A Scalable 5–15 Gbps, 14–75 mW Low-Power I/O Transceiver in 65 nm CMOS

Ganesh Balamurugan, *Member, IEEE*, Joseph Kennedy, *Member, IEEE*, Gaurab Banerjee, *Senior Member, IEEE*, James E. Jaussi, *Member, IEEE*, Mozhgan Mansuri, Frank O'Mahony, Bryan Casper, and Randy Mooney, *Member, IEEE*

***Abstract*—We present a scalable low-power I/O transceiver in 65 nm CMOS, capable of 5–15 Gbps operation over single-board and backplane FR4 channels with power efficiencies between 2.8–6.5 mW/Gbps. Nonlinear power–performance tradeoff is achieved by the use of scalable transceiver circuit blocks and joint optimization of the supply voltage, bias currents and driver power with data rate. Low-power operation is enabled by passive equalization through inductive link termination, active continuous-time RX equalization, global TX/RX clock distribution with on-die transmission lines, and low-noise offset-calibrated receivers.**

***Index Terms*—Electrical signaling, high-speed I/O, I/O power optimization, inductive termination, low-power equalization, low-power I/O, passive clock distribution, power-efficient links, scalable circuits.**

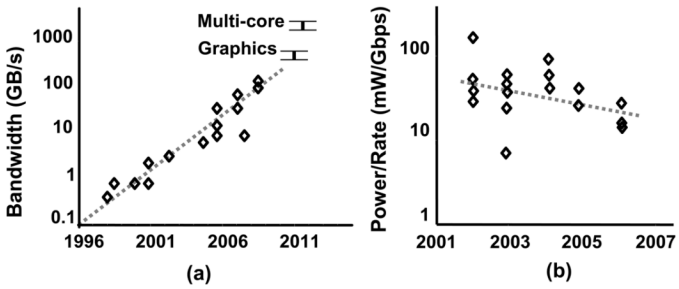

Fig. 1. Trends in (a) microprocessor I/O bandwidth demand, and (b) I/O power efficiency.

I. Introduction

OFF-CHIP bandwidth requirements in high-performance computing systems have increased dramatically in the last decade. Fig. 1(a) shows microprocessor I/O bandwidth demand approximately doubling every two years in recent times. However, the improvement in I/O power efficiency (usually expressed in mW/Gbps) has been relatively gradual [1] [see Fig. 1(b)]. Current 10 Gbps/pin I/O solutions operate with power efficiencies that are worse than 10 mW/Gbps. Future computing platforms that use multiple cores (e.g., [2]–[4]) and high-performance graphics will require several terabits-per-second of off-chip bandwidth to fully realize their on-chip computation capacity. With a power efficiency of 10 mW/Gbps, this entails an unacceptably large I/O power dissipation of 10 W. Thus, significant improvements in the power efficiency of multi-Gbps I/O links are needed to build future tera-scale computing systems.

Most I/O links published to date are fixed-rate designs that have at best a linear relationship between bandwidth and power. However, a *scalable* I/O transceiver, designed to have a *nonlinear* power–performance tradeoff (i.e., a linear decrease in data rate yields better than linear reduction in power) is more suitable in a system that uses intelligent power management, adapting system performance to workload demands [5]. With power playing a significant role in determining system cost, the use of such dynamic power management (DPM) techniques is becoming increasingly common. In order to fully realize the benefits of DPM, I/O circuits need to have favorable power–performance scaling characteristics, similar to core logic and memory circuits.

Several circuit techniques have been proposed to reduce I/O power dissipation. Notable among these are the use of: (a) offset-trimmed receivers to reduce transmit swing in [6]; (b) low swing voltage mode drivers to reduce driver power [1], [7], [8]; (c) inductive clock distribution [1], [9]; and (d) software-based CDR/link calibration [1]. An analysis of the tradeoffs involved in I/O transmitter power optimization was presented in [10]. Scalable links based predominantly on CMOS circuits have been described in [11], where a nonlinear tradeoff between power and performance is achieved by scaling the supply voltage to track data rate. In this paper, we propose techniques that enhance I/O scalability in two ways: (a) accommodate both low-swing current-mode (CML) and rail-to-rail CMOS circuits, and (b) ability to adapt performance to a variety of channels. The former is important as most multi-Gbps links today make extensive use of CML circuits due to their superior supply noise rejection characteristics, while the latter feature allows a single transceiver design to be used in several applications. We present a low-power I/O transceiver in 65 nm CMOS, suitable for parallel links and capable of 5–15 Gbps operation over single-board and backplane FR4 channels with power efficiencies between 2.8–6.5 mW/Gbps[1] [12]. The 3.6 mW/Gbps power efficiency achieved at 10 Gbps is >2x better than the best published fixed-rate implementation [9], while at lower data rates, the power efficiencies are comparable to state-of-the-art [1] (see Fig. 2, [1], [9], [12]–[15]). Transceiver power reduction is enabled by passive equalization

Manuscript received August 24, 2007; revised November 7, 2007.

G. Balamurugan, J. Kennedy, G. Banerjee, J. E. Jaussi, M. Mansuri, F. O'Mahony, B. Casper, and R. Mooney are with the Microprocessor Technology Labs, Intel Corporation, Hillsboro, OR 97124 USA (e-mail: ganesh.balamurugan@intel.com).

Digital Object Identifier 10.1109/JSSC.2008.917522

[1]These numbers reflect the measured power of the local transmitter and receiver and do not include the power due to global clock distribution. If the global clock distribution power is amortized across 20 transceivers, the resulting degradation in power efficiency is 7% at 10 Gbps.

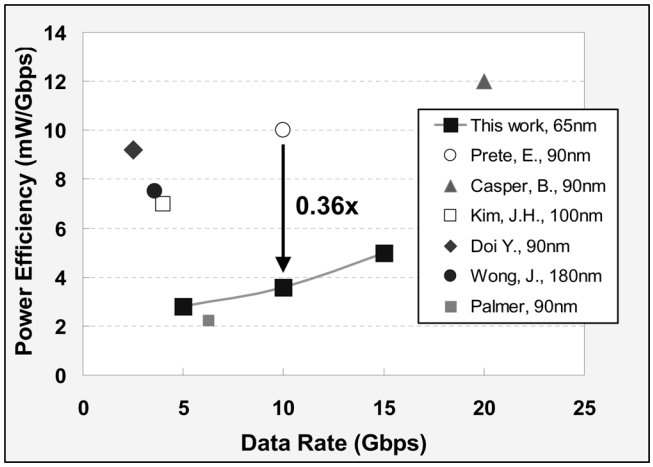

Fig. 2. I/O power efficiency comparison.

