A 16-Gb/s Backplane Transceiver With 12-Tap Current Integrating DFE and Dynamic Adaptation of Voltage Offset and Timing Drifts in 45-nm SOI CMOS Technology

Gautam R. Gangasani, Chun-Ming Hsu, *Senior Member, IEEE*, John F. Bulzacchelli, *Member, IEEE*, Sergey Rylov, *Member, IEEE*, Troy Beukema, David Freitas, William Kelly, Michael Shannon, Jieming Qi, Hui H. Xu, Joseph Natonio, Todd Rasmus, Jong-Ru Guo, Michael Wielgos, Jon Garlett, Michael A. Sorna, and Mounir Meghelli

***Abstract*—This paper presents a 16-Gb/s 45-nm SOI CMOS transceiver for multi-standard backplane applications. The receiver uses a 12-tap DFE with circuit refinements for supporting higher data rates. Both the receiver and the transmitter use dynamic adaptation to combat parameter drift due to changing supply voltage and temperature. A 3-tap FFE is included in the source-series-terminated driver. The combination of DFE and FFE permits error-free NRZ signaling at 16 Gb/s over channels exceeding 30 dB loss. The 8-port core with two PLLs is fully characterized for 16GFC and consumes 385 m W/link.**

***Index Terms*—DFE, duty cycle correction, dynamic adaptation, FFE, phase rotator linearity, serial links.**

I. Introduction

THE demand for higher bandwidth data transmission systems has brought transceiver design into >10 Gb/s data rates. Although transmission with advanced equalization techniques at >10 Gb/s has been demonstrated [1], [2], achieving a robust system at these speeds remains challenging. This paper describes a transceiver which dynamically adapts for several types of drifts during system operation and also features circuit innovations to improve performance at higher data rates.

The receiver decision feedback equalizer (RX DFE) can improve signal-to-noise ratio of the received signal over that achieved by linear equalizers [3]. However, DFE power dissipation reduces link energy efficiency. One way of reducing DFE power consumption is to use current integrating summers [4]. This paper presents improvements to the current integrating summer, which enable it to perform at higher data rates.

Offset calibration is needed in order to achieve high RX sensitivity. If this is done only at startup, then accuracy of the compensation is affected by subsequent supply voltage and temperature drift. Earlier architectures [3], [4] did offset calibration at startup but did not dynamically adapt for parameter drift during system operation. Fig. 1(a) shows the voltage offset drift in a receiver datapath after the supply and temperature have changed to the opposite extreme of those used during startup calibration. Fig. 1(b) fits a Gaussian curve to the simulated data, which is used to estimate the worst case drift with 3σ probability. With 3 mV/lsb, this drift could be around 25 mV, which would affect the receiver robustness by degrading its sensitivity. A periodic or dynamic calibration of the offset would track any drift due to supply and temperature variation. One structure [5] which permits dynamic calibration of drift in latches is the structure shown in Fig. 2. While the MUXes permit each latch to be swapped out for dynamic calibration, the extra MUX delay compromises critical timing paths of the DFE.

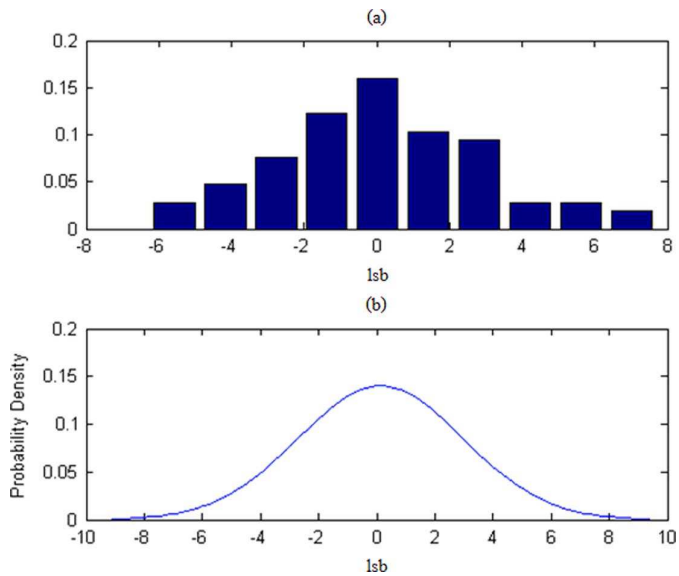

Fig. 1. (a) Simulated histogram of voltage offset drift after initial calibration upon startup. (b) Gaussian curve fitted to the simulated data.

This paper demonstrates a solution for dynamically adapting the offset in each of the DFE data paths, which does not affect critical DFE feedback loop timing. It also shows improvements made to phase rotators (feeding DFE clocks), which improve phase linearity, quadrature error and duty cycle in the face of

Manuscript received November 21, 2011; revised February 19, 2012; accepted March 28, 2012. Date of current version July 19, 2012. This paper was approved by Guest Editor Alvin Loke.

The authors are with the IBM Systems and Technology Group, Hopewell Junction, NY 12533 USA and with the IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: gautamg@us.ibm.com)

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

Digital Object Identifier 10.1109/JSSC.2012.2196313

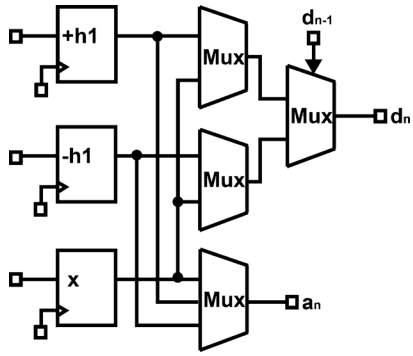

Fig. 2. MUXing structure of [5] which allows latches to be swapped out for dynamic calibration.

parameter drifts. The transmitter (TX) duty cycle is also calibrated dynamically during operation. The paper is organized as follows. Section II presents the architecture. Section III reveals the circuit refinements to improve DFE accuracy at high data rates, as well as different techniques to dynamically adapt for offset and timing drifts. Experimental results are described in Section IV, and Section V concludes with a summary.

II. Architecture

Fig. 3 shows the top-level architecture of the transceiver. The TX is similar to the one described in [6] and uses a source-series-terminated (SST) driver with 3-tap FFE. The RX input data path is similar to the one in [4], except that series inductive peaking is used to extend the bandwidth. The data path is implemented using body-contacted devices to avoid pattern dependent history effect in floating body devices [7]. While the high impedance of the body contact adds a small amount of ISI to the data path (with a duration of around a nanosecond), this ISI term is largely compensated by the DFE. A 12-tap DFE is chosen to set the equalization range at higher baud rates and to clear near reflections.

The DFE uses a half-rate architecture with single speculative tap. In [4], each half of the DFE had dedicated paths to detect data and to monitor the vertical and horizontal eye. Since the path detecting data bits and the path monitoring the eye could not be swapped in architecture [4], offset drift could not be dynamically adapted. The architecture in Fig. 3 has two banks in each half, namely "Bank A" and "Bank B". This architecture has the ability to swap any bank to be the one detecting data, thereby giving it the ability to dynamically adapt every path in the DFE for voltage offset drift.

Other system innovations highlighted in Fig. 3 and detailed later in the paper compensate for timing drifts. Among these are dynamic quadrature clock correction (DQCC) system to correct rotator duty cycle and I-Q separation, and duty cycle correction (DCC) to maintain TX duty cycle close to 50%.

III. Circuit and System Level Innovations

A. S/H and Current Integrating Summer for Higher Data Rates

Fig. 4 shows the circuit improvements in the current integrating summer. In the circuit described in [4], the kickback from the S/H can disturb the peak amplifier output, thereby affecting the sampled voltage in other paths. The isolation buffer in Fig. 4 is used to shield this kickback. Cross-coupled passgates (p3, p4) remove differential feed-through from the passgates (p1, p2) in their OFF state. A negative-C cell is added at the output of the isolation buffer to speed up its recovery from S/H kickback.

In [4], the summer kickback from the integrating nodes to the hold nodes led to gain loss. To avoid differential kickback and gain loss, a pair of neutralization capacitors (n1, n2) are added. The pfet reset switch (r1) is used to reduce differential reset errors at high speed. As described in [4], the summer calibrates its current dynamically. But, as g_m of the input differential pair does not scale proportionally to calibrated current, it leads to higher gain and poorer linearity at lower speeds. To mitigate this issue, this design uses switched load capacitors at the summer output to minimize current variation as a function of speed [8]. For higher sensitivity, a CML buffer is added at the summer output to stabilize the common mode presented to the master latch (ML) and to achieve extra gain. This CML buffer does not affect the timing margins per se, but reduces the effective summer integration window thereby increasing the calibrated summer current proportionally.

While use of current integrating summers removes RC settling time of resistive summers (e.g., [3]), feedback tap weights could have integration error (Fig. 5(a)) which depends on when the feedback tap arrives relative to the start of integration (i.e., setup time). As shown in Fig. 5(b), the tap structure used in [4] (also used for h6-h12 in Fig. 4) entails a significant setup time (20–30 ps) to meet the integration error target of 4%, which is problematic at high data rates. The cause of this setup time is the glitch on the tap tail node, created when the DFE feedback signal switches, and its slow decay time. In this design, a return-to-zero (RZ) structure is introduced in which clocked differential pairs steer current away from hp/hn differential pairs during reset. The RZ structure makes the glitch on the tail node independent of data pattern, effectively removing positive setup time requirement on feedback tap, as shown in Fig. 5(c). This structure is implemented for h2–h5 taps where timing margins are tight.

