A 10-Gb/s/ch, 0.6-pJ/bit/mm Power Scalable Rapid-ON/OFF Transceiver for On-Chip Energy Proportional Interconnects

Da Wei, *Student Member, IEEE*, Tejasvi Anand, *Member, IEEE*, Guanghua Shu, *Member, IEEE*, José E. Schutt-Ainé, *Fellow, IEEE*, and Pavan Kumar Hanumolu, *Member, IEEE*

***Abstract*— Modern multiprocessor system-on-chips employ network-on-chip (NoC) to efficiently connect different components together. NoCs need global and local interconnects to deliver high on-chip bandwidth and low communication latency to avoid being a performance bottleneck. They must also have high throughput density to reduce area occupied by wires. This paper presents techniques to implement power efficient transceivers for on-chip links that can achieve energy proportional operation. Conventional on-chip links optimized for best energy efficiency at peak data rate suffer from degraded energy efficiency under low utilization conditions. Dynamic voltage and frequency scaling and clock gating can partially alleviate this problem, but become ineffective in applications like mobile devices, where the data traffic can be very sporadic. In this paper, architecture and circuit techniques to improve energy efficiency under all utilization levels are presented. The proposed transceiver uses single-ended signaling with only 0.5 μm width and spacing and achieves 5-Gb/μm throughput density. Fast locking signaling and clocking circuits greatly reduce the power-ON time. Fabricated in 65-nm CMOS technology, the proposed 10-Gb/s transceiver achieves wake-up time in less than 17 ns. More than 125× effective data rate scaling (10 Gb/s to 80 Mb/s) is obtained with an energy efficiency degradation of only 1.6× (627 to 997 fJ/b/mm). When the supply voltage is scaled from 1 to 0.7 V, the peak data rate scales from 10 to 6 Gb/s and the power scalable range increases to 208× (10 Gb/s to 48 Mb/s) with the energy efficiency degradation of only 1.2× (627 to 753 fJ/b/mm).**

***Index Terms*— Burst mode, energy efficient, energy proportional, low power, multiplying delay-locked loop (MDLL), on-chip link, power scalable, rapid ON/OFF, serial link, transceiver.**

I. Introduction

NETWORK-ON-CHIP (NoC) is a vital part of the state-of-the-art multiprocessor system-on-chips. NoCs allow fast run-time reconfiguration and help meet design targets with minimum hardware resources [1]. NoCs must satisfy several requirements such as: 1) high throughput density (higher Gb/s/μm); 2) low power even while communicating across long distances (lower fJ/bit/mm); and 3) maintain high power efficiency under all workload scenarios. However, fulfilling these requirements is challenging. For instance, driving global NoC interconnects that are often very long (>5 mm) at high data rate results in excessive power consumption.

Low-swing signaling techniques [2], power efficient voltage-mode drivers [3], high sensitivity receivers that employ power efficient equalizers [4], and offset-canceled sense amplifiers [5] are used to lower power consumption. However, these techniques do not consider bursty and application-dependent nature of data traffic, and as a result, they are ineffective in scaling power in proportion to link utilization [6]. In other words, aforementioned low power techniques only help reduce active power and are ineffective for reducing idle power. In fact, on-chip links consume nearly the same amount of power regardless of whether the link is active or idle. Consequently, link energy efficiency severely degrades at lower utilization levels [7]. This is particularly the case in mobile applications, where the data traffic can be very sporadic (short active periods interspersed between long inactive periods). Interconnect power can be partially scaled by clock gating [8]; however, due to always-on power consumption of clock generating circuits, such solution suffers from low energy efficiency at low link utilization. Prior research [5], [8], [9] also suffers from low throughput density due to large spacing between lanes and wide wire width. Techniques, such as dynamic voltage and frequency scaling (DVFS), could potentially scale power in proportion to link utilization [10]. However, the time required to scale the supply voltage can be much longer than the data burst interval, which makes DVFS alone ineffective in such use cases.

It has been demonstrated that rapid ON/OFF is a promising way to reduce power consumption of serial links by almost 70% [11]–[16]. However, these examples are limited to off-chip interconnects, suffer from long power-ON time and low energy efficiency in always-on condition, and operate at low data rates. A high-speed power scalable rapid-ON/OFF on-chip interconnect does not exist.

In this paper, we present a complete source synchronous rapid-ON/OFF transceiver for on-chip interconnects capable of scaling its power down to near zero in accordance with utilization. Proposed 10-Gb/s prototype transceiver features a capacitive driver-based 3-tap FFE implemented in a 65-nm

Manuscript received April 19, 2017; revised September 20, 2017; accepted November 29, 2017. Date of publication December 28, 2017; date of current version February 21, 2018. This paper was approved by Guest Editor Dejan Markovic. This work was supported in part by Analog Devices and Silicon Labs. *(Corresponding author: Da Wei.)*

D. Wei, J. E. Schutt-Ainé, and P. K. Hanumolu are with the Department of Electrical and Computer Engineering, University of Illinois at Urbana–Champaign, Urbana, IL 61801 USA (e-mail: dawei1@illinois.edu).

T. Anand is with the Electrical Engineering and Computer Science Department, Oregon State University, Corvallis, OR 97331 USA (e-mail: anandt@eecs.oregonstate.edu).

G. Shu is with Oracle Labs, Belmont, CA 94002 USA (e-mail: guanghuashu@gmail.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.2017.2782082

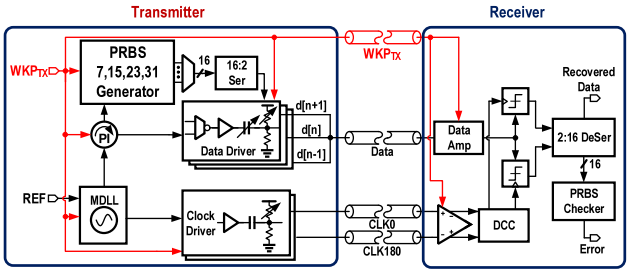

Fig. 1. Block diagram of the proposed source synchronous link.

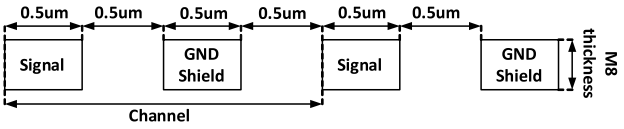

Fig. 2. Channel cross section.

CMOS process, and achieves 5-Gb/s/μm throughput density along with 627-fJ/b/mm energy efficiency (471 fJ/b/mm if the clocking power is amortized among nine data lanes). More importantly, its energy efficiency degrades to only 997 fJ/b/mm (768 fJ/b/mm assuming nine data lanes) even when the effective data rate is scaled by more than 100$\times$ to about 80 Mb/s.

Rest of this paper is organized as follows. Section II presents the proposed architecture. Sections III and IV describe circuit implementation of a transmitter and a receiver, respectively. Experimental results from the test chip are presented in Section V. Key contributions of this paper are summarized in Section VI. Channel analysis is presented in the Appendix.

II. Proposed Architecture

The block diagram of the proposed transceiver is shown in Fig. 1. The transmitter consists of parallel pseudorandom binary sequence (PRBS) generators of different lengths, a serializer that serializes 16 parallel streams of 625-Mb/s data into 10-Gb/s data, and a half-rate capacitively driven CMOS-based driver with 3-tap pre-emphasis that launches 10-Gb/s low-swing data signal onto a 0.5-μm-wide, 5-mm-long wire (cross section is shown in Fig. 2) using a 5-GHz low-jitter clock provided by a fast power-on-lock multiplying delay-locked loop (MDLL).

The clock is forwarded to the receiver on a dedicated channel using a differential capacitively driven output driver. Differential signaling helps to minimize duty cycle distortion, and because clock lane can be amortized among several data lanes, it incurs insignificant area or power overhead. A dedicated 7-bit phase interpolator (PI) is used in each transmitter to correct path mismatch between the data and CLK channels and ensure near optimal sampling of the data by the clock on the receive end.

