

Phase and Amplitude Pre-Emphasis Techniques for Low-Power Serial Links

James F. Buckwalter, *Member, IEEE*, Mounir Meghelli, Daniel J. Friedman, *Member, IEEE*, and Ali Hajimiri, *Member, IEEE*

Abstract—A novel approach to equalization of high-speed serial links combines both amplitude pre-emphasis to correct for intersymbol interference and phase pre-emphasis to compensate for deterministic jitter, in particular, data-dependent jitter. Phase pre-emphasis augments the performance of low power transmitters in bandwidth-limited channels. The transmitter circuit is implemented in a 90-nm bulk CMOS process and reduces power consumption by pushing CMOS static logic to the output stage, a 4:1 output multiplexer. The received signal jitter over a cable is reduced from 16.15 ps to 10.29 ps with only phase pre-emphasis at the transmitter. The jitter is reduced by 3.6 ps over an FR-4 backplane interconnect. A transmitter without phase pre-emphasis consumes 18 mW of power at 6 Gb/s and 600 mVpp output swing, a power budget of 3 mW/Gb/s, while a transmitter with phase pre-emphasis consumes 24 mW, a budget of 4 mW/Gb/s.

Index Terms—Equalization, high-speed serial links, intersymbol interference (ISI), data-dependent jitter, deterministic jitter, pre-emphasis, transmitter.

I. INTRODUCTION

EQUALIZATION of high-speed serial links has evolved to compensate intersymbol interference (ISI) caused by frequency-dependent attenuation and reflections found in interconnects. Pre-emphasis-based equalization in the transmitter and decision feedback equalization in the receiver figure prominently in overcoming signal degradation and improving the bit-error rate (BER) [1]–[3]. Currently, one challenge of equalization is minimizing power consumption while still improving signal integrity in the presence of attenuation and reflections. This design targets low power consumption while offering equalization appropriate for shorter (under 16 inches), less dispersive interconnects.

In this transmitter, we expand the notion of pre-emphasis beyond amplitude compensation of ISI and introduce phase pre-emphasis for compensating data-dependent jitter (DDJ). DDJ compensation has been demonstrated to increase the timing margins in clock and data recovery [4], [5]. Unlike random jitter (RJ), DDJ can be addressed by exploiting the relationship between the data sequence and the timing deviation [6], [7]. Earlier work on magnetic write heads proposed

compensating transition timing deviations caused by long sequences of ones and zeros through the addition of isolated pulses during data runs [8], [9].

Transmitters must drive enough power over lossy interconnects to meet minimum receiver sensitivity requirements. Finding new methods to lower power consumption in serial links has been explored through low-common-mode signaling [10] and through low-supply operation [11]. Amplitude pre-emphasis distorts the signal to compensate ISI introduced by the bandwidth limitations of the interconnect [12], [13]. The number of taps of amplitude pre-emphasis depends on the channel quality and bit rate. Further improvement of the signal integrity motivates the addition of phase pre-emphasis. Since DDJ compensation is implemented through adjusting the transition times of the data, it can be introduced without significant additional power consumption in the transmitter. The effectiveness of combining amplitude and phase pre-emphasis depends on the behavior of the channel response.

In Section II, we summarize the operation of one-tap amplitude pre-emphasis in the context of modern bandwidth limitations in transmission lines. Section III introduces the notion of phase pre-emphasis and discusses a general algorithm for reducing DDJ at the receiver. In Section IV, we discuss the implementation of a 6-Gb/s transmitter with amplitude and phase pre-emphasis. Finally, the hardware results are discussed in Section V and we illustrate the performance over both a bandwidth limited cable and a backplane interconnect.

II. AMPLITUDE PRE-EMPHASIS

Transmission lines are often limited by frequency-dependent skin-effect and dielectric losses. Skin effect arises from nonuniform electric fields in conductors. At high frequencies, the resistance of the wire increases and is accompanied by an effective internal inductance. The skin loss can be expressed as

$$G_{\text{skin}}(f) = e^{-(1+j)l\sqrt{\pi\mu\sigma f}} \quad (1)$$

where l is the wire length, μ is the permeability, and σ is the conductivity. Interestingly, the phase shift and amplitude attenuation are identical for skin effect. Additionally, the dielectric material of the transmission line causes loss.

$$G_{\text{dielectric}}(f) = e^{-l(\sqrt{\varepsilon_r} \tan \delta) f/c} \quad (2)$$

where ε_r is the dielectric constant, $\tan \delta$ is the loss tangent of the material, and c is the speed of light. A thorough discussion of the skin and dielectric losses in modern materials is given

Manuscript received October 19, 2005; revised February 16, 2006. This work was supported by the Maryland Procurement Office under Contract H98230-04-C-0920 and the National Science Foundation.

J. F. Buckwalter was with the California Institute of Technology, Pasadena, CA 91125 USA (e-mail: jamesb@caltech.edu).

A. Hajimiri is with the California Institute of Technology, Pasadena, CA 91125 USA.

M. Meghelli and D. J. Friedman are with IBM Research, Yorktown Heights, NY 10598 USA.

Digital Object Identifier 10.1109/JSSC.2006.874270

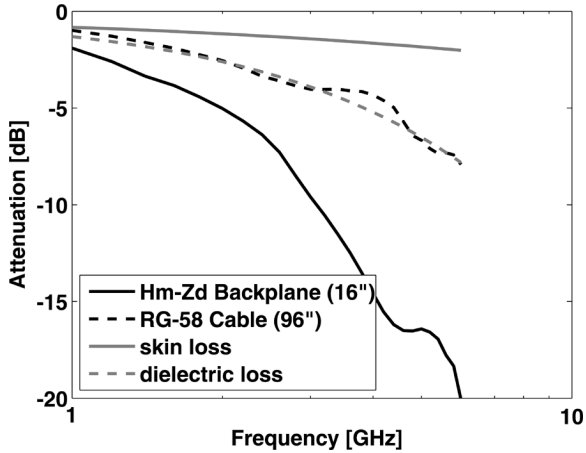


Fig. 1. Bandwidth of 16 inches FR-4 backplane channel with Hm-Zd connectors and 96 inches of RG-58 cable compared to analytical skin and dielectric loss.

by Deutsch [14]. Fig. 1 compares skin and dielectric loss to the measurement of 96 inches of RG-58 cable. At 3 GHz, one-half the bit rate of 6 Gb/s, the loss of the cable is roughly 4 dB. Recent equalizer circuits have been introduced to compensate for cable loss [15]. In high-speed backplanes, via stubs and connectors cause signal reflections at gigahertz frequencies. These reflections also cause dispersion and, consequently, ISI in the link. For modern backplanes, these mechanisms limit the bandwidth of interconnects to around 3 GHz, depending on length. Interconnect bandwidth decreases with longer length and the number of discontinuities. In Fig. 1, the frequency response of a Tyco backplane with Hm-Zd connectors rolls off much faster due to these reflections. The loss over this channel is close to 10 dB at 3 GHz.