B. Dynamic Adaptation of Offset and Timing Drifts

1) Dynamic Compensation of Offset Drift: The bank swapping architecture is detailed in Fig. 6 wherein "Bank A" can be used as a data path or as a path to measure eye metrics when "Bank B" is being used as the data path. This swapping allows the system to check for vertical eye symmetry in each path dynamically, hence adapting for offset drift. An advantage over the architecture of [5] is that no extra delay is added to the path wherein the previous bit selects the speculated data path. The clocks (Ca, Cb) to the different banks move independently of each other, so as to allow horizontal scan of the eye. After bank swap, both data and clock to the DFE pipeline (shift register) need to be swapped, using the 2:2 data and clock MUXes. The clocks to "Bank A" and "Bank B" are positioned optimally based on horizontal eye scans. Therefore, the system is not sensitive to the phase errors in the phase detection paths,

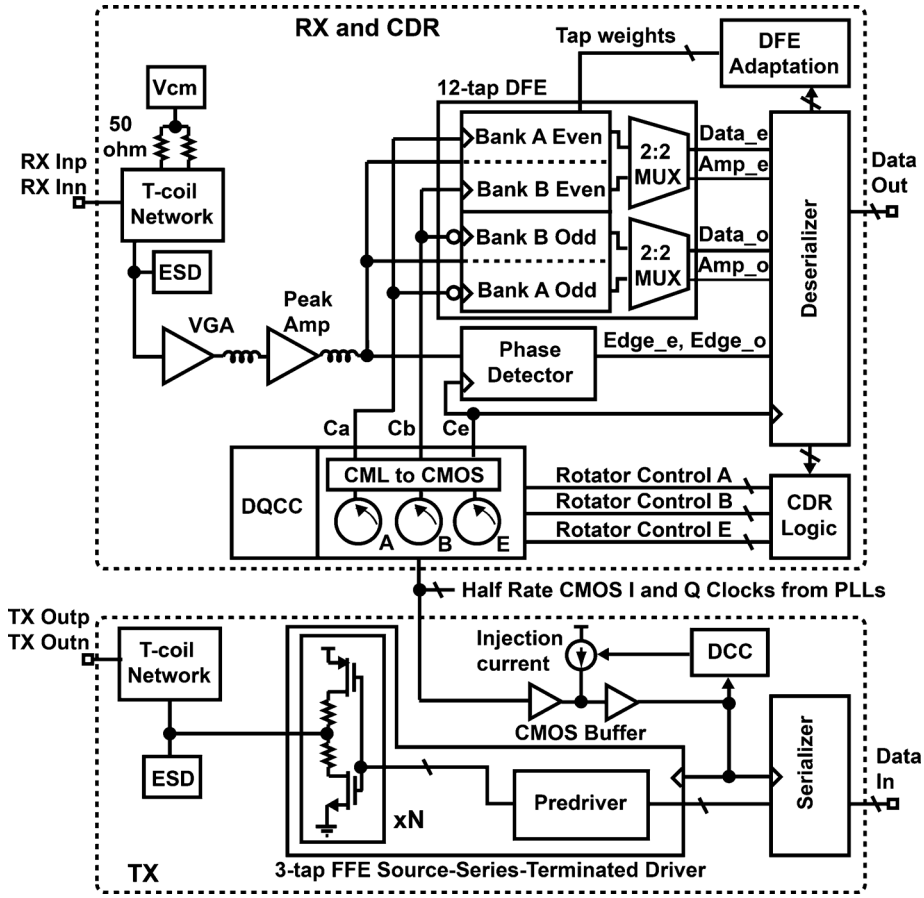

Fig. 3. Top-level transceiver architecture.

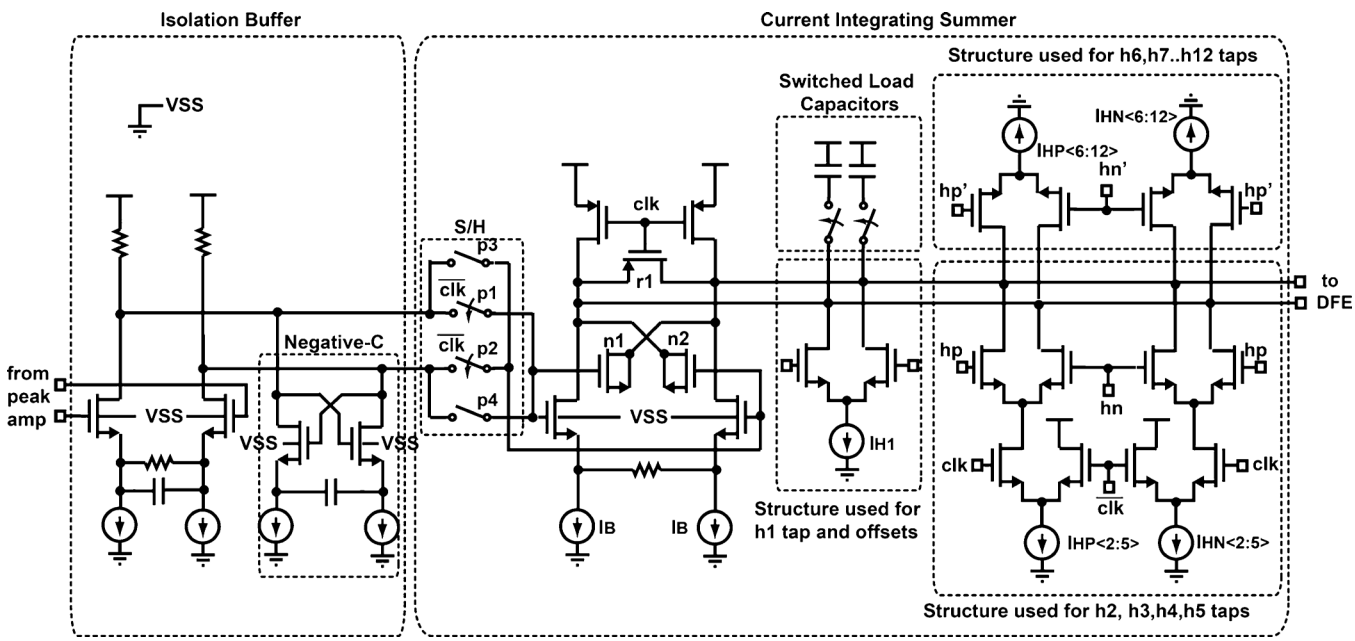

Fig. 4. Schematic of the current integrating summer with isolation buffer and S/H.

for which dynamic offset correction is not performed. "Bank A" and "Bank B", including their clocks, are simultaneously active for $<2\%$ of the time, thereby not affecting the power efficiency of the DFE substantially.

A key improvement to the DFE architecture in [4] significantly increases the sweep range of the clocks without stressing the DFE feedback timing. When "Bank A" clock is being swept late with respect to the "Bank B" clock to measure the eye width, and if its summer were to be getting dynamic feedback tap data from "Bank B", then the feedback tap transitions could occur before the "Bank A" summer integration has completed (hold violation), as shown in Fig. 7. To overcome this issue, each of the

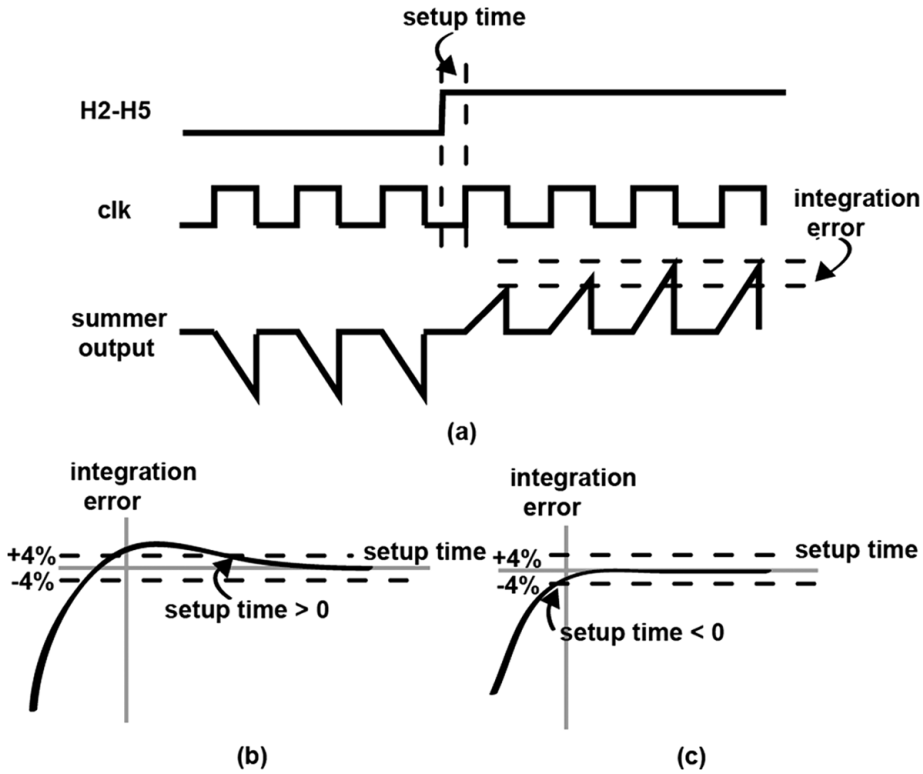

Fig. 5. (a) Tap weight integration error that occurs when feedback tap signals arrive at summer with insufficient setup time. (b) Integration error as a function of setup time without return-to-zero structure. (c) Integration error as a function of setup time with return-to-zero structure.

