A 6.25-Gb/s Binary Transceiver in 0.13-μm CMOS for Serial Data Transmission Across High Loss Legacy Backplane Channels

Robert Payne, *Member, IEEE*, Paul Landman, *Member, IEEE*, Bhavesh Bhakta, *Member, IEEE*, Sridhar Ramaswamy, *Member, IEEE*, Song Wu, *Member, IEEE*, John D. Powers, *Member, IEEE*, M. Ulvi Erdogan, *Member, IEEE*, Ah-Lyan Yee, *Member, IEEE*, Richard Gu, Lin Wu, *Member, IEEE*, Yiqun Xie, *Member, IEEE*, Bharadwaj Parthasarathy, *Member, IEEE*, Keith Brouse, *Member, IEEE*, Wahed Mohammed, Keerthi Heragu, Vikas Gupta, *Member, IEEE*, Lisa Dyson, and Wai Lee, *Member, IEEE*

***Abstract*—A transceiver capable of 6.25-Gb/s data transmission across legacy communications equipment backplanes is described. To achieve a bit error rate (BER) $< 10^{-15}$, transmit and receive equalization that can compensate up to 20 dB of channel loss is employed to remove intersymbol interference (ISI) resulting from finite channel bandwidth and reflections. The transmit feed-forward equalizer (FFE) uses a four-tap symbol-spaced programmable finite impulse response (FIR) filter followed by a 4-bit digital-to-analog converter (DAC) that drives a 50-Ω transmission line. The receiver uses a half-baud-rate adaptive decision feedback equalizer (DFE) that cancels the first four symbol-spaced taps of postcursor ISI without use of speculative techniques. Both the transmitter and receiver use an LC-oscillator-based phase-locked loop (PLL) to provide low jitter clocks. Techniques to minimize the complexity of the FIR and DFE implementations are described. The transceiver is designed to be integrated in a standard ASIC flow in a 0.13-μm digital CMOS technology. System measurements indicate the ability to transmit and recover data eyes that have been fully closed due to crosstalk and signal loss.**

***Index Terms*—Adaptive equalizers, current-mode logic, data communications, decision feedback equalizers, serial links, transceivers.**

I. Introduction

EVER increasing line data rates and integration are driving the need for higher bandwidth backplane data transmission in communications equipment. Systems vendors are reluctant to deploy new backplanes due to the high development cost as well as a widely installed base of legacy systems. Meeting these increased bandwidth requirements over legacy backplane channels requires new transceiver circuit technology.

Many legacy backplanes were designed for either the Gigabit Ethernet or the 10 Gigabit Ethernet extended attachment unit interface (XAUI) standards, leading to data rates in the 1.25–3.125-Gb/s range. These backplane channels typically include over 30 in of copper trace on a flame resistant 4 (FR-4) dielectric, multiple connectors, and several plated through-hole vias [Fig. 1(a)]. Although the resulting signal integrity was sufficient at these data rates, the need for higher throughput and increased port density is pushing backplane rates to 6.25 Gb/s. At this high data rate, the channel nonidealities result in signal loss and reflections as well as significant high-frequency crosstalk. Fig. 1(b) illustrates measurements of a typical legacy backplane channel with approximately 20 dB of loss at 3.125 GHz and crosstalk energy that actually exceeds the signal energy at slightly higher frequencies.

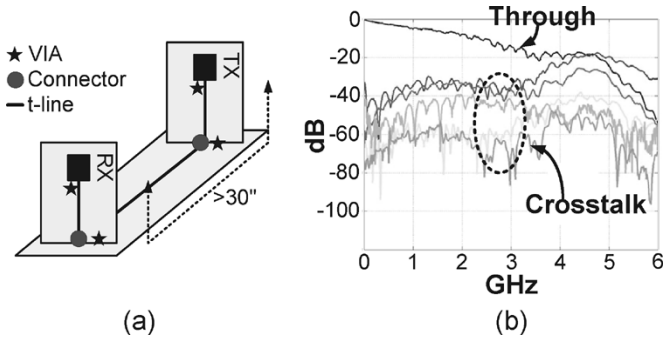

Fig. 1. Legacy backplane. (a) Physical channel. (b) Electrical characteristics.

Achieving a BER less than 10^{-15} across such band-limited channels requires equalization to flatten the channel response. Typical multi-gigabit-per-second equalizer solutions include transmitter preemphasis (or, more properly, deemphasis) to boost the ratio of high- to low-frequency signal energy sent from the transmitter [1], [2] or receiver feed-forward or linear equalizers that accomplish the same function in the receiver [3]. Multilevel signaling solutions such as four-level pulse amplitude modulation (PAM-4) can also be used to fully utilize the available bandwidth [4]–[7]. While effective for equalizing isolated channels, transmit deemphasis also proportionally increases high-frequency crosstalk, resulting in a decreased system signal-to-noise ratio (SNR). In addition, without a back channel, transmit equalization is typically not adaptive, resulting in suboptimal performance in time-varying channels where loss is dependent on factors such as temperature and humidity [8]. Handling these variations requires an adaptive equalizer approach, which is most readily accomplished in the receiver. Linear receive filters can flatten the channel response, but they do not discriminate crosstalk from the desired signal. As a result, the high-frequency SNR is unchanged, rendering these unsuitable for legacy backplanes where significant

Manuscript received April 7, 2005; revised July 25, 2005.

The authors are with Texas Instruments Incorporated, Dallas, TX 75243 USA (e-mail: b-payne@ti.com; landman@ti.com; lee@ti.com).

Digital Object Identifier 10.1109/JSSC.2005.856583

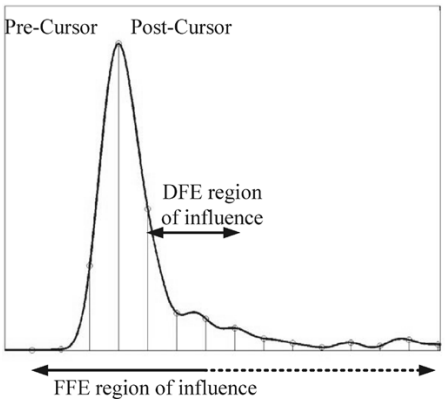

Fig. 2. Channel pulse response illustrating DFE and FFE regions of influence.

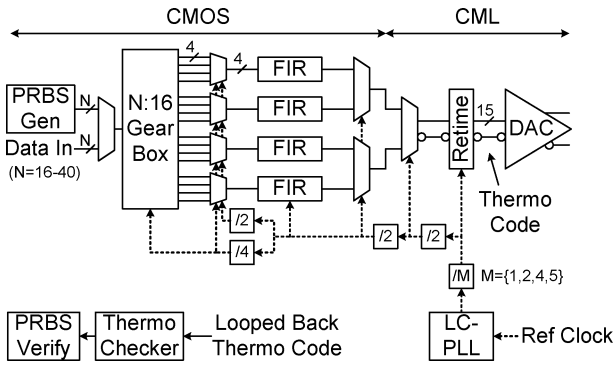

Fig. 3. Transmitter block diagram.

high-frequency crosstalk exists. Finally, although PAM-4 signaling reduces the required bandwidth, no signaling standards exist and the impact of reflections and crosstalk on PAM-4 solutions can be more severe compared to binary signaling, as suggested in [5]. Overcoming these limitations requires a combination of nonlinear adaptive receive equalization and increased transmit equalizer performance.

A DFE is well suited to the receive equalizer function as the slicer nonlinearity allows it to amplify the recovered signal while rejecting noise [9]. At multi-gigabit-per-second rates, the primary challenge is feeding back the decisions quickly enough to implement the first filter tap. Due to speed limitations, most multi-gigabit-per-second DFEs have employed speculative or loop-unfolding techniques [5], [6], [10]. These approaches relax the timing requirements of the first tap feedback by precomputing the equalized eye for either prior input data polarity, sampling both results, and choosing the proper result once the previous bit decision is known. Unfortunately, this also introduces unwanted loading in the critical signal and clock paths as well as complicates the associated clock and data recovery (CDR) circuit design [6]. This paper presents a DFE at 6.25 Gb/s where the first tap is fed back directly from the input slicers without speculation. The direct DFE architecture avoids the loading and CDR issues and provides a single straightforward equalization approach for all taps. In contrast, speculative methods require either the addition of a standard DFE for later taps, a combination of DFE and linear or feed-forward equalization, or an exponential complexity increase if speculation is applied to later taps.

