

IBM Research Report

A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology

**John F. Bulzacchelli, Mounir Meghelli, Sergey V. Rylov, Woogeun Rhee,
Alexander V. Rylyakov, Herschel A. Ainspan, Benjamin D. Parker,
Michael P. Beakes, Aichin Chung, Troy J. Beukema, Petar K. Pepeljugoski,
Lei Shan, Young H. Kwark, Sudhir Gowda, Daniel J. Friedman**
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598



Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology

John F. Bulzacchelli, Mounir Meghelli, Sergey V. Rylov, Woogeun Rhee, Alexander V. Rylyakov, Herschel A. Ainspan, Benjamin D. Parker, Michael P. Beakes, Aichin Chung, Troy J. Beukema, Petar K. Pepeljugoski, Lei Shan, Young H. Kwark, Sudhir Gowda, and Daniel J. Friedman

IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598 USA

Abstract – This paper presents a 90-nm CMOS 10-Gb/s transceiver for chip-to-chip communications. To mitigate the effects of channel loss and other impairments, a 5-tap decision feedback equalizer (DFE) is included in the receiver and a 4-tap baud-spaced feed-forward equalizer (FFE) in the transmitter. This combination of DFE and FFE permits error-free NRZ signaling over channels with losses exceeding 30 dB. Low jitter clocks for the transmitter and receiver are supplied by a PLL with LC VCO. Operation at 10-Gb/s with good power efficiency is achieved by using half-rate architectures in both transmitter and receiver. With the transmitter producing an output signal of 1200 mVppd, one transmitter/receiver pair and one PLL consume 300 mW. Design enhancements of a half-rate DFE employing one tap of speculative feedback and four taps of dynamic feedback allow its loop timing requirements to be met. Serial link experiments with a variety of test channels demonstrate the effectiveness of the FFE/DFE equalization.

Index Terms – Adaptive equalizer, feed-forward equalizer, decision-feedback equalizer, serial link, transceiver.

Sponsorship – This work was supported in part by MPO (Maryland Procurement Office) Contract H98230-04-C-0920.

Contact information:

John F. Bulzacchelli
IBM T. J. Watson Research Center Rm. 40-103
P. O. Box 218
Yorktown Heights, NY 10598 USA

Tel: 1-914-945-1058

Fax: 1-914-945-2141

E-mail: jfbulz@us.ibm.com

I. Introduction

The continuing growth in processing power of digital computing engines and the increasing demand for advanced network services are creating a need for higher bandwidth data transmission in systems such as servers and data communication routers. To meet this need, industry standards [1] are being developed which define the channel characteristics and I/O electrical specifications of short reach (~ 4-in, on-board) and long reach (~ 30-in+, intercard) serial links operating at data rates in excess of 10 Gb/s. While serial link transceivers in the 6 Gb/s range [2]-[4] are often intended to extend the bandwidth of “legacy” backplane channels, reliable operation above 10 Gb/s will require in many cases improved channel characteristics, so the standards above 10 Gb/s are primarily aimed at new (“greenfield”) backplane designs benefiting from improvements in board, connector, and chip-level package technologies.

Even with greenfield backplane designs, however, the need to remain price-competitive will discourage adoption of the most exotic (and expensive) board, connector, and package technologies. As in the recent past, advanced equalization capabilities in the I/O circuitry will be employed to compensate for the signal distortions of lower cost interconnect technologies. Optimizing cost tradeoffs at the system level requires knowledge of how much equalization is needed for a specific combination of board, connector, and package technologies.

In order to gain a more detailed understanding of these tradeoffs between interconnect and circuit technologies, a prototype of a complete 10 Gb/s serial link has been designed, fabricated, and tested. The project can be divided into two major efforts. One was the development of advanced packaging and board technologies, the design, modeling, and measurements of which are detailed in [5]. The other was the development of a 90-nm CMOS 10 Gb/s transceiver with 4-tap feed-forward equalizer (FFE) and 5-tap decision-feedback equalizer (DFE), the main topic of

this paper. This transceiver has equalization capabilities similar to those of the current generation of serializer/deserializer (SerDes) ASIC I/O core [2], which operates at 6.4 Gb/s.

The paper begins in Section II with a review of the backplane channel application and a discussion of the considerations made in defining the signaling and equalization capabilities of the 10 Gb/s transceiver. Section III describes architectural and circuit details of the transceiver components. Section IV discusses the features of the link demonstrator test chip and presents the results of several serial link experiments. A summary in Section V concludes the paper.

II. Background

A. *Backplane Channel Characteristics*

A typical backplane/line card application is sketched in Fig. 1(a). A long (30-in or more) transmission line on the backplane is used to transfer data from a processor or ASIC on one line card to a processor or ASIC on another line card. Several physical effects degrade signal integrity at data rates above a few Gb/s. Skin effect and dielectric losses of the transmission lines become severe at these data rates. Via stubs on the circuit boards and other impedance discontinuities associated with the chip packages and connectors cause reflections easily observed in the channel impulse response (Fig. 1(b)). In the frequency domain, these reflections cause notches which further degrade the channel frequency response (Fig. 1(c)). Since the transmitted signal is attenuated by loss, it is easily corrupted by crosstalk from other channels. Even for greenfield backplanes with improved board technology, the loss at 5 GHz (Nyquist frequency for 10 Gb/s data) may be 20-30 dB. With the channel adding so much loss and distortion to the signal, the data eye at the far end of the link (Fig. 1(a)) is completely closed, and advanced equalization is required to recover the transmitted bits.

B. *Signaling and Equalization Considerations*

One approach to increasing data rates in high-loss channels is the use of multilevel signaling such as four-level pulse amplitude modulation (PAM-4) [6], [7], [8]. Compared to binary non-

return-to-zero (NRZ) signaling at the same data rate, PAM-4 signaling reduces the baud rate (and therefore the required bandwidth) by a factor of two. On the downside, the additional voltage levels used in PAM-4 signaling decrease the level spacing (vertical eye height) by a factor of three (9.5 dB). These last two statements lead to the following rule-of-thumb [6]: if the loss difference between NRZ and PAM-4 Nyquist frequencies exceeds 10 dB, the SNR improvement due to baud rate reduction may exceed the 9.5 dB level spacing penalty, and PAM-4 signaling is likely to give better link performance. Adding a linear equalizer to flatten the channel response does not alter this basic analysis, as such a linear equalizer amplifies crosstalk and other high-frequency noises as much as the desired signal, leaving the high-frequency SNR unchanged.

On the other hand, adding a nonlinear equalizer in the form of a receive-side DFE does alter the analysis because the DFE is able to flatten the channel response without amplifying noise or crosstalk. A simple counter-example shows that the conventional rule-of-thumb is invalidated by the use of a DFE. Consider transmitting 12.5 Gb/s data over a channel with the impulse and frequency responses of a 25th-order Bessel filter, as plotted in Fig. 2. The cutoff frequency of the Bessel filter has been chosen so that the loss (36.5 dB) at 6.25 GHz is 28 dB greater than the loss (8.4 dB) at 3.125 GHz. According to the usual rule-of-thumb, this should be a clear-cut case for using PAM-4 over NRZ signaling. To test this assertion, the equalized eye diagrams for NRZ and PAM-4 signaling over this channel were calculated with the high-level link simulation tool described in [2]. Fig. 3 shows the results when the receiver includes a 2-tap DFE (in both NRZ and PAM-4 modes), with no other equalization (such as FFE) applied. In each of these diagrams, the DFE feedback signals are held at constant values across a two baud interval in order to allow clear inspection of the eye margins for the symbol being detected at $t=0$ ps. (Since the DFE feedback signals are only correct for this symbol, the equalized eye diagrams represented in this fashion are not periodic.) Because of the short duration of the Bessel filter impulse response, increasing the number of DFE taps improves the eye diagrams negligibly. In the simulations, the

receiver gain is set to unity so the heights of the vertical eye openings represent the voltage margins for sampling the data, referred back to the receiver input pin. The simulated NRZ eye is bigger than the PAM-4 eye both vertically (by 93%) and horizontally (by 20%).

An examination of how post-cursor cancellation by the DFE alters the frequency response of the channel reveals why the usual rule-of-thumb breaks down here. As shown in Fig. 4(a), the pulse response of the Bessel filter channel has only two significant pre-cursors and two significant post-cursors when sampled at 12.5 GHz. The pre-cursors, main cursor, and post-cursors can be treated as a discrete-time sequence (Fig. 4(b)). Since the latch of the receiver only samples its input at discrete times, this discrete-time sequence fully characterizes the channel in terms of its effect on the signal being detected. Taking a discrete-time Fourier transform (DFT) of this sequence yields the unequalized frequency response of the channel. Assuming that the DFE has enough taps and is accurately adapted, the DFE feedback cancels out the post-cursors of this discrete-time sequence. Therefore, the frequency response of the serial link after applying the DFE can be obtained by taking a DFT of the sequence comprising only the pre-cursors and main cursor. Applying this method to the Bessel filter example yields the frequency responses shown in Fig. 4(c). Because the post-cursors have the same polarity as the main cursor, their cancellation by the DFE does reduce the dc gain by a few dB. The DFE substantially flattens the channel response, though, so the equalized response exceeds the unequalized response at both 3.125 and 6.25 GHz. After equalization, the loss difference between 3.125 and 6.25 GHz is only 6.3 dB, so switching to PAM-4 signaling is not worth the 9.5 dB level spacing penalty.

