A 6.0-mW 10.0-Gb/s Receiver With Switched-Capacitor Summation DFE

Azita Emami-Neyestanak, *Member, IEEE*, Aida Varzaghani, John F. Bulzacchelli, *Member, IEEE*, Alexander Rylyakov, Chih-Kong Ken Yang, *Member, IEEE*, and Daniel J. Friedman, *Member, IEEE*

***Abstract*—A low-power receiver with a one-tap decision feedback equalization (DFE) was fabricated in 90-nm CMOS technology. The speculative equalization is performed using switched-capacitor-based addition at the front-end sample–hold circuit. In order to further reduce the power consumption, an analog multiplexer is used in the speculation technique implementation. A quarter-rate-clocking scheme facilitates the use of low-power front-end circuitry and CMOS clock buffers. The receiver was tested over channels with different levels of ISI. The signaling rate with BER $< 10^{-12}$ was significantly increased with the use of DFE for short- to medium-distance PCB traces. At 10-Gb/s data rate, the receiver consumes less than 6.0 mW from a 1.0-V supply. This includes the power consumed in all quarter-rate clock buffers, but not the power of a clock recovery loop. The input clock phase and the DFE taps are adjusted externally.**

***Index Terms*—Decision feedback equalization (DFE), interconnects, loop-unrolling, receiver, summation, switched capacitors.**

I. Introduction

MOST ADVANCED electronic systems today require complex architectures that consist of many integrated circuit (IC) chips or modules. The high-bandwidth communication between these ICs and modules is one of the most critical design issues that engineers face. Examples of such systems are computer servers and high-performance multiprocessor systems. Large numbers of high-speed inputs and outputs (IOs) are used to create efficient interfaces between different processors, memory units and other modules located at varying distances from each other. With the continuous scaling of feature sizes in chip manufacturing technology, the speed of on-chip data processing as well as the level of integration will continue to scale. In order to enhance the overall performance of the system, the bandwidth of the interconnections among chips needs to follow the same trend. To achieve this goal, scaling of data rate per IO as well as the number of IOs per chip is necessary. Although the increased switching speed of transistors allows faster transceiver electronics, the scaling of interconnect bandwidth has proven to be very difficult. Limitations imposed by the electrical channel (the signal path from the transmitter to the receiver) are increasing in significance as per IO data rates grow to 10 Gb/s and beyond. One contributor to this effect is that the dielectric and resistive losses of the printed-circuit-board (PCB) traces increase as the operation frequency increases. Such frequency dependent attenuation causes pulse dispersion and inter-symbol interference (ISI), ultimately degrading the signal-to-noise-ratio (SNR). In addition, reflections from the discontinuities in the signal path due to the connectors and via stubs generate more ISI and further reduce the SNR; as data rates increase, the effects of these discontinuities are magnified. A common approach in the design of high-speed serial links over long, bandwidth-limited channels is to use equalization techniques [1]–[8], including decision feedback equalization (DFE) at the receiver and feed-forward equalization (FFE) at the transmitter. As data rates approach 10 Gb/s, similar techniques can be used in parallel short-haul chip-to-chip interconnects to significantly enhance their performance. For parallel links with thousands of IOs per chip, however, the area and power consumption of on-chip transceivers must be very small. For this reason, only very compact and low-power equalization techniques are applicable to this design space.

In this paper, we introduce a low-power DFE receiver which is suitable for signaling over short to medium-length wire traces for chip-to-chip interconnections. The techniques used to reduce the power consumption are described in the course of the paper. These techniques include use of a switched-capacitor summation DFE, analog multiplexing, and quarter-rate clocking. In Section II, we discuss the basics of DFE design and its requirements. We show that the power consumption and performance of summer circuit required by the DFE has a large effect on the overall power and performance of the receiver. Also, we discuss the speculation technique (also known as loop-unrolling) used in many designs to meet the timing requirements of DFE receivers. Following this introduction, we will propose a switched-capacitor approach to implement the summer for the DFE and will discuss its design issues and performance. Next, we discuss the design of the multiplexer required by the speculative DFE, in this case implemented as an analog multiplexer to reduce overall DFE receiver power consumption and area. Section III focuses on the overall architecture of the receiver and the use of multiphase clocking. In Section IV, we present experimental results from the evaluation of a 90-nm bulk CMOS implementation of the low-power DFE receiver. The performance of the receiver in a variety of channels is discussed in this section, and it is shown

Manuscript received August 25, 2006; revised December 19, 2006. This work was supported in part by MPO Contract H98230-04-C-0920.

A. Emami-Neyestanak is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: azita@ee.columbia.edu).

A. Varzaghani and C.-K. K. Yang are with the University of California, Los Angeles, CA 90095 USA (e-mail: aida@ee.ucla.edu.edu; yang@ee.ucla.edu.edu).

J. F. Bulzacchelli, A. Rylyakov, and D. J. Friedman are with the IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: jfbulz@us.ibm.com; sasha@us. ibm.com; dfriedmn@us.ibm.com).

Digital Object Identifier 10.1109/JSSC.2007.892156

that the DFE significantly improves the performance of the receiver for signaling over short to medium distance PCB traces while consuming a small amount of additional power.

