

# A 16 Gb/s/Link, 64 GB/s Bidirectional Asymmetric Memory Interface

Haechang Lee, *Member, IEEE*, Kun-Yung Ken Chang, *Member, IEEE*, Jung-Hoon Chun, *Member, IEEE*, Ting Wu, *Member, IEEE*, Yohan Frans, *Member, IEEE*, Brian Leibowitz, *Member, IEEE*, Nhat Nguyen, T. J. Chin, Kambiz Kaviani, *Member, IEEE*, Jie Shen, Xudong Shi, Wendemagegnehu T. Beyene, *Senior Member, IEEE*, Simon Li, Reza Navid, *Member, IEEE*, Marko Aleksic, *Member, IEEE*, Fred S. Lee, *Member, IEEE*, Fredy Quan, Jared Zerbe, *Member, IEEE*, Rich Perego, and Fariborz Assaderaghi, *Senior Member, IEEE*

**Abstract**—This paper describes a bidirectional, differential, 16 Gb/s per link memory interface that includes a Controller and an emulated DRAM physical interface (PHY) designed in 65 nm CMOS. To achieve high data rate, the interface employs the following technology ingredients: asymmetric equalization, asymmetric timing calibration, asymmetric link margining, inductor based (LC) PLLs, multi-phase error correction, and a data dependent regulator. At 16 Gb/s, this interface achieves a unit-interval to inverter FO4 ratio of 2.8 (Controller) and 1.4 (DRAM) and operates in a channel with 15 dB loss at Nyquist. Under such bandwidth limitations on and off chip, the Controller and DRAM PHYs consume 13 mW/Gb/s and 8 mW/Gb/s, respectively. Using PRBS  $2^{11}-1$ , the link achieves a timing margin of 0.19 UI at a BER of  $1e-12$  for both read and write operations.

**Index Terms**—CMOS memory integrated circuits, DRAM chips, equalizers, high-speed integrated circuits, interconnections, inter-symbol interference, jitter, synchronization, transceivers.

## I. INTRODUCTION

SYSTEM memory bandwidth is one of the key limitations to high performance computing. Total memory bandwidth can be increased either by increasing the number of links or by increasing the per link data rate. However, there is a cost advantage to increasing the per link data rate to achieve this goal as it reduces the package size as well as the number of DRAMs and eases routing congestion that would otherwise require more metal layers in the package and PCB. Thus, the primary objective of this work is to build the fastest memory transceiver possible while constraining power efficiency (mW/Gb/s) such that thermal issues do not negate the cost savings achieved by reducing the number of links.

It is estimated that the total memory bandwidth required by graphics processors and game consoles will approach 1 TB/s in

2012 [1]. The interface in this work operates at 16 Gb/s/link in order to achieve this aggregate bandwidth with 512 links, a count that has already been demonstrated in state-of-the-art GPUs.

Fig. 1 illustrates such a 1 TB/s memory system [2]. It shows 16 DRAMs, each of which has 32 data links (DQ) and two Command/Address (C/A) links all operating at 16 Gb/s. The 32 DQs provide a combined data bandwidth of 64 GB/s to each DRAM.<sup>1</sup> Both C/A and DQ are point-to-point differential links. Point-to-point signaling provides good signal integrity by minimizing impedance discontinuities while differential signaling is robust against common-mode noise sources such as supply noise and ground bounce. Furthermore, differential signaling results in less crosstalk/EMI and simultaneous switching output noise than single-ended signaling. One difference between the C/A and DQ links is that the C/A is unidirectional whereas the DQ is bidirectional.

The key technical challenges to designing 16 Gb/s memory links discussed in [3] and [4] are: bandwidth limitation in the channel, bandwidth limitation on the silicon, jitter, skew, and power efficiency. To overcome the bandwidth limitation of the channel, an asymmetric equalization architecture that places all the equalization capability on the Controller is used. Since the transistors in the Controller are about twice as fast as the periphery devices in the DRAM in any given process generation, the power of the overall system can be reduced by placing equalization circuits for both read and write directions on the Controller. As a reference, the inverter fanout-of-4 (FO4) gate delay is 22 ps for the controller and 45 ps for the emulated DRAM in this work. To further alleviate the speed limitation of the silicon, half rate and quadrature clocking are employed on the Controller and DRAM, respectively, to reduce the on-chip clock rates. Since high frequency jitter is amplified when passing through a low pass channel [5], duty cycle correction and quadrature correction are employed in the Controller and DRAM respectively to reduce deterministic jitter (DJ). The speed limitation of the silicon exacerbates another type of jitter, namely power-supply induced jitter (PSIJ) which cannot be alleviated with multi-phase clocking. A data-dependent regulator powers the datapath and clock tree of the DRAM which would

Manuscript received August 25, 2008; revised October 22, 2008. Current version published March 25, 2009.

H. Lee, K.-Y. K. Chang, T. Wu, Y. Frans, B. Leibowitz, N. Nguyen, T. J. Chin, K. Kaviani, J. Shen, X. Shi, W. T. Beyene, S. Li, M. Aleksic, F. Quan, J. Zerbe, and R. Perego are with Rambus Inc., Los Altos, CA 94022 USA (e-mail: hlee@rambus.com).

J.-H. Chun was with Rambus Inc., Los Altos, CA 94022 USA, and is now with the Department of Semiconductor Systems Engineering, SungKyunKwan University, Suwon, Korea.

R. Navid was with Rambus Inc., Los Altos, CA 94022 USA, and is now with True Circuits Inc., Los Altos, CA 94022 USA.

F. Assaderaghi and F. S. Lee were with Rambus Inc., Los Altos, CA 94022 USA, and are now with SiTime Corporation, Sunnyvale, CA 94085 USA.

Digital Object Identifier 10.1109/JSSC.2009.2014199

<sup>1</sup>The two C/A links to each DRAM enables a 64-byte access granularity. A 32-byte granularity can be achieved simply by adding two more C/A links per DRAM. The  $2\times$  improvement in granularity would increase the area and power of the interface by only 6%.

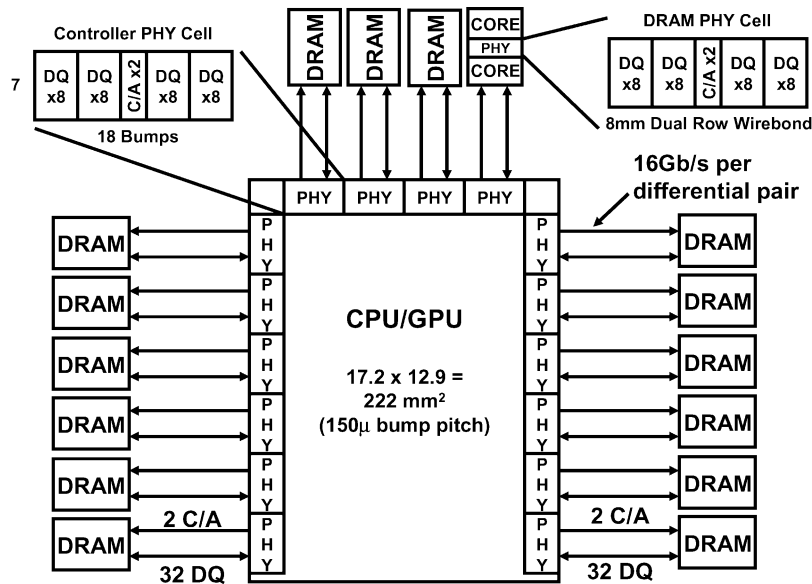


Fig. 1. 1 TB/s memory system using 16 high speed memory interfaces each providing 64 GB/s bandwidth.

otherwise have excessive PSIJ owing to the slower periphery transistors. Skew from trace length, clock distribution, and random transistor mismatches become problematic as the data rate increases since the bit time is reduced (62.5 ps at 16 Gb/s). Phase mixers with active calibration correct lane-to-lane skew in this system. Once again, to reduce the overall system power, an asymmetric clocking architecture [12] is employed where all timing adjustment circuits are implemented on the Controller. Finally, LC PLLs are used on both devices to optimize the random jitter (RJ) performance of the system.

