A 28-Gb/s 4-Tap FFE/15-Tap DFE Serial Link Transceiver in 32-nm SOI CMOS Technology

John F. Bulzacchelli, *Member, IEEE*, Christian Menolfi, *Member, IEEE*, Troy J. Beukema, *Member, IEEE*, Daniel W. Storaska, *Member, IEEE*, Jürgen Hertle, *Member, IEEE*, David R. Hanson, Ping-Hsuan Hsieh, *Member, IEEE*, Sergey V. Rylov, *Member, IEEE*, Daniel Furrer, Daniele Gardellini, *Member, IEEE*, Andrea Prati, Thomas Morf, *Senior Member, IEEE*, Vivek Sharma, *Member, IEEE*, Ram Kelkar, Herschel A. Ainspan, William R. Kelly, Leonard R. Chieco, Glenn A. Ritter, John A. Sorice, Jon D. Garlett, Robert Callan, Matthias Brändli, Peter Buchmann, Marcel Kossel, *Senior Member, IEEE*, Thomas Toifl, *Senior Member, IEEE*, and Daniel J. Friedman, *Member, IEEE*

***Abstract*—This paper presents a 28-Gb/s transceiver in 32-nm SOI CMOS technology for chip-to-chip communications over high-loss electrical channels such as backplanes. The equalization needed for such applications is provided by a 4-tap baud-spaced feed-forward equalizer (FFE) in the transmitter and a two-stage peaking amplifier and 15-tap decision-feedback equalizer (DFE) in the receiver. The transmitter employs a source-series terminated (SST) driver topology which doubles the speed of existing half-rate designs. The high-frequency boost provided by the peaking amplifier is enhanced by adopting a structure with capacitively coupled parallel input stages and active feedback. A capacitive level-shifting technique is introduced in the half-rate DFE which allows a single current-integrating summer to drive the four parallel paths used for speculating the first two DFE taps. Error-free signaling at 28 Gb/s is demonstrated with the transceiver over a channel with 35 dB loss at half-baud frequency. In a four-port core configuration, the power consumption at 28 Gb/s is 693 mW/lane.**

***Index Terms*—Active feedback, backplane, capacitive level shifter, chip-to-chip communications, current-integrating summer, decision-feedback equalizer (DFE), feed-forward equalizer (FFE), peaking amplifier, serial link, source-series terminated (SST) driver, transceiver.**

I. Introduction

WITH the proliferation of digital devices accessing advanced network services such as multimedia-on-demand and the predicted rise of cloud computing, the I/O bandwidth requirements of systems such as routers and servers are expected to grow rapidly. To expand I/O capacity, serial link data rates are now being pushed up to 25–28 Gb/s, as exemplified by recent and upcoming standards such as OIF CEI-25G-LR and CEI-28G-SR [1], 32GFC [2], IEEE 802.3bj [3], and InfiniBand EDR [4]. Such data rates represent a near doubling of the state-of-the-art for fully integrated backplane transceivers, which have been previously reported up to 16 Gb/s [5]–[7]. With technology scaling no longer providing large gains in raw device speed [8], significant design advances must be made to attain the desired data rates.

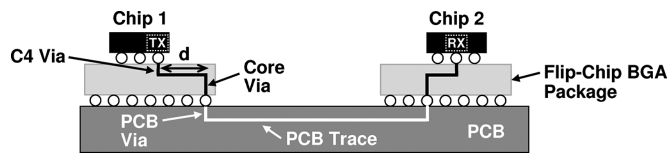

Fig. 1. Chip-to-chip serial link across PCB.

Adding to the design challenge is the difficulty of electrical channel characteristics at data rates approaching 30 Gb/s. For a 1–m-long printed circuit board (PCB) trace or backplane, the loss at half-baud frequency may exceed 30 or even 35 dB. A common practice in backplane transceiver design [5]–[7], [9], [10] is to employ a feed-forward equalizer (FFE) in the transmitter and a decision-feedback equalizer (DFE) in the receiver. To handle higher channel loss, the number of taps in the FFE and DFE can be increased, but at the cost of extra circuit power and area. A previous system-level study [11] of electrical links operating at 25 Gb/s showed that a 4-tap FFE provides close-to-optimal performance, while both vertical and horizontal eye openings benefit from increasing the number of DFE taps to at least 20. While the transceiver developed in this work [12], [13] does include a 4-tap FFE in its transmitter, the DFE in its receiver only has 15 taps. The number of required DFE taps is reduced in this design by including a wide-range (>10 dB) peaking amplifier in the receiver (a feature not assumed in the study of [11]). The linear equalization provided by the peaking amplifier helps compensate for intersymbol interference (ISI) outside the time span of the DFE. This usage of a linear equalizer to reduce the DFE tap requirements is conceptually similar to that described in [14], but the equalizer employed here has a more conventional response, with the gain peaked at high frequency rather than at

Manuscript received May 01, 2012; revised July 13, 2012; accepted July 27, 2012. Date of publication October 09, 2012; date of current version December 21, 2012. This paper was approved by Guest Editor Jack Kenney.

J. F. Bulzacchelli, T. J. Beukema, S. V. Rylov, H. A. Ainspan, and D. J. Friedman are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: jfbulz@us.ibm.com; troyb@us.ibm.com).

C. Menolfi, T. Morf, M. Brändli, P. Buchmann, M. Kossel, and T. Toifl are with the IBM Zurich Research Laboratory, Rüschlikon CH-8803, Switzerland.

D. W. Storaska, D. R. Hanson, W. R. Kelly, L. R. Chieco, G. A. Ritter, J. A. Sorice, J. D. Garlett, and R. Callan are with the IBM Systems and Technology Group (STG), Hopewell Junction, NY 12533 USA.

J. Hertle, D. Furrer, D. Gardellini, A. Prati, and V. Sharma are with Miromico AG, Zurich CH-8006, Switzerland.

P.-H. Hsieh was with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA. She is now with the Department of Electrical Engineering, National Tsing Hua University, Hsinchu 30013, Taiwan.

R. Kelkar is with IBM STG, Essex Junction, VT 05452 USA.

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

Digital Object Identifier 10.1109/JSSC.2012.2216414

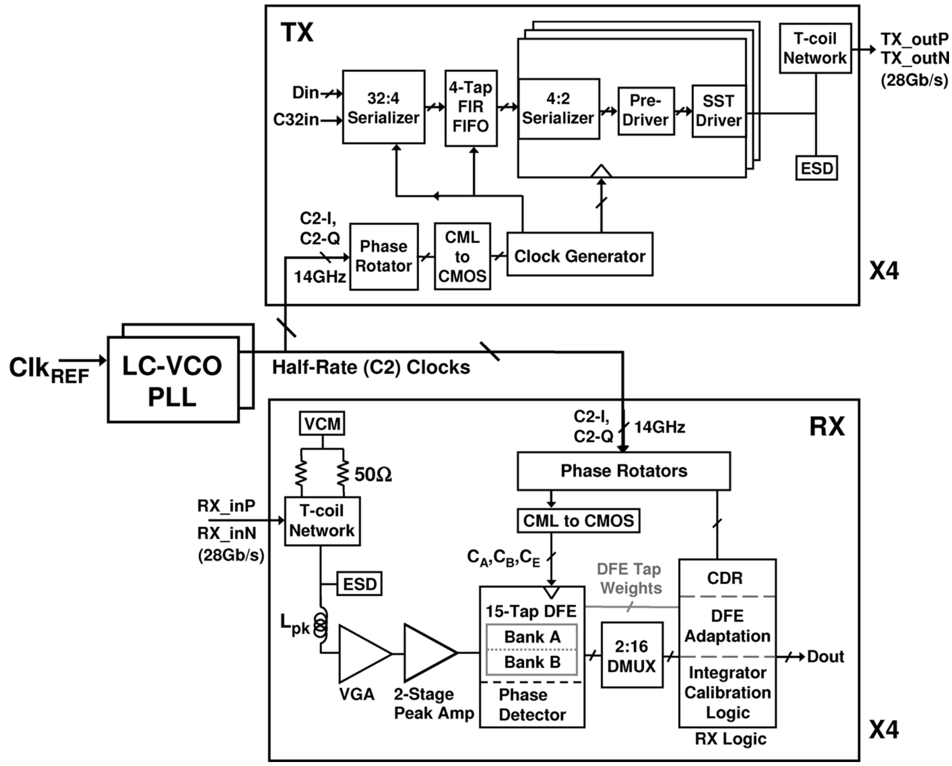

Fig. 2. Top-level architecture of four-port I/O transceiver.

low frequency. The high-frequency gain of the peaking amplifier reduces the amount of de-emphasis needed in the transmitter FFE. Using less de-emphasis in the transmitter and more linear equalization in the receiver increases the average signal level at the receiver input and helps reduce high-frequency jitter amplification by high-loss channels [15].

The need to compensate ISI arising from package escape reflections prevents one from reducing the number of DFE taps too much. Consider, for instance, the chip-to-chip link depicted in Fig. 1. Because of the impedance discontinuities introduced by the core via of the package, the solder ball, and the PCB via, some of the signal launched by the transmitter (TX) is reflected at the package-to-PCB interface. Due to imperfect output return loss, the transmitter does not completely absorb the reflection, and a reflected signal appears at the input of the receiver (RX). The delay T_r of this reflected signal (relative to that of the main cursor) equals $2d/c$, where d and c are the package trace length and wave velocity, respectively. Assuming a maximum package trace length of 25 mm, T_r may be as large as 450 ps with typical package materials, which corresponds to 12.6 unit intervals (UIs) at 28 Gb/s. This reflection (as well as the corresponding one inside the receiver package) can be effectively cancelled with a 15-tap DFE.

These system-level considerations require that a 28-Gb/s backplane transceiver have greater equalization capabilities than the previously reported transceivers operating at 14–16 Gb/s [5]–[7]. The design techniques used to implement a 28-Gb/s transceiver with such equalization performance in 32-nm silicon-on-insulator (SOI) CMOS technology are the major focus of this paper, which is organized as follows. Section II presents the architectures of the transmitters and receivers of the I/O core. Sections III and IV describe the circuit design details of the transmitter and receiver, respectively. Experimental results are discussed in Section V, and Section VI concludes with a summary.

II. Transceiver Architecture

Fig. 2 presents the top-level architecture of the transceiver, which is configured as a four-port I/O core. Two phase-locked loops (PLLs) with 2:1 dividers generate the half-rate (C2) clocks which are distributed to the four transmitters and four receivers. Each PLL includes two different LC voltage-controlled oscillators (VCOs) so that the oscillator frequency can be varied over a range of 14–28.05 GHz. The transmitters and receivers both employ half-rate architectures, which are described in the following subsections.

A. Transmitter

The transmitter consists of three main circuit blocks: a data path, a clock generator, and a segmented source-series terminated (SST) driver. The data path includes a 32:4 serializer and a shift register that produces time-delayed quarter-rate tap data streams for a baud-spaced 4-tap FFE. The tap data streams are then distributed to a set of weighted SST driver segments, which perform the final serialization to the data output. Asymmetric T-coils are used to compensate for driver output capacitance and parasitics of the electrostatic discharge (ESD) device (low-capacitance silicon-controlled rectifier) and to provide wideband impedance matching [16].

