

A 6.4-Gb/s CMOS SerDes Core With Feed-Forward and Decision-Feedback Equalization

Troy Beukema, Michael Sorna, Karl Selander, Steven Zier, Brian L. Ji, *Member, IEEE*, Phil Murfet, James Mason, *Senior Member, IEEE*, Woogeun Rhee, *Member, IEEE*, Herschel Ainspan, Benjamin Parker, and Michael Beakes, *Member, IEEE*

Abstract—A 4.9–6.4-Gb/s two-level SerDes ASIC I/O core employing a four-tap feed-forward equalizer (FFE) in the transmitter and a five-tap decision-feedback equalizer (DFE) in the receiver has been designed in 0.13- μm CMOS. The transmitter features a total jitter (TJ) of 35 ps p-p at 10^{-12} bit error rate (BER) and can output up to 1200 mVppd into a 100- Ω differential load. Low jitter is achieved through the use of an LC-tank-based VCO/PLL system that achieves a typical random jitter of 0.6 ps over a phase noise integration range from 6 MHz to 3.2 GHz. The receiver features a variable-gain amplifier (VGA) with gain ranging from -6 to $+10$ dB in ~ 1 dB steps, an analog peaking amplifier, and a continuously adapted DFE-based data slicer that uses a hybrid speculative/dynamic feedback architecture optimized for high-speed operation. The receiver system is designed to operate with a signal level ranging from 50 to 1200 mVppd. Error-free operation of the system has been demonstrated on lossy transmission line channels with over 32-dB loss at the Nyquist ($1/2$ Bd rate) frequency. The Tx/Rx pair with amortized PLL power consumes 290 mW of power from a 1.2-V supply while driving 600 mVppd and uses a die area of 0.79 mm².

Index Terms—Adaptive equalizers, analog equalization, decision-feedback equalization, high-speed I/O, transceivers.

I. INTRODUCTION

As improvements in silicon technology continue to advance the clock rates of processing cores, the data rate of I/O signals must be increased in step to realize maximum system-level performance. Industry standards [1] are being developed to define compliant channel and I/O electrical characteristics for operation at data rates from 6+ to 11+ Gb/s for both short-range (~ 4 -in, on-board) interchip links such as CPU memory applications and long-range backplane (~ 30 -in+, intercard) or coax (~ 10 m+) links that arise in systems such as scalable multiple-processor servers and high-speed routers/switches. The long-range/backplane applications are particularly challenging to realize robust high-speed I/O due to the combined effects of increased transmission line loss, crosstalk, and signal distortion arising from reflections that occur as data rates move into the microwave frequency range of operation and beyond.

To enable reliable operation on dispersive channels that produce significant intersymbol interference (ISI) at a given symbol (or baud) rate, the I/O core architecture can employ some form of line equalization [2]–[5]. A common approach to equalization for data rates up to 3–4 Gb/s is “feed-forward” equalization, or FFE, at the transmitter [2], [3], which predistorts the signal such that it is recovered at the receiver with a desired shape suitable for reliable data detection. Another form of equalizer is the “decision-feedback” equalizer, or DFE, which operates by subtracting the ISI arising from previously detected data symbols from the symbol currently being received [4], [7], [8]. The DFE operates as a nonlinear equalizer and can recover data that have been severely degraded by the distortion/noise arising from channel loss, reflections, and high-frequency crosstalk. All of these impairments can distort the signal beyond the capability of an FFE alone to equalize for reliable operation at 6-Gb/s+ data rates over backplane channels.

This paper describes the design of key elements of a CMOS I/O core intended for 6-Gb/s backplane interconnect and other “long-range” applications requiring up to several hundred I/Os per ASIC. The I/O core uses fixed transmitter FFE in combination with continuously adapted DFE in the receiver for line equalization. A system design overview including a description of FFE/DFE line equalization and the high-level I/O core architecture will be given. Next, the key circuit components of the I/O core that enable it to operate over lossy backplane channels will be described in detail. These components include the FFE-based transmitter equalizer, a linear receiver analog front-end with variable-gain amplifier (VGA), the DFE-based data slicer subsystem, and a low-jitter 4.9–6.4 GHz phase-locked loop (PLL). The performance of the I/O core as measured in laboratory tests on transmission line channels with high loss will also be summarized and compared with expected results from system simulations.

II. SYSTEM DESIGN

A. Overview of the Backplane Channel

A diagram of a typical backplane link is shown in Fig. 1(a). A processor or ASIC is mounted on a line card that plugs into a backplane. A transmission line on the backplane that might be 30–60 in long connects to another line card that holds a destination processor/ASIC. The transmission line introduces frequency-dependent loss, as shown in Fig. 1(c), that can become

Manuscript received April 12, 2005; revised July 25, 2005.

T. Beukema, W. Rhee, H. Ainspan, B. Parker, and M. Beakes are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: troyb@us.ibm.com).

M. Sorna, K. Selander, S. Zier, and B. L. Ji are with IBM Microelectronics, East Fishkill, NY 12533 USA.

P. Murfet and J. Mason are with IBM U.K., Hursley, Winchester SO21 2JN U.K.

Digital Object Identifier 10.1109/JSSC.2005.856584

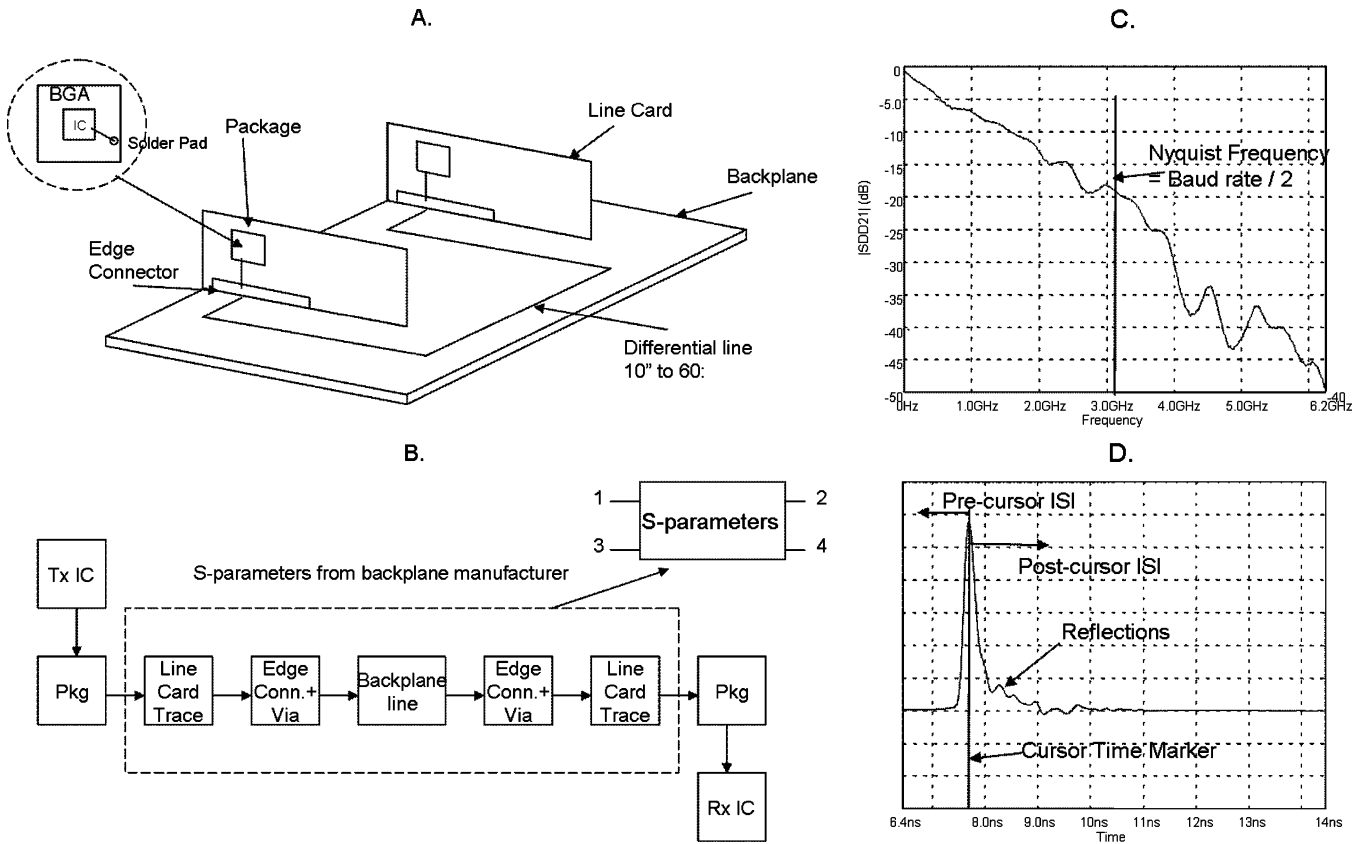


Fig. 1. Backplane channel environment. (a) Typical backplane/line card application. (b) Cascaded channel model. (c) Channel frequency response. (d) Channel impulse response.

significant in the high frequency range used in 6-Gb/s transmission and may also have stubs and other impedance discontinuities arising from the package and connectors that introduce nulls in the frequency response and corresponding time-domain reflection distortion, as shown in Fig. 1(d). Typical losses experienced at the 6-Gb/s Nyquist frequency (3 GHz) for FR-4 backplanes with lengths on the order of 30-in range from 20 dB to over 30 dB, presenting a significant challenge for the I/O equalization system.