II. Decision Feedback Equalization

In many electrical signaling channels, for data rates above 5.0 Gb/s, the losses in the wires and connectors, as well as reflections from via stubs and connectors, are significant. A common approach to remove ISI and enhance SNR is to use FFE, DFE, or both. These techniques help to compensate for post-cursor and pre-cursor ISI arising from the spread of a single pulse over time. The high level block diagram of a typical DFE is shown in Fig. 1. The ISI from previous bits is compensated by adjusting the DFE taps: $h1, h2, \ldots, hn$. The delay elements taking on values of unit bit-times can be implemented using latches and flip-flops. In the DFE, weighted versions of previous samples are added or subtracted to the main sample by a summer amplifier at the front-end. Here we assume that the tap weights are adjusted using an open-loop or a closed-loop adaptive technique. The dashed line in the figure indicates the critical path for the DFE, the arrival of feedback from the decision made on the previous bit in time to properly influence the decision on the current bit. The constraint imposed by this critical path is that the sum of the slicer latch delay (Clock-to-Q delay), the settling time of the summer, and the setup time of the latch needs to be less than one unit interval (UI) or bit-time. Satisfying such a timing requirement becomes a difficult problem at increasingly high data rates. The problem can be mitigated by the application of speculative techniques, known also as loop-unrolling [9]. A block diagram of speculative DFE and the new critical path associated with this approach is shown in Fig. 2. The figure shows a half rate architecture in which the data is received with both edges of a clock at a frequency half that of the data rate. In this case, the critical path is the arrival of feedback from the decision made two bits earlier in time to properly influence the decision on the current bit; the combination of speculation and half-rate architecture (parallelism) effectively allows the time available for the closure of this critical path to increase to two unit intervals. Note that while this approach relaxes timing constraints, the design of the analog summer and the multiplexer (for the speculative first tap) remains critical and challenging. In Sections II-A and B of this paper, we discuss the design of these components, including approaches taken to reduce the power consumption despite the relatively high target data rates.

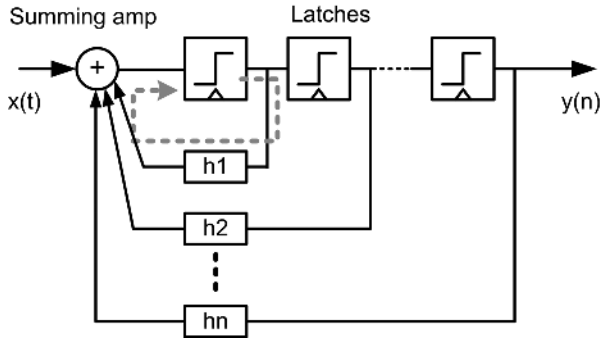

Fig. 1. Block diagram of a DFE receiver; critical path is shown with the dashed line.

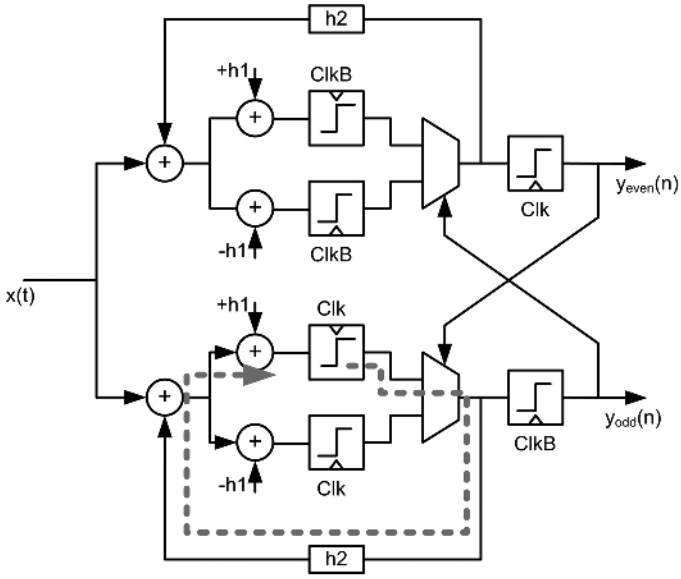

Fig. 2. Half-rate DFE receiver with speculation; new critical path is shown with the dashed line.

A. Summer Design

The first ISI post-cursor can be equalized by subtracting the estimated error from the main sample. This operation is commonly done by using an analog summer or by introducing an offset to the slicer [1]. As we discussed in the previous section, the delay of the summer needs to be much smaller than the overall timing budget of the feedback loop. At very high data-rates, the design of a linear and precise summer that meets this timing requirement is a very challenging task. Most summer circuits shown to date are based on current mode schemes [2], [3]. An example of this approach, a differential current mode summer, is shown in Fig. 3, including the main amplifier stage and the extra current branches for adding the DFE taps. Here these taps are represented by adjustable current sources that are switched to the appropriate leg of the main differential amplifier depending on the previous bit value. The summer is capacitively loaded by the input capacitance of the next stage latches, the parasitics of the summer devices themselves, and the wires used for these connections. In order to achieve a high SNR, the summer output needs to get close to its final value before it can be safely used by the next stage. Since in many systems the first ISI tap can be as high as half of the main tap, the summer output needs to be close to its final value at the time of sampling. Also, if the summer output is too far from its final value, the dependency of ISI cancellation on the clock jitter will increase and further degrades the SNR. Considering the trade off between SNR and power consumption, here we assume that the summer output needs to settle to more than 95% of its final value. The settling requirement could be less stringent, however it will not change the result of this analysis significantly. The settling requirement implies that the *RC* time constant of the output node should be much smaller than 2UI in a half-rate architecture. This requirement often dictates high power consumption in the summing amplifier. As an example, let us assume that our target data rate is 10 Gb/s and in order to relax timing requirements,

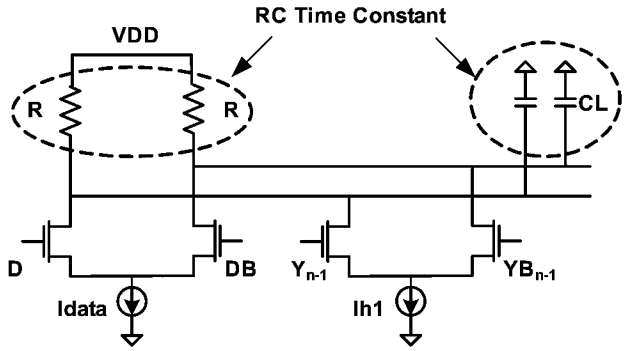

Fig. 3. Current-based summer.

speculation and multiphase clocking (at least two phases) are used. To allow ample timing budget for the digital circuits in the feedback loop (without having to power up the digital circuits excessively), three times the *RC* delay of the summer should be less than 50 ps. If the total capacitance of the output node is 25 fF, the load resistor of the summer must therefore be less than 667 Ω. For reliable operation when receiving signals with partially closed eyes, the differential signal amplitude at the output of the DFE summer needs to be on the order of 300 mV. To handle a 300 mV differential signal with high linearity, the DC voltage drops across the load resistors should be on the order of 0.5 V. Therefore, the bias current of the main differential amplifier needs to be about 1.5 mA. If a large DFE feedback term (e.g., 200 mV) is needed, the tail current which sets the first tap of the DFE can be as high as 0.3 mA. Therefore, the overall current of a one-tap summer stage implemented using the typical approach shown in Fig. 3 can be as high as 1.8 mA. Since a half-rate DFE with speculation (Fig. 2) employs four summers, the current draw of the summers alone may exceed 7 mA.

