

# A 40 Gb/s Serial Link Transceiver in 28 nm CMOS Technology

Reza Navid, *Senior Member, IEEE*, E-Hung Chen, *Member, IEEE*,  
 Masum Hossain, Brian Leibowitz, *Member, IEEE*,  
 Jihong Ren, Chuen-huei Adam Chou, Barry Daly, *Member, IEEE*, Marko Aleksić, *Member, IEEE*,  
 Bruce Su, *Member, IEEE*, Simon Li, Makarand Shirasgaonkar, *Member, IEEE*,  
 Fred Heaton, Jared Zerbe, and John Eble, *Member, IEEE*

**Abstract**—A 40 Gb/s serial link interface is presented that includes four lanes of transceiver optimized for chip-to-chip communication while compensating for 20 dB of channel loss. Transmit equalization consists of a 2-tap feed-forward equalizer (FFE) while receive equalization includes a 2-tap FFE using a transversal filter, a 3-stage continuous-time linear equalizer with active feedback, and discrete-time equalizers consisting of a 17-tap decision feedback equalizer (DFE) and a 3-tap sampled FFE. The receiver uses quarter-rate double integrate-and-hold sampling. The clock and data recovery (CDR) unit uses a split-path CDR/DFE design which facilitates wider bandwidth and lower jitter simultaneously. A phase detection scheme that filters out edges affected by residual inter-symbol interference allows recovering a low-jitter clock from a partially-equalized eye. A fractional-N PLL is implemented for frequency offset tracking. Combining these techniques, the digital CDR recovers a stable 10 GHz clock from an eye containing 0.8 UI<sub>p-p</sub> input jitter and achieves 1-10 MHz of tracking bandwidth. The transceiver achieves horizontal and vertical eye openings of 0.27 UI and 120 mV, respectively, at BER = 10<sup>-9</sup>. The quad SerDes is realized in 28 nm CMOS technology. Amortizing common blocks, it occupies 0.81 mm<sup>2</sup> per lane and achieves 23.2 mW/Gb/s power efficiency at 40 Gb/s.

**Index Terms**—Active feedback continuous-time linear equalizer, chip-to-chip communications, current-integrating DFE summer, decision feedback equalizer (DFE), distributed ESD protection structure, high-speed serial link (SerDes), receive-side feed-forward equalizer (RX-FFE), split-path clock and data recovery (split-path CDR), transversal filter, wireline transceiver.

Manuscript received August 27, 2014; revised November 04, 2014; accepted November 09, 2014. Date of publication December 18, 2014; date of current version March 24, 2015. This paper was approved by Guest Editor Jeffrey Gealow. This paper is an extended version of “A 40-Gb/s Serial Link Transceiver in 28-nm CMOS Technology” presented by the same authors at the 2014 Symposium on VLSI Circuits.

R. Navid, M. Aleksić, and S. Li are with Rambus Inc., Sunnyvale, CA 94089 USA (e-mail: reza@rambus.com).

E.-H. Chen was with Rambus Inc., Sunnyvale, CA 94089 USA. He is now with MediaTek Inc., Hsinchu 30078, Taiwan.

M. Hossain is with Rambus Inc., Sunnyvale, CA 94089 USA, and also with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB Canada T6G 2V4.

B. Leibowitz was with Rambus Inc., Sunnyvale, CA 94089 USA. He is now with Apple Inc., Cupertino, CA 95014 USA.

J. Ren was with Rambus Inc., Sunnyvale, CA 94089 USA. She is now with Altera Corporation, San Jose, CA 95134 USA.

C.-H. A. Chou was with Rambus Inc., Sunnyvale, CA 94089 USA. He is now with Xilinx Inc., San Jose, CA 95124 USA.

B. Daly, B. Su, F. Heaton, and J. Eble are with Rambus Inc., Chapel Hill, NC 27514 USA.

M. Shirasgaonkar is with Rambus Inc., Bangalore 560029, India.

J. Zerbe was with Rambus Inc., Sunnyvale, CA 94089 USA. He is now with Apple Inc., Cupertino, CA 95014 USA.

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

Digital Object Identifier 10.1109/JSSC.2014.2374176

## I. INTRODUCTION

OVER the past few decades, widespread adoption of data-intensive applications such as video streaming, cloud-based computing, and virtual presence devices has led to an explosive demand for data bandwidth. In order to satisfy this demand, I/O speeds of communication systems such as routers and backplane-based servers are growing rapidly. Industry groups focusing on next-generation Ethernet connectivity, such as IEEE's 400-Gbps Ethernet Standard Task Force (P802.3bs), aim to quadruple the bandwidth of the existing 100 Gbps standards to accommodate the growing demand. This trend is expected to receive continued support in coming years as data-intensive applications gain more popularity and mass adoption.

Due to cost considerations, electrical signaling over copper is still the preferred option for backplane-side communications. For 400 Gbps connectivity, the conventional solution is to use a multilane SerDes interface with well-controlled lane-to-lane skew and an aggregate bandwidth of 400 Gbps. Unfortunately, in order to achieve 400 Gbps using today's commercially available 25 Gbps interfaces, 16 such links need to work together on the backplane causing significant system complexity. To reduce this complexity, a significant effort is underway to extend the data rate of SerDes interfaces and reduce the number of required lanes. Although several building blocks are recently reported at 60 Gbps and beyond (e.g., [1], [2]), complete CMOS transceiver solutions over 30 Gbps are scarce and are often designed for relatively benign channels (e.g., [3]).

Significant challenges exist in designing high-speed CMOS transceivers at 10's of Gbps [3]–[5]. Arguably, the most serious difficulty is insufficient effective gain-bandwidth product of low-cost CMOS transistors. Parasitic wire capacitance and specially resistance (which, in modern technologies, are becoming more problematic) are mainly to blame. This deficiency has a specially pronounced effect in the design of receiver (RX) front-end where, in order to compensate for channel loss, continuous-time linear equalizers (CTLEs) need to provide significant gain up to the Nyquist frequency. Although inductors are helpful, their assistance is limited by issues such as area overhead, phase distortion, and added parasitics due to extra routing.

The design of RX samplers that follow CTLE stages is also complicated by bandwidth deficiency. Limited bandwidth

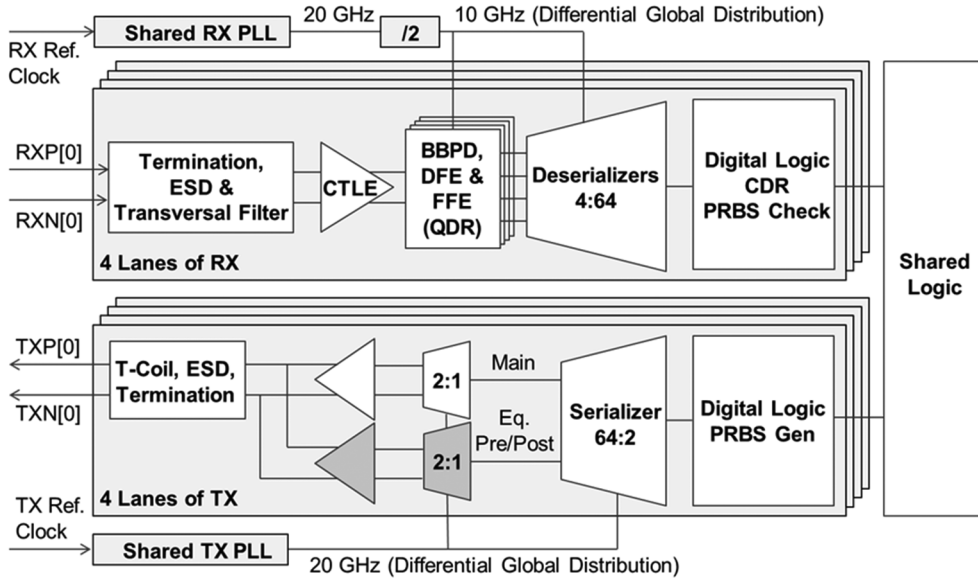


Fig. 1. Transceiver architecture: overall block diagram.

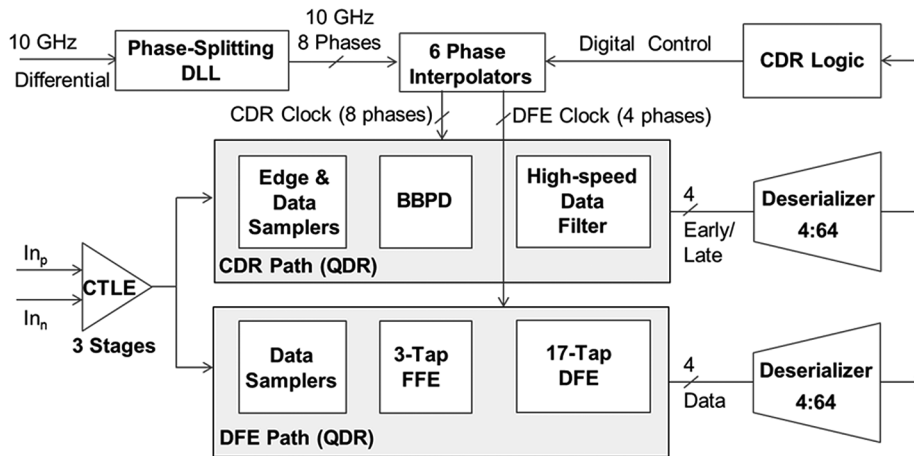


Fig. 2. Transceiver architecture: receiver details.

makes sampling apertures too wide for sampling the incoming signal at the desired rate. A combination of track-and-hold, integrate-and-hold, and interleaved sampling is sometimes used to alleviate the problem at the cost of extra power.