B. FFE and DFE Line Equalization

At the transmitter, FFE is normally implemented as a low-frequency de-emphasis process, which reduces the low-frequency signal envelope level in proportion to the attenuation experienced by the high-frequency (1010...) pattern in the channel. This reduces ISI at the expense of lowered sensitivity for the low-frequency components of the transmitted signal. The de-emphasis process also increases the relative strength of the high-frequency content of the transmitted signal, thereby increasing the proportion of total signal power coupled into an adjacent line through a high-pass parasitic coupling path.

The FFE de-emphasis effect can be analyzed with a simple two-tap FFE using one post-cursor tap. The difference equation and corresponding z transform of this equalizer with the tap values normalized to an absolute-value sum of 1 is given by

$$yt(n) = \frac{x(n) + kx(n-1)}{(1 + |k|)} \rightarrow YT(z) = \frac{X(z)(1 + kz^{-1})}{(1 + |k|)} \quad (1)$$

where $x(n)$ is the binary input data stream, $yt(n)$ is the predistorted output signal, and $YT(z)$ and $X(z)$ are corresponding z transforms. At dc, z is equal to 1, and at the half-baud frequency, z is equal to -1 . The value k in (1) is set to a negative value to achieve low-frequency de-emphasis. In this case, the high frequency gain $YT(-1)/X(-1) = 1$ (i.e., unity gain) and the low frequency gain is attenuated or de-emphasized by the factor

$$\frac{YT(1)}{X(1)} = \frac{(1 - |k|)}{(1 + |k|)} \quad (2)$$

where the value of k increases in magnitude in proportion to increased channel loss. This type of gain-normalized FFE is done because a practical transmitter cannot output a voltage beyond a given swing limit, requiring the sum of the absolute values of the FFE coefficients to be scaled such that the output is not driven into compression for all input data patterns.

A DFE enables significantly less de-emphasis to be used at the transmitter since it effectively compensates high-frequency loss by canceling ISI arising from bandwidth limitations in the channel at the receiver data slicer. The transfer function for a simple one-tap DFE can be described by

$$y_d(n) = y(n) + h_1 \text{sgn}(y_d(n-1)) \quad (3)$$

where $y(n)$ is the signal level being received, $y_d(n)$ is the DFE-corrected signal, h_1 is the first DFE feedback tap weight, and $\text{sgn}()$ produces a 1 for $y_d(n-1) \geq 0$ and -1 otherwise. For the case of zero-error detection, (3) can be analyzed in a linear sense by replacing $\text{sgn}(y_d(n-1))$ with the transmitted data

$x(n-1)$ and $y(n)$ with the convolution of sequence $x(n)$ with the channel response $h(n)$ (described here as a discrete-time channel for simplicity). The resulting difference equation and corresponding z transform are

$$\begin{aligned} y_d(n) &= (x(n)*h(n)) + h_1x(n-1) \rightarrow Y_d(z) \\ &= X(z)(H(z) + h_1z^{-1}). \end{aligned} \quad (4)$$

To illustrate the operation of the DFE, assume a low-pass channel-response function

$$H(z) = h(0) + h(1)z^{-1} \quad (5)$$

with $h(0)$ normalized to 1 and $h(1)$ positive, resulting in reduced gain at high frequencies. In this case

$$YD(z) = X(z)(1 + h(1)z^{-1} + h_1z^{-1}). \quad (6)$$

To cancel the line ISI, h_1 is set to the value $-h(1)$, which realizes a fully equalized channel $Y_d(z) = X(z)$. For severe channel loss cases, h_1 approaches the value 1, and the DFE equalizer effectively regenerates the high-frequency (1010...) sequence at the receiver with no low-frequency de-emphasis of the signal required in the channel. Therefore, DFE provides the dual benefits of maintaining a larger received low-frequency signal envelope with no proportional increase in crosstalk level or other effective high-frequency noise amplification. This benefit is only achieved if DFE has a first feedback tap; systems not using the first feedback tap [4] do not regenerate the high-frequency alternating 1/0 pattern signal level that suffers the highest attenuation in lossy channels. For this reason, a DFE with first feedback tap was used in the system design.

Although DFE can offer significant advantages, one problem with the DFE equalizer is the inability to equalize ISI arising from pre-cursor channel response. The pre-cursor channel response [Fig. 1(d)] results in ISI generated by future data bits not yet detected while a current symbol is being sliced. Highly dispersive low-bandwidth channels may have significant time duration of pre-cursor response that can be mitigated through use of an FFE with pre-cursor taps. For this reason, a pre-cursor tap was used in the transmitter FFE architecture. A second problem with DFE is that it can only compensate a fixed time span of ISI. In contrast, an FFE can compensate ISI arising from transmission line loss (i.e., unipolar post-cursor impulse response) over a very wide time span since the FFE filter response, unlike the DFE filter response, is convolved with the impulse response of the channel. Therefore, in a very low bandwidth transmission line channel, significant unipolar post-cursor ISI may fall outside the time span covered by the DFE and some level of pre-emphasis using FFE becomes required to lower the ISI magnitude past the time span covered by DFE feedback taps. Finally, use of a DFE in a configuration with normalized first feedback tap greater than 0.5 exposes the receiver system to the potential of error propagation on alternating 1010... patterns, which requires consideration in the design of the data coding/error control correction layer.

When FFE and DFE are used together optimally, the strengths of both can be brought to bear on the channel and maximum

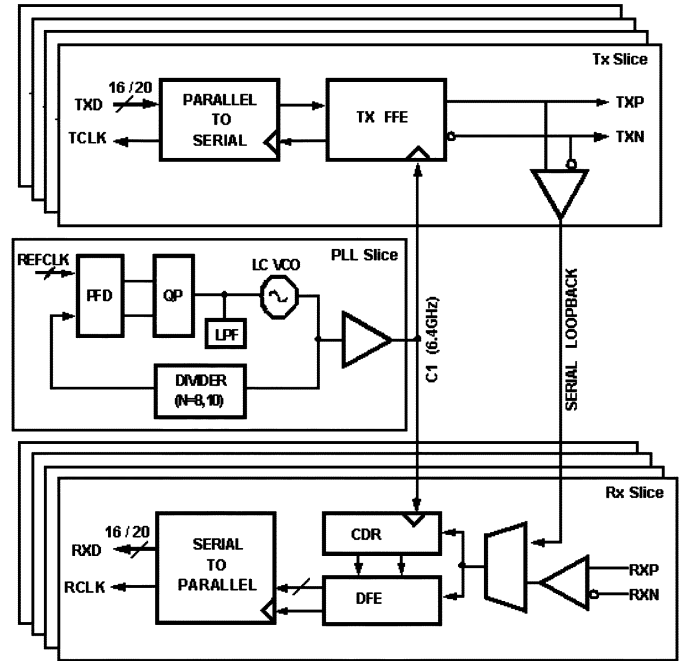


Fig. 2. SerDes core configuration.

performance obtained. The DFE with first feedback tap permits use of less low-frequency de-emphasis at the transmitter. This results in a larger received signal envelope with a lower proportion of energy at high frequencies, where channel reflections are normally larger and the signal/crosstalk ratio is smaller. A final advantage of receiver-based equalization using DFE and/or FFE [5] is the ability for the system to employ continuous adaptive equalization of its feedback taps to maintain the line equalization at optimum performance without the need for a reverse channel to a transmitter FFE.

C. I/O Core Architecture

A high-level block diagram of the I/O core subsystem is shown in Fig. 2. The I/O cores may be configured as a dual Tx/Rx pair or a Tx/Rx quad. A single PLL macro drives up to four Tx/Rx pairs and generates the reference 4.9- to 6.4-GHz clock for the system. Sharing the PLL slice between several Tx/Rx pairs is done to lower the power draw and die area of the I/O system. The receivers implement asynchronous local clock recovery using phase rotators to adjust the frequency and phase of the fixed-frequency PLL clock. The receiver deserializers can be configured to either 16- or 20-bit width. The transmit FFE has programmable tap settings that are normally set to fixed values (with optional power-up tap adaptation) while the DFE tap weights are continually updated using an on-chip adaptive equalization algorithm. The 50- Ω single-ended line terminations may be ac or dc coupled to the external line.