Amplitude pre-emphasis compensates the frequency-dependent loss. The block diagram for a one-tap feed-forward amplitude pre-emphasis scheme is shown in Fig. 2. The circuit consists of two cross-coupled drivers that compete to drive the line. The amplitude pre-emphasis driver transmits data delayed by one bit period, T . The amplitude pre-emphasis gain, G , is the ratio of current in the pre-emphasis driver and main driver, $I_{\text{preemp}}/I_{\text{bias}}$. The transfer function for this stage is expressed as

$$H_{\text{preemp}} = \frac{TX(s)}{X(s)} = I_{\text{bias}}(1 - Ge^{-sT}). \quad (3)$$

A first-order Pade approximation is used to express the time delay as a rational function, $e^{-sT} = (1 - s(T/2))/(1 + s(T/2))$. The transfer function for the scheme is written as

$$\begin{aligned} H_{\text{preemp}} &= \frac{TX(s)}{X(s)} = I_{\text{bias}} \left(1 - G \frac{1 - sT/2}{1 + sT/2} \right) \\ &= I_{\text{bias}}(1 - G) \frac{1 + s(T/2)(1 + G)/(1 - G)}{1 + s(T/2)}. \end{aligned} \quad (4)$$

This transfer function reveals that the amplitude pre-emphasis scheme reduces the DC gain. For this reason, amplitude pre-emphasis is often considered “de-emphasis.” The DC gain can be




Fig. 2. Block diagram for one-tap feed-forward amplitude pre-emphasis driver.

increased to maintain a constant voltage swing by increasing the bias current and power consumption. High-frequency amplification is introduced through a zero that depends on G . The transfer function has an additional pole at $1/T$. Other poles in an actual implementation limit the all-pass behavior. In Fig. 2, the operation of amplitude pre-emphasis creates four distinct transmit levels in the data eye. The higher levels, $\pm(1 + G)$, are transmitted when there is a data transition. Since the transmitter operates from a fixed voltage supply, a maximum peak power limits the swing of the higher levels. One-tap amplitude pre-emphasis tends to provide some compensation for the frequency-dependent losses, however, more robust approaches use multiple taps to implement the transmit finite-impulse response (FIR) filter at the expense of additional power consumption [1], [2], [13].

III. PHASE PRE-EMPHASIS

To improve the timing margins of the data eye, phase pre-emphasis is introduced to the transmitted signal to compensate the effects of data-dependent jitter. Data-dependent jitter, the result of ISI at the transition times, is a form of deterministic jitter (DJ) that limits the timing margins of the received data eye. In [6] and [7], the relationship between the link characteristic and the DDJ is established. In particular, it is shown analytically that there are multiple DDJ peaks related to when the previous transition occurred. If we assume that the channel is linear, the timing deviation of complementary signals, i.e., a 101 and 010 data sequences, is identical. In [5], we demonstrated the adjustment of the data transition timing at the receiver to compensate for the effect of DDJ. This work focuses on DDJ compensation in the transmitter where there are several implementation advantages.

Transitions that have occurred recently tend to strongly impact the DDJ. This is described in [6], where the DDJ resulting from a first-order response comprises diminishing contributions for each previous transition. In Fig. 3, the transitions of a first-order data eye are magnified. The timing deviations due to previous transitions are highlighted. For the k th previous transition, the transition between the $-k$ and $-k - 1$ bit, we denote the mean timing deviation as $t_{c,DDJ}^{(k)}$. For instance, $t_{c,DDJ}^{(1)}$ describes the DDJ contribution of the transition between previous and the penultimate bit. The figure illustrates that the DDJ contribution tapers off quickly. After the third transition, the contribution is small compared to the bit period. Phase pre-emphasis manipulates the data transition timing to neutralize DDJ.

The relationship between different data sequences and the threshold crossing time is illustrated in Fig. 4. A few data sequences are plotted to demonstrate the associated timing devia-

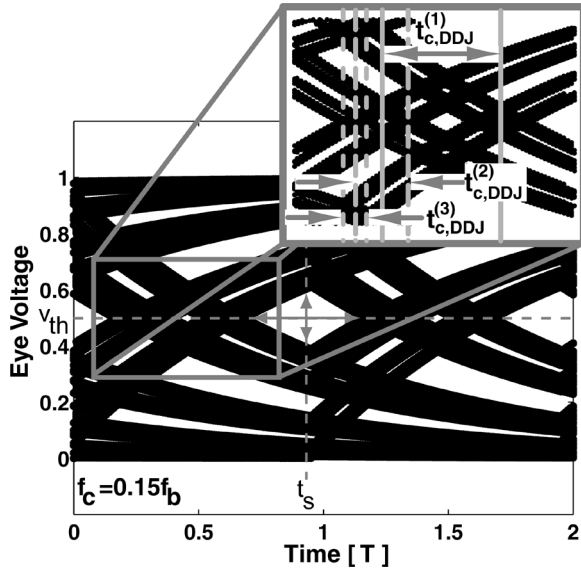


Fig. 3. Data eye showing the contribution to data-dependent jitter of previous transitions.