Crosstalk limits maximum achievable data rates and data throughput in on-chip interconnects [17], while crosstalk can be partially mitigated by implementing twisted routing [8]. Such twisting, however, requires two dedicated metal stacks for channels as opposed to one, which severely constraints lower level floor planning. In this paper, a single-ended data lane shielded by a ground lane is used to reduce crosstalk and achieve high throughput density. Simulations based on an electromagnetic field solver high frequency structure simulator (HFSS) indicate that shielding improves signal-to-crosstalk ratio by 3 dB compared with unshielded differential signaling for the target throughput density. In the simulation, the physical model of the channel together with the surrounding dielectric layers is built using HFSS, and s-parameters are extracted from the physical model by a full-wave solver. The crosstalk simulation is based on the fully extracted s-parameter model. Note that unshielded differential signaling also suffers from the same crosstalk problem because of the asymmetry between the aggressors present on either side of the differential traces (see the Appendix for detailed analysis). To obtain the same signal-to-crosstalk ratio as the proposed scheme, unshielded differential signal channel needs either increased spacing or additional shield channel, both of which lower throughput density.

The receiver consists of a broadband data amplifier, two half-rate charge-based samplers, and a 2:16 de-serializer. The samplers are clocked by the amplified and duty-cycle corrected forwarded clock. De-serializer generates 16 parallel streams of 625-Mb/s data from 10-Gb/s data. A parallel PRBS checker verifies the correctness of the recovered data and flags an Error signal if any bit is found to be in error.

Transceiver power is scaled in proportion to utilization by cycling the transceiver between active and sleep modes. Externally provided signal (WKP_{TX}) indicating inactive period (lack of traffic) puts the link in sleep mode during which all the transmitter, receiver, and clocking circuits are powered down to bring the link power to near zero (155 μW in the prototype). When WKP_{TX} signal is asserted high, the link exits the sleep state and enters the active state in the following sequence: 1) MDLL and PI bias are powered ON by WKP_{TX} signal that is synchronized to MDLL reference clock; 2) a line driver is powered ON and the common-mode voltage is restored; 3) WKP_{TX} is transmitted to the receiver where it is used to turn on amplifiers' bias circuitry; and 4) all the digital circuits are turned on once the MDLL reaches steady state. The power down process follows the above steps executed in the reverse order. All the circuits are designed, as described later, to turn on rapidly and minimize exit latency (less than 17 ns in the prototype). Near zero OFF-state power and very small exit latency help achieve linear power scaling across a very wide range of traffic profiles.

We note that an ideal rapid-ON/OFF link should consume zero OFF-state power and turn on from the idle state instantly. However, this is difficult to achieve in practice because of slow-settling circuits. For instance, clock generators implemented using phase-locked loops (PLLs) take a long time to reach steady state, thereby greatly increasing the time taken by the link to turn on.

III. Transmitter

A. Rapid-ON/OFF Clock Generator

The block diagram of a fast power-on-lock clock generator is shown in Fig. 3. Since PLLs suffer from sluggish phase locking [12], we employed a MDLL-based clock generator

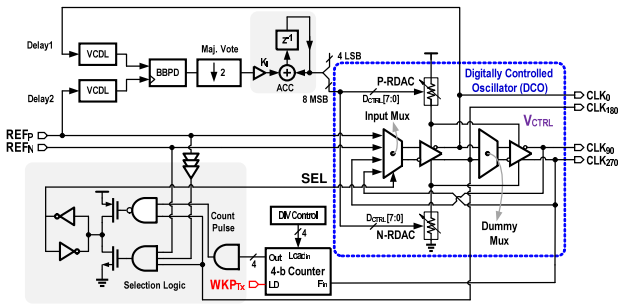

Fig. 3. Proposed MDLL architecture.

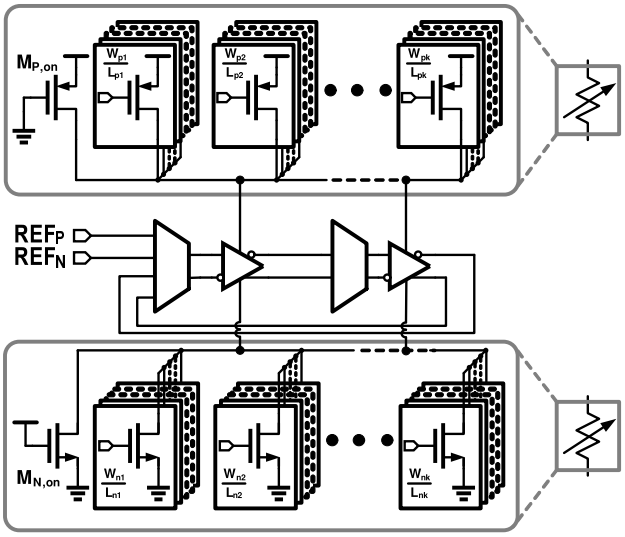

Fig. 4. rDAC schematic.

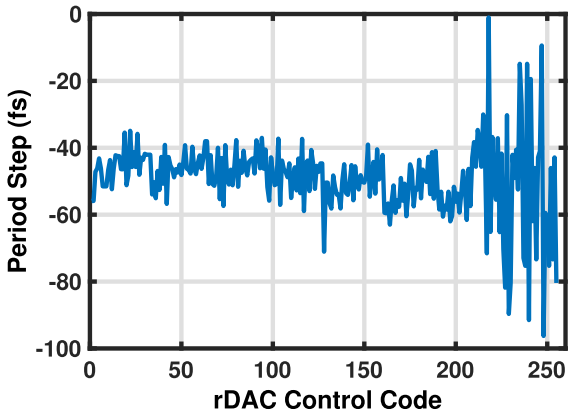

Fig. 5. Simulated MDLL period step versus code.

that is capable of almost instantaneous phase locking [11]. The proposed enhancements reduce MDLL power-on-lock to two reference cycles ($2T_{REF}$), as compared with three reference cycles for past fast-locking MDLLs [11]. The MDLL consists of a bang–bang phase detector (BBPD), a digital loop filter implemented using an accumulator, resistive digital-to-analog converter (DAC) (rDAC), voltage-controlled oscillator (VCO), and the selection logic. The proposed MDLL generates a 5-GHz output clock from a 312.5-MHz reference clock. BBPD measures the phase difference between reference and feedback clocks in the form of early/late signals that are decimated by a factor of 2 by the majority voter and accumulated by a 12-bit accumulator clocked at a reduced frequency of 156.25 MHz. To reduce dithering jitter resulting from bang–bang behavior of the feedback loop, four Least significant bits (LSBs) of the 12-bit accumulator output are dropped. Note that this also reduces the integral path bandwidth by 16×. The remaining eight most significant bits are fed to the rDAC, which controls the VCO output frequency. Since the MDLL frequency is controlled by an 8-bit Nyquist frequency DAC, the quantization noise is dictated by the DAC step size and dropping the four LSBs does not affect the quantization noise.

Replacement of the every 16th VCO output edge with reference clock edge is performed by selection logic. The two voltage-controlled delay lines (VCDLs) at the input of BBPD are used to cancel static phase offset resulting from mismatches between the reference and feedback paths. Without these two VCDLs, any delay mismatch will be translated into residual phase error and injected into the VCO loop. The periodically injected phase error at reference frequency will then result in a large reference spur. The wake-up signal, WKP_{TX}, re-timed by the reference clock powers ON/OFF the MDLL.