To motivate the architecture and circuits, the on and off chip bandwidth limitations are discussed in Section II. Section III discusses the asymmetric equalization architecture and the transmitter (TX) and receiver (RX) circuits. The clocking architecture and circuits are described in Section IV. The experimental results are given in Section V and the conclusions are made in Section VI.

## II. ON AND OFF CHIP BANDWIDTH LIMITATIONS

The interconnect between the Controller and DRAM consists of a flip-chip thin-core FBGA package (Controller) with substrate traces up to 20-mm long, a 3-inch FR4 PCB trace, and a wirebond FBGA package (DRAM) with a trace length of 8-mm. However, the channel attenuation is dominated by the capacitance of the Controller and DRAM devices which total about 1 pF each. This capacitance is due to the primary ESD diodes, pad metallization, metal routing between the pad and TX, and the diffusion of the TX. The combined RX gate and secondary ESD device capacitance of about 100 fF is shielded from the pad by a 100 Ω resistor, but adds an additional pole below 16 GHz at the receiver input that further attenuates the signal.

Fig. 2 shows the simulated channel response for the DQ links. The first graph shows the magnitude of  $S_{21}$  of the channel with 15 dB attenuation at the Nyquist frequency of 8 GHz. The second graph shows the single bit response (SBR) with 4 main

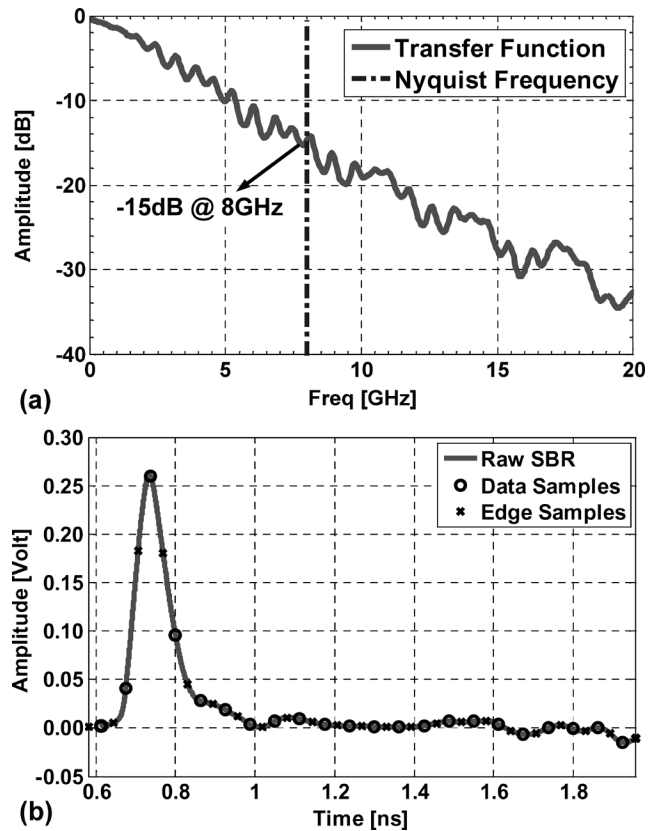


Fig. 2. Target DQ channel characteristic showing (a) the magnitude of  $S_{21}$  with 15 dB attenuation at Nyquist frequency and (b) single bit response with four main ISI components (input normalized to  $1 V_{ppd}$  in simulation).

ISI components. The SBR shows the need for equalization as the sum of the absolute value of the ISI components is almost equal to the magnitude of the main cursor, indicating a closed eye.

At 16 Gb/s, the on-chip bandwidth is also a limitation. In the CMOS process that the Controller and DRAM were designed

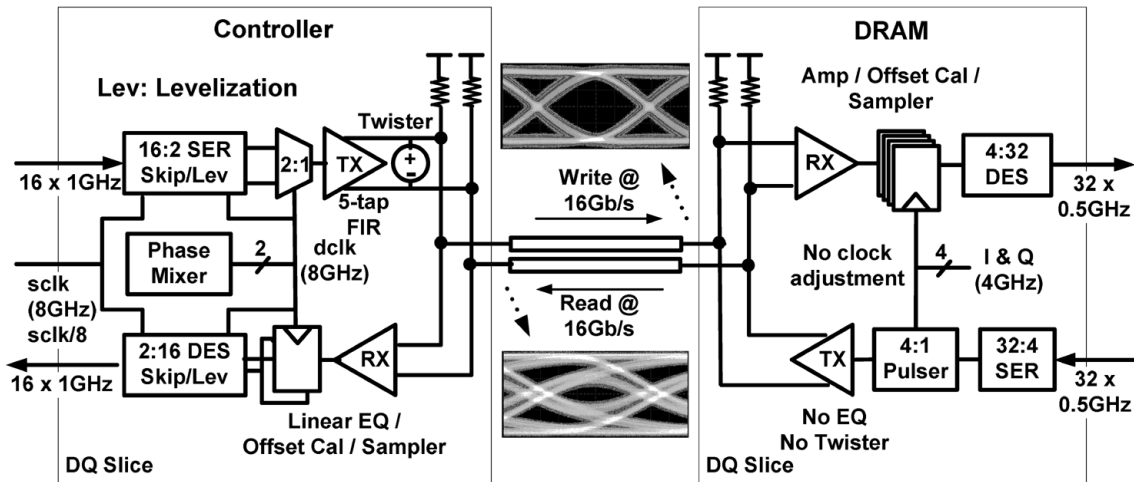


Fig. 3. Asymmetric link architecture places equalization, clocking, and link diagnostic complexity on the Controller for improved power efficiency.

in, the FO4 inverter delays are 22 ps and 45 ps, respectively, in the slow device, low voltage, high temperature corner. This translates into a low unit-interval to inverter FO4 ratio (UI/FO4) of 2.8 (Controller) and 1.4 (DRAM). Prior research showed that the minimum clock period to pass through a FO4 chain of inverters without amplitude reduction is 8-FO4 [6]. This means that the high speed portions of this design will require logic with effective FO less than 4. In fact, the required FO of the clock trees approaches 2 for double data rate at the Controller and quad data rate at the DRAM. The primary cost of low FO design is added power consumption. A related observation is that designing a given circuit in the Controller is more power efficient than designing it in the DRAM. This is true for analog circuits as well as logic, as the lower  $f_t$  and high thresholds in the DRAM process degrade amplifier power efficiency. This is the motivation behind the asymmetric memory link architecture, shown in Fig. 3, where the Controller contains more of the timing, equalization, and link diagnostic circuitry. The details will be covered in subsequent sections.

In addition to the slower devices, DRAM processes also have fewer metal layers (at most 3) which have higher resistivity and larger capacitance to the substrate than those for ASICs. This limits the options for on-chip inductors as well as global clock and power grid distribution.

On a final note, we want to make clear that the DRAM presented in this paper is an *emulated* DRAM. Since this memory link is designed for production in the 2010–2012 timeframe, the process does not exist yet. Hence, to validate the architecture and ideas set forth, a 65 nm CMOS process was degraded (by design constraints) to match the best estimates on the future DRAM process. For example, non minimum length and high threshold devices were exclusively used to match the threshold voltage and the FO4 of the future DRAM process and metal usage was restricted to three routing layers as described in [4]. In contrast, the ASIC process in the 2010–2012 timeframe will be faster than the 65 nm CMOS process used for the Controller. Hence, the power efficiency (mW/Gb/s) of the Controller will improve in the future process.

