VCO is about 500 MHz, assuming Q of 1.4 for the CMOS VCO [6].

The residual quadrature phase error due to random variation is detected by an on-die delay analyzer circuit [7]. The quadrature error corrector (QEC) circuit, consisting of two

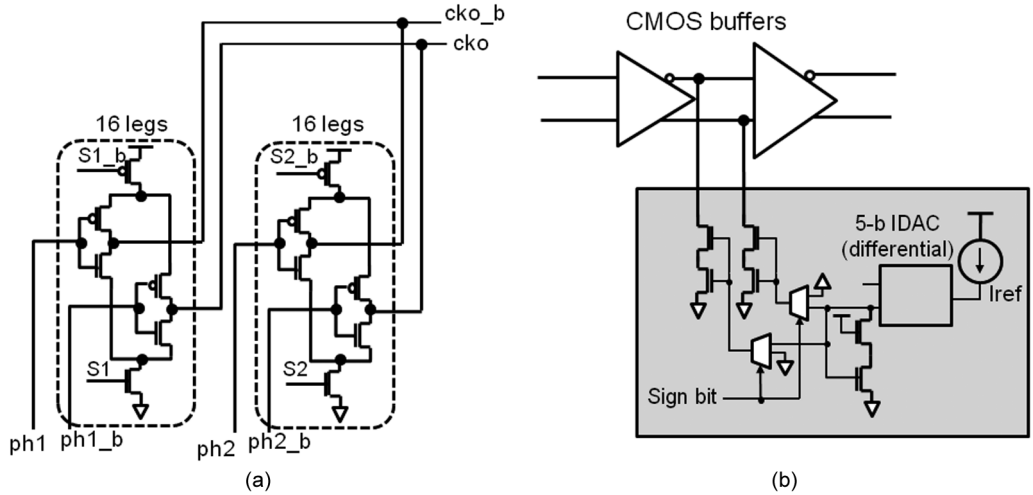

Fig. 9. (a) 16-bit thermometer coded PI with limited-range. (b) 5-bit binary coded DCC with 1-bit sign select.

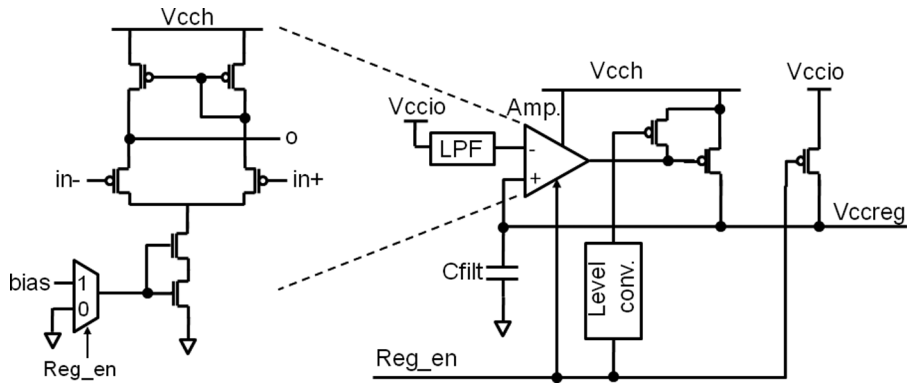

Fig. 10. LDO regulator.

phase interpolators (PIs) implemented inside the clock distributions for I and Q phases, corrects for the quadrature phase error. Each PI has limited range and rotates the phase between the two adjacent VCO clock phases, $\mathrm{ph1} = 0°$ (90°) and $\mathrm{ph2} = 30°$ (120°), by a 16-step thermometer code [Fig. 9(a)]. The duty-cycle error of the input clock is attenuated by the AC-coupled clock amplifier. The duty-cycle error at the clock amplifier and clock distribution is detected with a common duty-cycle detector (DCD), similar to the design in [1]. A duty-cycle corrector (DCC), consisting of a 5-bit current DAC, a sign bit select, and current mirrors, corrects for the detected error [Fig. 9(b)].

The quadrature clock phases, generated in the UPC, are passively distributed to right and left transmitter (or receiver) bundle clock circuits over 1200 μm unterminated TLs (Fig. 6). Unlike [8], the TLs are not terminated to eliminate static power consumption and enable CMOS bundle clocking without the need for level convertors. Reflections in the line for clock frequencies up to 4 GHz are not significant owing to relatively short distance TLs (approximately 1/20th of the clock wavelength) and lossy TLs. The clock TLs are differential coplanar in M7 with an explicit microstrip shield in M5. The explicit M5 shielding increases the loss by 0.6 dB which attenuates the reflection signal further. The total TL loss at 4 GHz is 2.5 dB. Two differential clock phases of the VCO drive the FWCK transmitter. The FWCK transmitter is composed of three parallel data transmitters with shorted outputs. Each data transmitter, composed of a pre-driver and driver, provides up to 500 $\mathrm{mV_{pp-diff}}$. This tunable FWCK transmitter provides larger output swing which overcomes channel loss and improves the received clock signal quality such as duty-cycle error. For the receiver port clock, the FWCK transmitter is disabled.

The UPC including FWCK transmitter, with the exception of the FWCK driver, operates at CMOS swing levels. While CMOS clocking circuits are highly scalable, they are more susceptible to supply noise than CML clocking circuits. To mitigate the supply noise induced jitter, optional LDO regulators can be enabled (Fig. 6). The simulation result indicates the LDO attenuates supply noise by $>$22 dB near the 200 MHz package resonance. If externally induced noise is minimal or the link operates at lower data rates, the LDOs are disabled to save power (Fig. 10). Enabling the LDO degrades the clock power efficiency by 50% and total link power efficiency by less than 4%, owing to port clock amortization.

B. Transmitter Bundle

The block diagram of transmitter bundle clocking and six transmitter lanes is shown in Fig. 11. The quadrature clock phases from the UPC are received by CMOS input buffers. A fully symmetric XOR multiplies the clock frequency to half rate

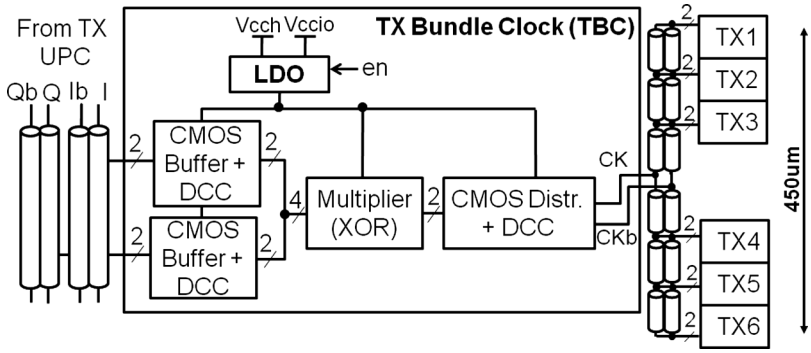

Fig. 11. Transmitter bundle block diagram: transmitter bundle clocking and transmitter lanes.