The clock generator produces the subrate clocks needed in the serializer stages and provides a mechanism for adjusting the duty cycles of the performance-critical half-rate clocks. For this

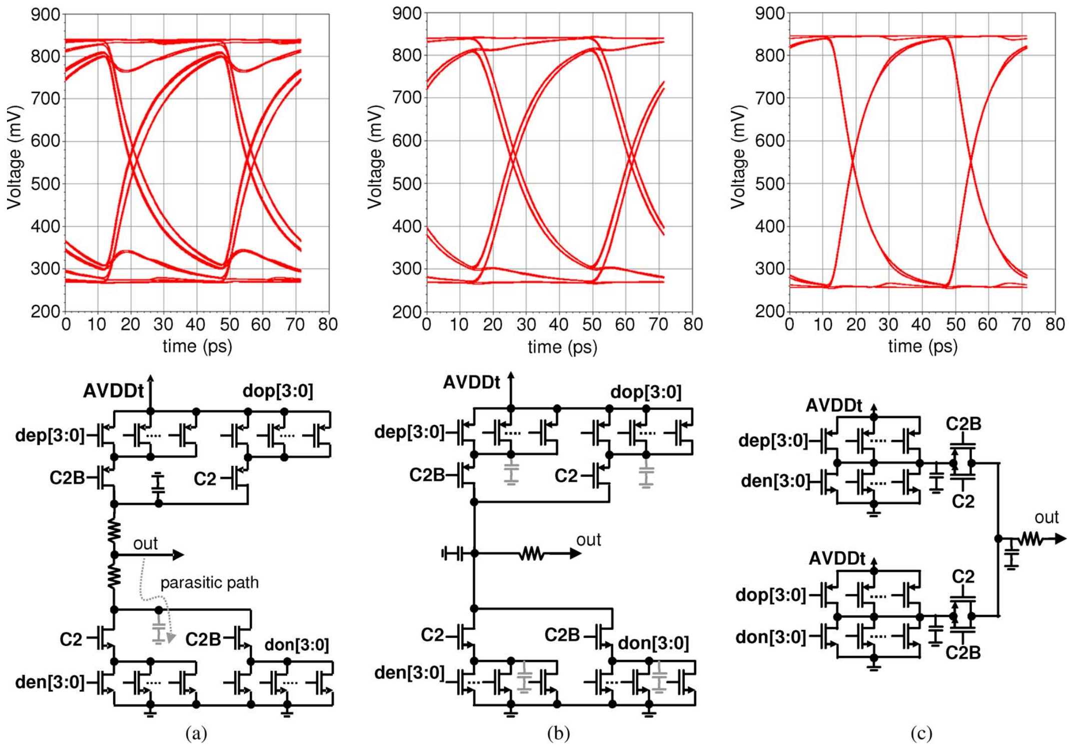

Fig. 3. Evolution of the SST driver speed optimization. (a) Original stacked structure. (b) Stacked structure with single shared linearization resistor. (c) Final structure.

prototype, the duty cycle control bits are set manually (after measuring the transmitter outputs with an oscilloscope), but adding closed-loop duty cycle correction (DCC), such as that described in [6], would be straightforward. Facilitated by the pseudodifferential structure of the SST drivers, an optional feature of the clock generator adds a variable delay between the half-rate clocks of the true and complement outputs (TX_outP, TX_outN) which can be used to compensate length mismatches in cable pairs or differential skew in long PCB traces [17]. As an experiment, this optional feature was implemented (as an open-loop adjustment) in a separate breakout test site of the transmitter but was not included in the fully integrated transceiver. Insertion of a current-mode logic (CML) phase rotator in the clock path allows the C32 clock for the data serializer to be aligned with a clock (C32in) forwarded from logic outside the I/O core. During initial link setup, the clock alignment is checked with a latch used as a bang-bang phase detector (not shown in the figure), and on-chip logic determines the best rotator setting for capturing the input data; once established, this setting is fixed during normal data transmission.

A separate supply strategy helps mitigate supply noise-induced jitter without needing on-chip voltage regulation. While the data path and SST drivers are powered from a data supply (AVDDt), the clock generation and distribution circuits are powered from a separate clock supply (CVDD). The two supplies have the same nominal value (1.05 V) and are kept separate up to the board level to minimize interference.

B. Receiver

The major functional blocks of the receiver are similar to those in [6], but their underlying circuits are extensively modified to support higher data rates. Inductive peaking is used heavily to extend the bandwidths of the variable gain amplifier (VGA) and peaking amplifier. Another inductor $L_{\rm pk}$ (actually, pair of inductors since the signals are differential) is placed in series with the VGA input to provide some fixed passive peaking (about 3–4 dB boost at 12.5–14 GHz), which helps compensate for package losses. The two-stage peaking amplifier provides up to 12 dB of gain boost at half-baud frequency.

The 15-tap DFE employs two redundant banks (A and B), each of which is realized as a half-rate structure. As in the design of [6], the two banks can be swapped between the functions of data detection and adaptation/calibration. CML-based phase rotators generate the half-rate clocks for the two DFE banks and the phase detector that provides edge samples for a digital clock and data recovery (CDR) loop. Each DFE bank clock (e.g, C_B for bank B) can be independently swept relative to the other clocks (e.g., C_A and C_E) to monitor the horizontal eye opening, and the information gained from such measurements is used to position the DFE bank clocks for optimal sampling of the equalized data eye. The system is not sensitive to static phase offsets between the data samples and the (non-DFE-equalized) edge samples [6]. In an analogous manner, the vertical eye opening is monitored for asymmetry, which is corrected by applying a compensating dc offset inside the VGA.

In contrast with the transmitter, the receiver features closed-loop DCC of clocks C_A, C_B, and C_E, based on the circuits presented in [6]. In particular, an offset-compensated comparator is used to detect the difference in the average voltages of a clock and its complement (near the end of the clock distribution, inside the DFE), and a low-bandwidth digital control loop adjusts the duty cycle (in a stage after the CML-to-CMOS converter) to

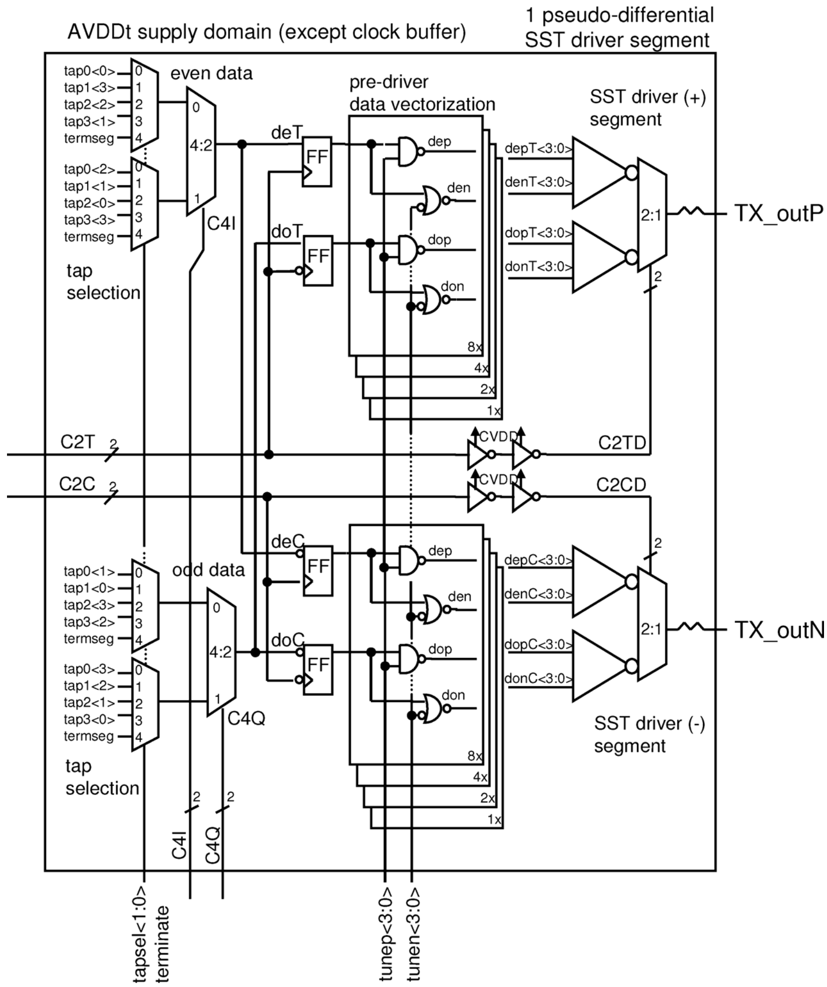

Fig. 4. SST driver segment with tap selection and pre-driver circuitry.

compensate for any duty cycle distortion (DCD) accumulated in the distribution. As the data-dependent supply current variations of the DFE are not as large as those of the SST driver in the transmitter, the receiver data path and clock circuits are powered from a single supply (AVDDr), with a nominal value of 1.05 V. Synthesized logic executes the algorithms used for CDR, DFE adaptation, and analog circuit calibrations and operates from the main digital supply (VDD) of the chip, with a nominal value of 0.85 V.

III. Transmitter Circuits

A. SST Driver

An important decision in the design of a half-rate transmitter is the location of the final 2:1 multiplexer (MUX) in the output signal chain. Placing a lower power MUX early in the chain, followed by a full-rate SST driver, is certainly attractive from a power perspective since the multiplexing half-rate clock does not need to be powered up to the final driver size, and the full-rate SST driver switches are typically smaller than the stacked switches of a multiplexed half-rate SST driver [18]. However, multiple full-rate buffer stages would be exposed to delay variations due to noise on the data supply, ISI, and floating-body effects in partially depleted SOI technologies [19]. To avoid degrading the output signal, a half-rate SST driver has been chosen for this design, in which the output timing is tightly controlled by a low-jitter half-rate clock.

Fig. 3 depicts the optimization steps which have been applied in doubling the speed of existing half-rate SST drivers [18]. Fig. 3(a) shows the original structure along with the associated 28-Gb/s eye diagram. The driver incorporates a stacked MUX that is selected by a complementary clock signal (C2/C2B) and driven with half-rate even (dep, den) and odd (dop, don) data streams. A variable data transistor width is used for driver impedance tuning. The corresponding eye suffers from limited slew rate, incomplete settling, and data-dependent jitter. The root cause of this degradation is parasitic capacitors within the driver stack which may become undriven and store data-dependent charges. As an example, consider the parasitic capacitor highlighted in gray in Fig. 3(a). During a pull-up operation, this capacitor is charged upwards relatively slowly through the pull-down resistor, and the current flowing through this parasitic path contributes to sluggish settling. This particular source of slow settling can be eliminated by converting the separate pull-up and pull-down resistors to a single shared resistor, as shown in Fig. 3(b). The parasitic capacitance behind the resistor still exists but is now always driven high or low actively.

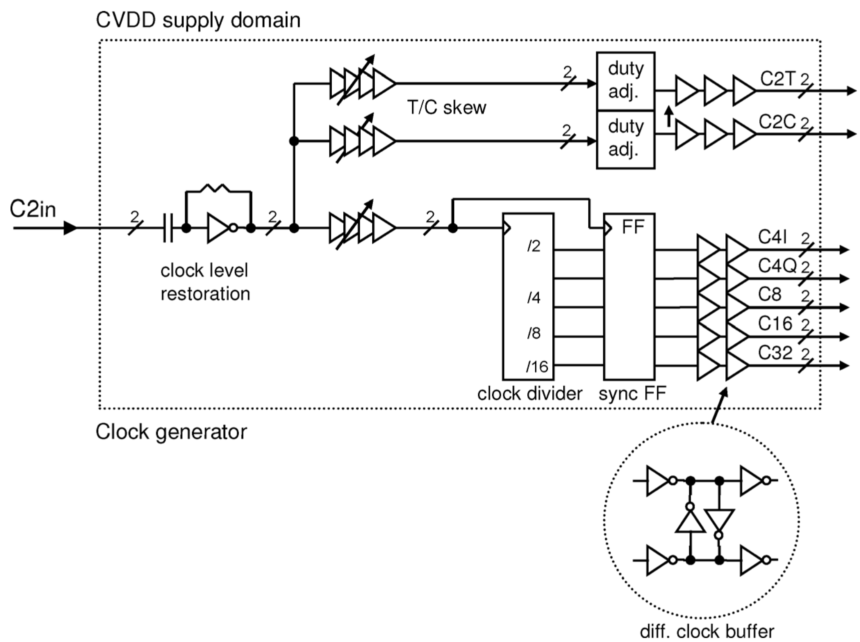

Fig. 5. Clock generator of transmitter.