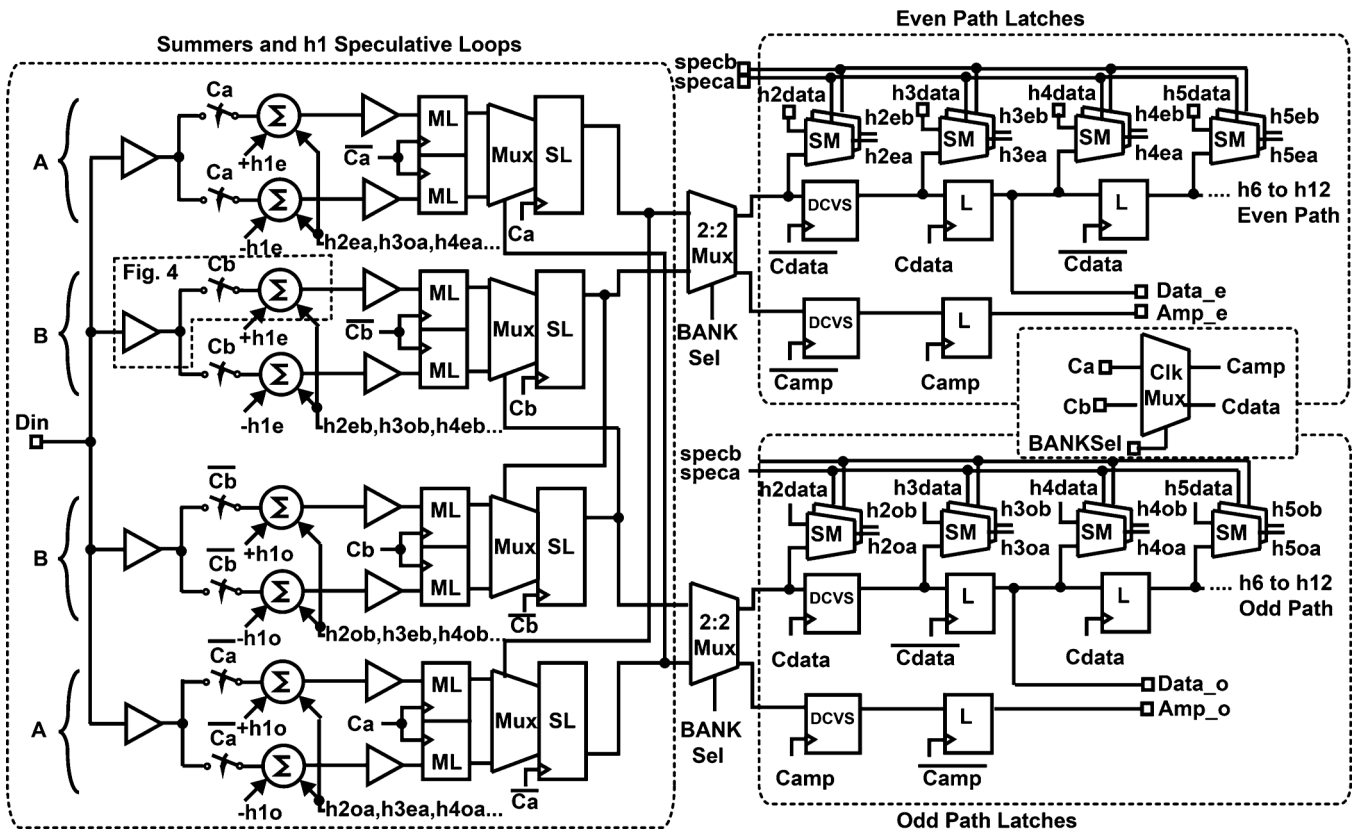

Fig. 6. 12-tap DFE architecture with bank swapping.

feedback taps has a speculative MUX (SM) in their paths, giving the control logic the ability to break the dynamic feedback path and insert static feedback bits. For instance, when "Bank A" is in horizontal eye scan mode, the dynamic feedback path to its summers is broken to avoid hold violations, and all active taps in "Bank A" are applied speculatively in the control logic.

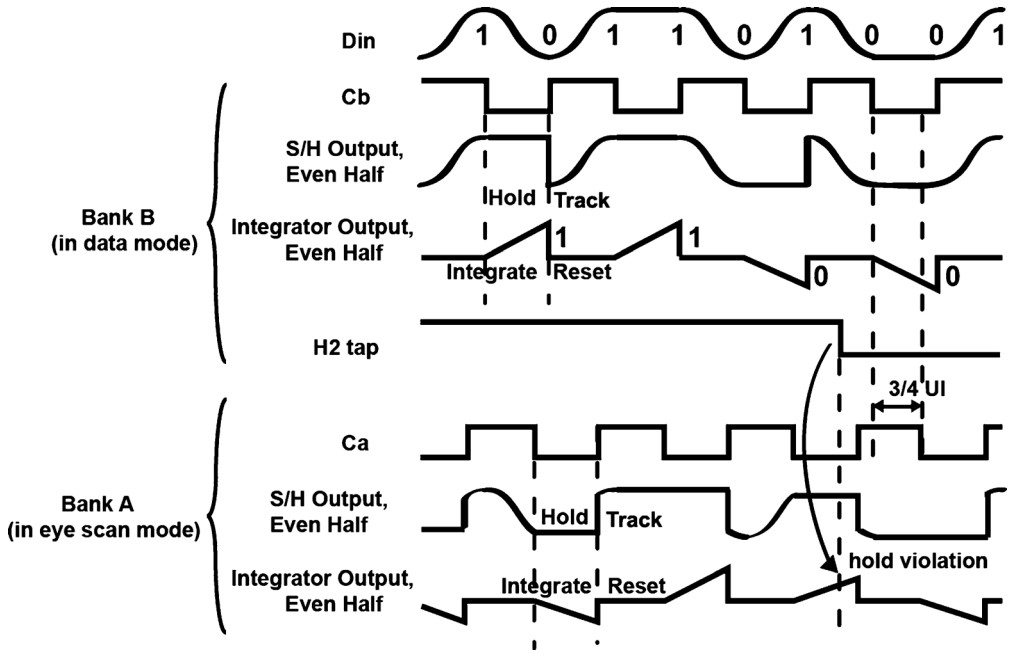

Fig. 7. Example of hold violation on feedback tap with Ca clock swept 0.75 UI later than Cb clock.

2) Octagonal Phase Rotator With Dynamic Quadrature Clock Correction (DQCC): Although the architecture of the CDR is essentially the same as that in [3], the innovations in this work include: (a) a phase rotator with an octagonal phase constellation ("octagonal rotator") to improve phase linearity and reduce the amplitude modulation error, (b) a two-stage clock-conditioning buffer with open-loop cancellation of both quadrature phase separation error and duty-cycle-distortion (DCD), (c) a dynamic-quadrature-clock-correction (DQCC) loop to finely compensate the residual quadrature error and DCD in the four-phase clock that feeds the rotator mixer(s).

In principle, a phase rotator with round constellation can be implemented. However, the need for good phase linearity and minor amplitude modulation often demands for higher circuit complexity. For example, although the work in [11] intends to realize 32-state round constellation, two 5-bit DACs are necessary to control the I and Q weights, respectively, while the realized constellation is only approximately an octagonal shape. Instead, in this paper, we aim at implementing the octagonal constellation with less complexity.

Our earlier phase rotator architecture [3] with a diamond-shaped phase constellation is shown in Fig. 8(a). The proposed octagonal phase rotator, as shown in Fig. 8(b), is derived from the former one. Both rotator architectures are based on a conventional four-quadrant CML mixer that interpolates between four quadrature clock phases (IP, IN, QP, QN) and is driven by four tail currents from one or more current-steering DACs (the latter are different in the two architectures). Specifically, the old architecture used just a single current steering DAC (shown in Fig. 9) that was connected to the mixer via two polarity switches to select which two out of four mixer tails are driven. This means that in the old architecture the non-selected tails in the mixer were fully off while the sum of interpolating weights applied to the remaining two tails was a constant thus yielding a characteristic diamond-shaped phase constellation, as shown in Fig. 10(a). In contrast, the new architecture uses three current-steering DACs: an IQ-DAC with two switches (similar to the old one) and two new ones (I-DAC and Q-DAC) that apply/steer their interpolation weights only within the IP, IN and QP, QN phase pairs respectively. The new I and Q DACs, when steered (one at a time), create vertical and horizontal segments in the phase constellation and thus, in combination with the 45° slanted segments from the IQ DAC, give it an overall octagonal shape as shown in Fig. 10(b). An immediate advantage of using an octagonal constellation is dramatic reduction of its amplitude modulation (AM) at different phase angles, from 3 dB in the diamond constellation to <1 dB in the octagonal one. This is important because large AM may lead to a phase error from AM-PM conversion in the latter gain stages. A more significant advantage of the octagonal rotator is its fundamentally better interpolation linearity that arises from avoidance of generating small interpolation weights with proportionally small tail currents. For example, instead of giving a clock phase I a small positive interpolation weight by sending a small tail current to the phase IP (and none to IN) we apply a pair of relatively large tail currents (with the desired positive difference) to both IP and IN. Because this approach causes much more stable biasing conditions of the interpolator quadrants, it minimizes their errors in both gain and propagation delay (that become the largest at very low bias) and as a consequence, the octagonal rotator in Fig. 8(b) works optimally with uniformly sized DAC elements in each of its three DACs. In contrast, our single-DAC rotator [3] required significantly non-uniform sizing of its DAC elements for best rotator linearity leading to much more complicated design.

The rotator has 64 phase states uniformly distributed on a phase circle, i.e., 16 states per quadrant. Each single-LSB step (i.e., transition between any two neighbor states) is achieved by changing the state of just one current-steering element in

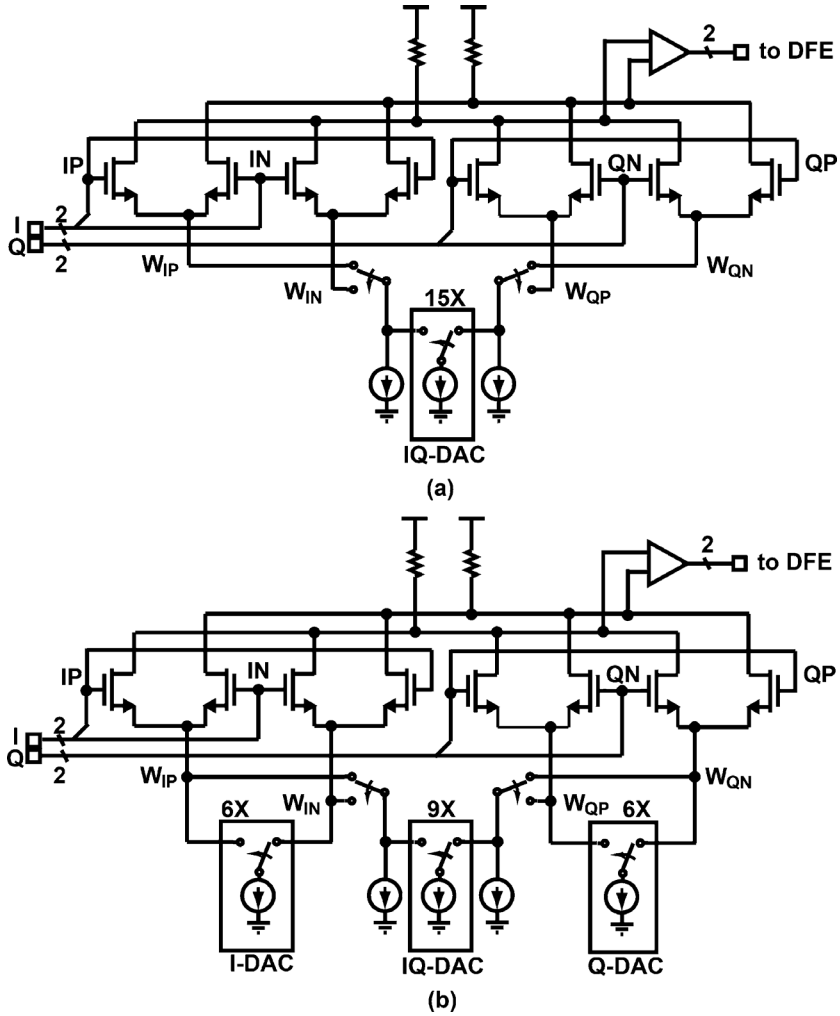

Fig. 8. Simplified schematics of (a) rotator with diamond-shaped constellation and (b) rotator with octagonal-shaped constellation.