Performance limitations also complicate the design of decision feedback equalizers (DFEs). A short unit interval (UI) makes it difficult to close timing on the DFE loop at high data rates. A conventional solution is to unroll the loop of the first DFE tap [6]. This technique relaxes DFE timing requirements at the cost of doubling the number of samplers and hence the load on RX front-end. Unfortunately, at very high speeds, even subsequent taps (i.e., second or third) might not meet timing, despite their more relaxed requirements. Although one can potentially unroll these taps as well, this approach would lead to an exponential growth of load on the RX front-end. Furthermore, unrolled DFE taps add to the latency of the succeeding direct-feedback DFE taps, making their timing worse. This may necessitate more loop unrolling and a cycle which is hard to break [7].

Performance limitations of low-cost digital CMOS circuits also make it difficult to satisfy bandwidth requirement of the clock and data recovery (CDR) unit. In many SerDes standards, CDR bandwidth requirement grows linearly with data rate. In a digital CDR (preferred choice because of portability and compactness), this requirement calls for increasing update rate of CDR logic. This becomes problematic when digital logic gates are not sufficiently fast.

Finally, parasitic capacitance at input/output pads, referred to as  $C_i$ , causes both bandwidth and channel discontinuity issues. In most cases, the value of  $C_i$  does not scale well with technology due to electro-static discharge (ESD) protection requirements. Excessive capacitance at the pad reduces bandwidth resulting in extra insertion loss and degradation of return loss ( $S_{11}$ ).

In addition to the above, short UI at high data rates makes the timing budget tight, leaving little room for clock dither, random jitter, supply-induced jitter, and duty cycle + quadrature phase errors. This paper presents a 40 Gbps CMOS transceiver and

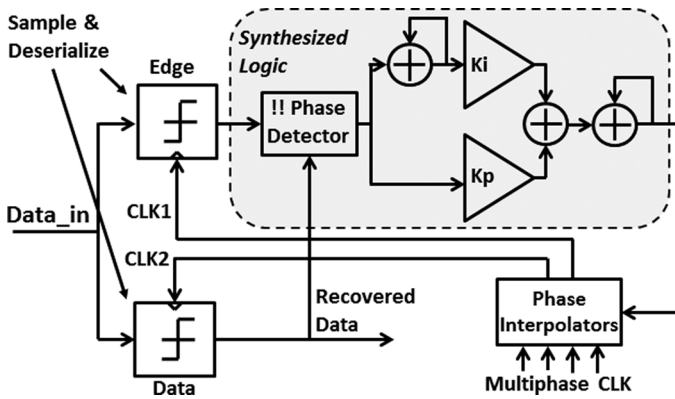


Fig. 3. Conventional second-order digital CDR.

discusses challenges and opportunities of such a design. Special attention is paid to the aforementioned issues. The organization of this paper is as follows. Section II presents an architectural overview. Sections III and IV discuss CDR architecture and channel equalization with a focus on RX which includes the majority of the equalization components. Sections V and VI present the details of circuit implementation and measurement results, respectively. Section VII provides concluding remarks.

## II. ARCHITECTURAL OVERVIEW

The proposed link is a duplex interface with four transceiver lanes. Fig. 1 shows the overall block diagram (also discussed in [8]). Each transmitter (TX) lane comprises a 64:2 serializer, a 2-tap feed-forward equalizer (FFE), and a 2:1 output multiplexer operating as a double data rate (DDR) driver. Each RX lane includes an analog front-end, a quad data rate (QDR) discrete-time equalizer with DFE and FFE components, a digital CDR and 4:64 deserializers. The serializer is 64:2 since the TX driver is DDR whereas the deserializer is 4:64 since RX front-end is QDR. ESD protection circuitry of the RX is embedded in an  $LC$  delay line to reduce  $C_i$  while providing tap points for a transversal filter. On-chip calibrated termination resistors are used on both TX and RX. Each transceiver lane also includes a built-in self-test (BIST) unit comprising PRBS generators and checkers of various lengths. A shared logic block controls the link and facilitates parallel loopback for test and characterization.

TX and RX have dedicated PLLs in this design. This helps to shorten the length of clock distribution networks and provides an additional degree of freedom for CDR design. On the TX side, a 20 GHz clock is distributed differentially and used by the DDR drivers. On the RX side, a lower-frequency 10 GHz clock is distributed differentially simplifying the RX clock network. As shown in Fig. 2, this globally distributed 10 GHz clock drives a local phase-splitting DLL in each lane which generates eight equally spaced phases at 10 GHz. These eight phases are then used by a group of phase interpolators, controlled by CDR logic, to generate RX clocks.

## III. CLOCK AND DATA RECOVERY

CDR bandwidth requirements of serial links are often linearly scaled with data rate. In a conventional bang-bang digital

CDR (Fig. 3) higher bandwidth can be achieved by increasing either the step size of phase interpolators or the update rate of CDR logic. Unfortunately both of these alternatives lead to challenging design issues. Increasing phase step size increases dithering jitter due to larger quantization noise while increasing update rate tightens timing constraints for any synthesized logic. Although advancements in technology are expected to provide faster devices and resolve this issue over time, the maximum operating frequency of logic blocks synthesized using a conventional synthesis flow does not seem to scale as fast as data rate making bandwidth requirements increasingly harder to satisfy. Frequency offset tracking requirements exacerbate this issue by necessitating continuous phase slewing which occupies some of the digital phase update slots. These slots can collide with phase updates leading to extra quantization noise when the digital phase interpolator is forced to perform a multi-step update. Another added difficulty is that the CDR is expected to satisfy these requirements while receiving a partially-equalized eye because it does not take full benefit of all of the link equalization components (e.g., unrolled DFE taps).

To address these issues, the proposed interface modifies the conventional digital CDR architecture. Fig. 4 shows the proposed CDR (also discussed in [9]) and highlights its three main features. First, the CDR and DFE paths are split from each other and each has its own samplers and clocks independently optimized for their respective functions. For the CDR path, the most important requirement is low decision latency. Here, jitter on the edge sampling clocks can be tolerated to some extent since it only indirectly affects link performance as long as the jitter is small compared to the transition distribution. For the DFE path, on the other hand, the data sampling clock must be very clean as its jitter should be small compared to the data eye opening. In order to satisfy both requirements, the CDR logic generates two sets of phase codes for CDR and DFE paths. The phase codes of the CDR clock are updated with the lowest latency and minimum averaging to achieve high tracking bandwidth, while those of the DFE clock are created with time-averaging to reduce dithering on data-sampling clocks. This helps to maintain good timing margin while still achieving a wide jitter tracking bandwidth with small phase step size and moderate CDR logic speed.

The second feature of the proposed CDR is a high-speed data filtering scheme. As shown in Fig. 4, this feature aims at filtering out edge information that is corrupted by residual intersymbol interference (ISI). The split-path design uses two data samplers in the CDR path with different threshold levels,  $+\alpha$  and  $-\alpha$ . Using these samples, the phase detector recognizes which transitions are severely affected by ISI (as shown in the inset of this figure) and eliminates them from the phase detection process. The choice of the decision level  $\alpha$  sets the tradeoff between recovered clock jitter and tracking bandwidth. If the threshold is set high, the phase detector sees a tight edge distribution resulting in lower recovered clock jitter but tracking bandwidth also drops since phase detector has fewer transitions to work with. The adaptive samplers described in Section IV-C can be used to adapt  $\alpha$ . Note that the elimination of edges with bad timing information helps to extract reliable phase information from a partially-equalized eye.