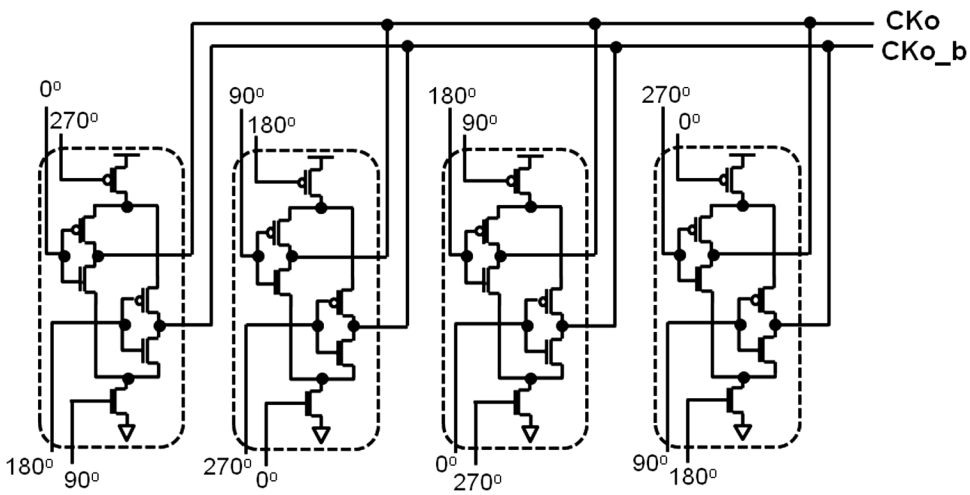

Fig. 12. Fully symmetric XOR to multiply transmitter clock frequency.

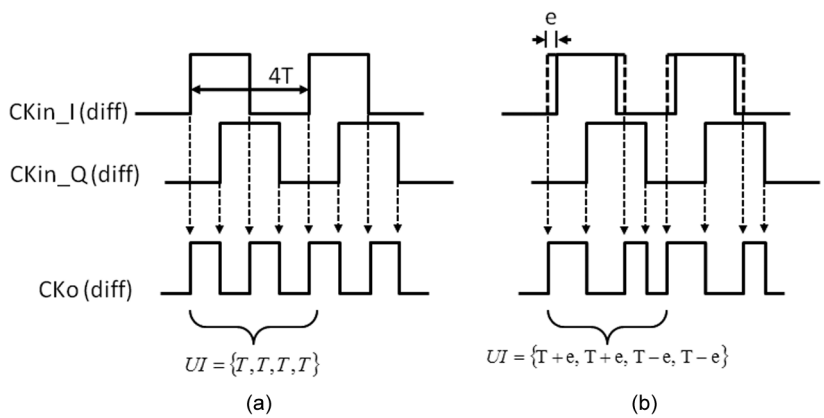

Fig. 13. Impact of duty-cycle error of input clock signals at the multiplier output clock. (a) Zero DCE. (b) Nonzero DCE.

(Fig. 12). This half-rate clock is distributed to six transmitter lanes over an unterminated 450-μm TL, similar to the port clock TL design shown in Fig. 6. The input clock signals to the multiplier suffer from both duty-cycle error and phase error between quadrature clock phases due to random variation of bundle input CMOS buffers and also port clock passive distribution. The quadrature phase error causes duty-cycle error at the output of the multiplier. This error can be corrected as long as the error is within the range of the following DCC circuit located in the bundle clock distribution. Nonetheless, excess duty-cycle error results in more serious clock signal integrity issues. On the other hand, the duty-cycle error present at the multiplier input is translated into nonuniform unit interval (UI) at the multiplier output as shown in Fig. 13. Unlike duty-cycle error which repeats every two UIs, quadrature error repeats every four UIs and thus cannot be corrected with the subsequent bundle distribution DCC. To mitigate this issue, the duty-cycle error of multiplier input signals is corrected by embedding the DCC inside the bundle CMOS input buffers.

The transmitter consists of a 2-to-1 serializer, pre-driver, and reconfigurable driver that includes a 3-tap TXLE. The output driver is divided into three segments—pre-cursor, cursor, and

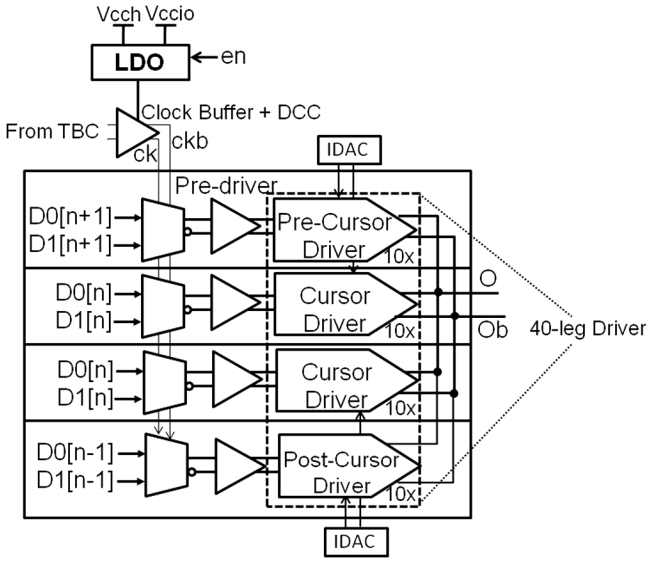

Fig. 14. Reconfigurable transmitter lane with 3-tap linear equalizer and optional LDO for lane clock.

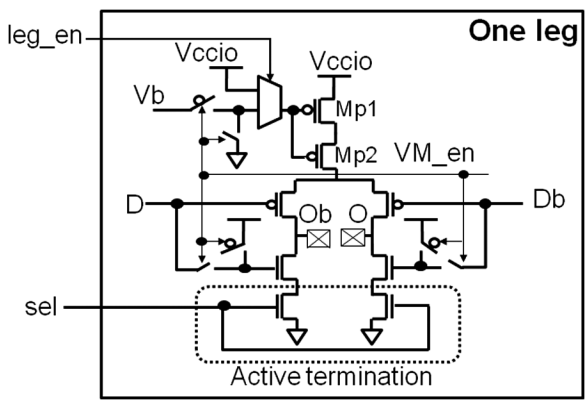

Fig. 15. One leg of reconfigurable CM or VM transmitter driver.

post-curser—to implement the three taps (Fig. 14). The polarity of taps is fixed to be negative, and, thus, it can equalize specific range of channel behavior. The advantage of this 3-tap segmented TXLE architecture is to reduce power overhead of the fully segmented equalizer [9] by enabling the optimal fanout in pre-drivers and avoiding the minimum device size limits, similar to [1]. The transmitter output stage is a reconfigurable CM or VM driver [10]. One leg of this 40-leg reconfigurable driver is shown in Fig. 15. This reconfigurable design allows trading off better PSRR and CMRR performance of a CM driver for better power performance of VM driver at different rates. At lower data rates, switching to VM driver saves power. When the driver is reconfigured as a CM driver, the TXLE coefficients are set with two 4-bit current DACs which control the bias (Vb) of tail current devices, Mp1 and Mp2, of pre-cursor, cursor, and post-curser segments. To reduce pad capacitance and improve bandwidth, digitally controllable NMOS devices in the triode region are embedded into the reconfigurable driver which set the transmitter termination ("sel" signal). For the VM driver, the TXLE coefficients are set by turning on/off the tail devices. The pull up and pull down active termination values are controlled independently by setting "leg_en" and "sel" signals, respectively. A local DCC circuit at the transmitter lane corrects for any residual duty-cycle error in the clock signal which drives the final 2-to-1 serializer (Fig. 14). Similar to the port clock, optional LDOs, located in the transmitter bundle and local clocking circuits at the transmitter lanes can be enabled to reject supply noise induced jitter (Figs. 11 and 14). Enabling the LDO degrades total link power efficiency by $<4\%$.