While the Bessel filter of the previous example is not a realistic model for a lossy backplane channel, post-cursor cancellation by a DFE also flattens the channel responses of more realistic transmission line models, with the result that NRZ signaling often has larger margins than PAM-4 signaling. Past research by Stojanovic [9] has demonstrated that even a 1-tap DFE is sufficient to provide NRZ signaling with better voltage margins than PAM-4 signaling (with no DFE)

when transmitting 6.25 Gb/s data over backplane channels with lengths of 3-in, 10-in, and 20-in. In the early phase of this project, NRZ signaling and PAM-4 signaling at 10-12.5 Gb/s were compared in high-level link simulations using S21 data of various backplane channels. For a large majority (>90%) of the channels surveyed, these simulations showed that if the receiver includes a 5-tap DFE (operational in both NRZ and PAM-4 modes), better margins can be obtained with NRZ signaling than with PAM-4 signaling. Since this study did not indicate compelling advantages for PAM-4 signaling in this application, the decision was made to develop an NRZ-only transceiver employing both linear (4-tap FFE) and nonlinear (5-tap DFE) equalization, a choice which is becoming increasingly popular [2]-[4], [10]. Linear equalization by the FFE complements the operation of the DFE by compensating for pre-cursor ISI, as well as post-cursor ISI outside the time span of the DFE. More discussion of the merits of a combined FFE/DFE system can be found in [2].

III. Transceiver Components

This section describes the implementation of the three major components of the I/O circuitry: transmitter, receiver, and PLL. The high-level functions of these components closely match those of the 130-nm CMOS 6.4 Gb/s SerDes core presented in [2]. The high-speed sections of these components are realized with resistor-loaded current mode logic (CML) circuits, which provide good common mode and power supply rejection, so the general circuit style is also similar to that used in the 6.4 Gb/s core. Here the circuits are designed and fabricated in a 90-nm CMOS technology which uses a strongly nitrated oxide to achieve low gate leakage [11]. In CML circuit applications, the gate leakage is negligible (e.g., 10 nA) over all process, voltage, and temperature (PVT) corners. An effect that cannot be neglected in this technology is the stress induced by shallow trench isolation (STI) [12], which can reduce the NMOS drain current by up to 30%. In circuits such as current mirrors where device matching is important for accurate

biasing, STI stress effects are mitigated by placing dummy transistors at each end of the active transistors.

While technology scaling from 130-nm to 90-nm reduces the gate delays of conventional static CMOS by about 20%, the enhancement in speed for CML circuits is more modest. Since the main core voltage of the chip is now lower (1.0 V instead of 1.2 V), the differential pairs and current sources of the CML circuits must be increased in size to maintain transistor operation in saturation [13]. The larger devices increase the loading on previous stages and hinder compact layouts, so the wiring parasitics are also greater. Consequently, the improvement in CML circuit speed is relatively small – in many cases, less than 10%. Assuming that the current levels and resistor values of the CML circuits are unchanged, the lower supply voltage does reduce power dissipation by 17%, but it is difficult to trade off this increased power efficiency for higher speed, as increasing the currents (and device sizes) of all the CML circuits yields much less than linear increases in speed as the technology limit is approached. Achieving a 56% improvement in operating frequency (from 6.4 Gb/s to 10 Gb/s) without a large increase in power dissipation requires significant design enhancements to the transceiver components beyond mapping them to the newer technology. These design enhancements are the main focus of the subsections to follow. Circuit building blocks which were not modified from the 6.4 Gb/s design (aside from device resizing) are not discussed here, as they were already covered in [2].

A. *Transmitter with 4-Tap FFE*

The half-rate architecture of the transmitter with baud-spaced 4-tap FFE is illustrated in Fig. 5. The transmitter receives a half-rate (C2) clock from the on-chip PLL. On the link demonstrator test chip, $\frac{1}{4}$ -rate data are externally delivered to the transmitter as single-ended signals. The 4:2 multiplexer (MUX) retimes these signals and generates two differential half-rate even and odd data streams. These streams are shifted by 1 unit interval (UI) relative to each other and then interleaved together with a MUX to form full-rate data for the first tap of the FFE.

Additional shifting and interleaving yield the delayed data for the remaining 3 taps of the FFE. After sign selection by exclusive OR (XOR) gates, these signals are amplified with pre-driver and driver stages, whose output currents are summed together in the line termination loads. The tap weights are programmed to fixed values with current digital-to-analog converters (DACs) that bias the tail currents of the output drivers. The FFE taps have been sized to maximum relative weights of 0.25, 1.0, 0.5, and 0.25 (with one pre-cursor and two post-cursors), with DAC resolutions of 4, 6, 5, and 4 bits, respectively. Most of the CML logic gates are powered off V_{DD} (nominally 1.0 V), which represents the main digital supply of a large ASIC. The pre-drivers and drivers are powered off V_{DDA} (nominally 1.2 V), which is the I/O supply to which the termination resistors are connected. A similar partition of supplies is used in the receiver. With the second tap set to maximum and the other taps powered down, the transmitter dissipates 70 mW and produces an amplitude of 1200 mV peak-to-peak differential (mVppd) into a 100 Ω differential load. A breakout test site of this transmitter is described in detail in [14].

The use of a half-rate architecture instead of the full-rate one described in [2] improves power efficiency, as the lower operating frequencies of the CML gates allow them to be scaled down in power. Because the even and odd data streams are shifted with simple latches (L) instead of the master-slave flip-flops required in a full-rate implementation, the increase in CML gate count is quite modest (despite the appearance of complexity in Fig. 5). Counting the circuits between the 4:2 MUX and the XOR gates (the only circuits where the half-rate and full-rate versions differ), the half-rate design employs 13 gates (nine half-rate latches and four 2:1 MUXes), while the full-rate design employs 12 gates (three half-rate latches, one 2:1 MUX, and eight full-rate latches). A study conducted for a 6.4 Gb/s design implemented in the same 90-nm CMOS technology showed that the half-rate latches could be scaled down in power by a factor of 2.5 relative to the full-rate ones, in which case the 13 gates of the half-rate design consume 25% less power than the 12 gates of the full-rate design. The power savings would be even greater at a

data rate of 10 Gb/s, as full-rate designs become especially power hungry near the technology limit [15]. Because the clock is distributed from the PLL to the transmitter at half-rate ($C2=5$ GHz) instead of full-rate ($C1=10$ GHz), additional power savings can be obtained in the sizing of the CML clock distribution buffers. The above-mentioned study for a 6.4 Gb/s design showed that a 3.2 GHz clock could be distributed with 33% less power than a 6.4 GHz clock. The power savings would again be greater for a 10 Gb/s design, as the losses of the on-chip wires increase with frequency, and the intrinsic bandwidth of the CML buffers becomes more of a limiting factor at 10 GHz. With a half-rate architecture, the duty cycle of the C2 clock needs to be close to 50% to minimize transmit duty cycle distortion (DCD). To prevent the accumulation of DCD, the C2 clock distribution path includes an ac-coupled CML buffer [16], which rejects the dc offsets of the previous stages.

B. Receiver with 5-Tap DFE and Digital CDR Loop

The receiver architecture is presented in Fig. 6. A T-coil network [17] is used for broadband compensation of the electrostatic discharge (ESD) diode capacitance, which improves input return loss (S_{11}) and front end bandwidth. To ensure linear operation of the DFE summing stages, a variable gain amplifier (VGA) regulates the data swing at the slicer input to about 600 mVppd. In the highest gain setting, the cascaded VGA and DFE summers achieve a minimum gain of 3 with a -3-dB bandwidth of at least 5 GHz across all PVT corners. The 5-tap DFE which equalizes and slices the data employs a half-rate architecture, so in-phase (I) and quadrature (Q) C2 clocks are shipped from the PLL to the receiver. Phase interpolation (PI) [18] by phase rotators controlled by a digital CDR loop generates the I and Q clocks used to sample the centers and edges of the data bits. In addition to detecting the data bits, the DFE block monitors the amplitude (Amp) of the equalized eye by comparing it with an expected target; this information is used in adapting the equalizer. After further demultiplexing, the even and odd data (as well as the Amp samples) from the DFE and the edge samples from the phase detector are

processed by DFE logic which performs continuous adaptation of the DFE and by CDR logic which keeps the data and edge sampling clocks properly aligned to the incoming data bits. As discussed in [2], the DFE tap weights are adapted with a sign-sign least-mean square (LMS) algorithm in which the sign errors of the Amp samples are correlated with the polarities of the data bits. The power dissipation of the receiver, including the DFE and CDR logic, is 130 mW at nominal PVT. This total does not include the $50\ \Omega$ line drivers used to monitor $\frac{1}{4}$ -rate data (D_0 - D_3) from the 2:8 demultiplexer.

High linearity with large (1200 mVppd) inputs is achieved by adopting a parallel amplifier architecture for the VGA [2], whose schematic is shown in Fig. 7(a). The received signal is split into full-amplitude and half-amplitude paths with resistive dividers in the input termination network (Fig. 6). The half-amplitude signal (AP2/AN2) is connected to a differential amplifier which is always active, while the full-amplitude signal (AP/AN) is connected to a differential amplifier which is only turned on at high gain settings (when the input signal is small). Thermometer-coded switched resistor networks are used as variable degeneration to adjust the VGA gain within each operating mode (FULL or HALF). Over all gain settings, the amplified data signal at the slicer input exhibits less than 1 dB of compression at a 600 mVppd amplitude. This topology is more suitable for low-voltage operation than the one used in the 6.4 Gb/s design (Fig. 7(b)), as the bias currents do not flow through these resistors, and increasing degeneration does not affect circuit headroom. A potential drawback of the topology chosen here is that the tail device capacitances may cause peaking at low gain settings. With careful device sizing, the peaking is held to less than 3 dB at the lowest gain setting. Another improvement is the addition of proportional-to-absolute-temperature (PTAT) biasing, which helps reduce gain variation over temperature by partially compensating the reduction in device transconductance at high temperature due to lower mobility. (Complete compensation of the transconductance would require a stronger than PTAT temperature dependence [19], but only partial compensation is

adequate here.) PTAT biasing is also employed in the DFE summing amplifiers. The power dissipation of the PTAT-biased circuitry is 18 mW at $T=25^{\circ}\text{C}$ and rises to 24 mW at $T=125^{\circ}\text{C}$. The cascaded dc gain of the VGA and DFE summers was measured for all 16 gain settings (Fig. 7 (c)) on a breakout test site; the results are plotted in Fig. 7(d). The measured gain range is 20 dB, with an average gain step of 1.3 dB and a maximum gain step of 2.0 dB.