However, a DFE alone is not the complete solution since it cannot cancel precursor intersymbol interference (ISI), and its postcursor cancellation capabilities are limited by filter length (see Fig. 2). Therefore, additional equalization is needed for optimum system performance over a variety of channels. Adding a feed-forward equalizer (FFE) to the transmitter addresses these concerns. Specifically, this design incorporates a programmable four-tap transmit finite impulse response (FIR) filter capable of canceling ISI beyond the region of influence of the DFE.

The remainder of the paper is organized as follows. Section II describes the transmit architecture including the FFE and output driver digital-to-analog converter (DAC). Section III covers the receiver and DFE architectures. Section IV presents the design for testability (DFT) techniques that enable high volume manufacturing of the transceiver circuits. Section V details the performance measurements of the transceiver circuits. Finally, Section VI concludes the paper.

II. Transmitter Design

A. Architecture

A block diagram of the transmitter is shown in Fig. 3, where a 4-bit digital FIR filter combined with a 4-bit DAC realizes the FFE function. An alternative implementation employs analog filters with parallel drivers that sums current contributions from each tap. This leads to large capacitive overheads, since, in its simplest form, the analog approach requires a separate driver for each tap. If each tap driver is capable of driving a full-scale output, the capacitive overhead factor can be as high as the number of filter taps beyond the cursor tap. In practice, designers restrict the supported filtering functions and share portions of the tap drivers to reduce this overhead. Even then, however, a capacitive overhead of 50% or more is common [5].

The digital approach eliminates the parasitic overhead of the analog approach and provides for highly flexible filter characteristics. The reduced capacitance at the driver output increases the output bandwidth and optimizes the return loss for minimal reflections. In particular, the measured differential return loss for this transmitter was less than -12 dB from 10 MHz to 3.125 GHz, and the common mode return loss was less than -12.9 dB over the same frequency range. This compares favorably to the requirements of the Optical Internetworking Forum (OIF) CEI-6G-LR specifications [11]. Moreover, the minimized parasitic capacitance helps the transmitter achieve a large-signal bandwidth of 6.3 GHz without inductive peaking, yielding a compact layout. This bandwidth is roughly twice the Nyquist frequency and is desirable to minimize ISI when transmitting data streams with broadband frequency content.

B. DAC Implementation

To minimize glitches and mismatch-induced timing jitter, the 4-bit DAC is implemented as a fully segmented array of 15 identical current-mode logic (CML) drivers, as shown in Fig. 4. Since each of the DAC segment drivers need only be capable of contributing a single least significant bit (LSB) of current (or 1/15th of the full-scale output), the output capacitance for all 15 DAC segments is equivalent to that of a single tap driver in the analog filtering approach.

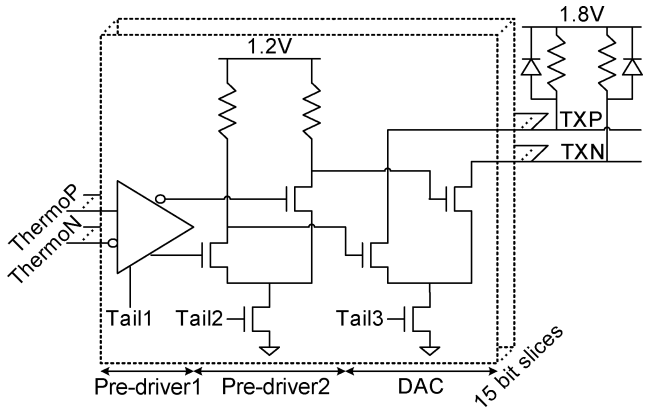

Fig. 4. Fully segmented Tx DAC circuit.

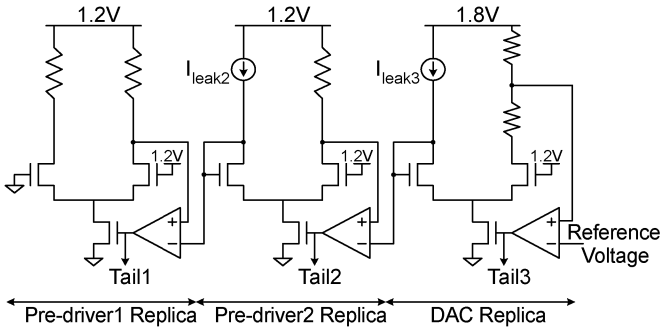

Fig. 5. Tx output driver and predriver V_{T} referenced biasing.

Linearity is a key specification for any DAC design. A differential nonlinearity (DNL) of 0.2 LSB and integral nonlinearity (INL) of 0.3 LSB were achieved by using passive loads and a 1.8-V termination voltage, which was necessary to keep the output transistors fully saturated during capacitively-coupled operation at the maximum output swing. The replica-bias circuit on the right side of Fig. 5 controls the output swing and can be programmed to output from 700 to 1200 mV peak-to-peak (p–p) differential by altering the reference voltage.

The DAC uses thin-oxide devices to maximize the performance. To address channel hot carrier (CHC) concerns arising from the 1.8-V termination voltage, channel lengths of greater than twice the allowed minimum were used. More detailed CHC simulations later revealed that using low-Vt devices raised the common-source voltage of the switching pair sufficiently to alleviate any CHC concerns for even minimum channel length devices. This fact was used to halve both the width and length of the switching devices in a later power-optimized design. This reduced the DAC input capacitance by a factor of 4 and cut the total transmit power in half.

The input to each DAC is driven by a resistor-loaded CML predriver with sufficient bandwidth such that no reactive loading component is required. As Fig. 4 illustrates, multiple predriver stages are required to drive the DAC. The later power-optimized design reduced the DAC input capacitance, allowing one of the predriver stages to be eliminated. Lower power CMOS predrivers were also considered but were eliminated due to the jitter introduced by their delay sensitivity to power supply variations.

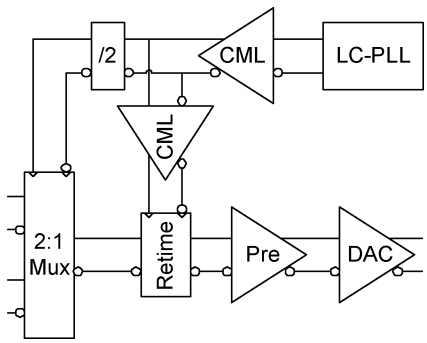

Fig. 6. Full-baud-rate transmitter output retiming.

The predrivers are biased using the threshold referenced biasing scheme of Fig. 5. The circuit sets the predriver amplitude just large enough to steer ~99% of the DAC tail current through the "on" side of the differential pair, leaving the "off" side device still saturated but approaching the onset of the cutoff region. This scenario results in optimal output waveforms with symmetrical rising and falling edges and minimal common-mode noise. To accomplish this, the "off" side device in the DAC replica is diode-connected, and a leakage current $\mathrm{I_{leak3}}$ of roughly 1% of the DAC tail current is applied. The resulting "off" side gate voltage is then approximately an nMOS threshold voltage above the common-source voltage, which is precisely what is required to achieve the desired current steering. A replica of predriver no. 2 in a control loop is used to generate the signal Tail2, which is used to bias all predriver segments. A similar scheme is used to control the swing of the first predriver stage. Controlling the CML swings in this manner keeps all devices saturated and enables the DAC to meet the desired swing, linearity, and signal integrity requirements over process, voltage, and temperature variations.

C. Clock Generation and Distribution

To eliminate deterministic jitter induced by duty cycle distortion in the transmit clock, the transmitter employs a full-baud retimer implemented as a CML master–slave flip-flop (see Fig. 6). The bit-rate clock required by the retimer is generated using an LC-VCO-based phase-locked loop (PLL) and a CML clock tree to minimize susceptibility to supply and substrate noise. CML has the added advantage that it draws a constant current and, therefore, minimizes noise injection onto the supplies. For these reasons, CML was used for all circuits that influence the transmitter jitter performance. Circuits prior to the final 2:1 mux do not impact jitter performance and were built with CMOS logic powered from a separate 1.2-V supply. A programmable divide by M circuit, not shown here, with $M = \{1, 2, 4, 5\}$ provides backward compatibility to lower speed standards such as Gigabit Ethernet or XAUI.