### III. EQUALIZATION ARCHITECTURE AND CIRCUITS

As shown in Fig. 3, the Controller provides equalization for both write and read operations. In the write direction, the channel is equalized using a 5-tap TX FIR.<sup>2</sup> In the read direction, an analog linear equalizer provides about 6 dB of emphasis at the Nyquist frequency of 8 GHz. As the write and read directions leverage different equalizers, some performance mismatch between them is expected, but they are designed to offer similar performance to prevent one direction from limiting the overall system margin. Based on signal integrity simulations, the nominal swing of the Controller (DC swing without equalization) is made 50% larger than the DRAM for this purpose. The DRAM TX output swing is set to 540 mV<sub>ppd</sub> ( $I_{TX} \sim 11$  mA) so that its high threshold transistors have sufficient saturation margin. The Controller swing is thus 800 mV<sub>ppd</sub> ( $I_{TX} \sim 16$  mA).

#### A. Transmit FIR and Link Diagnostic Circuits

A block diagram of the Controller TX is shown in Fig. 4. It has 5 FIR taps (1 pre, 1 main, 3 post) to equalize the SBR of Fig. 2. Each segment has a 2:1 multiplexing pre-driver that serializes the 8 Gb/s input into a full rate 16 Gb/s stream. The output drivers are open-drain differential pairs with digitally controlled tail current sources for tap weight adjustment.

The FIR circuits necessary for equalization are also used to evaluate the performance of the link. With minor reconfiguration enabled by digital registers, the Controller TX can add a DC differential voltage offset into the channel so that the voltage margin can be measured. This is important because while the write eye at the pad of the DRAM is relatively wide open due to the Controller TX-FIR, the read eye at the pad of the Controller will be severely degraded since the effect of equalization is only visible on chip at the output of the Controller RX. This is illustrated by the eye diagrams in Fig. 3.<sup>3</sup> Hence, it is not possible to verify the

<sup>2</sup>The DRAM RX amplifier has 3 dB nominal gain and can be configured to provide some equalization of on-chip losses via de-emphasis, but is not capable of equalizing the ISI shown in Fig. 2.

<sup>3</sup>The eye diagrams in Fig. 3 are for illustration only and are not the actual eyes at the pads.




Fig. 4. (a) Block diagram, and conceptual schematics of (b) normal write mode, (c) margining write mode, and (d) margining read mode of segmented Controller TX.

functionality of the equalizers by observing eye diagrams on an external scope. The ability to add a voltage offset combined with timing control from the phase mixers enables the Controller to map the effective BER eye diagram in voltage and time for the entire link, including the RX circuits themselves. There are two distinct features in our implementation compared to [7] and [8]: the first is that the voltage offset is added by the transmitter rather than the receiver and the second is that the Controller transmitter adds the voltage offset to margin the link for both the write and read operations. These features leverage both the asymmetry and the bidirectionality of the link. Fig. 4(b)–(d) shows how the Controller TX is reconfigured for the various modes of operation. Conceptually, each TX tap driver consists of 3 separate segments. The first segment consumes half the current and the sum of the other 2 segments makes up the other half. During normal write operation, all three segments are driven with data by their respective pre-drivers. When the write direction is margined, only the first segment (‘Data Segment’) is connected to its pre-driver and transmits data. The second segment (‘Offset Segment’) is hard switched to add a differential offset to the channel. The polarity of the offset is set by which transistor of the differential pair is connected to supply and which is connected to ground. The magnitude of the offset ( $I_{OS}$ ) is adjusted with the digital code to the tail current source. The inputs of the third differential pair (‘CM Segment’) are both tied to the supply to pull both output nodes down. The current of the CM segment ( $I_{CM}$ ) is set according to constraints (1) and (2) so that the output common-mode level is kept constant even as the offset is varied. Essentially, the Controller TX is adding a differential offset to its own signal on the channel. When margining the read direction, the DRAM TX is put in half swing mode and the Data Segment of the Controller

TX is now turned off. The Offset and CM segments are still enabled. As mentioned before, the swing of the DRAM is less than the Controller, so the current of the Offset and CM Segments are scaled by the same proportion to maintain the common mode of the signal at the input of the Controller RX. In this manner, the Controller TX adds a differential offset to the signal from the DRAM.

$$I_{TX} = I_{SW} + I_{OS} + I_{CM} \quad (1)$$

$$I_{SW} = I_{OS} + I_{CM} = I_{TX}/2 \quad (2)$$

### B. Controller Receive Equalizer

The Controller RX shown in Fig. 5(a) starts with a continuous time equalizer (i.e. linear equalizer). The RX then splits into two sampling paths (even and odd) to relieve the maximum on-chip clock frequency. The circuits are kept as small as possible to minimize power consumption, and the resulting random mismatch components are corrected via current summing at the sampler inputs. The simulated mismatch of the complete RX has a standard deviation of 12 mV referred to the input of the samplers where the correction takes place. The offset correction range is  $\pm 40$  mV with 3 mV resolution.

A linear equalizer is chosen over the more versatile decision feedback equalizer (DFE) to avoid having a feedback path of 3 FO4. Unlike the more prevalent source degenerated and resistor loaded topology [9], the linear equalizer in Fig. 5(b) uses active inductor loads [10] to provide emphasis in its transfer function. Hence, the DC gain is not changed when varying the peaking in the transfer function, as seen in Fig. 5(c). The active inductor topology was chosen as simulations showed that it provides 3 dB more emphasis than the source degenerated topology




Fig. 5. (a) Controller RX block diagram, (b) linear equalizer using active inductor loads, and (c) AC response in nominal corner.

when the input gate size, current consumption, and DC gain are kept equal.

The active inductor equalizer is not without its challenges. For  $V_{dd}$  referenced signaling, the active inductor loads need to be biased with a gate voltage higher than  $V_{dd}$  ( $V_{bh}$ ) in order to provide significant output swing. A good choice for  $V_{bh}$  is  $V_{dd} + V_{tn}$  such that the load device remains saturated and still acts as an inductor for all output voltages.  $V_{bh}$  is generated by a voltage boosting circuit similar to [11] with an added feedback loop to compensate for various losses from PVT and input clock conditions.

The transfer function of the equalizer can be expressed as (3).  $C_L$  is the load capacitance,  $g_m$  is the transconductance of the input device, and  $g_{m,L}$  is the transconductance of the load. To simplify the analysis, the finite output impedance of the devices as well as any parasitic capacitance between the gate of the active inductor and (AC) ground were ignored. However, care must be taken to minimize this parasitic capacitance as it reduces the effective inductance. The frequency of the zero can be lowered, and the peak gain increased, by either increasing  $R_G$  or  $C_G$ . Digitally controlled resistors are used for  $R_G$  to provide adjustable peaking of the AC response. NMOS accumulation mode capacitors were added to supplement the inherent  $C_{gs}$  of the load device to increase the peaking.

$$A_V(s) = -\frac{g_m}{g_{m,L}} \cdot \frac{s \cdot R_G \cdot C_G + 1}{s^2 \cdot \frac{R_G \cdot C_G \cdot C_L}{g_{m,L}} + s \cdot \frac{(C_L + C_G)}{g_{m,L}} + 1} \quad (3)$$

Fig. 5(c) shows a family of transfer function curves simulated in the nominal corner for the complete receiver. The DC gain is 2.5 dB at all settings, and the AC gain is controllable from 6–9 dB to provide up to 6.5 dB of emphasis.

