Autonomous Dual-Mode (PAM2/4) Serial Link Transceiver With Adaptive Equalization and Data Recovery

Vladimir Stojanović, *Member, IEEE*, Andrew Ho, *Member, IEEE*, Bruno W. Garlepp, *Member, IEEE*, Fred Chen, *Member, IEEE*, Jason Wei, *Member, IEEE*, Grace Tsang, *Member, IEEE*, Elad Alon, *Student Member, IEEE*, Ravi T. Kollipara, *Member, IEEE*, Carl W. Werner, *Member, IEEE*, Jared L. Zerbe, *Member, IEEE*, and Mark A. Horowitz, *Fellow, IEEE*

***Abstract*—This paper describes an adaptively equalized, dual-mode (PAM2 one-tap DFE/PAM4) 0.13 μm CMOS transceiver chip, and the techniques used to continuously adapt the link. Interestingly, with only minor modification the same hardware needed to implement a PAM4 system can be used to implement a PAM2 loop-unrolled single-tap decision-feedback equalization (DFE) receiver. Adaptive equalization using data-based update filtering allows continuous updates while minimizing the required sampler front-end hardware and significantly reduces the cost of implementation in multi-level signaling schemes. To allow the transmitter to adapt to the channel, the link uses common-mode signaling to create a back-channel communication path over the existing pair of wires. The design uses a three-level return-to-null signaling scheme which allows the receiver to simultaneously extract voltage and timing references and minimize the required receiver hardware. The measured results indicate that this back-channel achieves reliable communication without noticeable impact on the forward link for back-channel data rates of up to 16 Mb/s and swings of 20–100 mV.**

***Index Terms*—Adaptive equalization, back-channel, common-mode, data recovery, decision-feedback equalization (DFE), link.**

I. Introduction

HIGH-SPEED link rates are increasing to the point where they are running into the bandwidth limitation of the cables or backplanes. This bandwidth limitation is caused by dielectric loss, skin-effect and impedance discontinuities. Additionally, in many applications the wires within a system can have significantly different channel characteristics, as shown in Fig. 1. In these systems, achieving optimal performance for each link requires a flexible equalization/modulation solution that can adapt to the specific requirements of its channel [1].

This paper describes the design of a transceiver that is optimized to work with these diverse, bandwidth limited channels.

Manuscript received August 31, 2004; revised November 18, 2004.

V. Stojanović was with Rambus, Inc., Los Altos, CA 94022 USA and also with the Department of Electrical Engineering, Stanford University, CA 94305 USA. He is now with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: vlada@mit.edu).

A. Ho, B. W. Garlepp, F. Chen, J. Wei, G. Tsang, R. T. Kollipara, C. W. Werner, and J. L. Zerbe are with Rambus, Inc., Los Altos, CA 94022 USA.

E. Alon is with the Department of Electrical Engineering, Stanford University, CA 94305 USA, and also with Rambus, Inc., Los Altos, CA 94022 USA.

M. A. Horowitz is with the Department of Electrical Engineering, Stanford University, CA 94305 USA.

Digital Object Identifier 10.1109/JSSC.2004.842863

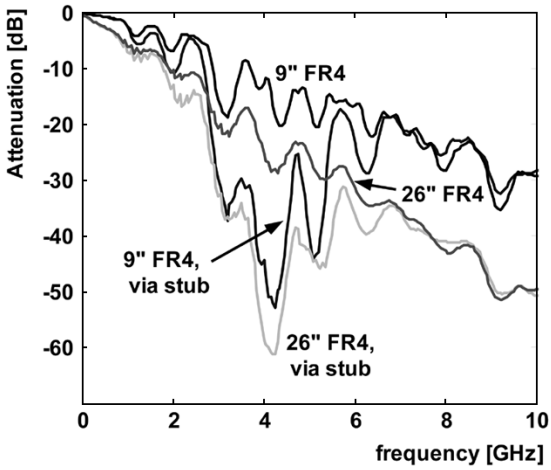

Fig. 1. Frequency response of different channels within the same backplane: FR4 material, 9″ and 26″ trace length, top and bottom routing layers.

The transceiver supports both one bit per symbol (PAM2) and two bits per symbol modulation (PAM4), as well as linear equalization, and a limited form of decision-feedback equalization (DFE). By re-using the circuits in the PAM4 receiver, we were able to seamlessly incorporate a PAM2 receiver with one tap of DFE using loop unrolling [2], without additional cost. To control the equalization values, the link is continually adapted to track changes in the channel and the transceiver chips. Thus the transceiver is able to track changes in channel characteristics due to variations in temperature and humidity, as well as adapt the configuration to the properties of its specific channel. Change in channel loss from 10 to 20 dB at 3 GHz due to operating conditions has been reported by Sheets and D'Ambrosia [3]. The adaptive equalization extends the previous work on multi-level equalizing transceivers [1], [4] by reducing the overhead for equalization to a single additional input sampler. That sampler has an adjustable threshold, and is used in a time-multiplexed fashion to obtain the information required for adaptation.

Like many high-speed links, a transmitter pre-emphasis filter [4], [5] is used in this design, which means that the link needs a back-channel to allow the transmitter to adapt to the received values. While there are many approaches to address this need, common-mode signaling [6] is used over the same wires that carry the differential forward channel, to create a completely self-contained link.

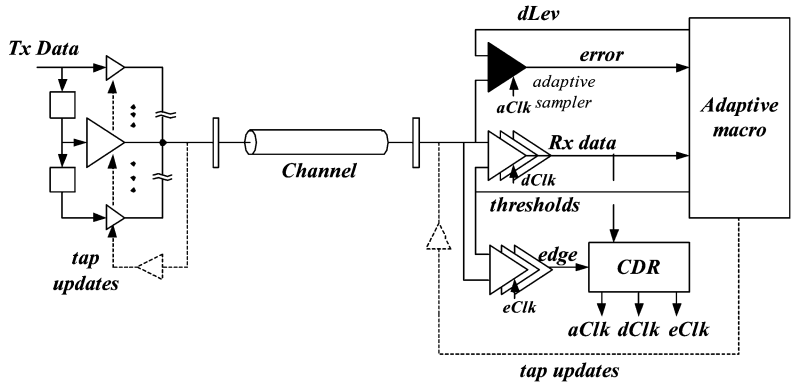

Fig. 2. Adaptive multi-level equalizing link architecture.

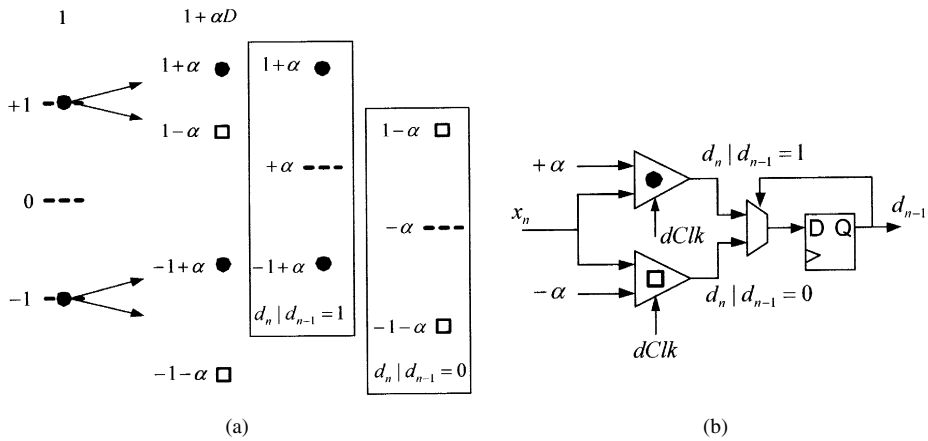

Fig. 3. One tap DFE using loop unrolling. (a) Transmitted PAM2 signal levels corrupted by ISI split to $\pm 1 \pm \alpha$ levels at the receiver, and can be recovered with two slicers offset by the amount of ISI $\pm\alpha$. (b) Practical implementation of the one-tap DFE using loop unrolling.

The next section describes the design of the forward differential part of the high-speed link. Since the hardware is an extension of a previous design [1], this section focuses on the additions made to enable both DFE and adaptive equalization. Having described the main link, Section III then reviews the design of the back-channel. It begins by modeling cross-talk between the two channels caused by channel and circuit imperfections, and then constrains the common-mode signaling so that it does not affect the quality of the forward path. The last part of that section describes a CMOS implementation of the tri-level signaling scheme that was selected. Section IV provides measured experimental results from this transceiver, which show adaptation in PAM2 mode, with and without loop-unrolled DFE.

II. Adaptive High-Speed Link Design

In order to explore techniques for automatic link configuration with minimal hardware overhead, an adaptive link was built, as extension of the design in [1]. The link, shown in Fig. 2, has both transmit pre-emphasis and feedback equalization and can operate in both PAM2 and PAM4 modes to efficiently combat ISI over various backplane channels. The major addition from the previous design is that the new design can automatically tune its parameters for optimal performance (e.g., equalization coefficients, receiver offsets and thresholds, choice of PAM2 or PAM4).

High-speed receivers often use four banks of samplers for each input. Two banks are used as data samplers, allowing the data rate to be twice the clock rate (commonly called double-data-rate or DDR), and an additional two banks are used to sample the input at the transitions (edge samplers) to facilitate timing recovery. For a PAM4 system, each bank requires three samplers to detect the four possible levels. In addition to these standard data slicers and edge samplers, the receiver has one extra sampler (the adaptive sampler in Fig. 2) used for monitoring the link performance [8]. This adaptive sampler has variable timing and voltage references and, as we will show later, is used for a number of tasks, such as monitoring performance during link operation, or providing the information necessary for adaptive equalization and link configuration.

One use of the adaptive sampler is to reduce the effective input offset of the samplers. Since the magnitude of the received signal is significantly attenuated due to channel rolloff and limited swing at the transmitter, the input offset of the receiver is a critical parameter. A multiplexing method is used where each of the edge or data samplers can be temporarily taken out of service without disrupting the flow of data, and calibrated offline. During the calibration period, the adaptive sampler takes the role of the sampler under calibration. Each sampler has a 6-bit dedicated offset canceling digital-to-analog converter (DAC), and a shared 8-bit DAC for threshold selection, while the adaptive sampler has a 9-bit DAC for adaptive threshold, dLev, setting.