Fig. 4. Phase pre-emphasis operation and falling-edge truth table for compensating DDJ.

tion and to illustrate the operation of the pre-emphasis scheme. Detection of each previous transition is used to compensate the current transition time. The detection of k th previous transition is denoted $X[k]$ for the current bit, $D[n]$, and is calculated from $X[k] = D[n-1] \oplus D[n-1-k]$, where \oplus is the XOR operator. From derivations in [5], the compensated transmit time is calculated through the following algorithm:

$$t_{comp}[n] = X[1]t_{c,DDJ}^{(1)} + X[2]t_{c,DDJ}^{(2)} + X[3]t_{c,DDJ}^{(3)} \quad (5)$$

For example, the 0010 sequence results in the fastest threshold crossing time for a first-order system. For this sequence, $X[1]$ and $X[2]$ are both 1, implying that we will introduce the largest delay to compensate the fast transition for this sequence. On the other hand, the 1110 sequence results in the slowest threshold

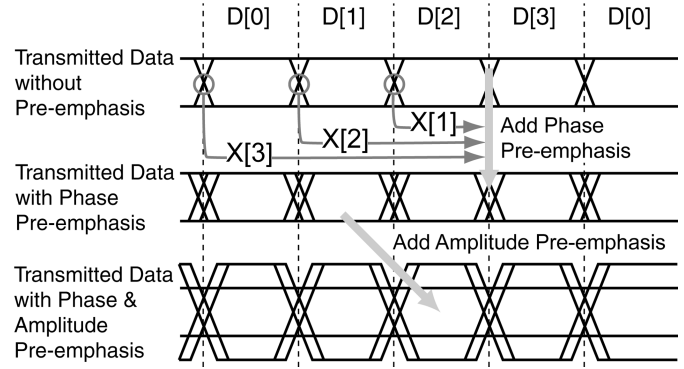


Fig. 5. Transmitted eye with phase and amplitude pre-emphasis.

crossing time. For this sequence, $X[1]$ and $X[2]$ will both be 0 and the transmitter will not introduce delay in the transmitted bit timing. In this example, we illustrate two 1110 curves corresponding to whether the initial condition, the unshown previous bit, is 0 or 1. Ideally, compensation is introduced for each previous transition until the difference in DDJ contributed by these initial conditions is negligible.

In Fig. 5, we demonstrate heuristically how phase pre-emphasis, along with amplitude pre-emphasis, introduces DDJ and ISI to the signal that is removed by the loss mechanisms in the serial link. The combination of both approaches can provide more signal integrity. In essence, the feed-forward amplitude pre-emphasis is a symbol-spaced FIR filter. The addition of phase pre-emphasis introduces an approximation for a half symbol-spaced FIR filter. While the use of amplitude pre-emphasis alone can provide some improvement in the DDJ, the symbol-spaced FIR filter cannot generally adjust the DDJ and the ISI to be simultaneously zero. Consequently, the use of half symbol-spaced FIR filters is essential to minimize both the DDJ and ISI.

Finally, for the purposes of this work, the coefficients of the equalizer are assumed to be adjusted *ad hoc*. The DDJ in simple linear time invariant (LTI) systems can be solved exactly to calculate the necessary timing compensation, but generally some pulse response characterization or equalizer adaptation is required to adjust both phase and amplitude pre-emphasis coefficients. The coefficients for the phase pre-emphasis could be compensated by sampling the timing deviation at the receiver for particular transmitted data sequences. For instance, in Fig. 4 the 0110 sequence introduces the delay $t_{c,DDJ}^{(2)}$ while the 1010 sequence introduces the delay $t_{c,DDJ}^{(1)}$. At the receiver, these particular sequences could be detected and a timing interval could be calculated from the mean transition timing. The coefficients would be transmitted back to the transmitter in a back-channel scheme [3].

IV. CIRCUIT IMPLEMENTATION

The architecture for the low-power phase and amplitude pre-emphasis transmitter is illustrated in Fig. 6. The design is based on a 4:1 output multiplexer that provides amplitude pre-emphasis, combinatorial logic for phase pre-emphasis, delay generation cells for controlling the clock edges of the multiplexer, and duty-cycle control for each clock phase, which is useful for

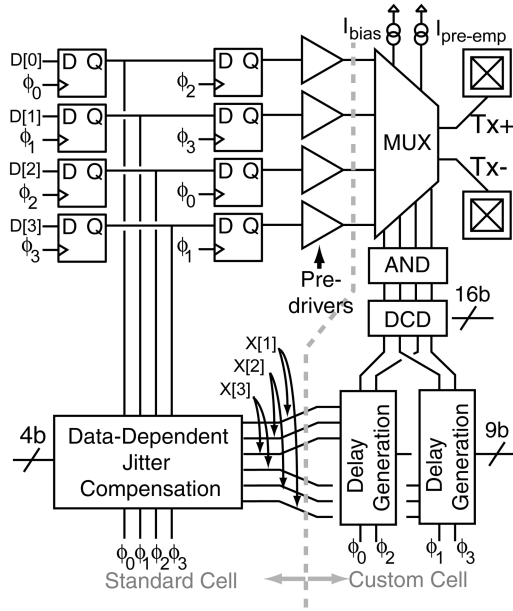


Fig. 6. Transmitter schematic showing DDJ computation from parallel data and delay generation in clock path. Clock phases are independently adjusted for DCD and ANDed at output multiplexer.

counteracting process variations. Each of the quarter-rate clock phases is ANDed with its neighbor to generate a 25% duty cycle clock at the 4:1 multiplexer.

The transmitter is implemented in IBM CMOS 9SF, a 90-nm bulk triple-well CMOS technology. Static logic in this technology operates over 2 Gb/s, and, consequently, can directly drive a 4:1 output multiplexer at serial rates faster than 8 Gb/s. In this implementation, we targeted 6 Gb/s operation. The output multiplexer is a current-mode logic (CML) stage which is linear, operates at a low supply voltage, and has a relatively fixed large-signal output impedance to avoid source reflections present in some low-power transmit schemes [10].