one of the three DACs or by flipping one of the two polarity switches in the IQ-DAC (the latter happens only at the quadrant boundaries). By doing so, the glitch in phase is minimized while the rotator position is switching. In contrast, the implementation in [11] suffers a more severe glitch problem because of the utilization of the binary DACs, in which multiple current elements are switching each time. The IQ-DAC has nine identical current-steering sections of unity size that steer current between its two outputs, plus two fixed sections of half-size permanently connected to the two outputs. Therefore, such DAC has 10 different output states with the output currents (9.5, 0.5), (8.5, 1.5),. . .(0.5, 9.5). The remaining two DACs are identical, each having 6 identical current-steering sections of half-unity size yielding seven output states: (3, 0), (2.5, 0.5),. . .(0, 3). Since these outputs are applied to the two opposite clock phases (e.g., IP and IN), they can be interpreted as seven effective weight values (integers from $+3$ to -3) applied to a single phase I. i.e., $3-0=3, 2.5-0.5=2, \ldots 0-3=-3$. By combining I-DAC, Q-DAC and IQ-DAC (with the related switches) one can create the octagonal constellation shown in Fig. 10(b). Since 21 $(9+6+6)$ current elements are used in total, this implementation is equivalent to roughly 4.5 bits, compared to the 6-bit implementation (two 5-bit DACs) in [11]. Its main operation principle is that the states within the vertical or horizontal segments (for example, vertical segment near phase $+$I) are obtained by applying the maximum possible weight to the dominant phase (in this case IP) using both IQ-DAC (gives IP weight of 9.5) and I-DAC (gives IP weight of $+3$) for a total IP weight of 12.5 units, while giving the opposite phase (IN) zero weight. All steering therefore is done with the two remaining phases QP, QN that both receive variable, non-zero weights from the Q-DAC, plus an extra 0.5 units of weight from the IQ-DAC that is applied to just one of them (which one depends on Q weight polarity). Overall, the non-dominant Q phase receives effective weights ranging from 3.5 to -3.5 in unity increments. The operation rule for the sloped segments (e.g., the one between $+$I and $-$Q) is first giving the maximum weight to the two dominant clock phases (IP, QN) from the I-DAC and Q-DAC respectively (3 units each) and then using IQ-DAC to steer its available weight between the same two phases. This causes overall weights applied to phases IP, QN to take 10 different states: (12.5, 3.5), (11.5, 4.5),. . .(3.5, 12.5) that form a 45° sloped segment. Since none of these weights is small, there is no concern about using zero weights on the remaining two phases IN, QP. Note that in the state of (12.5, -0.5) (Fig. 10(b)), the small weight of -0.5 is achieved by the difference of $\mathrm{W_{QP}}=1.5$ and $\mathrm{W_{QN}}=2$, so

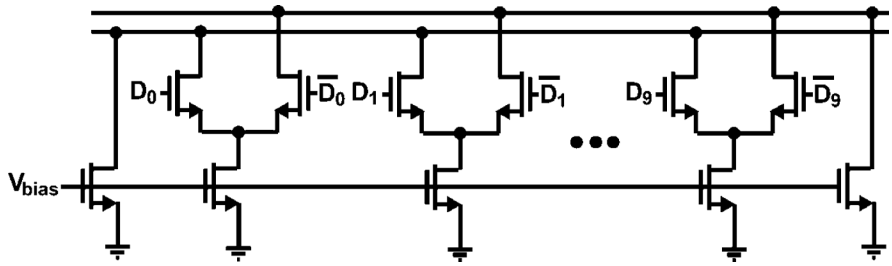

Fig. 9. Current-steering IQ DAC.

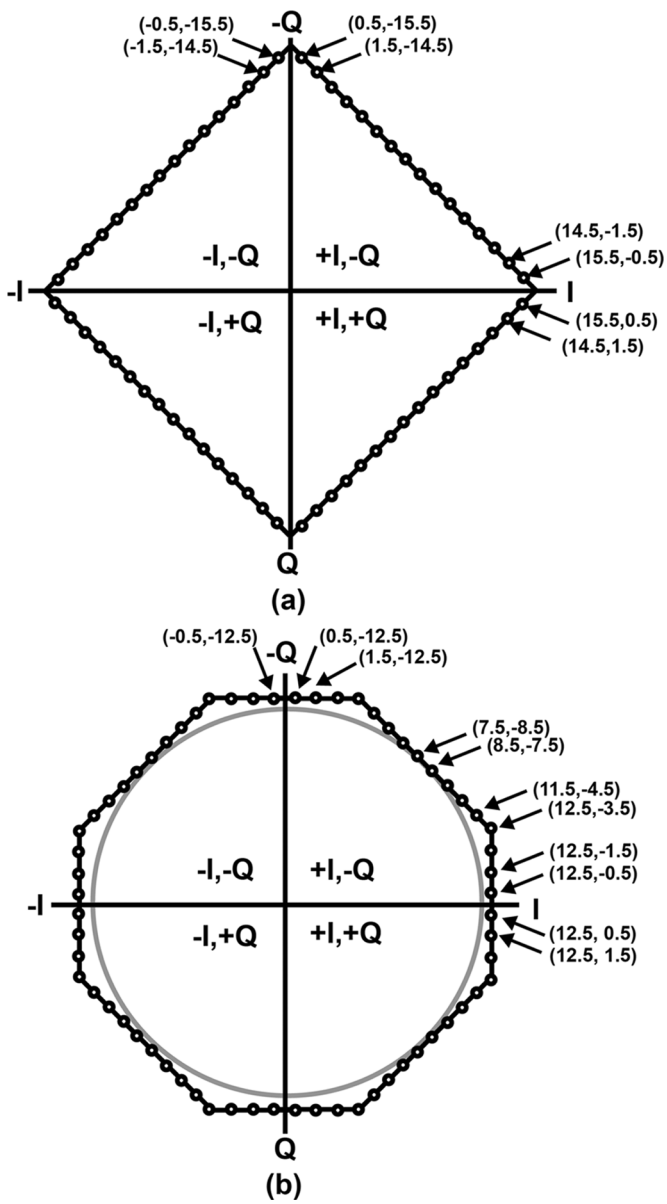

Fig. 10. Phase constellation of (a) diamond rotator and (b) octagonal rotator.

there is no small current used compared to the diamond-shape rotator, in which a weight of 0.5 is necessary in the state of (15.5, -0.5) (Fig. 10(a)). In the state of (12.5, -2.5), $\mathrm{W_{QP}} = 0.5$ and $\mathrm{W_{QN}} = 3$. Although a small weight of 0.5 is used in this state, the overall delay of the Q clock is dominated by $\mathrm{W_{QN}} = 3$, so using a small weight in this state is less of a concern in the octagonal rotator.

Fig. 11 illustrates the complete rotator used in this chip, including the clock-conditioning buffer for slew-rate control and dynamic quadrature clock correction described below. The quadrature clocks (3.5–8 GHz) shared among eight CDRs are distributed from two central PLLs. Two PLLs are built in this chip to allow different receivers to support different data rates simultaneously. Instead of CML clocks, CMOS clock distribution from PLLs is utilized since its power dissipation automatically scales down with frequency. A CMOS mux first selects a set of clocks, and a CMOS quadrature divider can extend the frequency range down to 1.75 GHz. Two CML 2:1 muxs then select either the original received clocks or the divided-by-2 version for the following clock-conditioning buffer, which is used to mitigate the possible sharp clock edges in fast process corners and low-frequency operation. The CML rotator outputs are buffered and converted back to CMOS levels before being distributed to DFE.

A linear phase interpolator needs its incoming clock phases to have overlapping edges in order to obtain uniform phase steps, so in the case of quadrature clock phases the clock therefore must be slewing at least 50% of the time. Therefore, a clock-conditioning buffer is necessary to ensure sufficiently slow clock edges across multiple operating conditions. While one can achieve slower edges by varying the power level of the clock buffers in the clock chain, an important alternative technique for conditioning of quadrature clock sets is to perform a fixed-weight clock interpolation $\mathrm{I''} = \mathrm{I} + \mathrm{Q}$ and $\mathrm{Q''} = \mathrm{Q} - \mathrm{I}$ that corresponds to a 45° phase rotation, as illustrated in Fig. 12(a) [9]. While this technique by itself does not directly produce the desired uniformly sloped edges and requires use of other bandwidth-limiting mechanisms; it does contribute positively to the task of optimal low-pass filtering of the clock waveform since it is equivalent to a single-tap low-pass transversal filter with a time constant of 1/4 clock period. The primary importance of this technique, however, comes from its ability to provide highly efficient open-loop correction of the quadrature phase separation error [9]. As shown in Fig. 12(b), assuming a linear operation, any quadrature phase separation at the input of this stage will be completely removed. Nevertheless, a single-stage implementation of this technique suffers from two drawbacks: a tendency to produce a flat spot in the middle of the clock transitions (fast corners), as shown in Fig. 12(c), and deterministic conversion of the quadrature phase mismatch into a comparable amplitude mismatch (slow

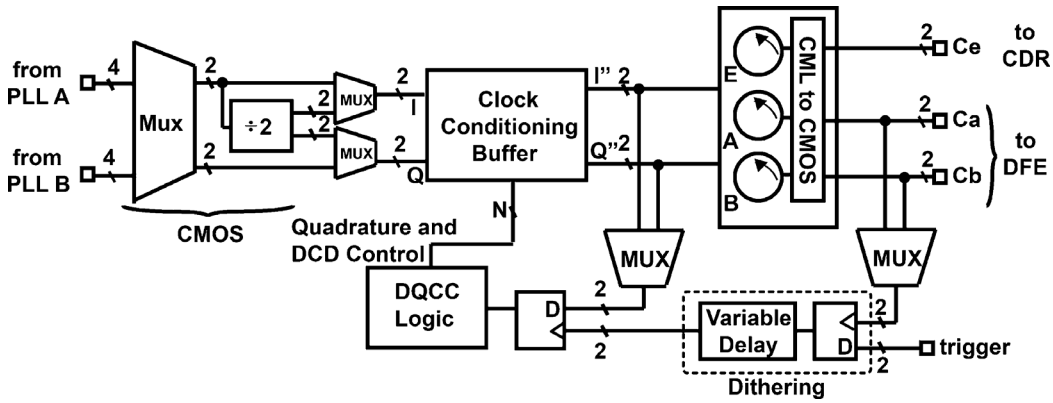

Fig. 11. Block diagram of complete rotator macro including DQCC circuits.