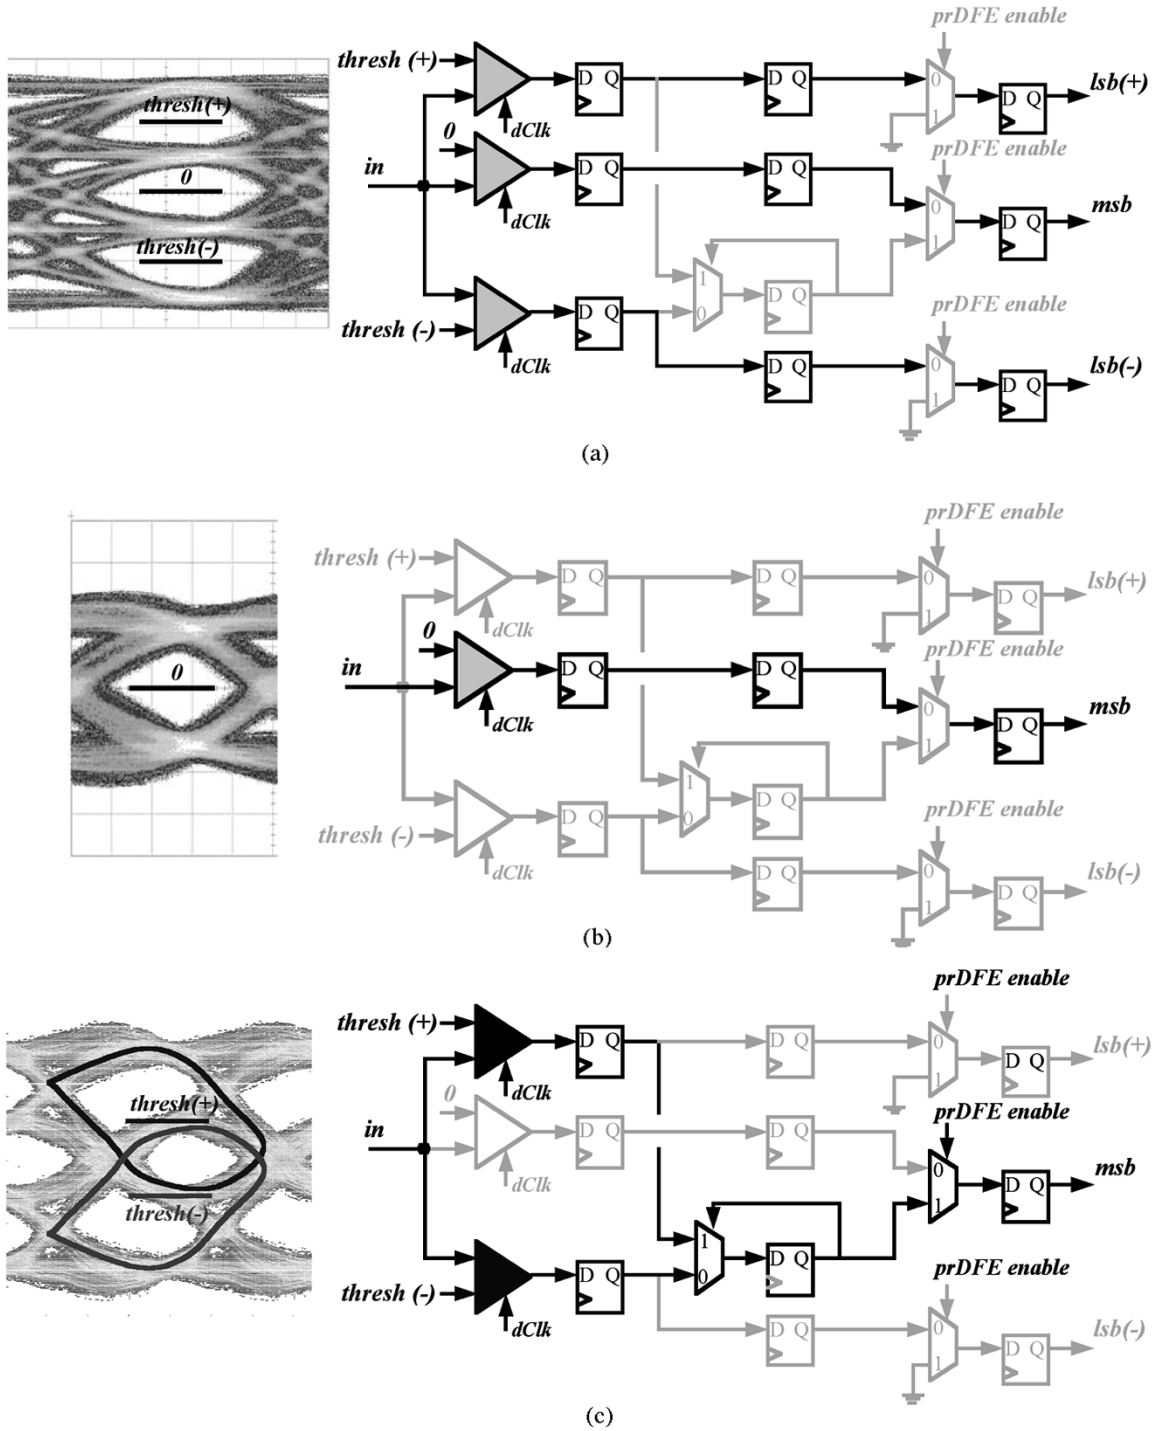

Fig. 4. Integration of PAM2 partial response DFE receiver with loop unrolling into PAM4 receiver by re-use of PAM4 *lsb* slicers: (a) PAM4 mode, (b) PAM2 mode, (c) PAM2 with one-tap loop unrolling.

The multiplexing method enables both swapping of the outputs of each of the samplers with the adaptive sampler and independent swapping of the adaptive twist DAC with the threshold DAC for a particular sampler. This enables the calibration of both the sampler offset and any residual threshold DAC errors.

A. Dual-Mode Implementation

Another change to this link from the previous design was to add one tap of DFE to increase the performance of the link in PAM2 mode by reusing some of the PAM4 hardware. In general, the first tap of a DFE filter is the most difficult to implement, since the circuit must sample the input, resolve its value, and then subtract a signal proportional to that value from the input in one symbol time. As shown by Kasturia [2] and more recently by Sohn [9], this tight timing constraint can be avoided by unrolling the critical loop and making two decisions each cycle. One comparator decides the input as if the previous output was a 1, and the other comparator decides the input as if the pre-

vious bit was a 0. Once the previous bit is known, the correct comparator output is selected, as shown in Fig. 3.

Instead of just one data sampler for PAM2 signaling, the receiver now has two samplers that are offset by $\pm\alpha$[1], anticipating the impact of the trailing ISI tap α, from a previously sent symbol of value of ± 1. While this presents significant overhead in a simple PAM2 receiver, the inactive samplers of the PAM4 link (when in PAM2 mode) are re-used to implement such a scheme with no additional front-end hardware overhead, as shown in Fig. 4.

Given that in such a signaling scheme the first tap of trailing ISI is not really physically cancelled in the channel but rather predicted by the receiver (based on previously received data), we are faced with two issues with respect to automatic link configuration. First, it is necessary to modify the adaptive algorithm for transmit pre-emphasis to tolerate one tap of trailing ISI, and second, the magnitude of this trailing ISI needs to be estimated such that data slicing and clock recovery can be robustly performed. The next section shows how this can easily be accomplished with a small modification to our adaptation algorithm.

B. Link Equalization

A sign-sign LMS (a derivative of the well-known least-mean square (LMS) algorithm [10]) is used to adapt the equalization taps since it is one of the simplest adaptive algorithms to implement. It creates the updates for the tap coefficients (w) based only on the sign of the data and the measured error

$$w_{n+1}^{k} = w_{n}^{k} + \Delta_{w}\mathrm{sign}(d_{n-k})\mathrm{sign}(e_{n}) \tag{1}$$

where n is the time instant, k is the tap index, d_n is the received data and e_n is the error of the received signal with respect to the desired data level, dLev.

1) Dual-Loop Adaptive Equalization: One issue in using sign-sign LMS for transmit pre-emphasis based equalization, which is often used in high-speed links, is that the ideal reference level dLev from which the error signal is created is unknown *a priori*. This problem arises because the peak output swing constraint in the transmitter forces the equalizer to attenuate the low-frequency components of the signal to match the loss of the signal at high frequencies [see Fig. 5(a)]. Thus, the amount of voltage swing available at the receiver depends on the frequency characteristics of the channel.

One of the solutions to alleviate this effect, proposed in our earlier work [11], was to introduce a variable gain element at the receiver (prior to slicer input), which was adjusted during adaptation such that constant reference levels are maintained in the data slicer. A more practical and power efficient approach for high-speed links is to adaptively adjust the reference level of the data slicer, rather than amplifying the signal. Thus, a second loop is created, which adjusts dLev to track the signal level using the following updates:

$$\mathrm{dLev}_{n+1} = \mathrm{dLev}_{n} - \Delta_{\mathrm{dLev}}\mathrm{sign}(e_{n}). \tag{2}$$

[1]In a differential system, it is highly likely that thresholds will be symmetric around zero, i.e., $\pm\alpha$. However, even in the case of asymmetry, as we will see later, since α is determined by locking to different signal levels, the algorithm can be easily extended to obtain positive and negative α separately at the expense of two α DACs and two registers.

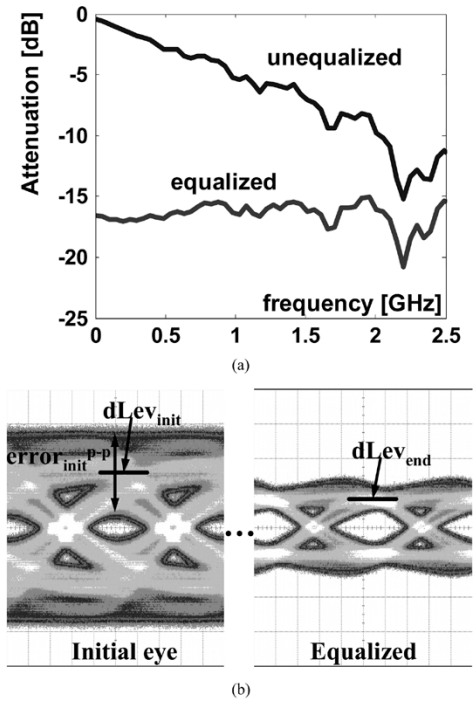

Fig. 5. Effect of peak voltage swing constraint on transmit pre-emphasis: (a) Frequency view. (b) Scaling of the dLev reference loop (2) in a dual-loop interaction with the equalizer loop. As the signal gets more equalized, scaling in the transmitter decreases the value of received signal and reference loop adjusts dLev accordingly.