#### IV. CLOCKING ARCHITECTURE AND CIRCUITS

To ensure sufficient timing margin at 16 Gb/s, high performance clocking circuits are necessary on both the Controller and DRAM. Therefore, LC-VCO based PLLs, which offer superior noise performance compared to ring-VCO based PLLs, are employed on both PHYs. Fewer and more resistive metal layers present challenges to designing an LC-VCO in a DRAM process. Solutions to these challenges are discussed in [4]. A low-cost 500 MHz reference clock from an external clock part is distributed in a controlled impedance fashion to the Controller and then to the DRAM. An asymmetric clocking architecture is employed (Fig. 3). The burden of phase adjustment, skip, and levelization are all placed on the Controller. The phase mixer adjusts the TX and RX clock for optimal BER performance and also compensates the skew between bit slices to ease routing on the board. The skip block handles the arbitrary phase clock domain crossing between  $sclk$  and  $dclk$  (phase mixer output). The purpose of levelization is to address skew greater than 1 UI by equalizing the latency of the bit slices. Together, the levelization and phase mixer enable an effective phase adjustment range of 64 UI for both the TX and RX. These features are also used to compensate C/A and DQ flight time differences as well as reference clock skew across the multiple DRAMs.

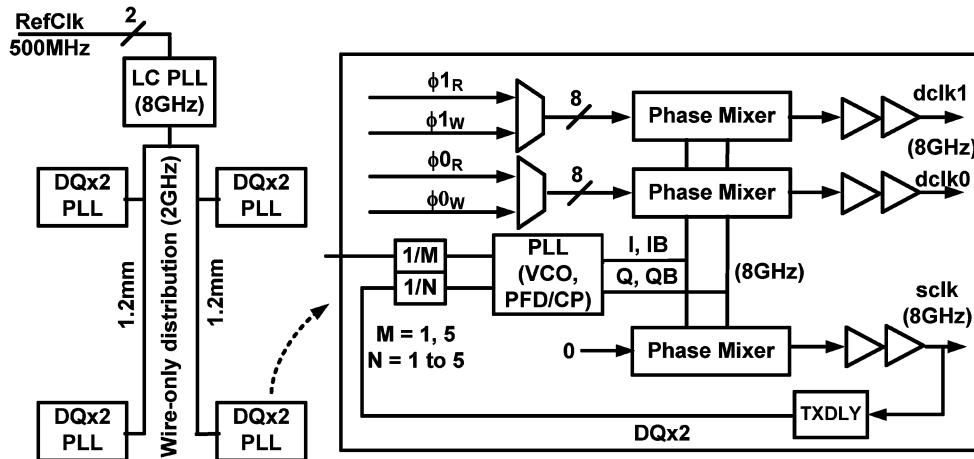


Fig. 6. Controller clocking architecture using a cascade of LC and ring-based PLLs. Each ring-based PLL is shared between two bitslices. Each bitslice has one phase mixer to correct skew in both the write and read directions.

### A. Controller Clocking Architecture and Circuits

The Controller clocking architecture is shown in Fig. 6. On the top of each byte, an LC-VCO based PLL generates an 8 GHz differential clock signal from the 500 MHz reference. This 8 GHz clock is divided down to 2 GHz before distribution to reduce the power consumption. Although this LC PLL results in a clean clock output, there are two challenges. First, the phase mixer that adjusts dclk requires quadrature 8 GHz phases to be interpolated. Second, a wide frequency range is desired to test the interface at different data rates. To address these, each byte level LC PLL is followed by local M/N ring-VCO based PLLs which generate multiple phases over a wide frequency range.

To reduce power consumption, each ring PLL is shared by two neighboring bit slices (DQx2). The DQx2 clocking circuits generate a dclk0 for bit0, a dclk1 for bit1, and a common sclk. To avoid the area and power consumption of two phase mixers per bit slice (a total of five for a DQx2 block), one phase mixer is time-multiplexed between the TX and RX of the same bit slice, since the two operate at different times in the same bidirectional I/O. Two phase codes are obtained from separate read and write operation timing calibrations and provided to the shared phase mixer through a multiplexer which is appropriately selected according to the data transaction taking place.

### B. DRAM Clocking Architecture and Circuits

The clocking architecture of the DRAM is shown in Fig. 7. One LC-PLL is used in each DQ byte to multiply the 500 MHz reference clock to 8 GHz. Because of the aforementioned transistor and global metal routing limitations, the 8 GHz PLL output is divided and distributed as quadrature 4 GHz clocks (two pairs of differential 4 GHz clocks I/Ib and Q/Qb) to the individual bit slices. These 4 GHz clocks are distributed in CML to minimize PSIJ. In each bit slice, these quadrature 4 GHz clocks are locally converted to CMOS levels to drive the front-end circuits of the TX and RX.

16 Gb/s operation in the DRAM is challenging even with quadrature rate operation. Of primary concern is PSIJ which is proportional to the delay of the circuit, and is not improved by

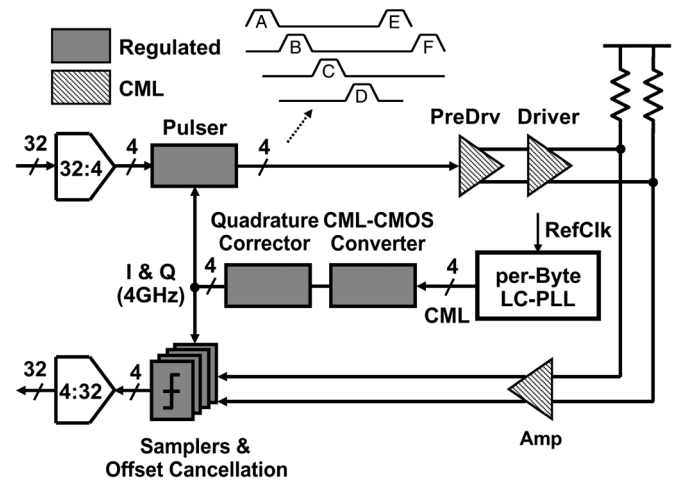


Fig. 7. DRAM PHY architecture showing clocking and supply regulation domain.

multiphase clocking. Since the relatively slow process leads to a large circuit delay, PSIJ is more severe on the DRAM. It is not just the clock tree that is of concern, but also the clocked TX and RX circuits that are closest to the pad. Variations in the clk-q delay of the pulser as well as the aperture time of the RX sampler equally hurt the performance of the link. These blocks are implemented in CMOS due to headroom limitations and better power efficiency than their CML counterparts. However, this comes at the cost of higher PSIJ, which is addressed by supply regulation of the front-end datapath circuits and local (bit slice level) clocking circuits, as shown in Fig. 7. This approach maintains a net power advantage over a full CML datapath implementation, even with the regulator overhead.

The primary challenge of the voltage regulator design is maintaining a stable output voltage despite the data-dependent transient load currents of the CMOS datapath circuits. A voltage regulator with a replica load has been shown to provide power and area efficient suppression of high frequency transient noise from the external supply [13]. We extend the idea by using a data-dependent replica (Fig. 8) that has a transient load current

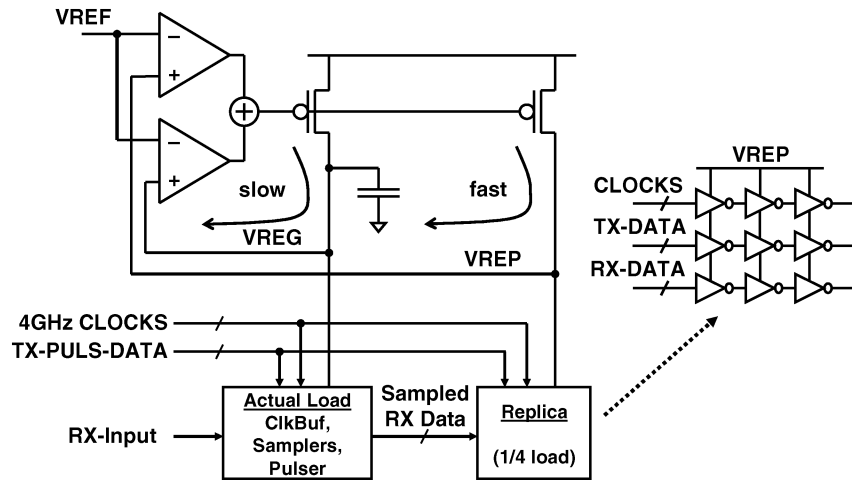


Fig. 8. Data dependent regulator reduces power supply induced jitter on DRAM.