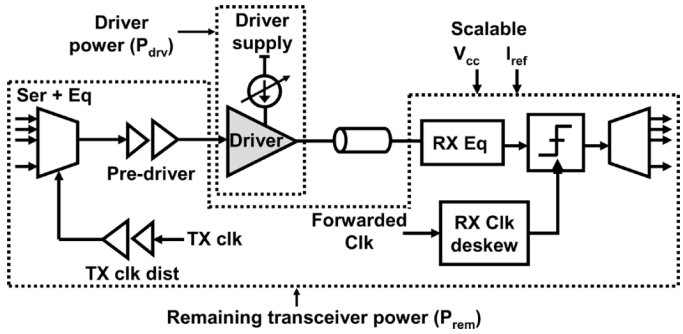

Fig. 3. Transceiver components in a forwarded clock link with dual supply scalable link architecture.

through inductive link termination, active continuous-time RX equalization, global TX/RX clock distribution with on-die transmission lines, and low-noise offset-calibrated receivers. Nonlinear power–performance scaling is obtained by jointly scaling the supply voltage, bias currents, and transmit power with data rate.

In Section II, we describe the scalable link architecture and techniques used to optimize power dissipated by the output driver and receiver clocking circuits - two of the dominant sources of power dissipation in an I/O link. Circuit implementation details are discussed in Section III, measurement results in Section IV, and conclusions in Section V.

II. Architecture

Several circuits contribute to the total power dissipation in a high-speed I/O transceiver. The key components of such a transceiver in a forwarded clock link are shown in Fig. 3. In the transmitter, power is dissipated by: (a) buffers that distribute the transmit clock to the serializer and TX equalizer; (b) TX data path consisting of the equalizer, serializer and pre-driver; and (c) the output driver. In the receiver, the sources of power dissipation are: (a) clock-to-data phase de-skew circuits to optimally sample the receive data; (b) receive-side equalizer to compensate for channel ISI; and (c) front-end circuits that amplify the data to CMOS levels. Below, we discuss architecture features that enable optimization of transceiver power for a range of channels and data rates.

A. Dual Supply Link

It is useful to examine the dependencies of power dissipation of various transceiver components shown in Fig. 3. The power dissipated in the output driver (the "signaling" power) $P_{\rm drv}$, is primarily a function of the channel response, receiver sensitivity and data rate. The remaining transceiver power $P_{\rm rem}$, is largely dependent on the data rate. Since these two transceiver components have different power dependencies, it is desirable to choose an architecture that allows us to independently optimize the two circuits, depending on the channel and data rate. The dual supply architecture shown in Fig. 3 allows us to decouple the optimization of $P_{\rm drv}$ and $P_{\rm rem}$. A dedicated supply voltage powers the output driver (along with global clock generation and biasing circuits), while a separate scalable supply $V_{\rm cc}$ that tracks the data rate, powers the remaining transceiver circuits. This architecture relies on an efficient voltage regulator like that described in [11] to reap the power benefits of voltage scaling. In addition, a scalable current reference biases the current-mode circuits in the transceiver (e.g., clock buffers and RX front-end). The determination of $V_{\rm cc}$ and the power optimization strategy is described in Section IV.

There are two ways to implement the output driver [16]: current-mode (CM) or voltage-mode (VM). While a VM driver can reduce $P_{\rm drv}$ by as much as 4x compared to a CM implementation for the same output swing [1], [8], a CM driver was chosen for the scalable link for two reasons:

a) In order to operate with optimal power efficiency over both the single board and backplane channels (see Section IV), the driver output swing needs to be variable from 100 mVpp to 1 Vpp. While efficient VM driver implementations exist for either low or high output swing [7], [16], it is difficult to design one that delivers the required range of output swing with controlled output impedance. By contrast, a CM driver output swing can be controlled simply by changing its current bias, without adversely affecting the effective termination impedance.

b) Even for a fixed output swing VM driver, the power savings are mitigated by the higher complexity of the pre-driver for impedance control, either by driver segmentation [8] or supply regulation [1], [7]. Adding transmit pre-emphasis capability to a VM driver [7] further diminishes the advantages of a VM implementation over a CM one.

Hence, in a scalable link, a CM implementation has significant advantages over a VM implementation as it allows independent control of output swing, driver impedance, and transmit pre-emphasis with little power overhead.

B. Driver Power Scaling and Optimization

As the driver power can be a significant part of total transceiver power, its scaling behavior affects the power scalability of the overall transceiver. We define the optimal driver power $P_{\rm drv,opt}$ as the minimum driver power required to satisfy certain voltage and timing margins at the receiver. Fig. 4(a) shows the variation of $P_{\rm drv,opt}$ as a function of data rate to achieve voltage and timing margins of 50 mV and 0.3 UI respectively for single

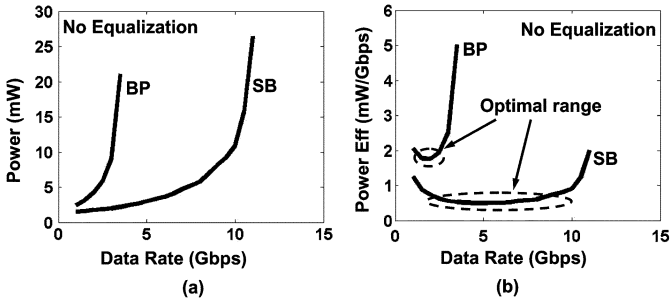

Fig. 4. Scaling behavior of (a) transmit power, and (b) transmit power efficiency.

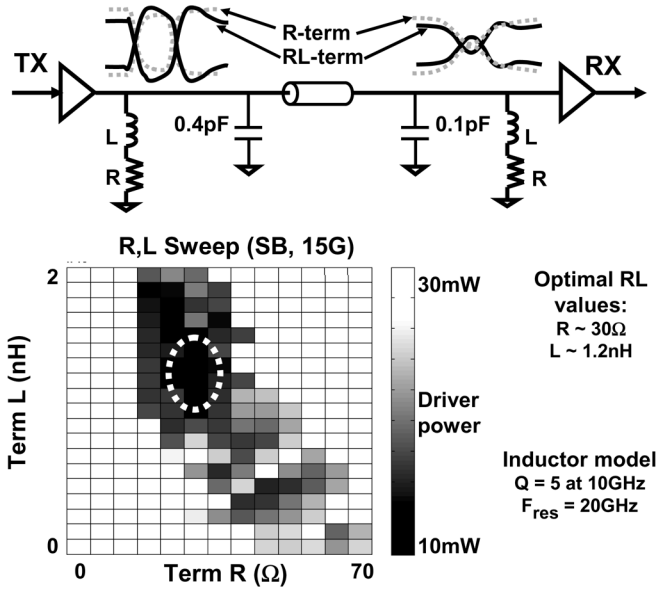

Fig. 5. Optimization of *RL* termination.

board (SB) and backplane (BP) channels. The non-ideal (frequency dependent) nature of the channels results in a strongly nonlinear scaling of $P_{\mathrm{drv,opt}}$ with data rate. A corollary of this nonlinearity is that for channels with inter-symbol interference (ISI), there exists a range of data rates over which it is possible to signal with optimal driver power efficiency [see Fig. 4(b)]. This optimal range is limited by the channel bandwidth, as seen by comparing the SB and BP results in Fig. 4(b). Channel equalization can significantly increase the effective channel bandwidth, and thus increase the optimal signaling range.