At each iteration, the adaptive sampler is adjusted using (2) to provide the error signal e_n for both the signal level (2) and equalizer tap (1) loops. The peak-to-peak error and dLev setting are shown in Fig. 5(b), for initial and final iteration of the algorithm.

In order to obtain the highest signal levels at the receiver, maintain transmit output peak swing constraint, and avoid the trivial stability point of both loops (at zero tap magnitudes and signal level), the proposed values of the equalization taps after every iteration (1) need to be rescaled such that the sum of their magnitudes always equals the maximum allowed by the peak swing constraint. A simple, implementation driven approximation of this rescaling modifies the update algorithm such that the update on the main tap is computed from the updates of the other taps and the peak constraint requirements, rather than using its own update information.

At first it might seem that one would need an error sampler for each data level, as proposed by Stonick *et al.* [4]. Instead, only one adaptive sampler is used and updates are performed only when data is received that corresponds to the signal level at which the adaptive sampler is located. The convergence time is traded-off for receiver simplicity since convergence is not a problem with multi-Gb/s data rates and slow channel changes.

2) Decision-Feedback Equalization Using Loop Unrolling: Similarly, the dual-loop adaptive framework can be extended directly to support feedback equalization using loop unrolling. Instead of filtering the error signal and loop updates (for both dLev and equalizer taps) by the bit values that form the current received symbol, data filtering can be applied with the current and past bit in order to lock the dLev to one of the four signal levels $(\pm 1, \pm\alpha)$, present in a one-tap DFE

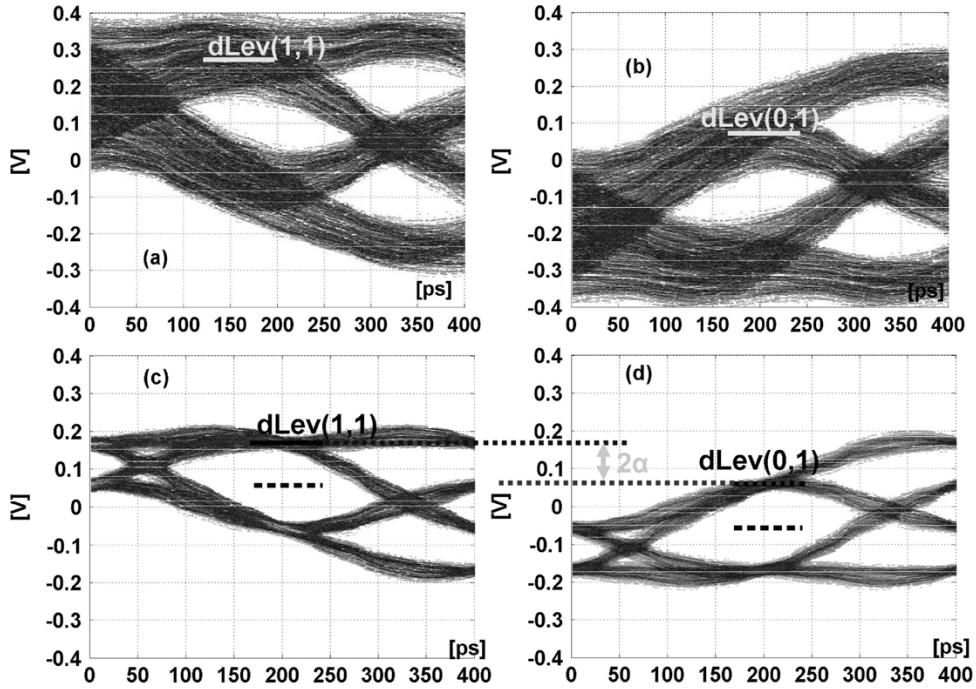

Fig. 6. Joint equalization and extraction of the trailing tap magnitude. Plots are based on simulation using the measured pulse response, Fig. 18, obtained with the adaptive sampler; symbol time is 200 ps: (a) Locking of dLev to (1, 1) level—eye as seen by the upper sampler in Fig. 3(b), (b) Locking to (0, 1) level- eye as seen by the lower sampler in Fig. 3(b), (c) Final locking point of dLev to (1, 1) level after equalization, (d) Final locking point of dLev to (0, 1) level after equalization. Sampler thresholds are offset by the extracted final magnitude of the trailing ISI $0.5 * (\mathrm{dLev}(1,1) - \mathrm{dLev}(0,1))$ (dashed line).

system. This filter is very similar to data filtering for PAM4 equalization. A similar algorithm, but without data-based update filtering, was proposed for one-tap DFE by Winters and Kasturia [12] and incurs significant sampler overhead.

Using just one adaptive sampler and data-based update filtering the size of the first trailing ISI tap (α) is estimated in an iterative manner. In the first phase, loop updates are filtered by the $(\mathrm{d_{n-1}}, \mathrm{d_n}) = (1,1)$ criterion to lock dLev to the $1+\alpha$ level, and in the second phase, updates are filtered by $(\mathrm{d_{n-1}}, \mathrm{d_n}) = (0,1)$ to lock to the $1-\alpha$ level. During these two phases, the equalizer only compensates for the error caused by the ISI taps other than the first trailing tap (α), as shown in Fig. 6. The value of the first trailing tap (α) is then extracted as half of the difference between dLev values in the first $(1+\alpha)$ and second $(1-\alpha)$ phase. Once the first trailing tap value is updated, the adaptation of the equalizer continues for a while at the dLev from the first phase, before the new first trailing tap value (α) is extracted.

In this manner, equalization, and locking phases one and two are interleaved such that the optimal value of first trailing ISI tap (α) is found at the point when all other ISI has been minimized by the transmit equalizer and long-latency feedback equalizer (reflection canceller [1]). This is necessary since the absolute magnitude of the main tap and first trailing ISI tap (α) change due to rescaling which maintains the peak power constraint in the transmitter.

C. Clock and Data Recovery for Loop Unrolling

Correcting for the trailing ISI at the data samplers unfortunately is not enough in a system with edge samplers. We have already seen in Fig. 6 that the presence of the trailing tap of ISI causes the received signal to have four levels, similar to PAM4, albeit nonuniformly separated. Thus the transitions from one level to another are guided by the values of the future, current and immediately preceding data bits, as shown in Fig. 7, and do not cross zero at the same point in time. They instead form two distinct modes or principal zero crossings, denoted by arrows in Fig. 7. In order to avoid this bi-modal behavior, one type of transition could be filtered out by filtering the edge crossings in the clock and data recovery (CDR) block. This is done in a way similar to that in PAM4 clock and data recovery, where edge-filtering is used to eliminate the edges that cause tri-modal zero crossing distributions [1]. Using this approach we can directly extend the PAM4 CDR filtering based on two-bit symbols, to partial response CDR filtering based on pairs of current and preceding bits.

Since edge filtering decreases the probability of CDR updates and puts additional constraints on first-order CDR loops in plesiochronous systems, additional samplers are used to record the minor transitions in PAM4 systems [1]. In PAM2 partial response mode of operation, these available samplers (used as *lsb* edge samplers in PAM4 mode) are offset by the amount of trailing ISI, aligning the edge slicing timing as shown by the left arrow and three dotted levels in Fig. 7. In this way, no transitions are lost and the rate of CDR updates is maximized.

The clock and data recovery front-end remains the same as in the PAM4 case (three edge samplers providing tentative early/late information). In PAM4 mode the transition filtering section uses $lsb(+)$, msb, and $lsb(-)$ data from the current symbol n and the previous symbol $n-1$, as shown in Fig. 8(a). In PAM2 partial response mode msb data from the current bit n, previous bit $n-1$ and additional bit of history $n-2$ are used, as shown in Fig. 8(b).

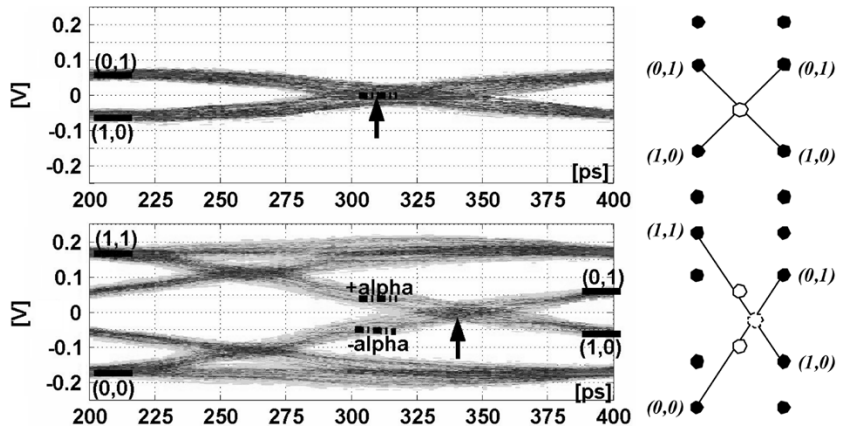

Fig. 7. Bi-modal transitions in $1+\alpha\mathrm{D}$ channel: first mode $(1,1)\rightarrow(1,0)$, and $(0,0)\rightarrow(0,1)$, second mode $(0,1)\rightarrow(1,0)$ and $(1,0)\rightarrow(0,1)$.