A major challenge in the design of a DFE is ensuring that the feedback signals have settled accurately at the slicer input before the next data decision is made. If a full-rate DFE architecture is used, the feedback loop delay (including the decision-making time of the slicer and the analog settling time of the DFE summing amplifiers) needs to be less than one UI or 100 ps at 10 Gb/s. This timing requirement is eased by using the hybrid speculative/dynamic feedback DFE architecture (Fig. 8) originally developed for the 6.4 Gb/s design [2]. Analog settling time requirements are eliminated for the first feedback tap (H1) by using two-path speculation or loop unrolling [20], while the half-rate clocking allows 2 UI (200 ps) for the H2-H5 dynamic feedback signals to be accurately established at the slicer inputs. Reducing the response times of the H2-H5 feedback taps to meet this 2 UI requirement without an excessive increase in power dissipation necessitated some improvements to the 6.4 Gb/s DFE.

The first improvement is the adoption of zero-skew clock distribution to all of the slicers and latches in the DFE. Though not discussed in [2], the latches (L) used to delay the data for the H3-H5 feedback taps are clocked in the 6.4 Gb/s design with a buffered version of the clock which triggers the decision-making slicers. This buffering reduces the load on the CML buffer supplying the C2 clock to the DFE, but the resulting skew between the slicer and latch clocks delays the H3-H5 feedback signals. In this design, such skew is eliminated by clocking all slicers and latches with the same clock signal, driven by a large CML buffer. The most important improvement, however, is the decrease in settling time of the analog summers. In accord with common practice [2][3], the H2-H5 feedback signals are added to the data input by pulling

weighted currents from the positive or negative output of a resistively loaded differential amplifier. Fast settling requires small RC time constants on these output nodes. Powering up circuits to reduce R is a very inefficient solution, as the larger devices and wider wires needed to handle the higher currents increase C significantly, so large increases in power yield only small speed improvements. In this work, most of the reduction in RC time constants was achieved with a new floorplan of the DFE summers which minimizes wiring parasitics on the critical nodes. Figure 9 compares the original and improved layout floorplans for the odd half of the DFE. (The even half is similar, except that the offset summer used for the amplitude monitor is replaced with a load-balancing dummy amplifier.) With lower capacitance on the H2-H5 summer outputs, 2% settling of the feedback signals is achieved within 2 UI.

The edge samples processed by the CDR loop are taken on the non-DFE equalized data signal. The CDR logic converts the data and edge samples into early and late signals which are digitally filtered to generate increment/decrement signals that control the phase rotators. The CDR has a tracking bandwidth of about 9 MHz and can handle frequency offsets up to +/- 4000 ppm. The digitally controlled phase rotators must be precise not to degrade the timing position of the recovered clock. The phase rotator (Fig. 10(a)) is driven by two differential C2 quadrature clock phases, I and Q. The circuit selects polarity of the phases (quadrant selection) and then interpolates between them to generate 16 phase positions within each quadrant for a total of 64 on a 360 degree (2 UI) circle. The interpolator uses a current-steering DAC which supplies tail currents to the differential pairs processing I and Q phases and having a common resistive load. The 64-point phase constellation of the rotator is diamond-shaped (Fig. 10(b)), reflecting constant total interpolator tail current. The current-steering DAC employs 15 switched cells plus two fixed (non-switched) cells of half-size to realize 16 different interpolation ratios ranging from 0.5:15.5 to 15.5:0.5. Rotator settling time is improved by never applying zero tail current to the interpolator branches. When the rotator steps across the quadrant boundary, the interpolation

ratio stays constant, so only the polarity of one input phase needs to be switched. The 15 steering DAC cells are not uniform; their relative sizing, with the largest cells being switched near the quadrant boundaries, is optimized for the most linear relationship between digital control code and rotator output phase. Generally, the need for non-uniform DAC sections arises from the non-circular (diamond) shape of the phase constellation. Uniform DAC sections produce uniform distribution of phase states on the diamond sides (as shown), but non-uniform distribution of phase angles, since angles near the middle of the quadrant exceed those near quadrant boundaries by a factor of 2. (Compare the angular widths of the two shaded wedges in Fig. 10(b).) However, simulations of the phase rotator circuit show that actual non-uniformity is much weaker, and the non-uniformity of the DAC sections has been designed accordingly. In fast process corners, the inputs to the rotator may be closer to square waves than sine waves. A pair of slew-rate-control buffers in front of the rotator core reshape these signals to be more sinusoidal to ensure adequate overlap of the clock edges being interpolated.

The phase rotator has been evaluated in a separate breakout site and demonstrated high linearity within a 2 to 6 GHz frequency range. Fig. 11 shows superimposed oscilloscope traces of all 64 rotator states and corresponding measurements of rotator step size and integral non-linearity at 2 GHz and 6 GHz. At both frequencies, the measured min-to-max step ratio of the rotator is better than 1:2.

C. PLL

The block diagram of the PLL used to supply C2 clocks to the transmitters and receivers is shown in Fig. 12. To achieve low phase noise, the PLL employs a band-switched LC voltage-controlled oscillator (VCO) operating at full-rate (10 GHz). The PLL runs off its own 1.8 V supply. The PLL output drives a CML divide-by-2 stage to generate a C2 clock with low DCD, which is then distributed to the receivers or transmitters. When used in the link demonstrator test chip, the PLL and clock distribution buffers dissipate 100 mW. Aside from the higher center

frequency, this PLL is similar to the one used in the 6.4 Gb/s core, whose details are covered in [2]. Table I presents performance data measured on two 10 GHz PLLs of identical design. Since the bandwidth (2 MHz) of the PLL's jitter transfer function is smaller than the tracking bandwidth (9 MHz) of the CDR in the receiver, the CDR is able to track most of the reference jitter not filtered by the PLL.

IV. Serial Link Experiments

A. *Link Demonstrator Test Chip*

The transceiver components were assembled into a link demonstrator test chip (Fig. 13) in order to conduct various serial link experiments. The chip was fabricated in a 90-nm CMOS process and then attached with controlled collapse chip connection (C4) technology to a flip-chip plastic ball grid array (FCPBGA) module, which is detailed in [5]. The test chip includes two receiver (Rx) pairs and two transmitter (Tx) pairs, with each pair being clocked either internally by a PLL or externally by a full-rate (10 GHz) sinusoidal source. External clocking allows testing at frequencies outside the PLL tuning range, but it was found that crosstalk from the 10 GHz external signal hindered serial link performance. All the results presented below have been obtained with PLL-based clocking. The chip is configured through a parallel port interface.

B. *Experimental Results with Single Chip Evaluation Board*

Mounting a single chip module on a socketed evaluation board (Fig. 14(a)) allows characterization of the individual transmitters and receivers, as well as serial link testing with a transmitter and receiver connected together in a loop-back configuration. The high-speed inputs and outputs of the transmitters and receivers are brought out to cables with SMP connectors. S-parameter measurements (Fig. 14(b)) show that the combined loss of module, evaluation board, and 24-in of cable is 4 dB at 5 GHz (8 dB in loop-back testing). The ISI due to this channel loss is readily observed in the transmit eye diagram at 10 Gb/s (top of Fig. 15). As shown in the

figure, the measured eye diagram is similar to that calculated with the link simulation tool from the S-parameter data. (The simulated eye does not include random jitter.) Setting the FFE to normalized coefficients of [0, 0.85, -0.15, 0] equalizes the channel (bottom of Fig. 15).

Receiver performance was studied with input data directly supplied from a pseudo-random bit sequence generator. As an indicator of operating margins at 10 Gb/s, Fig. 16 shows plots of the receiver bit error rate (BER) as a function of data sampling time (often referred to as “bathtub curves”). To make these measurements, the CDR loop is frozen, and the phase rotator generating the data sampling clock is manually swept over its positions (32 steps/UI). The DFE tap weights are also held at their previously adapted values. With the DFE off (zero tap weights), the horizontal eye openings at a BER of 10^{-9} are 56% and 50% for input amplitudes of 1200 mVppd and 200 mVppd, respectively. These horizontal eye openings increase to 68.75% and 62.5% upon application of the DFE, which helps equalize the channel loss of the setup.

It should be noted that the bathtub curves of Fig. 16 are pessimistic estimates of the receiver’s real performance. In normal operation, low frequency jitter from the receive PLL is tracked out by the CDR loop. With the CDR frozen, PLL jitter is not tracked out and directly impacts the estimates of horizontal eye opening. Fig. 17 shows a histogram of the phase rotator’s position in normal operation. The phase tracking provided by the CDR is a significant fraction of a UI, so the pessimism of the measured bathtub curves is substantial. All of the bathtub curves and estimates of horizontal eye opening presented in this paper are similarly pessimistic.

Connecting a transmitter’s outputs to a receiver’s inputs through a 16-in Tyco legacy backplane with HM-Zd edge connectors (loop-back configuration) is a demanding test of the transceiver’s equalization capabilities. S-parameter measurements (Fig. 18(a)) show that the combined loss of the single-chip evaluation board setup, the 16-in backplane, and the cables (12-in from evaluation board to backplane, and 12-in from backplane to evaluation board) is 33.5 dB at 5 GHz. The FFE tap weights are initially set to the values predicted by the high-level link

simulation tool but are then fine tuned empirically for best link performance. After the DFE adaptation has converged, the bathtub curve (Fig. 18(b)) is measured in the manner explained above. The horizontal eye opening of the equalized signal is 22% at a BER of 10^{-9} , and error-free operation is obtained at eye center.