The alternative summer design approach described in this paper is a charge/voltage-mode summer as shown in Fig. 4(a) (for simplicity, only the half-circuit of the differential structure is shown.) In this design, we extend the standard "sample and hold" front-end to a switched-capacitor summer. $\alpha V_{\rm ref}$ and $-\alpha V_{\rm ref}$ represent the DFE tap values for one and zero bits, respectively. The clocking of the switches is also shown in this figure. The input signal is sampled onto capacitor C_s through switch S1d. During the sampling period, switch S1 connects the right terminal of capacitor C_s to a fixed common-mode voltage $V_{\rm CM}$. Therefore, during the sampling phase the voltage across C_s is equal to $V_i - V_{\rm CM}$ and the output voltage $V_{\rm out}$ is equal to $V_{\rm CM}$. The third switch S1B connects the input (left) terminal of C_s to an adjustable voltage source that represents the DFE tap $(V_{\rm CM} + \alpha V_{\rm ref})$ weight. The switches S1, S1d, and S1B are turned ON with clock phases Ck1, Ck1d, and Ck1B, respectively. During the hold/equalize phase, switch S1B is turned ON. Fig. 4(c) shows the equivalent circuit during this phase, where C_i represents the sum of the input capacitance of the next stage and the parasitic capacitances of switch S1 and all associated wiring. Assuming that the initial voltages on C_i and C_s are $V_{\rm CM}$ and $V_i - V_{\rm CM}$, respectively, the new output voltage will be equal to

$$V_{\rm out} = \frac{C_s\left(2V_{\rm CM} - V_i + \alpha V_{\rm ref}\right) + C_i V_{\rm CM}}{C_s + C_i}. \qquad (1)$$

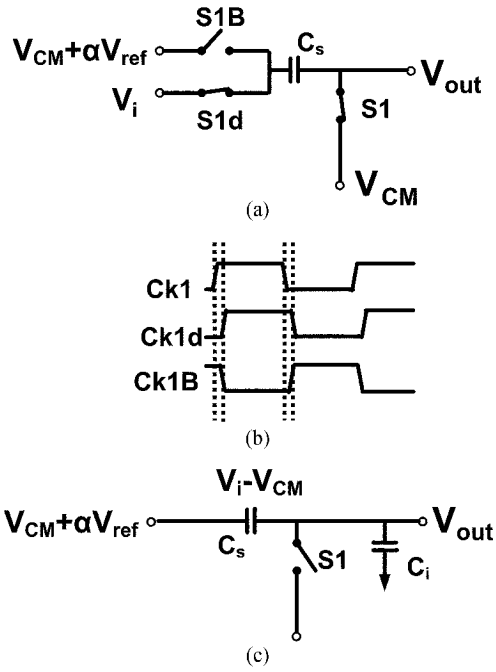

Fig. 4. Switched-capacitor summation. (a) Front-end design. (b) Clocking scheme. (c) Equivalent circuit in the equalization phase.

This equation shows that $\alpha V_{\rm ref}$ is directly subtracted from V_i and that the circuit operates as a summer. If we choose C_i and C_s so that $C_i \ll C_s$, then (1) can be simplified to $V_{\rm out} = 2V_{\rm CM} - V_i + \alpha V_{\rm ref}$. However, due to the charge sharing between C_i and C_s, the sampled voltage V_i is slightly attenuated with a factor of $k = C_s/(C_s + C_i)$. C_i is usually dominated by the input capacitance of the next stage and is set by the speed requirements of the receiver. We need to choose C_s to be large enough to avoid significant signal attenuation. On the other hand, the bandwidth of the sample–hold circuit is inversely proportional to R_sC_s, the time-constant of the sampler, where R_s is the ON resistance of switches S1 and S1d (in series). In order to minimize the power consumption and area of the receiver, capacitor C_s should not be very large.

In the 90-nm bulk CMOS implementation presented in this paper, C_s is a 20-fF lateral capacitor built from four metal levels, where a metal-to-metal capacitor was chosen to maximize linearity. The area of the sampling capacitors was minimized by using four layers of metal with interleaved fingers. Using the provided models, the performance of this capacitor was verified across process and temperature corners. Capacitor C_i is dominated by the input capacitance of the next stage multiplexer (MUX) and the parasitic capacitance of switch S1. The input transistors of the MUX are relatively small with $L =$ 80 nm and $W = 2$ μm. Switch S1 is also a small pMOS device with $L =$ 80 nm and $W = 5$ μm. With these sizes the parasitic capacitor C_i is estimated to be around 1.5 fF. For a power supply voltage of 1.0 V, the common-mode voltage $V_{\rm CM}$ is set to about 800–850 mV. All the switches were implemented as pMOS pass-gates with minimum channel lengths. These transistors were sized to meet the speed requirements of the receiver and are switched with full swing clock signals driven by CMOS inverters. As shown in the clock diagram in Fig. 4(b),

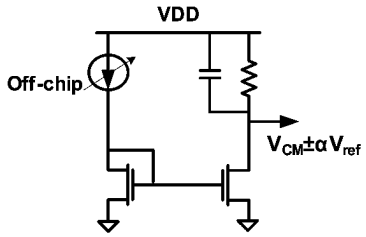

Fig. 5. Adjustable DFE tap weight.