A variety of equalization techniques have been demonstrated in high-speed I/O links: transmit pre-emphasis [17], decision feedback equalization (e.g., [18]), continuous-time [15], [19] and discrete-time [20] receive-side linear equalization. In any equalizer implemented using active circuits, the power consumed by the equalizer mitigates the reduction in driver power afforded by equalization. Hence it is desirable to achieve as much equalization as possible using passive elements, subject to implementation constraints. Fig. 5 shows the effect of replacing a conventional resistive termination at the transmitter and receiver with a series combination of a resistor and inductor (*RL*-term). The frequency dependent impedance of an *RL*-term emphasizes data transitions at the transmitter output, which can open the receiver eye when signaling over an ISI

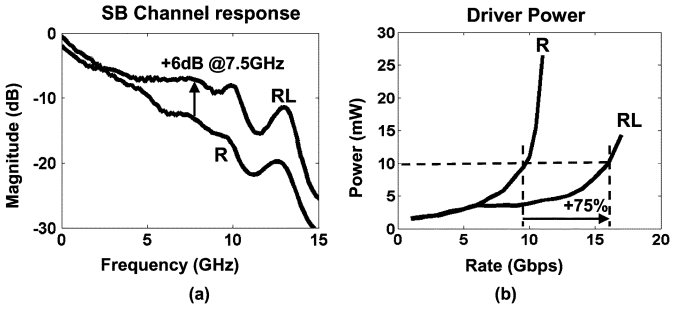

Fig. 6. Effect of *RL* termination on (a) SB channel response, and (b) TX driver power.

channel. This can be used to reduce driver power and extend signaling rates. However, to reap the benefits of an *RL*-term, it is essential to choose appropriate values for the resistance and inductance. Fig. 5 shows the dependence of driver power on the *R* and *L* values in an *RL*-term for 15 Gbps signaling over a SB channel. The optimal *R* value in the *RL*-term ($\sim 30\ \Omega$) is significantly smaller than 50 Ω, while the optimal inductance is approximately 1 nH and hence can be realized on-die. Note that the low-Q inductor contributes significant series resistance, resulting in an effective termination resistance larger than 30 Ω. An optimized *RL*-term can significantly enhance high-frequency transfer characteristics of the SB channel resulting in a flatter response as shown in Fig. 6(a). This can be exploited to reduce driver power, particularly at data rates higher than 7 Gbps [Fig. 6(b)]. For example, at 10 Gbps, an *RL*-term can reduce driver power by $>$2x with no additional power cost. Another benefit of an *RL*-term for forwarded clock systems is indicated by the frequency responses in Fig. 6(a). The better high-frequency transfer characteristics afforded by *RL*-term (as much as 6 dB improvement at 7.5 GHz) results in a smaller attenuation of the high-frequency forwarded clock. This reduces the number of clock pre-amplifier stages required in the receiver, resulting in lower jitter and power. An *RL*-term by itself is insufficient to equalize the lossier BP channel and support data rates up to 15 Gbps. Hence, it is augmented by a simple continuous-time linear equalizer [15] and a low-power 3-tap TX equalizer (described in Section III).

C. RX Clock De-Skew Power Optimization

In addition to the output driver, a significant source of power dissipation in forwarded clock links is RX clock de-skew circuitry [11]. Per-pin de-skew is commonly used in high-speed parallel links to compensate for clock-data channel mismatches. Clock de-skew is usually accomplished by using a combination of a voltage-controlled delay line (VCDL) and a phase interpolator. The VCDL is typically in a delay-locked loop (DLL) and generates multiple clock phases that are then selectively mixed by an interpolator to select the optimal sampling phase. A local DLL at each RX is preferable to a global DLL [11], as it avoids the power/jitter cost of distributing multiple clock phases to all RXs. The number of stages used to implement the VCDL (N_{stg}) has a strong impact on the clock de-skew power. As N_{stg} increases, the bandwidth per delay cell increases to maintain the total VCDL delay. While this improves the overall

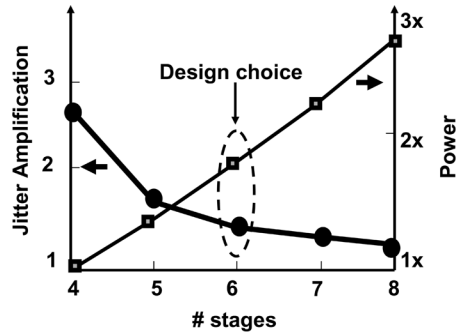

Fig. 7. Jitter amplification and RX clock power versus number of DLL delay stages.

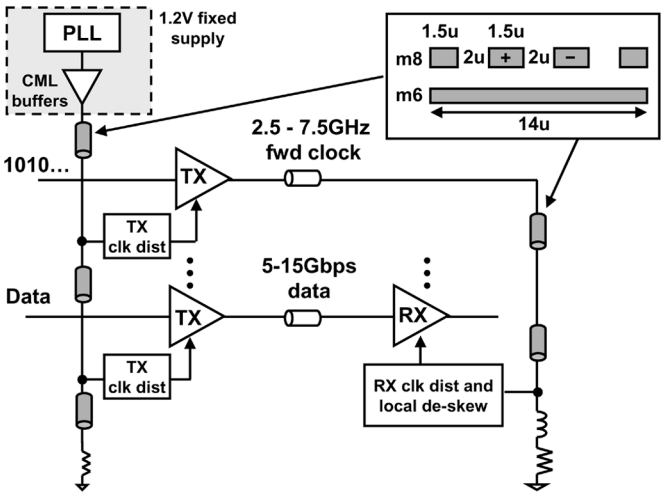

Fig. 8. Link architecture with global clock distribution and T-line cross section.

jitter transfer characteristics of the clock path, it comes at the expense of power. It is important to note that this tradeoff exists at *any* operating frequency (within the DLL lock range). This inter-relationship between $N_{\rm stg}$, power and jitter transfer for a VCDL based on symmetric load delay cells is shown in Fig. 7. Jitter transfer is quantified by jitter amplification—the ratio of output to input (rms) jitter for a white Gaussian input jitter. While there is a >2x reduction in jitter amplification when $N_{\rm stg}$ is increased from 4 to 6, there is only a marginal improvement beyond six stages. Hence, a six-stage DLL was chosen to provide an optimal compromise between power and jitter for the scalable transceiver.