D. FFE Implementation

The FFE is implemented using four parallel FIR filters operating at 1/4 line rate with outputs serialized by a tree-type multiplexer. The 4-bit DAC allows 5-bit filter coefficients (one sign and four magnitude) that can equalize any combination of

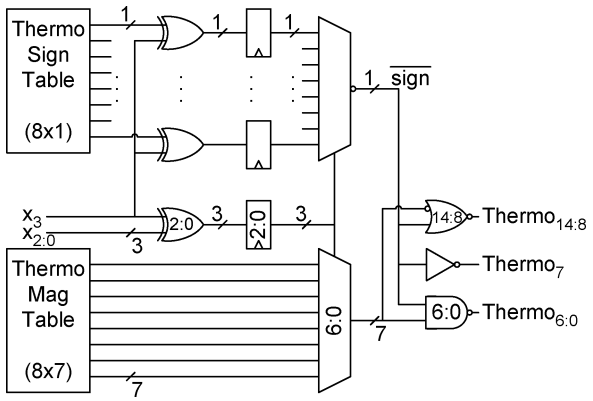

Fig. 7. Optimized look-up table implementation of Tx FFE.

precursor and postcursor taps to match the channel characteristics. A look-up table implementation (shown in Fig. 7) enables the design to operate at 1.5625 GHz. Four consecutive input bits form a table address that selects a precomputed filter output. With 15-bit thermometer-coded outputs, a direct implementation would require a $16 \times 15 = 240$ bit table. The table size is halved by recognizing that $y(x_3, x_2, x_1, x_0) = -y(-x_3, -x_2, -x_1, -x_0)$ where $x_i = x(n-i) \in \{-1, +1\}$ is a delayed input sample and y is the filter output. This allows storage of only the $y(-1, x_2, x_1, x_0)$ half of the table, while the other half is recovered by sign correcting the address bits and filter output based on x_3 [12]. The table size is nearly halved again by storing the filter output as one sign bit and seven magnitude bits from which the full 15-bit thermometer code is reconstructed. This leads to a manageable table size of $8 \times (1+7) = 64$ bits. While further compression is possible, it would make the translation to a thermometer code difficult at 1.5625 GHz.

Selecting the optimal filter coefficients for a particular channel can be difficult. The high degree of flexibility afforded by the digital filtering approach adds to this challenge. Although system models can help select the coefficients, optimizing to the actual channel including the transmitter and its package is superior. To address this, the transmitter incorporates a calibration mode that sends full-rate random codes to the DAC. This allows broadband characterization of the channel transfer function including the DAC, package, and backplane. The receiver package and input model (generated from return loss measurements) is added to this, completing the signal path. Optimal filter coefficients can then be derived mathematically from this end-to-end transfer function.

III. Receiver Design

A. Architecture and DFE Implementation

Figs. 8 and 9 illustrate the half-baud-rate receiver and direct feedback DFE. An LC-VCO-based PLL generates quadrature phases of a 6.25-GHz clock. A digital phase interpolator and CDR circuit sum weighted copies of these phases to generate data slicing clocks nominally centered in the recovered data eye. Following the interpolator, a programmable divider generates baud-rate clocks to support the native 6.25-Gb/s and legacy 3.125- and 1.25-Gb/s rates without requiring a

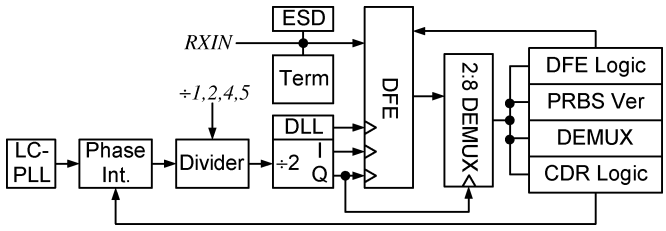

Fig. 8. Receiver block diagram.

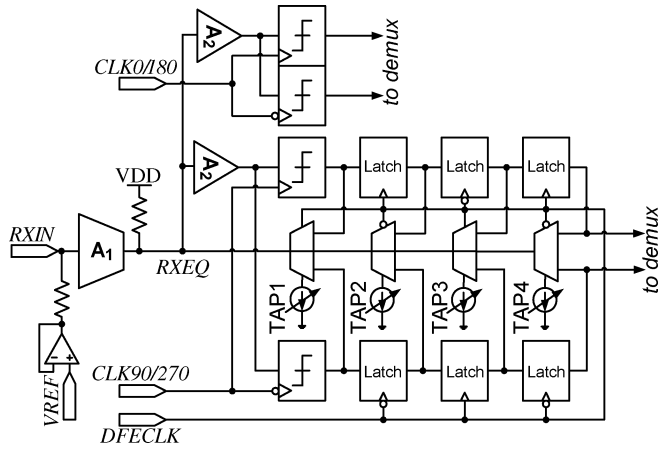

Fig. 9. Receiver analog front end and DFE.

wide tuning range PLL. Operating the interpolator at a fixed frequency simplifies its design and optimizes linearity at the critical highest frequency. A subsequent divide-by-two results in in-phase (CLK0 and CLK180) and quadrature (CLK90 and CLK270) half-baud-rate clocks that strobe the equalized data eye at each transition and eye center. An additional clock pair, DFECLKP/N, is also generated to retime the feedback of the decisions in the equalizer.

In Fig. 9 (shown single-ended for simplicity), the ac-coupled receiver input, RXIN, is terminated to a regulated common-mode point through an on-chip 50-Ω resistor. The low-impedance termination voltage provides better common-mode return loss over a wider frequency than a simple voltage divider solution and helps to reduce common-mode to differential conversion. Amplifier A_1 consists of a resistor loaded differential transconductor that buffers the incoming data prior to equalization. In addition to providing $\sim$6 dB gain, A_1 also isolates the equalized signal, RXEQ, from the channel and minimizes the parasitic capacitance at RXIN to improve the input return loss and bandwidth. The output current from transconductor A_1 and the weighted decision feedback currents are summed into a resistive load, producing the equalized signal RXEQ.

The DFE tap multiplication is provided by current-mode DACs whose LSB currents are referenced to the gain of transconductor A_1 to provide process-, voltage-, and temperature-independent input-referred ISI cancellation. Each LSB of current multiplied by the load resistor corresponds to approximately 5 mV of input-referred ISI. Referring the tap currents to the input gain reduces the range requirement for the DACs and the need for complex gain control of A_1. Statistical modeling [13] indicated a need for 5 bits of range for the first tap, 4 bits

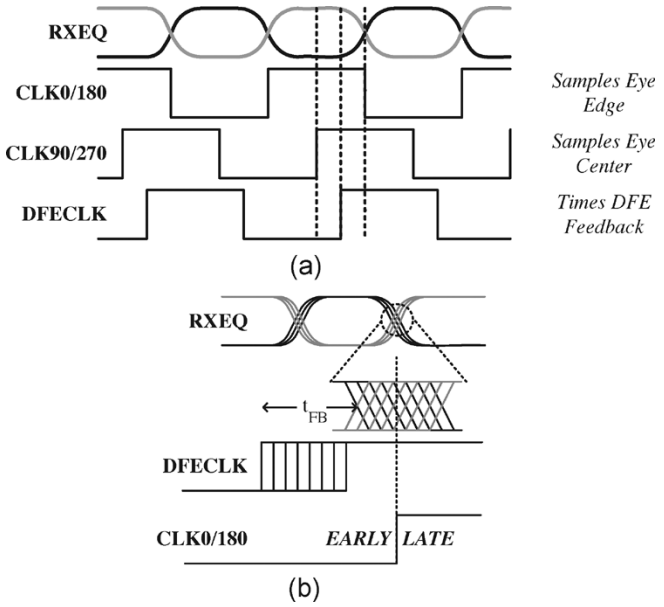

Fig. 10. (a) Receiver timing diagrams. (b) DFE feedback clock timing adjustment.