C. *Chip-to-Chip Link Experiments*

Directly soldering two modules on a board (Fig. 19) allowed serial link experiments from chip-to-chip. Boards fabricated in both conventional and advanced technologies were used to compare performance so that the benefits of improved interconnect could be assessed. The channels under test also differed in the trace length between chips and in the number and types of via-stubs. The first (#1) 10-in line was fabricated in conventional board technology, using Nelco 4000-13 material. The presence of two 3.8 mm via stubs along the line results in signal reflections, and the frequency response (Fig. 20(a)) exhibits a deep notch near 8 GHz. At 5 GHz, the channel loss is 12 dB. The second (#2) 10-in line was fabricated in an advanced technology, using lower loss APPE material. The sub-composite construction of this board technology (described more fully in [5]) allows the via stub length to be reduced to 1.8 mm. With the via stub resonant frequency more than doubled, the channel response (Fig. 20(b)) has no deep notches in the frequency range of interest, and the loss at 5 GHz is 10 dB. The 15-in and 20-in lines were also fabricated in this advanced technology, though the 15-in line was designed to be a difficult channel, with four 3.8 mm and two 1.8 mm stubs. The frequency responses of these channels are plotted in Fig. 20 (c) and (d).

The operating margins of these chip-to-chip links were evaluated for three cases of equalization: FFE only, DFE only, and both FFE and DFE. Figure 21 shows how the choice of equalization affects the link margins at 10 Gb/s. With the highest loss at 5 GHz and bad reflections, the 15-in line is the most difficult to equalize, and only a combination of FFE and DFE is able to achieve a low BER. The first 10-in line also requires both FFE and DFE for a low

BER. The other two lines (without bad reflections) are easier to equalize, and a low BER is achieved for all three cases of equalization. With both FFE and DFE employed, the horizontal eye openings approach 60%.

V. Summary

This paper has presented a 90-nm CMOS transceiver for chip-to-chip communications at 10 Gb/s. Using a 4-tap FFE and 5-tap DFE enables error-free NRZ signaling over channels with more than 30 dB of loss. Design features such as half-rate clocking of the transmitter and receiver improve power efficiency. One transmitter/receiver pair and one PLL consume 300 mW. The effectiveness of the FFE/DFE equalization has been studied in serial link experiments using different types of channels. These experiments not only demonstrate the performance of the transceiver but also highlight the importance of reducing reflections due to structures such as via stubs, as links with bad reflections are often more difficult to equalize than those with more loss. The use of more advanced board technologies is one clear solution to the problem.

Acknowledgment

The authors wish to thank M. Sorna, S. Zier, P. Metty and K. Heilmann from IBM Fishkill for valuable advice and assistance, and M. Oprysko and M. Soyuer from IBM Yorktown for technical and managerial support of this project.

References

- [1] *Common Electrical I/O (CEI) – Electrical and Jitter Interoperability Agreements for 6+ Gbps and 11+ Gbps I/O*, Optical Interconnect Forum, IA # OIF-CEI-02.0, Feb. 2005.
- [2] T. Beukema, M. Sorna, K. Selander, S. Zier, B. L. Ji, P. Murfet, J. Mason, W. Rhee, H. Ainspan, B. Parker, and M. Beakes, “A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization,” *IEEE J. Solid-State Circuits*, vol. 40, pp. 2633-2645, Dec. 2005.

- [3] R. Payne, P. Landman, B. Bhakta, S. Ramaswamy, S. Wu, J. D. Powers, M. U. Erdogan, A-L. Yee, R. Gu, L. Wu, Y. Xie, B. Parthasarathy, K. Brouse, W. Mohammed, K. Heragu, V. Gupta, L. Dyson, and W. Lee, "A 6.25-Gb/s binary transceiver in 0.13- μ m CMOS for serial data transmission across high loss legacy backplane channels," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2646-2657, Dec. 2005.
- [4] K. Krishna, D. A. Yokoyama-Martin, A. Caffee, C. Jones, M. Loikkanen, J. Parker, R. Segelken, J. L. Sonntag, J. Stonick, S. Titus, D. Weinlader, and S. Wolfer, "A multigigabit backplane transceiver core in 0.13- μ m CMOS with a power-efficient equalization architecture," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2658-2666, Dec. 2005.
- [5] L. Shan, Y. Kwark, P. Pepeljugoski, M. Meghelli, T. Beukema, C. Baks, J. Trehwella, and M. Ritter, "Design, analysis, and experimental verification of an equalized 10Gbps link," in *DesignCon 2006*, Santa Clara, CA, Feb. 2006.
- [6] 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, pp. 2121-2130, Dec. 2003.
- [7] R. Farjad-Rad, C.-K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.3- μ m CMOS 8-Gb/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 757-764, May 2000.
- [8] 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- μ m CMOS," *IEEE J. Solid-State Circuits*, vol. 38, pp. 436-443, Mar. 2003.
- [9] V. Stojanovic, "Channel-limited high-speed links: modeling, analysis and design," Ph.D. dissertation, Stanford University, Stanford, CA, 2004.
- [10] N. Krishnapura, M. Barazande-Pour, Q. Chaudhry, J. Khoury, K. Lakshmikumar, and A. Aggarwal, "A 5Gb/s NRZ transceiver with adaptive equalization for backplane transmission," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 60-61.

- [11] T. Schafbauer et al., "Integration of high-performance, low-leakage, and mixed signal features into a 100nm CMOS technology," in *Symp. VLSI Technology Dig. Tech. Papers*, Honolulu, HI, June 2002, pp. 62-63.
- [12] G. Scott, J. Lutze, M. Rubin, F. Nouri, and M. Manley, "NMOS drive current reduction caused by transistor layout and trench isolation induced stress," in *Int. Electron Devices Meeting Tech. Dig.*, Washington, DC, Dec. 1999, pp. 827-830.
- [13] M. Anis, M. Allam, and M. Elmasry, "Impact of technology scaling on CMOS logic styles," *IEEE Trans. Circuits Syst. II*, vol. 49, pp. 577-588, Aug. 2002.
- [14] A. Rylyakov and S. Rylov, "A low power 10 Gb/s serial link transmitter in 90-nm CMOS," in *IEEE Compound Semicond. IC Symp. (CSICS) Tech. Dig.*, Palm Springs, CA, Oct./Nov. 2005, pp. 189-191.
- [15] M. Meghelli, "A 43-Gb/s full-rate clock transmitter in 0.18- μ m SiGe BiCMOS technology," *IEEE J. Solid-State Circuits*, vol. 40, pp. 2046-2050, Oct. 2005.
- [16] C. Menolfi, T. Toifl, R. Reutemann, M. Ruegg, P. Buchmann, M. Kossel, T. Morf, and M. Schmatz, "A 25Gb/s PAM4 transmitter in 90nm CMOS SOI," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 72-73.
- [17] S. Galal and B. Razavi, "Broadband ESD protection circuits in CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, pp. 2334-2340, Dec. 2003.
- [18] S. Sidiropoulos and M. A. Horowitz, "A semidigital dual delay-locked loop," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1683-1692, Nov. 1997.
- [19] J. Chen and B. Shi, "Novel constant transconductance references and the comparisons with the traditional approach," in *Proc. Southwest Symp. Mixed-Signal Design*, Las Vegas, NV, Feb. 2003, pp. 104-107.
- [20] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, pp. 711-717, June 1991.

List of Table Captions:

Table I. Performance summary of 10 GHz PLL.

List of Figure Captions:

Fig. 1. Backplane channel characteristics. (a) Backplane/line card application. (b) Channel impulse response. (c) Channel frequency response.

Fig. 2. Impulse and frequency responses of Bessel filter channel.

Fig. 3. Simulated PAM-4 and NRZ eye diagrams for Bessel filter channel.

Fig. 4. Pre-cursors, main cursor, and post-cursors of Bessel filter channel. (a) Channel pulse response sampled at 12.5 GHz. (b) Discrete-time representation. (c) Channel frequency response before and after post-cursor cancellation by DFE.

Fig. 5. Transmitter with 4-tap FFE.

Fig. 6. Receiver block diagram.

Fig. 7. (a) Two-path VGA used in present work. (b) Two-path VGA described in [2]. (c) Gain control bit settings. (d) Measured receiver dc gain as function of VGA setting.

Fig. 8. Half-rate DFE architecture with first feedback tap realized by speculation.

Fig. 9. Layout floorplans of odd DFE half. (a) Original floorplan. (b) Improved floorplan.

Fig. 10. (a) Phase rotator schematic. (b) Phase constellation of rotator with uniform DAC sections.

Fig. 11. Measured phase rotator performance at frequencies of 2 GHz and 6 GHz. (a) Oscilloscope traces of rotator output. (b) Rotator step size and integral non-linearity (INL).

Fig. 12. PLL block diagram.

Fig. 13. Layouts of transceiver components and their arrangement on the link demonstrator test chip.

Fig. 14. (a) Photograph of single chip evaluation board. (b) S-parameter data showing combined loss of package (without socket), board, and 24-in of cable.

Fig. 15. Measured and simulated transmit eye diagrams at 10 Gb/s before (top) and after (bottom) applying one post-cursor of FFE.

Fig. 16. Measured bathtub curves of receiver.

Fig. 17. Histogram of phase rotator position with CDR active.

Fig. 18. Loop-back testing through 16-in Tyco legacy channel. (a) Channel frequency response. (b) Measured bathtub curve.


Fig. 19. Links used for chip-to-chip experiments.

Fig. 20. Frequency responses of channels used in chip-to-chip experiments. (a) 10-in (#1) line. (b) 10-in (#2) line. (c) 15-in line. (d) 20-in line.


Fig. 21. Measured horizontal eye openings of four chip-to-chip links with different equalizations applied.