C. Receiver Bundle

The block diagram of receiver bundle clocking and six receiver lanes is shown in Fig. 16. Quadrature phases of the port clock are distributed to receiver bundle clocking. Two 360-degree CMOS quadrature phase interpolators, shared across six lanes in a bundle, rotate the quadrature clock phases for optimum sampling of the 4-way interleaved receiver. The quadrature PI consists of four sets of 32-leg current-starved delay cells, each driven by one of the quadrature clock phases (Fig. 17). Two bits of coarse control $(\mathrm{sel\_0^\circ/90^\circ/180^\circ/270^\circ})$ select two out of four delay cells. Fine control is set by a 32-step thermometer code (sel/selb) which rotates the clock phase by 32 steps/UI. This quadrature PI eliminates preceding clock phase selection multiplexers and achieves a monotonic phase transfer [11]. However, the coarse separation of clock phases causes systematic nonlinearity at the output of quadrature PI. To mitigate the nonlinearity, the slew-rate of PI input clocks is set with current-starved delay cells. To track clock frequency/data rate, both slew-rate control and PI delay cells are replica biased to the port clock IL-VCO. Nonetheless, PI nonlinearity is inevitable due to both random and systematic variation. The on-die delay analyzer is used to characterize the PI phase transfer function [7] and thus, to calibrate the measured eye timing margin. The quadrature clock phases are distributed to receiver lanes over unterminated 450 um TLs, similar to the port clock TL design shown in Fig. 6. The quadrature phase and duty-cycle errors are corrected with QEC and DCC embedded into the receiver bundle clock distribution. The QEC, consisting of two limited-range PIs [Fig. 9(a)], corrects for the quadrature error by rotating input phases of the PIs [$\mathrm{ph1} = 0^\circ$ (90°) and $\mathrm{ph2} = 90^\circ$ (180°)].

The four-way interleaved receiver lane is composed of a continuous-time linear equalizer (CTLE) and four sets of offset-compensated data sampler slices (Fig. 18). A 6-bit current DAC compensates for self-offset of the CTLE and provides additional range for eye margining. Each sampler has an independent 6-bit current DAC for offset control. Tunable per-slice deskew (ΔT) compensates for phase/delay error among quadrature clock phases up to 1/4 UI. If the interconnect skew among six lanes in a bundle exceeds 1 mm, an optional per-lane PI circuit with limited range is enabled to correct for lane-to-lane skew [Fig. 9(a)]. The PI provides up to $\pm 1/2$UI phase deskew by rotating the phase between I and Q with a 16-step thermometer code at the cost of 10% additional receiver lane power. To detect quadrature phase errors in a lane or lane-to-lane skew, a software-based low-bandwidth CDR employing a receiver eye margining technique is used. Per-slice deskew (or per-lane PI in case of lane-to-lane skew) is swept until the optimum aggregate received eye for a bundle is achieved. This low-bandwidth CDR runs periodically to track

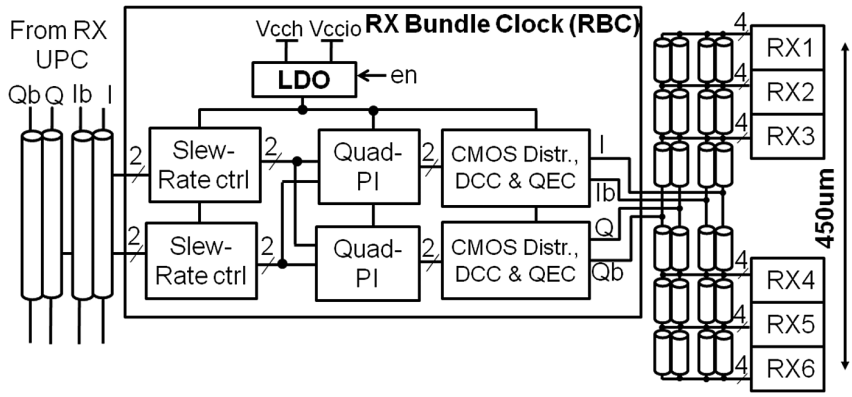

Fig. 16. Receiver bundle block diagram: receiver bundle clocking and receiver lanes.

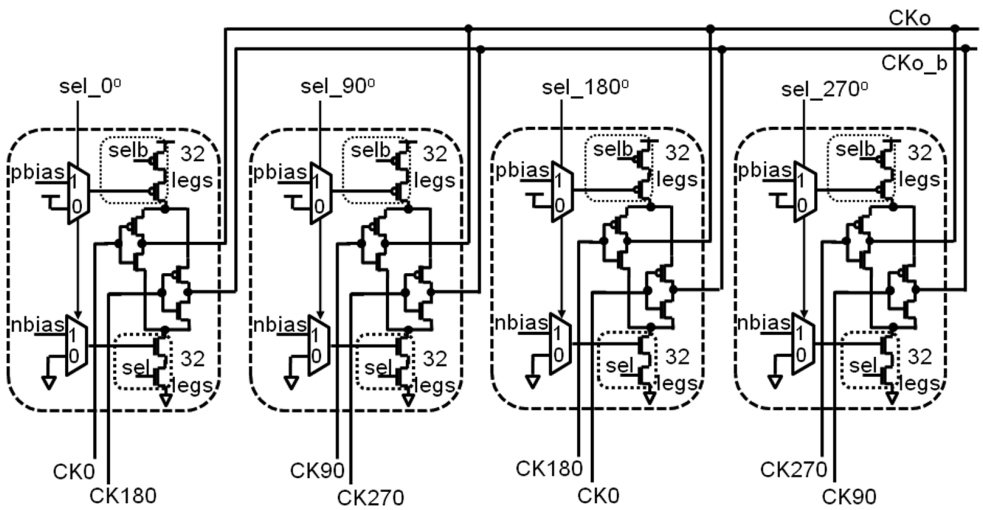

Fig. 17. 360° quadrature phase interpolator.

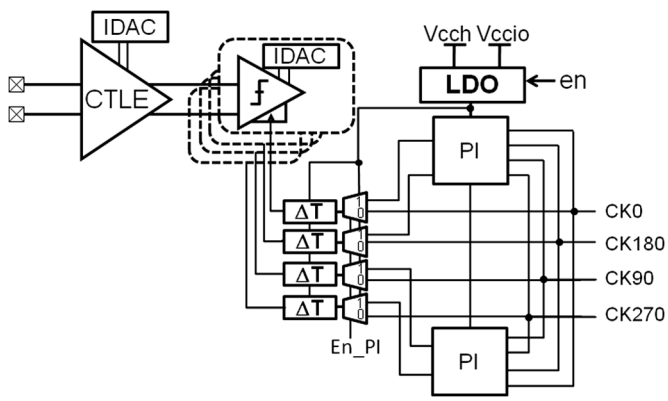

Fig. 18. Four-way interleaved receiver lane with CTLE, optional deskew, and LDO for lane clock.

voltage and temperature (VT) drift of clock with respect to data. Similar to the port clock and transmitter bundle, optional LDOs can be enabled to reject supply noise induced jitter at the receiver bundle clock and local clocking circuits at receiver lanes, at the cost of $<4\%$ of total link power efficiency (Figs. 16 and 18).