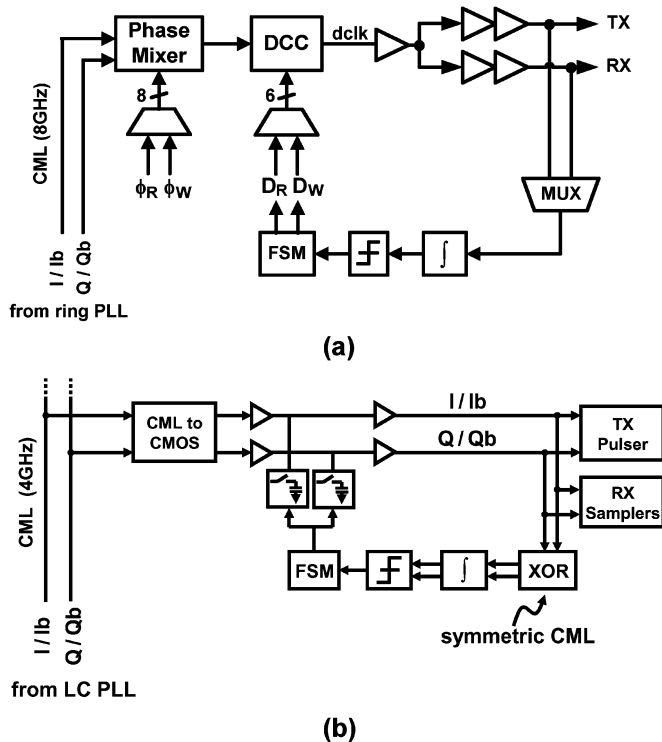


Fig. 9. (a) Duty cycle correction (DCC) circuits for the Controller and (b) quadrature error correction (QEC) circuits for the DRAM.

proportional to that of the front-end, providing fast feed-forward regulation of transient load currents. When a transient load step occurs, the small capacitance of the replica circuit allows the change in VREP to be detected very quickly. The amplifier can then quickly adjust the PMOS gate voltage to match the new load condition before any significant change occurs in the output voltage VREG. The regulated load includes local clock buffers, receive samplers, transmit pulser, and several flip-flops. The replica load is implemented with simple CMOS inverters which receive the same clock and data signals as the actual load, but are scaled to consume one quarter of the load current. The regulator is designed to power a 36 mA maximum load current with 300 MHz closed loop bandwidth, and an output

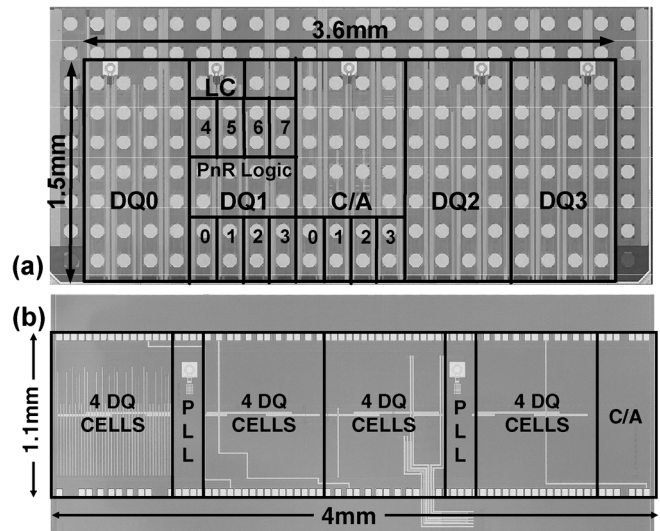


Fig. 10. Photomicrographs of (a)  $\times 32$  Controller PHY and (b)  $\times 16$  DRAM PHY.

capacitance of 100 pF. This regulator reduces the capacitance required to achieve a given transient load ripple by 4 $\times$ . The simulated worst case PSRR over all frequencies is -16 dB.

### C. Duty-Cycle and Quadrature Correction

High frequency jitter can be severely amplified by lossy channels [5]. Hence, both duty cycle error on the Controller and quadrature phase error on the DRAM must be minimized. These are achieved by using a duty cycle corrector (DCC) and a quadrature error corrector (QEC), respectively, as shown in Fig. 9(a) and (b).

Similar to the phase mixer on the Controller, a single digitally controlled DCC is time-multiplexed between the TX and the RX clocks. The clock duty cycle is adjusted toward its desired 50% value by using a duty cycle correction loop. This feedback loop depicted in Fig. 9(a) consists of the DCC to adjust clock duty cycle, a MUX to select the clock signal to be calibrated, an integrator followed by a comparator to detect the duty cycle error polarity, and an FSM that accordingly increments or decrements

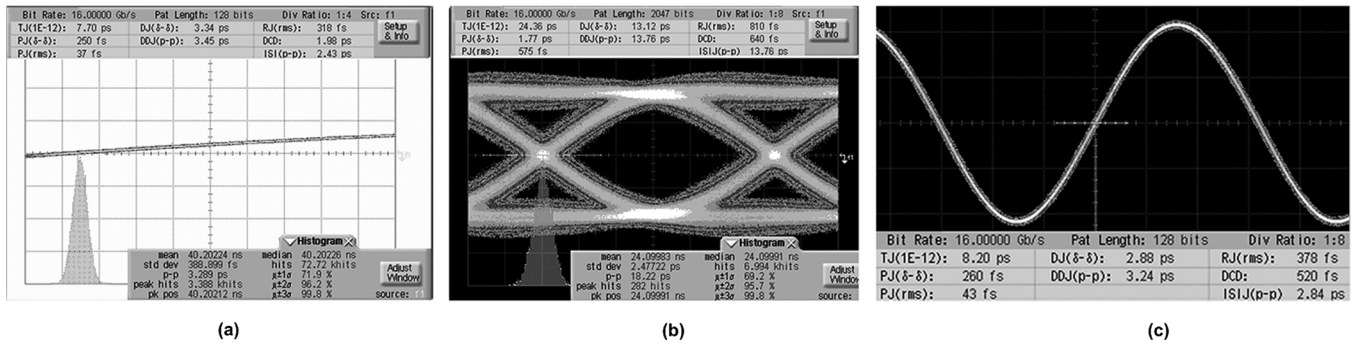


Fig. 11. Jitter characterization of (a) Controller LC PLL at 8 GHz, (b) Controller TX output at 16 Gb/s, and (c) DRAM TX output at 16 Gb/s.

the control code of the DCC. The duty cycle adjustment range is approximately  $\pm 0.08$  UI with an LSB of 2.5 mUI.

On the DRAM PHY, the accumulated quadrature phase errors from mismatches in the global distribution, CML-to-CMOS conversion, and local clock buffers are corrected by the feedback loop shown in Fig. 9(b). Quadrature error is detected near the front-end circuits by a symmetric CML XOR transconductor followed by an integrator. The integrator output is then sensed by a comparator whose output is accumulated by an FSM to adjust switched capacitor loads. The resolution is 20 mUI (1.25 ps at 16 Gb/s) and the range is  $\pm 0.07$  UI (4.5 ps). Using this scheme, quadrature error is limited by local device mismatches in the quadrature detector itself.