The VCO is implemented with a four-stage ring oscillator, where only two of the stages are controlled to reduce gain. A MUX selects VCO edge during normal operation and reference edge during edge injection. Output of the 8-bit accumulator tunes VCO frequency by varying the delay cell supply voltage through rDAC shown in Fig. 4. Nyquist-rate rDAC is chosen to achieve fast settling compared with current-mode or over-sampling DACs [18]. The rDAC is divided into eight banks, each of which containing parallel combination of MOS transistors [18]. Transistors within each bank are sized to achieve linear transfer characteristic. Thermometer coding is used to guarantee monotonicity. The control code versus VCO period step is shown in Fig. 5. The period resolution is about 100 fs and the tuning range in typical process corner is 4.9 to 5.2 GHz. While the process variation is compensated by tuning the load capacitor, the 300-MHz tuning range through rDAC allows MDLL to operate at 5 GHz across a wide range of temperatures and supply voltages. The worst case peak to peak jitter caused by quantization error is 1.6 ps.

While rDAC enables fast settling of the VCO frequency, it is susceptible to supply noise. Although on-chip decoupling capacitors can compensate for supply noise during run time, supply droops during a turn-on event can significantly affect the locking behavior of the MDLL and increase its settling time. While power-ON time of the MDLL can ideally be very small (about 1 T_{REF}) [11], settling transient of VCO control voltage can increase locking time to several T_{REF} values, as shown in Fig. 6. Since control voltage is pulled-up to V_{DD} during the OFF state, VCO runs faster than target frequency immediately after it is turned on, as depicted by region 1. Increased supply impedance due to bond wire inductance limits the supply current during the power-ON transient, which

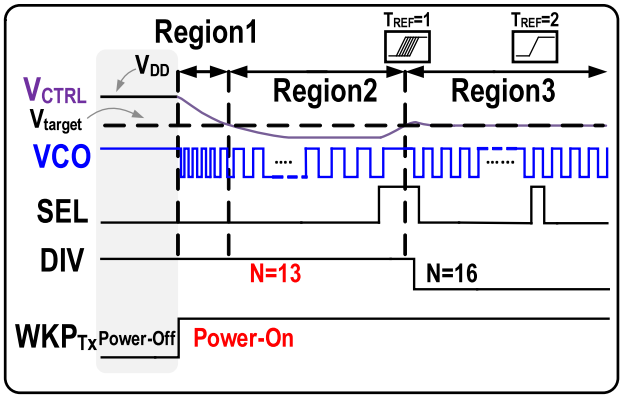

Fig. 6. MDLL control voltage (V_{ctrl}) settling behavior.

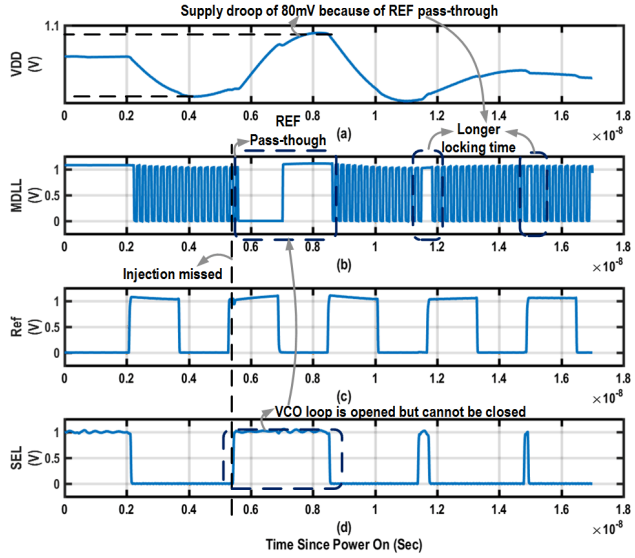

Fig. 7. Simulated MDLL power-ON behavior. (a) Vctrl settling waveform. (b) MDLL output settling waveform. (c) Reference clock waveform. (d) Selection logic output waveform.

causes internal supply voltage to droop due to discharging of the on-chip decoupling capacitor by the MDLL current (see Fig. 6). This large voltage droop (about 80 mV in the illustrated case in Fig. 7) slows down VCO, causing it to take longer than T_{REF} to complete $N = 16$ cycles. When VCO finally finishes the 16th cycle, VCO loop is broken by the selection logic and kept open until the next reference edge is injected, resulting in passing through of a full reference cycle. This phenomenon significantly increases the power-on-lock time. In order to reduce the power-on-lock time in the presence of supply droops, an architectural solution is proposed. A programmable divider is used to temporarily decrease N (to $N = 13$ or 14) during the power-ON process, which allows injection of the reference edge. This reduces (di/dt) noise and controls voltage ripple. Compared with the fixed divider ratio case, using a programmable divider reduced the supply droop by 50%, and increased timing margin of the selection signal to 0.5 ns. This enabled MDLL to lock within $2T_{\text{REF}}$, as shown in Fig. 8. Note that the temporary division ratio cannot be set too small as it will stop the oscillator too soon. This causes control voltage to go higher than the target value, which increases settling time. In addition, the temporary

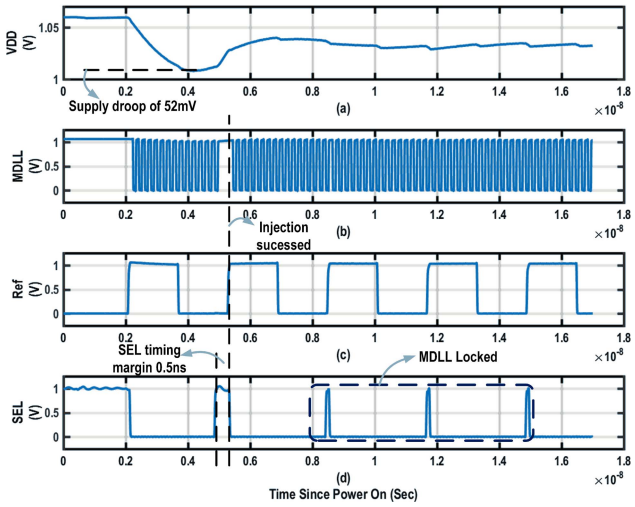

Fig. 8. Simulated MDLL power-ON behavior with a programmable divider. (a) Vctrl settling waveform. (b) MDLL output settling waveform. (c) Reference clock waveform. (d) Selection logic output waveform.

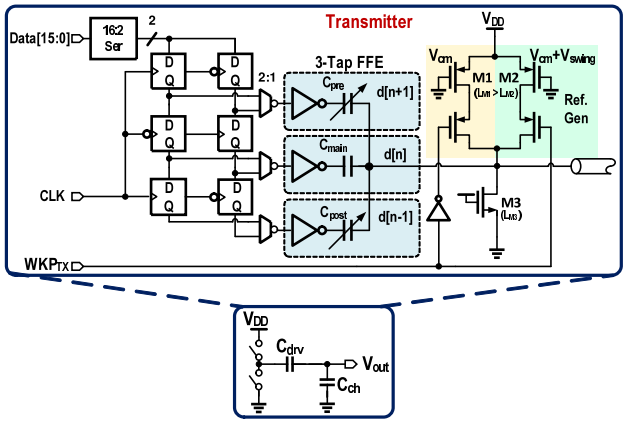

Fig. 9. Block diagram of the proposed line driver.

division ratio is only determined by the supply droop during fast-on process, which in turn is determined by the load current step. Since the load current step and the parasitic inductance do not significantly vary with temperature, supply voltage, or process, the same temporary division ratio can be used across process voltage and temperature and at all effective data rate conditions.

B. Line Driver With Capacitive Feed Forward Equalization (FFE)

A voltage-mode driver has better energy efficiency compared with a current-mode driver [12], [19]. Unfortunately, separate supply rail required to generate low-swing signals complicates the design of a voltage-mode driver. As the data rate significantly exceeds 3-dB bandwidth of the channel, transmitter side equalization becomes inevitable. Such equalization can be performed using a linear feed-forward equalization (FFE), implemented using a DAC [20]. However, such a DAC requires partitioning the line driver and pre-drivers depending on the desired FFE coefficient resolution, which increases the switching and signaling power [21].