IV. System-Level Duty-Cycle and Quadrature Error Correction

To achieve a low-power scalable link, the clocking circuits take full advantage of process scaling by employing small transistors. However, random device variation causes clock nonidealities such as duty-cycle and quadrature phase error. As discussed in Section III, the proposed clocking architecture employs local detector and corrector circuits at different levels of the clock hierarchy. The proposed hierarchical detection and correction mechanism prevents error accumulation within the clock hierarchy and ensures clock signal integrity. Additionally since the required range of error correction per DCC/QEC is smaller, the correction circuits can be designed with finer resolution. Nonetheless, due to variation limited resolution and accuracy of the detectors circuits, residual error exists following local detection and correction. This residual error can severely degrade clock quality and robustness. For instance, the duty-cycle error of the transmitter clock leads to high-frequency transmit jitter and subsequent ISI amplification [12]. Similarly, nonuniform UI spacing caused by duty-cycle error at the multiplier input clocks is amplified across the lossy channel.

To improve transmitter clock quality, we developed a "system-level" clock calibration scheme in which the existing receiver is used to measure the transmitter clock error, where

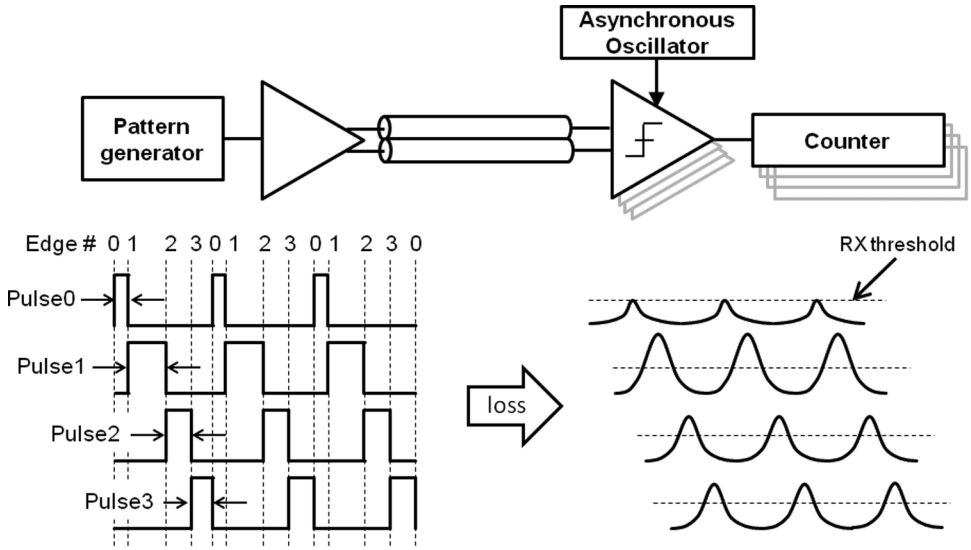

Fig. 19. Asynchronous sampling of "1000" pattern to achieve precise phase placement of TX symbols.

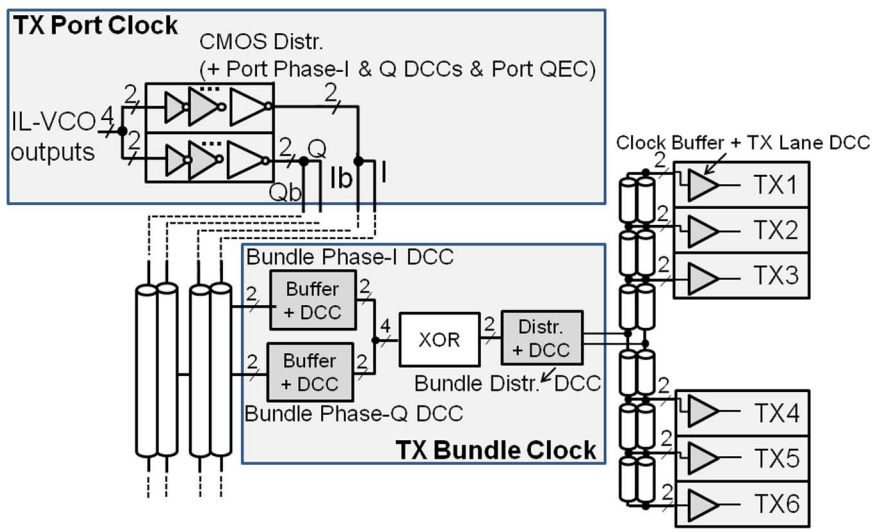

Fig. 20. Hierarchical DCC and QEC at transmitter port, bundle, and lane clock.

it has the most detrimental effect. The proposed system-level transmitter clock calibration takes advantage of the channel loss to amplify clock errors and thus, achieves better phase detection accuracy (Fig. 19). The calibration mode is enabled by programming a repeating pattern of "1000" using the pattern-generation capabilities of the transmitter. The receiver clock frequency is set to be asynchronous with respect to the transmit clock by disabling the injection lock at receiver. Asynchronously sampling the received pattern with a single data sampler ensures uniformly distributed time samples across the waveform similar to [7]. A counter following the sampler measures the number of '1's sampled versus the overall number of samples taken (referred to as "Distribution of '1's"). The relative placement of Edges 0, 1, 2, and 3 (or Pulse0, 1, 2 and 3) can be measured based on the sampled distribution of '1's by setting the repeating patterns to "1000," "0100," "0010," and "0001," respectively. Based on measured error, the clock calibration controls are tuned until edges 0–3 are equally spaced. To ensure consistent measurements between four different patterns and eliminate the possibility of sampler variation corrupting the calibration, the distribution counter of one of the four interleaved samplers is used. Additionally, all pre-emphasis is disabled so that only two adjacent edges are considered when measuring the width of the pulse. The receiver threshold level is adjusted based on the loss of the channel to ensure a strong, albeit somewhat nonlinear, correlation between the transmit-side pulse width and the measured received "1" distribution.

The hierarchical calibration controls for the transmitter clock are shown in Fig. 20. Given the multitude of DCC and QEC controls, it is crucial to calibrate in an order that produces consistent and optimal results, even in the presence of large variation. Table I summarizes the calibration sequence in which each control is tuned to optimize the listed target. At link power/wake up, the local transmitter clock calibration followed by the system-level calibration is performed. To reduce the system-level calibration time at each calibration step, the controls in the corresponding hierarchy and any "downstream" controls are set to

TABLE I
CLOCK CALIBRATION SEQUENCE

Calibration Control	Number of Instances	Optimization Target
Port QEC	1	Σ_{port}(Pulse0+Pulse2) = Σ_{port}(Pulse1+Pulse3)
Port phase-I DCC	1	Σ_{port}Pulse1 = Σ_{port}Pulse3
Port phase-Q DCC	1	Σ_{port}Pulse0 = Σ_{port}Pulse2
Bundle phase-I DCC (multiplier input)	6	Σ_{bundle}Pulse1 = Σ_{bundle}Pulse3
Bundle phase-Q DCC (multiplier input)	6	Σ_{bundle}Pulse0 = Σ_{bundle}Pulse2
Bundle distr. DCC	6	Σ_{bundle}(Pulse0+Pulse2) = Σ_{bundle}(Pulse1+Pulse3)
TX lane DCC	36	Σ_{lane}(Pulse0+Pulse2) = Σ_{lane}(Pulse1+Pulse3)