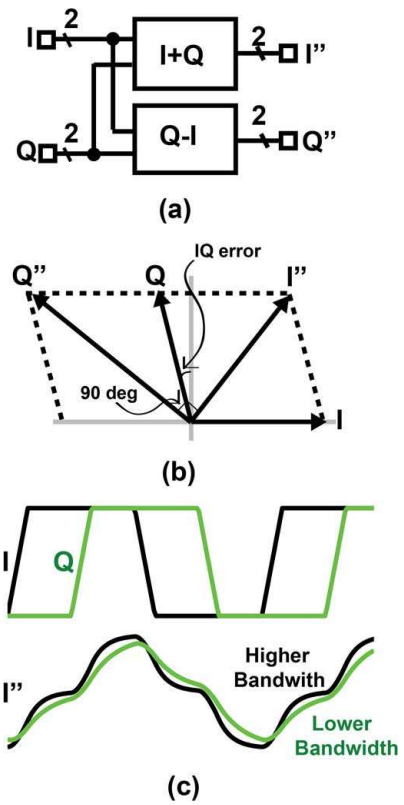

Fig. 12. (a) Functional block diagram of single-stage clock conditioning buffer. (b) Vector diagram of its input and output clocks. (c) Input and output transient waveforms.

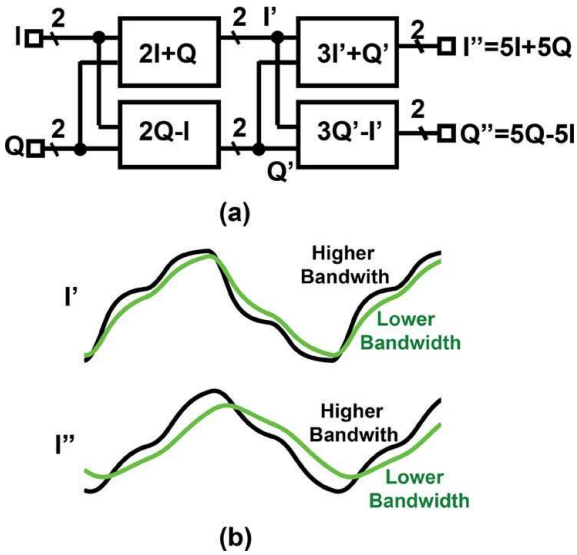

Fig. 13. (a) Functional block diagram of two-stage clock conditioning buffer. (b) Transient waveforms at outputs of first and second stages.

corners), as illustrated in Fig. 12(b). The former flat spot problem can be alleviated by bandwidth limiting the circuits as illustrated in Fig. 12(c), in which the example waveforms are obtained from behavior simulation with bandwidth of 3 and 1.5 times of the clock frequency. The latter mismatch, if directly applied to the main rotator interpolator, would lead to a rotator phase error with the same magnitude as if no quadrature error correction were performed. One way to mitigate substantially both of these problems is to substitute a single 45° fixed phase rotation stage with a cascade of two similar stages, as shown in Fig. 13(a), where each implements a part of overall rotation (26.6° + 18.4° in our implementation). While in the linear approximation the single-stage and the multiple-stage implementations produce exactly the same output, in the real circuits that have saturation behavior, single and dual stage circuits do differ significantly. For example, the amplitude mismatch created by the first stage can be attenuated substantially by saturation in the second one. Similarly, the waveform shape of the clock after a bandwidth-limited two-stage slew buffer no longer has a flat spot in the middle of the transition, as shown in Fig. 13(b) and has a general tendency to approximate a triangular wave, i.e., to slew well over 50% of the clock cycle. Notice that by setting the bandwidth in the one-stage and two-stage approaches to the same value, 1.5 times of the clock frequency in the lower bandwidth case, the two-stage approach provides better waveform at I". To achieve similar waveform with the one-stage approach, one should lower the bandwidth further, which attenuates the amplitude at I″. This two-stage circuit uses the first stage with a pair of 2:1 fixed interpolators (26.6° rotation) followed by a second stage with 3:1 fixed interpolators (18.4° rotation) for a total of 45°. An additional feature of this circuit is use of capacitively degenerated tails for blocking the dc component of the input waveforms that

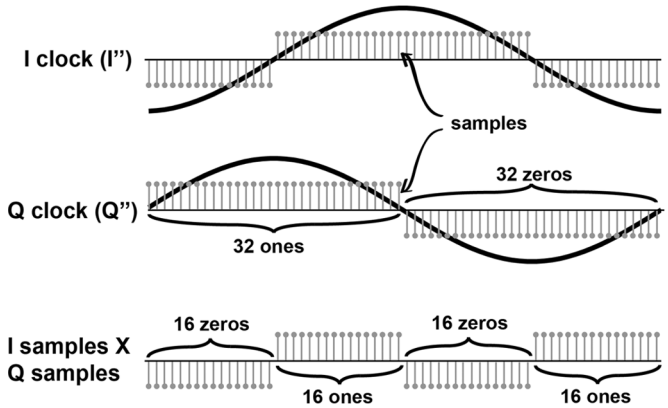

Fig. 14. Operating principle of DQCC. The DQCC circuit takes samples of the I and Q clocks at 64 positions within one clock cycle, from which duty cycle and quadrature information is extracted.

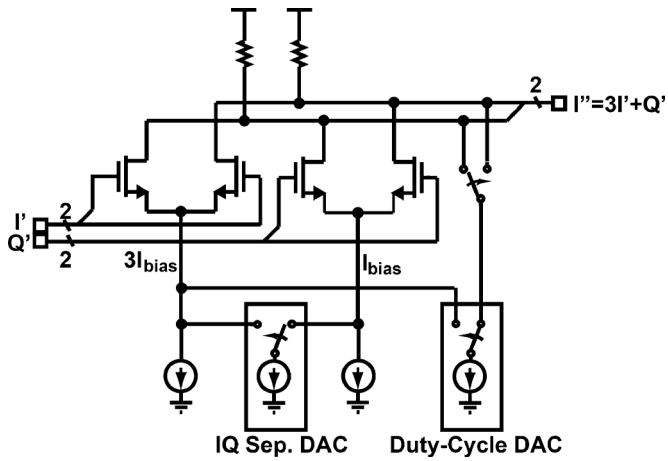

Fig. 15. DQCC actuators inside second stage of clock conditioning buffer.

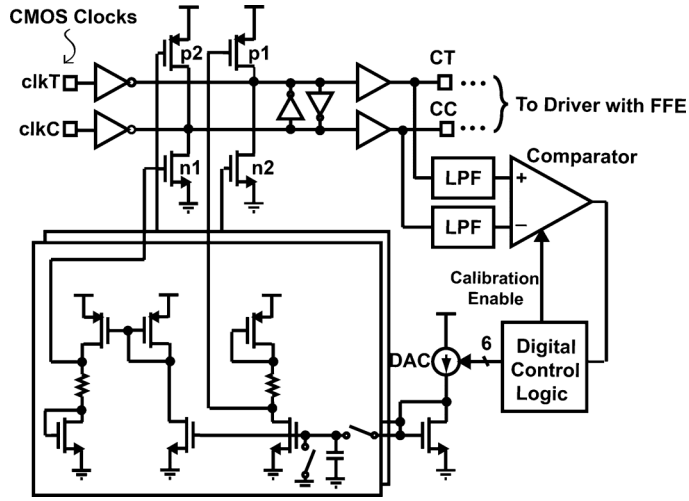

Fig. 16. Digital control loop to correct TX duty cycle.

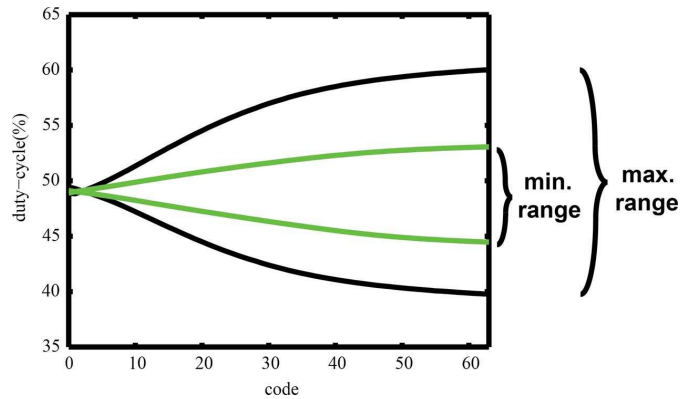

Fig. 17. Simulated maximum and minimum correction range of TX DCC control loop.

partially removes DCD error. The residual error terms that cannot be removed with open-loop techniques are addressed with a closed-loop Dynamic Quadrature Clock Correction (DQCC) circuit.