Capacitively driven line driver-based FFE is an excellent alternative to decouple the tradeoff between complexity and

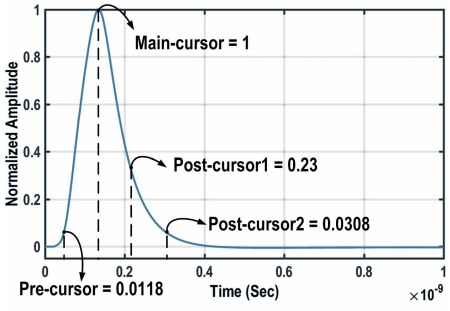

Fig. 10. Simulated channel pulse response.

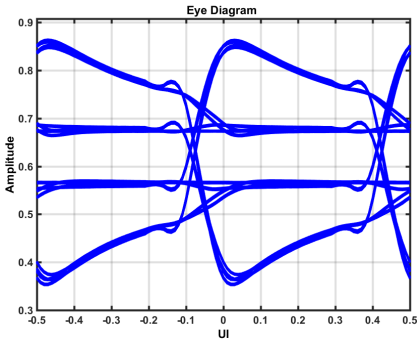

Fig. 11. Simulated transmitted eye diagram with FFE.

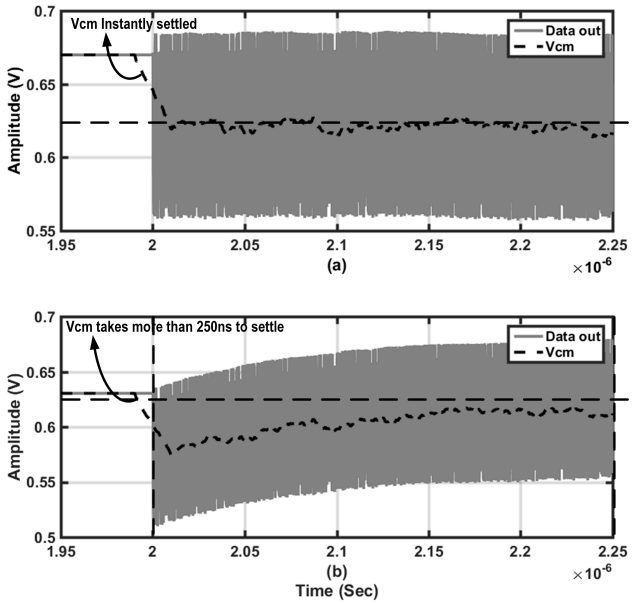

Fig. 12. Simulated V_{cm} settling during the power-ON process. (a) Channel common mode voltage with second dc line driver. (b) Channel common mode voltage without second dc line driver.

energy efficiency present in a voltage- and current-mode driver [8]. As shown in Fig. 9, output swing is governed only by the ratio of output capacitance to the channel capacitance. Therefore, output swing can be accurately programmed by adjusting the output capacitance. Also, the outputs of multiple drivers can be shorted without creating a short circuit between supply and ground. As shown in Fig. 9, a smaller auxiliary driver is shorted with the main driver. By setting the polarity of the auxiliary driver to be positive or negative, the induced charge on output capacitors can be enforced or partially canceled, respectively. Thus, a voltage-mode driver with feed forward equalization (FFE) can be implemented without using a secondary supply rail as follows:

$$\begin{aligned} V_{\text{out}} &= \frac{V_{\text{DD}}}{\frac{C_{\text{ch}}}{C_{\text{drv}}}+1} \\ &\approx V_{\text{DD}}\cdot\frac{C_{\text{drv}}}{C_{\text{ch}}} \quad \left(\text{since } \frac{C_{\text{ch}}}{C_{\text{drv}}} \gg 1\right) \\ &= V_{\text{DD}}\cdot\frac{d[n-1]\cdot C_{\text{post}}+d[n]\cdot C_{\text{main}}+d[n+1]\cdot C_{\text{pre}}}{C_{\text{ch}}}. \end{aligned} \quad (1)$$

Pulse response of the channel obtained using electromagnetic field solver HFSS is shown in Fig. 10. Equalizer coefficients of -0.184 and -0.0179 are used to suppress post-cursor inter-symbol interference. The simulated transmitted eye with 3-tap FFE is shown in Fig. 11. Peak distortion analysis indicates that FFE improves worst case vertical eye opening by more than 58% (from 45.7 to 72 mV).

C. DC Driver

Since neither transmitter nor receiver is terminated, a dedicated dc driver is needed to set the common-mode voltage V_{cm} of the data line. To this end, a transistor-based voltage divider is used to set the common voltage. As shown in Fig. 9, a long-channel PMOS in series with an NMOS transistor forms a voltage divider. Channel lengths of both the PMOS and NMOS transistors are made programmable with 3-bit resolution. The resistance is chosen to keep static current below 10 μA per branch. Because driver's output capacitor and the dc driver form a high-pass filter, its 3-dB bandwidth is designed to be low enough to support at least PRBS-31 data. It should be noted that when the transmitter enters sleep mode, output of the driver is pulled high and the dc line driver should maintain that voltage level. To achieve this, a second branch is added. When WKP_{TX} signal is asserted, the main branch is activated and the common-mode voltage of the channel is kept at V_{cm}; when it is de-asserted, the main branch is deactivated, and the auxiliary branch is activated to change the common-mode voltage of the channel to $V\text{cm} + V\text{swing}$. Fig. 12(a) and (b) shows the simulation of power-ON settling difference between with and without auxiliary branches, respectively. Turn-on time increases by at least 250 ns in the absence of an auxiliary branch.

IV. Receiver

A. Data Amplifier

The channel output is fed into the data amplifier that is implemented using a resistively loaded common source amplifier shown in Fig. 13. A load resistor is set to 500 Ω considering the tradeoff between gain, bandwidth, and power consumption and to also produce output common-mode voltage (800 mV) that is within the input common-mode range of the following sampler stage. Two 6-bit resolution programmable current sources with a full-scale current of 63 μA are attached to the output to perform offset cancellation. Monte Carlo simulations indicate an output referred offset of ±14.1 mV

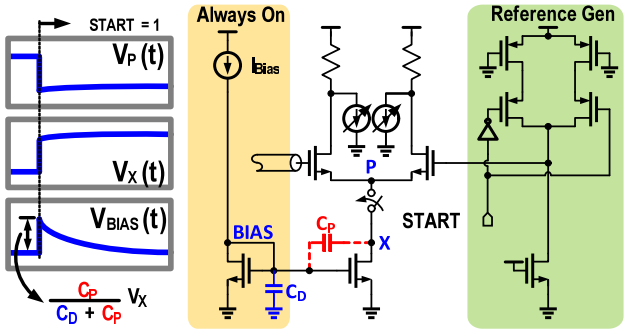

Fig. 13. Schematic of the proposed data amplifier.

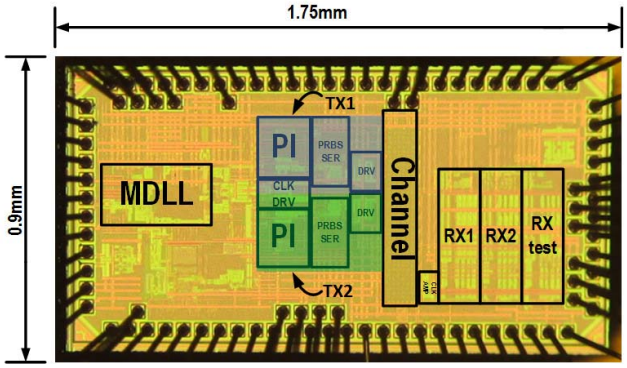

Fig. 14. Die photograph.

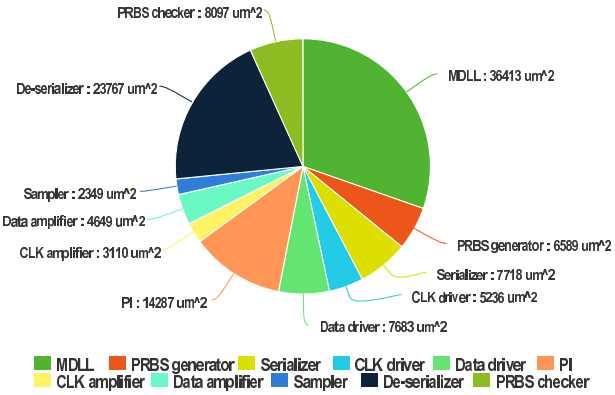

Fig. 15. Area breakdown.