The corresponding eye is substantially improved but still exhibits some data-dependent components, which are due to other parasitic capacitors [highlighted in gray in Fig. 3(b)]. One of these capacitors may become undriven when a clock transistor is turned off. In the final step of the optimization [Fig. 3(c)], the clock transistors (now operating as transmission gates) are relocated between the even/odd branches and the single shared resistor. The SST driver has effectively been transformed from a stacked MUX to a passgate MUX with programmable variable width inverters for the even and odd data. There are no undriven circuit nodes in this very simple structure. The clean data eye confirms the superior performance of the proposed circuit, which has been adopted here for the transmitter driver segments.

A single SST driver segment is shown in Fig. 4. Each driver segment is independently configurable as one of the four FFE taps or as a terminating static high or low segment, which is accomplished with a static tap selection MUX. After being converted to half-rate by 4:2 MUXes, the data streams are retimed to clocks C2T and C2C, which control the timing of the true and complement transmitter outputs (TX_outP, TX_outN), respectively. As an optional feature, C2T and C2C may be skewed by a programmable amount up to about $+/-20$ ps. The retimed data bits are then multiplied by pull-down $(\mathrm{tunen}\langle 3:0\rangle)$ and pull-up $(\mathrm{tunep}\langle 3:0\rangle)$ impedance tuning vectors in the pre-driver and delivered to the SST driver circuits. The complete driver is composed of 24 weighted SST driver segments. A driver segment weighting of 8×8, 4×4, 4×2 and 8×1 segments has been chosen, which results in an SST driver with 96 equivalent segments.

B. Clock Generator

The clock generator circuitry is shown in Fig. 5. An ac-coupled inverter with resistive feedback restores the incoming differential half-rate clock C2in to rail-to-rail levels and drives three clock paths. The first and second paths generate the differential clocks C2T and C2C mentioned above. The third path produces sub-rate clocks for the data serializer stages. The delay in each path may be varied by up to 20 ps, resulting in a maximum differential skew between C2T and C2C of $+/-20$ ps. The variable delay is implemented with current-starved buffers controlled in 13 programmable monotonic steps.

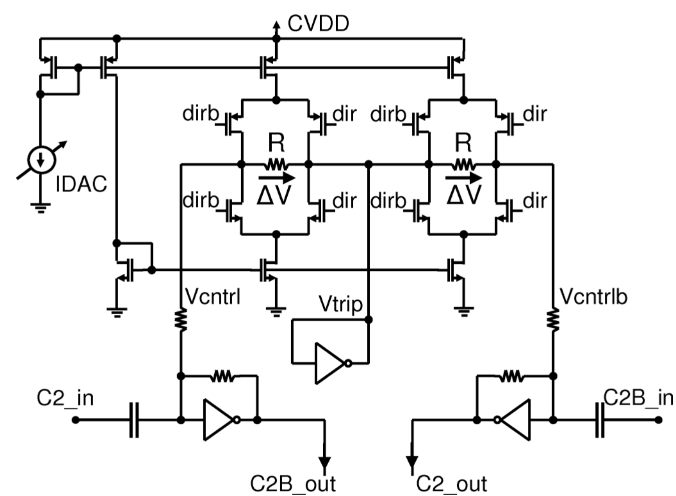

Fig. 6. Duty-cycle adjust circuit for half-rate clocks of transmitter.

Half-rate transmitter architectures are sensitive to DCD in the clock signals. A mismatch analysis of the clock paths has shown that an accurate duty cycle is not guaranteed at the maximum clock frequency (14 GHz). The clock generator includes circuits for adjusting the duty cycles of C2T and C2C. Fig. 6 shows the schematic of the duty-cycle adjust circuit, which is based on tuning the trip points of two ac-coupled inverters with resistive feedback. If the current digital-to-analog converter (IDAC) is set to zero, no currents are driven through the resistors R, $\Delta\mathrm{V} = 0$, and the inverters are biased at their natural trip point

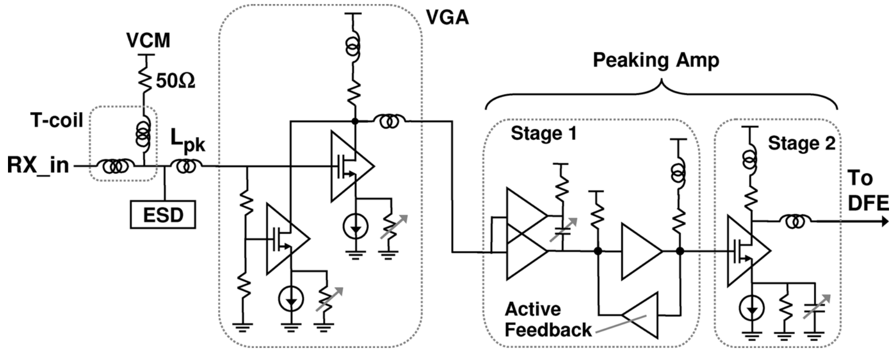

Fig. 7. Analog data path of receiver.

Vtrip (equilibrium point for 50% duty cycle), which is generated with a self-biased replica inverter. If the IDAC is set to a nonzero code, currents are driven through the resistors R, and voltages Vcntrl and Vcntrlb are moved away from Vtrip by ΔV and $-\Delta$V, respectively, which effects differential tuning of the output duty cycles (with a nominal range of $+/-6\%$ at the maximum clock frequency). The polarity of ΔV is set with control bits dir/dirb.

IV. Receiver Circuits

A. Analog Data Path

The key challenges in designing the analog data path of the receiver are extending its bandwidth and increasing the peaking available at half-baud frequency. Inductors are often employed for bandwidth extension of differential amplifiers [20] and CML circuits [21], and their usage here is illustrated in Fig. 7, which shows a single-ended representation of the analog data path. As in the amplifier of [20], both shunt and series inductors are used in broadening the bandwidths of the VGA and peaking amplifier. In the differential implementation, there are a total of 12 peaking inductors (not counting the T-coils). To save area, these inductors are realized as stacked spirals, as depicted in Fig. 8. As an example, each inductor $L_{\rm pk}$ (0.89 nH) is formed as a three-turn spiral on three metal levels, which fits within an area of 20 μm $\times$ 20 μm. With the spaces between inductors equal to at least half their linear dimensions, the electromagnetic coupling between inductors is weak enough [22] to have negligible effect on the frequency responses of the amplifier stages.

To accommodate a wide range of input signal levels, the VGA employs a parallel amplifier architecture [23] in which one differential amplifier receives the full input signal while another receives a resistively divided version. The second stage of the peaking amplifier employs a conventional zero-peaked topology with switched capacitive degeneration. A fundamental limitation of this topology is that its high-frequency gain cannot exceed the dc gain of a non-degenerated CML stage, and even that gain ($\sim$6 dB) cannot be obtained given bandwidth limitations. Better peaking is achieved in the first stage by adopting a structure with capacitively coupled parallel input stages and active feedback, whose operation is now explained.

Let each differential stage of the peaking amplifier be identified by its transconductance, as indicated in Fig. 9. The bias current consumed by each differential stage is also labeled in the figure. At low frequencies, input stage gm_{1B} is isolated from the rest of the circuit by capacitor C_c, and the active feedback structure operates like the broadband amplifier described in [24]. Applying standard feedback equations to stage 1 shows that its dc gain $A_{\rm dc}$ equals

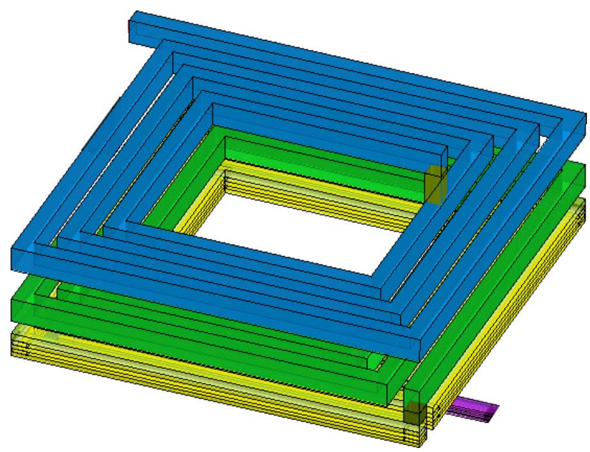

Fig. 8. Three-turn stacked spiral inductor on three metal levels.

$$A_{\rm dc} = \frac{gm_{1A}R_{1A}gm_2R_2}{1+gm_{\rm FB}R_{1A}gm_2R_2}. \tag{1}$$

The ratio of gm_{1A} to $gm_{\rm FB}$ is chosen so that $A_{\rm dc}$ is at least 0 dB when the circuit is simulated across all process, voltage, and temperature (PVT) corners. (As in [23], the use of proportional-to-absolute-temperature (PTAT) bias currents helps reduce the variation of device transconductance over temperature.)

At high frequencies, capacitor C_c couples together the outputs of the parallel input stages, and the extra input transconductance increases the voltage gain. Mathematically, a zero and pole are added to the transfer function so that (ignoring parasitic capacitances and the shunt inductor)

$$\frac{\mathrm{Vout1}(s)}{\mathrm{Vin}(s)} = A_{dc}\frac{\tau_1 s+1}{\tau_2 s+1} \tag{2}$$

where $\tau_1 > \tau_2$. While expressions for τ_1 and τ_2 can be derived, more insight into the essential advantage of this circuit is gained by examining the high-frequency gain limit $A_{\rm dc}(\tau_1/\tau_2)$. Since capacitor C_c can be considered a short at such high frequencies,

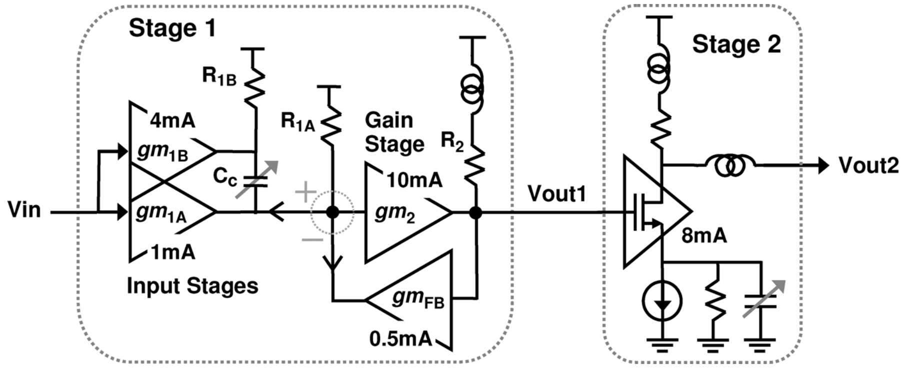

Fig. 9. Block diagram of peaking amplifier showing power allocation among stages.

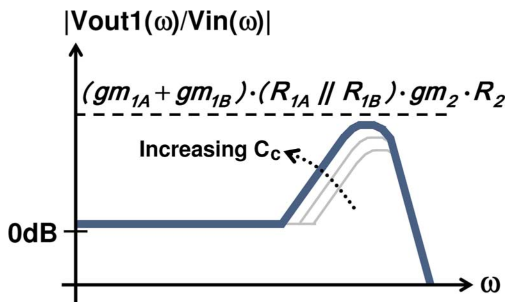

Fig. 10. Frequency response of peaking amplifier stage with active feedback (stage 1).

gm_{1A} and R_{1A} in (1) can be replaced by $gm_{1A} + gm_{1B}$ and $R_{1A} \| R_{1B}$, respectively, to yield the high-frequency gain limit

$$A_{\rm dc}\left(\frac{\tau_1}{\tau_2}\right) = \frac{(gm_{1A}+gm_{1B})(R_{1A}\|R_{1B})gm_2R_2}{1+gm_{\rm FB}(R_{1A}\|R_{1B})gm_2R_2}. \quad (3)$$

Because the value of R_{1B} is significantly lower than that of R_{1A}, the loop gain term in the denominator of (3) is much smaller than the corresponding term in (1) and is, in fact, less than unity. Therefore, (3) can be approximated as

$$A_{\rm dc}\left(\frac{\tau_1}{\tau_2}\right) \approx (gm_{1A}+gm_{1B})(R_{1A}\|R_{1B})gm_2R_2. \quad (4)$$

Thus, the high-frequency gain of stage 1 may approach the dc gain of *two* cascaded CML stages. Intuitively, with strong capacitive coupling, operation of stage 1 is effectively open loop because the high-power gm_{1B} stage overwhelms the feedback from the much weaker $gm_{\rm FB}$ stage. As depicted in Fig. 10, the peaking is adjusted by switching the value of the capacitor C_c.