The 4:1 output multiplexer provides advantages when considering amplitude pre-emphasis [13]. The output multiplexer schematic is illustrated in Fig. 7. Each bit is transmitted sequentially and is available for three additional bit periods during which the next bit must be set-up. Using two cross-coupled multiplexers, we can implement one-tap amplitude pre-emphasis. The first multiplexer transmits the original bit and the second multiplexer is cross-coupled to invert the bit during the following period. Sequential clock phases trigger the two multiplexers to provide one bit delay. Therefore, the 4:1 multiplexer adds amplitude pre-emphasis to the output data without requiring additional circuitry and power to latch and hold the data. The only cost to this scheme is the additional pre-driver current required to drive twice the capacitance at the input of the output stage. However, this pair of cross-coupled multiplexers also reduces the ISI and DDJ inherent in the output multiplexer architecture. The gate-drain capacitance of the output stage provides a parasitic path for energy coupling between the input and output as shown in Fig. 7. With the cross-coupled multiplexer, the gate-drain capacitance is neutralized. The accuracy of this neutralization is subject to the limits of process variations and mismatch for these large output

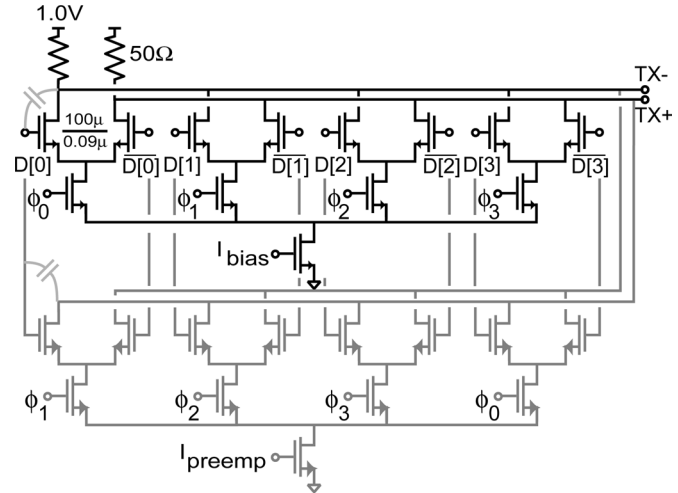


Fig. 7. Output multiplexer schematic. The primary 4:1 multiplexer, in black, transmits the current data bit, while the pre-emphasis multiplexer, in gray, is cross-coupled to transmit an inverted replica of the previous data bit.

multiplexer transistors ($W = 100 \mu\text{m}$) [16]. The output swing, ΔV , is designed for 600-mVpp differential, a tradeoff between the link sensitivity requirement and the headroom restrictions of a 1-V supply.

The phase pre-emphasis combinatorial logic is shown in Fig. 8. The stage is implemented with standard cell static logic. The transitions in the data are calculated with a cascade of XOR stages that implement the algorithm in (5). Each stage of XORs calculates a previous transition in the data. The transition calculations are stored with a D-flip-flop (DFF) at each stage. The phase pre-emphasis calculation for data-dependent jitter compensation results in three differential transition detection control bits for each of the two delay generation cells. Additionally, the combinatorial logic includes a 4-bit interface to provide for a sign adjustment in the DDJ compensation.

A 3-GHz quadrature differential clock generates the four phases used in a quarter rate architecture. The first and third clock phases are initially fully differential and control the timing of the first and third bits. The same holds for the second and fourth clock phases. Therefore, the transition detection output of the phase pre-emphasis combinatorial logic is multiplexed to a 3-GHz rate and controls the rising and falling edges of the clock respectively for the first and third bits. This process is illustrated in Fig. 9. For instance, the rising edge of one differential clock controls the output timing of the first bit while the falling edge of this clock controls the output timing of the third bit. Consequently, the transition multiplexer introduces different phase pre-emphasis for the first and third bits on the rising and falling edges individually. The multiplexers are modulated with the quadrature clock to ensure the appropriate setup time.

The transition multiplexer output signals switch two independent delay generation cells that handle each of the quadrature differential clocks. These delay generation cells, shown in Fig. 10, are designed with fully differential CML logic to benefit from the power supply rejection on the clocks. Each delay generation cell consists of a cascade of three 3-bit programmable delay cells. The delay generation is provided by




Fig. 8. Data-dependent jitter compensation schematic. The bank of XORs calculates when transitions occur in the data sequence. The MUX switches the transition detection based on which bit is being transmitted.

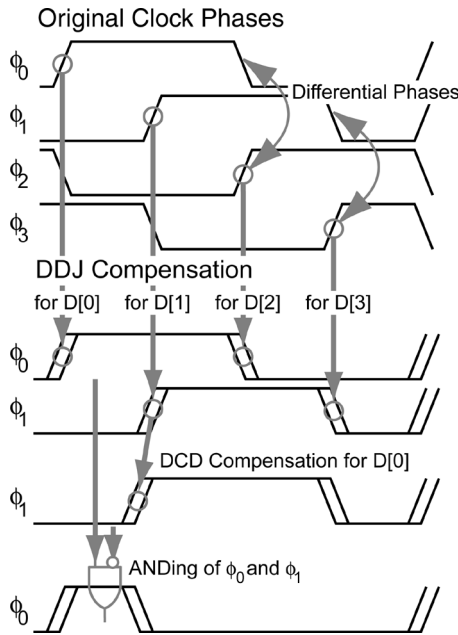


Fig. 9. Manipulation of the clock phases through DDJ and DCD compensation. ANDing before the output multiplexer provides seamless clock transition in spite of the introduction of deterministic jitter.

switching between two versions of the clock: one programmable delay and one nominal delay. Depending on the DDJ sign bit, the programmable delay is introduced when a transition is detected in the transmitted data. Each consecutive delay cell is used to handle the timing deviation corresponding to a previous transition. The programmable delay cell is designed to provide 3 ps of delay for each digital bit. Consequently, the maximum t_{comp} that can be implemented is 24 ps.