III. CIRCUIT DESIGNS

The I/O core is realized in a CMOS technology with a feature size of 0.13 μm using Cu/Al interconnects. The operating voltage range is from 1 to 1.6 V, with nominal V_{dd} of 1.2 V. A wide V_{dd} range of 1.0 to >1.6 V is achieved using resistor-




Fig. 3. Block diagram of transmitter.

loaded current mode logic (CML) circuits, exploiting their superior common mode rejection ratio (CMRR). In addition, a closed loop regulated current reference removes any dc power supply sensitivity due to current source channel length modulation.

The following sections expand the circuit design details of the transmitter, receiver VGA and peaking amp, DFE subsystem, and PLL.

A. Transmitter Section

A block diagram of the transmitter is shown in Fig. 3. The transmitter system is based on a full-rate clock architecture. This architecture results in inherently low transmit duty-cycle distortion (DCD). The transmitter realizes a four-tap FFE function with one pre-cursor and two post-cursor taps. The FFE tap weights have been sized to maximum relative weights of 0.25, 1.0, 0.5, and 0.25 for the pre-cursor through final tap settings to minimize diffusion capacitance load at the driver output. If a desired FFE tap weight profile cannot be contained within this fixed tap weight range, the main cursor tap can be backed off from full scale to accommodate the profile. Reduction of the main cursor tap magnitude results in a higher tap least significant bit (LSB) quantization level but expands the FFE tap coverage space to virtually any set of practical coefficients needed. The cursor tap adjustment range is 16, 64, 32, and 16 steps for the pre-cursor through post-cursor taps, with equal LSB weights for all taps. Each tap also has independent sign control.

The FFE tap weights can be either programmed to fixed values or optionally adapted on power-up using an up-channel link protocol. An up-channel receiver is used to receive tap inc/dec messages from a tap adaptation algorithm in the receiver in the case that power-up equalization is used. An automatic level control algorithm scales the transmit drive level so the peak-to-peak output voltage maintains a fixed programmed




Fig. 4. Transmit eye diagram at 6.4 Gb/s.

setting as tap values change. In effect, this algorithm realizes the FFE tap normalization (sum of absolute values of FFE taps = constant) described in Section II-B. This simplifies power control for the transmitter by automatically avoiding operating conditions that can overdrive or saturate the driver.

A CML latch first-in first-out (FIFO) that operates at full rate provides the data history signs that are XOR-ed with the sign of the tap weights. The outputs of the XOR gates control differential current switch drivers that generate pull-down current on either the + or - polarity of the line termination load. The magnitude of the pull-down current is programmed with current digital-to-analog converters (DACs) that bias the tail current of the current switches appropriately from "off" (0 current switched) to full tap weight possible (0.25, 1.0, 0.5, 0.25 relative weights).

The transmit driver system outputs a maximum of 1200 mVppd into a 100- Ω differential load. A typical 6.4-Gb/s transmission demonstrates a total jitter (TJ) of 35 ps p-p for BER < 10^{-12} as shown in Fig. 4. This eye diagram plot includes the effects of driver slew rate limit due to IC parasitics, ceramic package, test board trace (4 in), two connectors, and low-loss cabling (18 in) to the sampling scope. The total channel loss at 3.2 GHz for this eye trace is approximately 6 dB, requiring a normalized two-tap FFE setting of approximately [0.89, -0.11] to equalize the line. A PRBS-7 data pattern was used to generate the eye diagram.

B. Receiver System

The receiver system features a VGA, peaking amplifier, five-tap DFE, and analog phase-rotator-based clock-and-data recovery (CDR) loop (see [6] for further information on the analog phase rotator technique) as shown in Fig. 5. As opposed to the full-rate transmitter, the receiver is based on a half-rate clocking architecture. This design was chosen to both minimize power draw and enable a DFE feedback architecture to be realized with adequate settling time for high rate operation. The DFE system features continuous adaptive equalization to result in maximal vertical eye opening at the data slicing instant. An optional up-channel driver is implemented at the receiver input to enable training of transmitter FFE taps on power-up. The cascaded VGA, peaking amplifier, and DFE summation sections have been designed to achieve a minimum gain of 3 with 3-dB bandwidth of 3.2 GHz and 1-dB compression point of 300 mVpd. To help achieve these design specifications,




Fig. 5. Receiver block diagram.

extensive use of a custom genetic algorithm (GA)-based optimizer was made that assisted in tuning the circuit designs to meet prescribed performance targets over process, voltage, and temperature (PVT) corners.

C. Receiver VGA

The VGA enables linear (<1 -dB compression) operation on up to 1200 mVppd input signals. In order to improve linearity of the VGA for high input levels and to allow operation with high input signals at low supply voltages, the received signal is split into full-amplitude and half-amplitude paths by using a resistive divider in the 50- Ω termination network as shown in Fig. 5 and a parallel amplifier architecture [9].

The full- and half-signal data paths each drive a separate switched-gain amplifier as shown in Fig. 6. The gain is adjusted by setting the degeneration resistance of each amplifier to one of eight values using a thermometer-coded switched resistor network. At low values of gain, only the half-signal amplifier is active. As the gain is increased, the full-signal amplifier is enabled at the most significant bit (MSB) transition and works in parallel with the half-signal amplifier. When the full-signal amplifier is inactive, the common mode voltage at the output is maintained by steering its tail current into the half-signal amplifier. In this way, the overall VGA provides 16 gain steps with a targeted gain range from 0.5 to 3 with gain resolution of approximately 1 dB/step. The circuit is able to provide gain adjustment without corrupting data by enabling a low glitch transition between successive gain settings.

D. Receiver Peaking Amp

The VGA drives a second-stage peaking amplifier that can provide additional fixed gain and can also be used to introduce a programmable amount of high-frequency emphasis. The high-frequency emphasis (or peaking) effect is realized through reduction in the low-frequency gain relative to the high-frequency gain. The primary purpose of the peaking amplifier is to enable the receiver to equalize the Rx package loss (which ranges typically from 1 to 3 dB at the Nyquist frequency) so that a measured eye diagram at the package input will not have significant extra ISI at the detection latch sample point. This improves performance for standard compliance tests that may refer eye closure to the receiver package input point. The potential degradation added to the system by increasing the proportion of high-frequency noise with receiver peaking can be compensated by using less low frequency de-emphasis in the transmitter FFE tap settings.

A simplified diagram of the peaking amplifier is shown in Fig. 7. Similar to the VGA architecture, this design uses a parallel amplifier topology, but here both amplifiers have the same input, which is the output of the VGA. The overall amplifier response is a combination of the responses of these two amplifiers, similar to the approach described in [10]. The peaking level is adjusted by controlling the ratio of the tail currents to two amplifiers, one with a fixed 6-dB peak and one with no peaking. Applying no bias current to the 6-dB peaking amplifier and full bias to the flat-gain amplifier produces no peaking, conversely full bias on the 6-dB peaking amp and no bias on the second amplifier give maximum peaking. The 6-dB peaking section is implemented using an amplifier with a degeneration network