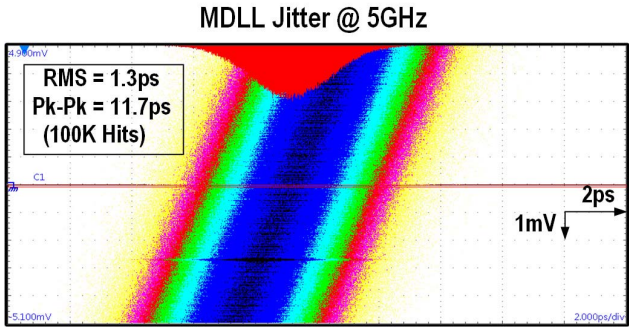

Fig. 16. Measured MDLL jitter at 5-GHz output frequency.

($\pm 3\sigma$), which mandates a resolution of 1-μA resolution to cancel the offset to within 0.5 mV. The data amplifiers are turned off during the idle state by using the switch inserted between the input pair and tail current source. The switch is sized to minimize kickback on to the bias node during the power-ON process. Because kickback is governed by the capacitive voltage divider formed by C_P and C_D, C_D is made large compared with C_P. Reference generator is implemented using a replica of the dc line driver of the transmitter. It is kept always on to reduce power-ON time; this adds about 10-μA idle current penalty per lane.

B. Clock Amplifier

The clock channel's output is fed into the clock amplifier, which has an architecture similar to that of the data amplifier. A source degeneration capacitor is added between two differential branches to reduce the impact of common-mode mismatch on the output clock duty cycle. Output of the clock amplifier is ac coupled using a 220-fF metal–oxide–metal capacitor and fed to a self-biased inverter that generates a rail-to-rail clock signal.

V. Measurement Results

The prototype transceiver is implemented in a 65-nm CMOS process; the die photograph is shown in Fig. 14 and the area breakdown is shown in Fig. 15. We note that the circuit area can be scaled with technology, but the channel footprint cannot in order to achieve the same channel frequency response. The die is packaged in a 10 mm $\times$ 10 mm 88-pin QFN plastic package. The critical supplies of MDLL, PI, and line drivers are carefully designed to avoid supply voltage droops caused by the large current step during a turn-on event. Due to high supply sensitivity of the MDLL, care was taken to reduce power supply inductance by dedicating four bond pads ($\sim$1-nH supply inductance) to MDLL supply, while all other supplies used two bond pads ($\sim$2-nH supply inductance). A large number of pads are used to reduce the impact of long bonding wires of the chosen QFN package. Off-chip decoupling capacitors and damping resistors are used to prevent supply ringing. The measured transceiver's performance in always-on and rapid-ON/OFF conditions is presented in Sections A and B.

A. Always-on Measurements

The measured output jitter of the MDLL operating at 5 GHz is shown in Fig. 16. The proposed MDLL achieves 1.3-ps rms jitter and 11.7-ps pk–pk jitter. The phase noise measurement result is shown in Fig. 17. Measured reference spur is 47.8 dB, and the largest in-band spur is 41.7 dB, as shown in Fig. 18. The BER bathtub curve of the transceiver operating at 10 Gb/s with PRBS31 data under two conditions, with and without FFE, is shown in Fig. 19. These data are obtained by synchronizing one sub-rate recovered data from 1:8 de-serializer with Tektronix BSA260C BERT. The sampling position of the data eye was altered by adjusting the transmitter PI code. The bathtub curve indicates that the FFE improves horizontal eye opening by 280%, from 0.1 to 0.38 UI. The PI code and FFE ratio are set to maximize eye opening for best BER and retained through entire measurement.

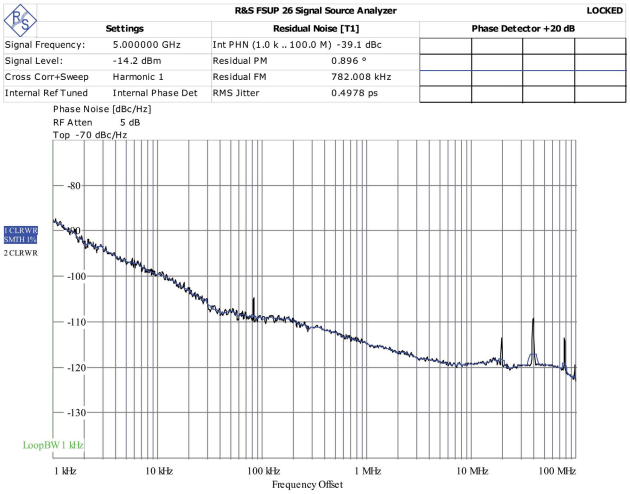

Fig. 17. Measured MDLL phase noise plot at 5 GHz.

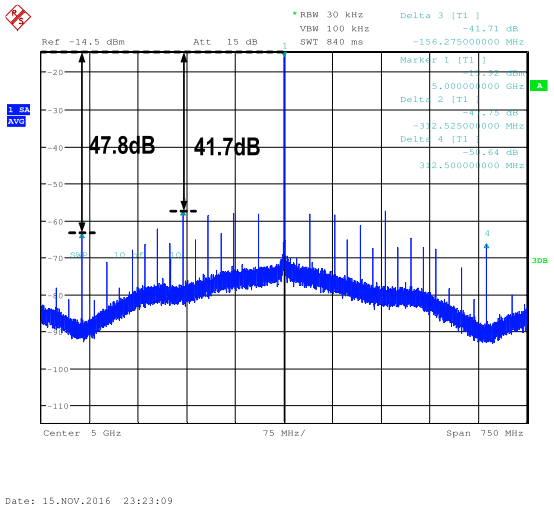

Fig. 18. Measured MDLL output voltage spectrum at 5 GHz.

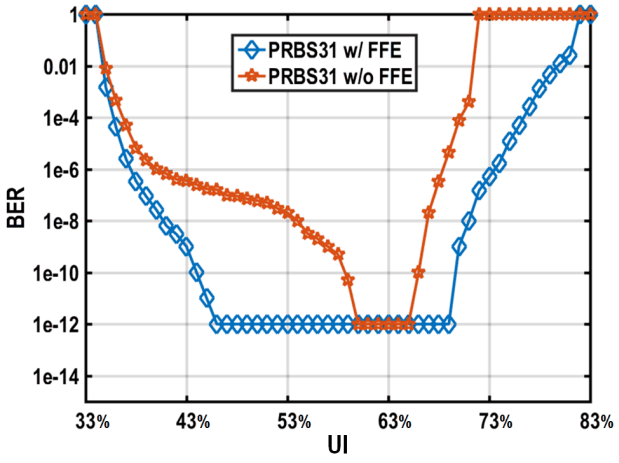

Fig. 19. Measured transceiver bathtub curve.

A similar method was used to obtain the statistical data eye as well. In addition to adjusting sampling instance, the sampler's reference voltage is also varied. Fig. 20 shows the receiver data eye with and without FFE. It confirms that the FFE increased horizontal opening by 280%, and vertical

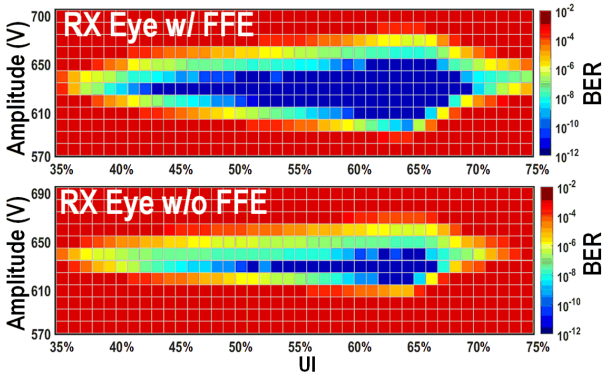

Fig. 20. Measured receiver eye diagram.