Fig. 11 shows the detailed implementation of the peaking amplifier. RC degeneration is employed in the gm_{1A}, gm_{1B}, and $gm_{\rm FB}$ stages for improved linearity and bandwidth extension. The $R_{\rm FB}C_{\rm FB}$ low-pass filters reduce the feedback factor at high frequencies for a small ($\sim$1 dB) enhancement of the maximum peaking. The value of capacitor C_c is set by thermometer-coded Peaking control bits. Except for inverted polarity, these same Peaking bits are used to switch the capacitive degeneration in stage 2. In addition, there are Un-Peaking control bits for reducing the peaking; when asserted, these bits connect differential resistances across the shunt inductors, thereby de-Qing them. By controlling both Peaking and Un-Peaking bits, 17 levels of peaking are obtained.

Extracted simulations were performed to study the performance of the analog data path (from input pad to peaking amplifier output). Fig. 12 presents the simulated frequency responses at slow, nominal, and fast PVT corners. The black curves show the effects of changing the Peaking bits, while the lighter gray curves show the effects of changing the Un-Peaking bits. At the slow corner, up to 11 dB of peaking is achieved at 12.5 GHz. Considerably higher peaking at 12.5 GHz is achieved at the nominal and fast corners (19 and 23 dB, respectively). Due to the parasitic resistances of the shunt inductors (stacked spirals), asserting the Un-Peaking bits reduces the differential load impedances of the peaking amplifier stages even at dc; these parasitic resistances are only a small fraction of the total load resistance, however, so the resulting modulation of dc gain is less than 0.5 dB [Fig. 12(a)]. While modeled in the simulations, the variations of these parasitic resistances are only a minor contributor to the overall PVT corner dependence of the peaking responses.

B. DFE

To relax DFE feedback timing requirements, the first two taps (H1 and H2) are realized speculatively (loop unrolled). Limiting the power and area consumed by high-speed circuits in four parallel speculative paths is a critical design challenge. Previous works [25], [26] have shown that DFE power consumption can be reduced with the use of current-integrating summers, but in such designs a separate summer was employed for each speculative path. This overhead quickly becomes excessive as more taps are speculated.

In principle, the dc offsets representing the H1 and H2 compensation can be added into the decision-making latches themselves [27], but inserting extra devices into a latch increases its internal parasitics, which is undesirable at these data rates. As shown in Fig. 13, which presents the block diagram of a DFE half (of one bank), the dc offsets in this design are stored across series capacitors placed between the output of a current-integrating summer and the CML buffers which stabilize the input common-mode presented to the latches. This capacitive level-shifting technique allows dc offsets to be added to the received data signal with good linearity and without com-

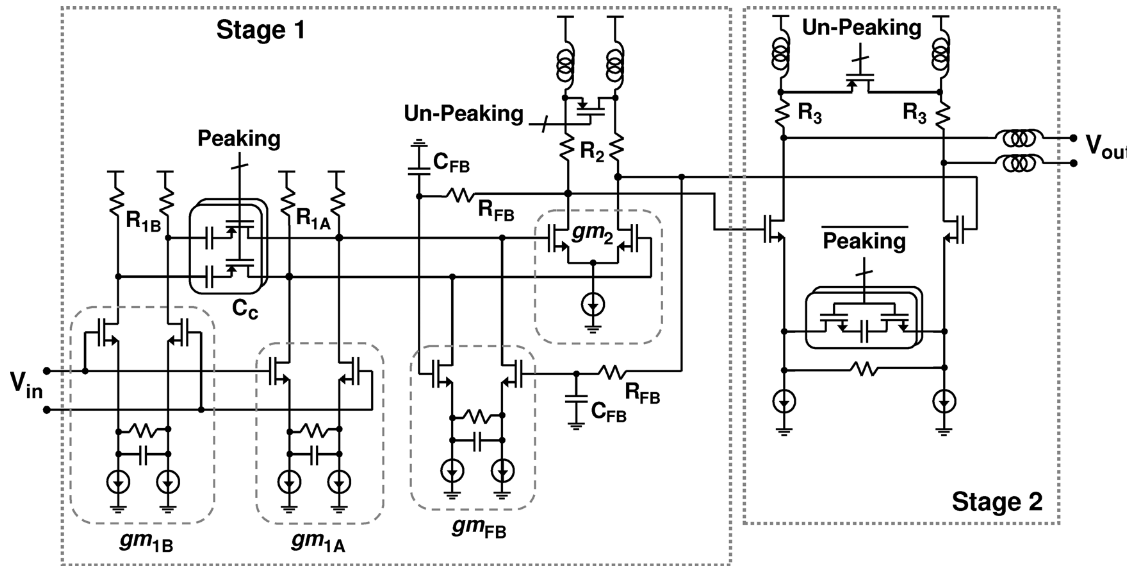

Fig. 11. Detailed schematic of peaking amplifier.

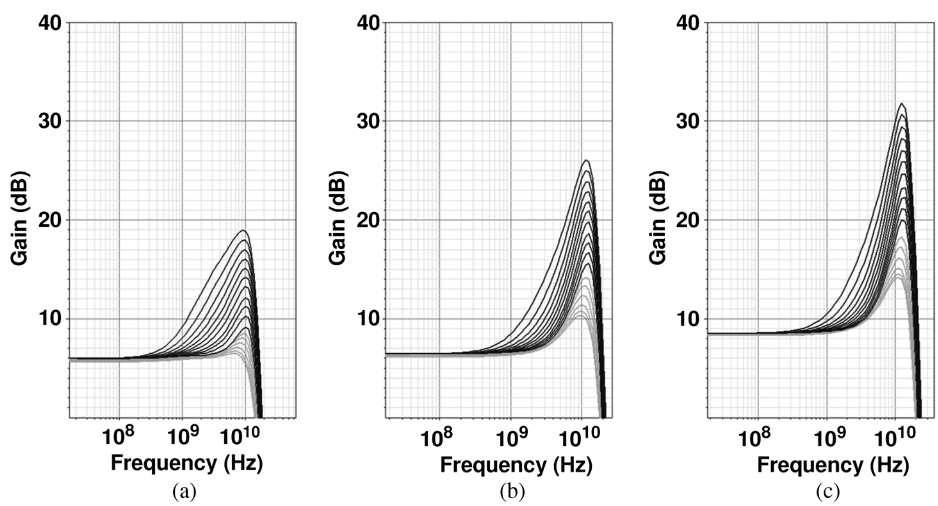

Fig. 12. Simulated frequency responses of analog data path with 17 different peaking settings. (a) Slow PVT corner. (b) Nominal PVT corner. (c) Fast PVT corner.

promising latch performance. Using a single summer to drive all four parallel paths eliminates potential mismatches between summers and saves area. Sense amplifiers producing rail-to-rail outputs are used as the decision-making latches, and the DFE feedback logic is implemented in domino and static CMOS circuitry. Domino MUXes are used to select the data decision from the speculative path with the correct H1 and H2 compensation. Static CMOS MUXes are inserted in the DFE feedback paths so that the control logic can apply static feedback bits (H1data, H2data, H3data) during operations such as eye monitoring [6].

While current-integrating summers offer good power efficiency, integrating the analog input signal for 1 UI introduces frequency-dependent loss amounting to 3.9 dB at half-baud frequency [28]. Such loss would be a significant penalty in a receiver intended to equalize high-loss channels. Fig. 14 shows two solutions for eliminating this loss penalty. In the sampled integrating amplifier [Fig. 14(a)], a passgate sample-and-hold (S/H) is placed in front of the amplifier so that it integrates a held signal. This completely eliminates the systematic loss of integration [28], but including a S/H has a couple of significant drawbacks. The kT/C noise of a low-capacitance sampler may degrade SNR, and kickback from the sampling switch disturbs the previous stage, which may have difficulty recovering by the next sampling interval (especially at these data rates). The S/H and its associated difficulties are eliminated in the peaked integrating amplifier [Fig. 14(b)]. In this approach, the input stage is peaked with an *RC* degeneration network, whose values are chosen to provide about 3.9 dB of peaking at half-baud frequency. Because the required *RC* time constant depends on the half-baud frequency, the degeneration capacitor must be switched to support different data rates. This peaked integrator approach has been adopted here for the DFE summers.

The schematic of the DFE summer is shown in Fig. 15(a). The H3-H15 tap circuits employ a return-to-zero (RZ) structure [6]. Because the glitches on the tap tail nodes occur every clock cycle and are independent of data pattern, this RZ structure generates accurate integration currents with virtually no positive setup time requirement on the DFE feedback signal. For the ith

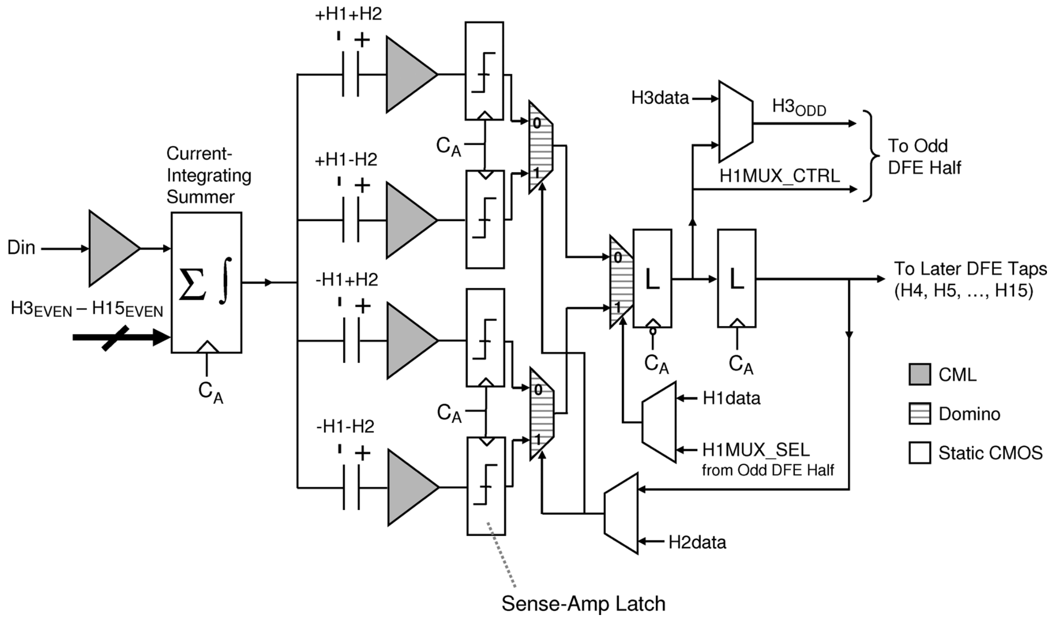

Fig. 13. Block diagram of even DFE half.

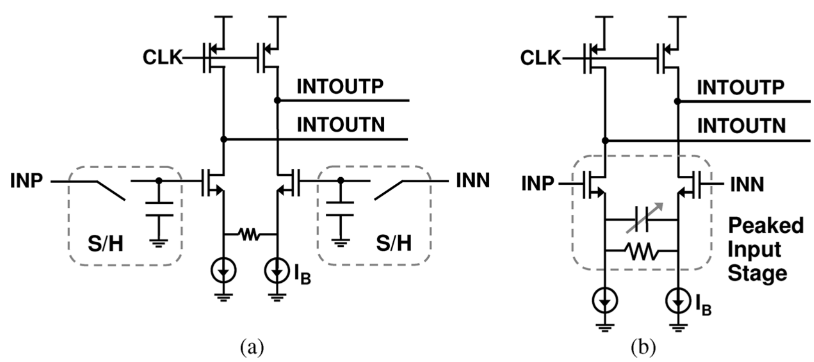

Fig. 14. Two solutions to integrator loss. (a) Sampled integrating amplifier. (b) Peaked integrating amplifier.