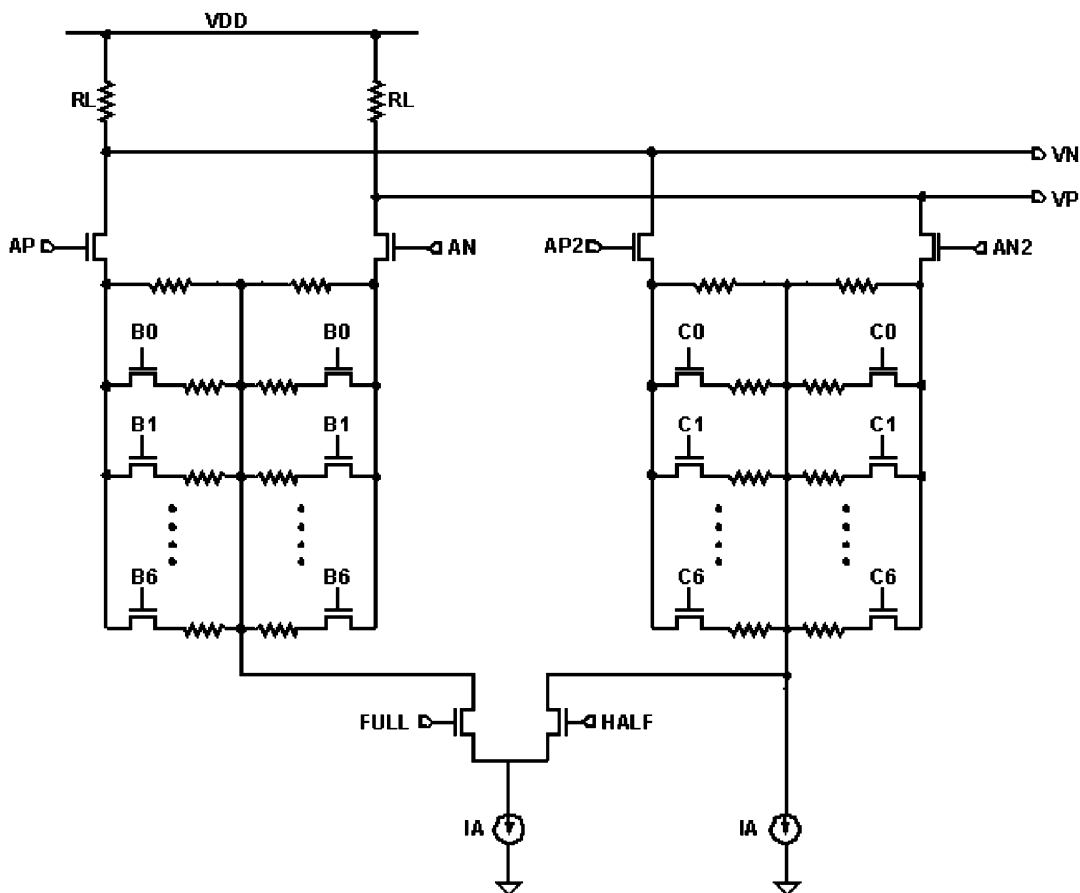


Fig. 6. Circuit diagram of two-path VGA.

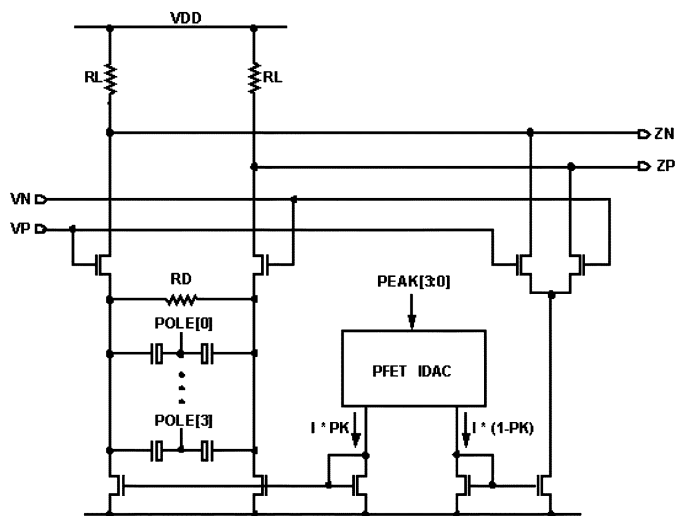


Fig. 7. Circuit diagram of peaking amp.

formed from a resistor and a shunt capacitor. As the degeneration impedance drops with increasing frequency, the amplifier gain increases proportionately until the bandwidth limitation of the amplifier causes the gain to roll off. This forms a peak in the gain/frequency characteristic. The frequency at which the peak occurs, referred to here as the pole position, is adjusted by varying the value of the degeneration capacitor. Both the peaking amount and the pole position can be varied in 16 steps.

E. DFE Summers and Data Slicer

An expanded block diagram of the receiver DFE system is shown in Fig. 8. The DFE sums a compensation value with the received signal as a function of prior sliced data decisions and associated tap weights. The DFE feedback architecture chosen uses a speculative approach for the first feedback tap (further description of this technique is found in [7] and [8]) with dynamic feedback for the remaining feedback taps. The speculative DFE feedback works by summing both + and - first feedback tap weights with the received signal. Both of these values are then sliced to a binary value. The correct slicer output is selected by the previously detected data bit value, which effectively results in a sum of the desired DFE feedback sign to the signal to cancel the ISI without the need to wait for an analog feedback waveform stabilization. Dynamic feedback is used for taps two through five, which does require analog feedback waveform stabilization. However, since the receiver design is based on a half-rate clocking architecture, the time interval for successive data samples is doubled, resulting in a relaxed settling time requirement for the dynamic feedback. This enables the DFE system to work at a 6.4-Gb/s data rate without pushing the limits of achievable analog waveform settling times in the 0.13- μm CMOS technology.

DFE correction is added to the signal by pulling weighed currents from either the +/− leg of a differential amplifier output using current switches as shown in Fig. 9, which is a block diagram of the dynamic feedback DFE summer. The tail current

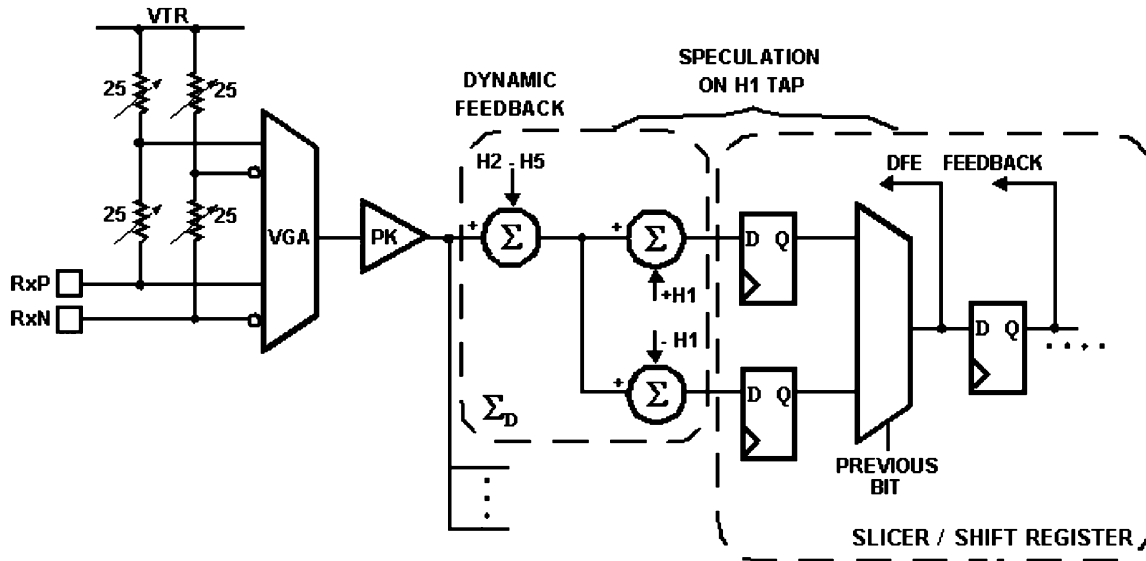


Fig. 8. DFE feedback system.

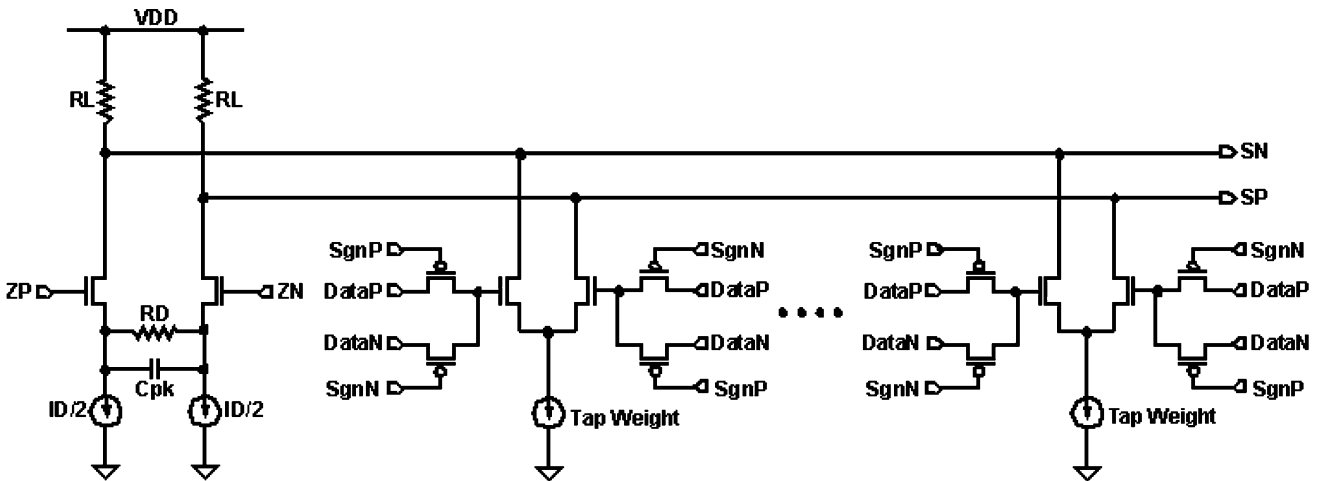


Fig. 9. Circuit diagram of multiple-tap DFE summer.

magnitudes in the current switches set the desired tap weight. Pass gates are used to XOR the data value with the tap weight sign for dynamic feedback taps $h_2 \dots h_5$. The same circuit structure is used for the speculative feedback DFE summers, with fixed data signs of + and - on each of the two feedback paths. A small amount of capacitive degeneration in the summer circuit is used to add high-frequency peaking to extend the bandwidth of the summers while minimizing power draw.