## V. MEASUREMENTS

The testchip was fabricated in TSMC 65 nm G+ technology. Fig. 10 shows the cell photomicrographs of a Controller interface with 4 C/A<sup>4</sup> and 32 DQ links as well as a DRAM interface with 1 C/A and 16 DQ links. Each Controller byte has one LC PLL, eight bit slices, and a dedicated place and route (PnR) logic block that includes pattern generators and checkers, timing margin measurement logic, termination calibration, phase mixer control, etc. On the DRAM interface, the PLL on the left provides high speed clocks to 8 neighboring DQ slices, while the PLL on the right provides clocks to the other 8 DQ slices as well as to the receiver in the C/A cell for command requests. The DRAM C/A cell also includes associated logic for command processing. Memory reads and writes are targeted to 512 kbits of on-chip SRAM to allow interface testing with various read and write activity patterns. Note that the C/A slice is placed to the side to emulate the signal routing and timing requirements for a 32-bit interface with C/A in the center, which would be approximately 8 mm wide.

The measured jitter performance of the Controller LC PLL, Controller TX output, and DRAM TX output are shown in Fig. 11. The measured random jitter (RJ) of the Controller LC-PLL output routed directly to pads is 318 fs rms and total jitter (TJ) is 7.7 ps at a BER of 1e-12. When transmitting a 2<sup>11</sup>-1 PRBS pattern at 16 Gb/s, the Controller TX waveform has RJ of 810 fs and TJ of 24.36 ps. Apart from the RJ contribution, the TJ increases because of power supply induced jitter (PSIJ) in the PLL/clock tree and transmitter, crosstalk from

<sup>4</sup>The extra C/A were routed off the board via SMP connectors to provide visibility of various analog/mixed signal components (e.g. TX-FIR, PLL).

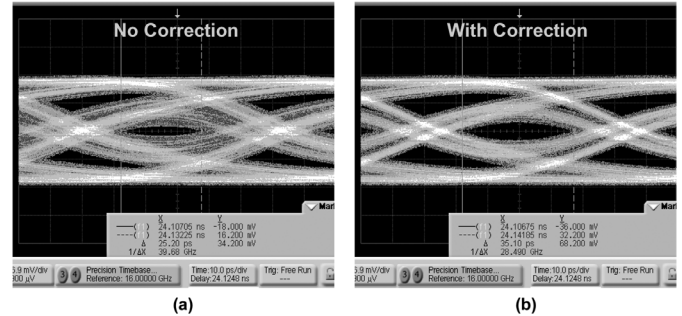


Fig. 12. DRAM TX eye (a) before and (b) after quadrature correction. Eye is improved from 34 mV  $\times$  25 ps to 68 mV  $\times$  35 ps.

adjacent bit slices, residual ISI, and reference feed-through. The DRAM TX output has a measured RJ of 378 fs and a TJ of 8.2 ps when transmitting a clock pattern. The DRAM TX has less jitter than the Controller TX since the LC drives the TX directly rather than passing through another ring-based PLL. The ring-based PLL provides flexibility to the Controller but increases the RJ and PSIJ.

The effect of jitter amplification and the need for duty cycle and quadrature (IQ) correction is shown in Fig. 12. The timing margin of the PRBS 2<sup>11</sup> DRAM TX eye increases by 10 ps with only a 3.75 ps clock correction. Thus, the loss in this channel, as evidenced by the significant amount of ISI, more than doubles the impact of the quadrature error on the eye size.

Fig. 13 shows measured BER bathtub curves for both write and read directions at 16 Gb/s with a PRBS 2<sup>11</sup>-1 data pattern. The measurement was taken with optimized TX/RX equalization coefficients and calibrated settings for RX voltage offset (both devices), for DCC (Controller), and for IQ correction (DRAM). Positive eye openings of 0.19 UI are observed at a BER of 1e-12 in both directions.

Fig. 14 shows the 16 Gb/s on-chip eyes measured using the asymmetric link diagnostic feature for both the write and read directions. This *in situ* measurement shows the *effective* eye of the system including the complete effects of the channel, equalization, jitter, and all TX and RX circuit limitations. Although the transmit swings are cut in half for these tests as mentioned earlier, both directions show timing margin greater than 0.3 UI and voltage margin of 25 mV<sub>ppd</sub> at a BER of 1e-9.

To test the robustness of the memory interface, the timing margin is measured across a byte. As shown in Fig. 15, both



TABLE I.  
PERFORMANCE SUMMARY OF CONTROLLER (CTRL) AND DRAM

Technology	TSMC 65nm G+ CMOS
Supply Voltage (VddA/Vdd/VddIO) CTRL & DRAM	1.1/1.1/1.2 & 1.0/1.0/1.2 V
CTRL Cell size (4 DQ bytes + C/A block)	3.6mm x 1.5mm
DRAM Cell size (2 DQ bytes + C/A block)	4mm x 1.1mm
Bandwidth per DQ and C/A link	16Gb/s
Total data bandwidth per x32 interface	64GB/s
Worst Case UI/FO4 Ratio CTRL & DRAM	2.8 & 1.4
Interface Power Efficiency CTRL & DRAM	13 & 8 mW/Gb/s
CTRL LC clock jitter ( $RJ_{rms}/TJ_{BER=10^{-12}}$ )	318fs / 7.7ps
CTRL Tx jitter ( $RJ_{rms}/TJ_{BER=10^{-12}}$ ) w/ PRBS, all links active	810fs / 24.4ps
DRAM Tx jitter ( $RJ_{rms}/TJ_{BER=10^{-12}}$ ) w/ 1010, all links active	378fs / 8.2ps
Read/write BER @ 16Gb/s (from overnight runs)	$< 10^{-14}$
Write timing margin at BER = $10^{-12}$	0.19 UI, 11.9ps
Read timing margin at BER = $10^{-12}$	0.19 UI, 11.9ps




Fig. 13. 16 Gb/s BER bathtub curves of both write and read operations using continuous  $2^{11} - 1$  PRBS data





Fig. 15. Timing margin measured across a byte at 16 Gb/s. Both (a) write and (b) read operations have timing margin greater than 0.5 UI at a BER of  $1e-3$ . This measurement includes the effect of bus turnaround.

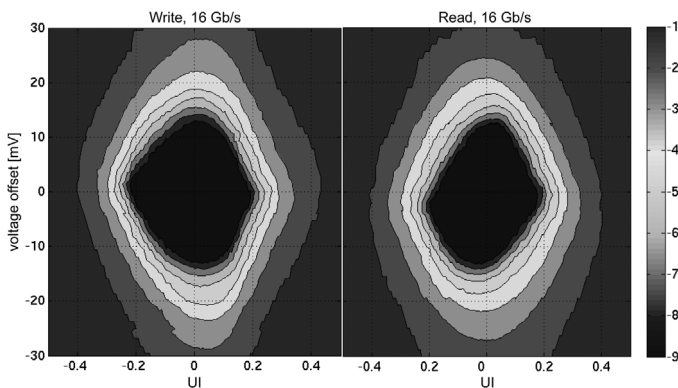


Fig. 14. On-chip eyes measured using the Controller TX to margin both the write and read directions at 16 Gb/s.

the write and read operations have timing margin greater than 0.5 UI at a BER of  $1e-3$ . The effect of bus turnaround (the act of powering on and off the TX or RX between transactions), which results in the worst case supply noise condition, is included in these measurements.




Fig. 16. Power breakdown at 16 Gb/s of Controller and DRAM.

Fig. 16 shows the breakdown of the power for both devices. The power efficiency of the interface cells are 13 mW/Gb/s for

the Controller and 8 mW/Gb/s for the DRAM. In both cases, clock related power is greater than 40% and is the most significant contributor, while the TX and RX combined account for about 25% of the power. Finally, Table I summarizes the achieved performance of the Controller and DRAM interfaces.

## VI. CONCLUSION

A 16 Gb/s/link memory interface has been demonstrated. An asymmetric memory architecture and circuit innovations allow the testchip to overcome the primary challenges of on and off chip bandwidth limitations while achieving 13 mW/Gb/s and 8 mW/Gb/s at the Controller and emulated DRAM, respectively. The jitter in the system was optimized using LC PLLs, duty cycle and quadrature error correction circuits, and a data dependent regulator to reduce PSIIJ. The technologies demonstrate the feasibility of a 1 TB/s memory system in the 2010–2012 timeframe.