DFE tap $(\mathrm{H}i)$, the sign of its coefficient can be either positive or negative depending on whether programmable tail current $I_{\mathrm{H}i\mathrm{P}}$ is greater or less than programmable tail current $I_{\mathrm{H}i\mathrm{N}}$. At high data rates, integration times are short, and integrator gain is reduced. Boosting integration currents can help restore gain but causes excessive common-mode drop on the summer output, which degrades linearity. This limitation is overcome by introducing a PMOS injector circuit which is capacitively coupled to the summer output nodes. As shown in the timing diagram of Fig. 15(b), nodes $\mathrm{INTOUTP_{PMOS}}$ and $\mathrm{INTOUTN_{PMOS}}$ are grounded during integrator reset. During the integration period, the NMOS reset switches inside the PMOS injector are shut off, and currents from sources $I_{\mathrm{B\_PMOS}}$ are driven (through coupling capacitors) into the summer output nodes, which raises their common-mode. A similar PMOS injector is discussed in [29]. As proposed in [25], a calibration circuit based on a replica integrator is used to set all of the summer bias currents (including $I_{\mathrm{B\_PMOS}}$) so that the desired output common-mode is obtained over process variations and different data rates.

The switches inside the box labeled Capacitive Level Shifters are used to establish the dc offset voltages stored across the series capacitors. Because the voltages stored on the capacitors are only modified slowly (on the time scale of DFE adaptation), the charging circuitry for the capacitors can afford to be relatively sluggish, so its switches are minimum size devices (for small parasitic loading), and its bias voltage generators have relatively high output impedances (for low power dissipation). It is important that data-dependent signals not modulate the voltages stored on the capacitors, for such errors could create ISI with a time duration (due to sluggish recharging) which exceeds the correction range of the 15-tap DFE.

During integrator reset, the left sides of the capacitors are pulled up to the supply. Because the bias currents (I_{B}) of the input stage are not shut off during integrator reset, the reset of nodes INTOUTP and INTOUTN may be incomplete. To prevent these data-dependent errors (e.g., 20 mV differential) from modulating the capacitor voltages, the left sides of the capacitors are pulled up to the supply by dedicated switches in an Enhanced Reset Circuit, which ensures proper nulling of the differential voltage between nodes VSWP and VSWN. After such nulling has occurred, the right sides of the capacitors are connected to bias voltages (VBP and VBN) representing the desired H1, H2, and offset compensation. As indicated in the timing diagram, intentional skew between the falling edges of clock signals CLK

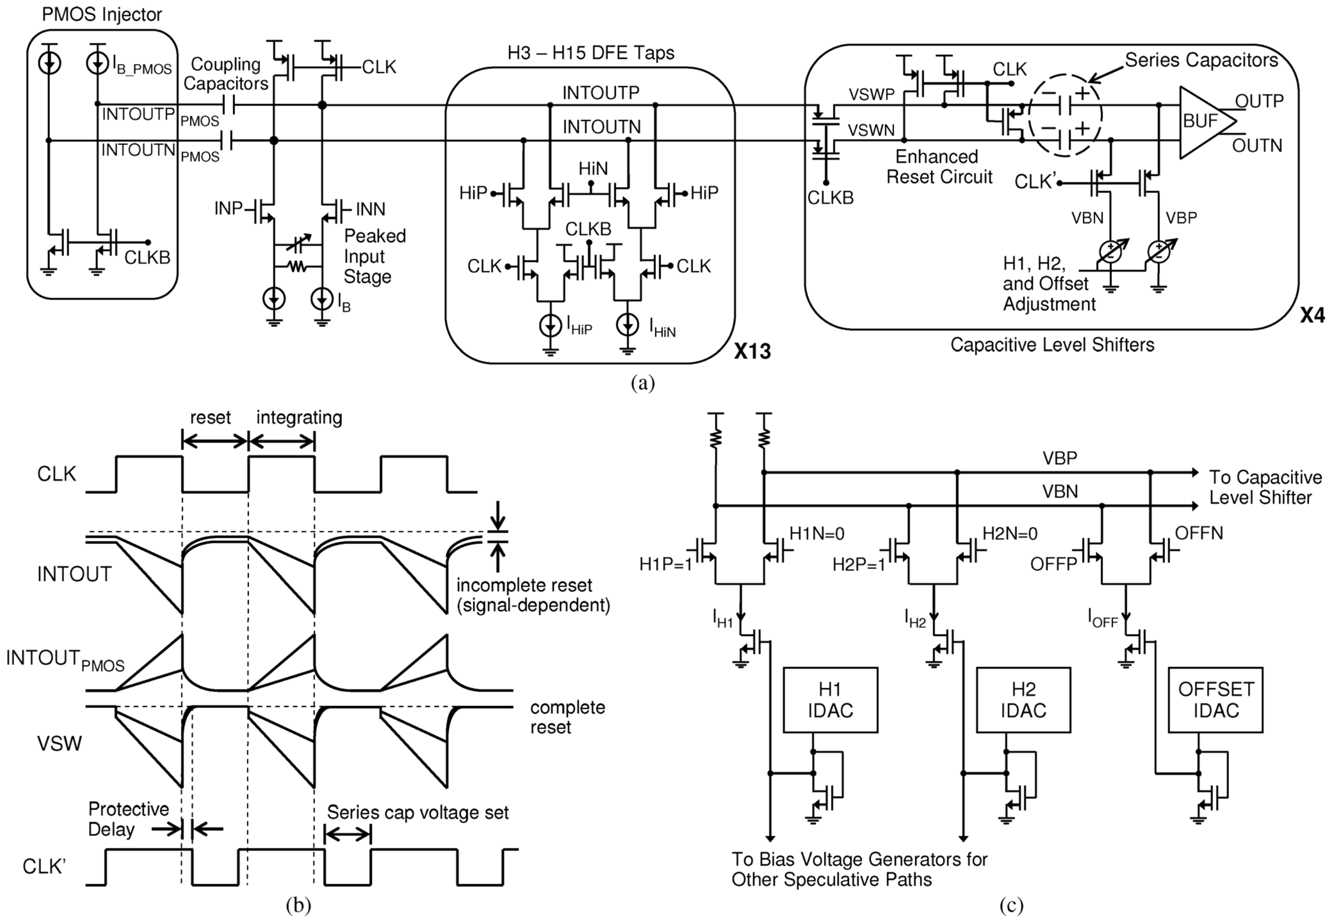

Fig. 15. Current-integrating summer of DFE. (a) Summer schematic. (b) Timing diagram. (c) Bias voltage generator for producing VBP and VBN.

and CLK' provides a protective delay against making these connections too early (avoiding data-dependent disturbances of bias voltages VBP and VBN). During the integration period, CLK' is high, so the voltages (or charges) stored on the series capacitors are held constant until the next charging cycle (leakage currents are negligible at data rates above a few Gb/s). As shown in Fig. 15(c), voltages VBP and VBN are generated across load resistors by summing together currents from the IDACs used to program the H1 tap, H2 tap, and offset compensation. This biasing arrangement allows the H1 and H2 IDACs to be shared among multiple speculative paths.

V. Experimental Results

Fig. 16 shows a micrograph of the four-port I/O core, which was fabricated in a 32-nm SOI CMOS process. With the PLL overhead amortized over four lanes, the area of a single transmitter/receiver pair is 0.81 mm^2. The test chip holding the four-port I/O core was attached with controlled collapse chip connection (C4) technology to a flip-chip plastic ball grid array (FCPBGA) package, which was then mounted on a socketed evaluation board.

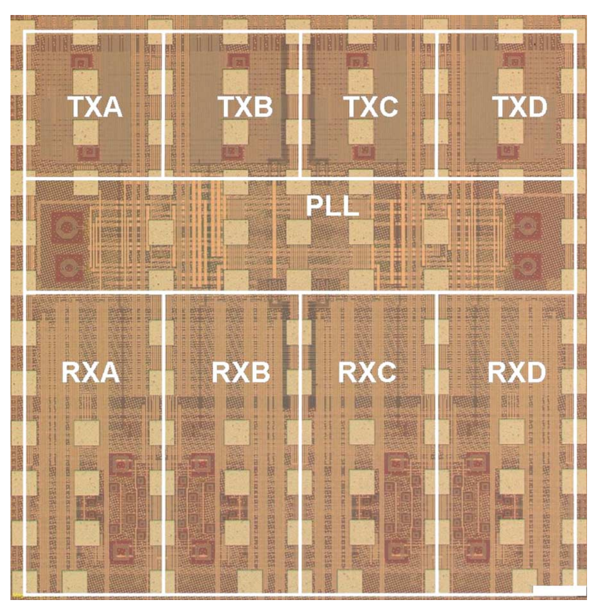

Fig. 16. Micrograph of four-port I/O core.

In addition to the fully integrated I/O core, a separate breakout test site containing one transmitter was built. The breakout test site was not packaged but was characterized on a wafer probe station with high-bandwidth probes. In this setup, a differential half-rate clock is provided externally from a low-noise clock synthesizer, and an on-chip programmable pattern generator supplies data to the transmitter. The characterization of the transmitter with the breakout test site is

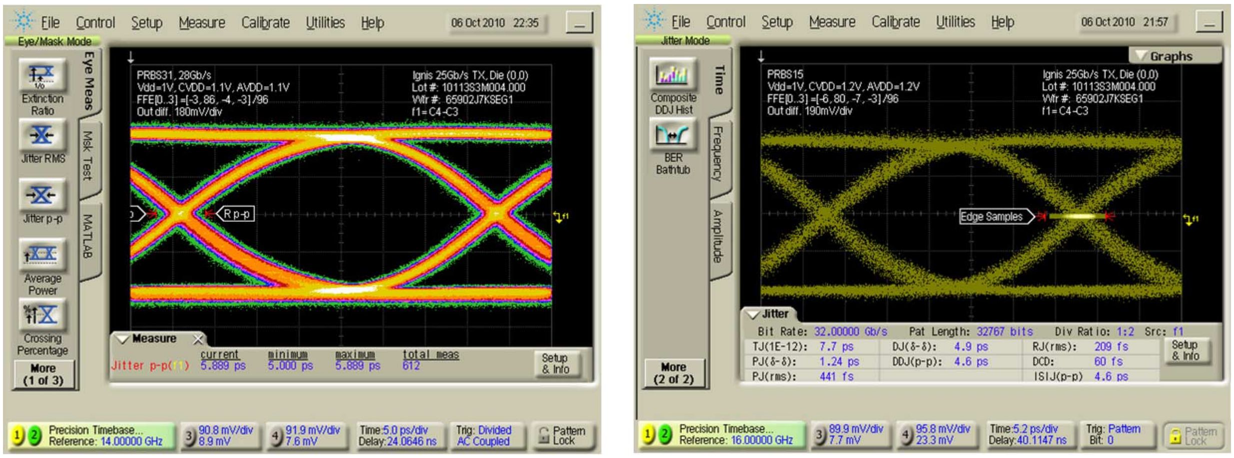

(a) (b)

Fig. 17. Measured differential output eye diagrams of transmitter on breakout test site. (a) 28-Gb/s PRBS31 data pattern $(\text{vertical scale} = 180\ \text{mV/div.})$. (b) 32-Gb/s PRBS15 data pattern $(\text{vertical scale} = 190\ \text{mV/div.})$.