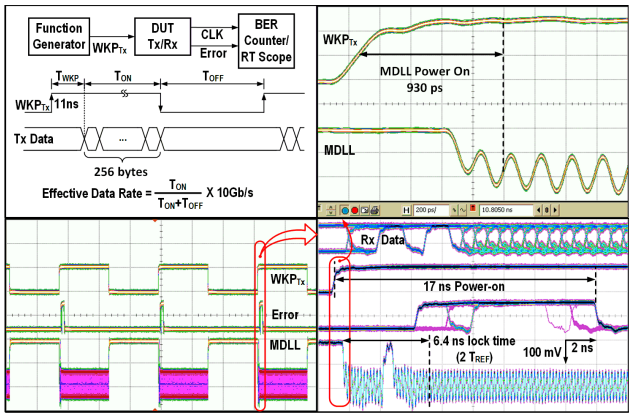

Fig. 21. Measured transceiver ON/OFF behavior.

opening by 133% (from 30 to 70 mV), thus proving the effectiveness of the proposed FFE scheme.

B. Rapid-ON/OFF Measurements

Prior to turn off, the pre-calibrated PI control code, the RX amplifier mismatch control code, the FFE ratio, and the MDLL frequency code are stored digitally during the OFF state. These control codes are restored at turn on so that the wake-up time is minimized. This assumes no temperature drift during the power-OFF state. However, if the temperature changes, then either the duration of OFF time must be limited or lookup table-based compensation [13] must be used.

Fig. 21 shows the measurement setup and the ON/OFF performance of the transceiver. PRBS31 data are transmitted through the on-chip channel, recovered by the receiver, and examined by the on-chip PRBS checker. Approximately, 3.6 billion ON/OFF transitions are captured, in these measurements. The error signal is generated by the PRBS checker and captured by the high-speed real-time scope that goes low within 17 ns (170 bits), which indicates the transceiver operates error free after 17 ns from the power-ON event. Of the 17-ns power-ON time, MDLL needs 1 ns to start and two reference injections of 6.4 ns to settle. This represents 33% improvement over prior work [11] thanks to the proposed programmable divider. The PRBS checker seeding latency is

TABLE I
PERFORMANCE COMPARISON OF THE PROPOSED TRANSCEIVER

	This work	JSSC'15 [8]	IITC'07 [24]
Link Data Rate (Gb/s)	**10**	10	10
Channel Length (mm)	**5**	6	5
Technology	**65nm**	65 nm	90 nm
Throughput Density (Gbit/s/um)	**5**	2.1	0.5
TX swing (V)	**0.45**	N/A	0.12
Equalization	**3-tap FFE**	Pre-emphasis	Passive
Fast On-Off Functionality	**Yes**	Yes	No
Off state power (mW)	**0.155**	4.5	N/A
Energy Efficiency;TX, RX only (fJ/b/mm)	**148**	48	54
Total Energy Efficiency; 1 Lane (fJ/b/mm)*	**627*****	N/A	N/A
Total Energy Efficiency; 9 Lane (fJ/b/mm)	**472****,***	255	N/A
Channel loss (5GHz) (dB)	**16.6**†	12	< 8
Energy efficiency per unit loss (fJ/b)††	**51.1**	96.5	N/A

* Total power including SerDes, Pattern generation, Clocking, TX and RX.
** Assuming there are 9 data lane to share the clocking power.
*** The energy efficiency is traded off for higher throughput density.
† Estimated channel loss.
†† Calculated by $\frac{\text{Energy/bit/Channel}}{\text{Channel loss in linear scale}}$.

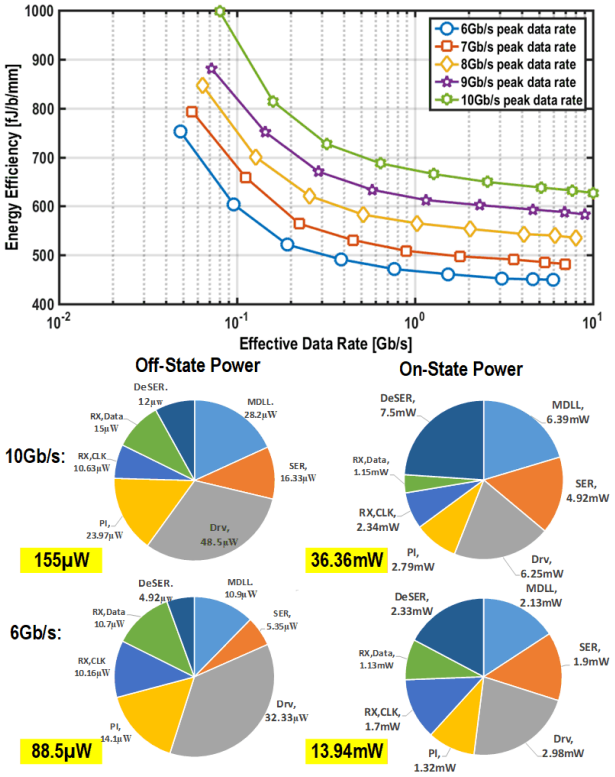

Fig. 22. Measured energy efficiency versus effective data rate and power breakdown in ON state and OFF state.

roughly 7–9 ns. We believe that the extra power-ON time is due to the slow common-mode settling of the channel.

The effective data rate is scaled by changing the duty cycle of the WKP_{TX} signal, and the resulting power-scalable behavior of the transceiver is shown in Fig. 22. More than 125× effective data rate scaling (10 Gb/s to 80 Mb/s) is achieved with an energy efficiency degradation of only 1.6× (627 fJ/b/mm to 997 fJ/b/m). When the supply voltage is scaled from 1 to 0.7 V, peak data rate ranges from 10 to 6 Gb/s and the power scalable range increases to 208× (10 Gb/s to 48 Mb/s) with energy efficiency degradation of only 1.2× (627 fJ/b/mm at 1 V to 753 fJ/b/mm at 0.7 V). The wake-up signal path contains only a few minimum-sized inverters operating at low frequency (several MHz), which consume negligible power.

The performance of the transceiver is compared with the state-of-the-art in Table I. To the best of our knowledge, the proposed transceiver achieves the smallest turn-on time, least OFF-state power, and almost constant energy efficiency.

Other solutions to further decrease the exit latency include keeping MDLL always on. In this case, exit latency can be reduced to sub 10 ns but the OFF-state power will increase from 155 uW to 6.4 mW. A 2× reduction in exit latency results in a 41× increase in the OFF-state power, which severely degrades the overall power efficiency.

VI. CONCLUSION

On-chip interconnects used in NoC applications must deliver high bandwidth with low communication latency and high power efficiency. Because of the bursty traffic, the utilization of these on-chip interconnects varies with application. Therefore, it is paramount to maintain excellent power efficiency across all utilization levels. However, the power efficiency of conventional interconnects severely degrades at low utilization levels, because they consume large power even when they are idle. This paper presents techniques to implement power efficient transceivers for on-chip interconnects that can eliminate idle power almost entirely, thus achieving energy proportional operation. Fabricated in 65-nm CMOS technology, the proposed 10-Gb/s transceiver delivers 5-Gb/s/um throughput density and power-on-lock time of less than 17 ns thanks to the programmable divider that allows MDLL to lock within 2 T_{ref}. In rapid-ON/OFF mode, more than 125× effective data rate scaling (10 Gb/s to 80 Mb/s) is achieved with an energy efficiency degradation of only 1.6× (627 fJ/b/mm to 997 fJ/b/m). When the supply voltage

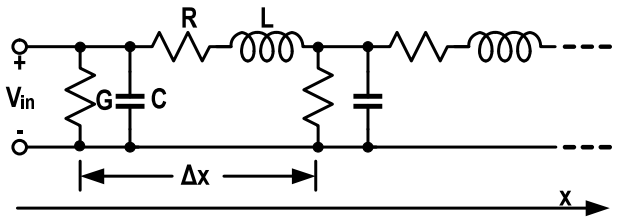

Fig. 23. General transmission line model.