**Phase Detection Scheme:**

$D_{n-1}$	$E_{n-0.5}$	$D_n$	Phase Info
$< -\alpha$	$> 0$	$> +\alpha$	Early
$< -\alpha$	$< 0$	$> +\alpha$	Late
$> +\alpha$	$> 0$	$< -\alpha$	Late
$> +\alpha$	$< 0$	$< -\alpha$	Early
Otherwise			No Info

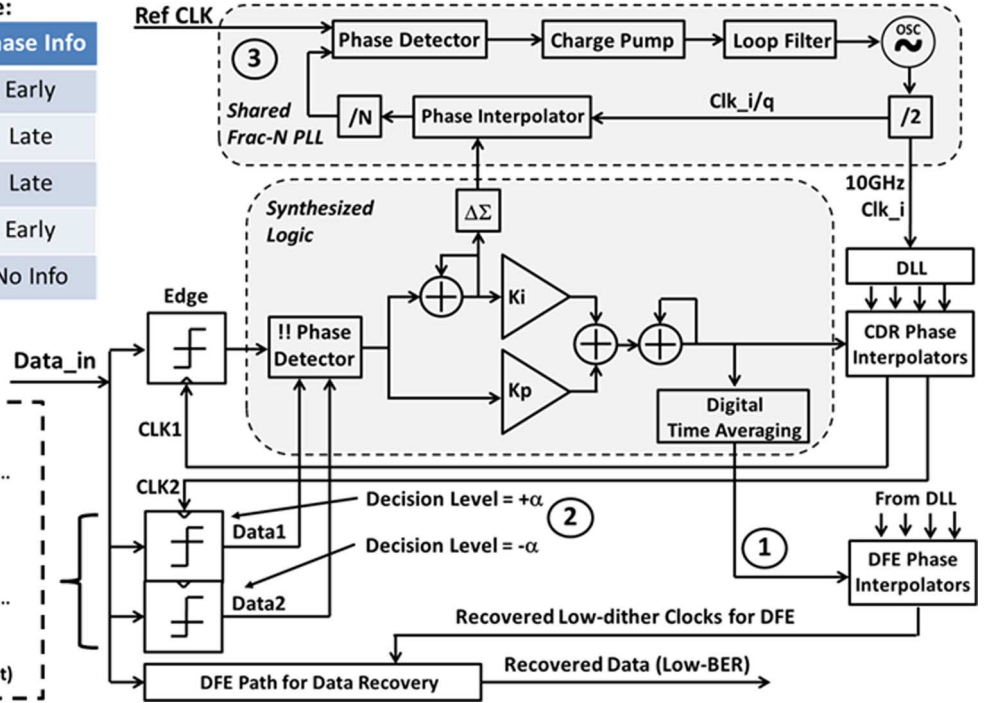



Fig. 4. Proposed CDR architecture with three main features highlighted: 1) Split-path design with additional filtering on DFE clock; 2) high-speed ISI filtering; and 3) fractional-N PLL.




Fig. 5. (a) Distributed ESD structure with transversal filter and (b) comparison of return loss.

Finally, the third feature of the CDR is a central fractional-N PLL that can achieve frequency offset tracking with minimum deterministic jitter. In this scheme, since the phase interpolator is inside the PLL loop, its quantization noise is filtered by the frequency response of the PLL resulting in lower dithering jitter [10]. A delta-sigma modulator is used in front of PLL's phase interpolator to shape its quantization noise and maximize filtering benefits [11].

#### IV. CHANNEL EQUALIZATION

Channel equalization is a key component of high-performance serial link interfaces operating in practical environments with channel loss. The most common configuration is to use TX FFE along with RX CTLE followed by DFE. However, a few tradeoffs exist when choosing a multicomponent equalization scheme to compensate a targeted loss profile. The proposed interface uses a 2-tap FFE on the TX side which is programmable for pre- or post-cursor ISI cancellation. On the RX side, a combination of CTLE, FFE and DFE is used.

##### A. FFE With Distributed ESD

Parasitic capacitance at RX input pad ( $C_i$ ) is often a significant source of high frequency loss and impedance discontinuity. This capacitance, of which ESD protection diode is a major component, can reach 700 fF. Along with the 50 Ohm termination resistance, this creates a pole around 5 GHz. This pole causes additional on-chip loss and creates an impedance discontinuity making  $S_{11}$  requirements hard to satisfy. A common remedy is to use a T-coil which can theoretically improve bandwidth by nearly a factor of two. Unfortunately, this is not sufficient at 40 Gbps as the pole frequency would still be significantly less than the Nyquist frequency.

An effective solution to this problem is to distribute the ESD device (and hence the parasitic capacitance,  $C_i$ ) and use on-chip inductors to emulate a transmission line [12]. As shown in Fig. 5(a), the main ESD diode is divided into a number of segments separated by series inductors  $L_{series}$ . Since ESD is a comparatively low frequency event, the impedances of these series inductors are negligible at frequencies relevant to an ESD event as long as their series resistance is well controlled. Therefore, with proper design, ESD performance can remain intact.

For a 3-segment distributed LC line, bandwidth and delay per stage are given by  $f_c = (3/2\pi) * (L_{series} * C_i)^{-0.5}$  and  $\Delta T = (L_{series} * C_i/3)^{0.5}$ , respectively. With  $C_i$  of 700 fF, choosing  $L_{series} = 400$  pH provides bandwidth extension to beyond 25 GHz and delay per stage of around 10 ps. This is a major improvement for  $S_{11}$ . Furthermore, since the delay is close to 0.5 UI at 40 Gb/s, one can use these delay elements to implement a transversal filter structure acting as a 2-tap FFE with a transfer function of  $y(t) = h_0x(t) + h_1x(t - \Delta T)$ . By

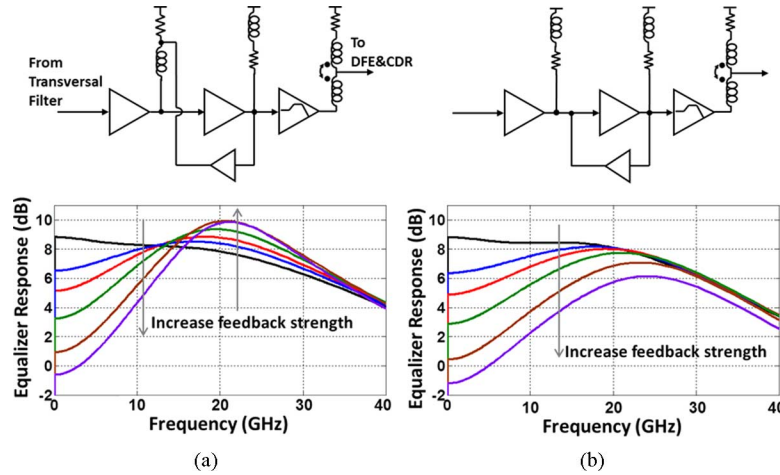


Fig. 6. (a) Proposed active-feedback CTLE compared with (b) the conventional architecture.

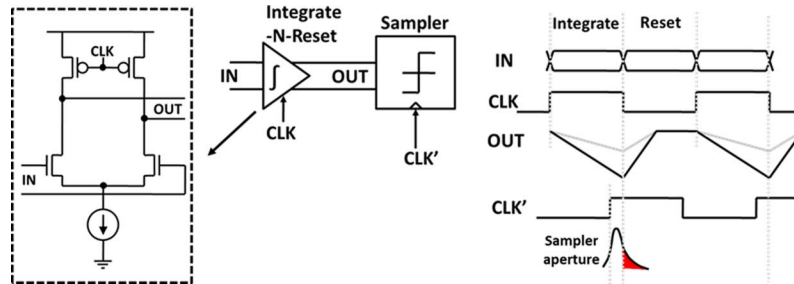


Fig. 7. Circuit and timing diagram of a current-integrating stage in half-rate operation.

adjusting the relative weight of  $h_0$  and  $h_1$ , this FFE can be used to cancel pre- or post-cursor ISI.

The link uses this concept to improve bandwidth and  $S_{11}$  and implement a 2-tap RX FFE. Fig. 5(b) compares the performance of the transversal filter to that of a T-coil along with experimental verification. Using the transversal filter,  $S_{11}$  of better than  $-10$  dB is achieved up to 35 GHz while a T-coil counterpart would extend only to 16 GHz. From a silicon area viewpoint, although the distributed ESD requires three inductors, these inductors are smaller in value ( $\sim 400$  pH) compared with the inductance required for a T-coil ( $\sim 750$  pH). Thus area overhead is moderate and the penalty is justified by improved bandwidth, better  $S_{11}$  and the additional RX FFE tap.

### B. Continuous-Time Linear Equalizer

The RX front-end uses a CTLE with  $+10$  dB of peaking at 20 GHz to help with channel loss equalization. High frequency peaking is achieved by incurring 3 dB of loss at DC and furnishing 7 dB of high-frequency gain. To achieve this target, a two-stage active feedback amplifier is used that employs source degeneration and inductive peaking. The final stage of the CTLE consists of a peaking buffer that uses an integrated T-coil to drive the parallel load of CDR and DFE paths.

Fig. 6(a) shows the CTLE design. Unlike conventional designs [Fig. 6(b)], the proposed solution connects the output of the active feedback amplifier to the far end of the peaking inductor. The advantage of this choice is twofold. First, by taking advantage of Miller effect, the effective inductance is increased

and larger peaking is achieved. Second, the parasitic capacitance of the feedback amplifier is isolated from the feed-forward path leading to bandwidth improvements. To quantify these effects, Fig. 6 compares simulated frequency response of the proposed topology with that of a conventional design. The proposed topology achieves up to 3 dB larger peaking at the same frequency at no additional cost to power or area.