The output buffer of the delay generation cell is designed as a bandpass buffer stage to reject low-frequency noise. Since the phase pre-emphasis scheme is implemented by modulating the clock phases as opposed to modulating the data edges as in [4], bandpass buffering can be used to pass only frequency

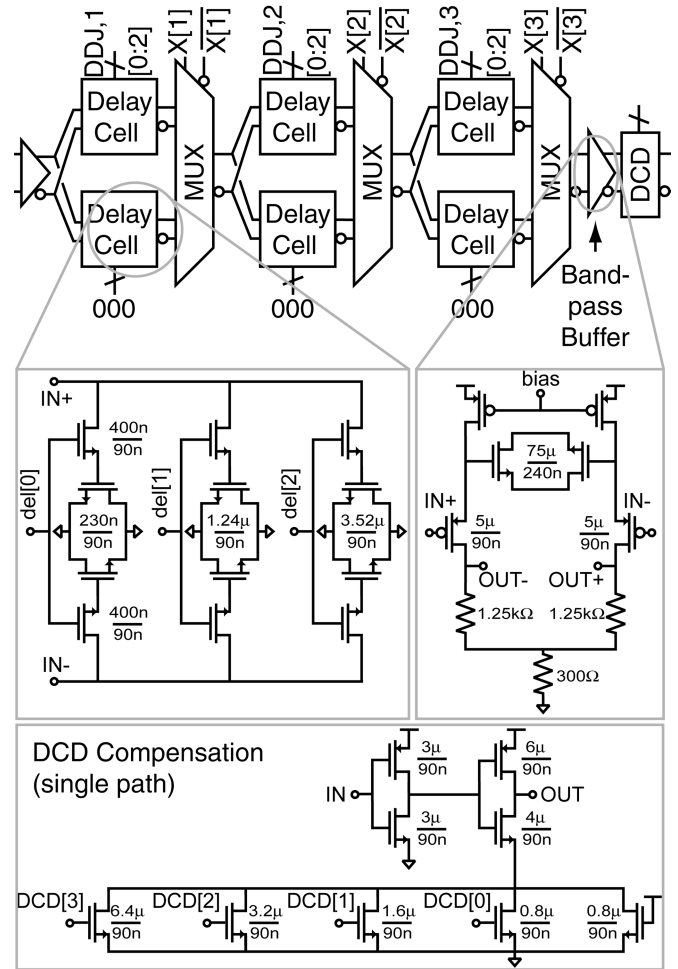


Fig. 10. Delay generation schematic. Detailed schematics are presented for the delay cell, clock amplifier, and duty-cycle distortion (DCD) compensation circuits.

content around the 3-GHz clock. The bandwidth of the band-pass must meet the phase modulation requirements for the clock. A source-degenerated pMOS driver is implemented to provide low-frequency noise rejection.

At this point, the fully differential clocks are split into four different clock phases and duty-cycle distortion (DCD) is compensated. This is an additional source of DJ and must be eliminated in 4:1 multiplexers. DCD control is implemented through four individual pathways, each with four control bits as illustrated in Fig. 6. Independent DCD control allows adjustment for process and transistor matching variations that influence the duty cycle of each data path [13]. Since the DCD is a static error while DDJ compensation is a dynamic adjustment, it will not interfere with the DDJ adjustment on the clock. The compensation of DCD for one clock phase is illustrated in Fig. 9. The DCD circuit is based on a current starved inverter that is digitally programmable as shown in Fig. 10.

Finally, neighboring clock phases are ANDed together. This ideally converts each clock phase to a duty cycle of one-quarter the bit period. Notably, ANDing neighboring clock phases is useful for the phase pre-emphasis scheme. As shown in Fig. 9, the DDJ timing compensation is introduced to the clock phase that ends the transmission of the current bit. This phase also

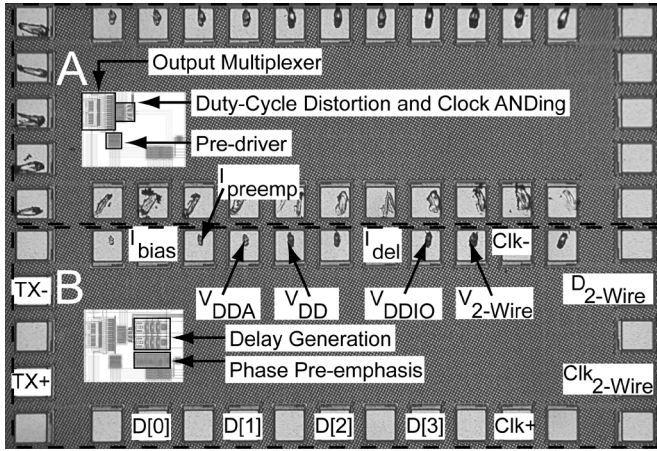


Fig. 11. Chip microphotograph of two breakout sites. Top site does not offer phase pre-emphasis.

triggers the beginning of the consecutive bit. Hence, the clock hand-off is seamless, avoiding any clock overlap issues that might otherwise arise from introducing timing variations on four different clock phases.

V. RESULTS

Two break-out sites are shown in Fig. 11. The top break-out is the entire output transmitter without the phase pre-emphasis capability. The bottom site is the entire transmitter with the phase pre-emphasis capability. The total area of the transmitters is 240 μm by 150 μm , roughly the area required for two pads. The transmitters also include a two-wire interface for programming the DCD and DDJ delay cells.

The individual operation of the amplitude and phase pre-emphasis is demonstrated in Fig. 12. The first set of eyes show amplitude pre-emphasis with increasing pre-emphasis current at 6 Gb/s. The pre-emphasis gain is set through a reference current that is recorded for each data eye. As shown in the figure, the minimum swing decreases while the maximum swing increases with amplitude pre-emphasis current. The second row of data eyes show the use of phase pre-emphasis. The DDJ introduced to the signal depends on the digital code used to program the delay generation scheme. The first transition, $t_{c,DDJ}^{(1)}$ is compensated and the two resulting threshold crossing times are observed from the jitter histogram. The separation between the peaks increases with the code value. The variation in the total jitter is measured as a function of the pre-emphasis code in Fig. 13 at 6 Gb/s. The DDJ component is normalized out of this total jitter and the time delay compensation can be calculated from the DDJ peak separation [5]. The linearity of the time delay compensation is assessed from the slope of this curve as a function of the digital code. Finally, the bottom row of data eyes demonstrates the operation to 10 Gb/s of the transmitter without phase pre-emphasis. The DDJ compensation is limited to 6 Gb/s due to the bandpass amplifier in the delay generation circuit. The use of amplitude pre-emphasis opens the data eye slightly at 10 Gb/s to counteract the output bandwidth of the transmitter stage. Four consecutive eyes are shown to demonstrate the relative DCD matching between each of the four paths in the multiplexer.