S1 is turned OFF slightly earlier than S1d. By disconnecting and floating one of the plates of capacitor C_s, the overall charge of this capacitor stays almost constant. Therefore, the sampling duration of the input signal ends when switch S1 is turned OFF, and the switching time of S1d, which depends on the input signal, is not critical. Moreover, by fixing the charge of C_s, the signal-dependent charge injection from switch S1d also does not disturb the sampled value. This helps to make the charge injection and sampling time signal-independent [10]. In this design, the delay between CK1 and CK1d is approximately 10 ps and is set by adjusting the sizing of CMOS inverters in the clock buffer chains. The minimum value of this delay is set by the speed of pMOS switches. The maximum delay is limited by the data rate and timing budget for equalization and regeneration. The functionality of this design was verified through simulations across a wide range of temperatures and different process corners. This delay can be designed to be adjustable to compensate for process variation.

Note that α can be adjusted for a specific channel, signal amplitude, and data rate. Fig. 5 shows how α is set in this design. The current source shown in this figure is externally adjusted, and a 1-pF bypass capacitor filters out noise and the kickback from the switches. In systems where the channel properties can change over time, on-chip adaptive DFEs are designed to optimize the tap values continuously [1], [2].

B. Analog Multiplexer

As described above, speculation technique is used in this design to enhance the timing margin of the front-end to enable operation at higher data rates. A common approach for such designs is to fully resolve the outcomes for both possible values (one and zero) of the previous bit to digital levels and then use a digital multiplexer to choose the correct value. In such an approach, two clocked latches are required to store the digital levels. By contrast, the design described in this paper uses an analog multiplexer to select one of the two analog voltages directly at the output of the sampler/summer front-end. The power consumption associated with latches is reduced since a single latch is embedded into the multiplexer [11] as shown in Fig. 6. At low input voltages, the analog MUX must have enough gain and speed to switch the embedded latch within the timing budget of the feedback loop. Thus, careful design of the MUX is critical. Since the MUX is clocked with full-swing clock signals, the clock switch transistors are operating in the linear region, and the stacking of the transistors is easier given the limited

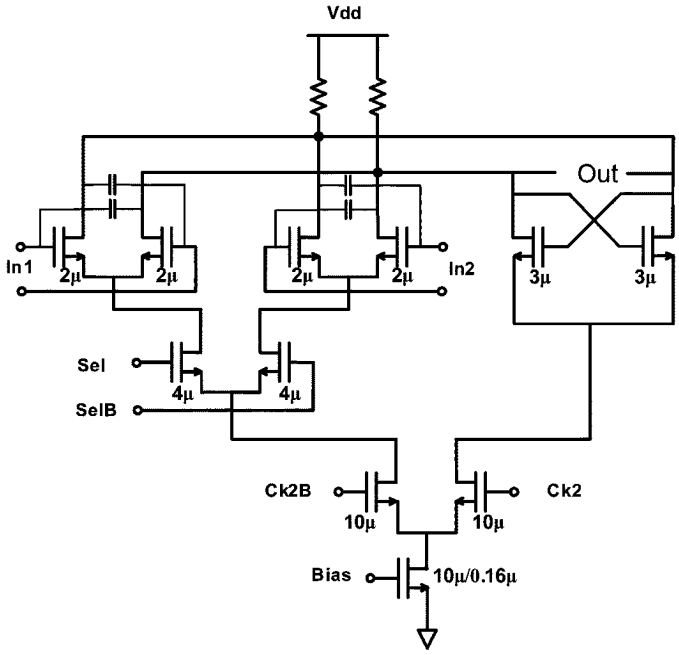

Fig. 6. Analog MUX and embedded latch. Widths of transistors are shown. The channel length of all transistors except the bias is 80 nm.

voltage supply. In order to cancel the kick-back from the latch onto the sampling nodes, small metal capacitors are used to cross-couple the output nodes to the input nodes of the MUX.

The tail transistor sets the bias current of the MUX in a current-mirror configuration simplified as Bias in the figure. This transistor as well as the Sel transistors are not in the delay critical path and are chosen to be relatively large to further relax the voltage headroom constraint of this stage.

III. Receiver Architecture

In many systems, clock generation and distribution at a frequency equal to the data-rate are costly in terms of power consumption and design effort. In such full rate designs, the front-end circuitry needs to run at the data rate, and the design of the equalization feedback loop is very challenging. Direct de-multiplexing at the front-end has been used to relax the design requirements of the clocking and the slicers. This approach also helps to save power in the following digital de-multiplexing stages; in most systems, the data stream needs to be de-multiplexed to lower rates for the next stage data processing and digital blocks regardless of whether initial data recovery is done at full rate or not. In a de-multiplexing front-end, multiple clock phases that run at a fraction of the data rate with phase spacings equal to the bit-time are used to sample the data. Each clock phase drives one of the parallel branches, each of which consists of the slicer and following latches. In this design, we use a 1:4 de-multiplexing scheme, where four equally-spaced phases of quarter-rate clock are used to sample the data, allowing the clock buffers and the four parallel front-end slicers to operate at a frequency of only one quarter that of the data rate. The block diagram of this quarter-rate architecture is shown in Fig. 7. The sampling clock