### C. RX DFE and Discrete-Time Sampled FFE

The design of a DFE in a high-speed link is challenging due to its feedback timing constraint. Sampler latency and the delay of the summing node are two major components of feedback timing in a typical implementation. While a direct-feedback DFE design is already reported at 66 Gb/s in [2], sampler sensitivity is sacrificed in order to close timing. To alleviate this issue, a loop-unrolled architecture can be used to relax the latency requirement of the sampler and eliminate the summing node [13], [14]. However, such an implementation only favors low tap counts (usually less than 3) due to the tradeoff between tap number and power/hardware penalty. In order to implement a DFE with more than 10 taps, an analog summing stage is necessary in order to combine the incoming waveform with the weighted feedback signals.

An analog summation stage based on current integration can save substantial power compared to a conventional resistive-based CML summer [15], [16]. Here, integrating currents onto load capacitors with no DC resistance provides the desired result at the end of a clock phase with no exponential settling tail.

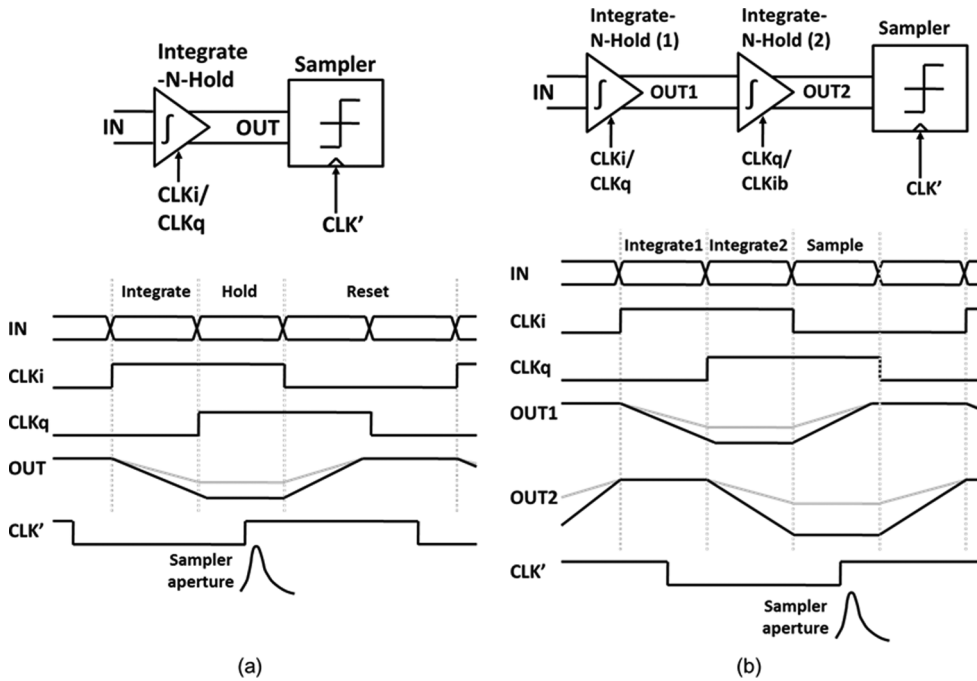


Fig. 8. Integrate-and-hold in quarter-rate operation. (a) Single integrate-and-hold. (b) Double integrate and hold.

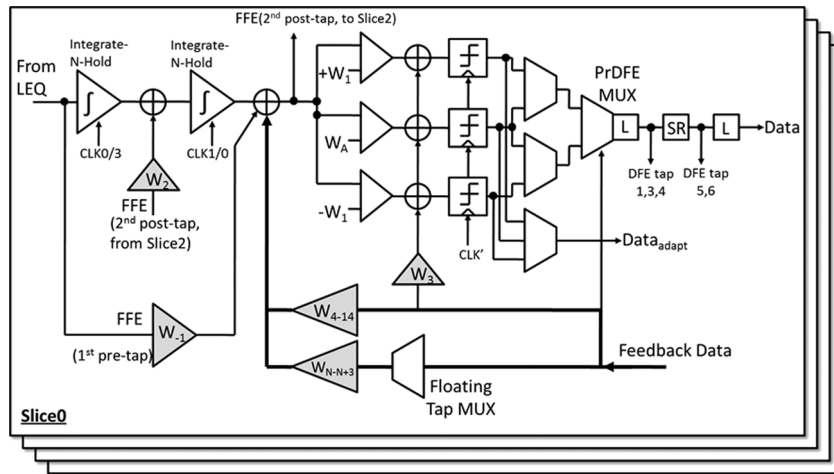


Fig. 9. Hybrid discrete-time equalizer with 3-tap sampled FFE and 17-tap DFE.

Fig. 7 shows this scheme in half-rate operation. In this case, the signal is integrated for 1 UI when the clock is low. The output is then reset when the clock goes high. For optimum operation, the sampling time of the following sampler should be chosen carefully in order to maximize the integration gain (as late as possible in the integration phase) while preventing from sampling incorrect information in the succeeding reset phase. In practice, finite sampling aperture creates design difficulty when the duration of the UI becomes too short.

A QDR architecture can be used to relax this timing difficulty. To realize an integrating summer in QDR operation, track-and-hold circuitry is usually required in order to retain the signal for longer than one UI when the quarter-rate clocks are used in the integrating stage [17]. Unfortunately, the implementation of an explicit 40-GS/s sample-and-hold circuitry is costly. To address

this difficulty, this work uses double integrate-and-hold stages similar to [18].

The operation of a QDR integrated-and-hold sampling scheme is shown in Fig. 8. Fig. 8(a) shows this operation for a single integrate-and-hold design. In this case, the integrating amplifier integrates the incoming signal over a period of 1 UI. Subsequently, the output of the amplifier is held for 1 UI to facilitate sampling with a finite sampling aperture. Finally, the integrator is reset and prepared for the next cycle during the remaining two UIs of the QDR cycle. Fig. 8(b) shows the double integrating-and-hold operation in QDR operation where integrate-and-hold stages are cascaded. The second stage operates similar to the first stage (described above) except that its clock inputs are delayed by 1 UI. This provides extra gain while the timing requirement of the sampler is still the same as

TABLE I  
SUMMARY OF EQUALIZATION COMPONENTS

Location	Type	Number of Taps	Programmable	Notes
Transmitter	FFE	2	Pre/Post	Set to pre
Receiver	Transversal Filter	2	Pre/Post	Set to post
Receiver	CTLE	3 Stages	-	-
Receiver	FFE	3	-	1 <sup>st</sup> pre + 2 <sup>nd</sup> post
Receiver	Fixed DFE	13	-	1 <sup>st</sup> post + 3 <sup>rd</sup> through 14 <sup>th</sup> post
Receiver	Floating DFE	4	15 <sup>th</sup> to 30 <sup>th</sup> tap	-

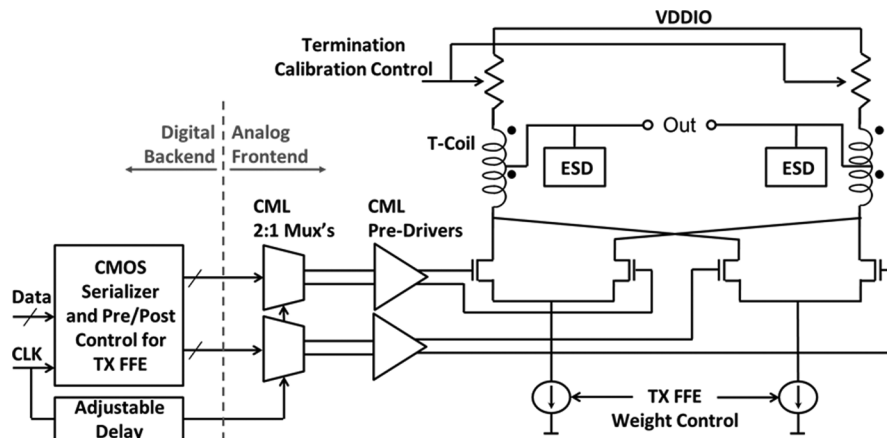


Fig. 10. Design details of the current mode transmitter.




Fig. 11. Circuit schematic of the CTLE.

before. This link uses double integrate-and-hold stages to take advantage of the extra gain. The power overhead of this extra stage is around 2 mW per lane and area overhead is moderate since no inductor peaking is used.

An extra advantage of the double integrate-and-hold scheme is that the integrating stages provide delayed versions of the incoming signal that can be used to form FFE taps in RX. Fig. 9 shows the overall architecture of such hybrid discrete-time equalizer combining DFE and FFE. Although the use of a DFE loop-unrolling technique relaxes timing

requirements for 1st post-tap, the feedback design for the 2nd post-tap is still very challenging. Instead of using a 2-tap loop-unrolling DFE, we take advantage of the integrate-and-hold stages to provide a 2-UI sampled-delay for FFE operation. As shown in Fig. 9, the signal is delayed by 2 UI and fed to the output of the first integrate-and-hold stage in the other interleaving path to perform a second post tap cancellation. Using the same technique, the input signal is also fed forwarded to the output of the second integrate-and-hold stage for 1st pre-cursor cancellation.