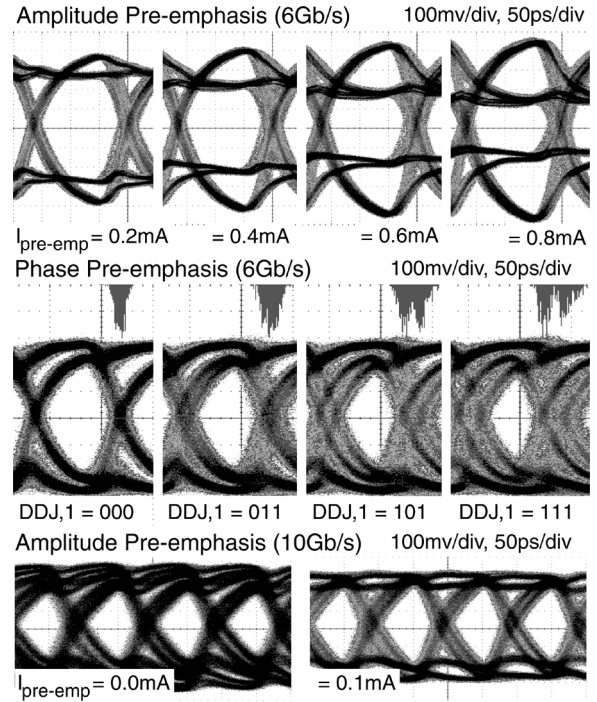


Fig. 12. Data eyes at 6 Gb/s and 10 Gb/s demonstrating amplitude and phase pre-emphasis. The first row shows the amplitude pre-emphasis at 6 Gb/s as a function of pre-emphasis current. The second row illustrates the phase pre-emphasis as a function of the compensation code. Finally, the last row shows the operation of the transmitter at 10 Gb/s.

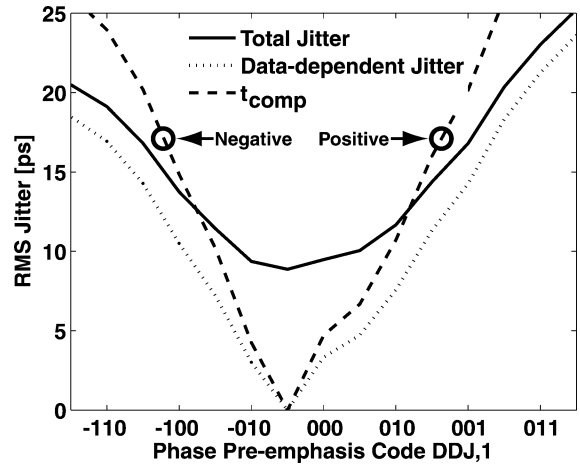


Fig. 13. Total and data-dependent jitter versus phase pre-emphasis codes for the first previous transition. The variation in DDJ can be used to calculate the time delay variation.

To test the transmitter featuring phase and amplitude pre-emphasis, a $2^{31} - 1$ pseudo-random bit sequence at 6 Gb/s was passed through two test channels. The first channel, 96 inches of RG-58 cable, has 4 dB of loss at 3 GHz. Three data eyes are illustrated in Fig. 14. The first is the eye without any compensation. The second eye is compensated using the first-transition phase pre-emphasis. The rms jitter reduces from 16.15 ps to 11.06 ps when the DDJ code is 011. The jitter statistics are collected over 5 k points. The third eye includes first and second transition phase pre-emphasis and the rms jitter drops to 10.29 ps with the DDJ code for the second transition set to 001. The RJ

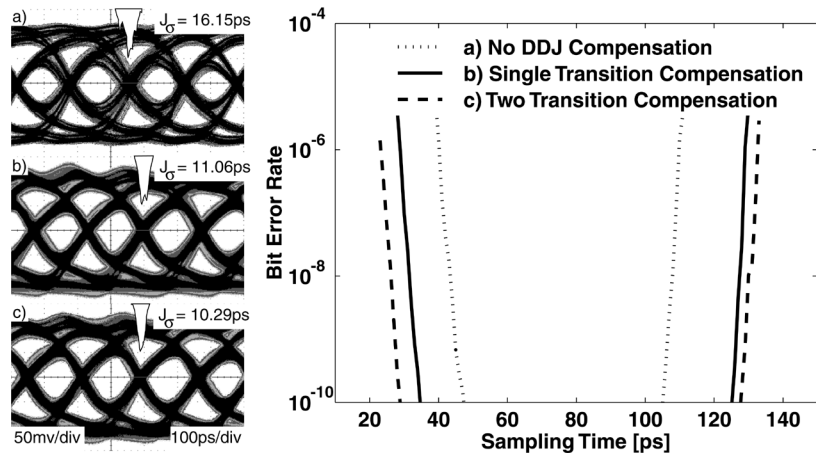


Fig. 14. Performance on 96 inches of RG-58 cable. The uncompensated eye is shown in (a) while phase pre-emphasis is introduced in (b) and (c).

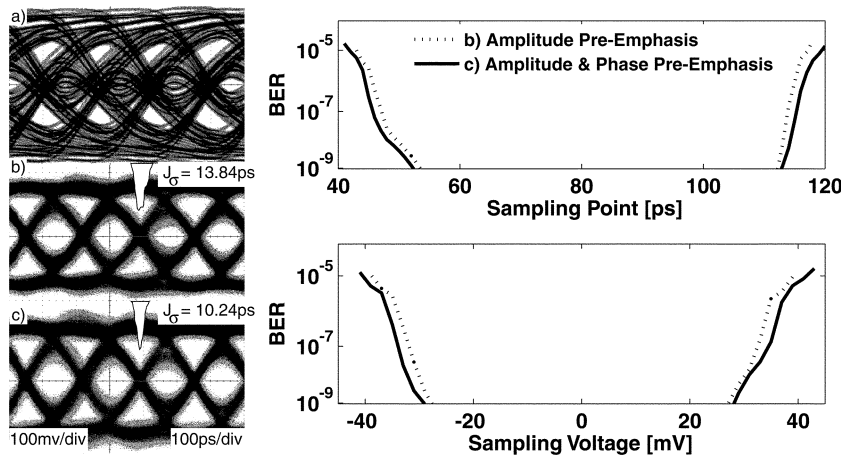


Fig. 15. Performance on 16 inches of FR-4 backplane with connectors. The uncompensated eye is closed in (a). In (b), amplitude pre-emphasis opens the eye and phase pre-emphasis improves the timing margins in (c).

is measured from a periodic pattern and has a jitter of 8.06 ps. The BER bathtub curve demonstrates that at 10^{-9} BER the timing margin increases from 62 ps to 95 ps with first transition phase pre-emphasis and compensating the second transition opens the bathtub by an additional 6 ps.