## ACKNOWLEDGMENT

The authors are grateful to M. Ludwig, the late Bob Nickell, G. Holst, D. Olarte, A. Chen, M. Christiansen, and C.-Y. Kuo for their layout support, C. Madden, N. Kim, G. Yip, R. Schmitt, and X. Qi for their measurement support.

## REFERENCES

- [1] S. Woo, "Computing trends and applications driving memory performance," in *Rambus Developer Forum—Japan*, Nov. 28, 2007.
- [2] R. Perego, "Architectural considerations for next-generation memory systems," in *Rambus Developer Forum—Japan*, Nov. 28, 2007.
- [3] K. Chang *et al.*, "A 16 Gb/s/link, 64 GB/s bidirectional asymmetric memory interface cell," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2008, pp. 126–127.
- [4] N. Nguyen *et al.*, "A 16-Gb/s differential I/O cell with 380 fs RJ in an emulated 40 nm DRAM process," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2008, pp. 128–129.
- [5] S. Chaudhuri, W. Anderson, J. Bryan, J. McCall, and S. Dabral, "Jitter amplification characterization of passive clock channels at 6.4 and 9.6 Gb/s," in *Proc. IEEE Electrical Performance of Electronic Packaging Conf.*, Oct. 2006, pp. 21–24.
- [6] A. Chandrakasan, W. Bowhill, and F. Fox, *Design of High-Performance Microprocessor Circuits*. Piscataway, NJ: IEEE Press, 2001, p. 402.
- [7] J. L. Zerbe, C. W. Werner, V. Stojanovic, F. Chen, J. Wei, G. Tsang, D. Kim, W. F. Stonecypher, A. Ho, T. P. Thrush, R. T. Kollipara, M. A. Horowitz, and K. S. Donnelly, "Equalization and clock recovery for a 2.5–10-Gb/s 2-PAM/4-PAM backplane transceiver cell," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2121–2130, Dec. 2003.
- [8] G. Balamurugan, J. Jaussi, D. R. Johnson, B. Casper, A. Martin, J. Kennedy, R. Mooney, and N. Shanbhag, "Receiver adaptation and system characterization of an 8 Gbps source-synchronous I/O link using on-die circuits in 0.13  $\mu\text{m}$  CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2004, pp. 356–359.
- [9] R. Farjad-Rad *et al.*, "0.622–8.0 Gbps 150 mW serial IO macrocell with fully flexible preemphasis and equalization," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2003, pp. 63–66.
- [10] E. Säckinger and W. C. Fischer, "A 3-GHz 32-dB CMOS limiting amplifier for SONET OC-48 receivers," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 1884–1888, Dec. 2000.
- [11] P. Favrat, P. Deval, and M. J. Declercq, "A high-efficiency CMOS voltage doubler," *IEEE J. Solid-State Circuits*, vol. 33, no. 3, pp. 410–416, Mar. 1998.
- [12] K. Chang, W. Ellersick, S. Chuang, S. Sidiropoulos, and M. Horowitz, "A 2 Gb/s/pin CMOS asymmetric serial link," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 1998, pp. 216–217.
- [13] E. Alon, J. Kim, S. Pamarti, K. Chang, and M. Horowitz, "Replica compensated linear regulators for supply-regulated phase-locked loops," *IEEE J. Solid-State Circuits*, vol. 41, no. 2, pp. 413–424, Feb. 2006.



**Haechang (Hae-Chang) Lee** (M'06) received the B.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1998 and 2007, respectively.

He is currently a Principal Engineer at Rambus Inc. developing technology for high-speed and low-power memory interfaces. Since joining Rambus in 2005, he has worked on CDR, PLL, and equalizer circuits for a backplane SerDes, chip-to-chip processor bus, and memory interface. Prior to joining Rambus, he worked at LGC Wireless

(a startup) and Agilent Technologies. He is an author or coauthor of more than ten technical papers.

Dr. Lee is a member of Phi Beta Kappa and Tau Beta Pi.



**Kun-Yung Ken Chang** (S'93–M'99) received the B.S. degree in electrical engineering from National Taiwan University, Taipei, Taiwan, in 1990, the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994 and 1999, respectively.


He is currently a Senior Engineering Manager with Rambus Inc., responsible for the next-generation high-speed low-power memory interface. His research interests include high-speed mixed-signal CMOS circuit design, transmitter and receiver design, PLL/DLL design, circuit noise analysis, and signal integrity analysis.

Dr. Chang is currently on the technical program committee for the VLSI Circuits Symposium and the IEEE Custom Integrated Circuits Conference.



**Jung-Hoon Chun** (M'06) received the B.S. and M.S. degrees in electrical engineering from Seoul National University, Seoul, Korea, in 1998 and 2000, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 2006.

He is an Assistant Professor in the Department of Semiconductor Systems Engineering, SungKyunKwan University, Suwon, Korea. From 2000 to 2001, he was with Samsung Electronics, Kiheung, Korea, where he developed BiCMOS RF front-end IC for wireless communication. From 2006 to 2008, he was with Rambus Inc. where he was involved with high-speed interface design. His current research includes advanced on-chip ESD protection circuit and I/O, RF/analog circuit, high-speed serial links, and high power devices.



**Ting Wu** (M'07) received the B.E. degree in electronic engineering from Tsinghua University, Beijing, China, in 2000, the M.Phil. degree in electrical and electronic engineering from the Hong Kong University of Science and Technology, Hong Kong, in 2003, and the Ph.D. degree in electrical and computer engineering from Oregon State University, Corvallis, OR, in 2007.

From February 2003 to August 2003, he was a Mixed-Signal Designer with Via Technologies Inc., Beijing, China, where he designed serializer for PCIX. During summer 2004, he was with Broadcom Corporation, Irvine, CA, working on CMOS circuits for high-speed serial link applications. In February 2007, he joined Rambus Inc., Los Altos, CA, as a Senior Member of Technical Staff. His current research interests are in CMOS integrated circuits and systems for wired and wireless communications.



**Yohan Frans** (M'02) received the B.S. degree in electrical engineering from Bandung Institute of Technology, Indonesia, in 1995, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2001.

He is a Senior Principal Engineer at Rambus Inc. His current interests include clocking circuits and architectures, high-speed serial links, and memory interfaces.



**Brian Leibowitz** (M'05) received the B.S. degree in electrical engineering and the Edwin H. Armstrong Award from Columbia University in 1998. In 2004, he received the Ph.D. degree in electrical engineering from the University of California at Berkeley, where 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 a Principal Engineer at Rambus Inc. Since joining Rambus in 2004, he has worked on equalization and mixed-signal circuit design for a variety of high-speed serial links and memory interfaces.




**Nhat Nguyen** received the B.S. degree with highest honors in computer engineering, computer science, and mathematics from Portland State University, Portland, OR, in 1986, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley, in 1987 and 1991, respectively.

From 1991 to 1995, he was with Hewlett Packard working on RF circuits in standard silicon bipolar technology. He was with MicroUnity in 1996 working on analog circuits for graphic processors. Since August 1996, he has been with Rambus

working on high-speed I/O interfaces for memory, networking, and computing applications. He is presently a Senior Design Manager. He has published one book chapter and several technical papers. He holds 15 patents. His current research interest includes PLL design, clocking architecture, high-speed and low-power serial links, and wideband circuits.

Dr. Nguyen was a past Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS.



**T. J. Chin** (M'96) was born in Taipei, Taiwan. He received the B.S. degree from Tamkang University, Taipei, Taiwan, in 1991, and the M.S. degree in electrical engineering from San Jose State University, San Jose, CA, in 1994.