In addition to FFE and 1st tap loop-unrolled DFE, a direct-feedback loop is implemented to cancel third to 14th post-cursor ISIs. In order to relax timing requirement of the third tap, it is applied at the kick-back isolation buffer that precedes the sampler. For the rest of the taps, summing happens at the output of the integrate-and-hold stages. An auxiliary sampler is also placed in parallel with loop-unrolled data sampler in each interleaving path. It can be used as the error sampler for DFE adaptation, eye height/width monitoring, and sampler swapping for live offset correction [7]. In addition to the fixed-position DFE/FFE taps which cover up to 14th post cursor ISI, floating taps are implemented to cover any consecutive four taps ranging from 15th to 30th post tap position to correct possible long tail or channel reflection effects. A minimum-BER adaptation engine [19] adapts DFE tap values for optimum performance.

Table I shows a summary of equalization components. The majority of the equalization is done on the RX side using DFE. This choice leads to superior signal to noise ratio (SNR) because, in contrast to linear equalizers (FFE and CTLE), DFE does not amplify high frequency noise along with signal content. This advantage is significant at high data rates where noise is integrated over a wide bandwidth. The replacement of the 2nd tap DFE with RX FFE reduces this advantage to some extent but saves hardware, and breaks an undesirable design cycle where adding more loop-unrolled taps would have led to more latency for the direct-feedback DFE taps potentially necessitating even more loop unrolling. Skipping a DFE tap helps to break this cycle.

## V. CIRCUIT IMPLEMENTATIONS

The interface is implemented in a 28 nm CMOS process. Fig. 10 shows the schematic of the TX data path which is partitioned into a CMOS back-end and a CML front-end. The back-end circuits include a 64:2 serializer, and an FIR filter that creates two 2-bit data streams for the main and pre/post taps. The CML front-end consists of the final 2:1 stage, a predriver, and a driver. The final multiplexer stage uses inductive peaking. Due to the area constraints, the predriver does not use inductors, which resulted in it becoming the bandwidth bottleneck. Output driver uses a T-coil which doubles its bandwidth. In order to achieve large output swing, the driver is supplied from an elevated power supply (1.1 V vs. 0.85 V), while the stacked device structure ensures that the thin-oxide devices do not experience stress in normal operation. The data handoff from the back-end to the front-end has the tightest timing in the TX data path. The adjustable delay in Fig. 10 is inserted to meet this timing. The delay is adjusted during initialization to maximize the timing margin at this critical clock domain crossing.

Receiver CTLE schematic is shown in Fig. 11. Conventional source-degenerated stages with inductive peaking are used in this design. Similar stages are used for transversal filter amplifiers ( $h_0$  and  $h_1$  in Fig. 5). We use the tail current value to control the gain of those stages, hence adjusting the coefficients of the transversal filter FFE.

In this 28-nm CMOS process, the sampling aperture of the highest performance sampler design proved to be too wide for a

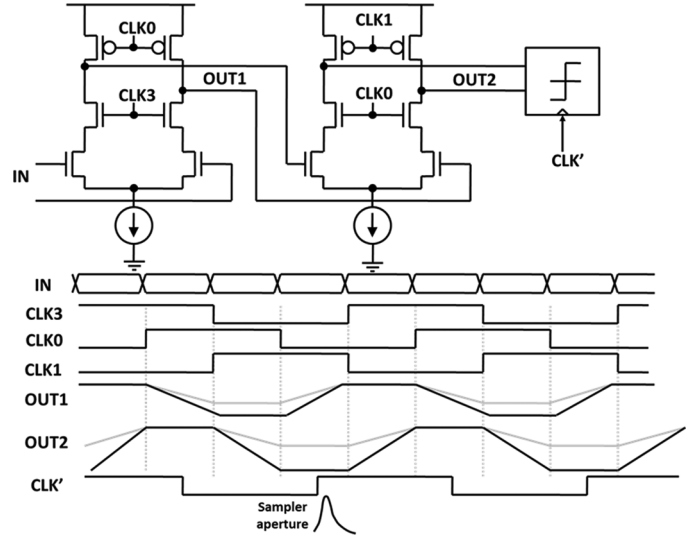


Fig. 12. Circuit and timing diagram of double integrate-and-hold stages in QDR operation.

UI of 25 ps. To alleviate the issue, QDR sampling is used. Unfortunately, while relaxing timing requirements, this approach increases the number of samplers, and hence the load on RX front-end. After accounting for split-path CDR/DFE, loop-unrolled DFE and, extra samplers for DFE adaptation and CDR high-speed data filtering, RX front-end needs to drive 24 samplers. The T-coil-based peaking buffer of the CTLE output is designed to drive this load.

Fig. 12 shows the circuit implementation and timing diagram of the double integrate-and-hold QDR DFE. Two quarter-rate clocks, CLK0 and CLK3, are used in the first stage. When CLK0 and CLK3 are both high, the signal is integrated at the output node for 1 UI. The output is then held at a constant value when CLK3 goes low (cascoded NMOSs turn off). In the final phase, when CLK0 is low, the output is reset to the supply. Thus each cycle of the QDR operation is divided into three phases: 1) integrate for 1 UI; 2) hold for another 1 UI; and 3) reset in the remaining 2 UIs. The succeeding integrate-and-hold stage operates in the same manner. In order to cover a wide operating frequency with a well-defined common mode voltage, a replica integrating summer path is used to set the bias current through the digital control logic [16].

Identical central PLLs (block diagram shown in Fig. 4) are used for TX and RX. Transmitter PLL uses a constant code for its phase interpolator. Each PLL includes an LC VCO which is biased using a local constant-gm bias generator through the center tap of its inductor (Fig. 13). Since inductor center tap voltage rises only to the common-mode voltage level of the output, this choice makes efficient use of voltage headroom for supply isolation. The constant-gm bias generator uses a calibrated resistor ( $R$  in Fig. 13) that tracks on-chip termination resistors.

The samplers in the CDR path use a simpler design than the DFE samplers with no integrating summers. This is because lower SNR is tolerable in this path. Therefore, although more samplers are used in the proposed architecture, the total power



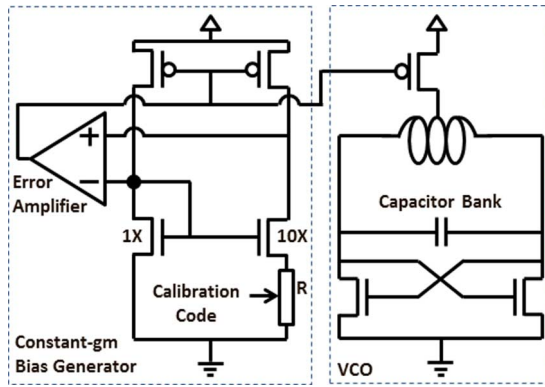


Fig. 13. Circuit schematic of the VCO and corresponding constant-gm bias generator.




Fig. 14. Chip micrograph of the quad SerDes.

consumption is comparable to a conventional design which replicates the data path (and integrating stages) for edge sampling. There is still a moderate power overhead due to extra phase interpolators.

Fig. 14 shows the chip micrograph of the quad interface occupying  $0.81 \text{ mm}^2$  per lane (amortizing common blocks). The extensive use of inductors is clearly visible. In addition to area overhead, these inductors create two extra challenges. First, they create ‘holes’ in the power grid which can impact power integrity. Second, larger silicon area translates to longer clock/data networks and thus higher power consumption. Nevertheless, the inductors are integral parts of the design at this speed. The challenges they raise were addressed through early stage macro-model simulation of the power grid and floor plan optimization to minimize incurred penalties.

## VI. EXPERIMENTAL RESULTS

The performance of the interface is tested using a Megtron 6 characterization board. Fig. 15 shows the measurement setup. It includes a device-under-test (DUT) channel and several chip-to-chip (C2C) channels of different lengths and DC-/AC-coupled configurations. The DUT channel is a benign channel (about 5 db of loss at 20 GHz excluding on-chip loss) for jitter and TX characterization. The C2C channels are for characterizing TX-RX communication. Link TX-RX characterization results were measured using the longest C2C channel which is 12 inches long and AC-coupled presenting 20 dB of loss at 20 GHz as can be seen in Fig. 16.

Fig. 17 shows the clock characterization results using the DUT channel when the TX is configured to send a 20 GHz




Fig. 15. Characterization setup.