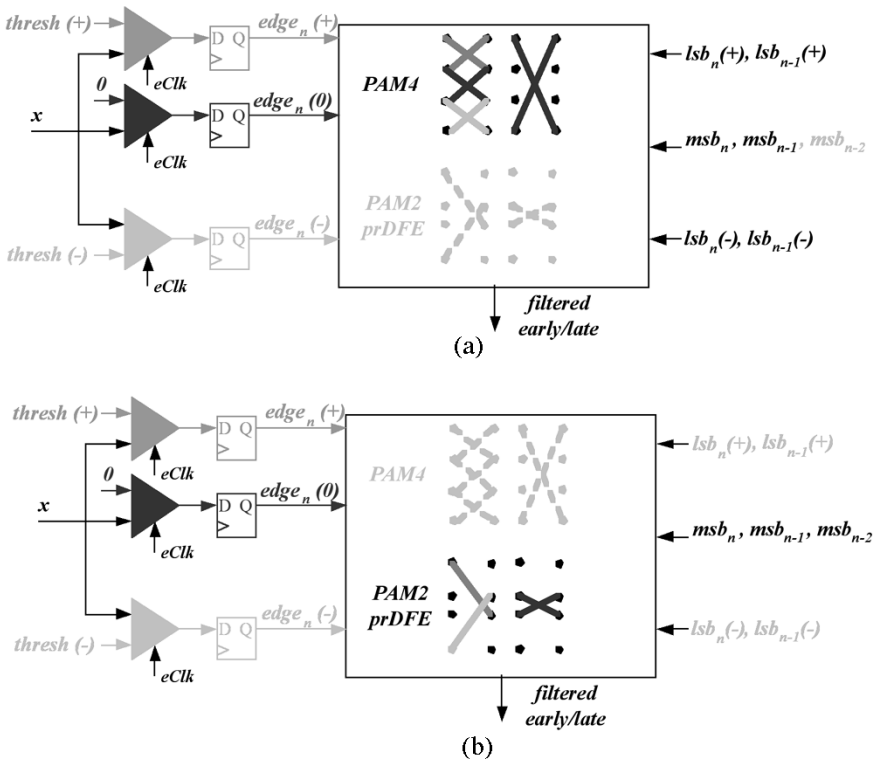

Fig. 8. Generation of early/late updates in 2x oversampling CDR loop: (a) PAM4 mode; (b) PAM2 with partial response DFE mode.

III. Back-Channel

To adapt the transmitter pre-emphasis filter the error information has to be communicated back to the transmitter. This can be achieved by using the orthogonality between differential and common-mode signals to provide a separate pathway on the same physical channel. The system block diagram with differential forward channel and common-mode back-channel is shown in Fig. 9. While in theory, the differential and common-mode signals are orthogonal, in practice, inevitable nonidealities lead to signal coupling between modes, causing signal integrity degradation in both domains. The transfer functions s_{21} and s_{34} are the channel response of the differential forward channel and common-mode back-channel, respectively. The crosstalk of the differential transmitter to the common-mode receiver is represented by the s_{31} transfer function. The coupling of the back-channel transmitter to the forward channel receiver corresponds to the s_{24} transfer function. Since it was imperative that the back-channel not affect the forward channel's performance, this coupling is examined next.

A. Effect of Common-Mode Signal on Differential Signal

Differential signals generated by uncorrelated common-mode transmission appear as noise to the differential receiver, degrading the signal integrity of the forward channel. This common-mode-to-differential conversion is, in part, a result of the mismatch in the passive components (connectors, vias, backplane traces, etc.) and can be simply analyzed as crosstalk. Measured s_{24} of a typical $26''$ FR4 backplane bottom trace and

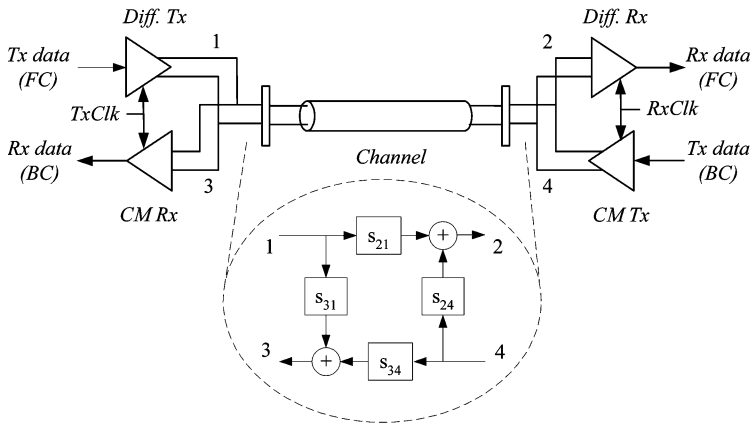

Fig. 9. Simultaneous common-mode and differential signaling system block diagram.

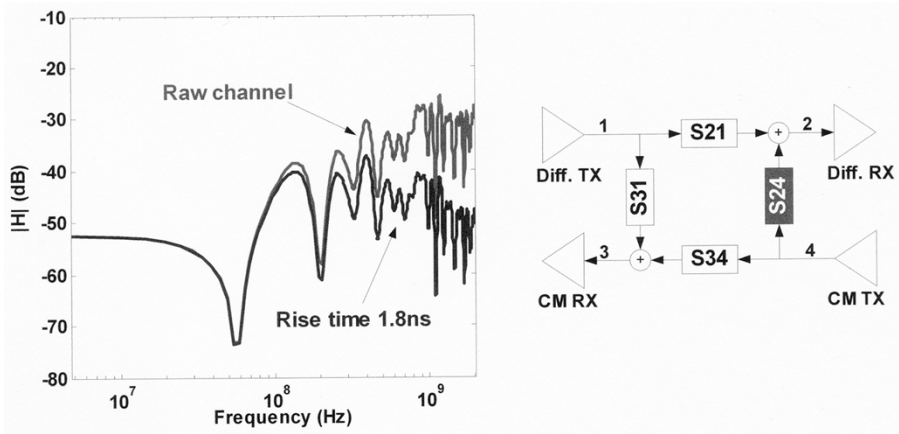

Fig. 10. Measured common-mode transmitter to differential receiver crosstalk transfer function (s_{24}) of a typical 26″ FR4 backplane bottom trace and linecards with HSD connectors, with and without increased rise time in the backchannel transmitter.

linecards with two connectors (Fig. 10) shows the significant increase in this crosstalk beyond 100 MHz.

To minimize the impact of this crosstalk on the forward link, the bandwidth of the back-channel is limited by increasing the rise time of the common-mode transmitter. The crosstalk response with the implemented increased rise time is overlaid in Fig. 10. Slower rise time attenuates crosstalk due to mismatch in passive components for frequencies outside of the back-channel bandwidth. Since this crosstalk is a proportional noise source, the back-channel swing can be decreased to further mitigate the crosstalk injected onto the forward channel.

Another cause of forward link margin degradation is the effect of the common-mode signal on the differential receiver's offset and sensitivity. Fig. 11(a) shows the measured offset over a range of common-mode voltages for the differential receiver. Static offset can be removed with one-time digital offset cancellation. Dynamic changes in offset directly reduce the voltage margin of the differential receiver, unless the receiver includes dynamic offset cancellation. Fig. 11(b) presents measured differential receiver sensitivity as a function of the receiver common-mode voltage. Reduction in receiver sensitivity due to common-mode variations will also decrease the voltage margin of the forward link. To mitigate this degradation, the common-mode signal magnitude and range should be chosen so that the corresponding range of offset and sensitivity changes is within the noise budget. With a fixed forward link noise budget, back-channel signal magnitude is maximized by choosing the operating point to be in the plateaus of Fig. 11(a) and (b).

Additionally, the common-mode signal can also "steal" the headroom from the differential transmitter outputs. To enable back-channel signaling transmit driver and receiver pre-amplifier share an adjustable supply that can be increased from 1 V to 1.3 V, separately from the core link supply of 1 V. This increased driver supply comfortably compensates for the back-channel swing without violating headroom constraints of the transmit driver and also enables larger differential signal swing for channels with high loss.

To satisfy the tight noise budget in high performance applications, the back-channel signal is limited to a very small swing relative to the forward signal. Our design has a programmable peak-to-peak common-mode signal swing of 12–100 mV to allow tradeoffs between forward-channel and back-channel reliability.

B. Effect of Differential Signal on Common-Mode Signal

Differential-to-common-mode conversion can severely degrade the signal integrity of the low-swing back-channel. This conversion is again a result of the mismatch in the passive components of the link. The s_{31} crosstalk transfer function is shown in Fig. 12. Since the bandwidth of the back-channel

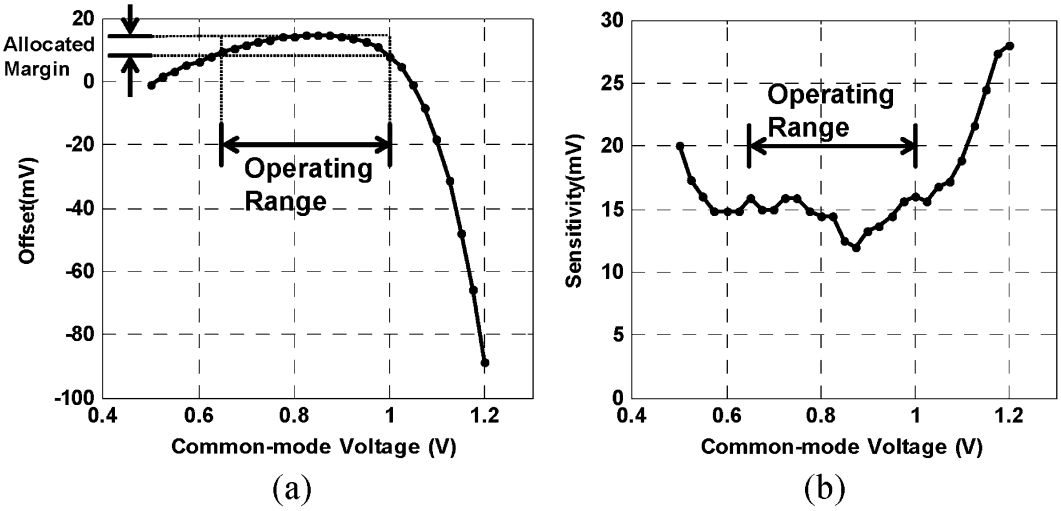

Fig. 11. (a) Measured differential receiver offset versus common-mode voltage. (b) Measured differential receiver sensitivity versus common-mode voltage.

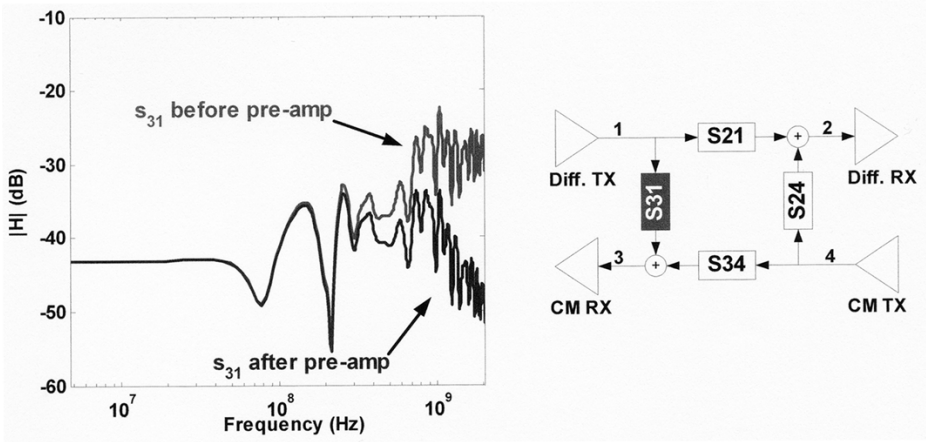

Fig. 12. Measured crosstalk of differential transmitter to common-mode receiver (s_{31}) of a typical 26″ FR4 backplane bottom trace and linecards with HSD connectors, before and after a two-pole low-pass preamp filter at 650 MHz in the backchannel receiver.

is below the knee of the high-pass nature of the cross-talk, a band-limited pre-amplifier in the back-channel receiver greatly reduces the effect of this noise.