Table I: Performance Summary of 10 GHz PLL


| PLL Parameter | TxPLL | RxPLL |
|---|--------------|--------------|
| Min freq (GHz) | 8.98 | 8.96 |
| Max freq (GHz) | 13.54 | 13.47 |
| Mean freq (GHz) | 11.26 | 11.22 |
| Lock range (GHz) | 4.56 | 4.52 |
| | +/-20.2% | +/-20.1% |
| Fine tune hold range | 5.8% | 5.8% |
| Phase noise @ 10 MHz offset (dBc/Hz, quarter rate clock) | -117.8 | -117.7 |
| Jitter, 1 MHz-100 MHz (ps rms) | 1.5 | 1.4 |
| Jitter, $f_c/1667$ -100 MHz (ps rms) | 0.64 | 0.64 |
| -3 dB bandwidth of jitter transfer function (MHz) | 2 | 2 |



(a)



(b)



(c)

Fig. 1. Backplane channel characteristics. (a) Backplane/line card application. (b) Channel impulse response. (c) Channel frequency response.


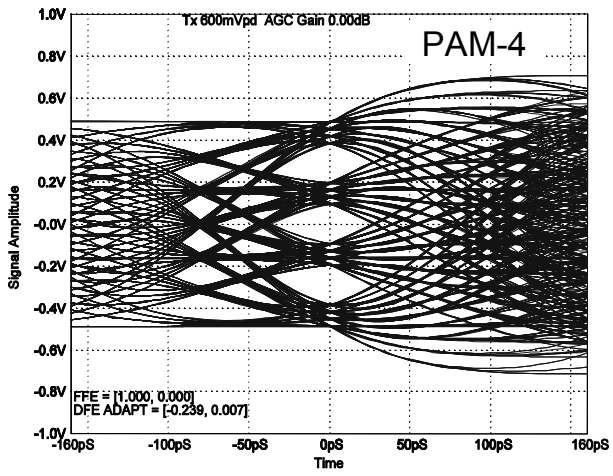
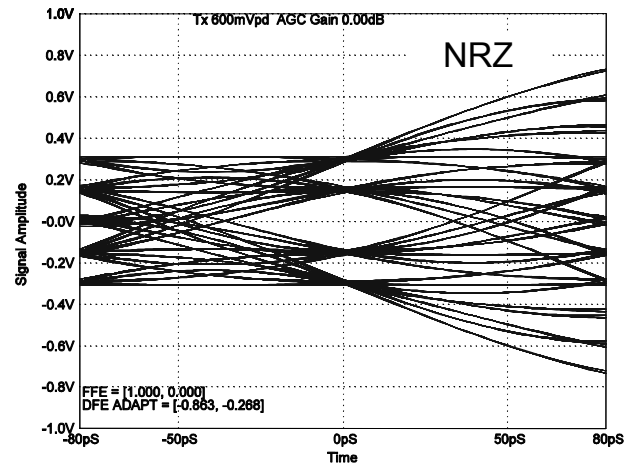


Fig. 2. Impulse and frequency responses of Bessel filter channel.



Vert. Eye = 137 mV, Hor. Eye = 50 ps



Vert. Eye = 265 mV, Hor. Eye = 60 ps

Fig. 3. Simulated PAM-4 and NRZ eye diagrams for Bessel filter channel.

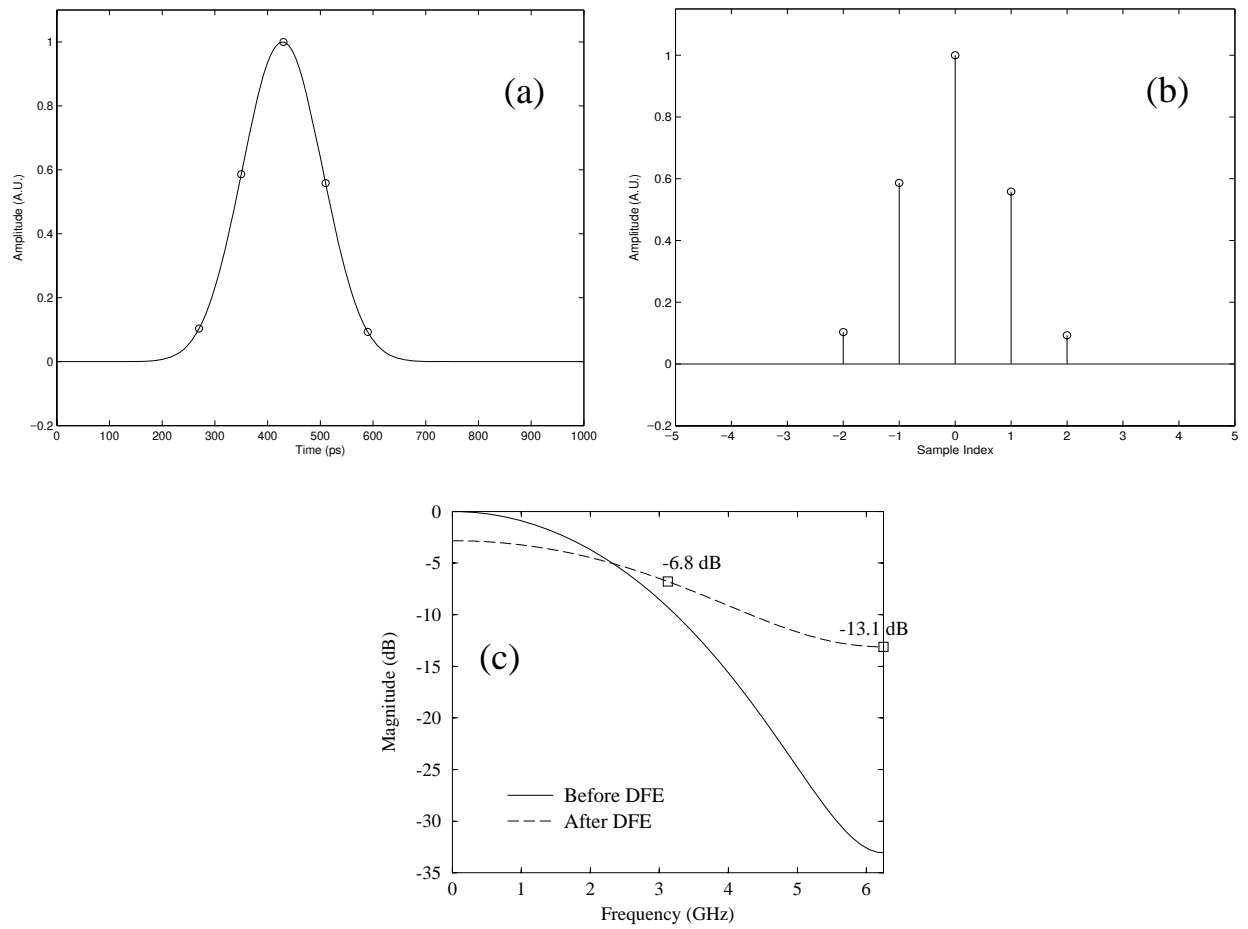


Fig. 4. Pre-cursors, main cursor, and post-cursors of Bessel filter channel. (a) Channel pulse response sampled at 12.5 GHz. (b) Discrete-time representation. (c) Channel frequency response before and after post-cursor cancellation by DFE.

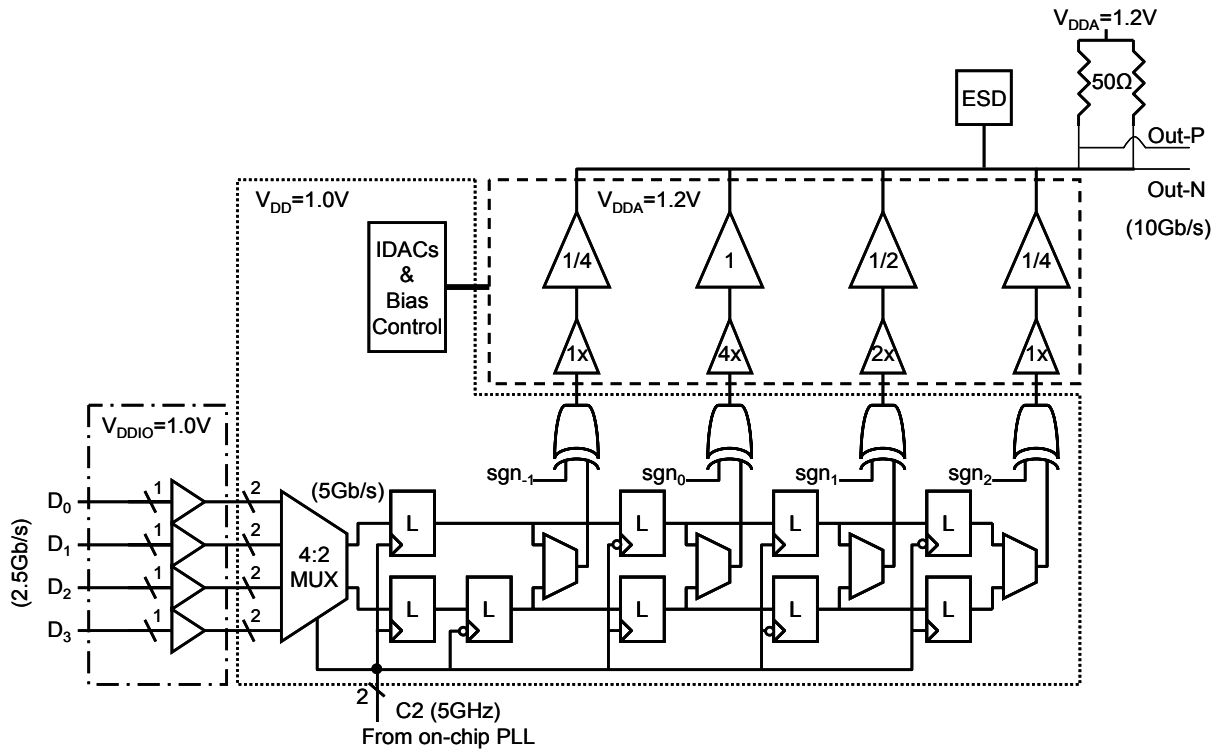


Fig. 5. Transmitter with 4-tap FFE.