Fig. 16. Experimental characterization result of the chip to chip channel.

periodic pattern. Out of band spot phase noise is  $-128 \text{ dBc/Hz}$  at 100 MHz and integrated jitter from 10 MHz to 10 GHz is 162 fs. The lower limit of integration is set by the CDR bandwidth. Using the same channel, the TX performance is characterized with PRBS data. As shown in Fig. 18, the TX eye diagram is completely closed when equalization is off (upper left) and partially open when TX FFE is on (upper right). The graphs at the bottom show a 4.7 dB reduction in height due to de-emphasis equalization. In this measurement, TX FFE is set to post-cursor ISI cancellation since it is larger than pre-cursor ISI for this channel.

Fig. 19 shows jitter tolerance of the CDR with and without the flexibilities provided by the split path CDR/DFE. Improvement of jitter tolerance by almost a factor of two at 30 MHz is due to the fact that the proposed CDR can achieve large bandwidth without causing undesirable peaking in CDR frequency response. In the inset of this figure, the effective eyes at the input of the CDR and DFE paths are also shown which confirms the ability of the proposed CDR to recover a stable clock from a partially-equalized eye with small voltage and timing margins.

The ESD protection performance of the transversal filter is characterized using the human body model that confirms sufficient protection up to 1.9 kV. Charged device model and machine model test results are not available unfortunately. Fig. 20 shows the measured RX eye for a 40 Gbps PRBS pattern sent by the proposed TX via the previously described

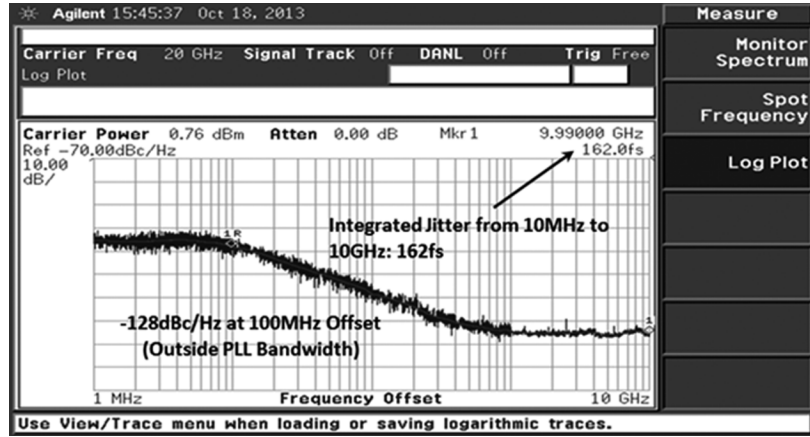


Fig. 17. Jitter characterization showing spot phase noise and integrated jitter. Lower limit of jitter integration (10 MHz) is set by CDR bandwidth.

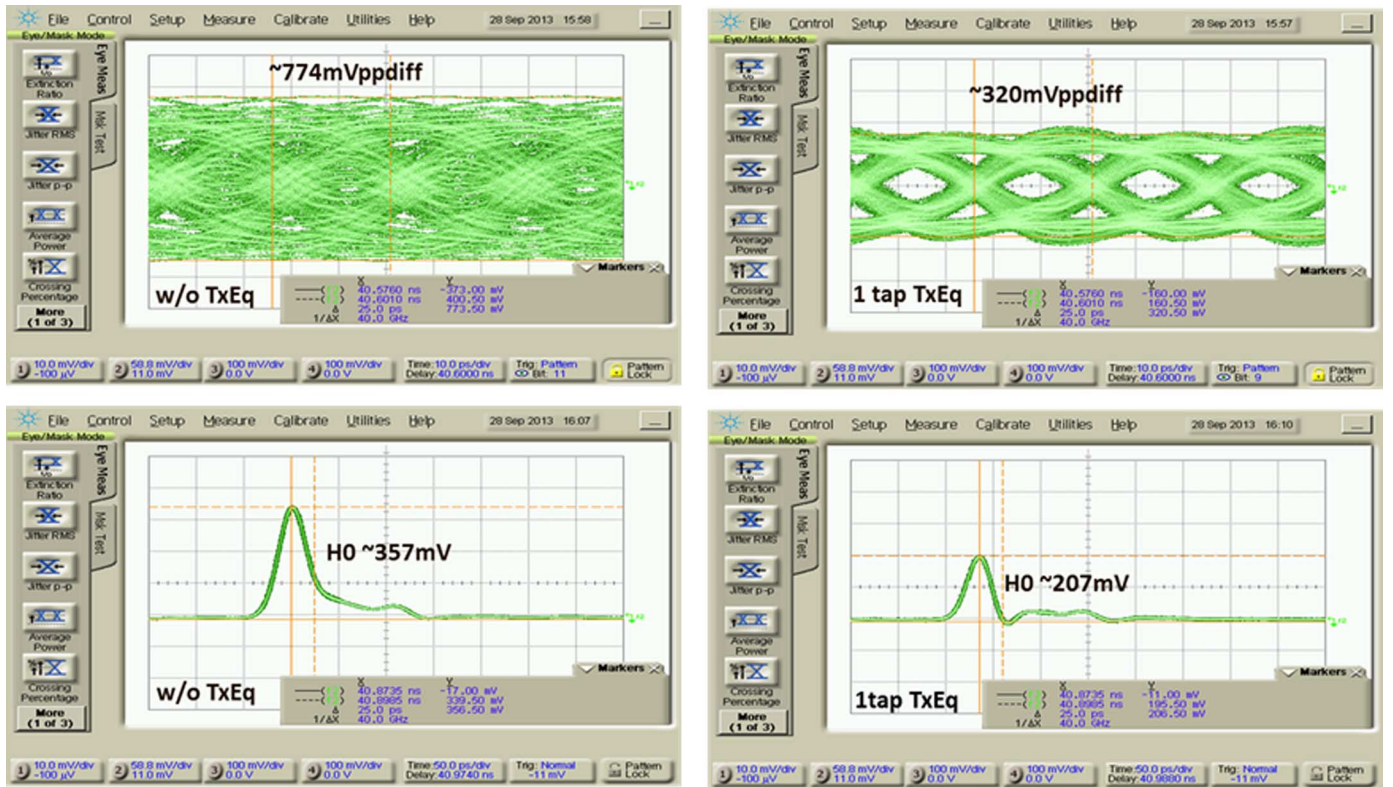


Fig. 18. 40 Gbps TX eye and single-bit response over the DUT channel.

C2C channel. An on-chip eye monitor captures internal RX eye to observe the effect of on-chip loss and RX equalization. On-chip loss due to  $C_i$  (TX + RX) is estimated to be 8 dB at 20 GHz bringing total loss to 28 dB. Only the first 10 DFE taps are activated while the rest (including floating taps) are disabled since they are not necessary for this relatively-smooth, chip-to-chip channel. At  $\text{BER} = 10^{-9}$ , horizontal and vertical eye openings are 0.27 UI and 120 mV, respectively. When extrapolated to  $\text{BER} = 10^{-15}$ , eye openings are 0.15 UI and 45 mV. This result is achieved without using any error correction schemes which could further improve these margins or be traded for extra channel loss. Table II shows performance summary and comparison with prior art. Compared to [4], [5], and [20], the link achieves 40% higher data rate occupying

similar area. Compared to [21] it operates at the same speed in nearly half of the area (probably due to RX FFE inductors in [21]). The link achieves power efficiency of 23.2 mW/Gb/s, which is comparable with the previously published work with comparable equalization at lower speeds [4], [20]. A power consumption breakdown graph is presented in Fig. 21 which shows that the majority of this power is used by the RX due to extensive RX equalization.

## VII. CONCLUSION

Architectural innovation and careful implementation must go hand in hand when designing a high-speed serial link for maximum performance in low-cost CMOS processes. We present a fully-integrated quad serial link transceiver in a

TABLE II  
PERFORMANCE SUMMARY AND COMPARISON

	[4] JSSC 2012	[5] ISSCC 2012	[20] JSSC 2014	[21] JSSC 2012	<b>This Work</b>
<b>Technology</b>	32 nm SOI CMOS	40 nm CMOS	28 nm CMOS	65 nm CMOS	<b>28 nm CMOS</b>
<b>Data Rate</b>	28 Gb/s	28 Gb/s	28 Gb/s	40 Gb/s	<b>40 Gb/s</b>
<b>TX Equalization</b>	3-tap FFE	3-tap FFE	3-tap FFE	5-tap FFE	<b>2-tap FFE</b>
<b>RX Equalization</b>	CTLE and 15-tap DFE	CTLE	CTLE and 14-tap DFE	3-tap FFE	<b>2-tap Transversal Filter, CTLE, 3-tap Sampled FFE, and 17-tap DFE</b>
<b>Area/Lane</b>	0.81 mm <sup>2</sup>	0.9 mm <sup>2</sup>	0.83 mm <sup>2</sup>	1.8 mm <sup>2</sup>	<b>0.81 mm<sup>2</sup></b>
<b>Random Jitter (rms)</b>	250 fs	150 fs *	250 fs	319 fs	<b>170 fs**</b>
<b>Channel Loss</b>	35 dB **	13 dB	34 dB	19 dB	<b>20 dB **+</b>
<b>Power/Lane</b>	693 mW	225 mW	560 mW	655 mW	<b>927 mW</b>
<b>Energy Efficiency</b>	24.8 mW/Gb/s	8 mW/Gb/s	20 mW/Gb/s	16.4 mW/Gb/s	<b>23.2 mW/Gb/s</b>