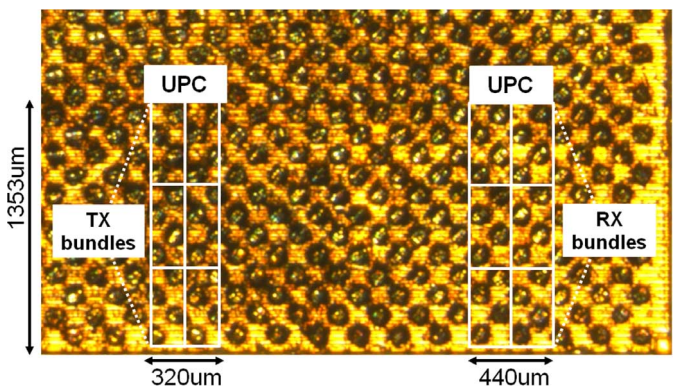

Fig. 21. Die photograph and link floorplan.

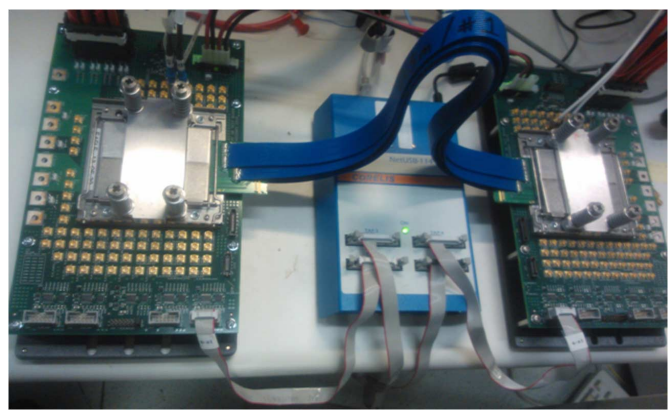

Fig. 22. Link measurement setup with 50-cm micro-twinax cable.

the value based on local calibration. Then, the error between the target and current value is calculated by the on-die distribution counter supplemented with off-chip software control. The calibration starts with adapting the Port QEC control to set the measured average (among 36 lanes) of "Pulse0 + Pulse2" to be equal to "Pulse1 + Pulse3". Once this goal is achieved, the calibration algorithm proceeds to the next level of hierarchical control where Port DCCs for both I and Q phases are tuned. The calibration continues until all transmitter clock controls at port, bundles and lanes are calibrated. Once the system-level calibration is complete, the local calibration at the transmitter lanes is activated periodically to correct for any error due to VT.

For the receiver clock at link power/wake up, the local clock calibration at the receiver port and bundle clock hierarchy is performed. Then, per-slice deskew (or per-lane PI in case of lane-to-lane skew) for six receiver lanes in a bundle are tuned periodically for an optimum aggregate bundle received eye.

V. MEASUREMENT RESULTS

One port of the link, including 72 data lanes and two FWCK signals, was implemented in a 32-nm low-power CMOS process (Fig. 21). Fig. 22 shows the measurement setup for the one-port interface between two packaged dies, connected by 50-cm-long micro-twinax ribbon cables. The setup includes power supplies, scan control interface, and two differential transmitter input clock signals. The measurements are automated through software which controls scan. The

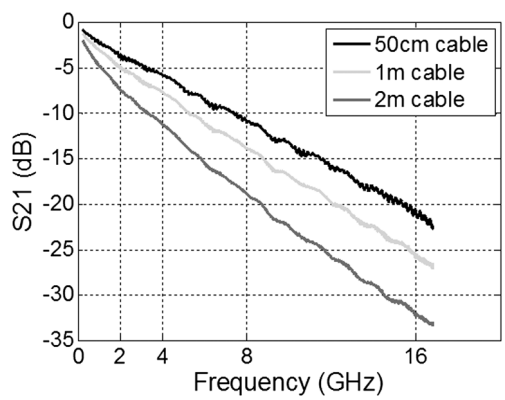

Fig. 23. Channel response, including package routings, connectors, pad capacitance, and 50-cm, 1-m, or 2-m-long micro-twinax ribbon cables.

channel response, including package routings, connectors, pad capacitance, and micro-twinax ribbon cables ranging from 50 cm to 2 m long, is shown in Fig. 23. At the 8-GHz Nyquist frequency, the loss is 11, 14, and 19 dB, respectively. The on-die TL increases the loss by 2 dB at 8 GHz. The total loss at the 8-GHz Nyquist frequency is 13 dB across 50-cm-long cables. Measurement results indicate the receiver front-end bandwidth limitation which increases the loss by 4 to 5 dB at the 8-GHz Nyquist frequency. Using 50-cm micro-twinax ribbon cables, the link was tested at per-lane data rates from

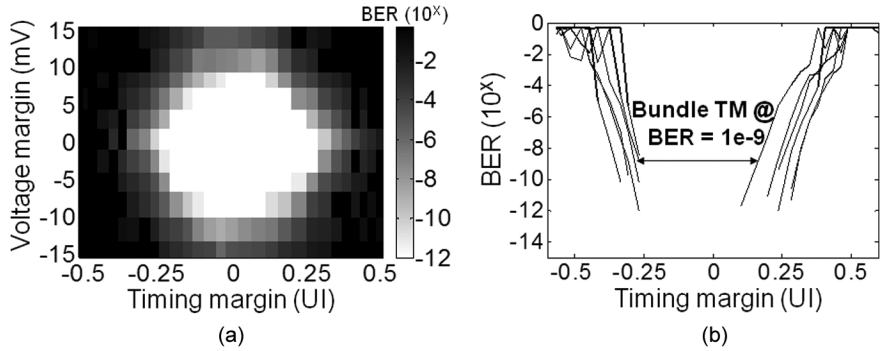

Fig. 24. (a) Measured eye and (b) bundle timing margin at 16 Gb/s.

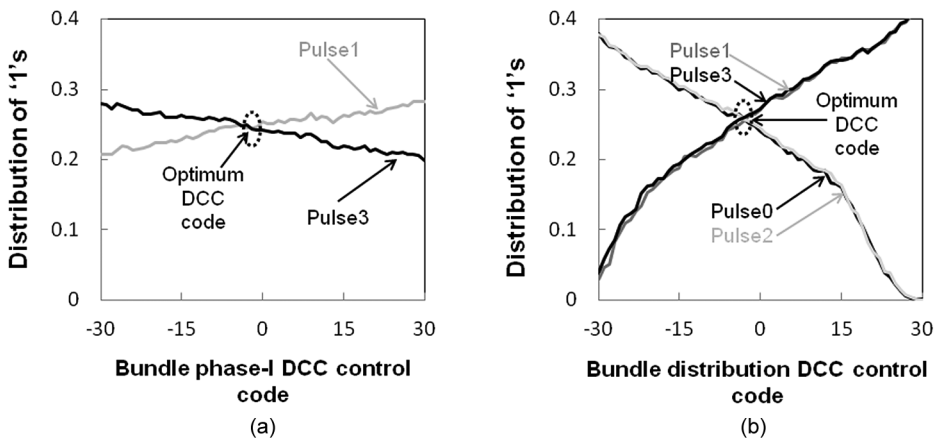

Fig. 25. System-level transmitter clock calibration measurements. (a) Distribution of 1's versus bundle phase-I DCC control code. (b) Distribution of 1's versus bundle distribution DCC control code.