III. Implementation

A top-level block diagram of the scalable parallel link is shown in Fig. 8. A fixed 1.2 V global supply powers the output drivers, global PLL and bias generation circuits. The remaining circuits use a scalable supply that tracks the data rate. The scalable supply in the prototype link was set by an external voltage source—a full-blown implementation would use an efficient voltage regulator such as the one described in [11]. A half-rate clock is distributed to all the transmitters and is also forwarded to all the receivers using a transmitter configured to send alternating data.

A. Global Clocking

Clocks are globally distributed using on-chip differential transmission lines constructed from top-level metal layers (Fig. 8). The global TX clock distribution network minimizes jitter due to supply noise by using CML buffers with low delay sensitivity while maintaining a high bandwidth and low total buffer latency. This passive distribution provides a more favorable power–jitter tradeoff as compared to more active distribution techniques [21]. It is also broadband in nature (in contrast to other low-power distribution techniques like tuned *LC* lines [1], [9]), making it suitable for a scalable link where a wide range of clock frequencies (2.5–7.5 GHz) needs to be supported. While the global transmission line in the TX is resistively terminated, at the RX side, we use an *RL*-term to resonate out part of the pad capacitance and increases RX clock amplitude. The maximum length of on-die transmission lines in the test-chip is $\sim$2 mm. However, since the estimated loss of these lines is $\sim$1 dB/mm, it is possible to scale this distribution to at least 6 mm. Several receivers and transmitters tap the low swing global distribution to generate local clocks, with duty-cycle correction and level conversion performed at each tap. Each receiver has local de-skew circuits to optimally sample the incoming data.

B. Scalable Circuits

The scalable transceiver is implemented using a combination of low-swing CML and rail-to-rail CMOS circuits. In order to realize the desired nonlinear power–performance tradeoff, both these circuits and their interfaces need to scale appropriately with data rate. For CMOS circuits, adjusting the supply voltage to track data rate yields the required nonlinear power–performance scaling [11]. Extending this scaling to CML circuits requires the scaling of their bias currents as well. A replica-biased symmetric load amplifier [22] [Fig. 9(a)] was chosen to implement all CML circuits in the transceiver since it exhibits the desired gain/bandwidth scaling. For such an amplifier, while its bandwidth varies approximately linearly with bias current, the gain remains fairly constant. This is because the reduced transconductance at lower bias currents is compensated for by the increased resistance of the active load. Where CML-to-CMOS conversion is required, as in the case of low-swing TX/RX clocks, a simple level converter based on a differential to single-ended converter and a pseudo-differential amplifier is used [Fig. 9(b)]. The duty-cycle penalty incurred by this conversion is compensated using local duty-cycle correction (DCC) circuits (described in Section III-C).

C. Transmitter

The transmitter (Fig. 10) employs a 2:1 input-multiplexed architecture [6] with a current-mode output driver. The CML output stage is powered by the fixed supply while the rest of the transmitter uses a scalable supply that tracks the data rate. As mentioned in Section II, this was done to facilitate independent optimization of the transmit power and the remaining transceiver power. The driver current is adjustable over a wide range (0–24 mA) in order to accommodate lossy backplane channels that need TX output swings of $\sim$1 Vpp. The transmitter is terminated by a series combination of a resistor (30 Ω nominal)

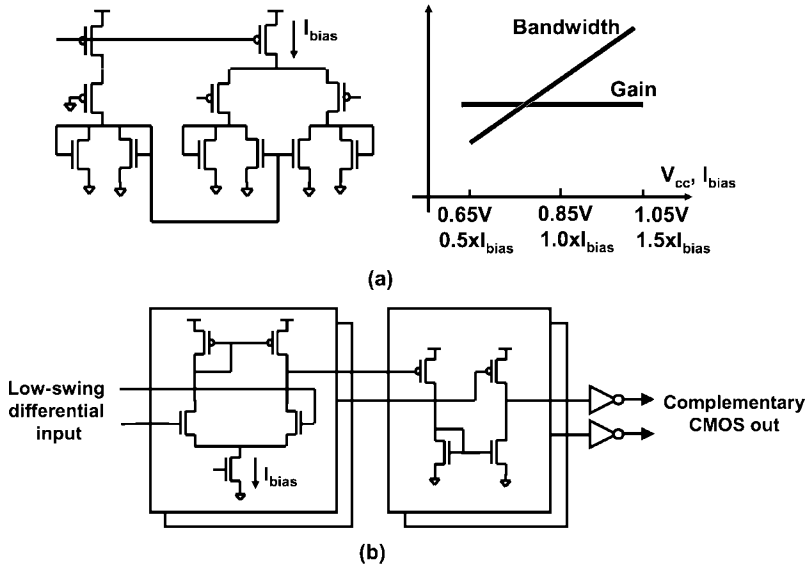

Fig. 9. (a) Symmetric load CML amplifier and scaling behavior. (b) CML-to-CMOS level converter.

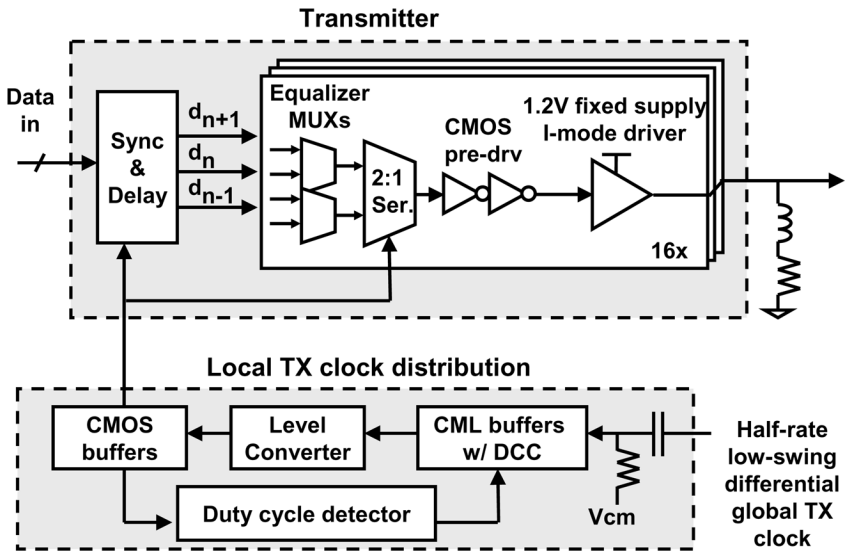

Fig. 10. Transmitter and local TX clock distribution.

and an inductor (1.1 nH) to deliver passive equalization at no additional power cost. The inductor in the *RL*-term is implemented differentially in top-level metal in a 65 μm $\times$ 65 μm area. It is designed to achieve Q $\sim$ 5 at 7.5 GHz. The TX output stage is segmented into 16 legs to implement a 3-tap TX pre-emphasis filter with 4-bit precision. This filter provides $\sim$ 12 dB of equalization in addition to *RL*-term and RX equalizer (see Section III-D). This is required to enable 15 Gbps I/O over the BP channel. Each output leg is composed of a CML driver, a CMOS pre-driver, and a 2:1 serializer preceded by multiplexers to implement TX pre-emphasis. A synchronizer and delay chain generate the 3 data-streams necessary to realize a 3-tap linear equalizer.