Feedback tap weights are generated using a sign-error-driven adaptation algorithm. The sign error is generated using a separate sampling path, shown in Fig. 5, which samples the sum of an adapted dc offset value with the received signal to determine if the received signal is greater or less than the mean received signal at the data sampling time. A long-term correlation of sign error to the data polarity at a given delayed bit position indicates that there is ISI arising from this data bit position at the sampling time. Integration of the feedback tap weights in the direction opposite the sign of the long-term correlation results in eventual convergence of the tap weights to realize minimum residual ISI contribution and maximum eye opening at the data

sample point. The feedback tap weights are continually adapted so that the system is capable of tracking slow changes in the channel that may arise due to temperature or voltage fluctuations in the system.

F. PLL Core and Clocking System

A block diagram of the PLL system is shown in Fig. 10. The PLL is responsible for multiplying a reference clock by 8/10/16/20 to the system baud rate of 4.9–6.4 GHz. To achieve a low-power programmable 8/10/16/20 divider, a 4/5 dual-modulus divider is employed [11]. The charge pump current is digitally controllable and programmed automatically as a function of the selected division ratio to maintain constant PLL bandwidth [12]. The PLL employs an LC-tank-based voltage-controlled oscillator (VCO) design with 16 overlapping coarse-tune bands to lower the VCO tuning sensitivity and minimize jitter while covering the 4.9–6.4-GHz band. A linear amplifier converts the differential loop filter outputs to a single-ended voltage to control the fine-tuned varactor. An additional high-order pole is embedded in the linear amplifier design to further filter out

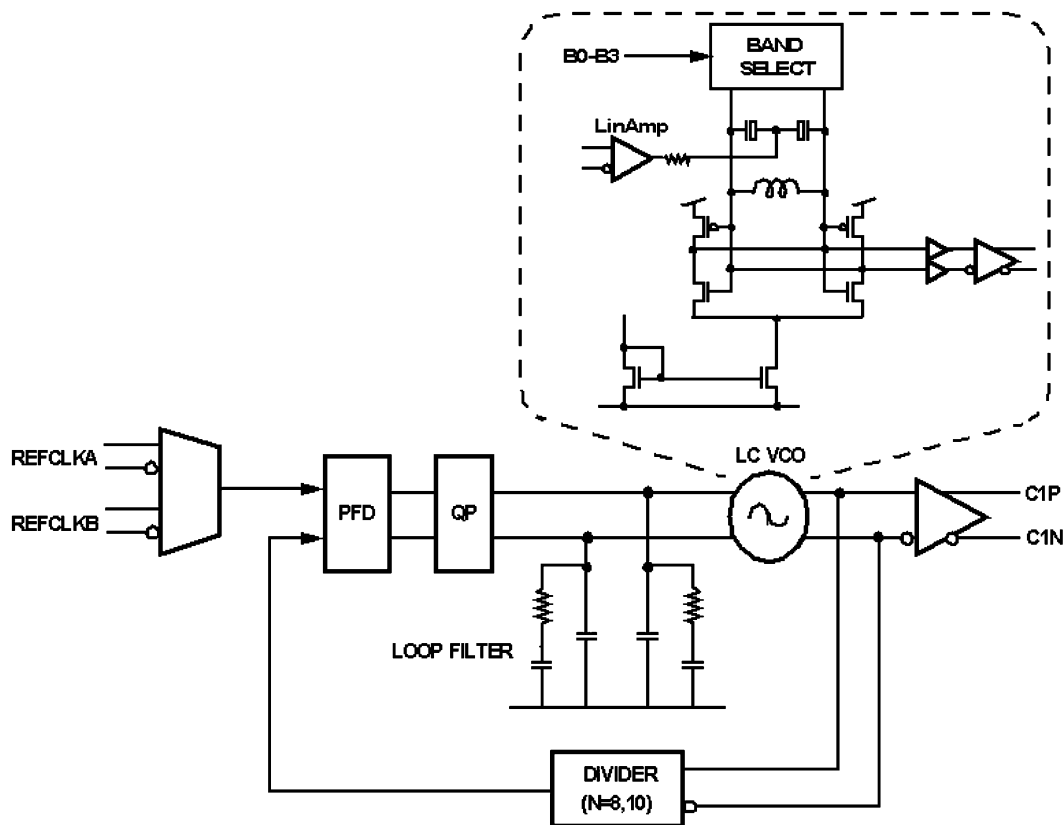


Fig. 10. PLL macro with LC VCO.

high frequency noise, which offers a type-II fourth-order PLL system.

The VCO design, shown in Fig. 10, comprises a pair of cross-coupled CMOS inverters that generate a negative resistance to compensate for the losses of the resonant tank formed by an accumulation-mode varactor and 1.8-nH spiral inductor [13]. The CMOS inverter N/P ratio is optimized for minimum phase noise [14]. The varactor is comprised of a digitally switched binary-weighted array of four varactors for band select of the VCO at startup, in parallel with a fine-tuning varactor controlled by the PLL. The varactor channel length of $0.25 \mu\text{m}$ is chosen to achieve wide tuning range and high Q for low VCO phase noise. The inductor is a symmetric spiral using the two $2\times$ thick upper metal layers shunted together with vias to achieve a simulated $Q = 8$ at 6 GHz. A series resistor after the linear amplifier provides isolation between the VCO switching nodes and the linear amplifier. A final key component of the VCO system is a power-up coarse calibration algorithm that picks the optimum coarse-tune band (frozen after power up) to maintain VCO operation over worst-case operating condition (supply voltage/temperature) variations.

Spectral plots of the phase noise characteristic of the PLL are given in Fig. 11 for VCO frequencies of 4.5 and 6.25 GHz. While the receiver CDR can track jitter components of the phase noise that are < 6 MHz, the system is susceptible to performance degradation from the jitter with frequency components from 6 MHz and beyond. In this frequency range, the noise drops 20 dB/decade and hits a noise floor at 100 MHz of approximately -140 dBc/Hz. Integration of the phase noise

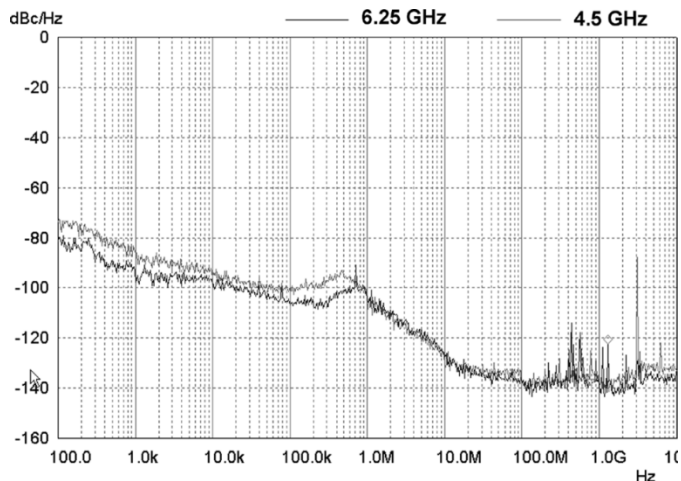


Fig. 11. PLL closed-loop phase-noise characterization.

characteristic from 6 MHz to 3.2 GHz results in a phase jitter that has been measured to fall within the range of 0.4 to 0.7 ps root mean square (rms) on nominal parts at room temperature. The jitter has also been verified to remain below 1 ps rms over process/temperature corners. The variation of this jitter is attributed dominantly to VCO gain variation, with best phase noise/lowest jitter achieved at band settings and process/temperature corners corresponding to lowest VCO gain. The VCO gain is largest in the center of each of the 16 tuning bands and at the high frequency band settings.