2 to 16 Gb/s while scaling I/O power supply from 0.6 to 1.08 V, respectively. At any given data rate, the best 32 lanes out of 36 lanes (in each link direction) were chosen based on eye margining while the other four lanes are disabled. In other words, the link redundancy enabled better performance/power efficiency by disabling the lanes with marginal performance. As a result, a total of 64 data lanes operated simultaneously with $\mathrm{BER} < 10^{-12}$. The measured BER eye and timing margin (TM) were calibrated by measuring the PI phase characteristic with 250-fs accuracy using the on-die delay analyzer circuit [7]. Fig. 24(a) and (b) shows the measured BER eye and bundle TM at 16 Gb/s per-lane data rate, respectively.

Fig. 25 demonstrates two measured examples of system-level transmitter clock calibration. As the DCC digital control code is swept, the widths of "Pulse0" through "Pulse3" are measured using distribution of 1's. Fig. 25(a) shows the UI pulse widths of "Pulse1" and "Pulse3" as the bundle phase-I DCC control code is swept from -31 to $+31$. The optimum code for the bundle phase-I DCC is achieved when the error between the two pulse widths is minimum. Fig. 25(b) demonstrates changing the multiplier output duty cycle controlled by bundle distribution DCC control code. The optimum DCC code is achieved when the error between "Pulse0 + Pulse2" and "Pulse1 + Pulse3" is minimum.

The measured power consumption from a 1.08-V supply was 2.7 W for a total aggregate bandwidth of 1.024 Tb/s, corresponding to a power efficiency of 2.6 pJ/bit. By aggressively

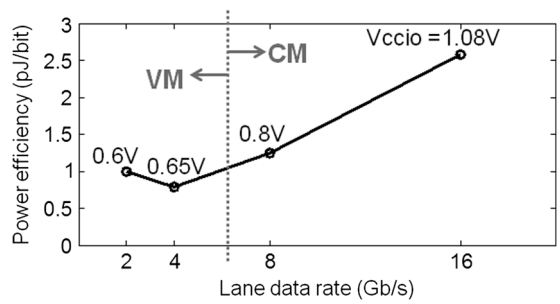

Fig. 26. Power efficiency versus lane data rate (across 50-cm micro-twinax cables).

scaling supply from 1.08 to 0.6 V and reconfiguring the transmitter to VM mode at rates $\leq$4 Gb/s, the link power efficiency improves to 1.25, 0.8, and 1 pJ/bit at 8, 4, and 2 Gb/s, respectively (Fig. 26). Table II summarizes link power breakdown at different data rates across a 50-cm link. When the transmitter is reconfigured to CM mode at 16 Gb/s, 45% of the power is dissipated in the transmitter, 30% in receiver, and 25% in clocking circuits, including port and bundle clocks. While the transmitter is reconfigured as a VM driver at 2 Gb/s, power dissipated by the transmitter reduces to 31% of the overall power. On the other hand, the percentage of total power dissipated by receiver increases to 45%, mainly due to performance degradation of the receiver front-end at a lower supply voltage. The key metrics of

TABLE II
POWER BREAKDOWN AT DIFFERENT DATA RATES (ACROSS 50-cm MICRO-TWINAX CABLES)

Data rate (Gb/s)	2	4	8	16
TX driver (pJ/bit)	0.23	0.16	0.35	0.7
TX pre-driver & 2-to-1 serializer (pJ/bit)	0.08	0.11	0.16	0.46
RX CTLE (pJ/bit)	0.22	0.14	0.18	0.27
RX samplers (pJ/bit)	0.23	0.19	0.27	0.52
TBC (pJ/bit)	0.07	0.05	0.08	0.19
RBC (pJ/bit)	0.07	0.07	0.1	0.23
UPC+ TX FWCK (pJ/bit)	0.09	0.07	0.11	0.21
Total (pJ/bit)	**0.99**	**0.79**	**1.25**	**2.58**

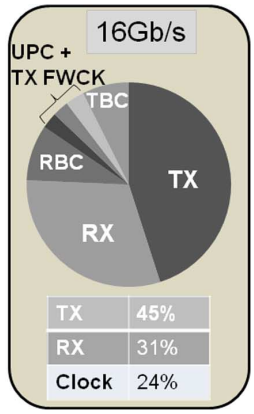

TABLE III
KEY LINK METRICS AT 16 Gb/s ACROSS 50-cm MICRO-TWINAX CABLES

TX swing	360mV$_{pp\text{-}diff}$
TX jitter	600fs-rms
RX input-referred noise	1.2mV-rms
TX pad cap	380fF
RX pad cap	350fF
Active area (mm²/lane)	Total = 0.039 TX = 0.014 RX = 0.020 Port clock = 0.005

the link, measured at 16 Gb/s per-lane data rate across 50-cm micro-twinax cables, are tabulated in Table III. The transmitter output swing is set to 360 $\mathrm{mV_{pp-diff}}$. The measured transmitter jitter and the receiver input-referred noise were measured at 600 fs-rms and 1.2 mV-rms, respectively. The measured transmitter and receiver pad capacitance were 380 and 350 fF, respectively. The link was also measured over longer cable lengths of 1 and 2 m. Across the 1-m channel, the maximum per-lane data rate of 16 Gb/s was transferred with $\mathrm{BER} < 10^{-12}$. The measured power efficiency was 3.2 pJ/bit from a 1.15-V I/O power supply, resulting in a 23% degradation compared with the measured power efficiency across a 50-cm channel. For 2-m cables, the maximum per-lane data rate was measured to be 14 Gb/s with 3.6-pJ/bit power efficiency from a 1.15-V supply. To overcome higher loss of longer channels, the transmitter swing was increased from 360 $\mathrm{mV_{pp-diff}}$ (for a 50-cm cable) to 500 $\mathrm{mV_{pp-diff}}$ (for a 1- and 2-m cables). The elevated I/O power supply of 1.15 V provided 80 mV more headroom at the transmitter driver and improved the receiver front-end bandwidth. The receiver per-lane PIs were disabled for all of the measurements. The measurements show that LDOs are not required to meet the link performance targets in this test platform. Enabling all LDOs, other than the local transceiver LDO circuits, degrades the link power efficiency by 10%. To verify the LDO functionality, we injected supply noise at different frequencies and measured the link TM. The measurement results with LDOs enabled indicate that TM at 16 Gb/s degrades by 6% and 11% when 7 and 14 mV-pp of supply noise at the package resonance of 200 MHz are injected, respectively. Disabling the LDOs, TM at 16 Gb/s degrades by more than 40% in the presence of 7-mV-pp supply noise at 200 MHz. Table IV summarizes the link performance across a 50-cm channel and compares against other links. We have demonstrated a dense I/O with nonlinear power efficiency improvement of >3X as the per-lane data rate is reduced by 4X from 16 to 4 Gb/s. The power efficiency of the proposed link at 16 Gb/s is 22% and 40% lower than [13] and [14], respectively.[1] Compared with recently reported aggregate bandwidths of 512 and 470 Gb/s, power is 6$\times$ lower than [15] and 10% better than [1] at longer channel length. The link area efficiency is 2.8X better than the reported link in [1].