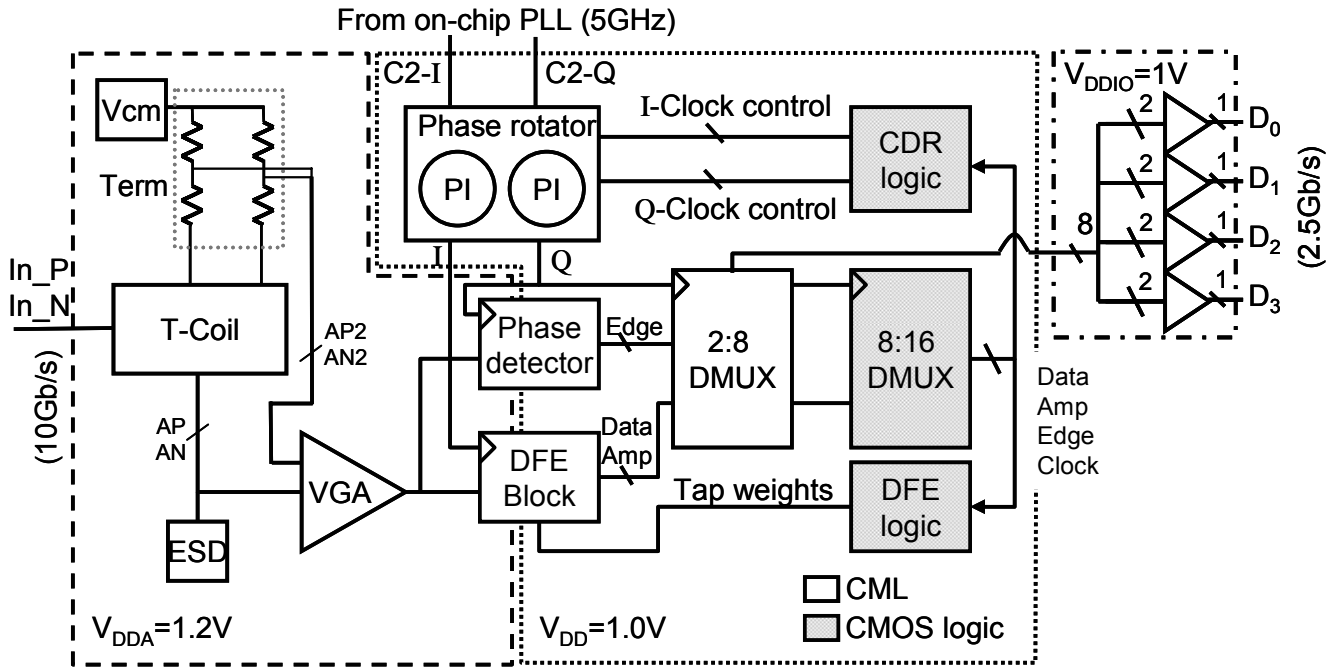




Fig. 6. Receiver block diagram.



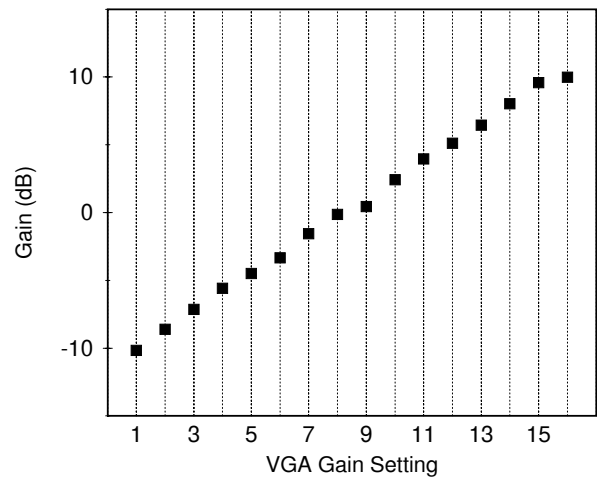
(a)



(b)

| VGA Gain Setting | FULL (HALF) | B0...B6 | C0...C6 |
|------------------|-------------|---------|---------|
| 1 | 0 | 0000000 | 0000000 |
| 2 | 0 | 0000000 | 1000000 |
| 3 | 0 | 0000000 | 1100000 |
| 4 | 0 | 0000000 | 1110000 |
| 5 | 0 | 0000000 | 1111000 |
| 6 | 0 | 0000000 | 1111100 |
| 7 | 0 | 0000000 | 1111110 |
| 8 | 0 | 0000000 | 1111111 |
| 9 | 1 | 0000000 | 1111111 |
| 10 | 1 | 1000000 | 1111111 |
| 11 | 1 | 1100000 | 1111111 |
| 12 | 1 | 1110000 | 1111111 |
| 13 | 1 | 1111000 | 1111111 |
| 14 | 1 | 1111100 | 1111111 |
| 15 | 1 | 1111110 | 1111111 |
| 16 | 1 | 1111111 | 1111111 |

(c)



(d)

Fig. 7. (a) Two-path VGA used in present work. (b) Two-path VGA described in [2]. (c) Gain control bit settings. (d) Measured receiver dc gain as function of VGA setting.

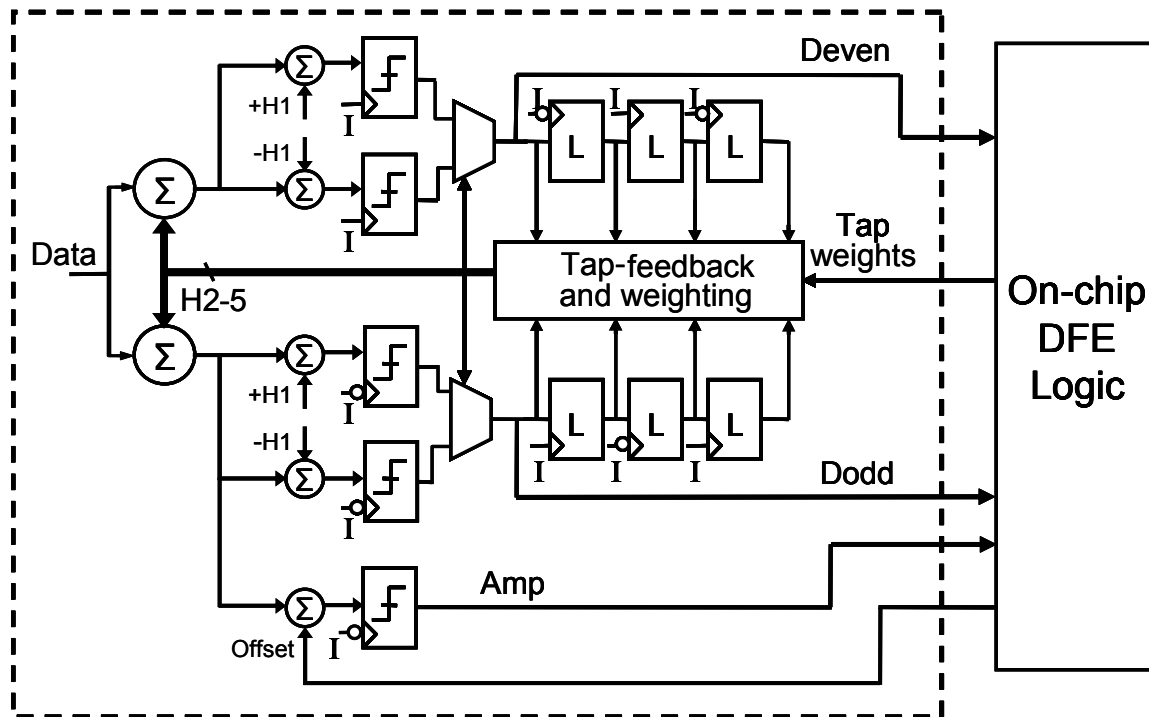


Fig. 8. Half-rate DFE architecture with first feedback tap realized by speculation.

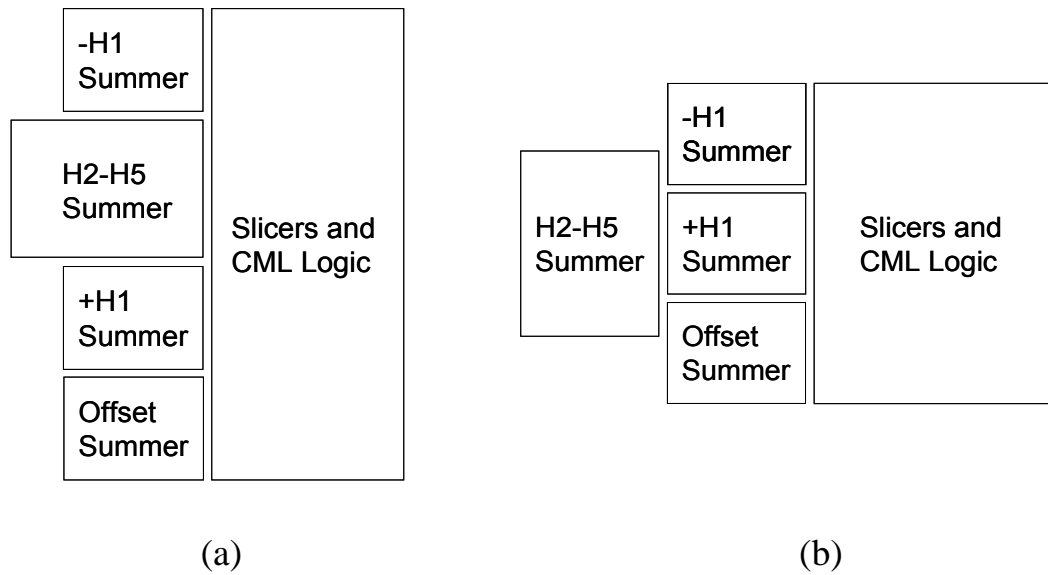


Fig. 9. Layout floorplans of odd DFE half. (a) Original floorplan. (b) Improved floorplan.