The second channel, a 16 inch FR-4 backplane interconnect with Tyco Hm-Zd connectors, has 10 dB of loss at 3 GHz. The frequency-dependent attenuation in this case is reflected in the closed data eye at 6 Gb/s shown in Fig. 15. Amplitude pre-emphasis is used exclusively in the first data eye. With the data eye open, phase pre-emphasis is used to open the eye further. The rms jitter reduces from 13.84 ps to 10.24 ps using first transition phase pre-emphasis. The change in the rms jitter is demonstrated as a function of the pre-emphasis (DDJ,1) code in Table I. Clearly, DDJ,1 = 010 minimizes the rms jitter in the data eye. The improvement at 10^{-9} BER is shown across sampling times and voltage thresholds. Notably, the phase pre-emphasis opens the data eye slightly in the time domain and voltage domain.

The transmitter nominally operates at 1.0 V but can operate from 0.8 to 1.2 V, offering a tradeoff between power consumption and performance. The transmitter without phase pre-emphasis consumes a minimum of 18 mW of power at 6 Gb/s with a 600 mVpp swing, giving a power budget of 3 mW/Gb/s. The

TABLE I
MEASURED RMS JITTER ACROSS BACKPLANE VERSUS
PHASE PRE-EMPHASIS CODES IN PICOSECONDS

| Code | -100 | -011 | -010 | -001 | 000 | 001 | 010 | 011 | 100 |
|--------|------|------|------|------|------|------|------|------|------|
| Jitter | 21.5 | 19.9 | 17.3 | 14.8 | 13.8 | 10.9 | 10.2 | 11.3 | 13.9 |

output multiplexer draws 12 mA while the data buffers and clock driver, including DCD control and phase ANDing, consume the remaining 6 mA from a 1-V supply. The power consumption of the two-wire interface and clock generation are not included in this power budget. Amplitude pre-emphasis increases the power budget through the current drawn through the amplitude pre-emphasis multiplexer. The transmitter with phase pre-emphasis consumes a minimum of 24 mW of power at 6 Gb/s, giving a power budget of 4 mW/Gb/s. This additional power is consumed primarily in the custom designed CML delay generation stages and the phase pre-emphasis combinatorial logic. The power consumption is scanned with the bias current in Fig. 16 to show the achievable differential signal swing and power consumption. The voltage swing is demonstrated on the right axis and the power consumption is on the left axis. As expected the

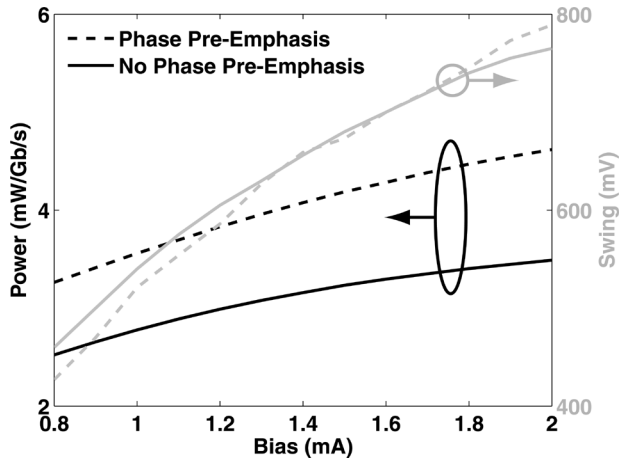


Fig. 16. Power consumption per bit rate as a function of bias current.

voltage swing tracks the power consumption for both implementations of the transmitter. The transmitter with phase pre-emphasis demands roughly 1 mW/Gb/s more over the entire output swing range. However, the power consumption of the phase pre-emphasis scheme is not fundamentally limited to 1 mW/Gb/s and migration of the delay generation from CML logic to static logic should realize additional power advantages.

VI. CONCLUSION

This work describes a novel equalization technique for amplitude and phase pre-emphasis in bandwidth limited interconnects. Phase pre-emphasis is introduced to compensate for datadependent jitter introduced over the channel. Combining amplitude and phase pre-emphasis gives the flexibility to tailor the signal integrity of the data eye. The implementation of this amplitude and phase pre-emphasis transmitter is demonstrated in 90-nm CMOS. The architecture builds upon a 4:1 multiplexer that allows for efficient implementation of amplitude pre-emphasis. The transmitter consumes between 18–24 mW of power at 6 Gb/s, giving a power budget of 3–4 mW/Gb/s/Channel. The transmitter operation is demonstrated over 96 inches of cable as well as a 16 inch backplane interconnect with connectors. The use of both pre-emphasis schemes to open the data eye is demonstrated.

ACKNOWLEDGMENT

The authors appreciate discussions with Drs. A. Rylyakov, S. Rylov, J. Bulzachelli, S. Gowda, M. Soyuer, and M. Oprysko regarding the direction and implementation of this project.

REFERENCES

- [1] J. Stonick *et al.*, "An adaptive PAM-4 5-Gb/s backplane transceiver in 0.25 μm CMOS," *IEEE J. Solid-State Circuits*, vol. 38, no. 3, pp. 436–443, Mar. 2003.
- [2] J. Zerbe *et al.*, "Equalization and clock recovery for a 2.5–10-Gb/s 2-PAM/4-PAM backplane transceiver cell," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2121–2130, Dec. 2003.
- [3] V. Stojanovic *et al.*, "Autonomous dual mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.