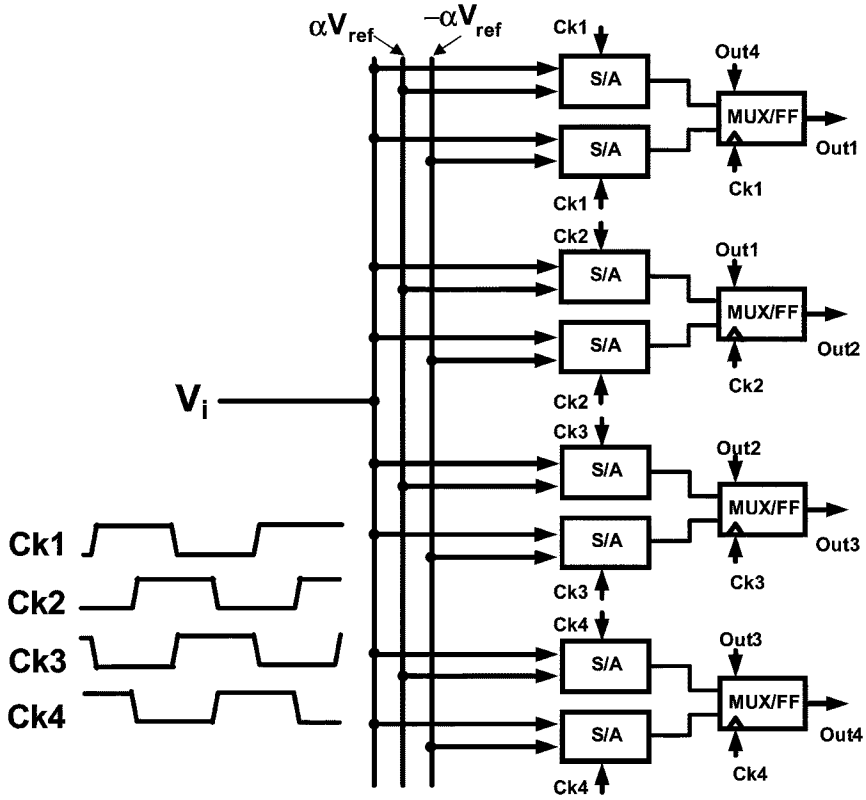

Fig. 7. Quarter-rate receiver architecture.

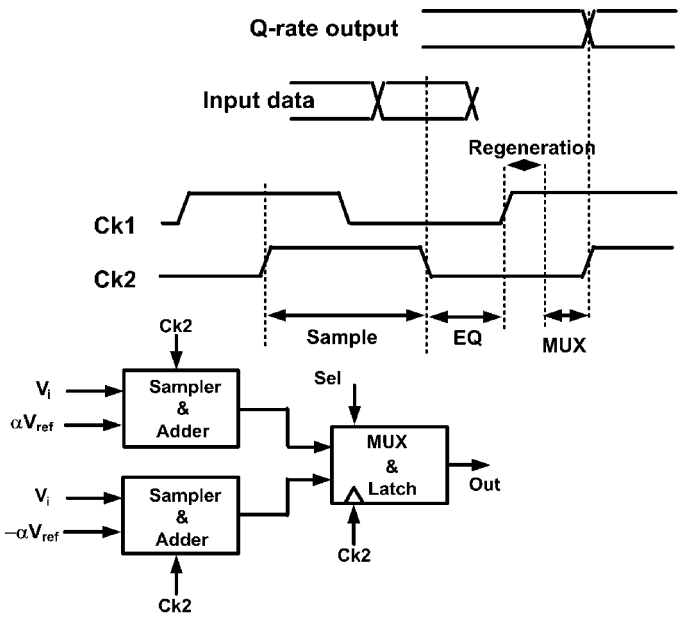

Fig. 8. Timing of the front-end receiver.

phases are Ck1 to Ck4. Note that the input and output signals are all differential in the actual implementation.

Fig. 8 illustrates the timing associated with one of the front-end branches that is triggered by Ck2. The input signal sampling is done with the falling edge of Ck2, when the speculative equalization starts. The next stage MUX is activated when Ck2 is low, and the final latch is triggered with the next rising edge of Ck2. The Sel signal for the MUX is the resolved previous bit, which is the output of the adjacent branch, triggered by Ck1. The delay from the rising edge of Ck1 to the arrival of the Sel signal is shown as "Regeneration" time. The sum of regeneration delay and MUX delay must be less than a bit-time. The equalization also must be completed in one bit-time plus "Regeneration" time. These timing requirements as well as the sampler bandwidth set the maximum receiver data rate.

The technique used for the generation of quarter-rate clock phases is shown in Fig. 9. An external differential full-rate clock signal is buffered by on-chip current-mode logic (CML) clock receivers and buffers. This clock is then sent to a CML differential clock divider, which consists of two divide-by-two stages. The quarter-rate CML clocks are then converted to full-swing CMOS levels using CML-to-CMOS converters shown in Fig. 10(a). The duty cycle and noise performance of these CML-to-CMOS converters is of critical concern in the design of this block. The duty cycles of the CMOS clocks are corrected by using cross-coupled inverters, as shown in Fig. 10(b).

The correction of the clock duty cycle is illustrated in Fig. 11. The names of the signals correspond to the labels in Fig. 10(b). To make the duty cycle errors more visible in the plots, the pMOS transistors of the CML-to-CMOS converter have been deliberately undersized, and the circuit has been simulated under difficult operating conditions: low voltage (VDD $=$ 0.8 V), high temperature ($T = 125\,^{\circ}$C), and slow process corner with PMOS/NMOS skew. After two stages of cross-coupled inverters, the duty cycle is improved from 40.7% to 49.5%.

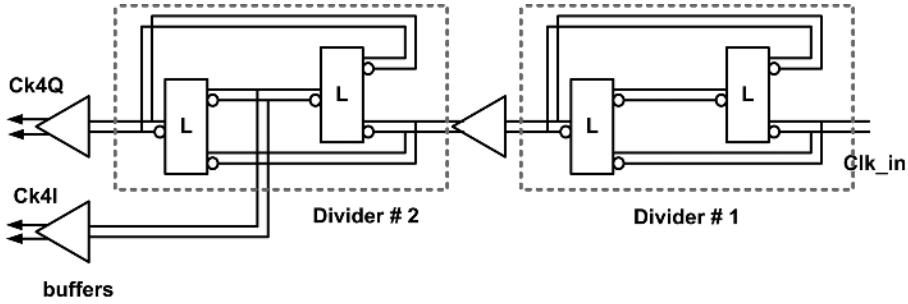

Fig. 9. Quarter-rate clock generation from the input full-rate clock.