In addition to the crosstalk caused by mode conversion, the differential transmitter directly generates common-mode noise. Fig. 13(a) shows the measured frequency response of the common-mode glitch caused by the asymmetrical rising and falling edges of each leg in a differential transmitter. The glitch happens at every forward channel bit transition, Fig. 13(b), with most of its energy concentrated in the frequency range outside that of the back-channel bandwidth. Like s_{31}, the frequency content of the glitch that is in the band of the common-mode receiver directly interferes with the back-channel. The low-pass filter added to reduce crosstalk also reduces the amount of differential transmitter inflicted common-mode glitch noise seen by the back-channel receiver.

As the back-channel signal magnitude is restricted to minimize the interference with the forward link and to contain electro-magnetic interference, the total common-mode noise, even after the low-pass filter, can be on the same order as the back-channel swing. Fortunately, since the back-channel is only used to adapt the transmitter's configuration, error rates as large as 30% are acceptable (see Section IV).

C. Signaling Scheme

To meet the constraints of robustly extracting timing and voltage references in a noisy environment a tri-level return-to-null signaling system is chosen, as shown in Fig. 14. Since each bit has two transitions, timing recovery is simplified. Simple edge detection, achieved by differentiating the oversampled digital output, provides both data extraction and synchronization. The oversampling rate is great enough to provide high tolerance to frequency offset variations. In our design, an oversampling ratio of 20x is used. The differentiator is implemented with a voltage tracking loop at the receiver (see Fig. 16); the voltage reference is extracted by using the oversampled digital output as feedback to track the common-mode signal via an integrator. Edges are detected when the number of 1's (for a rising edge) or 0's (for a falling edge) from the oversampled common-mode signal in a fixed size window exceeds a configurable threshold. As an additional benefit, offset in the common-mode receiver and slow variations in common-mode (e.g., baseline wander in ac-coupled systems) are cancelled in the process of signal tracking.

D. Implementation

To minimize its impact on the forward differential link, a common-mode transmitter must provide both high-output

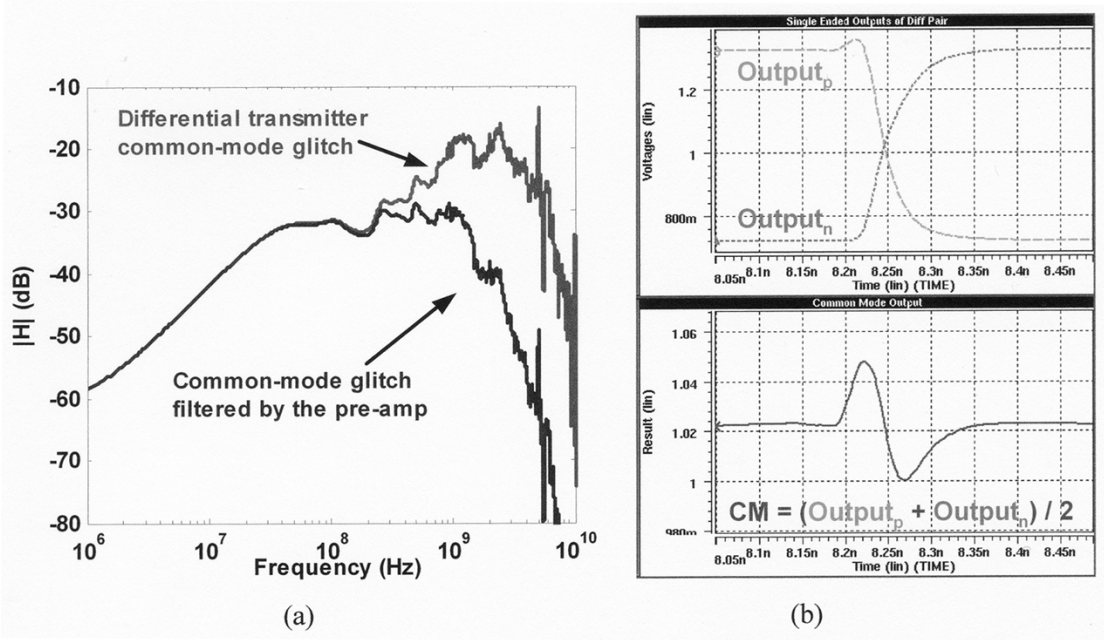

Fig. 13. (a) Measured frequency responses of the common-mode glitch generated by the differential transmitter before and after a two-pole low-pass preamp filter at 650 MHz. (b) Common-mode glitch in time domain due to asymmetric rise and fall times at the transmitter.

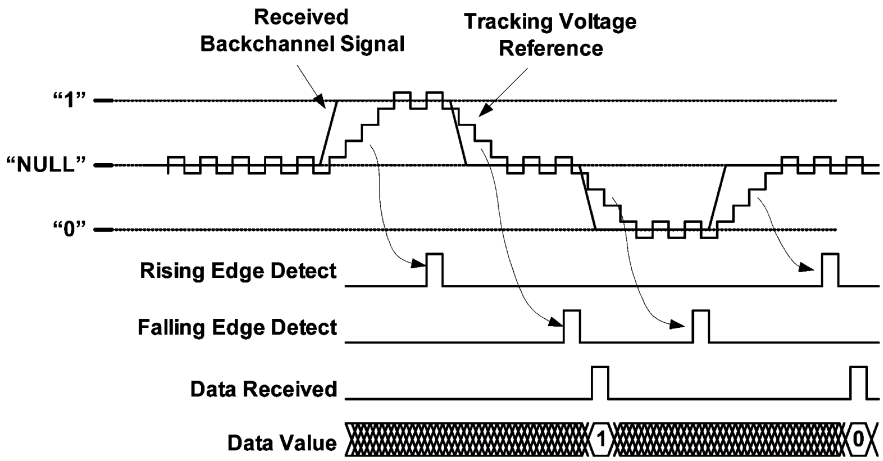

Fig. 14. Oversampled, edge-encoded three-level signaling with tracking receiver (three characters are denoted as a "1", "0", and "null"). Sliding window of 10 decides a "1" on 7 ups, and a "0" on 7 downs.

impedance and low parasitic capacitance. As shown in Fig. 15, a pair of carefully matched cascode structures is used to reduce channel length modulation so that the difference in current drawn from PADp and PADn is as small as possible in the presence of the received differential signal. The relatively small cascode devices isolate the large parasitic capacitance of the bottom current source devices from the pads. Rise time control is implemented directly from the digital data signal using an undersized driver.

In the back-channel receiver (Fig. 16), a resistive divider connected to both inputs generates the common-mode voltage. To prevent degradation of the differential receiver termination, a pair of high-valued, carefully matched on-chip resistors is used. A two-stage differential preamplifier with a gain of 3 acts as a two-pole low-pass filter at 650 MHz to reduce high-frequency noise from the differential transmitter.

To provide flexible back-channel bandwidth, the tracking receiver employs a digital integrator. The integrator consists of an 8-bit up-down counter and a DAC. The digital integrator provides good visibility and testability. It enables flexible back-channel data rate: as long as the DAC settling time does not extend beyond the sampling period, the tracking rate is solely determined by the clock that drives the integrator and the step size of the DAC. The digital integrator requires a wide-range, high-resolution, linear DAC to accurately track the low-swing common-mode signal. The wide range (up to 250 mV) is required for compatibility with ac-coupled and dc-coupled systems.

In this work, a pre-existing DAC circuit was re-used for the back-channel receiver. This DAC is the performance limiter in the back-channel system because of its large differential nonlinearity (DNL). This DNL results in an irrecoverable loss of voltage margin for the back-channel. The use of one coarse DAC to track the natural common-mode level set by the differential transmitter, and a second fine DAC to track the common-mode

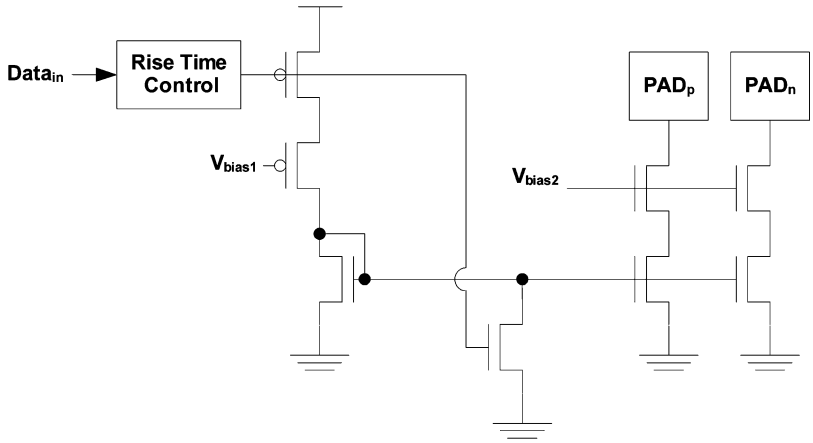

Fig. 15. Cascoded common-mode transmitter minimizes channel length modulation and common-mode-to-differential crosstalk.