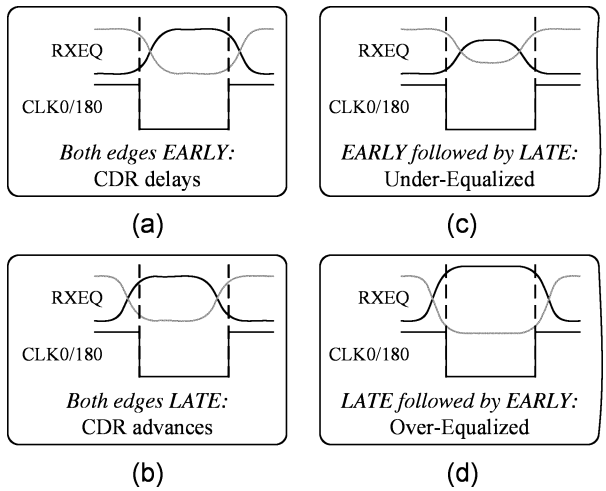

Fig. 11. Transition sampling based CDR and DFE update. (a) and (b) CDR update. (c) and (d) DFE tap coefficients update.

plus sign for the second, and 3 bits plus sign for the third and fourth taps.

The timing of the feedback is controlled by an additional clock, DFECLK, which alternately latches and selects the outputs of the two recovered half-baud-rate slicer outputs. The DFECLK delay with respect to CLK0 and CLK180, the transition sampling clocks, is set using a programmable delay stage using a technique similar to that in [5]. After initial reset, the delay is adjusted to set the crossing point of the fed back differential data to be coincident with the desired crossing point of the equalized data (Fig. 10). During the delay calibration period, input buffer A_1 is disabled and the DFE tap coefficients are fixed to $\mathrm{TAP}\{1,2,3,4\} = \{n,0,0,0\}$, which generates a repeating $1010\ldots$ pattern at the sense amplifier inputs. By sampling the resulting "eye" using CLK0 and CLK180, it is determined whether the fed back data is early or late with respect to the edge sampling clocks and the ideal position is found using a linear search technique. This delay control is required by the DFE tap coefficient update algorithm that is described in the next section. Although subsequent shifts in operating voltage and temperature can impact the optimum delay, these small changes (since process variations are already cancelled) have little simulated or measured impact on the tap coefficient convergence.

Before the input slicers, an additional pair of amplifiers, labeled A_2, is included. These amplifiers limit the equalized signal, improve the overall bandwidth by distributing the gain, aid in the common-mode rejection of the sense amplifiers, and isolate any charge kickback generated among the sense amplifiers.

B. DFE Tap Adaptation

A common approach to DFE tap adaptation measures the equalized eye height and minimizes the error between the actual and expected eye heights. Unfortunately, this requires additional hardware to accurately sense the eye height at high data rates. This overhead is avoided by reusing existing circuits already required by the CDR and exploits the property that minimizing the differential crossing point jitter also maximizes the center eye height [13]. The binary CDR already oversamples the receiver input data at the symbol center and edge. This provides information about both the eye width and recovered clock position that is used to update both the CDR and DFE tap coefficients.

Fig. 11 illustrates the CDR and DFE adaptation criteria. In (a) and (b), the two cases where both transition sampling clocks are either early or late with respect to the eye edges are illustrated. If both are early, the CDR delays the output phase of all clocks. If both are late, the CDR advances the output phase of all clocks. In (c) and (d), the cases for the DFE tap coefficient adaptation are illustrated. If the leading edge is early while the trailing edge is late, then the eye width is too narrow and the eye underequalized. If the leading edge is late while the trailing edge is early, the eye is too wide and is overequalized. This provides error information to a sign–sign least mean square (LMS) algorithm that uses this and the prior data decisions to determine which previous bits contributed to the residual ISI. The tap weights are then updated according to the following equation: $\mathrm{dfe}_k^{j+1} = \mathrm{dfe}_k^j - \lambda\,\mathrm{sgn}\{z_{m-1/2}(a_{m-k}+a_{m-k-1})\}$, where dfe_k^j is the kth DFE tap at time j, λ is a programmable variable setting the magnitude of the tap updates, $z_{m-1/2}$ is the transition slicer decision leading the current data decision a_m, and a_{m-k} and a_{m-k-1} are taken from the history of data decisions. The coefficients are then digitally low pass filtered to reduce the impact of noise and jitter on the tap updates and reduce any interaction between convergence of the DFE coefficients and the relatively high bandwidth CDR loop. A more comprehensive description of the DFE tap adaptation can be found in [13].

C. DFE Critical Timing Path

Fig. 12 illustrates the DFE critical timing path. Since the DFE adapts based on the equalized signal transitions, when properly equalized, their differential crossings should be coincident with the transition sampling clock edges. This requires control of the timing of DFECLK since improperly equalized edges lead to

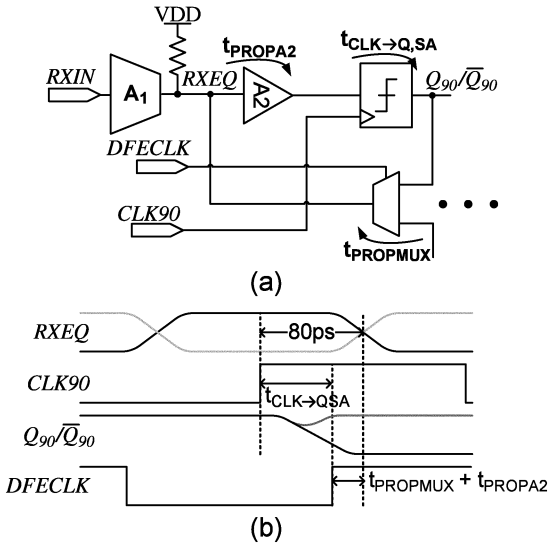

Fig. 12. DFE critical path timing. (a) Delay contributors. (b) Timing diagram.

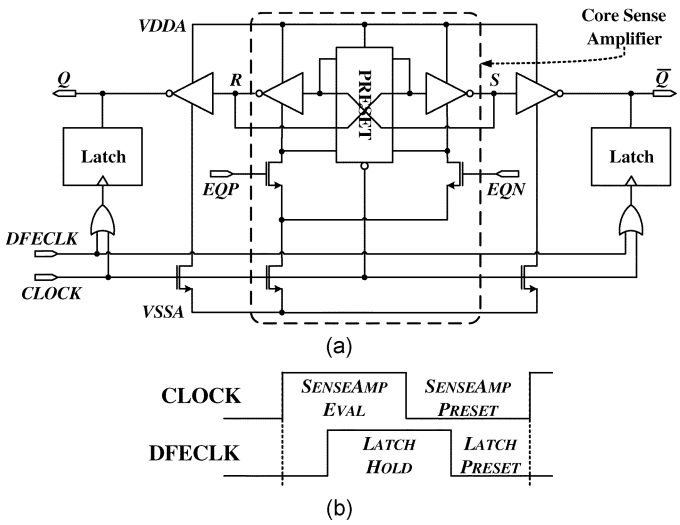

Fig. 13. Sense amplifier. (a) Schematic. (b) Timing diagram.

suboptimal DFE convergence and a higher BER. To properly converge, the feedback must settle to 50% of its final value half of a unit interval (UI) after the sample, or 80 ps in the case of a 6.25-Gb/s data stream. If not, the DFE tap update engine may falsely sense that the eye is underequalized, resulting in suboptimal convergence. This is most difficult for the first tap since the sense amplifiers must resolve input signal swings as small as 20 mV peak differential within this time window. The 80 ps critical path timing includes the propagation delays of amplifier A_2 and the tap feedback mux as well as the resolution time of the sense amplifier $(t_{\mathrm{CLK}\rightarrow Q,\mathrm{SA}})$. Therefore, the maximum $t_{\mathrm{CLK}\rightarrow Q,\mathrm{SA}}$ delay must be less than 80 ps minus the propagation delays of A_2 and the tap feedback mux.

Since $t_{\mathrm{CLK}\rightarrow Q,\mathrm{SA}}$ dominates the timing and is a strong function of the sense amplifier input amplitude, adding amplifier A_2 actually improves the system timing margins. For very small input amplitudes and possibly varying common-modes (due to crosstalk noise and different DFE tap magnitudes), the regenerative gain of the sense amplifier consumes significant time. Although A_2 adds delay in the critical timing path, the reduction in $t_{\mathrm{CLK}\rightarrow Q,\mathrm{SA}}$ due to its gain and better controlled common-mode at the sense amplifier input more than compensates.

D. Sense Amplifier Design

To accommodate the DFE timing, the sense amplifier and the first tap feedback latch are combined (Fig. 13). Most high-speed sense amplifier designs [14]–[16] utilize a core sense amplifier that generates a pulse according to the input data polarity followed by a set–reset (SR) latch to capture the result. However, even in an optimized design such as [15], the added latch delay is still too great to meet the critical timing path.