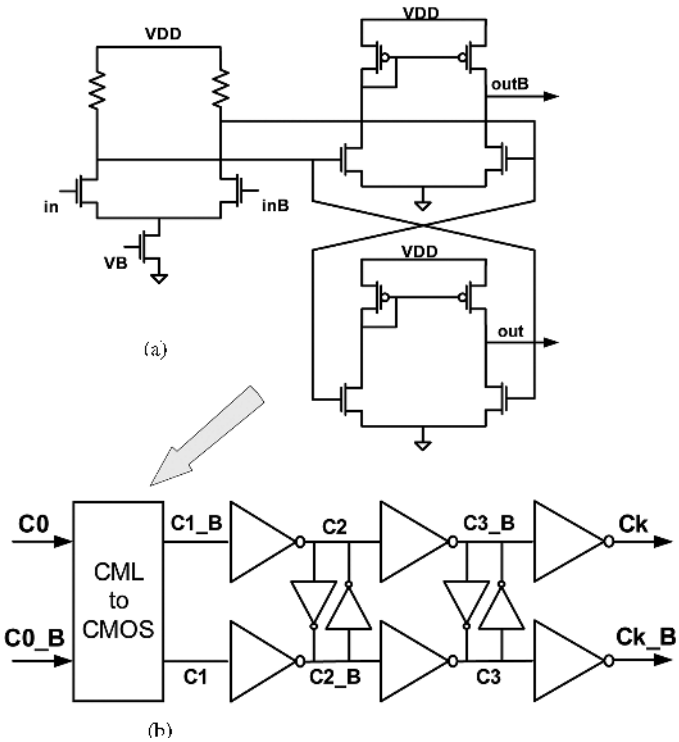

Fig. 10. CMOS clock generation. (a) CML to CMOS conversion. (b) Duty-cycle correction.

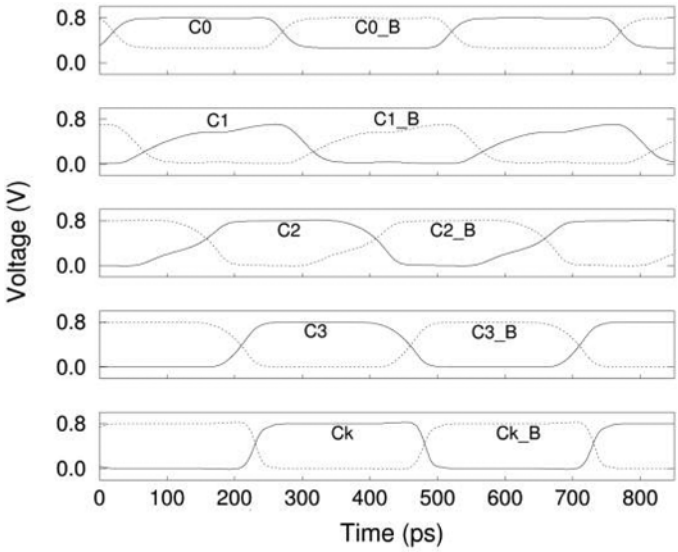

Fig. 11. Duty-cycle correction simulation results.

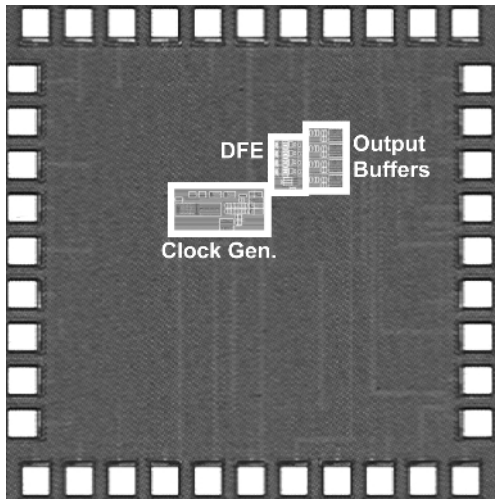

Fig. 12. Die microphotograph with layouts of major blocks superimposed.

IV. Measurement Results

The receiver was fabricated in IBM 90-nm bulk CMOS technology (Fig. 12) and was tested using high-frequency probe cards with a 1.0-V supply used to power the circuit. The test set-up is shown in Fig. 13.

The DFE receiver's performance was examined over channels with different amounts of ISI. For the initial characterization of the design, we used short, high-quality coaxial cables, thus implementing a low-ISI test channel. In this case, the receiver operates error-free at more than 10 Gb/s using PRBS31 patterns, with a minimum input amplitude of ± 40 mV. Note that offset compensation techniques to cancel the summer and slicer offsets are not incorporated into this design. In the second test, a moderate ISI test channel was realized using 5-inch, 4-mil transmission lines on PCB with two 2-mm through vias (no stub) at each ends of lines. The PCB used for this test has a total of 26 metal layers with relatively low-loss APPE dielectric, and the vias are 250 μm in diameter. The measured frequency response of this channel is shown in Fig. 14. The overall channel in this test has more than 6 dB loss at 5.0 GHz. With the DFE off and using a PRBS7 data pattern, the bit-error-rate (BER) was more than 10^{-9} across the eye at 10 Gb/s. Turning on the DFE and optimizing its coefficient enabled the BER to improve to less than 10^{-12} for data rates up to 11 Gb/s, with an eye-opening of more than 11 ps for BER $< 10^{-9}$. With the DFE on and a PRBS31

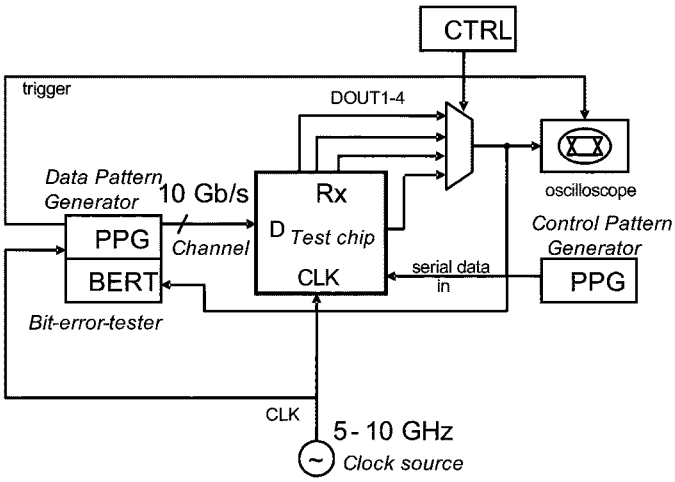

Fig. 13. Receiver test set-up.