As shown in Fig. 11, the sensing portion of DQCC circuit employs an offset compensated latch that samples the I and Q clocks at the conditioning buffer output at 64 positions within one clock cycle. The clock used for sampling is supplied from the idle rotator that is not currently used in the receiver data path such that its phase position can be swept freely. More specifically, a lower-speed trigger signal is first synchronized to the rotator edge, passed through a variable delay line, and then serves as the sampling clock. Fig. 14 illustrates the samples when the variable delay is fixed, assuming 50% duty cycles and zero quadrature separation error. In this condition, one should observe $32 + 1$ s and $32 - 1$ s from the 64 I or Q samples. Also, after multiplying the I samples by the corresponding Q samples, one should observe $16 + 1$ s followed by $16 - 1$ s, and so on. Therefore, the algorithm can determine if there are quadrature and DCD errors and compensate them using DAC-driven actuators built into the conditioning buffer. Since 64 samples per clock period cannot provide sufficient accuracy, the 64-step variable delay line is used to dither the sampling clock edge to improve accuracy (Fig. 11). The key is that the delay line step (approximately 1.2 ps) should not be sub-integer of a rotator step (1.95 ps for 16 Gb/s), otherwise the effective resolution degrades. Therefore, no calibration of the delay line is necessary. As a result, there are 64×64 sampling positions within each clock cycle. In addition, we also ensure that the delay between I'' and the sampling latch is more than four rotator steps for de-correlation.

Fig. 15 illustrates the implementation of the DQCC actuators inside the second stage of the clock conditioning buffer. The phase difference between the I and Q clocks can be adjusted by changing the interpolation ratio in this stage using a steer-current DAC. When more current sections are supplied to the I pair on the I path, resulting in mixing ratio of $(3 + \Delta)\mathrm{I}' + (1 - \Delta)\mathrm{Q}'$, more current sections are also supplied to the I pair on the Q path, resulting in mixing ratio of $(3 - \Delta)\mathrm{Q}' - (1 + \Delta)\mathrm{I}'$, such that the I'' clock and Q'' clock are shifted toward opposite directions. Another steer-current DAC injects current to either the positive or negative leg of the outputs to adjust the duty-cycle. The unused current sections are recycled to the primary differential pair to boost the gain.

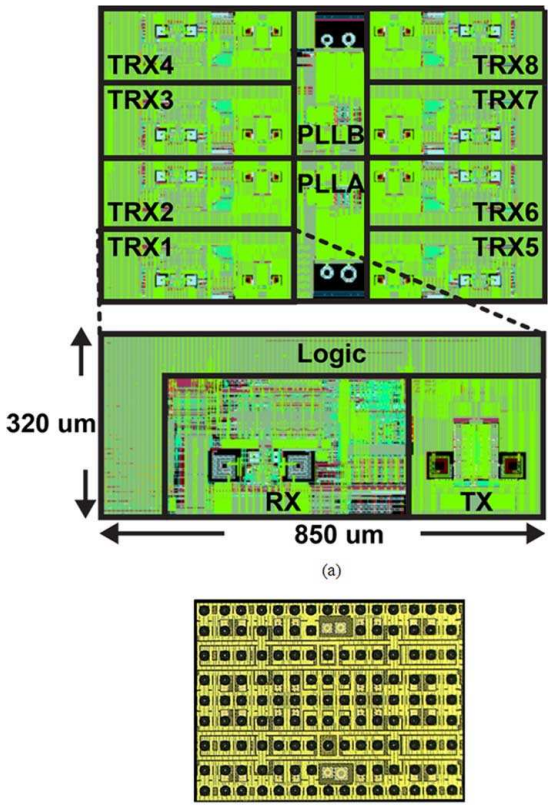

Fig. 18. (a) Layout of 8-port core with two LC VCO-based PLLs and eight transceivers. (b) Fabricated chip.

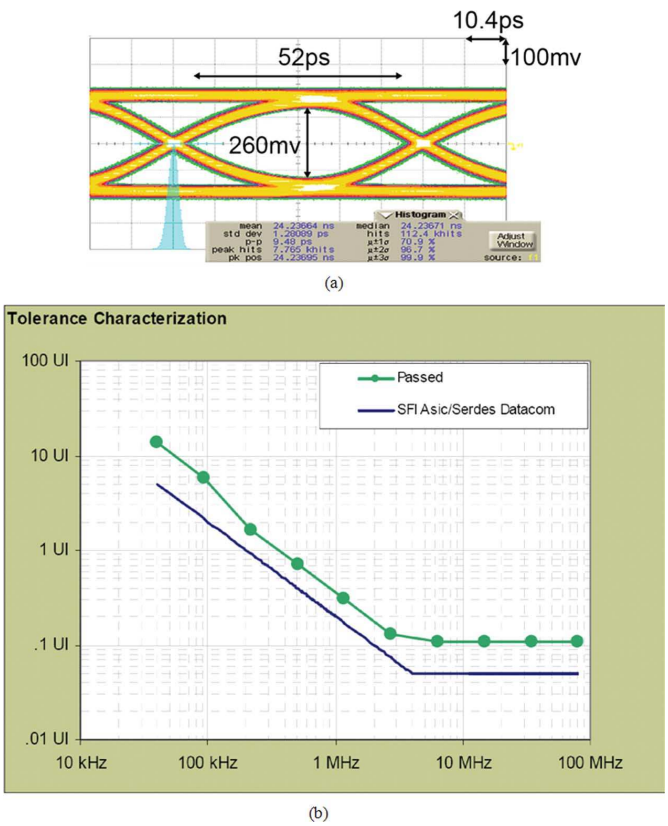

Fig. 19. (a) TX output eye diagram with 16-Gb/s PRBS31 pattern. (b) Rx SJ tolerance from a 16GFC jitter test at 11 Gbps.

3) TX Duty Cycle Correction: The mismatch-induced DCD of the TX changes with supply and temperature. In this design

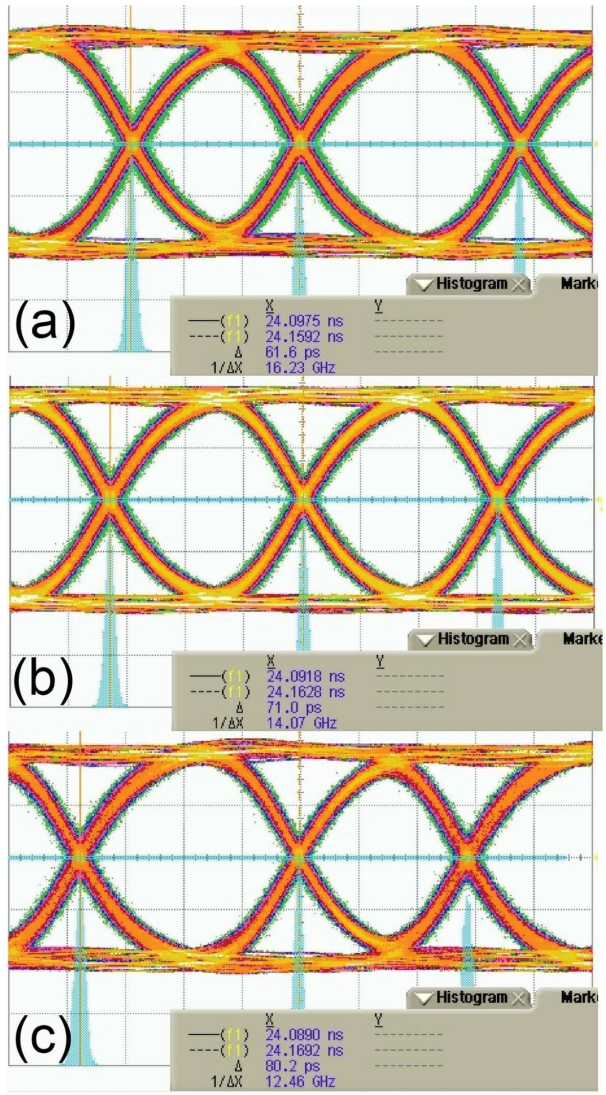

Fig. 20. Measurement of TX DCC operation. (a) Output duty cycle forced to 43% with TX DCC in open loop configuration. (b) Output duty cycle converged to 50% with closed loop running. (c) Output duty cycle forced to 56% with TX DCC in open loop configuration.

the TX DCD is dynamically corrected using a low bandwidth digital control loop. As shown in Fig. 16, the DCD of the complementary CMOS clocks is corrected using differential current sources with programmable gate controls (p1, p2, n1, n2). While the edges of the complementary clocks are aligned by the cross-coupled buffers, the average voltages of the two clocks are sent to an auto-zeroed comparator. The logic controller uses the comparator output to determine the direction (p1/n1 or p2/n2 ON) and amount of current injection.

Fig. 17 shows the minimum and maximum simulated DCD correction ranges over PVT variations for the DCC loop. The range is smallest at fast process corners, but that is not a cause for concern because the duty cycle distortion is also smaller at such corners.

IV. Measurements

An 8-port core with two LC VCO-based PLLs is fabricated in a 45-nm SOI CMOS process. Fig. 18 shows the layout of the

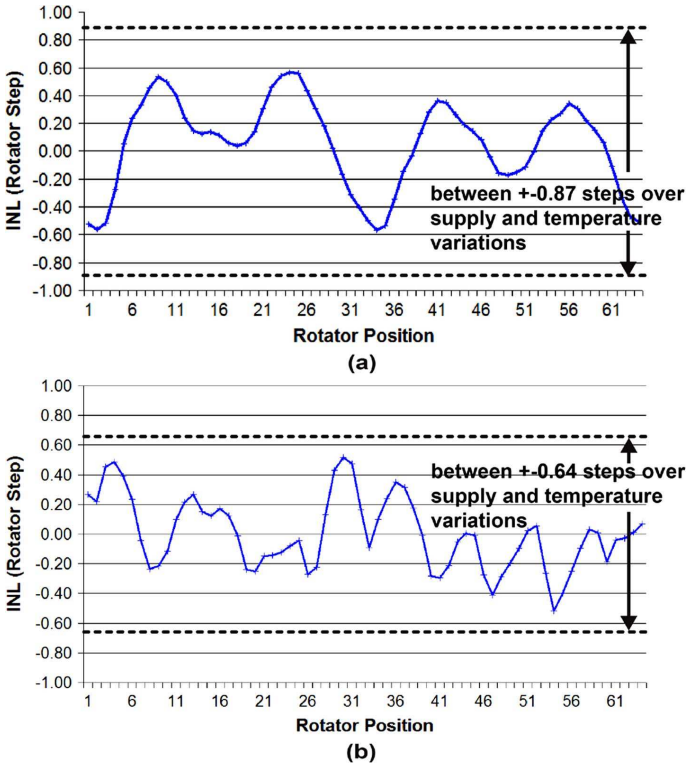

Fig. 21. Measured integral nonlinearity (INL) of phase transfer characteristic for (a) previous diamond rotator and (b) proposed octagonal rotator.