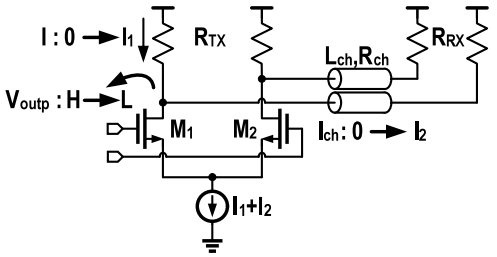

Fig. 24. General transceiver model.

is scaled from 1 to 0.7 V, peak data rate ranges from 10 to 6 Gb/s and the power scalable range increases to 208× (10 Gb/s to 48 Mb/s) with energy efficiency degradation of only 1.2× (627 fJ/b/mm at 1 V to 753 fJ/b/mm at 0.7 V).

APPENDIX
CHANNEL ANALYSIS

A. RC-Dominant Channel

A basic transmission line model of an on-chip wire is shown in Fig. 23. Each distributed section contains a series indicator (L) and resistor (R) that come from the wire and a parallel resistor (G) and capacitor (C) between wire and ground that models the dielectric loss and parasitic capacitance to ground. The voltage along the transmission line should satisfy the following equation at time t:

$$\frac{\partial v}{\partial x} = -\mathrm{Ri} - L\frac{\partial i}{\partial t}. \tag{2}$$

To validate RC-dominant channel assumption, $|\mathrm{Ri}| \gg |L(\partial i/\partial t)|$ should be true. Consider a general current-mode driver case in Fig. 24, when output transits from high to low, M1 turns on, pulling V_{outp} to V_L, and sinking I_1 from supply through R_{TX} and I_2 through channel. Assuming the transition time of 20% of a bit period T, right after the transition, it can be written $\partial i = I_2$ and $\partial t = 0.2$ T. Furthermore, the condition becomes

$$\begin{aligned} R_{\mathrm{ch}} I_2 &\gg L_{\mathrm{ch}} \frac{I_2}{0.2T} \\ T &\gg \frac{5L_{\mathrm{ch}}}{R_{\mathrm{ch}}} \\ D &\ll \frac{R_{\mathrm{ch}}}{5L_{\mathrm{ch}}} \end{aligned} \tag{3}$$

where D is the data rate. Even though the proposed channel is not terminated, the bit period is much smaller than the channel bandwidth, and thus constant current charging/discharging during bit transition can be assumed. From the electrical-magnetic solver, R_{ch} and L_{ch} can be extracted and their values are 191.57 Ω and 2.72 nH, respectively. According to (3), the maximum data rate when RC-dominant assumption is still valid is 14.19 Gb/s, 42% higher than the target data rate of the proposed transceiver, and thus the RC-dominant assumption is valid. Also, frequency response of the channel is extracted by full-wave simulation. The channel loss at 5 GHz is about 16.6 dB, which indicates that reflected wave will be attenuated by 33.2 dB before reaching the far end.

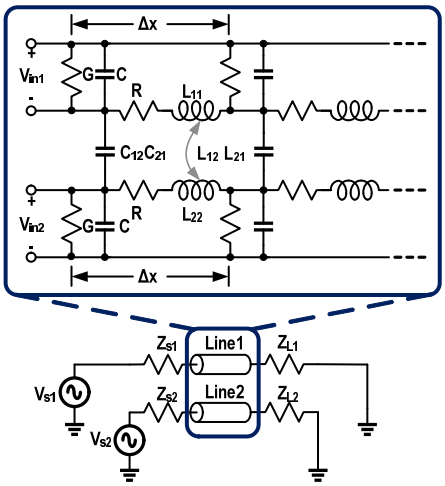

Fig. 25. General coupled transmission line.

B. Crosstalk Analysis

To understand crosstalk, a generic coupled line system is examined (shown in Fig. 25). In addition to the RLGC shown in Fig. 23, mutual capacitance and inductance between two lines are also included. Further assuming, line 1 is the aggressor with periodic excitation V_{s1}, line 2 is the victim line with no excitation, both lines have large attenuation thus reflection is near zero, both lines are terminated with the even-mode impedance ($Z_s = Z_e$) at near end and odd-mode impedance ($Z_L = Z_d$) at far end, then the voltage at near end becomes

$$\begin{aligned} V_1(0) &= \frac{V_{s1}}{4} + \frac{V_{s1} Z_d}{2(Z_d + Z_e)} \\ V_2(0) &= \frac{V_{s1}}{4} - \frac{V_{s1} Z_d}{2(Z_d + Z_e)}. \end{aligned} \tag{4}$$

For the proposed channel, the RLGC parameters of two adjacent channels can be obtained through an electrical-magnetic solver. The simulation indicates that when compared with the case without the ground shielding, the crosstalk between shield single-ended lines is reduced by 75.4%.

The case of differential signaling can be studied by the same way. The channel configuration with the same throughput density is shown in Fig. 26. In this case, line 2 is the victim of line 1, while line 3 is the victim of line 2. The direct coupling between line 1 and line 3 is weak and ignored for easy analysis (not ignored in the simulation). Again, the simulation shows that the crosstalk between unshielded differential lines is 229% more than the proposed channel, indicating that the single-ended signaling with shielding is the better choice in this case.

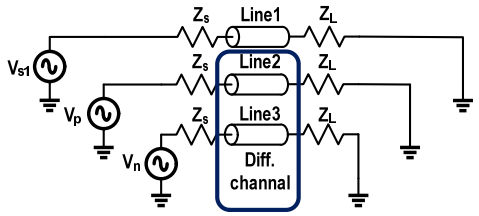

Fig. 26. Differential channel configuration.

Full-wave simulation indicates that at worst case frequency (5 GHz), the signal-to-crosstalk ratio of single-ended signaling is 3 dB better than the signal-to-crosstalk ratio of unshielded differential signaling.

REFERENCES

[1] R. Marculescu, U. Y. Ogras, L.-S. Peh, N. E. Jerger, and Y. Hoskote, "Outstanding research problems in NoC design: System, microarchitecture, and circuit perspectives," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 28, no. 1, pp. 3–21, Jan. 2009.

[2] K. Lee *et al.*, "A 51 mW 1.6 GHz on-chip network for low-power heterogeneous SoC platform," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, vol. 1, Feb. 2004, pp. 152–518.

[3] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. V. Tuijl, and B. Nauta, "Low-power, high-speed transceivers for network-on-chip communication," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 17, no. 1, pp. 12–21, Jan. 2009.

[4] B. Kim and V. Stojanovic, "A 4 Gb/s/ch 356 fJ/b 10 mm equalized on-chip interconnect with nonlinear charge-injecting transmit filter and transimpedance receiver in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2009, pp. 66–67.

[5] S.-K. Lee, S.-H. Lee, D. Sylvester, D. Blaauw, and J.-Y. Sim, "A 95 fJ/b current-mode transceiver for 10 mm on-chip interconnect," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 262–263.

[6] A. E. Kiasari, Z. Lu, and A. Jantsch, "An analytical latency model for networks-on-chip," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 21, no. 1, pp. 113–123, Jan. 2013.

[7] T. Anand, A. Elshazly, M. Talegaonkar, B. Young, and P. K. Hanumolu, "A 5 Gb/s, 10 ns power-on-time, 36 μW off-state power, fast power-on transmitter for energy proportional links," *IEEE J. Solid-State Circuits*, vol. 49, no. 10, pp. 2243–2258, Oct. 2014.

[8] S. Höppner *et al.*, "An energy efficient multi-Gbit/s NoC transceiver architecture with combined AC/DC drivers and stoppable clocking in 65 nm and 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 3, pp. 749–762, Mar. 2015.