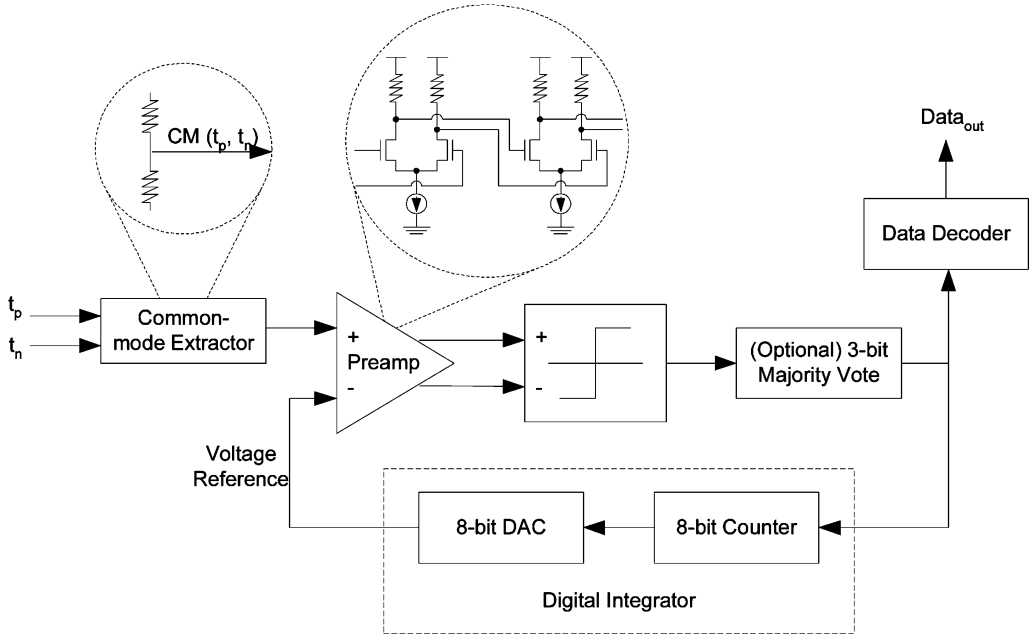

Fig. 16. Common-mode receiver block diagram. Two-stage preamp amplifies the common-mode signal and low-pass filters the high-frequency common-mode noise. Digital integrator provides both reference level extraction and offset cancellation.

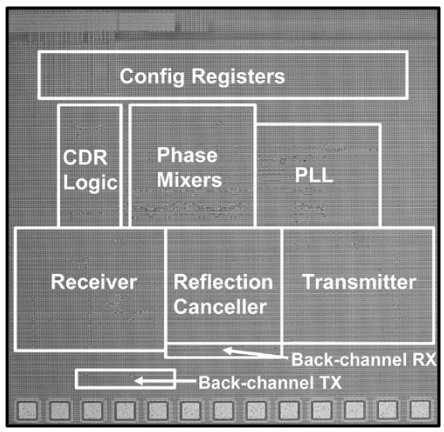

Fig. 17. Die photograph.

signal could relax the design constraints of a single wide-range, high-resolution DAC.

Since the back-channel is a noisy signaling environment, a parity-check bit is added to each transmitted packet. This additional bit allows the detection of single-edge slips, which are common in edge-based systems.

In our design, the back-channel operating frequency is tied to the frequency of the ASIC core system clock, which can be anywhere from 25 MHz to 640 MHz. Given that 20x oversampling is used in the back-channel receiver together with return-to-null signaling scheme, the frequency range of back-channel signals is from 0.65 MHz to 16 MHz. Although the measurements shown in the next section correspond to the dc-coupled operation, we establish here the ac-coupling and load capacitance guidelines to enable back-channel signal propagation as well as 6 dB of common signal return loss (starting at 100 MHz). An ac-coupling capacitance of 10–100 nF is large enough to pass a 0.65 MHz signal. A common signal terminating capacitance of 60–80 pF at the differential forward link receiver (back-channel transmitter) has high enough impedance at 16 MHz (so as not

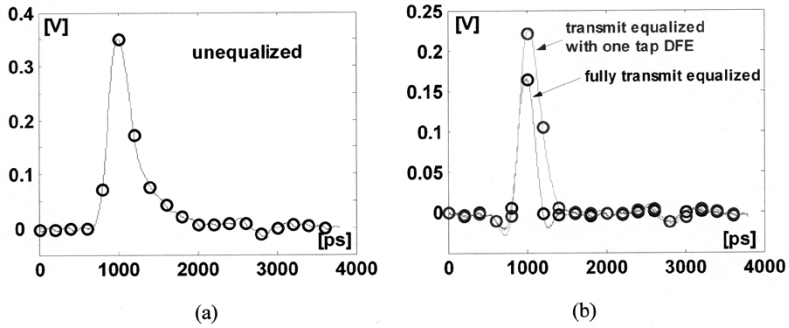

Fig. 18. *E-scope* [8] of the pulse response: (a) unequalized, (b) transmit equalized with one tap DFE and fully transmit equalized. Dots indicate symbol spaced sample points (symbol time is 200 ps).

to short the back-channel common signal), and low enough impedance at 100 MHz (the starting frequency for common signaling return loss specification) to provide 6 dB of common signal return loss.

IV. Experimental Results

A transceiver chip was designed and fabricated in a 0.13 μm CMOS process to investigate the dual-mode adaptive forward link and the common-mode back-channel. Operating at a supply of 1 V, the forward link achieves a data rate range of 1–10 Gb/s by combining both PAM2 and PAM4 modulation, while the back-channel is set to data rate of 1–16 Mb/s with signal swing range of 12–100 mV. The aggregate worst-case cost of signaling on the forward link is 40 mW/Gb/s.

Examining the die photograph in Fig. 17, it is clear that large area is dedicated to long-latency feedback equalization (reflection canceller [1]). The transmitter and receiver areas are also increased due to equalization. The synchronization sub-systems (CDR and PLL) consume a little more than half of the total link power. As can be noted from Fig. 17, the area (and also the power) of the back-channel is negligible when compared to that of the forward link.

A. Forward Link

The adaptive sampler can scan out the pulse response of the whole channel as seen by the receiver, including any bandwidth limitations in the receiver. Fig. 18 illustrates the pulse responses before and after equalization. The pulse response equalized for one tap DFE at 5 Gb/s, 26″ FR4 channel, is about 60 mV (40%) larger than the fully equalized pulse, due to the peak output power constraint in the transmitter.

This difference is also visible on the bit-error rate (BER) versus noise margin curves, shown in Fig. 19. The steep slope of the BER versus noise margin curves suggests that random noise components (jitter and voltage thermal noise) are relatively small and that ISI is the most dominant error term.

It is interesting to observe the shape of the equalized eye in a loop unrolling DFE scheme. Again, using the adaptive sampler to monitor the signal, the statistical eye can be measured, as shown in Fig. 20. In this plot we show the positive-conditioned eye. While not as symmetric as a fully equalized PAM2 eye, the eye for one-tap loop-unrolled DFE is actually slightly more robust to jitter.

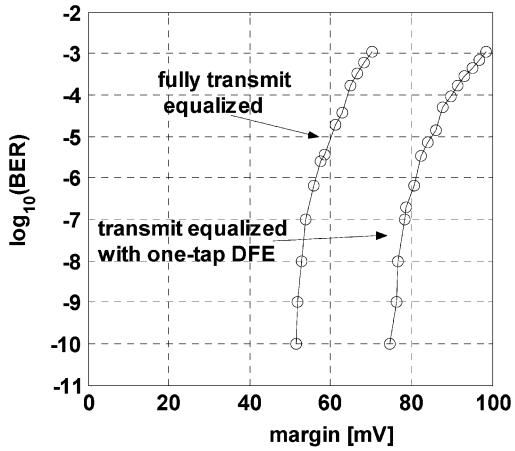

Fig. 19. Comparison of bit-error rate (BER) versus receiver noise margin for fully transmit equalized link and transmit equalized with one tap DFE.

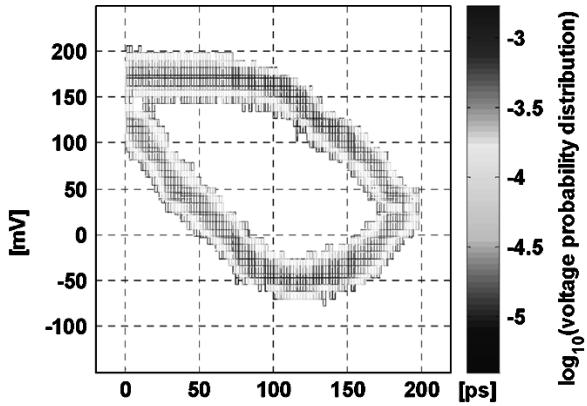

Fig. 20. Statistical shmoo of the eye diagram as presented to the positive *lsb* sampler for one tap DFE.

Measured peak-to-peak jitter from the 2.5 GHz recovered clock shows that CDR dither decreases from 14 ps to 5 ps when one tap DFE is used instead of full transmit pre-emphasis. The tri-modal edge distribution is partially avoided in the one-tap DFE scheme since the first post tap of the transmit pre- emphasis is not significantly engaged. Inherent PLL jitter was 26 ps peak-to-peak.

In Fig. 21, the convergence of the dual-loop adaptive algorithm is shown. Fig. 21(a) represents the learning curve of the reference level (dLev) loop, and Fig. 21(b) shows the convergence of the equalizer taps. Loop updates are filtered on received data being high (since in this case the adaptive sampler tracks the

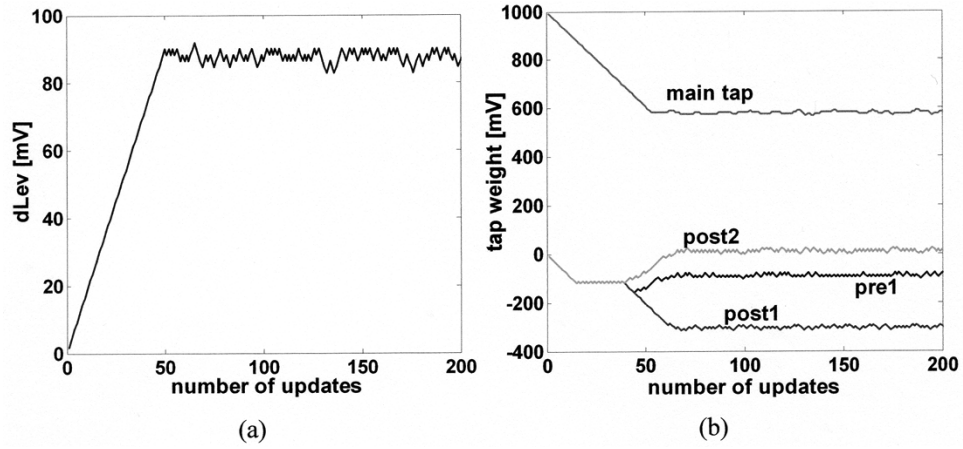

Fig. 21. Dual-loop convergence. (a) Reference level (dLev) loop. (b) Equalizer loop. PAM2 at 5 Gb/s over 20″ FR4.