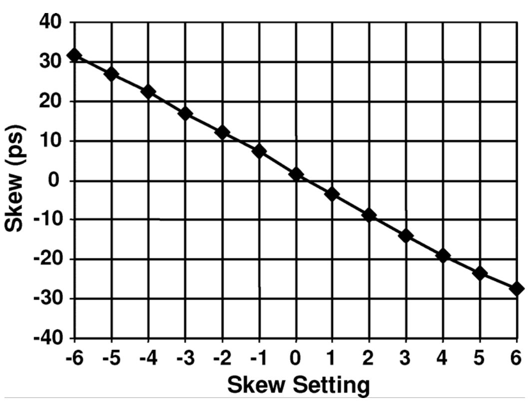

Fig. 18. Measured skew between true and complement outputs of transmitter as function of skew setting.

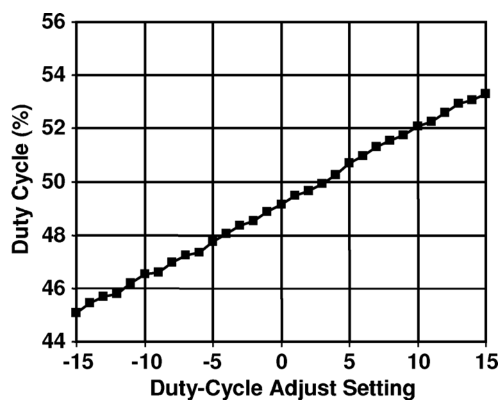

Fig. 19. Measured duty cycle at transmitter output as function of duty-cycle adjust setting with 25-Gb/s 0-1-0-1 sequence.

discussed next, and then the measurement results of the fully integrated I/O core are presented.

A. Transmitter Measurements With Breakout Test Site

Fig. 17(a) shows the measured differential data eye of a PRBS31 pattern at 28 Gb/s with a 1.1-V supply voltage. Measured peak-to-peak (p-p) jitter is below 6 ps. To overcome cable losses, the FFE tap coefficients have been set to $[-3, 86, -4, -3]/96$. Taking into account the de-emphasis factor of 76/96, the transmitter output amplitude is 1.046 V peak-to-peak differential (Vppd), which is close to the 1.1-V supply voltage (the ideal amplitude for an impedance-matched voltage-mode driver). The measured power consumption of the transmitter is 217 mW at 28 Gb/s with a 1.1-V supply. An output eye diagram with a 32-Gb/s PRBS15 data pattern is displayed in Fig. 17(b). With a 1.2-V supply voltage and FFE tap coefficients of $[-6, 80, -7, -3]/96$, the transmitter output amplitude is 1.14 Vppd. Total jitter (TJ) is extrapolated to be 7.7 ps p-p at a bit error rate (BER) of 10^{-12}, of which 4.6 ps p-p stems from ISI.

Fig. 18 shows the measured differential skew between the true and complementary outputs as a function of the digital skew setting. Differential skew up to $+/-28$ ps can be compensated, corresponding to a maximum cable length difference of about 5 mm. The measured duty cycle of a 25-Gb/s 0-1-0-1 sequence as a function of the digital adjustment setting is shown in Fig. 19. The tuning range is about $+/-3.5\%$, which is on the low end of that predicted by corner simulations but still sufficient to cover the DCD due to device mismatches.

B. Measurements With Four-Port I/O Core

Fig. 20 shows a 28-Gb/s differential output eye diagram generated by a transmitter on the fully integrated I/O core. Even with FFE de-emphasis, the ISI at eye center is visibly greater than in the breakout test site measurement [Fig. 17(a)]. This small loss of eye quality is accurately predicted by a link simulation tool with S-parameter models for the package, evaluation card, and cabling. On the other hand, the measured random jitter (RJ) at the transmitter output is 450 fs rms, about twice that predicted in circuit simulations. Such jitter is not a fundamental limitation of the LC-VCO-based PLLs in this technology, as significantly lower RJ ($\sim$250 fs rms) has been recently achieved with an updated version of the PLL (including layout refinements).

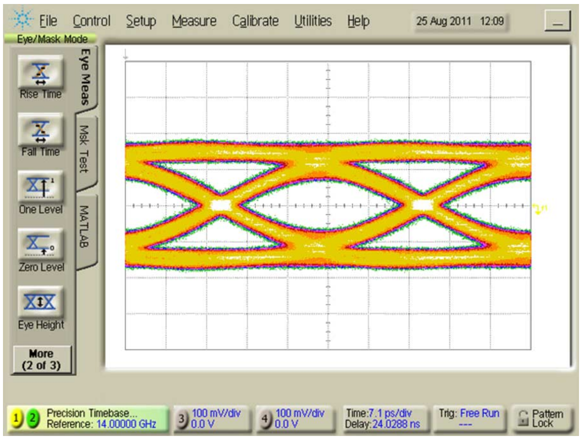

Fig. 20. Measured 28-Gb/s differential output eye diagram of transmitter on fully integrated I/O core (vertical scale $=$ 200 mV/div.). The data pattern is PRBS31.

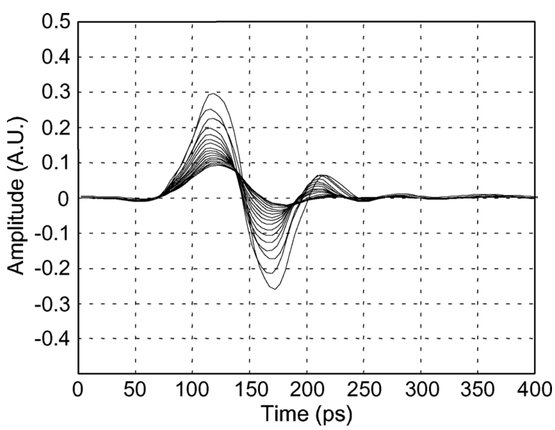

(a)

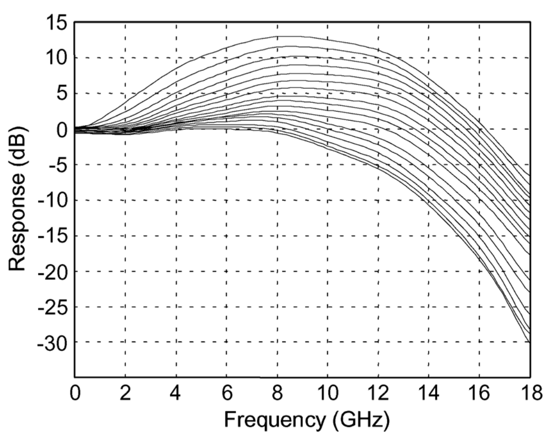

(b)

Fig. 21. Measured receiver responses with 17 different peaking settings. (a) Time-domain responses to single "one" bit at 28 Gb/s. (b) Derived frequency responses.

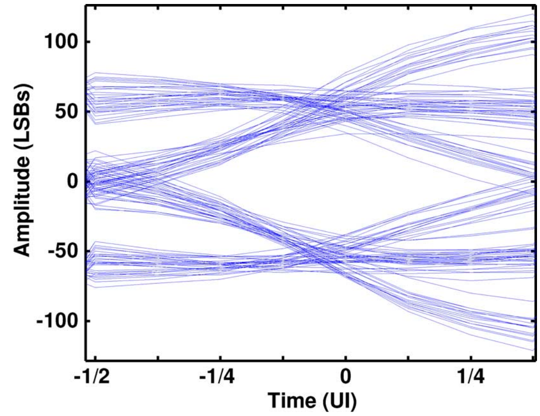

Fig. 22. Measured internal eye of receiver demonstrating equalization of 38-dB loss channel at 25 Gb/s.

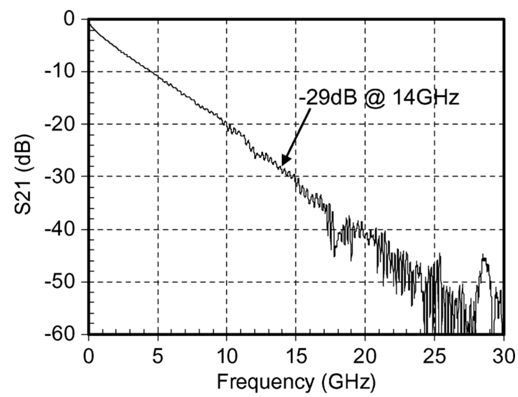

(a)

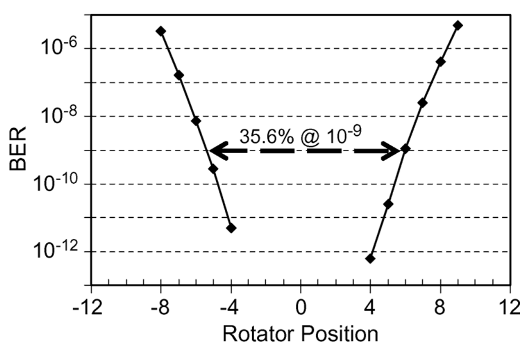

(b)

Fig. 23. Equalization experiment with test channel including 15-in trace on PCB. (a) S-parameters of 15-in PCB trace, interconnect cables, and evaluation card. (b) Equalized bathtub curve with 28-Gb/s PRBS31 data pattern.

Receiver characteristics have been studied by applying clean data to its inputs. Oscilloscope measurements of data signals transmitted across calibration traces on the evaluation card are used to set the FFE tap coefficients so that the loss of a short channel (including transmitter package) is equalized. Loss of the receiver package, however, is not corrected for. With (almost) clean data and the VGA set to maximum gain, the measured input sensitivity of the receiver at 28 Gb/s is 15 mVppd at a BER of 10^{-9}. An internal eye monitor of the receiver is used to measure its transient responses. Fig. 21(a) shows the receiver responses to a single "one" bit at 28 Gb/s with 17 different settings of the peaking amplifier. Fourier transforms can be used to derive the frequency responses of the receiver. Since each response in Fig. 21(a) is the convolution of the receiver impulse response with a 1-UI-wide pulse, its Fourier transform $R(f) = H(f)P(f)$, where $H(f)$ is the receiver frequency response, and $P(f) = \sin(\pi \bullet f \bullet \mathrm{UI})/(\pi \bullet f \bullet \mathrm{UI})$. Calculating $R(f)$ for each response in Fig. 21(a) and solving for $H(f)$ yields the frequency responses shown in Fig. 21(b). The maximum

TABLE I
PERFORMANCE SUMMARY FOR COMPLETE TRANSCEIVER

Technology	IBM 32-nm SOI CMOS
Data Rate	14–28.05 Gb/s
TX Equalization	4-Tap FFE
RX Equalization	Two-Stage Peaking Amp and 15-Tap DFE
TX Output Swing	950 mVppd
Random Jitter of TX Output @ 28 Gb/s	450 fs RMS
RX Input Sensitivity @ 28 Gb/s	15 mVppd (for BER=10^{-9})
Horizontal Eye Opening @ BER=10^{-9}, 35 dB Loss Channel	35.6% (28-Gb/s PRBS31 Data Pattern)
Area/Lane (TX+RX+PLL/4)	0.81 mm^2
Supply Voltages	1.2 V (for PLL), 1.05 V (for TX and RX), 0.85 V (VDD)
Power/Lane (TX+RX+PLL/4) @ 28 Gb/s	693 mW

peaking is about 11 dB at 12.5 GHz and 7 dB at 14 GHz. Accounting for the receiver package loss, the maximum peaking at 14 GHz is close to 10 dB.

The internal eye monitor is also used to measure the equalized eye of the receiver (i.e., after DFE taps are applied). Fig. 22 shows the equalized eye in a 25-Gb/s experiment in which the channel loss is 33 dB (38 dB with transmitter and receiver package losses). With a least significant bit (LSB) value of about 2.5 mV, the vertical eye opening exceeds 150 mVppd.