The globally distributed differential low-swing TX clock is tapped locally at each transmitter via AC coupling capacitors (Fig. 10). This enables a smooth transition between the fixed and scalable supply domains with optimal input common-mode for the CML buffers in the local TX clock distribution. Device mismatches in the CML/CMOS clock distribution circuits and CML-to-CMOS level conversion can lead to clock duty-cycle error. Since duty-cycle error is a manifestation of high-frequency jitter, it can have a significant detrimental impact on link performance, if left uncorrected [23]. Hence, each TX has local duty-cycle correction (DCC) and detection (DCD) circuits that are designed to achieve less than 2% residual duty-cycle error. The method of DCC and DCD is similar to that described in [24]. DCC is accomplished by using digitally controlled variable offset amplifiers, which introduce a separation between the DC values of a low-swing differential signal pair. DCD circuits detect the difference between the average values of complementary CMOS clock signals to drive the DCC loop. Following DCC, a combination of CML and CMOS buffers amplify the low-swing clock to CMOS levels to clock the 2:1 serializer and equalizer.

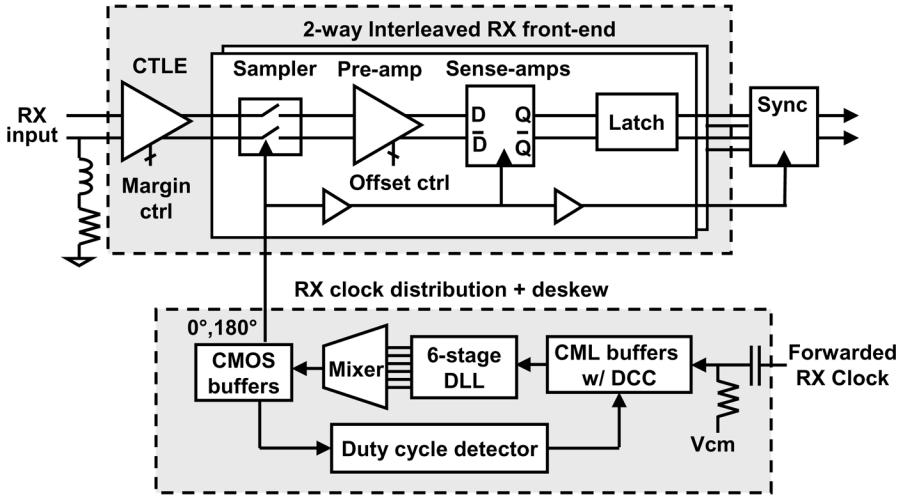

Fig. 11. Receiver and RX clock distribution.

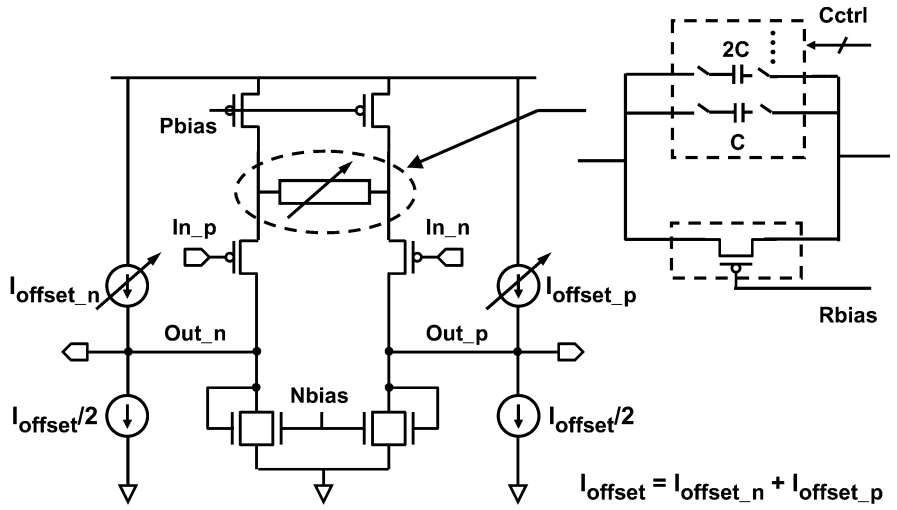

Fig. 12. CTLE with wide range offset control for link margining.

D. Receiver

The receiver (Fig. 11) consists of a continuous-time linear equalizer (CTLE), a 2-way interleaved receive front-end (RXFE) and local de-skew circuits to optimally sample the receive data—all powered by the scalable supply. The 2-way interleaved architecture eliminates the need to generate multi-phase (>2) sampling clocks. Similar to the TX output, the RX input is terminated using an *RL*-term. The CTLE (Fig. 12), designed to provide a high-frequency boost of up to 12 dB relative to DC, is realized by source degeneration of a symmetric load amplifier using a parallel *R-C* circuit [15]. Tunable *R* and *C* values in the degeneration network allow control of the magnitude of peaking and the zero frequency. Link voltage margining is enabled by injecting a positive or negative differential current at the CTLE output. Usually, this entails a tradeoff between achievable offset range and amplifier performance. However, using source *and* sink offset currents avoids sacrificing headroom for margining range. In order to realize output offset currents up to $I_{\rm offset}$, tunable offset currents totaling $I_{\rm offset}$ are sourced into the amplifier outputs as shown in Fig. 12. In addition, static sink offset currents equal to $I_{\rm offset}/2$ are also injected to zero offset currents for the nominal setting. This technique enables the use of large offset currents to realize up to $\pm$200 mV differential offset (input-referred) without perturbing the nominal operating point of the amplifier and unduly compromising gain. During normal operation, the margining circuits are turned off and the CTLE consumes 7% of the total RX power. Turning on the margining circuits during link calibration doubles the CTLE power.

Following the CTLE, the data is sampled by nMOS switches and fed to the RXFE. Each leg of the RXFE consists of a symmetric load based pre-amplifier followed by regenerative stages. The offsets in the pre-amplifier and regenerative stages are compensated by adding a differential current at the pre-amplifier output. Since the magnitude of the differential offset current is $<$10% of the pre-amp bias current, offset compensation is achieved with little impact on power or performance. The regenerative stage following the pre-amplifier consists of two clocked sense-amplifiers followed by two unclocked latches [25]. The RXFE consumes 11% of the overall RX power during normal link operation.