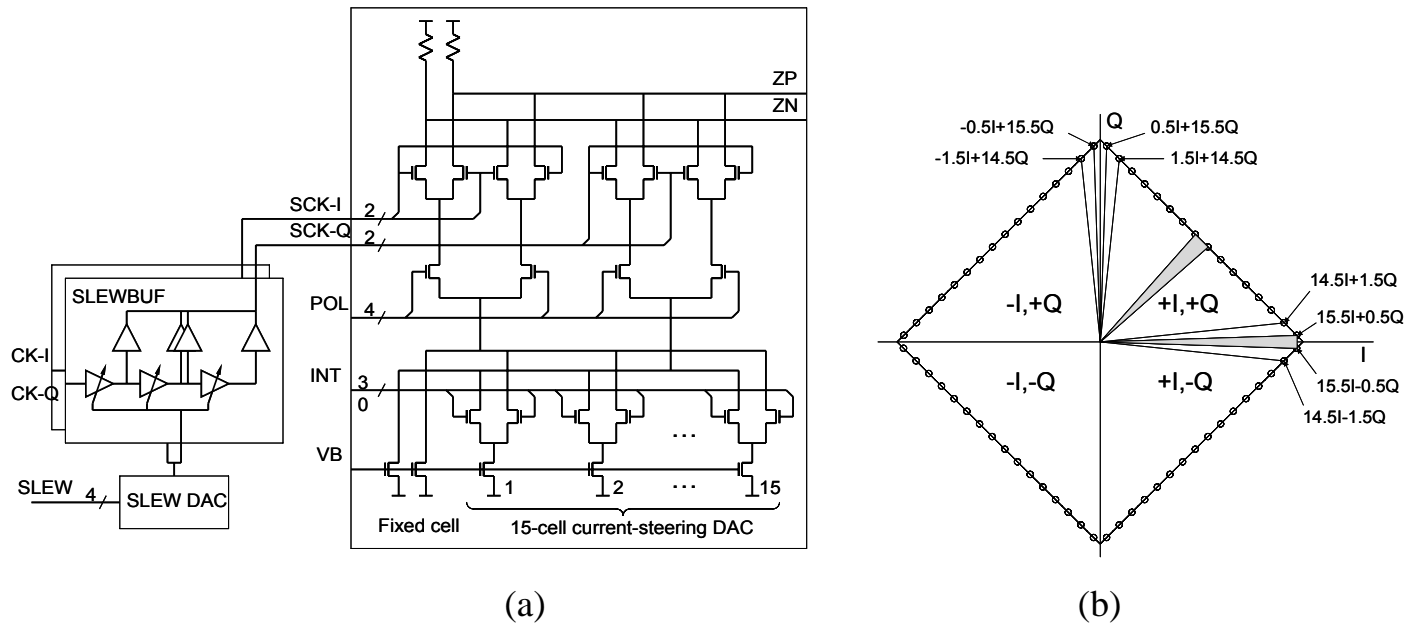


Fig. 10. (a) Phase rotator schematic. (b) Phase constellation of rotator with uniform DAC sections.

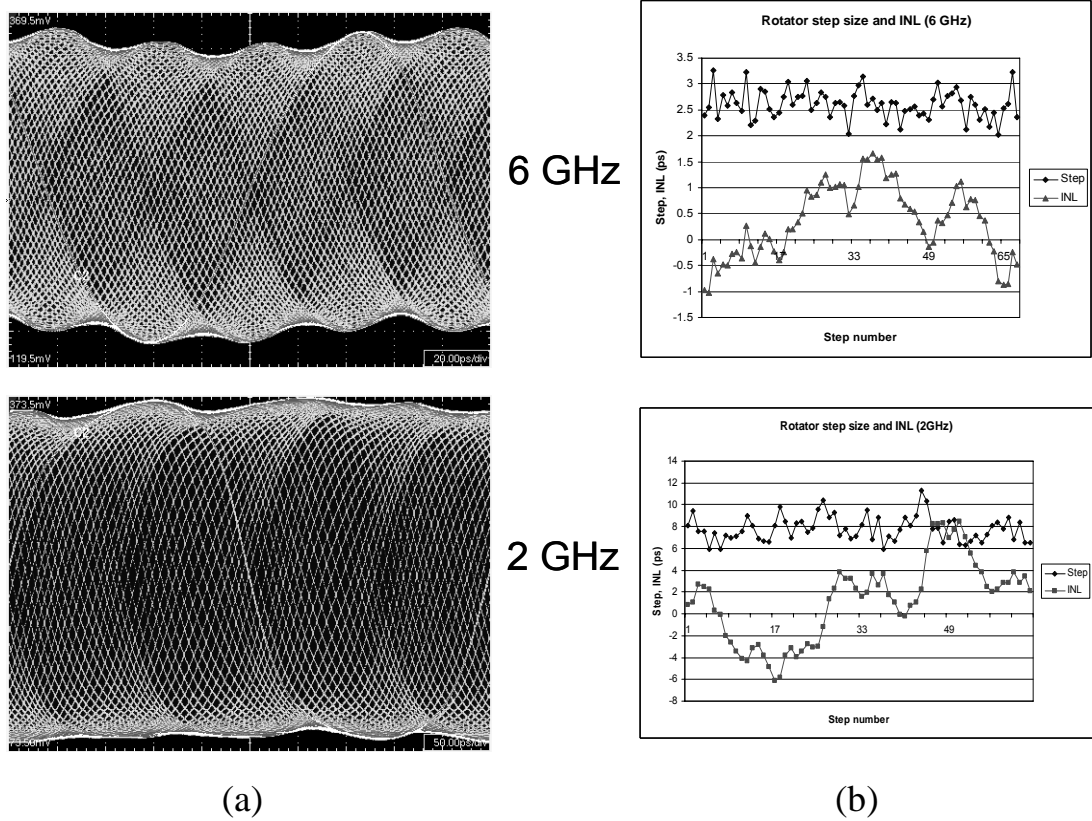


Fig. 11. Measured phase rotator performance at frequencies of 2 GHz and 6 GHz. (a) Oscilloscope traces of rotator output. (b) Rotator step size and integral non-linearity (INL).

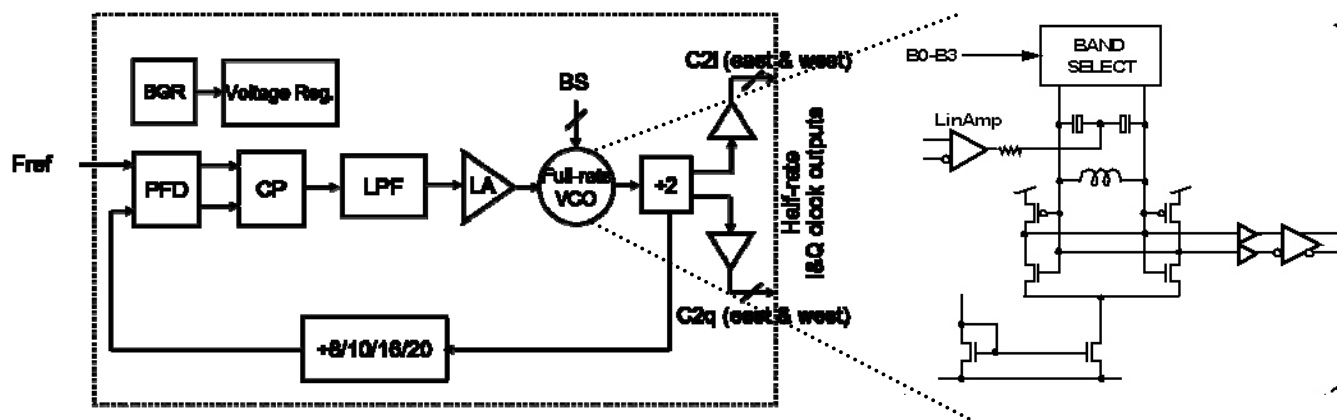


Fig. 12. PLL block diagram.