To satisfy the speed and latch timing requirements of the DFE, the SR outputs of the sense amplifier are buffered by a pair of clocked inverters and parallel hold latches (Fig. 13) and directly processed in the DFE. Hysteresis is minimized by precharging and shorting all internal differential nodes of the sense amplifier and secondary latches. This also reduces the impact of device mismatch on the input offset.

During the precharge state, the clocked inverters isolate the sense amplifier from the output latches and allow a full UI of precharge time, minimizing hysteresis while using modest device sizes. The inverters also reduce the load seen by the core sense amplifier, provide the drive strength needed to charge and discharge the large feedback mux capacitance, and distribute the gain to minimize the overall delay of the sense amplifier. The parallel latches hold the decision until it is no longer needed in the feedback loop and are reset using a combination of the sampling clock and DFECLK.

Optimizing the critical timing path circuits required extensive statistical simulation. High-level system modeling indicated that a 20-mV peak (input referred) differential sensitivity was required to meet a 10^{-15} BER in a worst case legacy backplane channel. In this implementation, the primary determinant of the receiver sensitivity is the resolution time of the sense amplifier. The impact of both device mismatch and process variations was analyzed using Monte Carlo analysis. Fig. 14(a) plots the fraction of Monte Carlo simulations meeting a 60-ps resolution time as a function of the sense amplifier differential and common-mode input voltages. Fig. 14(b) plots the input-referred receiver sensitivity based on statistical simulations of the entire analog front end. In this case, the minimum input swing required to meet the critical 80 ps feedback path timing was found for each statistical model set. Based on this form of analysis, it was shown that the 20-mV sensitivity requirements could be met in volume production.

E. Tap Feedback Mux Design

The tap feedback muxes from Fig. 9 are expanded in Fig. 15. A pair of current-mode DACs controls the tail current to set each dfe_k coefficient weight. Splitting the feedback into two separate CML muxes with the unselected current being shunted to VDD eased the layout of the critical equalized node. It also improved the signal integrity of the fed back signals by further isolating the decisions between the two data sense amplifiers.

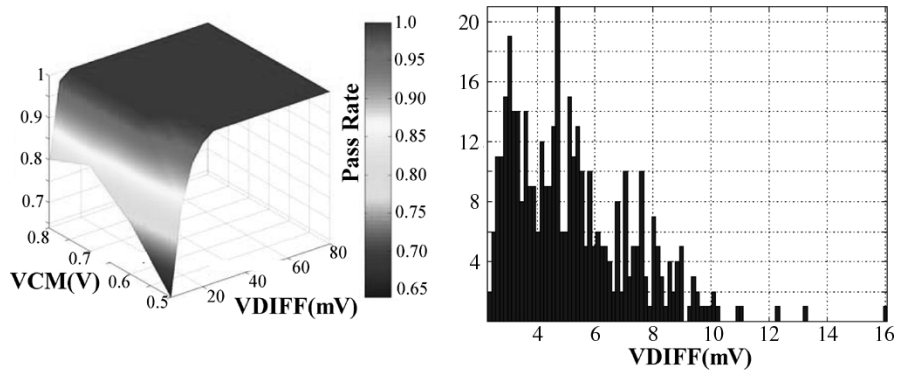

Fig. 14. (a) Statistical simulation results of sense amplifier $V_{\rm DIFF}$ and $V_{\rm CM}$ needed to meet $t_{\rm CLK\rightarrow Q,SA} < 60$ ps. (b) $V_{\rm DIFF}$ at bumps needed to satisfy DFE timing.

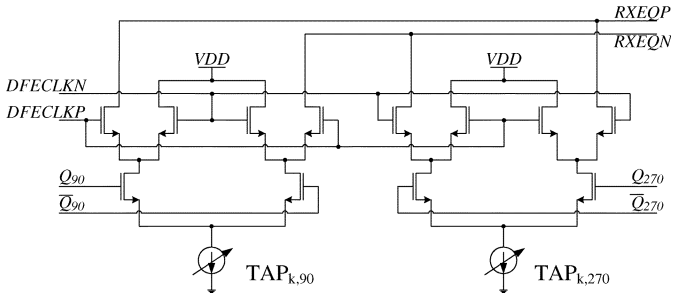

Fig. 15. DFE tap feedback mux schematic.

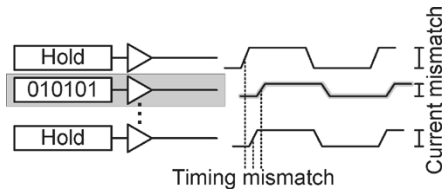

Fig. 16. Tx thermometer code toggle test mode.

F. Clock Generation and Distribution

Although CML clocks are well-suited to high data rate transceivers due to their supply noise immunity, sense amplifiers require full-swing clocks to achieve low-resolution times. As a result, the differential clocks from the phase interpolators and the dividers are converted to rail-to-rail voltage levels for the receiver sense amplifiers. Each clock has a duty cycle correction loop to minimize the impact of the half-baud-rate architecture.

IV. Test and Calibration Features

A. Transmitter Test and Characterization Features

The transmitter incorporates a pattern generator that can be used for built-in self test (BIST) in a production environment and bench testing. The generator supports clock patterns and pseudorandom bit sequences (PRBS) of length 2^7-1 and $2^{23}-1$. In addition to these traditional transceiver test features, the presence of a DAC in the transmitter mandates additional test and characterization capabilities.

For example, the transmitter features a DAC ramp mode for characterizing INL, DNL, glitch area, and settling time. In this mode, the pattern generator presents the DAC with triangle waves of alternately increasing and decreasing DAC codes. In another mode, the pattern generator sends random codes to the DAC. This feature is useful for gauging linearity and for characterizing the glitch impulse area and settling time for nonadjacent DAC codes.

The transmitter also offers a direct-access mode in which the FIR filters are bypassed allowing externally generated patterns to be fed straight to the DAC using 32 bits of the parallel transmit data interface. This mode is useful for applying patterns and code sequences not supported by the pattern generator such as the sinusoidal waveforms needed for single-tone and multitone intermodulation distortion (IMD) tests. It also provides support for spurious-free dynamic range (SFDR) and signal-to-noise-and-distortion (SINAD) measurements.

A DAC "thermo-toggle" mode is also available to allow characterization of timing and current mismatch between the DAC segments. In this mode, as illustrated in Fig. 16, all but one of the thermometer-coded DAC segment inputs are held at a constant value, while the input to the remaining segment is toggled, producing a clock pattern with nominal amplitude of one LSB. By repeating this procedure for each of the DAC segments and comparing the edge positions and amplitudes of the resulting waveforms, timing and current mismatch statistics can be gathered for each of the DAC segments.

B. Receiver Built-In Sensitivity Measurements

In addition to pattern verifiers that complement the transmit PRBS generators, more thorough characterization and verification of the receiver analog front end is needed in a production environment. Guaranteeing a 10^{-15} BER means verifying that the 20-mV sensitivity requirement is met for all devices. This sensitivity depends not only on the input offset, which is dominated by amplifiers A_1 and A_2 in Fig. 9, but also on the critical path timing at 6.25 Gb/s. While low-cost automated test hardware can test for offsets, they generally cannot generate small-swing high-speed data to exercise the critical timing paths and test at-speed sensitivity.