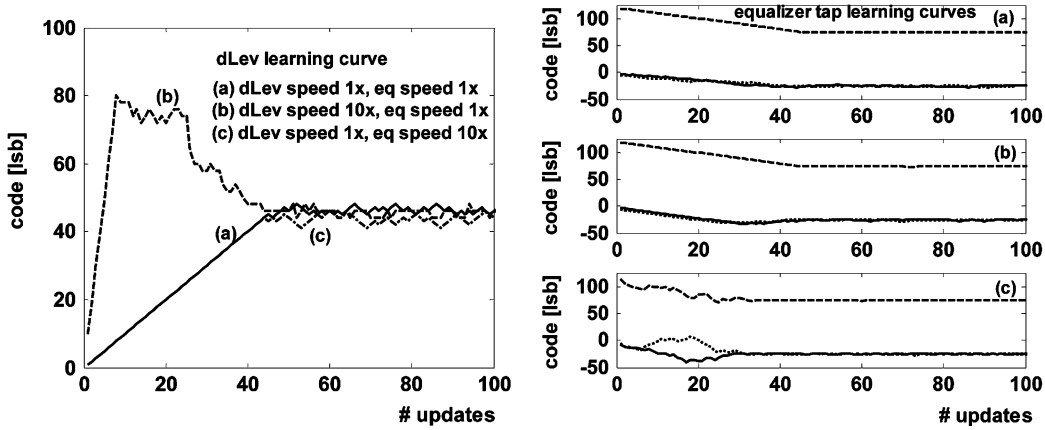

Fig. 22. Dual-loop adaptive learning curves for different speeds of the dLev and equalizer tap loops, PAM2 at 5 Gb/s over 20″ FR4.

positive signal level) and then block averaged by 127 to smooth the sign-sign gradient estimate. It only takes about 50 updates to lock to the signal level. After the reference level has settled, the taps quickly converge to their optimal values.

One of the most important issues in the dual-loop adaptive algorithm is the balance of the update rates for the equalizer and reference level loops. Our measurements show that the equalization algorithm is stable for a relatively wide range of update speeds of one loop with respect to another (Fig. 22).

B. Back-Channel

Fig. 23 shows the common-mode back-channel transmitted waveform and the relative tracking rate of the digital tracking loop, probed at the back-channel receiver pads. The digital loop tracks the transmitted back-channel pulses in approximately ten clock counts. The dithering at each plateau shows the bit boundaries as seen by the back-channel receiver. Depending on the common-mode of the forward link, the back-channel receiver may be positioned at the location of high DNL in the tracking DAC, as shown in Fig. 23(b), resulting in additional loss in signal margin. However, as we will show later, we were still able to operate the back-channel under these conditions with satisfactory performance.

The common-mode back-channel swing can be adjusted to provide a tradeoff between forward channel signal integrity and back-channel noise immunity. Fig. 24 shows the packet error and drop rates of the back-channel at different peak-to-peak common-mode swings. A packet is dropped when a parity error or an error in synchronization is detected (i.e., when the time between any two consecutive received bits in a packet exceeds a set threshold).

While packet error and drop rates are important for reliable transmission of various configuration and link information data, it is interesting to note that theoretically for adaptive algorithm convergence it is only necessary to have an error rate of less than 50%. The packet drop rate only influences the total adaptation time. For example, Fig. 24 shows that packet drop rate decreases nearly two orders of magnitude when back-channel swing is increased from 18.75 mV to 25 mV, so adaptation time improves rapidly. To illustrate this effect, the convergence of transmit pre-emphasis taps is shown in Fig. 25 as a function of the number of sent updates for these two back-channel swing amplitudes. These plots indicate that effective adaptation is still possible even at very low back-channel amplitude. Additionally, reliable communication for other configuration needs can be achieved for back-channel swings of 50 mV and higher, as indicated in Fig. 24.

Even at the largest available back-channel swing of 100 mV, there was no measurable impact on the noise margin of the forward channel at a target BER of 10^{-12}. This indicates that the noise induced by the back-channel on the forward channel is less than 2 mV, which is the minimum change in margin detectable

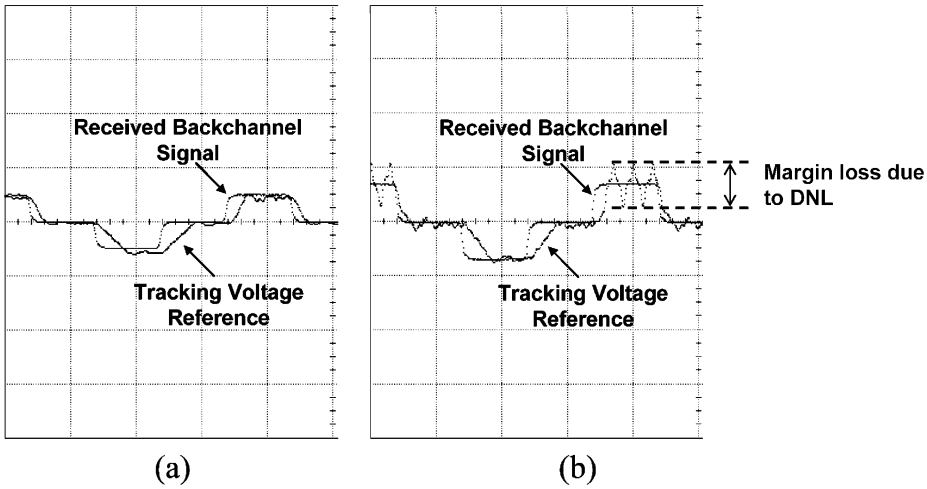

Fig. 23. Scope captured common-mode backchannel transmitted waveform overlaid with the tracking signal of the common-mode back-channel receiver. (a) Biased to avoid DAC DNL. (b) Biased at the DAC DNL.

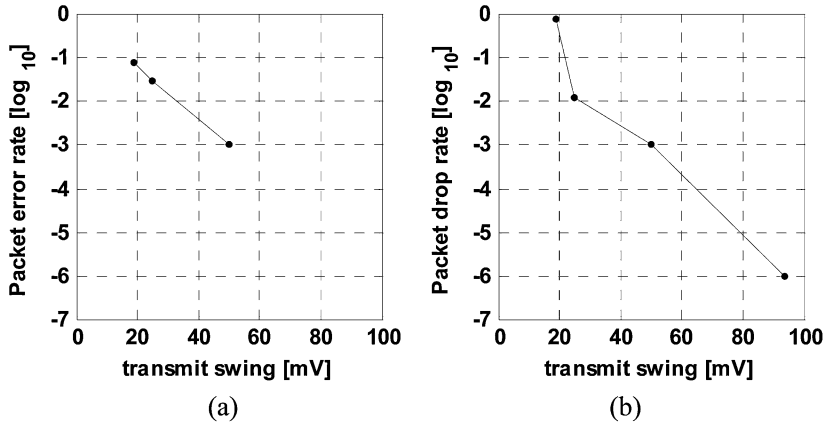

Fig. 24. (a) Packet error rate and (b) packet drop rate as a function of peak-to-peak common-mode swing.

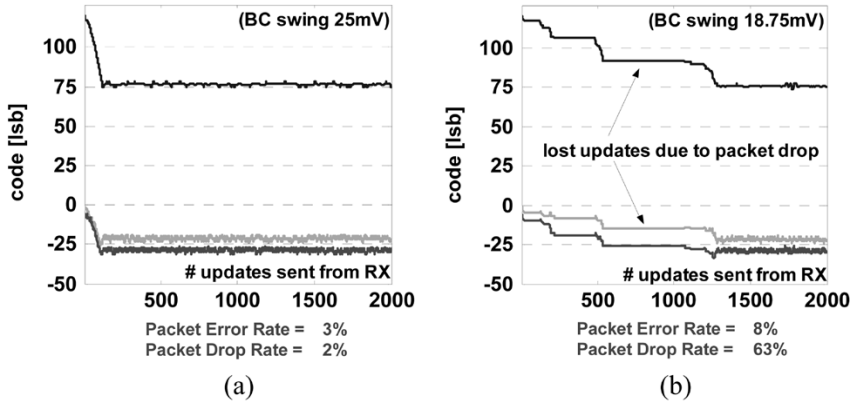

Fig. 25. Transmit pre-emphasis tap convergence: (a) Back-channel swing of 25 mV with packet drop rate of 2% and packet error rate of 3%. (b) Back-channel swing of 18.75 mV with packet error rate of 8% and packet drop rate of 63%.

by our measurement setup. Fig. 26 illustrates the PAM2 and PAM4 forward link shmoos when the back-channel is turned on and off.

V. Conclusion

It is possible to integrate a PAM2 one-tap DFE into a PAM4 receiver with minimal additional hardware by leveraging the multi-level aspects of the partial response signals in loop-unrolled DFE. Clock and data recovery techniques for these partial response signals are derived from standard multi-level edge filtering schemes. Adaptive equalization can also be added to a transceiver for a small hardware cost. The key is to first modify the popular sign-sign LMS procedure to enable adaptation under peak voltage swing constraint in the transmitter and then to incorporate data filtering methods. The data filtering enables adaptation using a single monitoring sampler even in multi-level schemes like PAM4 and loop-unrolled PAM2 DFE.

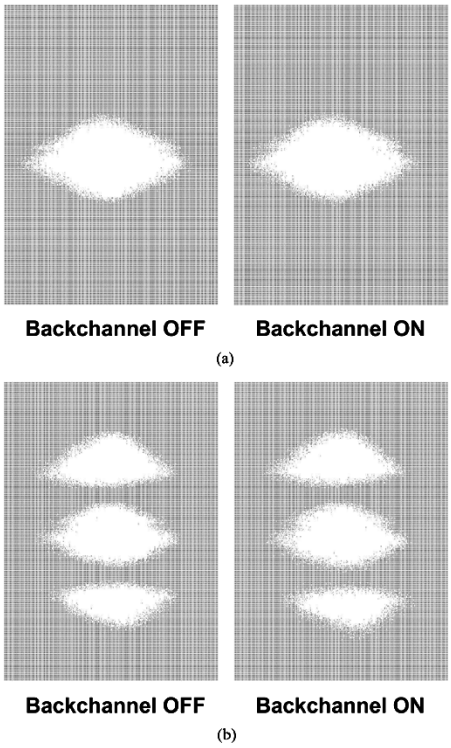

Fig. 26. Effect of back-channel on forward link, *e-scope*. (a) Shmoo of PAM2 on forward link. (b) Shmoo of PAM4 on forward link.