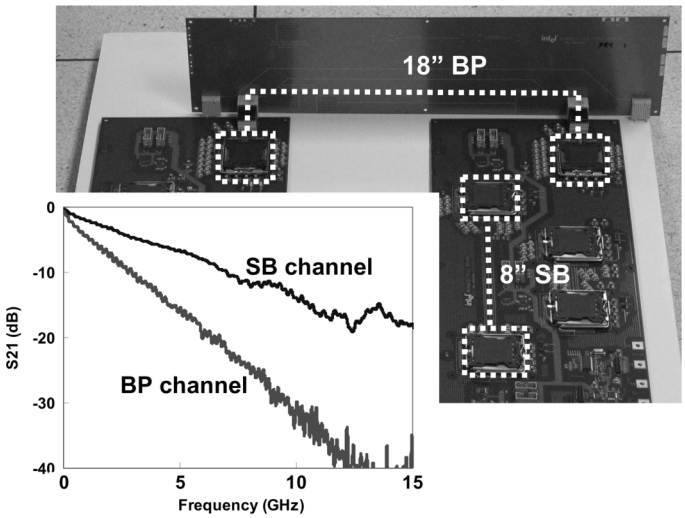

Fig. 13. Single board and backplane channels with measured responses (excludes pad capacitance).

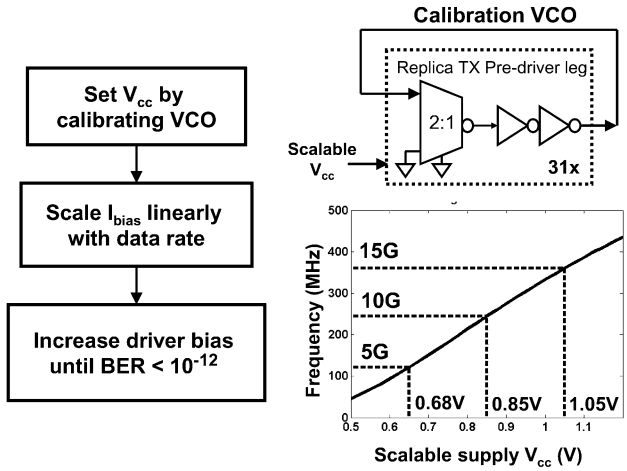

Fig. 14. Transceiver power optimization.

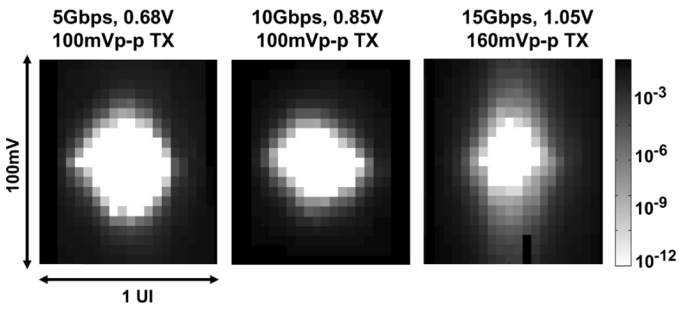

Fig. 15. Eye diagrams after link power optimization.

The global forwarded clock is tapped at each RX via an AC-coupler (Fig. 11) to level shift the input clock and allow the use of nMOS pre-amplifying stages for lower power. The AC coupler also rejects sub-GHz common-mode noise and reduces clock duty-cycle error. A 6-stage pre-amplifier (constructed from replicas of the DLL delay cell) ensures adequate signal swing at the input to the voltage-controlled delay line (VCDL). The first pre-amplifier stage also includes DCC circuits, similar to those in the TX. The DLL uses a 6-stage VCDL to provide an optimal tradeoff between power and jitter transfer for reasons discussed in Section II. A low-swing Gilbert-cell phase detector based on the DLL delay cell eliminates the need for level converters within the DLL. The DLL bandwidth is 400 MHz with a 7.5 GHz input clock. Taps from the VCDL feed a CML MUX based interpolator to provide 5° sampling phase resolution.

IV. Measurement Results

The transceiver test-chip is implemented in a 1.2 V, 65 nm CMOS technology. The small active TX (RX) area of 0.033 mm^2 (0.055 mm^2) makes this transceiver suitable for highly parallel links. The scalable transceiver was characterized over two channels shown in Fig. 13: (a) 8" FR4 microstrip with socket-Ts and packages at TX/RX (SB channel), and (b) 18" FR4 stripline over a backplane with two connectors (BP channel). For each channel, we optimized the transceiver power using the three-step process described in Fig. 14. First, the scalable supply voltage is set using the frequency-to-voltage transfer characteristics of a calibration VCO (CVCO). This VCO is a 31-stage ring oscillator in which each stage is constructed from a replica of the TX output stage (2:1 serializer + CMOS pre-driver). The TX output stage was chosen to construct the CVCO since it is one of the most bandwidth-critical stages in the link that needs to support the full rate data. The scalable supply voltage is chosen to achieve a delay of 0.7UI per stage of the CVCO. Circuit simulations indicated that meeting this delay criterion ensured adequate supply noise sensitivity and signal edge rates. For example at 10 Gbps, 0.7UI/stage implies a CVCO frequency of 230 MHz and hence a scalable supply voltage of 0.85 V (see Fig. 14). After setting the scalable supply voltage, the bias current reference for all current-mode circuits is scaled linearly with data rate to reflect the linear dependence of the bandwidth of a symmetric-load amplifier on its bias current. A reference value for the CML bias current scaling is obtained at a single data rate by bypassing the DLL loop and determining the delay cell bias current necessary to achieve the required bandwidth. With the transmitter configured to send PRBS data, receiver clock de-skew is performed by sweeping the sampling clock phase over 1UI and determining the center of the eye using on-die error counters and off-chip control software. Finally, the bias current (and hence output swing) of the current mode driver is increased until the target BER of 10^{-12} is met. Fig. 15 shows the receiver eye diagrams after power optimization at 5/10/15 Gbps for the SB channel, measured using on-die link margining circuits [15]. The scalable supply voltage varies from 0.68–1.05 V for 5–15 Gbps operation, while the transmit swing is in the range 100 mV–160mVpp. For the SB channel, the *RL*-term and CTLE are sufficient to equalize the channel and hence transmit pre-emphasis was disabled in this configuration. The eye diagrams in Fig. 15 indicate that the optimized transmit swing ensures a ±20 mV voltage margin at all three data rates. While there is a slight degradation in timing margin when the data rate is increased from 10 to 15 Gbps, the eyes at the three data rates are fairly similar, indicating that we have traded off excess performance margins to optimize transceiver power.