VI. SUMMARY

We developed a scalable 64-lane per-port parallel link, targeting high performance graphics and computing systems. The top-side package interconnect enabled this dense and low-loss electrical link that transfers data across a channel length ranging from 50 cm to 2 m. Codesigning of circuits and on-die and off-chip interconnects allowed the amortization of port and bundle clocking circuits and decoupling of I/O circuitry from C4 bumps; thus, saving power, area, and cost. CMOS clocking with optional LDOs and reconfigurable CM/VM transmitter driver enabled scalable per-lane data rate from 2 to 16 Gb/s and nonlinear power efficiency scaling at lower data rates. At the maximum per-lane data rate of 16 Gb/s, the link provides 1.024 Tb/s of bandwidth across a 50-cm channel while consuming 2.7 W from a 1.08-V power supply. For bandwidth demand 1/4 of the peak bandwidth, the proposed link consumes slightly over 0.2 W while operating at a 4-Gb/s per-lane data rate. At longer channel lengths of 1 and 2 m, the link provides up to 1.024 and 0.896 Tb/s of bandwidth, respectively, while consuming approximately 3.3 W from a 1.15-V power supply. The *in situ* measurement and system-level calibration of the critical link metrics such as transmitter duty-cycle error improved the link performance and power. Link redundancy enabled design robustness by mitigating the impact of circuit and channel defects.

ACKNOWLEDGMENT

The authors thank B. Horine, M. Leddige, V. Baca, T. Boyd, C. Beall, H. Do, K. Aygun, C. Chiu, K. Krutsinger, L. Avery, S. Spangler, V. P, R. Lupien, D. Finan, D. Jenkins, P. Stolt, A. Wilson, T. Nguyen, H. Wilson, A. Panikkar, D. Klowden,

[1]Excluding the power consumption of the serializer, deserializer, global PLL, and clock distribution of the references.

TABLE IV
PERFORMANCE COMPARISON SUMMARY

	[16]	[1]	[15]	[13]	[14]	This work (across 50cm cables)	
Total bandwidth (Gb/s)	75	470	512	128	256/320	256	1024
Lane data rate (Gb/s)	12.5	10	16	16	16/20	4	16
Power eff. (pJ/bit)	0.98	1.4/1.51	13/8[c]	4.06	5.3/6.1	0.79	2.58
Loss @ Nyquist (dB)	12	15[a]/19[a]	15	15	12/15	5[a]	13[a]
		8/12				4	11
Area eff. (mm^2/Gb/s)	0.02	0.007	0.027	0.01	0.012/ 0.01	0.01	0.0025
		0.017[b]				0.034[b]	0.0085[b]
CMOS technology (nm)	65	45	65	40	40	32	

[a] Includes on-die transmission line loss.
[b] Includes C4 bump area.
[c] Controller/DRAM.

M. Zeid, A. Sethuram, K. Vakil, V. De, R. Forand, W-H. Wang, G. Taylor, and M. Haycock.

REFERENCES

[1] F. O'Mahony, J. E. Jaussi, J. Kennedy, G. Balamurugan, M. Mansuri, C. Roberts, S. Shekhar, R. Mooney, and B. Casper, "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] A. Mirzaei, M. E. Heidari, R. Bagheri, S. Chehrazi, and A. A. Abidi, "The quadrature LC oscillator: A complete portrait based on injection locking," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1916–1932, Sep. 2007.

[3] B. Kleveland, X. Qi, L. Madden, T. Furusawa, R. W. Dutton, M. A. Horowitz, and S. Simon Wong, "High-frequency characterization of on-chip digital interconnects," *IEEE J. Solid-State Circuits*, vol. 37, no. 6, pp. 716–725, Jun. 2002.

[4] P. Kinget, R. Melville, D. Long, and V. Gopinathan, "An injection-locking scheme for precision quadrature generation," *IEEE J. Solid-State Circuits*, vol. 37, no. 7, pp. 845–851, Jul. 2002.

[5] S. Shekhar, M. Mansuri, F. O'Mahony, G. Balamurugan, J. E. Jaussi, J. Kennedy, D. J. Allstot, R. Mooney, and B. Casper, "Strong injection locking in low-Q LC oscillators: Modeling and application in a forwarded-clock I/O receiver," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 8, pp. 1818–1829, Aug. 2009.

[6] B. Razavi, "A study of phase noise in CMOS oscillators," *IEEE J. Solid-State Circuits*, vol. 31, no. 3, pp. 331–343, Mar. 1996.

[7] M. Mansuri, B. Casper, and F. O'Mahony, "An on-die all-digital delay measurement circuit with 250 fs accuracy," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2012, pp. 98–99.

[8] F. O'Mahony, M. Mansuri, B. Casper, J. E. Jaussi, and R. Mooney, "A low-power PLL and repeaterless clock distribution network for a 20 Gb/s link," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2006, pp. 29–30.

[9] B. Casper, J. E. Jaussi, F. O'Mahony, M. Mansuri, K. Canagasaby, J. Kennedy, E. Yeung, and R. Mooney, "A 20 Gb/s forwarded clock transceiver in 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 90–91.

[10] R. Bashirullah, W. Liu, R. Cavin, III, and D. Edwards, "A 16 Gb/s adaptive bandwidth on-chip bus based on hybrid current/voltage mode signaling," *IEEE J. Solid-State Circuits*, vol. 41, no. 2, pp. 461–473, Feb. 2006.

[11] C. Kromer, G. Sialm, C. Menolfi, M. Schmatz, F. Ellinger, and H. Jackel, "A 25-Gb/s CDR in 90-nm CMOS for high-density interconnects," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2921–2929, Dec. 2006.

[12] 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.

[13] A. Amirkhani, K. Kaviani, A. Abbasfar, F. Shuaeb, W. Beyene, C. Hishino, C. Madden, K. Chang, and C. Yuan, "A 4.1 pJ/b 16 Gb/s coded differential bidirectional parallel electrical link," in *IEEE ISSCC Dig. Tech. Papers*, 2012, pp. 138–139.

[14] K. Kaviani, T. Wu, J. Wei, A. Amirkhany, J. Shen, T. J. Chin, C. Thakkar, W. T. Beyene, N. Chan, C. Chen, B. R. Chuang, D. Dressler, V. P. Gadde, M. Hekmat, E. Ho, C. Huang, P. Le, Mahabaleshwara, C. Madden, N. K. Mishra, L. Raghavan, K. Saito, R. Schmitt, D. Secker, X. Shi, S. Fazeel, G. S. Srinivas, S. Zhang, C. Tran, A. Vaidyanath, K. Vyas, M. Jain, K.-Y. K. Chang, and X. Yuan, "A tri-modal 20-Gb/s link differential/DDR3/GDDR5 memory interface," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 926–937, Apr. 2012.

[15] H. Lee, K.-Y. K. Chang, J.-H. Chun, T. Wu, Y. Frans, B. Leibowitz, N. Nguyen, T. J. Chin, K. Kaviani, J. Shen, X. Shi, W. T. Beyene, S. Li, R. Navid, M. Aleksic, F. S. Lee, F. Quan, J. Zerbe, R. Perego, and F. Assaderaghi, "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.