Finally, Fig. 23 presents the results of a 28-Gb/s equalization experiment with a test channel including a 15-in trace on PCB (Megtron 6 material with HTE4P foil), 3.7 in of evaluation card traces, and interconnect cables (12 in from evaluation card to PCB and 12 in from PCB to evaluation card) with mini-SMP connectors. S-parameter measurements [Fig. 23(a)] of this test channel show a loss of 29 dB at 14 GHz; the losses of the transmitter and receiver packages bring the total to 35 dB. Fig. 23(b) shows the equalized bathtub curve with a 28-Gb/s PRBS31 data pattern. The horizontal eye opening is 35.6% at a BER of 10^{-9}, and operation is error-free $(\mathrm{BER} < 10^{-13})$ at eye center. The measured power consumption is 693 mW per lane (211 mW for transmitter, 392 mW for receiver, and 90 mW for amortized PLL). (This experiment was conducted at nominal temperature and supply voltages, and the process split of the test chip was also close to nominal.) The use of known power management schemes could reduce this power but was not exercised for this prototype. As an example, both DFE banks were always powered up during the experiments. If one of the DFE banks were shut off when it is not needed, a conservative estimate of the power savings would be 40 mW. The performance of the integrated transceiver is summarized in Table I.

VI. CONCLUSION

This paper has presented a 4-tap FFE/15-tap DFE transceiver in 32-nm SOI CMOS technology with a maximum data rate of 28 Gb/s, which is almost two times higher than that of other fully integrated backplane transceivers published to date. Key circuit techniques have been developed to achieve such data rates. The proposed SST driver topology eliminates the main speed bottlenecks of previous half-rate designs. The peaking amplifier based on an active feedback structure provides greater high-frequency gain than a conventional zero-peaked differential amplifier. The use of capacitive level-shifters facilitates efficient implementation of DFE architectures with multiple speculative taps. The equalization performance of the transceiver at 28 Gb/s has been demonstrated with error-free operation over a channel with 35-dB loss.

ACKNOWLEDGMENT

The authors would like to thank all of those from IBM Research, IBM STG, and Miromico who helped support this work. A partial list of contributors includes R. Denzer, M. Ruegg, J. Reddy, S. Tyagi, P. Carlile, B. Voegeli, A. Malladi, Z. Jin, V. Moy, M. Wielgos, R. Tompkins, J. Boettler, P. Kelly, H. Park, S. Duffield, H. Xu, G. Gangasani, W. Queen, A. Brouillette, J. Rockrohr, E. Tressler, K. Heilmann, P. Elakkumanan, K. Kramer, R. Reutemann, J. Tierno, and M. Soyuer.

REFERENCES

[1] Common electrical I/O (CEI)—Electrical and jitter interoperability agreements for 6G+ bps, 11G+ bps and 25G+ bps I/O. Optical Internetworking Forum, Sep. 2011 [Online]. Available: http://www.oiforum.com/public/documents/OIF_CEI_03.0.pdf

[2] Fibre Channel Solutions Guide Book 2010. Fibre Channel Industry Association (FCIA) [Online]. Available: http://www.fibrechannel.org/documents

[3] J. D'Ambrosia, IEEE 802.3WG Closing Plenary Report, IEEE P802.3bj 100 Gb/s Backplane and Copper Cable Task Force, Mar. 2012 [Online]. Available: http://www.ieee802.org/3/minutes/mar12/0312_bj_close_report.pdf

[4] InfiniBand Roadmap. InfiniBand Trade Association (IBTA) [Online]. Available: http://www.infinibandta.org/content/pages.php?pg=technology_overview

[5] 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*, Feb. 2011, pp. 350–351.

[6] G. R. Gangasani *et al.*, "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," in *Proc. IEEE CICC*, Sep. 2011.

[7] F. Zhong *et al.*, "A 1.0625 $\sim$ 14.025 Gb/s multi-media transceiver with full-rate source-series-terminated transmit driver and floating-tap decision-feedback equalizer in 40 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 46, no. 12, pp. 3126–3139, Dec. 2011.

[8] 2011 Edition, International Technology Roadmap for Semiconductors (ITRS). [Online]. Available: http://www.itrs.net/Links/2011ITRS/Home2011.htm

[9] T. Beukema, M. Sorna, K. Selander, S. Zier, B. L. Ji, P. Murfet, J. Mason, W. Rhee, H. Ainspan, B. Parker, and M. Beakes, "A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2633–2645, Dec. 2005.

[10] R. Payne *et al.*, "A 6.25-Gb/s binary transceiver in 0.13-μm CMOS for serial data transmission across high loss legacy backplane channels," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2646–2657, Dec. 2005.

[11] D. G. Kam *et al.*, "Is 25 Gb/s on-board signaling viable?," *IEEE Trans. Adv. Packag.*, vol. 32, no. 5, pp. 328–344, May 2009.

[12] J. Bulzacchelli *et al.*, "A 28 Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32 nm SOI CMOS technology," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 324–325.

[13] C. Menolfi, J. Hertle, T. Toifl, T. Morf, D. Gardellini, M. Braendli, P. Buchmann, and M. Kossel, "A 28 Gb/s source-series terminated TX in 32 nm CMOS SOI," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 334–335.

[14] C. Zhong, C. Y. Liu, W. Jin, A. Malipatil, G. Tang, and F. Y. Zhong, "Design considerations in high speed SerDes (25 Gbps)," in *Proc. IEC DesignCon*, Feb. 2008.

[15] T. Toifl *et al.*, "A 2.6 mW/Gbps 12.5 Gbps RX with 8-tap switched-capacitor DFE in 32 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 4, pp. 897–910, Apr. 2012.

[16] M. Kossel, C. Menolfi, J. Weiss, P. Buchmann, G. von Bueren, L. Rodoni, T. Morf, T. Toifl, and M. Schmatz, "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.

[17] S. McMorrow and C. Heard, "The impact of PCB laminate weave on the electrical performance of differential signaling at multi-gigabit data rates," in *Proc. IEC DesignCon*, Feb. 2005.

[18] C. Menolfi, T. Toifl, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. Schmatz, "A 16 Gb/s source-series terminated transmitter in 65 nm CMOS SOI," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 446–447.

[19] K. Bernstein and N. J. Rohrer, *SOI Circuit Design Concepts*. Boston, MA: Kluwer, 2000, ch. 3–4.

[20] S. Galal and B. Razavi, "40-Gb/s amplifier and ESD protection circuit in 0.18- μm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2389–2396, Dec. 2004.

[21] S. Kaeriyama *et al.*, "A 40 Gb/s multi-data-rate CMOS transmitter and receiver chipset with SFI-5 interface for optical transmission systems," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3568–3579, Dec. 2009.

[22] J. N. Burghartz, D. C. Edelstein, M. Soyuer, H. A. Ainspan, and K. A. Jenkins, "RF circuit design aspects of spiral inductors on silicon," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2028–2034, Dec. 1998.

[23] 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.

[24] S. Galal and B. Razavi, "10-Gb/s limiting amplifier and laser/modulator driver in 0.18- μm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2138–2146, Dec. 2003.

[25] M. Park, J. Bulzacchelli, M. Beakes, and D. Friedman, "A 7 Gb/s 9.3 mW 2-tap current-integrating DFE receiver," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 230–231.

[26] 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*, Feb. 2009, pp. 180–181.

[27] D. Z. Turker, A. Rylyakov, D. Friedman, S. Gowda, and E. Sánchez-Sinencio, "A 19 Gb/s 38 mW 1-tap speculative DFE receiver in 90 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2009, pp. 216–217.

[28] T. O. Dickson, J. F. Bulzacchelli, and D. J. Friedman, "A 12-Gb/s 11-mW half-rate sampled 5-tap decision feedback equalizer with current-integrating summers in 45-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1298–1305, Apr. 2009.

[29] A. Agrawal, J. Bulzacchelli, T. Dickson, Y. Liu, J. Tierno, and D. Friedman, "A 19 Gb/s serial link receiver with both 4-tap FFE and 5-tap DFE functions in 45 nm SOI CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 134–135.

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 from the Massachusetts Institute of Technology (MIT), Cambridge, in 1990, 1990, and 2003, respectively, all in electrical engineering.

From 1988 to 1990, he was a cooperative student with 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 he has focused on the design of mixed-signal CMOS circuits for high-speed data communications as well as high-performance integrated voltage regulators. He also maintains strong interest in the design of circuits in more exploratory technologies. He holds 13 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 corecipient of the Beatrice Winner Award for Editorial Excellence at the 2009 ISSCC, coauthored the IEEE JOURNAL OF SOLID-STATE CIRCUITS article awarded Best Paper of 2009, and was a corecipient of the 2011 IEEE Custom Integrated Circuits Conference (CICC) Best Regular Paper Award.

Christian Menolfi (S'97–M'99) was born in St. Gallen, Switzerland, in 1967. He received the Dipl. Ing. degree and the Ph.D. degree in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zürich, in 1993 and 2000, respectively.

From 1993 to 2000, he was with the Integrated Systems Laboratory, ETH Zürich, as a Research Assistant, where he worked on highly sensitive CMOS VLSI data acquisition circuits for silicon-based microsensors. Since September 2000, he has been with the IBM Zürich Research Laboratory, Rüschlikon, Switzerland, where he has been involved with multi-gigabit low-power communication circuits in advanced CMOS technologies.

Troy J. Beukema (M'07) 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, Yorktown Heights, NY, where he has focused on both wireless and wireline system designs in the Communication Technologies 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 data path system designs for backplane I/Os operating from 6 to 28 Gb/s spanning CMOS technology generations from 130 to 32 nm. His ongoing research interests include both ultrahigh-data-rate SerDes and low-power ADC-based I/O core designs. He holds 14 U.S. patents.

Mr. Beukema was a corecipient of the Lewis Winner Award for Outstanding Paper at the 2006 IEEE International Solid-State Circuits Conference (ISSCC), the 2008 DesignCon Paper Award, and the 2011 IEEE Custom Integrated Circuits Conference (CICC) Best Regular Paper Award.

Daniel W. Storaska (S'95–M'00) received the B.S. degree in physics from Miami University, Oxford, OH, in 1995, the B.S. degree in electrical engineering from Columbia University, New York, NY, in 1997, and the M.E. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2002.

Mr. Storaska joined IBM's DRAM Development Alliance, a joint venture between IBM, Siemens, and Toshiba, in 1997. He was a member of the design teams which successfully developed the world's smallest 1 Gb DRAM, as well as one with the world's fastest random access cycle time. In 2001, he joined the IBM High-Speed Serial (HSS) design organization, Hopewell Junction, NY, where he has been developing a variety of characterization techniques and circuits required for high-speed serial communications, from models for studying signal integrity to methods for calibrating PLLs. His current interests are in low-power protocols. He holds 18 U.S. patents.

Jürgen Hertle (S'00–M'05) received the Dipl. Ing. degree (with honors) in electrical engineering from the University of Erlangen, Erlangen, Germany, in 1997, and the Ph.D. degree in technical science from the Swiss Federal Institute of Technology (ETH), Zürich, in 2004.

In 2002, he joined ACP AG, Zollikon, Switzerland, where he designed a folding and interpolating ADC as part of his Ph.D. thesis. In 2004 he joined Photonfocus AG, Lachen, Switzerland, where he was responsible for the project management and design of high-speed CMOS image sensors for industrial applications. Since 2007, he has been with Miromico AG, Zürich, Switzerland, where he has contributed to the design of various high-performance analog circuits, including ADCs, power management units, and transmitters and receivers for several high-speed serial links.

David R. Hanson received the B.S.E.E. degree from the Rochester Institute of Technology, Rochester, NY, and the M.S.E.E. degree from Syracuse University, Syracuse, NY.

He joined IBM in 1985 and was previously engaged in the design of two generations of high-speed cache SRAM, three generations of DRAM, and two generations of low-power embedded DRAM. He is currently a Senior Engineer in IBM's East Fishkill, NY facility, where he is concentrating upon mixed-signal transmitter designs for serial links operating up to 28 Gbps. He holds 31 U.S. patents.

Ping-Hsuan Hsieh (S'02–M'11) was born in Taipei, Taiwan. She received the B.S. degree from National Taiwan University, Taipei, Taiwan, in 2001, and the M.S. and Ph.D. degrees from the University of California, Los Angeles, in 2004 and 2009, respectively, all in electrical engineering.