8-port core and the fabricated chip, with each transceiver occupying an area of 850 μm $\times$ 320 μm. Measurements are made to characterize individual transceiver components and to validate the new features in this design. Fig. 19(a) shows the TX output eye at 16 Gb/s; total output jitter is 9.5 $\mathrm{ps_{pp}}$. A full characterization gives the total TX output jitter, with extrapolation down to $\mathrm{BER} < 10^{-12}$, to be 12.69 $\mathrm{ps_{pp}}$. Fig. 19(b) shows the SJ tolerance from a 16GFC jitter test at 11 Gb/s. The TX-DCC loop converges to close to 50% DCD (Fig. 20(b)), while the open loop range is measured to be 43% to 56%, as shown in Fig. 20(a) and (c), well within the simulated bounds shown in Fig. 17. Fig. 21 compares the phase rotator nonlinearity of the proposed octagonal rotator with DQCC (measured with typical hardware at nominal temperature and supply voltage) with that of a previous diamond-shaped rotator, which was implemented in 65 nm bulk CMOS process for a 10.3125 Gb/s link. The INL improves from $\pm$0.57 to $\pm$0.52 rotator steps. An even larger improvement is observed over supply and temperature variations. As indicated by dashed lines in the figure, the worst case INL over supply and temperature variation improves from $\pm$0.87 to $\pm$0.64, an improvement of 26%. The receiver is characterized by measuring its step response with an internal eye monitor. Differentiating the step responses yields the impulse responses for eight different peaking amplifier settings, as shown in Fig. 22. The first positive lobe and the first negative lobe become larger in magnitude at higher peaking settings. Since these lobes are 1 UI apart, they boost the 1010 data pattern. The residual voltage bump marked in the figure is due to S/H kickback which has not completely decayed; since S/H is clocked at half rate this residual voltage acts as h2 ISI term. This modest h2 term is easily compensated by the DFE.

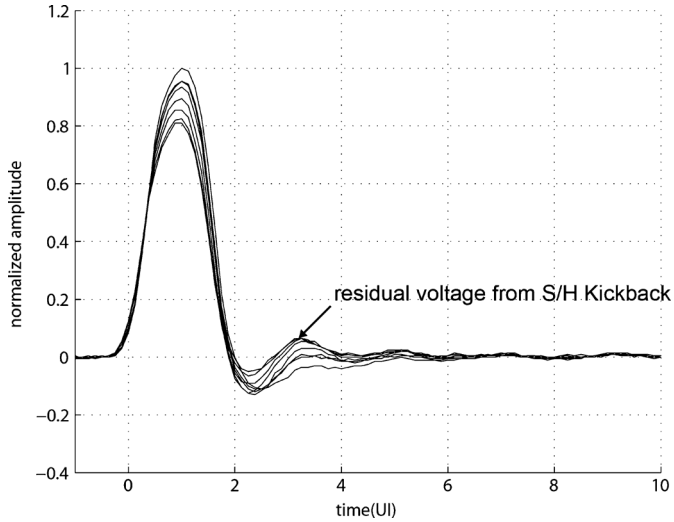

Fig. 22. Measured receiver impulse responses with eight different peaking settings. At a data rate of 14.025 Gb/s, 1 UI = 71.3 ps.

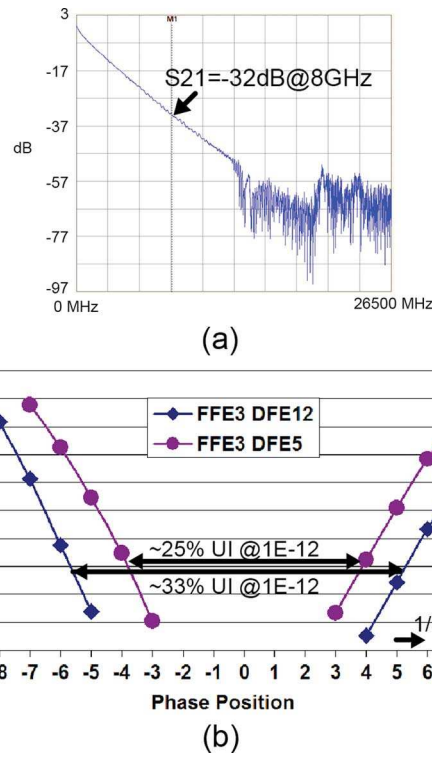

Fig. 23. (a) Frequency response of 40″ Nelco test card channel. (b) Equalized bathtub curve in 12-tap and 5-tap DFE modes.

Equalization at 16-Gb/s is evaluated over a 40″ Nelco test card channel with 32 dB loss at 8 GHz (Fig. 23(a)). In this equalization experiment, the 3-tap FFE of the TX is used; for sake of comparison, the DFE is operated in both 5 and 12-tap DFE modes. The RX recovers error-free ($\mathrm{BER} < 10^{-15}$) 16-Gb/s PRBS31 data in asynchronous mode with 200 ppm frequency offset; the horizontal eye opening at $\mathrm{BER} < 10^{-12}$ in 5-tap and 12-tap DFE modes is 25% and 33%, respectively (Fig. 23(b)).

The core is fully characterized for 16GFC (14.025 Gb/s) specification [10]. Fig. 24 shows the PLL phase noise measured at the TX output; integrated noise is 303 $\mathrm{fs_{rms}}$ after being filtered according to the 16GFC specification. The RX is compliant for jitter tolerance. Worst case power consumption with amortized

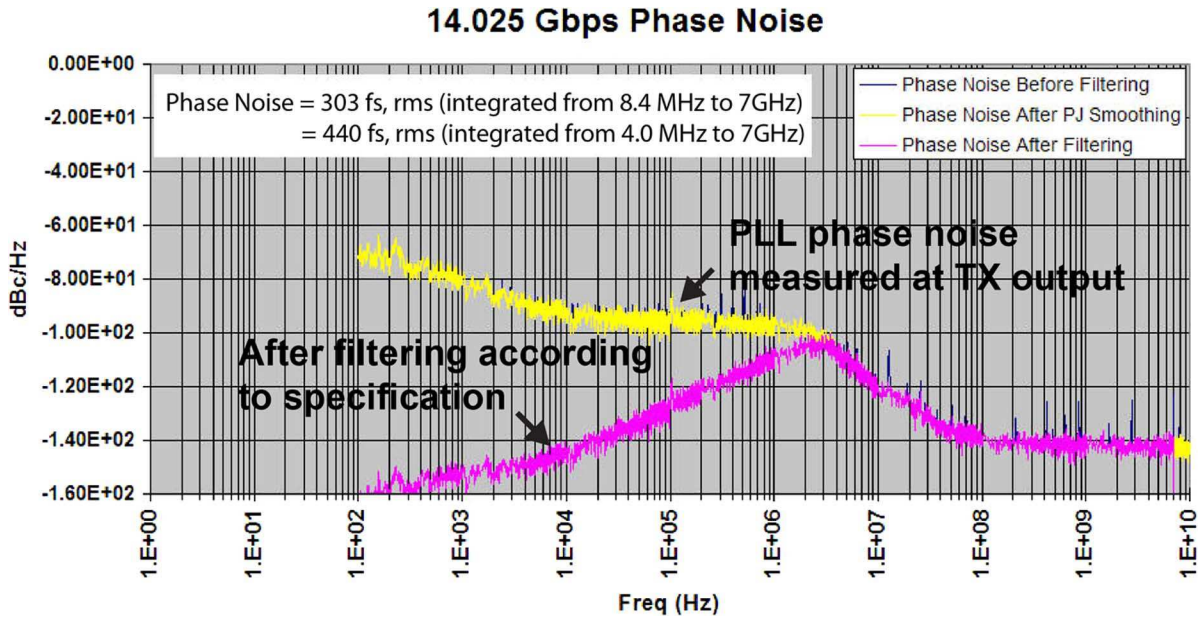

Fig. 24. PLL phase noise measured at TX output.

PLL power is 385 mW/link; this comes out to 27.5 pJ/bit, compared to 29.2 pJ/bit in [1] and 14.7 pJ/bit in [2].

V. Conclusion

This paper describes a transceiver architecture which dynamically compensates offset drifts in the DFE, quadrature and duty cycle drifts in the RX rotator clocks, and TX output duty cycle in order to achieve robust operation at >14-Gb/s data rates. It also details circuit refinements to the current integrating summers for improved performance. The transceiver operates error-free at 16 Gb/s when equalizing a channel with 32 dB loss.

Acknowledgment

The authors thank J. Rockrohr, P. Owczarski, M. Chhabra, R. Callan, A. Davies, P. Elakkumanan, and K. Kramer for help and support.

References

[1] S. Quan *et al.*, "A 1.0625-to-14.025 Gb/s multimedia transceiver with full-rate source-series-terminated transmit driver and floating-tap decision-feedback equalizer in 40 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2011, pp. 348–349.

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

[3] J. F. Bulzacchelli *et al.*, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.

[4] J. F. Bulzacchelli *et al.*, "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*, 2009, pp. 368–369.

[5] 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*, 2007, pp. 228–229.

[6] M. Kossel *et al.*, "A T-coil-enhanced 8.5 Gb/s high-swing SST transmitter in 65 nm bulk CMOS with <-16 dB return loss over 10 GHz bandwidth," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2905–2920, Dec. 2008.

[7] G. G. Shahidi *et al.*, "Device and circuit design issues in SOI technology," in *Proc. IEEE CICC*, 1999, pp. 339–346.

[8] L. Chen, X. Zhang, and F. Spagna, "A scalable 3.6-to-5.2 mW 5-to-10 Gb/s 4-tap DFE in 32 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2009, pp. 180–181.

[9] K. Yamaguchi, M. Fukaishi, T. Sakamoto, N. Akiyama, and K. Nakamura, "2.5 GHz 4-phase clock generator with scalable and no feedback loop architecture," in *IEEE ISSCC Dig. Tech. Papers*, 2001, pp. 398–399.

[10] 16 G Fibre Channel Spec. [Online]. Available: http://www.t11.org/t11/hps.nsf/smemhome

[11] H. Wang and A. Hajimiri, "A wideband CMOS linear digital phase rotator," in *Proc. IEEE CICC*, 2007, pp. 671–674.