[9] J.-S. Seo, R. Ho, J. Lexau, M. Dayringer, D. Sylvester, and D. Blaauw, "High-bandwidth and low-energy on-chip signaling with adaptive pre-emphasis in 90 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2010, pp. 182–183.

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

[11] T. Anand, M. Talegaonkar, A. Elshazly, B. Young, and P. K. Hanumolu, "A 2.5 GHz 2.2 mW/25 μW on/off-state power 2 ps_{rms}-long-term-jitter digital clock multiplier with 3-reference-cycles power-on time," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 256–257.

[12] B. Leibowitz *et al.*, "A 4.3 GB/s mobile memory interface with power-efficient bandwidth scaling," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 889–898, Apr. 2010.

[13] T. Anand, M. Talegaonkar, A. Elkholy, S. Saxena, A. Elshazly, and P. K. Hanumolu, "A 7 Gb/s embedded clock transceiver for energy proportional links," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3101–3119, Dec. 2015.

[14] J. Zerbe *et al.*, "A 5.6 Gb/s 2.4 mW/Gb/s bidirectional link with 8 ns power-on," in *Symp. VLSI Circuits, Dig. Tech. Papers*, Jun. 2011, pp. 82–83.

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

[16] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. A. Maestro, "IEEE 802.3az: The road to energy efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, no. 11, pp. 50–56, Nov. 2010.

[17] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "Optimally-placed twists in global on-chip differential interconnects," in *Proc. 31st Eur. Solid-State Circuits Conf. (ESSCIRC)*, Sep. 2005, pp. 475–478.

[18] M. Talegaonkar, A. Elshazly, K. Reddy, P. Prabha, T. Anand, and P. K. Hanumolu, "An 8 Gb/s–64 Mb/s, 2.3–4.2 mW/Gb/s burst-mode transmitter in 90 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 10, pp. 2228–2242, Oct. 2014.

[19] A. K. Joy *et al.*, "Analog-DFE-based 16 Gb/s SerDes in 40 nm CMOS that operates across 34 dB loss channels at Nyquist with a baud rate CDR and 1.2 Vpp voltage-mode driver," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2011, pp. 350–351.

[20] E.-J. Choi, S. Kim, Y.-K. Jeong, K.-W. Kwon, and J.-H. Chun, "A low-swing AC− and DC− coupled voltage-mode driver with pre-emphasis," in *Proc. IEEE 54th Int. Midwest Symp. Circuits Syst. (MWSCAS),* Aug. 2011, pp. 1–4.

[21] Y. Lu, K. Jung, Y. Hidaka, and E. Alon, "Design and analysis of energy-efficient reconfigurable pre-emphasis voltage-mode transmitters," *IEEE J. Solid-State Circuits*, vol. 48, no. 8, pp. 1898–1909, Aug. 2013.

[22] H. Ito, J. Seita, T. Ishii, H. Sugita, K. Okada, and K. Masu, "A low-latency and high-power-efficient on-chip LVDS transmission line interconnect for an RC interconnect alternative," in *Proc. IEEE Int. Interconnect Technol. Conf.*, Jun. 2007, pp. 193–195.

Da Wei (S'15) received the B.S. (Hons.) and M.S. degrees in electrical and computer engineering from the University of Illinois at Urbana–Champaign, Urbana, IL, USA, in 2012 and 2014, respectively, where he is currently pursuing the Ph.D. degree in electrical and computer engineering.

In 2014, he was an IC Design Engineer Intern with the Finisar Corporation IC Engineering Office, Champaign, IL, USA. In 2014, he was a Mixed-Signal Validation Engineer Intern with NVIDIA, Santa Clara, CA, USA. In 2017, he was a Research Engineer Intern with Samsung Semiconductor, Inc., San Jose, CA, USA. His current research interests include energy-efficient on-chip wireline communication systems and fast transient response power converters.

Tejasvi Anand (S'12–M'15) received the M.Tech. degree (Hons.) in electronics design and technology from the Indian Institute of Science, Bengaluru, India, in 2008, and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana–Champaign, Urbana, IL, USA, in 2015.

From 2008 to 2010, he was an Analog Design Engineer with Cosmic Circuits (now Cadence), Bengaluru. He was with the IBM Thomas. J. Watson Research Center, Yorktown Heights, NY, USA, in 2015. He is currently an Assistant Professor with the Department of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA. His current research interests include wireline communication, energy harvesters, and frequency synthesizers with an emphasis on energy efficiency.

He was a recipient of the 2014–2015 IEEE Solid-State Circuits Society Predoctoral Achievement Award, the 2015 Broadcom Foundation University Research Competition Award, the 2015 M. E. Van Valkenburg Graduate Research Award from the University of Illinois at Urbana–Champaign, the 2013 Analog Devices Outstanding Student Designer Award, the 2009 CEDT Design (Gold) Medal from the Indian Institute of Science, Bengaluru.

Guanghua Shu (S'10–M'16) received the B.Eng. degree from Beijing Jiaotong University, Beijing, China, the M.S. degree in microelectronics from Fudan University, Shanghai, China, and the Ph.D. degree from the Department of Electrical and Computer Engineering, University of Illinois at Urbana–Champaign, Urbana, IL, USA.

In 2014, he was a Research Intern with Xilinx, San Jose, CA, USA, with a focus on developing power and area-efficient parallel link architectures. He was involved in 56-Gb/s wireline receivers (both electrical and optical) with the Mixed-Signal Communication IC Design Group, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA, in 2014 and 2015. He is currently a Research Staff with Oracle Labs, Belmont, CA, USA. His current research interests include energy-efficient wireline communication systems, clocking circuits, power converters, and hardware accelerations for efficient computing systems.

He was a recipient of the Dissertation Completion Fellowship from 2015 to 2016 from the University of Illinois at Urbana–Champaign and the IEEE SSCS Predoctoral Achievement Award from 2014 to 2015. He serves as a reviewer for the IEEE JOURNAL OF SOLID-STATE CIRCUITS, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I & II, the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS, and the International Symposium on Circuits and Systems.

José E. Schutt-Ainé (S'86–M'86–SM'98–F'07) received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1981, and the M.S. and Ph.D. degrees from the University of Illinois at Urbana–Champaign (UIUC), Urbana, IL, USA, in 1984 and 1988, respectively.

From 1981 to 1983, he was with the Hewlett-Packard Technology Center, Santa Rosa, CA, USA, as an Application Engineer, where he was involved in research on microwave transistors and high-frequency circuits. In 1988, he joined the Electrical and Computer Engineering Department, University of Illinois at Urbana–Champaign, as a member of Electromagnetics and Coordinated Science Laboratories, where he is currently involved in research on signal integrity for high-speed digital and high-frequency applications. He is a consultant for several corporations. His current research interests include the study of signal integrity and the generation of computer-aided design tools for high-speed digital systems.

Dr. Schutt-Ainé was a recipient of several research awards, including the 1991 National Science Foundation (NSF) MRI Award, the NASA Faculty Award for Research in 1992, the NSF MCAA Award in 1996, and the UIUC-NCSA Faculty Fellow Award in 2000. He has received several publication awards, including the IEEE Electrical Design of Advanced Packaging and Systems Symposium-2013 Best Paper and the IEEE Electrical Performance of Electronic Packaging-2014 Best Paper. He is currently serving as a Co-Editor-in-Chief of the IEEE TRANSACTIONS ON COMPONENTS, PACKAGING AND MANUFACTURING TECHNOLOGY.

Pavan Kumar Hanumolu (S'99–M'07) received the Ph.D. degree from the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA, in 2006.

He was with Oregon State University until 2013. He is currently a Professor with the Department of Electrical and Computer Engineering and a Research Professor with the Coordinated Science Laboratory, University of Illinois at Urbana–Champaign, Champaign, IL, USA. His current research interests include energy-efficient integrated circuit implementation of analog and digital signal processing, sensor interfaces, wireline communication systems, and power conversion.