The receiver divides the full-rate clock from the PLL by 2 to produce half-rate in-phase and quadrature (I/Q) outputs as

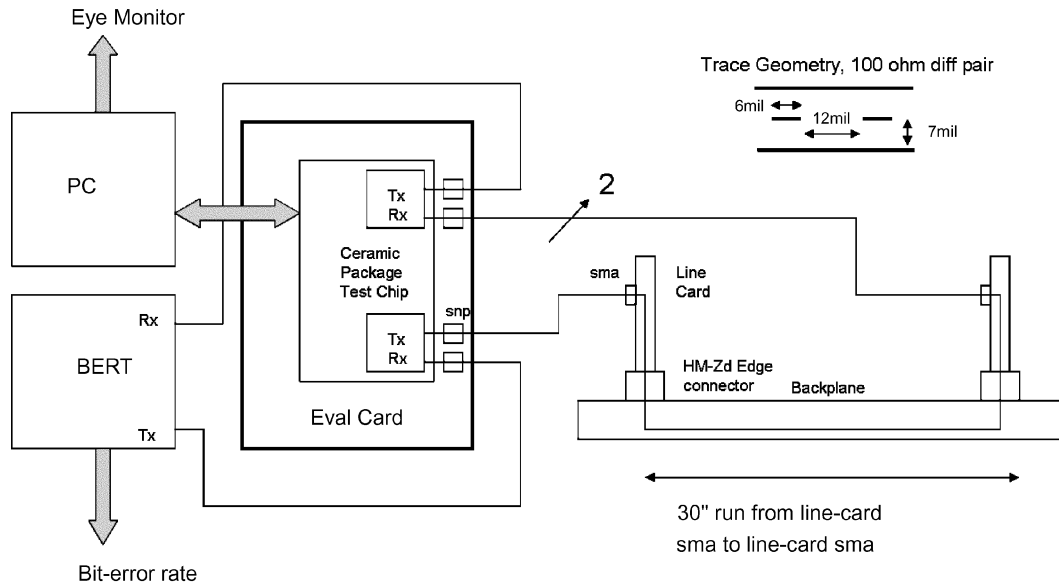


Fig. 12. Performance evaluation test setup.

shown in Fig. 5. The I/Q signals drive two four-quadrant phase rotators that produce data and edge latch clocks with phase controllable from 0° to 360° . For simplicity, only one edge/data latch is illustrated in Fig. 5, although it is understood that the half-rate design requires data and edge latches for both true and complementary phases of the rotator outputs. The edge latch paths have a “dummy” summer (summer with a “0” input) inserted to match the delay of the data path DFE summers, as illustrated in Fig. 5. An early-late digital CDR algorithm is used, which has a tracking bandwidth of approximately 6 MHz. The CDR algorithm continually updates the rotator phase shifts at a rate of $1/8$ the operating baud rate to maintain the desired data sampling point in the presence of frequency offset or periodic low-frequency jitter on the received data.

IV. SYSTEM PERFORMANCE RESULTS

A. Evaluation Card and Performance Measurement Setup

To evaluate system performance on application channels, a custom test board has been designed that allows connection of a test chip containing the I/O core to external channels through high-speed cable connectors. The test board interfaces to a PC, which enables configuration and monitoring of internal registers of the I/O core. A block diagram of a typical test setup configured in a loop-back test is shown in Fig. 12. Either an external bit-error-rate test (BERT) or a built-in self-test (BIST) BERT that generates PRBS-7 sequences can be used to evaluate the bit error rate (BER) performance of a link. Functionality has been designed into the core to enable read-out of the vertical inner-eye opening at the sampling latch from internal registers. Use of this feature enables generation of the inner-eye opening contour, which provides useful information for both evaluation of system performance and correlation to link model simulation results.

B. High-Level Link Simulation

A custom high-level link simulation tool was developed for the project. The link simulation tool builds a composite end-to-end channel response by concatenating S-parameter descriptions of IC models (predominantly transmit and receive termination electronics with associated device/wiring/ESD parasitics) with S-parameter descriptions of the IC package and application channel, as shown in Fig. 1(b). Statistical analysis of the simulated signal at the receiver detection latch enables estimation of the operating BER margin (typically $< 10^{-15}$ BER with no error correction coding in the backplane environment) for different link configurations. Behavioral simulation is useful for providing both performance prediction using S-parameter-based link models and a reference to compare circuit simulations and hardware performance against to verify expected operation.

C. Hardware Performance Tests

Two channels are considered here to test the equalization performance of the serializer/deserializer (SerDes) design: One 30-in backplane interconnect with 25-dB loss at the Nyquist frequency using Rx DFE-only equalization, and a second 70-in test board channel with ~ 32 -dB loss using Tx FFE + Rx DFE equalization. The receiver peaking amp is set to approximately 1 dB fixed pre-emphasis level for all hardware tests.

The 30-in evaluation backplane has two paddle cards attached using HM-zd connectors and is constructed of Nelco material with 6-mil striplines for the high-speed data paths. The frequency response of this link is shown in Fig. 13(a). The channel has approximately 25 dB of loss at the Nyquist frequency of 3.125 GHz for a data rate of 6.25 Gb/s. As shown in the eye diagram plot in Fig. 13(b), this amount of channel loss results in an eye with no discernable opening at the receiver input when transmit FFE is not used. The receiver DFE system has demonstrated blind (i.e., no special training pattern required for the