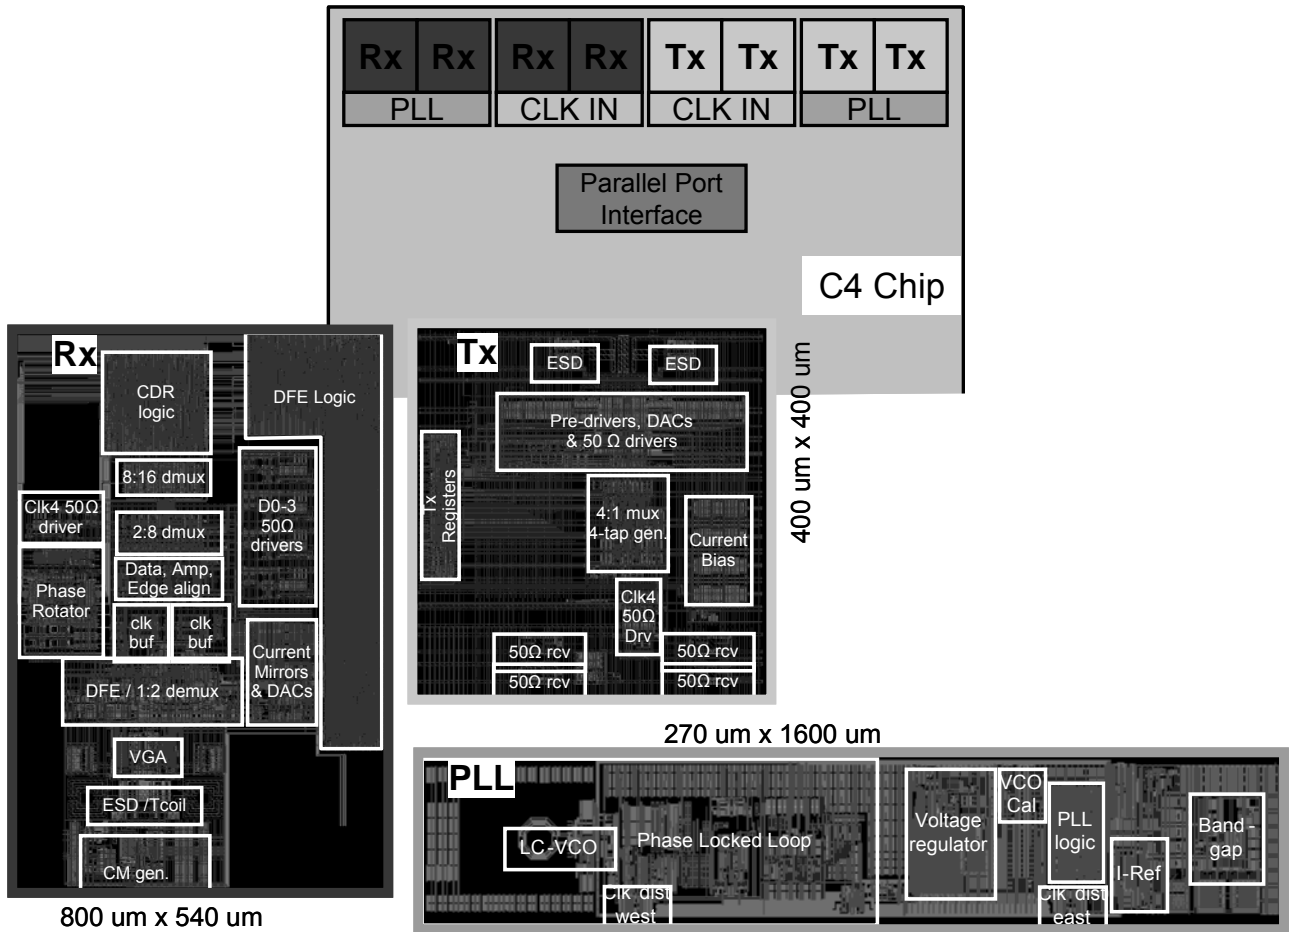

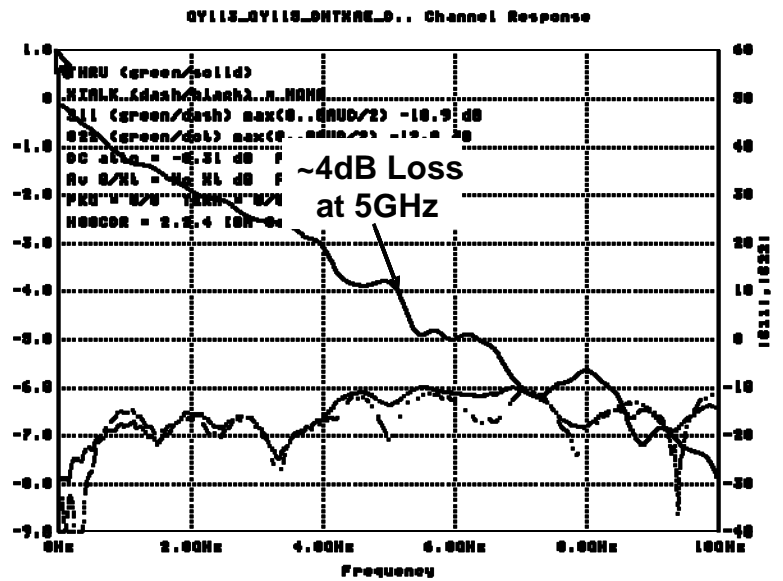


Fig. 13. Layouts of transceiver components and their arrangement on the link demonstrator test chip.



(a)



(b)

Fig. 14. (a) Photograph of single chip evaluation board. (b) S-parameter data showing combined loss of package (without socket), board, and 24-in of cable.




Fig. 15. Measured and simulated transmit eye diagrams at 10 Gb/s before (top) and after (bottom) applying one post-cursor of FFE.




Fig. 16. Measured bathtub curves of receiver.

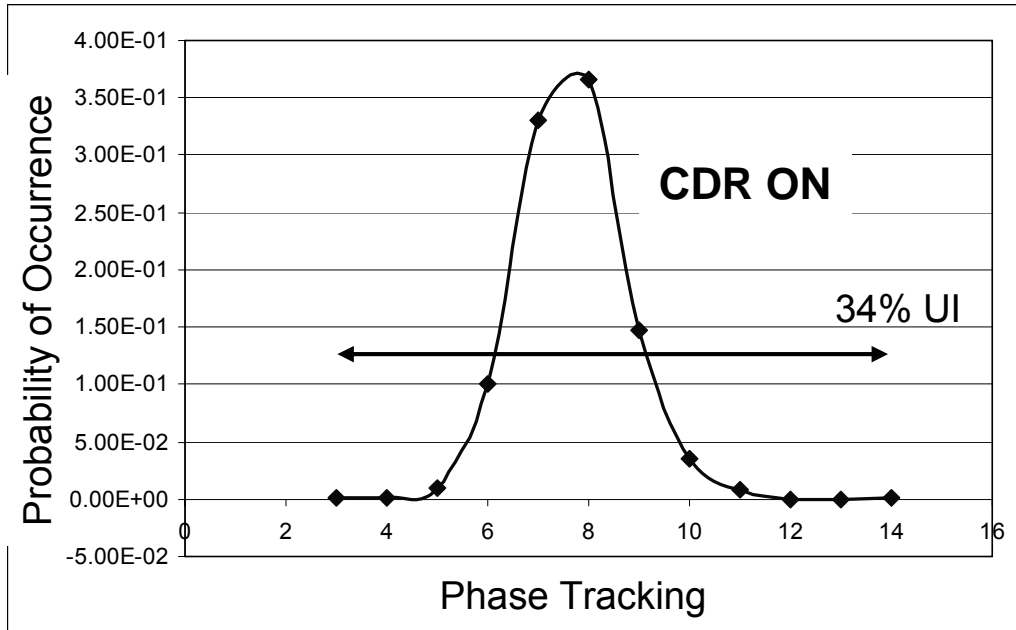

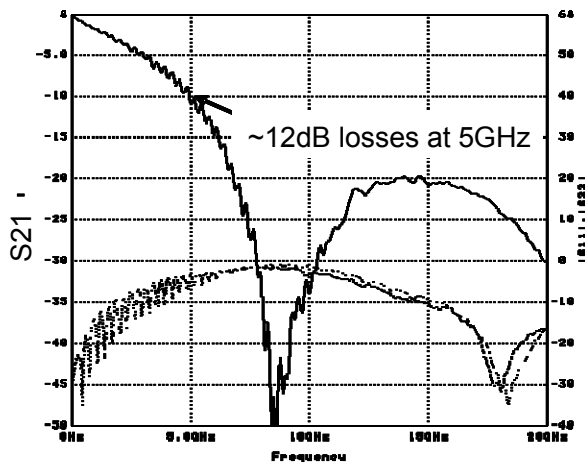


Fig. 17. Histogram of phase rotator position with CDR active.




| Trace Length | 5GHz losses (Tx module + board trace + Rx module) | Number of vias 3.8mm via stub / 1.8mm via stub / 1.8mm via through |
|--------------|--|---|
| 10" (#1) | 12dB | 2 / 0 / 0 |
| 10" (#2) | 10dB | 0 / 2 / 0 |
| 15" | 25dB | 4 / 2 / 0 |
| 20" | 15dB | 0 / 0 / 2 |

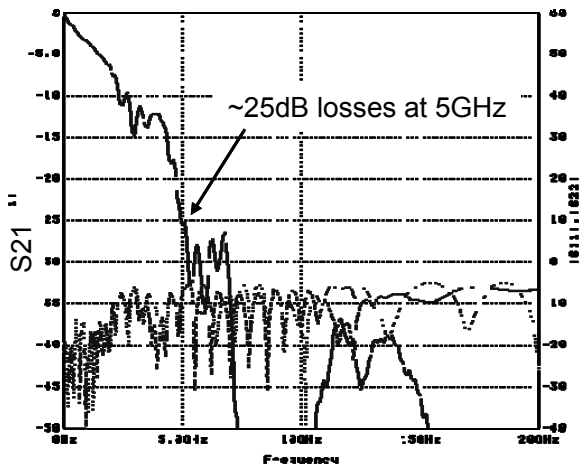
Fig. 19. Links used for chip-to-chip experiments.



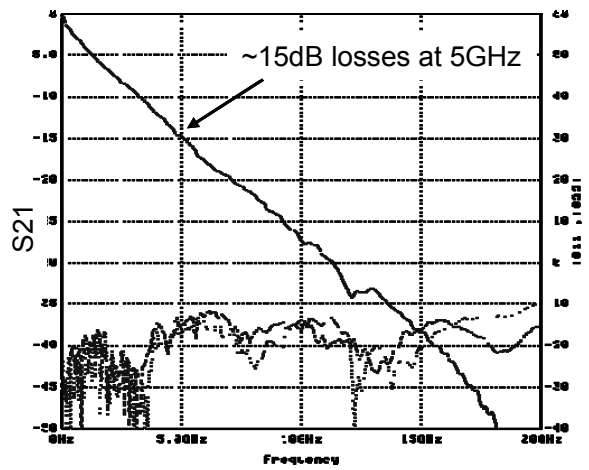
(a)



(b)



(c)



(d)

Fig. 20. Frequency responses of channels used in chip-to-chip experiments. (a) 10-in (#1) line. (b) 10-in (#2) line. (c) 15-in line. (d) 20-in line.




Fig. 21. Measured horizontal eye openings of four chip-to-chip links with different equalizations applied.