- [4] J. Buckwalter and A. Hajimiri, "A 10 Gb/s data-dependent jitter equalizer," *Proc. IEEE Custom Integrated Circuits Conf.*, pp. 39–42, Oct. 2004.
- [5] J. F. Buckwalter and A. Hajimiri, "Analysis and equalization of data-dependent jitter," *IEEE J. Solid-State Circuits*, vol. 41, no. 3, pp. 607–620, Mar. 2006.
- [6] J. F. Buckwalter, B. Analui, and A. Hajimiri, "Predicting data-dependent jitter," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 51, no. 9, pp. 453–457, Sep. 2004.
- [7] B. Analui, J. F. Buckwalter, and A. Hajimiri, "Data-dependent jitter in serial communications," *IEEE Trans. Microw. Theory Tech.*, vol. 53, no. 11, pp. 3388–3397, Nov. 2005.
- [8] M. I. Behr and N. S. Blessum, "Technique for reducing effects of pulse crowding in magnetic recording," *IEEE Trans. Magn.*, vol. 8, no. 3, pp. 608–611, Sep. 1972.
- [9] R. C. Schneider, "Write equalization in high-linear-density magnetic recording," *IBM J. Res. Dev.*, vol. 29, pp. 563–568, Nov. 1985.
- [10] K.-L. J. Wong *et al.*, "A 27-mW, 3.6-Gb/s I/O transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 602–612, Apr. 2004.
- [11] J. Kim and M. Horowitz, "Adaptive supply serial links with sub-1-V operation and per-pin clock recovery," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1403–1413, Nov. 2002.
- [12] W. J. Dally and J. Poulton, "Transmitter equalization for 4 Gb/s signaling," in *Proc. Hot Interconnects Symp.*, Aug. 1996, pp. 29–39.
- [13] R. Farjad-Rad, C.-K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4- μm , 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 580–585, May 1999.
- [14] A. Deutsch, "Electrical characteristics of interconnections for high-performance systems," *Proc. IEEE*, vol. 86, no. 2, pp. 315–357, Feb. 1998.
- [15] G. E. Zhang and M. E. Green, "A 10 Gb/s BiCMOS adaptive equalizer," *IEEE J. Solid-State Circuits*, vol. 40, no. 11, pp. 2132–2140, Nov. 2005.
- [16] P. Kinget, "Device mismatch and tradeoffs in the design of analog circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 6, pp. 1212–1224, Jun. 2005.



James F. Buckwalter (S'01–M'06) received the B.S. degree in electrical engineering from the California Institute of Technology (Caltech), Pasadena, in 1999. He received the M.S. degree in electrical engineering in 2001 from the University of California at Santa Barbara, where he developed coupled phase-locked loops for microwave beam-steering under the supervision of Prof. Robert York. He returned to Caltech to complete the Ph.D. degree with Prof. Ali Hajimiri. His thesis is focused on deterministic timing jitter in high-speed communi-

cation and testing.

He was a Research Scientist at Telcordia Technologies from 1999 to 2000 where he worked on rate-agile, burst-mode electronics under a Next-Generation Internet DARPA project. He worked at the IBM T. J. Watson Research Center in Yorktown Heights, NY, during the summer of 2004.

Dr. Buckwalter was awarded an IBM Ph.D. fellowship in 2004.



Mounir Meghelli was born in Oran, Algeria, in 1969. He received the M.S. degree in electronics and automatics from the University of Paris XI, Paris, France, in 1992 and the Engineering degree in telecommunication from the ENST-Paris in 1994. From 1994 to 1998, he was with the France Telecom Research Center, CNET-Bagneux, as a Ph.D. student working on the design of high-speed ICs.

Since 1998, he has been with the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he has been involved with the design of SiGe BiCMOS and CMOS high-speed ICs.



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

After completing postdoctoral work in image sensor design, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, initially working on analog circuit design for field-powered radio frequency ID tags and later for high-speed serial data communication. He now manages a team of circuit designers working on serial data communication and wireless applications.



Ali Hajimiri (S'95–M'99) received the B.S. degree in electronics engineering from the Sharif University of Technology, Tehran, Iran, and the M.S. and Ph.D. degrees in electrical engineering from the Stanford University, Stanford, CA, in 1996 and 1998, respectively.

He was a Design Engineer with Philips Semiconductors, where he worked on a BiCMOS chipset for GSM and cellular units from 1993 to 1994. In 1995, he was with Sun Microsystems, where he worked on the UltraSPARC microprocessor's cache RAM de-

sign methodology. During the summer of 1997, he was with Lucent Technologies (Bell Labs), Murray Hill, NJ, where he investigated low-phase-noise integrated oscillators. In 1998, he joined the Faculty of the California Institute of Technology, Pasadena, where he is an Associate Professor of electrical engineering and the Director of the Microelectronics Laboratory. His research interests are high-speed and RF integrated circuits. He is the author of *The Design of Low Noise Oscillators* (Kluwer, 1999) and holds several U.S. and European patents. He is a cofounder of Axiom Microdevices Inc.

Dr. Hajimiri is an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and a member of the Technical Program Committee of the IEEE International Solid-State Circuits Conference (ISSCC). He has also served as an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—PART II, a member of the Technical Program Committees of the International Conference on Computer Aided Design, a Guest Editor of the IEEE TRANSACTIONS ON MICROWAVE THEORY AND TECHNIQUES, and a member of the Guest Editorial Board of the Transactions of the Institute of Electronics, Information and Communication Engineers of Japan (IEICE). He was selected to the top 100 innovators (TR100) list in 2004 and is a Fellow of Okawa Foundation. He is a recipient of the Teaching and Mentoring award at Caltech. He was the Gold medal winner of the National Physics Competition and the Bronze Medal winner of the 21st International Physics Olympiad, Groningen, The Netherlands. He was a co-recipient of the ISSCC 1998 Jack Kilby Outstanding Paper Award, two times co-recipient of the IEEE Custom Integrated Circuits Conference's best paper awards, and a three times winner of the IBM faculty partnership award, as well as a National Science Foundation CAREER award.