She was an Intern with Texas Instruments, Dallas, TX, in the summers 2004–2006, where she was involved in a pilot study to investigate the feasibility and limitations of traditional analog phase-locked loop architectures. From 2009 to 2011, she was with the IBM T.J. Watson Research Center, Yorktown Heights, NY, where she worked on the design of mixed-signal integrated circuits for high-speed serial data communication. In 2011, she joined the Electrical Engineering Department of National Tsing Hua University, Hsinchu, Taiwan, where she is currently an Assistant Professor. Her research interests focus on mixed-signal integrated circuit designs for high-speed electrical data communications, clocking and synchronization systems, and energy-harvesting systems for wireless sensor networks and machine-to-machine applications.

Sergey V. Rylov (M'05) received the M.S. and Ph.D. degrees in physics from Moscow State University, Moscow, Russia, in 1984 and 1989, respectively.

From 1987 to 1991, he was a Research Scientist with the Laboratory of Cryoelectronics, Moscow State University, Moscow, Russia, in the field of superconducting Josephson junction microelectronic circuits and physically reversible computers. From 1991 to 1998, he was with HYPRES, Inc., Elmsford, NY, working as a designer of superconducting digital and analog circuits, including high-resolution and flash ADCs, single-flux-quantum logic devices, and analog amplifiers using dc SQUIDs. In 1998 he became a Research Staff Member with the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he designs high-speed mixed-signal circuits for multi-gigabit/s CMOS communications ICs, with a particular emphasis on design of clock phase rotators. He holds 16 U.S. patents.

Dr. Rylov was a corecipient of the 2011 IEEE Custom Integrated Circuits Conference (CICC) Best Regular Paper Award.

Daniel Furrer, photograph and biography not available at the time of publication.

Daniele Gardellini (M'01) received the M.Sc. degree in electrical engineering from the Politecnico di Milano, Milan, Italy, in 2000.

From 2000 to 2005, he was with Accent, working as an IC analog design Engineer on the development of analog and mixed-signal sensor interface circuits. Between 2005 and 2009, he was with Infineon Technologies AT, where he was involved in the development of hard-disk read/write-channels, PCI-e and SATA serial interfaces, and various sensor analog front-ends, in particular for optical applications. In 2009, he joined Miromico AG, Zurich, Switzerland, where his focus is on full-custom analog and mixed-signal IC design, in particular for high-speed SerDes applications and high accuracy analog functions.

Andrea Prati received the Dipl. Ing. degree in microengineering from the Swiss Federal Institute of Technology in Lausanne (EPFL), Lausanne, Switzerland, in 2005.

Since 2005, he has been with Miromico AG, Zurich, Switzerland, where he has been involved in the design of analog and mixed-signal circuits, in particular for high-speed serial links.

Thomas Morf (S'89–M'96–SM'09) was born on April 4, 1961, in Zürich, Switzerland. He received the B.S. degree from the Zürich University of Applied Science, Switzerland, in 1987, the M.S. degree in electrical and computer engineering from the University of California, Santa Barbara, in 1991, and the Ph.D. degree from the Swiss Federal Institute of Technology, Zürich, Switzerland, in 1996.

From 1989 to 1991, he was a Research Assistant with the University of California, Santa Barbara, performing research in the field of active microwave inductors and digital GaAs circuits. In 1991 he joined the Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, where in his doctoral work he investigated circuit design and processing for high-speed optical links on GaAs using epitaxial lift-off techniques. In 1996, he transferred to the Electronics Laboratory of ETH, where he led a research group in the area of InP-HBT circuit design and technology. In 1999 he joined the IBM Zürich Research Laboratory, Rüschlikon, Switzerland. His current research interests include ESD circuit protection, all aspects of electrical and optical high-speed high-density interconnects, THz antennas and detectors, and high-speed and microwave circuit design. He has coauthored more than 80 papers.

Vivek Sharma (S'99–M'04) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology (IIT), Kanpur, India, in 2000, and the M.S. degree in electrical engineering and computer sciences from Oregon State University, Corvallis, in 2003.

From 2000 to 2001, he was with ST Microelectronics-India, where he worked on the design of charge-pump PLLs. From 2001 to 2004, he was a Research Assistant with Oregon State University, Corvallis, where he conducted investigations into portable biosensors and low-power A/D converters for environmental monitoring applications. In the summer of 2003, he was with Berner Fachhochschule, Burgdorf, Switzerland, developing algorithms for real-time processing of music signals. From 2004 to 2009, he was with austriamicrosystems AG, Rapperswil, Switzerland, developing ASICs with a focus on sensor interfaces. Since 2009, he has been with Miromico AG, Zurich, Switzerland, where his focus is on high-speed I/O design. He holds two patents with one pending.

Ram Kelkar received the B.S. and M.S. degrees in electrical engineering from the Indian Institute of Technology in 1974 and 1976, respectively, in and the Ph.D. degree in electrical engineering in power electronics from Virginia Tech, Blacksburg, VA, in 1982.

From 1982 to 1984, he was with AT&T Bell Laboratories, Parsippany, NJ, working in the area of power electronics. He joined IBM in 1984 and worked until 1989 in IBM Power Systems. He is currently with IBM Microelectronics, Essex Junction, VT, carrying out development of macros focusing on PLLs, drivers, and references. He has developed macros in eight technologies, and IBM has shipped over four million units of his designs. He holds 30 U.S. patents and has authored and coauthored 27 publications in refereed IEEE journals and conferences.

Herschel A. Ainspan received the B.S. and M.S. degrees in electrical engineering from Columbia University, New York, in 1989 and 1991, respectively.

In 1989 he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he has been involved in the design of mixed-signal and RF IC's for high-speed data communications.

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

In 1982 he joined IBM, Poughkeepsie, NY, as an Eectrical 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. Since 2001, he has been a Design Engineer working on various high-speed SerDes interfaces, with a primary focus on development and implementation of advanced equalization architectures and related logic algorithms.

Mr. Kelly was a corecipient of the 2011 IEEE Custom Integrated Circuits Conference (CICC) Best Regular Paper Award. He holds four U.S. patents.

Leonard R. Chieco received the B.S. degree in electrical and computer engineering from Clarkson College, Potsdam, NY, in 1981.

Upon graduation, he joined IBM, Poughkeepsie, NY, where he was a member of the IO Processor design team for System 3090. In 1994, he transferred to IBM, East Fishkill, NY, where he has been involved in Phy and Protocol transmitter logic architecture and design for IEEE 1394, FibreChannel, and SATA standards. Most recently, he has been working on IBM's High-Speed Serial I/O cores.

Glenn A. Ritter, photograph and biography not available at the time of publication.

John A. Sorice received the A.A.S degree from Westchester Community College, Valhalla, NY, in 1983.

In 1984, he joined IBM Microelectronics Division, Hopewell Junction, NY, where he has been involved in mask design for custom interface products, RF products, GPS products and packaging, and high-speed data communications. His primary role is currently Lead Mask Designer for 28 Gb/s transmitters and DFE receivers.

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

He is a Senior Engineer with IBM, Hopewell Junction, NY, and, since 1995, has worked on high-speed SerDes analog circuit design and characterization for various communication standards. He holds four U.S. patents.

Mr. Garlett was a corecipient of the 2011 IEEE Custom Integrated Circuits Conference (CICC) Best Regular Paper Award.

Robert Callan received the B.S. degree from the University of Pennsylvania, Philadelphia, in 2007, and the M.S. degree from the University of Southern California, Los Angeles, in 2008, both in electrical engineering.

Since 2008, he has worked in the IBM Microelectronics High-Speed Serial testing group in East Fishkill, NY, and San Jose, CA, focusing on circuit characterization, FFE training, and on-chip signal integrity diagnostics.

Matthias Brändli received the Dipl. Ing. degree in electrical engineering from the Swiss Federal Institute of Technology, Zürich, Switzerland, in 1997.

From 1998 to 2001, he was with the Integrated Systems Laboratory, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, working on deep-submicron technology VLSI design challenges, digital video image processing for biomedical applications, and testability of CMOS circuits. In 2001, he joined the Microelectronics Design Center of ETH Zürich, where he was involved in numerous digital and mixed-signal ASIC design projects, worked on EDA design automation, and contributed to teaching. In 2008 he joined the IBM Zürich Research Laboratory, Rüschlikon, Switzerland, where he has been working on multi-gigabit/s, low-power communication circuits in advanced CMOS technologies.

Peter Buchmann was born in Zürich, Switzerland, in 1953. He received the diploma in experimental physics and Ph.D. degree in physics from the Federal Institute of Technology, Zürich, Switzerland, in 1978 and 1987, respectively.

From 1978 to 1981, he was involved in surface physics studies. From 1981 to 1985, he was working in the field of integrated optics in the group of Applied Research at the Federal Institute of Technology, Zürich, Switzerland. He was engaged in the technology, design, and characterization of III–V semiconductor waveguide devices, electrooptic modulators, and switches. In 1985, he joined the IBM Zürich Research Laboratory, Rüschlikon, Switzerland, where he has been engaged in MESFET technology and in the process technology of III–V semiconductor lasers. In particular, he was involved in research on dry-etching techniques and opto-electronic integration. Since 1994, he has been involved in the design and implementation of VLSI chips for communication applications in the field of ATM, SONET/SDH, and network processors. His most recent work includes circuit design for high-speed I/O and link technology.

Dr. Buchmann is a member of the Swiss Physical Society.

Marcel Kossel (S'99–M'02–SM'09) received the Dipl. Ing. and Ph.D. degrees in electrical engineering from the Swiss Federal Institute of Technology, Zürich, Switzerland, in 1997 and 2000, respectively.

In 2001 he joined the IBM Zürich Research Laboratory, Rüschlikon, Switzerland, where he has been involved in analog circuit design for high-speed serial links. His research interests include analog circuit design and RF measurement techniques. He has also performed research in the field of microwave tagging systems and radio-frequency identification systems.

Thomas Toifl (S'97–M'99–SM'09) received the Dipl.-Ing. degree and Ph.D. degree (with highest honors) from the Vienna University of Technology, Austria, in 1995 and 1999, respectively.

In 1996, he joined the Microelectronics Group of the European Research Center for Particle Physics (CERN), Geneva, Switzerland, where he developed radiation-hard circuits for detector synchronization and data transmission, which were integrated in the four particle detector systems of the new Large Hadron Collider (LHC). In 2001, he joined the IBM Zürich Research Laboratory, Rüchlikon, Switzerland, where he has been working on multigigabit per second, low-power communication circuits in advanced CMOS technologies. In that area, he has authored or coauthored fourteen patents and numerous technical publications. Since July 2008, he has managed the I/O Link technology group at the IBM Zürich Research Laboratory.

Dr. Toifl received the Beatrice Winner Award for Editorial Excellence at the 2005 IEEE International Solid-State Circuits Conference.

Daniel J. Friedman (S'91–M'92) received the Ph.D. degree in engineering science from Harvard University, Cambridge, MA, in 1992.

After completing consulting work at MIT Lincoln Labs and postdoctoral work at Harvard in image sensor design, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1994. His initial work at IBM was the design of analog circuits and air interface protocols for field-powered RFID tags. In 1999, he joined the mixed-signal communications IC design group and turned his attention to analog circuit design for high-speed serializer/deserializer macros. He managed the mixed-signal team from 2000 to 2009, focusing efforts on serial data communication and clock synthesis applications. In 2009, he became manager of the communication circuits and systems group, adding responsibility for teams in millimeter-wave wireless and digital communications IC design. He has authored or coauthored more than 30 technical papers in circuit topics including serial links, PLLs, RFID, and imagers. He holds more than 40 patents. His current research interests include high-speed I/O design, PLL design, and circuit/system approaches to enabling new computing paradigms.

Dr. Friedman was a corecipient of the Beatrice Winner Award for Editorial Excellence at the 2009 ISSCC and the 2009 JSSC Best Paper Award given in 2011. He has been a member of the ISSCC international technical program committee since 2008; he has served as the Wireline subcommittee chair from ISSCC 2012 to the present.