One solution exploits the DFE feedback loop to generate at-speed data patterns of varying amplitude, detect a pattern signature on the slow-speed parallel interface, and report the results to the test program. To turn the DFE into a pattern

TABLE I
SUMMARY OF PERFORMANCE OF COMPLETE TRANSCEIVER

Technology	0.13µm CMOS with 7 metal layers	
Supply Voltage	1.2V for Tx, Rx, and PLL 1.8V Tx DAC output termination	
Die size	3.0mm × 7.9mm	
Package	361-pin organic flip-chip BGA	
Tx/Rx channel area	0.24mm^2	
Mux path	16-3.125Gbps Rx → 8-6.250Gbps Tx	
Demux path	8-6.250Gbps Rx → 16-3.125Gbps Tx	
Power (chip)	7W	
PLL Technology	Quadrature LC-VCO	
PLL Jitter	<0.6ps rms (10kHz-1GHz) (62.5MHz reference clock)	
ESD protection	2kV HBM	
Tx Summary	Power	438mW (221mW for power-optimized version)
	Output swing	700mV→1200mV p-p differential
	Equalizer	4 tap FIR with 5 bit DAC @ 6.25GHz
	Near-end jitter	16ps p-p, 2^{31}-1 PRBS, 6.25 Gbps
	Far-end jitter	55ps p-p (for 36" FR-4 + 2 connectors)
	Return loss	<-12 dB from DC to 3.125GHz
Rx Summary	Power	210mW
	Sensitivity	<10mV p-p differential
	Equalizer	DFE, {5,4,3,3} bit DAC for taps{1,2,3,4}
	EQ Range	Up to 20 dB loss at 3.125GHz
	BER	<10^{-15} with worst-case NEXT
	Return loss	<-13 dB from DC to 3.125GHz

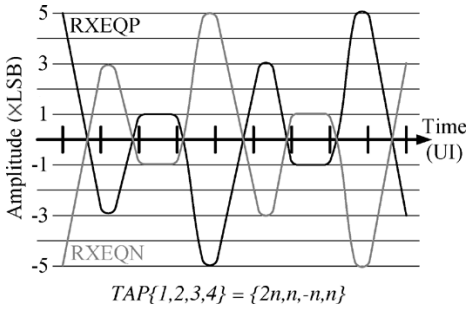

Fig. 17. DFE sensitivity test mode waveforms.

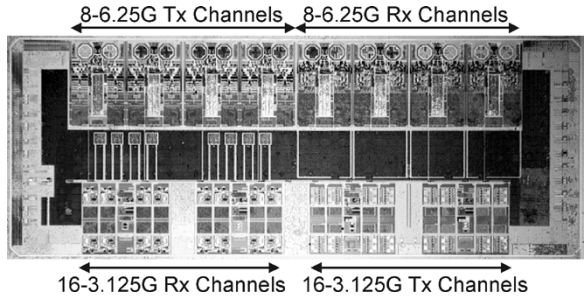

Fig. 18. Die photograph of complete mux/demux test chip.

generator, the receiver inputs are shorted to the common mode and the DFE tap coefficients are programmed to a fixed ratio. This places the DFE in a mode where it generates a known self-propagating data pattern. For example, setting $\mathrm{TAP}\{1,2,3,4\} = n \times \{2,1,-1,1\}$ results in a repeating 01101001 pattern at the sense amplifier inputs with an amplitude that varies between $n \times \mathrm{LSB}$ to $5 \times n \times \mathrm{LSB}$, as illustrated in Fig. 17. Next, n is decreased until a failure is detected on the demultiplexed data by a low-speed pattern verifier. The value of n indicates the resulting sensitivity in units of the DFE tap LSB for each receiver.

V. EXPERIMENTAL RESULTS

To validate transceiver performance, a mux/demux chip was built that integrated eight 6.25 Gb/s Tx/Rx pairs and sixteen 3.125 Gb/s Tx/Rx pairs. Fig. 18 shows a micrograph of the die. The 3.0 × 7.9 mm chip was fabricated in a 1.2-V 0.13-μm CMOS technology with seven metal layers and dissipates 7 W in a 361-pin organic flip-chip ball grid array (BGA) package. Table I summarizes the performance of the complete transceiver.

The PLL phase noise, measured using a spectrum analyzer with the Tx sending a 3.125-GHz clock pattern (1010...), was less than 0.6 ps root mean square (rms) integrated over a band of offset frequencies extending from 10 kHz to 1 GHz. The measurement was taken with a 100× PLL multiplication factor with a reference clock frequency of 62.5 MHz.

Fig. 19 shows the near-end eye diagram for a $2^{31}-1$ PRBS pattern superimposed on the OIF CEI-6G-LR mask [11]. The measured near-end total jitter was 16 ps p–p, including 8 ps of deterministic jitter. Fig. 20 shows the far-end eye diagram demonstrating the impact of the FFE on a legacy backplane channel consisting of 36 in of FR-4 with two connectors that exhibited a loss of 21.3 dB at the Nyquist frequency. The far-end jitter measured 55 ps p–p for this channel. Without the FFE, the far-end eye is completely closed.

Fig. 21(a) illustrates the setup used to duplicate the legacy backplane environment and test the capabilities of the receive equalizer. The output of a transmitter was connected to a worst case [in terms of loss and near-end crosstalk (NEXT)] system channel through a daughter card. The transmitter sent a $2^{31}-1$ PRBS pattern with a 1200-mV p–p differential output swing. All transmit equalization was disabled to determine the performance of the adaptive Rx equalizer alone. A bit error rate tester (BERT) pattern generator acting as a NEXT aggressor generated

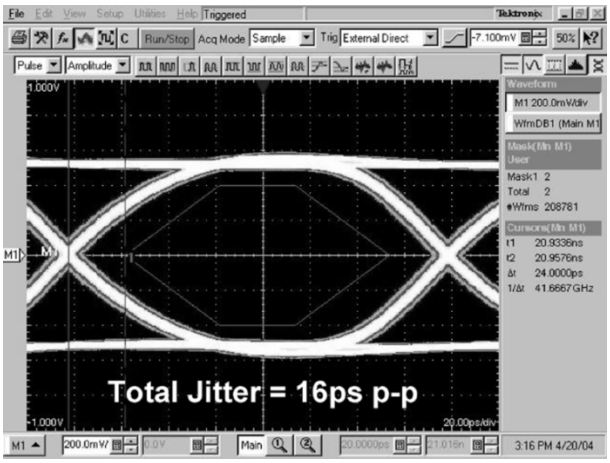

Fig. 19. Tx near-end eye diagram with CEI-6G-LR eye mask at 6.25 Gb/s.

Fig. 20. Tx eye diagram after 36 in of FR-4 and two connectors at 6.25 Gb/s.

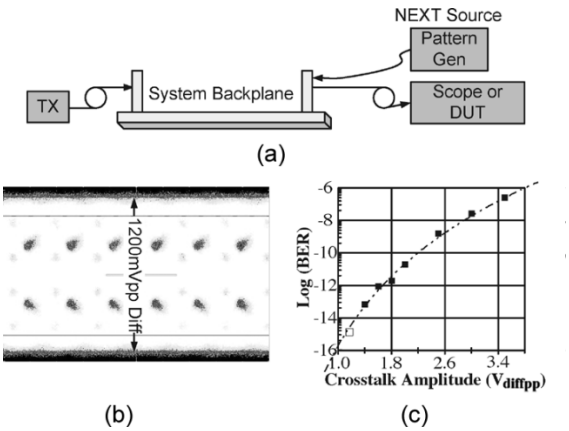

Fig. 21. (a) Rx equalizer test setup. (b) Received eye diagram at 6.25 Gb/s. (c) Impact of NEXT amplitude.

a 1200-mV p–p differential $2^{23}-1$ PRBS pattern at a slight frequency offset from the channel under test into the worst case NEXT channel. Due to the extremely sharp edges generated by the BERT transmitter, high-frequency crosstalk is greater than that in a real system. Fig. 21(b) plots the resulting fully closed eye at the receiver input. Under these conditions, a BER of less than 10^{-15} has been demonstrated, with the receiver adapting the CDR and DFE tap coefficients to the channel.

Next, the amplitude of the crosstalk aggressor was varied to judge the impact of increased crosstalk on the measured BER. Fig 21(c) plots the result, where no errors were seen for crosstalk amplitudes of 1200 mV p–p differential and below.

The receiver sensitivity BIST described in the previous section was used to make over 33 000 measurements of the Rx sensitivity as process, voltage, and temperature were varied. The results are plotted in Fig. 22, where more than 95% of the measurements indicate 10-mV peak input-referred differential sensitivity or better. In actual production test, limits would be set to guarantee that all shipped devices exceed the 20-mV required sensitivity.

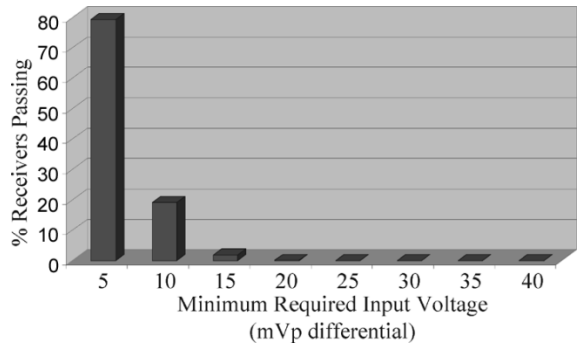

Fig. 22. Histogram of Rx built-in sensitivity measurement results.