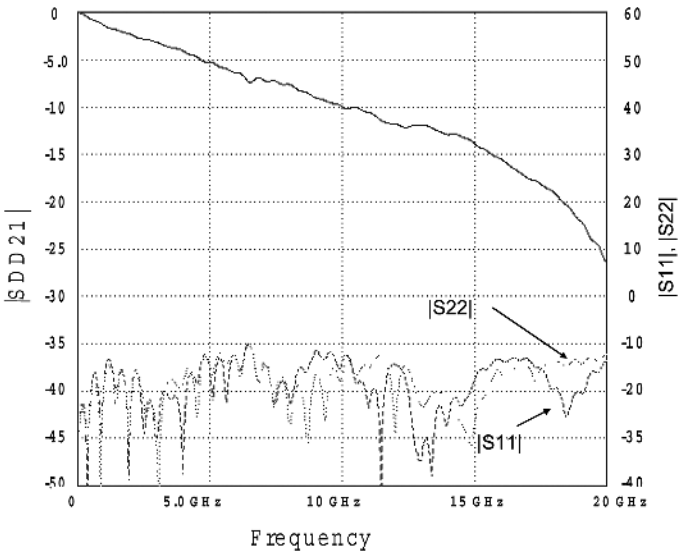

Fig. 14. Frequency response of a 5″ PCB net with vias.

data pattern, the receiver achieved a BER $< 10^{-12}$ at 9.0 Gb/s, again with the DFE coefficient set to its optimum value. The performance difference for the two data patterns does not necessarily indicate a circuit problem. Our system-level link simulation proves that even with ideal circuits, recovering PRBS31 data is more difficult. In the simulations, a reduction of more than 10% in the data rate was necessary to achieve the same BER as PRBS7 for PRBS31 pattern. At higher data rates the single DFE tap is not adequate, and the BER increases due to the residual ISI. In the third test, a high ISI test channel was realized using a 16″ Tyco channel with high levels of reflections and attenuation. The eye diagram of the received data after passing through this channel at 6.0 Gb/s is shown in Fig. 15(a). This channel has more than 11 dB loss at 3 GHz. Again by properly calibrating the DFE coefficient, a BER $< 10^{-12}$ was obtained at 5.0 Gb/s and 6.0 Gb/s with PRBS31 and PRBS7 data patterns, respectively. The output signal after the first stage latch is shown in Fig. 15(b).

Table I summarizes the performance of the low-power DFE receiver described in this paper. At a 1.0-V supply voltage and a data rate of 10 Gb/s, the receiver and clock buffers consume

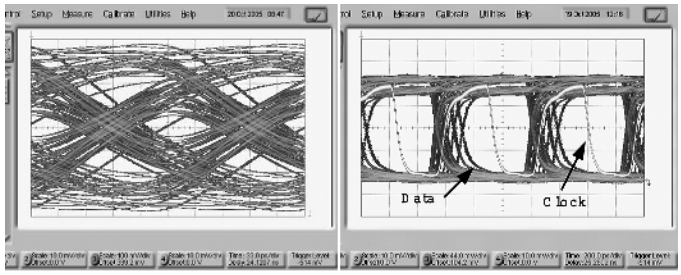

Fig. 15. 6.0-Gb/s data over 16″ Tyco. (a) Input signal. (b) Output of first-stage latch and quarter rate clock (44 mV per division).

TABLE I
RECEIVER PERFORMANCE SUMMARY

Technology	90nm CMOS
Supply voltage	1.0 V
Data rate: cable, 5" PCB, 16" Tyco	10Gb/s, 9.0Gb/s, 6.0Gb/s
Power consumption	6.0mW @ 10Gb/s
Sensitivity (no offset compensation)	± 40 mV differential
Receiver area	70 μm × 150 μm

6.0 mW; at a data rate of 6.0 Gb/s they consume 5.0 mW. At 10 Gb/s, our simulation shows that about 4 mW of power is consumed in the MUX and Latches and about 2 mW in the clocking circuitry and switches. Note that the MUX/Latch stage is effectively a CML-type circuit and thus its power does not scale with frequency in this design, while the power in the CMOS clock buffers for switching the summer and the MUX/latch scales with frequency; this explains the observed partial scaling of power consumption with data rate. The overall area of the receiver and clock buffers is less than 70 μm × 150 μm.

V. CONCLUSION

A one-tap DFE receiver with speculation is designed and fabricated in 90-nm bulk CMOS technology. This receiver is suitable for channels with moderate levels of ISI, mostly due to attenuation not reflections. The simple, low-power DFE can significantly enhance the data rate over short/medium-length channels. In this design, high power efficiency (0.6 mW/Gb/s) is achieved by using switched-capacitor summers, analog multiplexers, and quarter-rate clocking. Unlike current-based summers, the switched capacitor summer does not require a high bias current. Power is consumed in clocking of the switches, which are implemented using small transistors. By multiplexing analog instead of digital signals, the number of digital latches is minimized. In most systems, the buffering of high frequency clocks requires high-power CML drivers. In this design, quarter-rate clocking allows the use of CMOS clock buffers and relieves the speed requirements of the front-end circuits. Additional power is saved in the ensuing digital blocks since the resolved data is already at quarter-rate. The DFE control in this design is implemented by using adjustable external current sources.