\* Integrated up to 1GHz

\*\* Integrated from 10MHz (CDR bandwidth) to 10GHz

\*+ Includes package

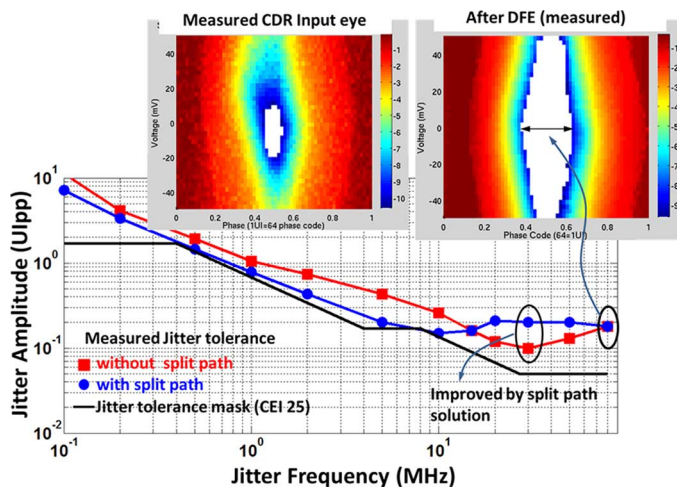


Fig. 19. Characterization results of CDR jitter tolerance.




Fig. 20. Two-dimensional internal eye diagram at RX sampler showing link performance at 40 Gbps using a PRBS31 pattern.

28 nm CMOS process operating at 40 Gb/s per lane using NRZ signaling. With ratio of UI to ‘fan-out of 4’ delay close to one, conventional design choices are often inadequate due to speed limitations. To improve RX bandwidth and return loss, we use a distributed ESD device forming a transversal

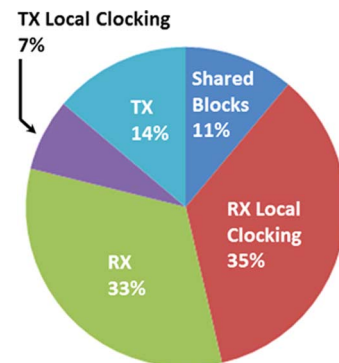


Fig. 21. Power breakdown of the link.

filter and a 2-tap RX FFE. To relax timing requirements for data samplers, a QDR double-integrating hybrid discrete-time equalizer (with FFE and DFE components) is implemented. To improve CDR bandwidth, a split-path CDR/DFE architecture is introduced with ISI edge elimination. Inductors are used extensively to improve bandwidth. While instrumental to the design, some of these choices result in unfavorable power consumption and silicon area. The transceiver demonstrates a 40 Gb/s error free operation ( $BER < 10^{-15}$ ) over a channel with 20 dB loss while achieving 23.2 mW/Gb/s power efficiency and occupying 0.81 mm<sup>2</sup> area per lane. As process technologies advance, certain process-related challenges such as bandwidth, DFE loop delay and sampler aperture may improve helping to reduce power and area. Other factors, such as return loss and CDR bandwidth, may not improve with process scaling as their underlying causes (e.g.,  $C_i$  and the speed of synthesized logic) do not scale as rapidly as data rate. Architectural innovations and/or revisions of system requirements (e.g., ESD protection) may be necessary in future generations to facilitate the design of lower-power and more compact NRZ transceivers at required data rates.

## REFERENCES

- [1] P.-C. Chiang, H.-W. Hung, H.-Y. Chu, G.-S. Chen, and J. Lee, "60 Gb/s NRZ and PAM4 transmitters for 400 GbE in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2014, pp. 42–43.
- [2] Y. Lu and E. Alon, "Design techniques for a 66 Gb/s 46 mW 3-tap decision feedback equalizer in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3243–3257, Dec. 2013.
- [3] J.-K. Kim, J. Kim, G. Kim, and D.-K. Jeong, "A fully integrated 0.13- $\mu$ m CMOS 40-Gb/s serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 44, no. 5, pp. 1510–1521, May 2009.
- [4] J. F. Bulzacchelli *et al.*, "A 28-Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 3232–3248, Dec. 2012.
- [5] M. Harwood *et al.*, "A 225 mW 28 Gb/s SerDes in 40 nm CMOS with 13 dB of analog equalization for 100 GBASE-LR4 and optical transport lane 4.4 applications," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 326–327.
- [6] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 6, pp. 711–717, Jun. 1991.
- [7] B. S. Leibowitz *et al.*, "A 7.5 Gb/s 10-tap DFE receiver with first tap partial response, spectrally gated adaptation, and 2nd-order data-filtered CDR," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 228–229.
- [8] E.-H. Chen *et al.*, "A 40-Gb/s serial link transceiver in 28-nm CMOS technology," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2014.
- [9] M. Hossain *et al.*, "A 4  $\times$  40 Gb/s quad-lane CDR with shared frequency tracking and data dependent jitter filtering," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2014.
- [10] P. Larsson, "A 2–1600-MHz CMOS clock recovery PLL with low-V<sub>dd</sub> Capability," *IEEE J. Solid-State Circuits*, vol. 34, no. 12, pp. 1951–1960, Dec. 1999.
- [11] P. K. Hanumolu, G.-Y. Wei, and U.-K. Moon, "A wide-tracking range clock and data recovery circuit," *IEEE J. Solid-State Circuits*, vol. 43, no. 2, pp. 425–439, Feb. 2008.
- [12] C. Ito, K. Banerjee, and R. W. Dutton, "Analysis and design of distributed ESD protection circuits for high-speed mixed-signal and RF ICs," *IEEE Trans. Electron Device*, vol. 49, no. 8, pp. 1444–1454, Aug. 2002.
- [13] J. E. Proesel and T. O. Dickson, "A 20-Gb/s, 0.66-pJ/bit serial receiver with 2-stage continuous-time linear equalizer and 1-tap decision feedback equalizer in 45 nm SOI CMOS," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2011, pp. 206–207.
- [14] K. Jung, A. Amirkhany, and K. Kaviani, "A 0.94 mW/Gb/s 22 Gb/s 2-tap partial-response DFE receiver in 40 nm LP CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2013, pp. 42–43.
- [15] J. F. Bulzacchelli, T. O. Dickson, Z. T. Deniz, H. A. Ainspan, B. D. Parker, M. P. Beakes, S. V. Rylov, and D. J. Friedman, "A 78 mW 11.1 Gb/s 5-tap DFE receiver with digitally calibrated current-integrating summers in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 368–369.
- [16] T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 12-Gb/s 11-mW half-rate sampled 5-tap decision feedback equalizer with current-integrating summers in 45-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1298–1305, Apr. 2009.
- [17] T. Toifl, M. Ruegg, R. Inti, C. Menolfi, M. Brändli, M. Kossel, P. Buchmann, P. A. Francese, and T. Morf, "A 3.1 mW/Gbps 30 Gbps quarter-rate triple-speculation 15-tap SC-DFE RX data path in 32 nm CMOS," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2012, pp. 102–103.
- [18] H.-J. Chi, J.-S. Lee, S.-H. Jeon, S.-J. Bae, Y.-S. Sohn, J.-Y. Sim, and H.-J. Park, "A single-loop SS-LMS algorithm with single-ended integrating DFE receiver for multi-drop DRAM interface," *IEEE J. Solid-State Circuits*, vol. 46, no. 9, pp. 2053–2063, Sep. 2011.
- [19] E.-H. Chen *et al.*, "Near-optimal equalizer and timing adaptation for I/O links using a BER-based metric," *IEEE J. Solid-State Circuits*, vol. 43, no. 9, pp. 2144–2156, Sep. 2008.
- [20] H. Kimura *et al.*, "A 28 Gb/s 560 mW multi-standard SerDes with single-stage analog front-end and 14-tap decision feedback equalizer in 28 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 1–13, 2014.
- [21] M.-S. Chen *et al.*, "A fully-integrated 40 Gb/s transceiver in 65 nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 47, no. 3, pp. 627–640, 2012.



**Reza Navid** (S'00–M'05–SM'11) received the B.Sc. degree from the University of Tehran, Tehran, Iran, in 1996, the M.Sc. degree from Sharif University of Technology, Tehran, Iran, in 1998, and the Ph.D. degree from Stanford University, Stanford, CA, USA, in 2005, all in electrical engineering.