VI. Conclusion

This paper presented the design and measurement results of a transceiver solution designed in a standard 0.13-μm CMOS process capable of communicating data at a 6.25-Gb/s data rate across legacy system backplanes originally designed for 1.25-Gb/s rates. The transmit equalizer used a 4-bit FIR filter and a fully segmented 4-bit DAC that both maximized the FFE tuning flexibility and minimized the output capacitance. A BER of less than 10^{-15} was achieved using a DFE capable of equalizing up to 20 dB of channel loss without amplifying the dominant high-frequency crosstalk noise. The receiver DFE operated from half-baud-rate recovered clocks and was capable of directly correcting the first symbol-spaced ISI tap without speculation. The receiver also used a low-overhead adaptation technique for the equalizer tap coefficients based on eye transition sampling. The transmitter and receiver both included BIST and measurement circuitry to aid in characterization and economical production testing.

Acknowledgment

The authors would like to thank B. Rothbauer, B. Dahl, A. Bhandal, and B. Barrie for their support in evaluation, modeling, and test system design. They also want to acknowledge J. Milton and the TI Internet Infrastructure design team for their contribution of the 3.125-Gb/s transceivers used in the mux/demux test chip.

REFERENCES

[1] A. Fiedler, R. Mactaggart, J. Welch, and S. Krishnan, "A 1.0625 Gbps transceiver with 2×-oversampling and transmit signal pre-emphasis," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 1997, pp. 238–239.

[2] R. Farjad-Rad, C.-K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4-μm CMOS 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 580–585, May 1999.

[3] G. Zhang, P. Chaudhari, and M. M. Green, "A BiCMOS 10 Gb/s adaptive cable equalizer," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2004, pp. 482–541.

[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] J. L. Zerbe, C. W. Werner, V. Stojanovic, F. Chen, J. Wei, G. Tsang, D. Kim, W. F. Stonecypher, A. Ho, T. P. Thrush, R. T. Kollipara, M. A. Horowitz, and K. S. 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.

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

[7] J. Sonntag, J. Stonick, J. Gorecki, B. Beale, B. Check, G. Xue-Mei, J. Guiliano, K. Lee, B. Lefferts, D. Martin, U.-K. Moon, A. Sengir, S. Titus, W. G.-Y. Wei, D. Weinlader, and Y. Yang, "An adaptive PAM-4 5 Gb/s backplane transceiver in 0.25 μm CMOS," in *Proc. IEEE Custom Integrated Circuits Conf.*, San Jose, CA, May 2002, pp. 363–366.

[8] J. Zerbe, Q. Lin, C. Werner, V. Stojanovic, A. Ho, and R. Kollipara, "Comparison of adaptive and nonadaptive equalization techniques in high performance backplanes over temperature, humidity, and impedance variations," in *DesignCon*, Santa Clara, CA, 2005.

[9] J. G. Proakis, *Digital Communications*, 3rd ed. New York: McGraw-Hill, 1995.

[10] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 5, pp. 711–717, Jun. 1991.

[11] Optical Internetworking Forum, "Common Electrical I/O (CEI)—Electrical and Jitter Interoperability Agreements for 6G+bps and 11G+bps I/O OIF-CEI-01.0,", Dec. 2004.

[12] W. Dally and J. Poulton, "Transmitter equalization for 4-Gbps signaling," *IEEE Micro*, vol. 17, no. 1, pp. 48–56, Jan./Feb. 1997.

[13] S. Wu, S. Ramaswamy, B. Bhakta, P. Landman, R. Payne, V. Gupta, B. Parthasarathy, S. Deshpande, and W. Lee, "Design of a 6.25 Gbps backplane SerDes with TOP-down design methodology," in *DesignCon*, Santa Clara, CA, 2004.

[14] M. Matsui, H. Hara, Y. Uetani, L. Kim, T. Nagamatsu, Y. Watanabe, A. Chiba, K. Matsuda, and T. Sakurai, "A 200 MHz 13 mm^2 2-D DCT macrocell using sense-amplifying pipeline flip-flop scheme," *IEEE J. Solid-State Circuits*, vol. 29, no. 12, pp. 1482–1490, Dec. 1994.

[15] B. Nikolic, V. G. Oklobdzija, V. Stojanovic, W. Jia, J. K. Chiu, and M. M. Leung, "Improved sense-amplifier-based flip-flop: Design and measurements," *IEEE J. Solid-State Circuits*, vol. 35, no. 6, pp. 876–884, Jun. 2000.

[16] J. C. Kim, Y. C. Jang, and H. J. Park, "CMOS sense amplifier-based flip-flop with two N-C^{2}MOS output latches," in *Electron. Lett.*, vol. 36, Mar. 16, 2000, pp. 498–500.

[17] B. Bhakta, S. Wu, P. Hanish, S. Hubbins, I. Hosagrahar, B. Dahl, and H. Liang, "Characterization and production test for CEI-6G-LR compliant 6.25 Gbps SerDes," in *DesignCon*, Santa Clara, CA, 2005.

[18] R. Payne, B. Bhakta, S. Ramaswamy, S. Wu, J. Powers, P. Landman, U. Erdogan, A. Yee, R. Gu, L. Wu, Y. Xie, B. Parthasarathy, K. Brouse, W. Mohammed, K. Heragu, V. Gupta, L. Dyson, and W. Lee, "A 6.25Gb/s binary adaptive DFE with first post-cursor tap cancellation for serial backplane communications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 68–69.

[19] P. Landman, K. Brouse, V. Gupta, S. Wu, R. Payne, U. Erdogan, R. Gu, A. Yee, B. Parthasarathy, S. Ramaswamy, B. Bhakta, W. Mohammed, J. Powers, Y. Xie, L. Wu, L. Dyson, K. Heragu, and W. Lee, "A transmit architecture with 4-tap feedforward equalization for 6.25/12.5 Gb/s serial backplane communications," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 66–67.

Robert Payne (M'96) received the B.S. and M.S. degrees in electrical engineering from the Georgia Institute of Technology, Atlanta, in 1994 and 1996, respectively.

He is a Distinguished Member of Technical Staff in the High Performance Analog Division of Texas Instruments, Inc. (TI), Dallas, TX. In 1995, he worked with Intel Corporation characterizing and developing ESD protection structures. He joined TI in 1996. He is currently working on high-speed/high-resolution analog-to-digital converters. His interests are in input/outputs (I/Os) and equalization, clock and data recovery, data conversion, and phase-locked loop design.

Mr. Payne is currently a member of the Wireline program committee for the ISSCC.

Paul Landman (S'92–M'95) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley, in 1989, 1991, and 1994, respectively.

His research focused on low-power digital design techniques and tools with an emphasis on digital signal processing (DSP) applications. In 1994, he joined Texas Instruments, Inc., Dallas, TX, where he worked on low-power and high-performance DSP design. He is currently a Distinguished Member of Technical Staff specializing in high-speed serial link technology.

Bhavesh Bhakta (S'95–M'96) received the B.S. and M.S. degrees in electrical engineering from the University of California, Irvine, in 1993 and 1996, respectively.

In 1993, he joined Silicon Systems, later acquired by Texas Instruments. Inc. (TI), Dallas, TX, developing high-performance mixed-signal circuits for disk drive read channels. He is currently a Member of Group Technical Staff with TI, working on high-speed serial link designs. His circuit design experience includes high-performance phase-locked loop (PLL), clock and data recovery (CDR), and equalizer designs in deep submicron complementary metal oxide semiconductor (CMOS).

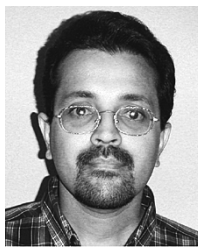

Sridhar Ramaswamy (M'88) received the B.Tech. (Hons.) degree from the Indian Institute of Technology, Kharagpur, India, in 1990, the M.S. degree from the University of Massachusetts, Amherst, in 1992, and the Ph.D. degree from the University of Illinois, Urbana-Champaign, in 1996, all in electrical and computer engineering.

He joined Texas Instruments, Inc., Dallas, TX, in 1996, where he is now a Senior Member of Technical Staff, working on issues related to large-scale integration of high-speed serial links. He has published over 25 papers in conferences and journals related to device modeling and circuit design.