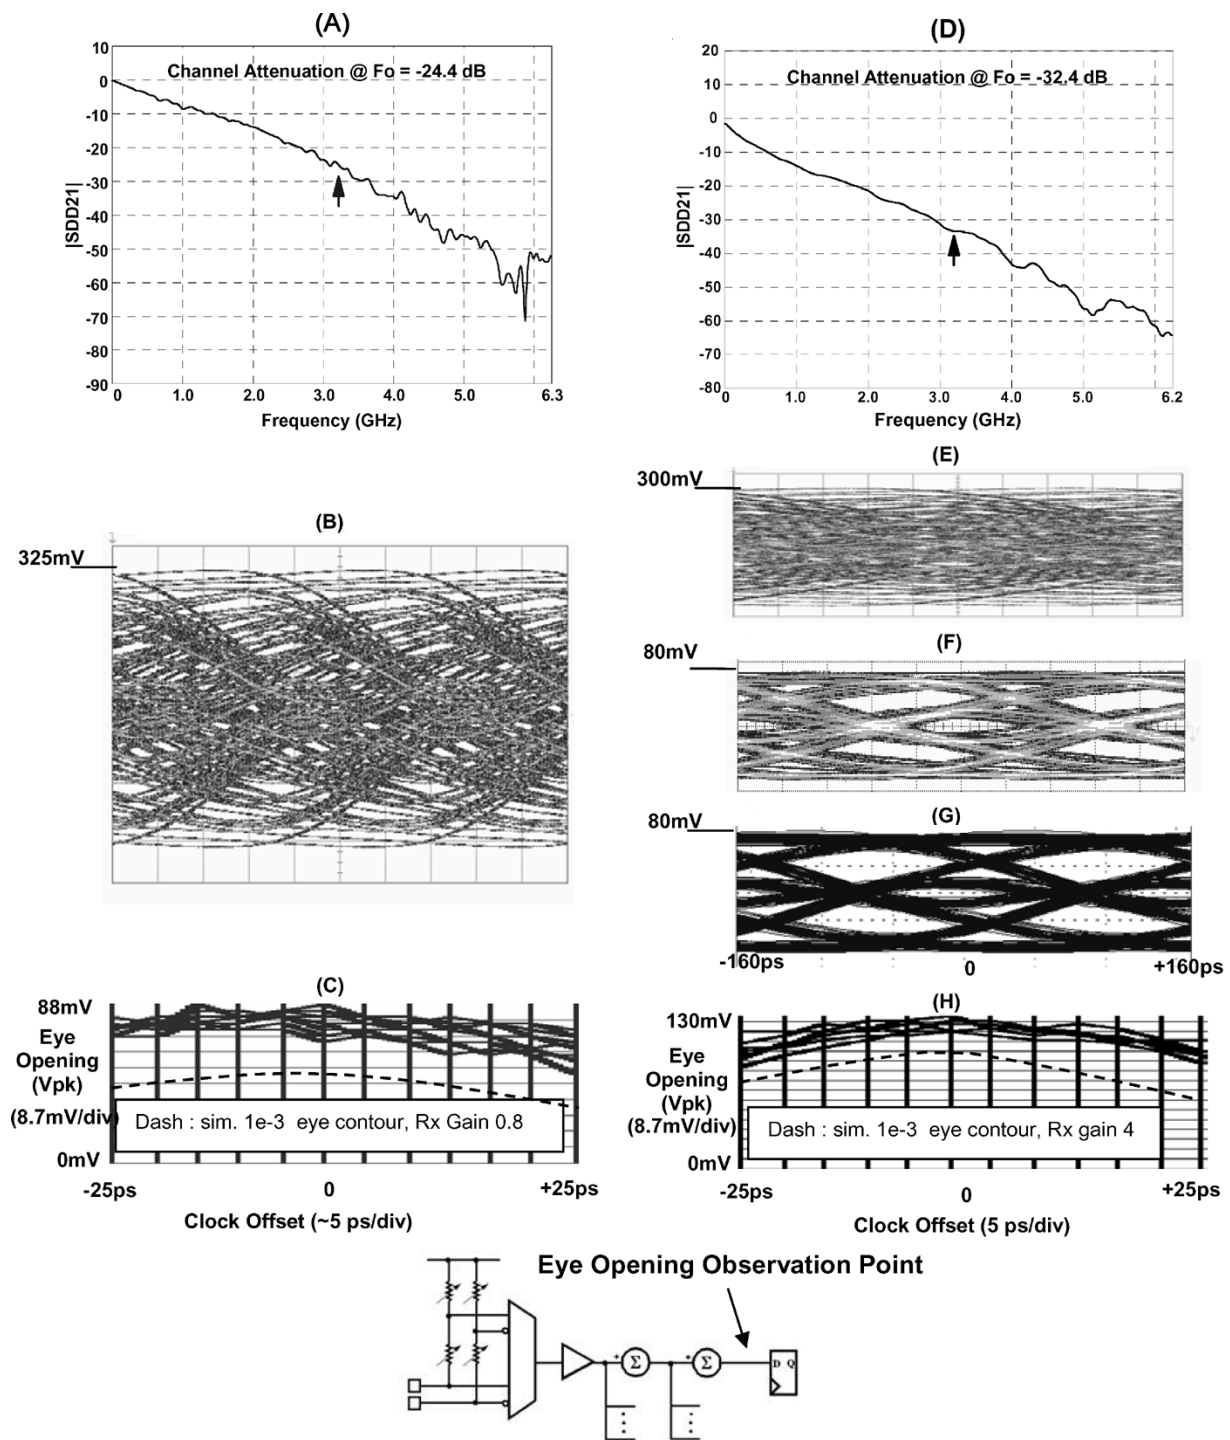


Fig. 13. Measured eye diagrams of 30-in backplane and 70-in Nelco test card channel. (a) 30-in backplane channel, two HM-zd connectors. (b) Measured 6.25-Gb/s eye at Rx input, no Tx FFE. (c) DFE5 only equalization 30-in backplane. (d) 70-in Nelco channel + interconnect cables. (e) Unequalized eye diagram: 6.25 Gb/s. (f) Measured eye at Rx package input with Tx FFE on. (g) Simulated eye at package input. (h) FFE4-DFE5 equalization 70-in Nelco T-line.

adaptation) equalization of this closed eye response to achieve low error rate operation.

A diagram showing a plot of the DFE-corrected vertical eye level at the data sampling latch versus sampling clock offset for this case is shown in Fig. 13(c). In this diagram, the vertical axis scale has been corrected to 8.75 mV/division compared to the incorrect 10 mV/division scale that was shown in [15]. The measured eye contours have a confidence of roughly 10^{-3} , i.e., approximately 999/1000 measured samples are greater or

equal to the shown levels. It can be seen from the figure that the vertical eye opening is maintained at a level >50 mVp for over ± 25 ps away from the nominal clock sampling point, indicating a solid error-free data recovery that is also verified by BERT. Due to time limitations, $< 10^{-15}$ BER operation was not verified in hardware, although the system ran error free in this configuration over time intervals corresponding to 10^{-12} BER confidence. The part has also been verified to run error free over one weekend time frame ($\sim 10^{15}$ total bits) in other tests, as-

certaining an achievable BER floor below 10^{-14} . A plot of the simulated vertical eye opening contour at a confidence of 10^{-3} is also shown superimposed on the measurement data with a dashed line in Fig. 13(c). The simulated eye opening was generated with an Rx gain of 0.8. The measured results show a similar shape to the simulated eye opening, but the measured signal level is larger, most likely due to a different Rx gain factor in the hardware system.

The second test case uses a channel chosen to stress the maximum equalization capability of the FFE/DFE combination. System simulations of the design indicate that the FFE4/DFE5 combination is capable of equalizing channels to $< 10^{-15}$ BER with loss up to 30–35 dB at the half-baud frequency. For hardware proof, a 70-in transmission line channel was built from test boards constructed with Nelco 4000-13 material. As shown in Fig. 13(d), this channel has 32-dB loss at the half-baud frequency for 6.25-Gb/s data rate. The unequalized eye at the channel output is shown in Fig. 13(e). In this case, both FFE and DFE must be used together to equalize the channel to low-error operation. Using FFE-only equalization, the hardware ran at approximately 10^{-4} BER, indicating the 32-dB loss channel is beyond the reach of the FFE equalizer to compensate. This result agreed well with the system link simulation, which predicted a BER floor of 7.2×10^{-5} under this condition. The DFE alone was also unable to handle this line due to excess ISI beyond the time span of its five feedback taps. However, with both FFE and DFE active, the system ran with no detected errors, with a horizontal eye opening (or jitter tolerance) of 34% unit interval (UI) at 10^{-15} BER predicted using the link simulator with a fixed maximum receiver gain of 3.

A plot of the measured eye at the receiver package input after FFE4 equalization is shown in Fig. 13(f). The measured results agree well with an eye diagram generated using the link simulator, shown in Fig. 13(g). The received signal has a low-frequency envelope level of about 80 mVp as a result of the transmitter de-emphasis, which has dropped the low frequency signal envelope 11 dB from the unequalized [Fig. 13(e)] 300 mVp level. A plot of the DFE-corrected vertical eye opening at the receiver's data sampling latch versus sample clock offset is given in Fig. 13(h). The eye opening is in the range of 100 mVp for ± 25 -ps clock offset from nominal sample time, indicating good compensation of the line ISI. The expected 10^{-3} confidence vertical eye found from simulations is shown as a dashed line in Fig. 13(h) for a receiver gain of 4. The simulated result is slightly below the hardware measurements, again indicating a mismatch in receiver gain setting between the hardware system and the simulation. Although the hardware has been designed to achieve a minimum receiver gain of 3, the system is capable of receive gain greater than 3 under nominal conditions to accommodate gain loss at PVT corners. Comparison of the eye contours in Fig. 13(c) and (h) shows that FFE4 + DFE5 equalization of a line with 32-dB loss at the half-baud frequency runs with significantly more margin than DFE5-only equalization of a line with 25-dB loss, demonstrating the value of the combination of FFE and DFE equalization together for best system performance.




Fig. 14. Die microphotograph of four-port Tx/Rx core die size is 0.79 mm^2 .

V. SUMMARY AND CONCLUSION

This paper has described the design of key components of a 4.9–6.4-Gb/s CMOS I/O core realized in $0.13\text{-}\mu\text{m}$ CMOS technology. A single Tx/Rx I/O pair draws approximately 290 mW from a 1.2-V supply while consuming a die area of 0.79 mm^2 . Power draw and area include amortized PLL power/area (1/4 PLL area and power for the four-port core). A die photo of a four-port Tx/Rx core with PLL slice is shown in Fig. 14.

The I/O core has demonstrated performance close to predictions from a system-level S-parameter-based simulation program. Highlights of performance include demonstrated decision-feedback equalizer (DFE)-only blind equalization of a 30-in backplane (25 dB loss at half-baud frequency) with two paddle cards at 6.25 Gb/s. Expected performance was also verified on a 70-in Nelco test card channel with over 32-dB loss at the half-baud frequency using feed-forward equalizer (FFE)/DFE equalization. The receiver system is based on a linear analog front-end with variable-gain amplifier (VGA), analog peaking amp, and hybrid speculative/dynamic feedback DFE architecture optimized for high-speed decision-feedback settling. The base architecture of the described I/O core is expected to be extendable to 12 Gb/s and beyond, enabling future I/O systems to keep up with the ever increasing demands of high-speed CMOS processing cores used in modern switch, routers, and computing systems.

ACKNOWLEDGMENT

A partial list of contributors to the system definition, design, layout, fabrication, and testing of the design described in this paper include J. Abler, J. Natonio, L. Chieco, W. Kelly, H. Xu, H. Camara, D. Storaska, P. Metty, M. Cordrey-Gale, G. Nicholls, L. Hsu, G. Gangasani, A. Mulgrav, and D. Friedman. Many more not listed here from the IBM Microelectronics High Speed Serial group, the IBM Yorktown Communications Technology group, and IBM U.K. Hursley have also made contributions to the design. Special thanks are also extended to Dr. M. Soyuer for his encouragement of this paper.