His interests include technology development for wireline communication systems and associated mixed-signal circuits. In 2005 he joined Rambus Inc., Sunnyvale, CA, USA, working on high-speed integrated circuits for chip-to-chip and memory link applications. From 2008 to 2010, he was with True Circuits Inc., Los Altos, CA, USA, where he focused on low-phase-noise timing circuits for communication and data conversion systems. In 2010, he returned to Rambus Inc., where he is now a Senior Principal of Engineering. Since his return, he has served in various capacities including product definition and architecture development for high-speed SerDes interfaces. He has also served as a design lead focusing on technology development for power-efficient wireline and memory links. His current interests include architecture development for high-performance and low-power serial links targeting chip-to-chip and backplane applications.

Dr. Navid serves on the technical program committee of VLSI Circuits Symposium.



**E-Hung Chen** (S'05–M'12) was born in Taipei, Taiwan. He received the B.S. degree from National Taiwan University, Taipei, Taiwan, in 2002, and the M.S. and Ph.D. degrees from University of California at Los Angeles (UCLA), in 2008 and 2011, respectively, all in electrical engineering.

From 2011 to 2014, he was with Rambus Inc., Sunnyvale, CA, USA, where he worked on mixed-signal circuit design and equalization for high-speed signaling. In 2014, he joined MediaTek Inc., Hsinchu, Taiwan, working on high-speed serial

transceivers for data center networking.



**Masum Hossain** received the B.Sc. degree from the Bangladesh University of Engineering and Technology, in 2002, the M.Sc. degree from Queen's University, in 2005, and the Ph.D. degree from the University of Toronto, Toronto, ON, Canada, in 2010.

He joined the faculty of the University of Alberta Department of Electrical and Computer Engineering in winter 2013. Before returning to academia, he has spent several years in industrial research. From 2008 to 2010, he was with Gennum Corporation in the

Analog and Mixed Signal division where he focused on the development of world's highest capacity and most power efficient cross point router solution. Following that, he joined Rambus Lab as a Senior Member of Technical Staff, where he focused on advanced equalization and clock recovery techniques for high-speed interfaces.

Dr. Hossain won the best student paper award at the 2008 IEEE Custom Integrated Circuits (CICC) Conference. He also won *Analog Device's* outstanding student designer award in 2010.




**Brian Leibowitz** (S'97–M'05) received the B.S. degree from Columbia University, New York, NY, USA, in 1998, and the Ph.D. degree from the University of California, Berkeley, CA, USA, in 2004, both in electrical engineering.

In his post-graduate work, he developed a fully integrated CMOS imaging receiver for free-space optical communication. His graduate studies at Berkeley were supported by a fellowship from the Fannie and John Hertz Foundation. He is currently an Analog IP Designer with Apple Inc., Cupertino,

CA, USA, was previously with Rambus Inc., where he worked on mixed-signal architecture and circuit design for a variety of high-speed serial links and memory interfaces.

Dr. Leibowitz was the recipient of the 1998 Edwin H. Armstrong Award at Columbia University.



**Jihong Ren** received the Ph.D. degree in computer science from the University of British Columbia, Vancouver, BC, Canada, in 2006, where she worked on optimal equalization for chip-to-chip high-speed buses.

She is currently a Senior Manager with Altera, San Jose, CA, USA, managing the SerDes system architecture group. Prior to Altera, she was with Rambus Inc. Her research interests include equalization, timing-recovery, adaptation algorithms, signal integrity, and power integrity.



**Chuen-huei Adam Chou** received the B.S. degree in electrical engineering from National Tsing-Hua University, Hsin-Chu, Taiwan, in 1991, and the M.S. degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 1994.

From 1995 to 1998, he was with Toshiba America Electronic Component, where he worked on USB and HSTL transceiver macro. From 1998 to 2013, he was with Rambus, Inc., where he worked on high-speed chip-to-chip and memory interfaces. He is now with Xilinx, Inc., San Jose, CA, USA. His current interests

include analog mixed-signal circuit design, high-speed serial links, memory interfaces and analog-digital co-simulation.



**Barry Daly** (M'98) received the B.Eng. degree in electronic engineering from the Cork Institute of Technology, Cork, Ireland, in 1996.

He is a Senior Design Engineer with Rambus Inc., Chapel Hill, NC, USA, working on high-speed serial link and memory bus design. Since joining Rambus in 2002, he has worked on mixed-signal circuit design for high speed chip-to-chip communications and related clocking circuits




**Marko Aleksić** (S'01–M'07) received the Dipl.Ing. degree from the University of Belgrade, Belgrade, Serbia, in 2000, and the M.S. and Ph.D. degrees from the University of California, Davis, CA, USA, in 2004 and 2006, respectively, all in electrical engineering.

In 2000 and 2001, he was a Visiting Researcher with the Advanced Computer Systems Engineering Lab at UC Davis, where he worked on high-performance low-power clocked storage elements, and circuit optimization. In 2002, he joined Solid-State Circuits Research Lab at UC Davis, where he worked on the design and noise analysis of mixed-signal circuits. He is currently with Rambus, Inc., Sunnyvale, CA, USA, where he works on high-speed memory interfaces.



**Bruce (Hsuan-Jung) Su** (S'06–M'13) was born in Taitung, Taiwan. He received the B.S. degree from National Tsing Hua University, Hsinchu, Taiwan, in 2001, and the M.S. and Ph.D. degrees from North Carolina State University, USA, in 2006 and 2012, respectively.

In 2010, he joined Rambus Inc., Chapel Hill, NC, USA, where he is currently engaged in low-power high-performance memory interface design. His interests include CMOS analog circuit design, PLL, and package design for high-speed links.



**Simon Li** received the B.S. degree in electrical engineering from the University of California, Berkeley, CA, USA, in 1992.

From 1992 to 1996, he was with HAL Computer Systems, Campbell, CA, USA, working on the SPARC CPU design. From 1996 to 1997, he was a Logic Designer with the MIPS division of SGI Corporation, Mountain View, CA, USA, working on high performance MIP microprocessors. In 1997, he joined Rambus Inc., Los Altos, CA, USA, where he has been working on high-speed memory interface

and serial link design.



**Makarand Shirasgaonkar** (M'05) received the M.Tech. degree in electrical engineering science from the National Institute of Technology Nagpur, India, in 2005.

Currently, he is a Principal Engineer with Rambus Inc., Bangalore, India, where he has worked since 2007 on PLL/DLL/clocking circuits & LDO design for high speed serial links. Prior to Rambus, he was with Qualcomm Logic working on PLL circuit design from 2005 to 2007.



**Fred Heaton** received the B.S. degree in electrical engineering from Brown University, Providence, RI, USA, in 1980, and the M.S. degree in computer engineering from Carnegie-Mellon University, Pittsburgh, PA, USA, in 1982.

Over the course of his career, he has been involved in circuit and chip design spanning a wide range of areas including RISC microprocessors (Bell Laboratories), SIMD devices (MCNC), experimental high-performance graphics systems (Division Inc & Hewlett Packard), as well as high-speed network encryptors (MCNC & Secant Communications). In 2000, Fred joined Velio Communications developing gigabit signaling systems. Since 2004, he has been with Rambus Inc., Sunnyvale, CA, USA, involved in the development of experimental and high-end SERDES products.



**Jared Zerbe** was born in New York, NY, USA, in 1965. He received the B.S. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1987.

From 1987 to 1992 he worked at VLSI Technology and MIPS Computer Systems, where he designed high-performance CPU floating-point blocks. In 1992, he joined Rambus, Inc., Mountain View, CA, USA, where for over 20 years he specialized in the design of high-speed I/O, PLL/DLL clock-recovery, and data-synchronization circuits. He became a

Rambus distinguished inventor with over 100 issued US patents. He has taught courses at both Berkeley and Stanford in high-speed I/O design and authored or coauthored over 40 IEEE conference and journal papers including forums & tutorials at ISSCC and VLSI and 1994 ISSCC best paper. In February 2013 he joined Apple Inc., Cupertino, CA, USA, where he is currently an AMS Architect.

Mr. Zerbe served on the program committee for DesignCon and VLSI Circuits Symposium from 2010–2013 and was an associate editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS until 2014.



**John Eble** (S'93–M'99) was born in Metairie, LA, USA, in 1971. He received the B.Cmp.E., M.S.E.E, and Ph.D. degrees in electrical engineering from Georgia Institute of Technology, Atlanta, GA, USA, in 1993, 1994, and 1998, respectively.

From 1998 to 2001 he worked on the EV7 high-speed I/O circuits in the Alpha Microprocessor Development Group, Compaq Computer Corporation, Shrewsbury, MA. From 2001 to 2003, he was with Velio Communications as a circuit designer. In 2003 he joined Rambus Inc. where he has since specialized

in the design of high-speed SERDES cells and next-generation memory signaling and clocking architectures. He has authored or coauthored over 30 technical publications and over five patents and has contributed a book chapter on off-chip signaling. At Georgia Tech, he developed the original version of the Generic System Simulator (GENESYS) and received the Best Paper Award at the 1997 International ASIC Conference. He is currently the Director of US Design where his teams are focused on advanced development of high-performance & low-power memory and serial link interface technologies.