Gautam R. Gangasani grew up in Bombay, India, and received his undergraduate degree from the Indian Institute of Technology, Bombay. He also holds an M.Phil. degree in electrical engineering from Columbia University, New York, and M.S. degrees in computer science as well as electronics and computer control systems, from Wayne State University, Detroit, MI.

His prior work experience was in developing plant automation control systems, and in ASIC design. Presently, he is working in the high speed serial IO development group at IBM Microelectronics, Fishkill, NY, developing serial link receivers ranging in speeds from 6 Gbps to 28 Gbps in CMOS technology. He holds five U.S. patents and has authored seven technical publications. His current interests include high speed decision feedback equalization, oscillators, phase-locked loops, phase interpolation, clock and data recovery, and low power high speed mixed signal CMOS circuits.

Chun-Ming Hsu received the B.S. and M.S. degrees in electrical engineering from the National Taiwan University, Taipei, Taiwan, in 1997 and 1999, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2008.

His research interest is in mixed-signal and RF circuits for wireless and wireline communications. From 1997 to 1999, He began his research career by exploring low-voltage PLL and DLL in the National Taiwan University. From 1999 to 2004, he worked at the Industrial Technology Research Institute, Hsinchu, Taiwan, as a circuit design engineer for GSM and WCDMA RF transceivers and was later promoted to management positions. While he was at the Massachusetts Institute of Technology from 2004 to 2008, he focused on implementation of high-performance digital PLL. His Ph.D. thesis demonstrated a 3.6 GHz frequency synthesizer with a 500 kHz BW and <300 fs jitter. Since 2008,

he has been with IBM, East Fishkill, NY, where he works on clock-and-data recovery and decision-feedback equalization in IBM's 45 nm and 32 nm high-speed serial links for backplane and optical applications.

Dr. Hsu was co-recipient of the Best Student Paper Award (2nd place) in the 2008 IEEE Radio Frequency Integrated Circuits (RFIC) Symposium. He has been a technical program committee member in RFIC Symposium since 2011.

John F. Bulzacchelli (S'92–M'02) was born in New York, NY, in 1966. He received the S.B., S.M., and Ph.D. degrees in electrical engineering, all from the Massachusetts Institute of Technology (MIT), Cambridge, in 1990, 1990, and 2003, respectively.

From 1988 to 1990 he was a co-op student at Analog Devices, Wilmington, MA, where he invented a new type of delay-and-phase-locked loop for high-speed clock recovery. From 1992 to 2002 he conducted his doctoral research at the IBM T. J. Watson Research Center, Yorktown Heights, NY, in a joint study program between IBM and MIT. In his doctoral work, he designed and demonstrated a superconducting bandpass delta-sigma modulator for direct A/D conversion of multi-GHz RF signals. In 2003 he became a Research Staff Member at this same IBM location, where his primary job is the design of mixed-signal CMOS circuits for high-speed data communications. He also maintains strong interest in the design of circuits in more exploratory technologies. He holds eleven U.S. patents.

Dr. Bulzacchelli received the Jack Kilby Award for Outstanding Student Paper at the 2002 IEEE International Solid-State Circuits Conference (ISSCC), was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 ISSCC, and co-authored the IEEE JOURNAL OF SOLID-STATE CIRCUITS article awarded Best Paper of 2009.

Sergey Rylov received the M.S. and Ph.D. degrees in physics from Moscow State University, Moscow, Russia, in 1984 and 1987, respectively.

Until 1991, he worked as a Research Scientist with the Laboratory of Cryoelectronics, Moscow State University. His research was in the field of superconducting Josephson junction microelectronics, particularly single-flux-quantum digital devices, analog-to-digital converters (ADCs) and physically reversible computers. From 1991 to 1998 he was with HYPRES, Inc., where he successfully designed many high-performance superconducting digital and analog devices, including high-resolution and flash ADCs, single-flux-quantum logic devices and analog amplifiers using DC SQUIDs. Since 1998, he has been with the IBM T. J. Watson Research Center, where he currently works on circuit design of high-speed digital and mixed-signal devices for CMOS communications ICs.

Troy Beukema received the B.S.E.E. and M.S.E.E. degrees from Michigan Technological University, Houghton, in 1984 and 1988, respectively.

Since 1996, he has been with IBM Research in Yorktown Heights, NY, where he has focused on both wireless and wireline system designs in the Communication Technology department. His most recent activity at IBM is in the area of high data rate SERDES system design and analysis. He has developed line equalization, clock-and-data recovery, and analog datapath system designs for backplane I/Os operating from 6 to 25 Gb/s spanning CMOS technology generations from 130 nm to 32 nm. His ongoing research interests include both ultra high data rate SERDES and low-power ADC based I/O core designs.

David Freitas received the B.S. degree in electrical engineering from San Jose State University, San Jose, CA, in 1980, and the M.S. degree in electrical engineering from the University of California at Davis in 1983.

In 1983, he joined IBM, San Jose, CA, where he worked on analog circuit designs for disk drives. Since 2004, he has worked on high speed SerDes analog circuit designs for various communication standards.

William R. Kelly was born in Jersey City, NJ, in 1957. He received the B.S. degree in electrical engineering from Rochester Institute of Technology, Rochester, NY, in 1981.

In 1982, he joined IBM in Poughkeepsie, NY, as an engineer. From 1982 to 1990, he worked on several projects related to optoelectronics and fiber optic communications. From 1990 to 1993, he held management positions in microprocessor engineering and optoelectronics engineering groups. From 1993 to 2001, he worked on projects developing personal computers, network adapters, video-conferencing systems, and storage applications. From 2001 to the present, he has worked as a design engineer for various high speed SERDES interfaces, with a primary focus on development and implementation of advanced equalization architectures and related logic algorithms. He holds four U.S. patents.

Michael Shannon received the B.S. degree in electrical engineering from the University of New Hampshire in 1997, and the M.S. degree in electrical engineering from the University of Minnesota in 2002.

He is a senior engineer with IBM working in the area of high speed SerDes mixed signal circuit development. He was previously employed by Micron Technology, Mathstar, and JDS Uniphase as a mixed signal integrated circuit designer. In these positions he worked in IC development for high speed serial links and optical data communications.

Jieming Qi received the B.S. degree from Beijing Institute of Technology, Beijing, China, and the Ph.D. degree in electrical engineering from University of Texas at Austin in 1997.

In 1998, he joined IBM and has worked on phase-locked loops and other high speed circuitry for high performance microprocessors. Since 2007, he has been working on design of high speed CMOS SerDes cores. He has over 60 patents issued.

Hui H. Xu received the B.S. and M.S. degrees in electrical engineering from Polytechnic University, Brooklyn, NY, in 2001.

In July 2001, he joined IBM Microelectronics, Hopewell Junction, NY, where he has worked on high speed SerDes analog circuit designs for various communication standards.

Joseph Natonio received the B.S. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY.

He is an advisory engineer at IBM and since 1999 has worked in the ASIC packaging group, SiGe group, as well as in the serial IO development group. He holds 13 patents.

Todd M. Rasmus received the B.S.E.E. degree from the South Dakota School of Mines and Technology, Rapid City, in December 1983.

In 1984 he joined IBM in Rochester, MN, as a Quality Engineer in the area of switched-mode power generation. In 1985, he relocated to IBM's Networking Division in Research Triangle Park, NC, as a Power Systems design engineer where he developed switched-mode power supplies for IBM PCs. From 1987 to 1999, he developed various analog systems in the realms of integrated video

and high-speed telephony. Since 1999, he has been involved in the design of mixed-signal CMOS circuits for high-speed data communications. He holds 21 U.S. patents.

Jong-Ru Guo received the M.S. degree from National Tsing-Hua University, Hsinchu, Taiwan, in 1996 and Ph.D. degree from Rensselaer Polytechnic Institute, Troy, NY, in 2005. He worked in SiGe high speed field programmable gate array running in giga-hertz range while pursuing his Ph.D. degree.

After receiving his degree, he joined IBM High Speed Serial-link (HSS) team in 2005 and worked on the mixed-signal circuit designs. During his stay with IBM, he received several patents with IBM and published some journal articles with the research team in Rensselaer. His research interests include the high-speed circuits and reconfigurable analog and digital systems.

Michael Wielgos received the B.S. degree in electrical engineering from SUNY, New Paltz, NY.

He is an advisory engineer at IBM and is currently working on high speed IO characterization for various communication standards.

Jon Garlett received the B.S. degree in electrical engineering from Northwestern University, Evanston, IL, in 1983.

He is a senior engineer at IBM and since 1995 has worked on high speed SerDes analog circuit design and characterization for various communication standards.

Michael A. Sorna received the B.S. degree in electrical engineering from the University of Pittsburgh, Pittsburgh, PA, in 1984, and the M.Eng. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1986.

He has been working at the IBM East Fishkill facility since 1984 and appointed IBM Distinguished Engineer in 2010. From 1984 to 1993, he worked in the Test Equipment Engineering Group, developing high-speed wafer-level production test equipment. From 1993 to the present, his responsibility centered around high speed SerDes/chip circuit design, in which a wide variety of standards were addressed. In addition, he leads IBM's DDR Memory PHY development group. His design experience includes both high speed analog and digital areas in bipolar (including SiGe), CMOS (22 nm–0.8 μm), and GaAs technologies. He is currently responsible for IBM ASIC SerDes and DDR Memory PHY R&D, including circuit design, architecture, and lab characterization. He holds over 50 patents and has authored numerous technical publications and presentations, primarily in the field of SerDes/Memory PHY circuits and BIST techniques.

Mounir Meghelli received the M.S. degree in electronics and automatics from the University of Paris XI, France, in 1992 and the Engineering degree in telecommunication from the ENST-Paris, France, in 1994. He received the Ph.D. degree from the University of Paris VI after a four-year research program with the CNET France Telecom Research Center, working on the design of high-speed ICs for optical communications in GaAs and InP HBT technologies.

From 1998 to 2005, he joined the IBM T. J. Watson Research Center as a Research Staff Member working on the design of high frequency ICs in SiGe BiCMOS and CMOS technologies for wireline and wireless applications. He is currently a Senior Technical Staff Member with the IBM Server and Technology Group leading the design of advanced serial links for storage, networking and server applications. He is author and coauthor of 45 technical publications and hold 11 U.S. patents covering a broad spectrum of subjects including low power serial links, PLLs, and advanced equalization.