TABLE I
MEASUREMENTS SUMMARY

Data rate	15G	10G	5G
Scalable supply	1.05V	0.85V	0.68V
Bias current scaling factor	1x	0.67x	0.33x
Rx input-referred noise (rms)	<1mV	<1mV	<1mV
Link timing uncertainty (rms)	1ps / 0.015UI	1.2ps / 0.012UI	2.6ps / 0.013UI
TX output swing (differential)	160mVpp	100mVpp	100mVpp
TX Power	34mW	17mW	7.5mW
RX Power	41mW	19mW	6mW
Transceiver power	75mW	36mW	13.5mW
Power Eff (mW/Gbps)	5	3.6	2.7

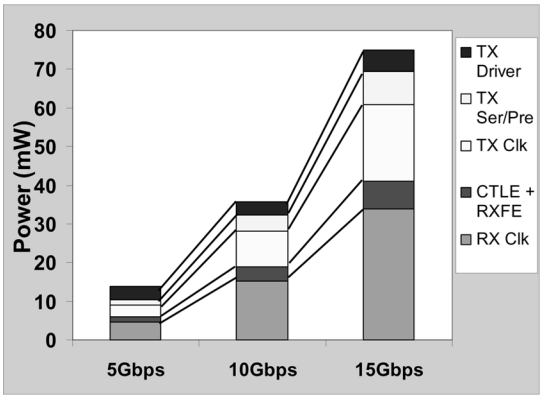

Fig. 16. Measured transceiver power breakdown.

The characteristics of the scalable transceiver are summarized in Table I. The receiver input-referred noise arising from thermal noise in receiver front-end devices and termination remains below 1 mV rms even as we scale the supply voltage and bias currents with data rate. Link timing uncertainty [15], which is a measure of the differential jitter between clock and data in a forwarded clock link, is extracted by sampling the edges in an alternating data pattern. As seen from Table I, this timing uncertainty remains approximately 0.015UI from 5–15 Gbps. After power optimization, the total transceiver power is almost evenly divided between the transmitter and receiver. The power efficiency (assuming ideal voltage regulation) improves from 5.0 mW/Gbps at 15 Gbps to 2.7 mW/Gbps at 5 Gbps, indicating a nonlinear relationship between power and data rate—a 3x reduction in data rate results in ∼5x decrease in power. Fig. 16 shows the measured power breakdown among various transceiver components at 5/10/15 Gbps over the SB channel after power optimization. It is clear that when signaling with optimal driver power, clocking circuits dominate the overall transceiver power, accounting for 56%–71% of total power for 5–15 Gbps operation.

The effect of the channel on transceiver power is shown in Fig. 17. Since the BP channel has significantly more ISI than the SB channel (Fig. 13), the 3-tap transmit pre-emphasis filter is enabled in addition to *RL*-term and CTLE. The transmit power also needs to be increased to compensate for greater channel

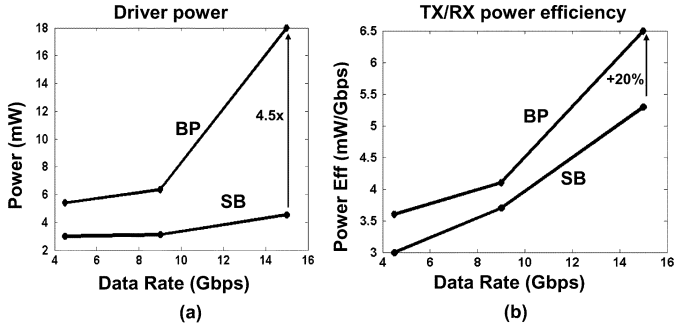

Fig. 17. Effect of channel on (a) driver power, and (b) transceiver power efficiency.

loss in the BP case. Driver power and transmit pre-emphasis were jointly optimized by determining the minimum driver power required for each pre-emphasis setting. As shown in Fig. 17, choosing the tap weights in this manner results in 4.5x increase in TX swing (to ∼720 mVpp) and 20% degradation in overall transceiver power efficiency compared to the SB channel at 15 Gbps.

V. CONCLUSION

A 5–15 Gbps scalable low-power I/O transceiver suitable for parallel links has been implemented in 65 nm CMOS. The transceiver was characterized over two FR4 channels—single-board and backplane—with very different frequency responses. The measured power efficiencies range from 2.7–6.5 mW/Gbps depending on the channel and data rate. Nonlinear power–performance tradeoff is realized by using scalable circuits and a dual supply link architecture. Several techniques were used to optimize transceiver power: joint systematic optimization of supply voltage, bias currents and transmit power, low-power equalization techniques like inductive termination and CTLE, low-power/jitter global clock distribution using on-die transmission lines, and sensitive offset-calibrated receivers that enable low transmit swings.

ACKNOWLEDGMENT

The authors thank H. Wilson, J. Howard, D. Finan, D. Jenkins, L. Peake, V. Baca, K. Kumar, and H. Nguyen for prototype assistance, J. Collias, C. Roberts and D. Casper for interconnect and test support, and C. Sweeney and Intel Fab 24 for silicon fabrication.

REFERENCES

[1] R. Palmer *et al.*, "A 14 mW 6.25 Gb/s transceiver in 90 nm CMOS for serial chip-to-chip communications," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 440–614.
[2] U. M. Nawathe *et al.*, "An 8-core 64-thread 64b power-efficient SPARC SoC," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 108–590.
[3] S. Vangal *et al.*, "An 80-tile 1.28TFLOPS network-on-chip in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 98–589.
[4] J. Dorsey *et al.*, "An integrated quad-core Opteron processor," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 102–103.
[5] Y. Yoshida *et al.*, "A 4320MIPS four-processor core SMP/AMP with individually managed clock frequency for low power consumption," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 100–590.