REFERENCES

[1] V. Stojanovic, A. Ho, B. Garlepp, F. Chen, J. Wei, E. Alon, C. Werner, J. Zerbe, and M. A. Horowitz, "Adaptive equalization and data recovery in a dual-mode (PAM2/4) serial link transceiver," in *Symp. VLSI Circuits 2004 Dig. Tech. Papers*, 2004, pp. 348–351.

[2] 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.

[3] R. Payne, P. Landman, B. Bhakta, S. Ramaswamy, S. Wu, J. D. Powers, M. U. Erdogan, A.-L. Yee, R. Gu, L. Wu, Y. Xie, B. Parthasarathy, K. Brouse, W. Mohammed, K. Heragu, V. Gupta, L. Dyson, and W. Lee, "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.

[4] K. Krishna, D. A. Yokoyama-Martin, S. Wolfer, C. Jones, M. Loikkanen, J. Parker, R. Segelken, J. L. Sonntag, J. Stonick, S. Titus, and D. Weinlader, "A 0.6 to 9.6 Gb/s binary backplane transceiever core in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2005, pp. 64–65.

[5] A. Varzaghani and C.-K. K. Yang, "A 6-GSamples/s multi-level decision-feedback-equalizer embedded in a 4-bit time-interleaved pipeline A/D converter," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 935–944, Apr. 2006.

[6] M. Callicotte, J. Little, H. Takatori, K. Dyer, and C.-H. Lee, "A 12.5 Gb/s single-chip transceiver for UTP cables in 0.13 μm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 86–87.

[7] R. Farjad-Rad, C.-K. K. Yang, and M. Horowitz, "A 0.3-μm CMOS 8-Gb/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, no. 5, pp. 757–764, May 2000.

[8] A. Fiedler *et al.*, "A 1.0625 Gb/s transceiver with 2X oversampling and transmit pre-emphasis," in *IEEE ISSCC Dig. Tech. Papers*, 1997, pp. 238–239.

[9] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 6, pp. 711–717, Jun. 1991.

[10] G. M. Haller and B. A. Wooley, "A 700-MHz switched-capacitor analog waveform sampling circuit," *IEEE J. Solid-State Circuits*, vol. 29, no. 4, pp. 500–508, Apr. 1994.

[11] S.-J. Bae, H.-J. Chi, Y.-S. Sohn, and H.-J. Park, "A 2 Gb/s 2-tap DFE receiver for multi-drop single-ended signaling systems with reduced noise," in *IEEE ISSCC Dig. Tech. Papers*, 2004, pp. 244–245.

Azita Emami-Neyestanak (S'97–M'04) was born in Naein, Iran. She received the B.S degree with honors in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1996, and the M.S and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1999 and 2004, respectively.

She has been with Columbia University, New York, NY, as an Assistant Professor in the Department of Electrical Engineering since July 2006. She was a Research Staff Member at IBM T. J. Watson Research Center, Yorktown Heights, NY, from 2004 to 2006. Her current research areas are VLSI systems, and high-performance mixed-signal integrated circuits, with the focus on high-speed and low-power optical and electrical interconnects, synchronization, and clocking.

Aida Varzaghani was born in Iran in 1978. She received the B.S. and M.S. degrees (both with honors) in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1999 and 2001, respectively. Since 2001, she has been working toward the Ph.D. degree in integrated circuits and systems at the University of California, Los Angeles.

From 1999 to 2001, she was with Emad Semiconductor Company as an Analog Circuit Designer. She designed a low-power binary receiver for IBM/Yorktown Heights in summer 2004. Her current research interests include high-speed, large-bandwidth A/D converters, equalization techniques for very high-speed serial I/O links, switched capacitor and switched opamp A/D converters.

Ms. Varzaghani was the recipient of a UCLA fellowship and dissertation year fellowship for Fall 2001 and academic year of 2006–2007, respectively.

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.

Dr. Bulzacchelli received the Jack Kilby Award for Outstanding Student Paper at the 2002 IEEE International Solid-State Circuits Conference. He holds two U.S. patents.

Alexander Rylyakov received the M.S. degree in physics from the Moscow Institute of Physics and Technology, Moscow, Russia, in 1989, and the Ph.D. degree in physics from the State University of New York at Stony Brook in 1997.

From 1994 to 1999, he worked in the Department of Physics at SUNY Stony Brook on the design and testing of high-speed (up to 770 GHz) digital integrated circuits based on a superconductor Josephson junction technology. In 1999, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, as a Research Staff Member, working on the design and testing of full-custom digital and mixed signal integrated circuits for serial communications (up to 80 Gb/s data rates and up to 100 GHz clock rates), using a broad spectrum of CMOS and SiGe technologies.

Chih-Kong Ken Yang (S'94–M'98) was born in Taipei, Taiwan, R.O.C. He received the B.S. and M.S. degrees in 1992 and the Ph.D. degree in 1998 from Stanford University, Stanford, CA, all in electrical engineering.

He has been with the University of California at Los Angeles, as an Assistant Professor and Associate Professor in 1999 and 2004, respectively. His current research area is high-performance mixed-mode circuit design for VLSI systems such as clock generation, high-performance signaling, low-power digital design, and analog-to-digital conversion.

Dr. Yang is the recipient of the IBM Faculty Development Fellowship in 2003–2005 and the 2003 Northrup-Grumman Outstanding Teaching Award.

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 Thomas 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 turned his focus to analog circuit design for high-speed serial data communication. Since June 2000, he has managed a team of circuit designers working on serial data communication, wireless, and PLL applications. In addition to circuits papers regarding serial links, he has published articles on imagers and RFID, and he holds more than 20 patents. His current research interests include high-speed I/O design, PLL design, and circuit/system approaches for variability compensation.