From 1995 to 1997, he was with S3 Inc. working on RAMDAC and analog/digital PLLs. Since 1997, he has been with Rambus Inc., Los Altos, CA, working on high-speed I/O circuit design.




**Kambiz Kaviani** (S'97–M'03) received the B.S. degree from Sharif University of Technology, Tehran, Iran, in 1994, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, both in electrical engineering, in 1999 and 2003, respectively.

He is currently a Principal Engineer at Rambus Inc., Los Altos, CA, working on high-performance VLSI embedded memory interface technologies. During the summer of 1998, he was with RadioLAN Corporation, Sunnyvale, CA, where he was involved in the design and development of next generation 5


GHz wireless local area networks (WLAN). At Stanford, he was focusing on high-speed and high-precision VLSI data converters and capacitive macro-machined ultrasonic transducers (CMUT) developments. His main interests include high performance CMOS mixed-signal and RF integrated circuits and systems.

Dr. Kaviani was a co-recipient of the 2002 Outstanding Paper Award of the IEEE Ultrasonic, Ferroelectrics, and Frequency Control Society. He was a member of the Iranian team in the 21st international Physics Olympiad, Groningen, The Netherlands.



**Jie Shen** received the B.S. and Ph.D. degrees in electrical engineering from Tsinghua University, Beijing, China, in 1992 and 1998, respectively.


He is a Principal Engineer at Rambus Inc. His current interests include high-speed memory interface and design methodology.



**Xudong Shi** received the B.S. and M.S. degrees in precision instruments from Shanghai Jiao Tong University, Shanghai, China, and the M.S. degree in electrical engineering from Georgia Institute of Technology, Atlanta, GA.


In 1994, he joined Level One Communications Inc., Sacramento, CA, where he designed T1/E1 transceiver AFEs. From 1996 to 1998, he was with Rockwell Semiconductor Systems, Newport Beach, CA, where he designed Ethernet and xDSL transceiver AFEs. From 1998 to 2002, he was with

Lucent Microelectronics (Agere Systems Inc.), most recently as a Technical Manager and DMTS, where he led design and development of HomePNA, Gigabit Ethernet and CMOS low-power 10-Gb/s SONET/SDH transceiver AFEs. Since 2002, he has been with Rambus Inc., Los Altos, CA, where he is currently a Principal Engineer designing mixed signal circuits for high-speed interfaces such as that in the Sony Playstation-3 Cell processor. He holds four U.S. patents.



**Wendemagegnehu (Wendem) T. Beyene** (M'88–SM'05) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer engineering from Columbia University and the University of Illinois at Urbana-Champaign.

In the past, he was with IBM, Hewlett-Packard Company, and Agilent Technologies, where he worked on design and electrical characterization of advanced multilayer packages and analog and RF circuit simulation tools. He is currently with Rambus as a Senior Principal Engineer responsible for signal integrity of high-speed systems. His current interests include signals and systems in general and efficient simulation and optimization of deterministic and stochastic systems in high-speed links in particular.



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

He is a Principal Design Engineer at Rambus Inc. His current interests include high-speed and low-power digital design.



**Reza Navid** (M'05) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Tehran, Sharif University of Technology, and Stanford University, in 1996, 1998, and 2005, respectively.

From 1998 to 2000, he was with ParsElectric MFG Corporation, Tehran, Iran, where he worked on RF circuits for TV tuners. From January 2000 to May 2001, he was with the University of Michigan, Ann Arbor, as a researcher working on MEMS for wireless communications. He spent the summer of 2001 with Maxim Integrated Products, Hillsboro, OR, where he worked on a 3.125 Gbit/s high-speed communication receiver front-end. From 2005 to 2008, he was with Rambus Inc., Los Altos, CA, working on high-speed integrated circuits for chip-to-chip and memory link applications. He is now with True Circuits Inc., Los Altos, CA. His current research interests include low-phase-noise timing circuits for various communication and data conversion systems.



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

In 2000 and 2001, he was a visiting researcher at 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., Los Altos, CA, where he works on high-speed memory interfaces.



**Fred S. Lee** (M'07) received the B.S., M.Eng., and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, in 2002 and 2007. His doctoral thesis was on "Energy Efficient Ultra-Wideband Radio Transceiver Architectures and Receiver Circuits".

He is currently with SiTime, Sunnyvale, CA, developing quartz-less timing references and mixed-signal/RF chips based on MEMS resonators. From 2007 to 2008, he was with Rambus Inc., Los Altos, CA, focusing on mixed-signal multi-GHz wireline and 60 GHz wireless transceiver circuits and systems.

While at MIT, Dr. Lee received many awards with his colleagues, including first place at the Autonomous LEGO Robot Competition in 2001, the ISLPED Low Power Design Contest in 2002, the DAC/ISSCC Student Design Contest Award in 2004, and the Jack Kilby Best Student Paper Award at ISSCC in 2007.



**Fredy Quan** received the B.S. degree from Texas A&M University in 1994 and the M.S. degree from North Carolina State University in 1996, both in electrical engineering.

He is a Senior Member of Technical Staff II at Rambus Inc. From 1996 to 2006, he was with Cadence Design Systems, Cary, NC. His current interests include high-speed memory interface and digital design methodology and automation.



**Jared Zerbe** (M'90) graduated from Stanford University, Stanford, CA, in 1987. From 1987 to 1992, he worked at VLSI Technology and MIPS Computer Systems. In 1992, he joined Rambus Inc. where he has since specialized in the design of high-speed I/O, PLL/DLL clock-recovery, and equalization and data-synchronization circuits. He has authored multiple papers and patents in the area of high-speed data transmission and clocking and has taught courses at Berkeley and Stanford in high-speed link design. He is currently a Technical

Director at Rambus where he is focused on development of future signaling technologies.



**Rich Perego** received the B.S. degree in electrical engineering with highest honors from Georgia Institute of Technology, Atlanta, GA, in 1989.

Since joining Rambus Inc. in 1999, he was the principal architect for the XDR memory system and the Terabyte Bandwidth Initiative. He is currently a Technical Director and Distinguished Inventor at Rambus and holds 49 U.S. patents. Prior to joining Rambus, he held key technical positions at 3dfx Interactive, Cyrix Corporation, AST Research, and Hughes Aircraft.

He is the recipient of the 2008 DesignCon DesignVision award, EDN Magazine's 2005 Innovation of the Year in the Digital SOC IP category, and was a finalist for EDN Magazine's 2005 Innovator of the Year.



**Fariborz (Fari) Assaderaghi** (M'93–SM'01) received the B.S. degree in electrical engineering from SDSU, graduating *summa cum laude* and University Valedictorian in 1989. After an internship with AT&T Bell Labs, he attended UC Berkeley where he received the M.S. and Ph.D. degrees in electrical engineering in 1992 and 1994, respectively.

After graduation, he joined HP Labs, Palo Alto, CA. From 1995 to 2001, he was with IBM T. J. Watson Research Center in New York where he co-developed the first commercial version of CMOS

SOI for VLSI. The technology became the workhorse for IBM's high-end microprocessors and was used in Apple's processors, as well as Sony Playsta-

tion-3 Cell processor. He was the technology director for Silicon Wave Inc., a San Diego based company developing Bluetooth products, from 2001 to 2003. Silicon Wave was later acquired by RF Micro Devices Corporation. From 2003 till 2008, he was with Rambus, Los Altos, CA, most recently as the Senior Director of Engineering. His groups were responsible for developing very high speed chip-to-chip data communication circuits, including all high-speed links for the PS3 Cell processor, TI DLP, and the Terabyte/sec initiative. Since February 2008, he has been Vice President of Engineering at SiTime, a fabless integrated circuit company developing silicon timing, clock, and RF chips, which incorporate MEMS timing reference devices inside standard silicon electronic chips. He has contributed to two book chapters and to more than 80 technical papers. He holds 47 patents, with more than 20 pending.