[6] M.-J.E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1591–1599, Nov. 2000.
[7] K.-L. J. Wong, H. Hatamkhani, M. Mansuri, and C.-K. K. Yang, "A 27-mW 3.6-Gb/s I/O transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 602–612, Apr. 2004.
[8] C. Menolfi, T. Toifl, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. Schmatz, "A 16 Gb/s source-series terminated transmitter in 65 nm CMOS SOI," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2007, pp. 446–614.
[9] E. Prete *et al.*, "A 100 mW 9.6 Gb/s transceiver in 90 nm CMOS for next-generation memory interfaces," in *IEEE Int. Solid-State Circuits Dig. Tech. Papers*, 2006, pp. 253–262.
[10] H. Hatamkhani, F. Lambrecht, V. Stojanovic, and C.-K. Ken Yang, "Power-centric design of high-speed I/Os," in *Proc. 43rd ACM/IEEE Design Automation Conf.*, Jul. 24–28, 2006, pp. 867–872.
[11] G. Wei *et al.*, "A variable-frequency parallel I/O interface with adaptive power-supply regulation," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1600–1610, Nov. 2000.
[12] G. Balamurugan, J. Kennedy, G. Banerjee, J. Jaussi, M. Mansuri, F. O'Mahony, B. Casper, and R. Mooney, "A scalable 5–15 Gbps, 14–75 mW low power I/O transceiver in 65 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 270–271.
[13] J.-H. Kim *et al.*, "A 4-Gb/s/pin low-power memory I/O interface using 4-level simultaneous bi-directional signaling," *IEEE J. Solid-State Circuits*, vol. 35, no. 1, pp. 89–101, Jan. 2005.
[14] Y. Doi *et al.*, "A 0.8–1.3 V 16-channel 2.5 Gb/s high-speed serial transceiver in a 90 nm standard CMOS process," in *Proc. IEEE Custom Integrated Circuits Conf.*, 2005, pp. 131–134.
[15] B. Casper, J. 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 B," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2006, pp. 263–272.
[16] W. J. Dally and J. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
[17] W. J. Dally and J. Poulton, "Transmitter equalization for 4 Gb/s signaling," in *Proc. Hot Interconnect*, Aug. 1996, pp. 29–39.
[18] V. Stojanovic *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.
[19] S. Gondi *et al.*, "A 10 Gb/s CMOS adaptive equalizer for backplane applications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2005, vol. 1, pp. 328–601.
[20] J. E. Jaussi, G. Balamurugan, D. R. Johnson, B. Casper, A. Martin, J. Kennedy, N. Shanbhag, and R. Mooney, "8-Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation, and clock de-skew," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 80–88, Jan. 2005.
[21] F. O'Mahony, M. Mansuri, B. Casper, J. E. Jaussi, and R. Mooney, "A low-jitter PLL and repeaterless clock distribution network for a 20 Gb/s link," in *2006 Symp. VLSI Circuits Dig. Tech. Papers*, 2006, p. 29.
[22] J. G. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1723–1732, Nov. 1996.
[23] G. Balamurugan and N. Shanbhag, "Modeling and mitigation of jitter in multi-Gbps source-synchronous I/O links," in *Proc. 21st Int. Conf. Computer Design*, Oct. 13–15, 2003, pp. 254–260.
[24] T. Lee *et al.*, "A 2.5 V CMOS delay-locked loop for an 18 Mbit, 500 MB/s RAM," *IEEE J. Solid-State Circuits*, vol. 29, no. 12, pp. 1491–1496, Dec. 1994.
[25] C.-K. K. Yang *et al.*, "A serial-link transceiver based on 8-GSamples/s A/D and D/A converters in 0.25-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1684–1692, Nov. 2001.

Ganesh Balamurugan (S'99–M'04) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 1996, and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 2004.

Since 2004, he has been with Intel's Circuit Research Lab in Hillsboro, OR. His research interests include low-power equalization and clock recovery in high-speed links, I/O link modeling and noise-tolerant digital system design.

Joseph Kennedy (S'88-M'91) received the B.S. degree in electrical and computer engineering from Oregon State University, Corvallis, in 1991.

He is a Senior Circuits Researcher with Intel's Circuits Research Labs in Hillsboro, OR. 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. His recent focus has been on low-power scalable-rate CMOS I/O transceivers. 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.

Gaurab Banerjee (S'98–M'00–SM'05) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Indian Institute of Technology, Kharagpur, Auburn University, Auburn, AL, and the University of Washington, Seattle, in 1997, 1999, and 2006, respectively.

In 1999, he joined Intel Corporation in Hillsboro, OR, to design analog and mixed-signal circuits for the first Pentium4 microprocessor. In 2001, he moved to Intel research and development to work on CMOS based analog, mixed-signal and RF circuits for wireless and wire-line communication. Since April 2007, he is with the Analog/RF design group of Qualcomm Inc. in Austin, TX, where he is currently a Staff Engineer, working on RFIC design for mobile broadcast video applications. His research/design interests are in analog, mixed-signal and RF circuits, data converters and phase locked loops. He has published more than 15 papers on semiconductor devices and circuits and has received three patents (with several more pending). He has also served as an industrial liaison for several university research projects, sponsored by Intel and the S.R.C.

Dr. Banerjee is a National Talent Search Scholar of India and a Senior Member of IEEE. He is currently an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART I.

James E. Jaussi (M'01) received the B.S. and M.S. degrees in electrical engineering from Brigham Young University, Provo, UT, in 2000. He is currently working towards the Ph.D. degree in electrical engineering at Oregon State University, Corvallis.

For the past seven years, he has worked for Intel Laboratories, Hillsboro, OR. His main focus is research, design and characterization of high-speed CMOS transceivers and mixed-signal circuits, with an emphasis in receiver equalization and clock recovery architectures.

Mozhgan Mansuri was born in Tehran, Iran. She 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 the University of California, Los Angeles, 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. In 2003, she joined Intel Corporation, Hillsboro, OR. Her research interests include low-power low-jitter clock synthesis/recovery circuits (PLL and DLL) and low-power high-speed I/O links.

Frank O'Mahony received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1997, 2000, and 2004, respectively. His doctoral research focused on resonant clock distribution techniques for high-performance microprocessors.

Since 2003, he has been with Intel's Circuit Research Lab in Hillsboro, OR. His research interests include high-speed and low-power data links, clock generation and distribution, and design techniques for low-noise, variation-tolerant clocking and signaling circuits.

Dr. O'Mahony received the 2003 Jack Kilby Award for Outstanding Student Paper at ISSCC.

Bryan Casper received the M.S. degree in electrical engineering from Brigham Young University, Provo, UT.

He is currently leading the high-speed signaling team of Intel's Circuit Research Lab, based in Hillsboro, OR. In 1998, he joined the Performance Microprocessor Division of Intel Corporation and worked on the development of Pentium and Xeon processors. Since 2000, he has been a circuit researcher responsible for the research, design, validation and characterization of high-speed mixed-signal circuits and I/O systems.

Randy Mooney (M'88) received the M.S. degree in electrical engineering from Brigham Young University, Provo, UT.

He currently is an Intel Fellow and director of I/O research in Intel's Corporate Technology Group. He was with the Standard Products Division of Signetics Corporation, Orem, UT, from 1980 to 1992, where he developed products in Bipolar, CMOS and BiCMOS technologies, concentrating on components for use in bus driving applications. In 1992, he joined the Supercomputer Systems Division of Intel Corporation, and worked on development of interconnect components for parallel processor communications. He was responsible for the development of signaling technology for these components, and developed a method of simultaneous bidirectional signaling that was used for the Intel Teraflops Supercomputer. His current work is focused on multi-gigabit, differential, serial stream based interfaces and pt-to-pt memory interfaces.