[16] K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, E. Suzuki, T. Takemoto, F. Yuki, and T. Saito, "A 12.3 mW 12.5 Gb/s complete transceiver in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 45, no. 12, pp. 2838–2849, Dec. 2010.

Mozhgan Mansuri (S'00–M'03) received the B.S. and M.S. degrees in electronics engineering from Sharif University of Technology, Tehran, Iran, in 1995 and 1997, respectively, and the Ph.D. degree in electrical engineering from University of California, Los Angeles, CA, USA, in 2003.

She was a Design Engineer with Kavoshgaran Company, Tehran, Iran, where she worked on the design of 46–49-MHz cordless and 900-MHz cellular phones from 1997 to 1999. Since 2003, she has been with Electrical Signaling Research of Intel Labs, Hillsboro, OR, USA. Her research interests include low-power low-jitter clock synthesis/recovery circuits (PLLs and DLLs), high-speed, low-power I/O links, and variation-tolerant clocking circuits.

Dr. Mansuri was the recipient of the 2010 Transactions on Circuits and Systems Darlington Best Paper Award.

James E. Jaussi (M'01) received the B.S. and M.S. degrees in electrical engineering from Brigham Young University, Provo, UT, USA, in 2000.

In January 2001, he joined Intel's Circuit Research Laboratory, Hillsboro, OR, USA. His research interests include high-speed CMOS transceivers and low-loss interconnect. He has worked with I/O industry standard efforts on link jitter budging, system clocking architectures and channel equalization.

Mr. Jaussi currently co-chairs the MIPI Alliance M-PHY Electrical Sub-Group. He was the recipient of the 2010 TRANSACTIONS ON CIRCUITS AND SYSTEMS Darlington Best Paper Award and the 2010 Outstanding Speaker Award at Intel's Developer Forum.

Joseph T. Kennedy (S'88–M'91) received the B.S. degree in electrical and computer engineering from Oregon State University, Corvallis, OR, USA, in 1991.

He is a Senior Circuits Researcher with Electrical Signaling Research, Intel Labs, Hillsboro, OR, USA. Since joining Intel in 1995, his responsibilities have included all aspects of research and design of high-speed mixed-signal circuits and I/O systems. Prior to joining Intel, he spent four years with Lattice Semiconductor, where he worked as a lead circuit designer developing data-path circuits and I/O interfaces for electrically programmable logic components.

Tzu-Chien Hsueh (S'06–M'10) received the B.S. and M.S. degrees from National Taiwan University, Taipei, Taiwan, in 1999 and 2001, respectively, and the Ph.D. degree from the University of California, Los Angeles, CA, USA, in 2010, all in electrical engineering.

From 2001 to 2006, he was a Design Engineer with Silicon Integrated Systems and Chip Advanced Technology Corporations, Hsinchu, Taiwan, where he worked on Delta-Sigma A/D converters and Gigabit Ethernet I/O systems. Since 2010, he has been with Electrical Signaling Research, Intel Labs, Hillsboro, OR, USA. His research interests include high-speed mixed signal circuits, I/O links, and data converters.

Sudip Shekhar (S'00–M'10) received the B.Tech. degree (Hons.) in electrical and computer engineering from the Indian Institute of Technology, Kharagpur, India, in 2003, and the M.S. and Ph.D. degree in electrical engineering from the University of Washington, Seattle, WA, USA, in 2005 and 2008, respectively.

He is currently a Research Scientist with Circuits and Systems Research, Intel Corporation, Hillsboro, OR, USA. His research interests include circuits for high-speed electrical and optical I/O interfaces, and wireless transceivers.

Dr. Shekhar was a recipient of the IEEE Solid-State Society Predoctoral Fellowship (2007–2008), Intel Foundation Ph.D. Fellowships (2006–2008), Analog Devices Outstanding Student Designer Award (2007), and the IEEE TRANSACTIONS OF CIRCUITS AND SYSTEMS Darlington Best Paper Award (2010).

Ganesh Balamurugan (M'04) received the B.Tech degree from Indian Institute of Technology, Madras, India, in 1996, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, Urbana, IL, USA, in 2004, both in electrical engineering.

Since 2004, he has been with the Intel's Circuit Research Lab, Hillsboro, OR, USA, working on high-speed wireline signaling. His research interests include low-power equalization and clock recovery in high-speed I/O, reconfigurable transceivers, on-die calibration circuits, and link modeling.

Frank O'Mahony (S'99–M'03) received the B.S., M.S., and Ph.D. degrees from Stanford University, Stanford, CA, USA, all in electrical engineering.

He currently leads the I/O Circuit Technology group within Advanced Design, Intel Labs, Hillsboro, OR, USA, where he is a Principal Engineer. His group develops the first wireline I/O circuits for each new CMOS process technology at Intel. From 2003 until 2011, he was a member of the Signaling Research group in Intel's Circuit Research Lab. His research interests include high-speed and low-power transceivers, clock generation and distribution, equalization, analog circuit scaling, and on-die measurement techniques.

Dr. O'Mahony is a member of the ISSCC Wireline Subcommittee and an Associate Editor for TCAS-I. He is a past recipient of the ISSCC Jack Kilby Award and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS Darlington Best Paper Award.

Clark Roberts received the A.S.E.E.T. degree from Mt. Hood Community College, Gresham, OR, USA, in 1985.

While pursuing his education, he held several technical positions with Tektronix Inc., Beaverton, OR, USA, from 1978 to 1992. In 1992, he joined Intel Corporation's Supercomputer Systems Division, where he worked on the research and design of clock distribution systems for massively parallel supercomputers. He worked on the system clock design for the ASCI Red TeraFlop Supercomputer (Intel, DOE and Sandia). From 1996 to 2000, he was a Signal Integrity Engineer with Intel Corporation's Microprocessor Division, focusing on clock distribution for Intel microprocessor package and motherboard designs. In 2001, he joined the Circuit Research area of Intel Labs, Hillsboro, OR, USA. He has been working on physical prototype design and measurement systems for high-speed IO and terascale microprocessor research.

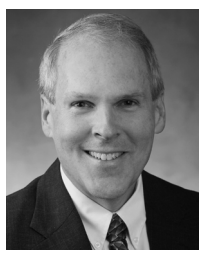

Randy Mooney (M'88) received the M.S.E.E. degree from Brigham Young University, Provo, UT, USA.

He is an Intel Fellow and Director of I/O Research with Intel's Circuit and Systems Lab, part of Intel Labs, Hillsboro, OR, USA. He joined Intel is 1992 and is responsible for circuit and interconnect research enabling bandwidth solutions for our platforms. His work focuses on multi-gigabit, chip-to-chip connections on microprocessor platforms, as well as new memory technologies. He has authored several technical papers and is listed as an inventor or coinventor on more than 70 patents. Prior to joining Intel, he developed products in bipolar, CMOS, and BiCMOS technologies for Signetics Corporation.

Bryan Casper (S'96–M'98) is a Senior Principal Engineering with Intel's Circuit Research Lab, Hillsboro, OR, USA. He leads Intel's high-speed signaling research group responsible for the development of next-generation high-speed mixed signal circuits and I/O systems. In 1998, he joined the Performance Microprocessor Division of Intel Corporation and contributed to the development of the Intel Pentium and Xeon processors. Since 2000, he has been a Circuit Researcher, contributing to the development of I/O self-test technology, signaling analysis methods, high-speed I/O circuit architectures, and multiple I/O standards.