REFERENCES

- [1] *Common electrical I/O (CEI)—Electrical and jitter interoperability agreement for 6+ Gbps and 11+ Gbps I/O*, Optical Interconnect Forum—OIF-CEI-01.0, Dec. 2004.
- [2] R. Gu, J. Tran, H.-C. Lin, A.-L. Yee, and M. Izzard, "A 0.5–3.5 Gb/s low power low-jitter serial data CMOS transceiver," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 1999, pp. 352–353.
- [3] J. T. Stonic, G.-Y. Wei, J. L. Sonntag, and D. K. Weindler, "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.
- [4] 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," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 2003, pp. 80–81.
- [5] J. E. Jaussi, G. Balamurugan, D. Johnson, B. Casper, A. Martin, J. Kennedy, R. Mooney, and N. Shanbhag, "An 8 Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation and clock deskew," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 2004, pp. 246–247.
- [6] D. Zheng, X. Jin, E. Cheung, M. Rana, G. Song, Y. Jiang, Y.-H. Sutu, and B. Wu, "A quad 3.125 Gb/s/channel transceiver with analog phase rotators," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 2002, pp. 70–71.
- [7] 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.
- [8] V. Stojanovic, A. Ho, B. Garlepp, F. Chen, J. Wei, E. Alon, C. Werner, J. Zerbe, and M. A. Horowitz, "Adaptive equalization and data recovery in a dual-Mode (PAM2/4) serial link transceiver," in *Symp. VLSI Circuits Dig. Tech. Papers* Honolulu, HI, Jun. 2004, pp. 348–351.
- [9] "Micro-architecture innovation in A/D converters," in *R. J. Brewer Oxford Analog Signal Processing Conf. Proc.*, 2002.
- [10] A. J. Baker, "An adaptive cable equalizer for serial digital video rates to 400 Mb/s," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 1996, pp. 174–175.
- [11] W. Egan, *Frequency Synthesis by Phase Lock*. New York: Wiley, 1990.
- [12] I. Novof *et al.*, "Fully integrated CMOS phase-locked loop with 15 to 240 MHz locking range and ± 50 ps jitter," *IEEE J. Solid-State Circuits*, vol. 30, no. 11, pp. 1259–1266, Nov. 1995.
- [13] H. Ainspan and J. Plouchart, "A comparison of MOS varactors in fully integrated CMOS LC VCOs at 5 and 7 GHz," in *Proc. Eur. Solid-State Circuits Conf. (ESSCIRC)*, Stockholm, Sweden, Sep. 2000, pp. 448–451.
- [14] A. Hajimiri and T. H. Lee, "Design issues in CMOS differential LC oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 717–724, May 1999.
- [15] M. Sorna *et al.*, "A 6.4 Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 62–63.



Troy Beukema received the B.S.E.E. and M.S.E.E. degrees from the Michigan Technological University, Houghton, in 1984 and 1988, respectively.

From 1984 to 1988, he was an R&D Engineer at Hewlett-Packard, focusing on communications test equipment. He joined Motorola Communications Sector, Schaumburg, IL, in 1989 and contributed to the development of digital cellular wireless systems with focus on digital signal processing algorithm design and implementation. In 1996, he joined IBM, Yorktown Heights, NY, where he is presently a

Research Staff Member concentrating on the areas of radio architecture and modulation for 60-GHz wireless systems and system designs for adaptive equalization of 6–12 Gb/s wireline serial I/O links. He has written a serial link simulation program for use in performance analysis and development of high-speed I/O systems that are in use at several IBM sites.



Michael Sorna received the B.S. degree in electrical engineering from the University of Pittsburgh, Pittsburgh, PA, in 1984, and the M.Eng. degree in electrical engineering from Rensselaer Polytechnic Institute, Troy, NY, in 1986.

Since 1984, he has been working at the IBM East Fishkill, NY, facility and was appointed Senior Technical Staff Member in 2003. From 1984 to 1993, he worked in the Test Equipment Engineering Group, developing high-speed wafer-level production test equipment. From 1993 to 1998, his responsibility centered around high-speed SerDes/chip circuit design, in which a wide variety of standards was addressed: IEEE 1394, fiber channel, gigabit Ethernet, Infiniband, and SONET (OC-192 and OC-768). His design experience includes both analog and digital areas in bipolar (including SiGe), CMOS (0.09–0.8 μm), and GaAs technologies. His recent achievements include leading a team to develop the first IBM OC-768 (43 Gb/s) SONET-compliant chip set. He is currently responsible for IBM ASIC SerDes R&D, including circuit design, architecture, and lab characterization. He has authored numerous technical publications and presentations, primarily in the field of SerDes circuits and BIST techniques.

Karl Selander received the B.E. degree in electrical engineering from Stevens Institute of Technology, Hoboken, NJ, in 1991, and the M.S. degree in computer engineering from Syracuse University, Syracuse, NY, in 1995.

In 1991, he joined IBM, East Fishkill, NY, where he is currently an Advisory Engineer and has worked on the design of high-performance packages and high-speed communications circuits. In addition to leading development teams, he has recently been involved in the architecture and design of high-speed CMOS SerDes cores. He holds a number of patents in this area as well.



Steven Zier received the B.E. degree in electrical engineering from the State University of New York, Stony Brook, in 1983, and the M.S. degree in computer science from Rensselaer Polytechnic Institute, Troy, NY, in 1993.

He joined IBM Microelectronics, Hopewell Junction, NY, in 1983. From 1983 to 1988, he was a Designer for low-cost bipolar gate array libraries. From 1988 to 1994, he was a Designer for mainframe BiCMOS/CMOS SRAM. From 1994 to 2000, he was a Designer for Flash/DRAM. And since 2000,

he has been working on SiGe BiCMOS/CMOS serial communications (in the 1–40 Gb/s range). His latest research interests include a genetic algorithm tool he has written that sizes circuits using PowerSpice. This tool has been used extensively for Flash/DRAM and serial communications work by himself and others at various sites.



Brian L. Ji (M'94) received the B.S. degree from the University of Science and Technology of China, Hefei, in 1984, and the Ph.D. degree in physics from Harvard University, Cambridge, MA, in 1991.

From 1991 to 1994, he was a Research Scientist at the State University of New York at Stony Brook, where he studied nanofabrication, single electron memory/logic devices, and superconducting devices. In 1995, he was a Visiting Scientist in physical sciences at the IBM T. J. Watson Research Center, Yorktown Heights, NY. He joined IBM Microelec-

tronics, Hopewell Junction, NY, in 1996, where he worked on several projects in VLSI circuit design, test, and product definitions, including 256-Mb, 512-Mb, 1-Gb DRAMs, and the logic-based embedded memory. Since 2001, he has been involved in high-speed serial link product design.



Phil Murfet received the First Class Honors degree from Loughborough University, Leicestershire, U.K., in 1972.

In 1972, he joined IBM, working in Hursley Laboratories, U.K. He has also worked in Burlington, VT, and Rochester, MN. He has worked on the design of silicon for computer display, magnetic disk files, and communication links. At present, he leads a small team that specializes in analog front-end design for fast serial links.



James Mason (M'83–SM'02) received the First Class B.Sc. (Honors) degree from the University of Birmingham, U.K., in 1982, and the M.B.A. with merit degree from the University of Southampton, U.K., in 1995.

Between 1978 and 1985, he held various engineering positions at Lucas Industries, principally working on engine management systems and semiconductor device design. In 1985, he joined IBM U.K. Laboratories and has worked on bipolar and CMOS integrated circuit design for video, disk

drive, and communication applications. He also has electronic product design and manufacturing experience in computer display, storage subsystem, and fiber-optic products.

Mr. Mason is a member of the Institution of Electrical Engineers and a Chartered Electrical Engineer.



Woogeun Rhee (S'93–A'98–M'00) received the B.S. degree in electronics engineering from the Seoul National University, Seoul, Korea, in 1991, the M.S. degree in electrical engineering from the University of California, Los Angeles, CA, in 1993, and the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, in 2001.

From 1997 to 2001, he was with Conexant Systems, Newport Beach, CA, where he was a Principal Engineer in the Wireless Communication Division.

Since 2001, he has been a Research Staff Member at IBM T. J. Watson Research Center, Yorktown Heights, NY. His current interests are in phase-locked loops and clock-and-data recovery circuits for high-speed I/O interfaces, and in low-power RF circuits with emphasis on frequency synthesizers for wireless communications.

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

In 1989, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he has been involved in the design of mixed-signal and RF integrated circuits for high-speed data communications.

Benjamin Parker received the B.S. degree in physics from Bowdoin College, Brunswick, ME, in 1979, and the M.S. degree in physics from Brown University, Providence, RI, in 1981. His graduate work dealt with the optical properties of adsorbed layers on metal surfaces.

In 1986, he joined the GaAs group at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he worked on the characterization of III–V semiconductor systems. In 1991, he joined the Mixed-Signal Communications IC Design Group, working on the design and verification of high-speed serial communication links.



Michael Beakes (S'79–M'80) is a Senior Engineer at the IBM T. J. Watson Research Center, Yorktown Heights, NY. Throughout his 25-year career at IBM, he has been involved in custom circuit design, design automation, and education. He is also a Mentor and Advisor for grade school and high school technology and robotics programs. His educational background includes physics, electrical engineering, and computer engineering.