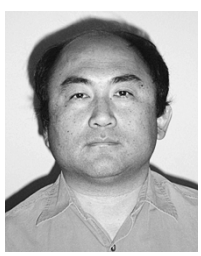

Song Wu (M'92) received the B.S. degree from Beijing University, Beijing, China, in 1986, the M.S. degree from Tsinghua University, Beijing, China, in 1988, and the Ph.D. degree from Rensselaer Polytechnic Institute, Troy, NY, in 1994.

He joined Texas Instruments, Inc., Dallas, TX, in 1995 and is now a Senior Member of Technical Staff. He has received 11 U.S. patents and published 20 technical articles.

John D. Powers (M'96) was born in Andalusia, AL, on March 4, 1976. He received the B.S.E.E. and M.S.E.E. degrees both from the Georgia Institute of Technology, Atlanta, GA, in 1998 and 2000 respectively.

He joined Texas Instruments, Inc. (TI), Dallas, TX, as a co-op in 1999. He has worked as a Design Engineer for TI since 2000.

M. Ulvi Erdogan (M'97) was born in Amasya, Turkey. He received the B.S. degree in electrical engineering and physics from Bogazici University, Istanbul, Turkey, in 1989, and the M.S. and Ph.D. degrees in electrical engineering from North Carolina State University, Raleigh, in 1991 and 1995, respectively.

In 1996, he joined Texas Instruments, Inc., Dallas, TX, where he worked on device and simulation program with integrated circuit emphasis (SPICE) modeling and circuit design. Until recently, he has been working on high-speed serial link interfaces. His current interests are in the design of high-speed circuits.

Ah-Lyan Yee (M'88) received the B.S.E.E. degree from the University of Texas at Arlington in 1985 and the M.S.E.E. degree from Southern Methodist University, Dallas, TX, in 1990.

Since 1997, he has been working on complementary metal oxide semiconductor (CMOS) serial link products for high-speed chip-to-chip and backplane applications. He joined Texas Instruments, Inc. (TI), Dallas, in 1983 as a Laboratory Technician in the Semiconductor Process and Design Center, involved in CMOS process and design development. From 1986 to 1996, he worked on various projects; bipolar CMOS (BiCMOS) gate array base cell and library development for TI application-specific integrated circuit (ASIC), silicon-on-insulator static random access memory (SOI SRAM), and CMOS serial link technology in the Digital Signal Processing R&D Center.

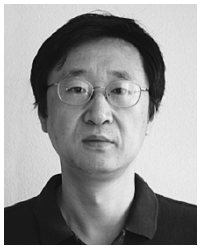

Richard Gu received the B.Sc. degree from Fudan University, Shanghai, China, and the Ph.D. degree in electrical engineering from the University of Waterloo, Waterloo, ON, Canada, in 1995.

In 1995, he joined Texas Instruments, Inc., Dallas, TX, where he is a Senior Member of Technical Staff. His major interests are in the area of communication systems and circuit design. He has 12 U.S. patents issued or pending and has published one technical book and 15 technical papers.

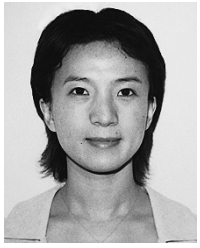

Lin Wu (M'00) received the B.S. and M.S. degrees from Tsinghua University, Beijing, China, in 1994 and 1996, respectively, and the Ph.D. degree from Iowa State University, Ames, in 2000, all in electrical engineering.

Since 1998, she has been working as an Analog IC Designer at Texas Instruments, Inc., Dallas, TX, first in the High-Speed Data Converter R&D Group, and then in the Serial Link Product Group for low jitter phase-locked loop (PLL) and clock tree design. She was promoted to a Member of Group Technical Staff in 2002.

Yiqun Xie (M'98) was born in China in 1971. He went to the University of Science and Technology of China, Hefei, Anhui, China, and received the M.S. and Ph.D. degrees both from University of California, Berkeley, in 1994 and 1998, respectively. His dissertation in the field of high-speed integrated circuit (IC) design.

He has been an Analog IC Designer since 2002 at Texas Instruments, Inc., Dallas, TX. Before 2002, he worked as an Analog IC Designer at Level One Communications, Intel, and Maxim Integrated Products.

Bharadwaj Parthasarathy (M'91) received the B.Tech. degree from the University of Mysore in 1990, and the M.S. degree from the University of South Florida in 1993, both in electrical and computer engineering.

He joined Texas Instruments, Inc., Dallas, TX, in 1996, where he is a Member of Group Technical Staff, working on design/integration of high-speed serial links with special interest in clock and data recovery circuits. He has published over 10 papers in conferences and journals related to circuit design.

Keith Brouse (M'04) received the B.S. and M.S. degrees from the Georgia Institute of Technology, Atlanta, in 2000 and 2002, respectively, both in electrical and computer engineering.

He joined Texas Instruments, Inc., Dallas, TX, in 2002, where he is a mixed-signal circuit designer. He is currently working on the design and integration of high-speed serial links. He has published three papers in conferences and journals related to circuit design.

Wahed Mohammed received the M.S. degree from Arizona State University, Tempe, in electrical engineering.

He is currently with Texas Instruments Incorporated, Dallas, TX, as a Mixed Signal Verification Engineer, where he is responsible for IC verification for high-speed mixed-signal circuits.

Keerthi Heragu received the B.E. degree in electronics and communication from the University of Mysore in 1992 and the Ph.D. degree in computer engineering from the University of Illinois at Urbana-Champaign in 1998.

He has since been with Texas Instruments, Inc., Dallas, TX, where he is currently managing a team that designs double data rate (DDR) interfaces used in application-specific integrated circuit (ASIC) chips.

Vikas Gupta (S'92–M'96) received the B.E. degree in electronics engineering from the University of Mumbai, Mumbai, India, in 1991, and the M.S. degree in electrical engineering from the University of Texas in 1995.

He joined Texas Instruments, Inc. (TI), Dallas, TX, in 1995 as a Reliability Engineer working on wafer-level reliability and later on ESD/latch-up related process development issues at the Silicon Technology Development Center. In 2000, he joined the Internet Infrastructure Business Unit working on the analog front-end and ESD strategy for high-speed serializer/deserializer (SERDES). He is now the 90-nm complementary metal oxide semiconductor (CMOS) platform manager, driving cross-functional teams to deliver a competitive solution (process, design, package, and test) to support TI's market segment strategies. He holds three patents on ESD-related process issues.

Mr. Vikas received the "Best Paper Award" at the 1998 EOS/ESD Symposium. He is in the Technical Program Committee of the EOS/ESD Symposium and is a TI representative to the Standards Committee of the ESD Association.

Lisa Dyson was born in Fort Wayne, IN, and graduated from ITT in Fort Wayne in 1984.

She joined Texas Instruments Inc., Dallas, TX, in 1984, working in test and design for the Final Test Department and the DISEG Division. In 1985, she transferred into the TI DSP Research and Development Division and worked in the development of gate arrays, 100Base-T2 physical layer design, multi-purpose ADC, DPLL, and multi-channel serial links. She continued the serial link development for the ASIC Division of TI and is currently working in the High Performance Analog division as a Digital Design Engineer, designing high-speed digital logic for serial interface designs.

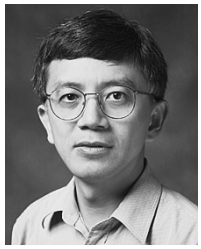

Wai Lee (S'90–M'83–S'83–M'88) received the S.B., S.M., and Ph.D. degrees in electrical engineering all from Massachusetts Institute of Technology, Cambridge, in 1983, 1986, and 1988, respectively.

He is currently a Fellow with Texas Instruments, Inc. (TI), Dallas, TX, and the US Design Manager in the Serial Link Products group in the High Performance Analog Division. Prior to TI, he was a Research Staff Member at IBM T. J. Watson Research Center, working on SiGe HBT and complementary metal oxide semiconductor (CMOS) microprocessors. He joined TI in 1993 to work on low-power and high-performance digital signal processing (DSP) design. He has since held design management positions in the DSP research and development (R&D) and application-specific integrated circuit (ASIC) Internet Infrastructure Business Units before his current assignment. He and his design team have been working on serial link products since 1998. He has authored and coauthored more than 30 journal articles.

Dr. Lee has been serving on the technical program committee of the Symposium on VLSI Circuits in the last 4 years and is currently the Publicity Chair. He has also served in the program committee for ISSCC.