Self-contained link re-configuration and adaptation of transmit pre-emphasis has been further enabled by using the common-mode signaling for the reverse direction data flow. Exploration of the signal integrity issues and common-mode impact on differential receiver circuits shows that common-mode back-channels with sub-50-MHz bandwidth, and swings as large as 100 mV, can be implemented with minimal impact on the error rate of the forward link. The mostly digital implementation allows for flexibility in operation and robustness to different sources of noise present in such a system.

Taken together these techniques enable a single, hardware-efficient link cell design to operate autonomously, extending its data rates to 5–10 Gb/s over a variety of channels.

ACKNOWLEDGMENT

The authors would like to acknowledge the help and support of the entire Rambus RaserX design team. V. Stojanović, A. Ho, and E. Alon would like to thank I. Stojanović, M. Ho, and M. Lee for their wholehearted support.

REFERENCES

[1] J. Zerbe, C. Werner, V. Stojanović, F. Chen, J. Wei, G. Tsang, D. Kim, W. Stonecypher, A. Ho, T. Thrush, R. Kollipara, M. Horowitz, and K. Donnelly, "Equalization and clock recovery for a 2.5–10 Gb/s 2-PAM/4-PAM backplane transceiver cell," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2121–2130, Dec. 2003.

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

[3] G. Sheets and J. D'Ambrosia, "The impact of environmental conditions on channel performance," presented at the DesignCon 2004, Santa Clara, CA.

[4] J. T. Stonick, G.-Y. Wei, J. L. Sonntag, and D. K. Weinlader, "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.

[5] V. Stojanović, 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 *Proc. IEEE Symp. VLSI Circuits*, Jun. 2004, pp. 348–351.

[6] T. Gabara, "Phantom mode signaling in VLSI systems," presented at the 19th Conf. Advanced Research in VLSI, Salt Lake City, UT, Mar. 2001.

[7] A. Ho, V. Stojanović, F. Chen, C. Werner, G. Tsang, E. Alon, R. Kollipara, J. Zerbe, and M. A. Horowitz, "Common-mode backchannel signaling system for differential high-speed links," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2004, pp. 352–355.

[8] J. L. Zerbe, P. S. Chau, C. W. Werner, T. P. Thrush, H. J. Liaw, B. W. Garlepp, and K. S. Donnelly, "1.6 Gb/s/pin 4-PAM signaling and circuits for a multidrop bus," *IEEE J. Solid-State Circuits*, vol. 35, no. 5, pp. 752–760, May 2001.

[9] Y.-S. Sohn, S.-J. Bae, H.-J. Park, C.-H. Kim, and S.-I. Cho, "A 2.2 Gbps CMOS look-ahead DFE receiver for multidrop channel with pin-to-pin time skew compensation," in *Proc. IEEE Custom Integrated Circuits Conf.*, Sep. 2003, pp. 473–476.

[10] B. Widrow, J. M. McCool, M. G. Larimore, and C. R. Johnson Jr., "Stationary and nonstationary learning characteristics of the LMS adaptive filter," *Proc. IEEE*, vol. 64, no. 8, pp. 1151–1162, Aug. 1976.

[11] V. Stojanović, G. Ginis, and M. A. Horowitz, "Transmit pre-emphasis for high-speed time-division-multiplexed serial link transceiver," in *Proc. IEEE Int. Conf. Communications*, May 2002, pp. 1934–1939.

[12] J. H. Winters and S. Kasturia, "Adaptive nonlinear cancellation for high-speed fiber-optic systems," *IEEE J. Lightwave Technol.*, vol. 10, no. 7, pp. 971–977, Jul. 1992.

Vladimir Stojanović (S'96–M'05) was born in Kragujevac, Serbia, Yugoslavia. He received the Dipl. Ing. degree in electrical engineering from the University of Belgrade, Serbia, in 1998, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2005, respectively.

He was also with Rambus, Inc., Los Altos, CA, from 2001 to 2004 and a Visiting Scholar with the Advanced Computer Systems Engineering Laboratory, Department of Electrical and Computer Engineering, University of California, Davis, in 1997-1998. He is currently an Assistant Professor in the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA. His current research interests include design, modeling, and optimization of integrated systems, from standard VLSI blocks to CMOS-based electrical and optical interfaces. He is also interested in design and implementation of digital communication techniques in high-speed interfaces and high-speed mixed-signal IC design.

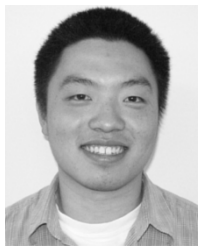

Andrew Ho (M'99) was born in Taipei, Taiwan, R.O.C., in 1979. He received the B.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2001 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2002.

In 2002, he joined Rambus Inc., Mountain View, CA, working in area of high-speed signaling and I/O design and test.

Bruno W. Garlepp (M'97) was born in Bahia, Brazil, in 1970. He received the B.S.E.E. degree from the University of California, Los Angeles, in 1993, and the M.S.E.E. degree from Stanford University, Stanford, CA, in 1995.

In 1993, he joined the Hughes Aircraft Advanced Circuits Technology Center, Torrance, CA, where he designed high-precision analog ICs for A/D applications and RF circuits for wideband communications applications. In 1996, he joined Rambus Inc., Mountain View, CA, where he designed high-speed CMOS clocking and I/O circuits for synchronous chip-to-chip interfaces. In 2000, he joined Silicon Laboratories, Austin, TX, where he designed high performance CDR and clock synthesis ICs for SONET applications. In 2003, he returned to Rambus, Los Altos, CA, where he designs multi-GHz, multi-mode signaling interfaces for chip-to-chip communications.

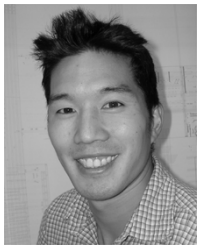

Fred Chen (M'00) was born in Wichita, KS, in 1975. He received the B.S. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 1997, and the M.S. degree in electrical engineering from the University of California at Berkeley in 2000.

In 1997 he joined Motorola, Libertyville, IL where he worked on discrete RF design for CDMA cell phones. In 2000, he joined Rambus, Inc., Mountain View, CA, where he has worked on the design of high-speed I/O and equalization circuits.

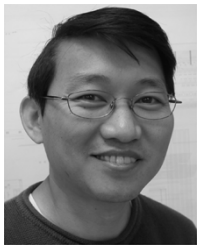

Jason Wei (M'00) was born in Taipei, Taiwan, R.O.C., and received the B.S. degree from National Cheng-Kung University, Tainan, Taiwan, in 1985 and the M.S. degree in electrical engineering from San Jose State University, San Jose, CA, in 1989.

From 1989 to 1994, he was with Raytheon and OKI Semiconductor working on emitter couple logic and analog/digital PLL. In 1994 he joined the technical staff at Rambus Inc., Mountain View, CA, where he has designed high-speed CMOS PLL circuits for clock recovery and data synchronization, and high-speed I/O circuits.

Grace Tsang (M'01) recceived the B.S.E.E. degree from the Massachusetts Institute of Technology, Cambridge, in 1984, and the M.S.E.E. degree from the California Institute of Technology, Pasadena, in 1991

She designed high-speed bipolar circuits at Tektronix until 1989. During 1991–1995, she worked at Western Digital on HDD read channel ICs. Since 1995, she has been with Rambus Inc. working on chip-chip interfaces and clock recovery.

Elad Alon (S'02) was born in Haifa, Israel, in 1979. He received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 2001 and 2002, respectively, and is currently working toward the Ph.D. degree at Stanford University.

He was a visiting researcher at Hewlett Packard Company, Ft. Collins, CO, in the summer of 2003, where he implemented supply noise measurement circuits on an Itanium processor. He has also been a visiting researcher at Rambus Inc., Los Altos, CA, since 2003, where he has worked on supply noise measurement and high-speed signaling circuits. His research interests include efficient on-chip power supply regulation and distribution, noise measurement techniques, high-speed interface design, and applications of optimization and nonlinear control to high-speed mixed-signal circuits.

Ravi T. Kollipara (M'88) is a Principal Engineer at Rambus Inc., Mountain View, CA, responsible for the signal integrity of the high-speed serial link channels. His responsibilities include design and development of models for packages, line cards, backplanes, connectors, traces and vias, and performing simulations for system level voltage and timing budgets and jitter characterization.

Carl W. Werner (M'97) was born in Chicago, IL, in 1962. He received the B.S. degree in electrical engineering from the University of Illinois at Urbana/Champaign in 1984.

He was with Siliconix Inc. from 1984 to 1988, and National Semiconductor from 1988 to 1997 where he worked on CMOS, bipolar and BiCMOS, analog and mixed-signal integrated circuit design. He holds several U.S. and foreign patents. In 1999 he joined the technical staff at Rambus Inc., Mountain View, CA, where he presently manages a team focused on high-speed circuit design and test.

Jared L. Zerbe (M'90) was born in New York City, NY, in 1965. He received the B.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1987.

In 1987 he joined VLSI Technology, Inc. where he worked on custom and semi-custom ASIC design. In 1989, he joined MIPS Computer Systems, where he designed high-performance CPU floating-point blocks. In 1992, he joined Rambus, Inc., Mountain View, CA, where he has since specialized in the design of high-speed I/O, PLL/DLL clock-recovery, and data-synchronization circuits. He has authored multiple papers and patents in the area of high-speed clocking and data transmission. He currently leads a design group focused on high-speed backplane serial links.

Mark A. Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founder's Professor of Electrical Engineering and Computer Science at Stanford University. His research area is in digital system design, and he has led a number of processor designs including MIPS-X, one of the first processors to include an on-chip instruction cache, TORCH, a statically scheduled, superscalar processor that supported speculative execution, and FLASH, a flexible DSM machine. He has also worked in a number of other chip design areas including high-speed and low-power memory design, high-bandwidth interfaces, and fast floating point. In 1990 he took leave from Stanford to help start Rambus Inc., Los Altos, CA, a company designing high-bandwidth chip interface technology. His current research includes multiprocessor design, low power circuits, memory design, and high-speed links.

Dr. Horowitz received the Presidential Young Investigator Award and an IBM Faculty development award in 1985. In 1993, he received the Best Paper Award at the IEEE International Solid-State Circuits Conference.