A 14-mW 6.25-Gb/s Transceiver in 90-nm CMOS

John Poulton, *Senior Member, IEEE*, Robert Palmer, Andrew M. Fuller, Trey Greer, *Member, IEEE*, John Eyles, William J. Dally, *Fellow, IEEE*, and Mark Horowitz, *Fellow, IEEE*

***Abstract*—This paper describes a 6.25-Gb/s 14-mW transceiver in 90-nm CMOS for chip-to-chip applications. The transceiver employs a number of features for reducing power consumption, including a shared *LC*-PLL clock multiplier, an inductor-loaded resonant clock distribution network, a low- and programmable-swing voltage-mode transmitter, software-controlled clock and data recovery (CDR) and adaptive equalization within the receiver, and a novel PLL-based phase rotator for the CDR. The design can operate with channel attenuation of −15 dB or greater at a bit-error rate of 10^{-15} or less, while consuming less than 2.25 mW/Gb/s per transceiver.**

***Index Terms*—Low power, transceiver.**

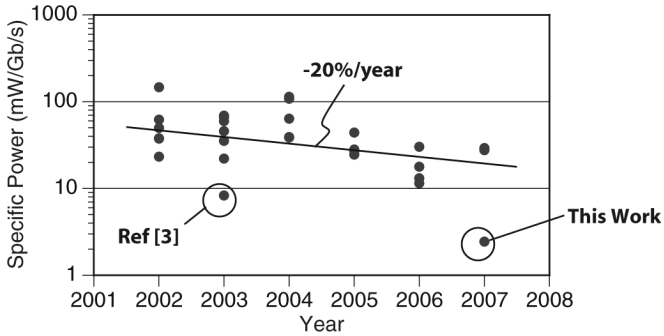

Fig. 1. Specific power of recently published serial links.

I. Introduction

PRESENT-DAY computing systems require very high off-chip communications bandwidth, and high-speed serial links for chip-to-chip interconnect are now ubiquitous. Many of these links have channels that present modest attenuation, crosstalk, and reflections. Present-day chip-to-chip serial links, however, have generally evolved from backplane transceivers that must deal with much more difficult channels, and they often dissipate far more power than necessary for short-haul chip-to-chip links, typically about 20 mW/Gb/s.

Few published designs, for example, [1]–[4], have explored the design space for chip-to-chip links in which minimizing power is the primary goal of the design. The specific power (W/Gb/s) of various recently published serial links is shown in Fig. 1. While the transceivers listed span a fairly wide range of applications and specific power and include both standards-based and noncompliant designs, the trend line shows that power efficiency is improving at about 20% per year. The transceiver described in this paper is about an order of magnitude more power-efficient than most contemporary links and represents a significant improvement over the recent trend.

In this paper, we present a transceiver in 90-nm CMOS designed to meet the typical needs of short-haul chip-to-chip interconnect while dissipating very low power. It is not intended to comply with an existing standard, since few present-day standards are suitable for this problem domain. The transceiver uses differential signaling since this method offers a 2× power advantage over single-ended signaling. It employs various power reduction techniques, including an *LC*-based PLL for reference clock multiplication and jitter reduction, a resonant clock distribution network, an on-chip-regulated low-swing voltage-mode transmitter, a low-power PLL-based phase rotator for the receiver clocks, and software-based CDR and adaptive equalization. To validate these ideas, we implemented a four-link test chip, each link operating at 6.25 Gb/s.

Section II describes the overall design of the transceiver and clock multiplier. Details of the transmitter are described in Section III, the receiver is described in Section IV, and the clock multiplier PLL and resonant clock distribution are detailed in Section V. Section VI summarizes the measurement results from the test chip.

II. Transceiver Overview

Fig. 2 shows the overall design of the transceiver. A clock multiplier module contains an *LC*-based PLL that multiplies a 195.3125 MHz reference clock up to 3.125 GHz for distribution to the array of four transceivers (only one transmitter and receiver are shown). The coarse frequency of the oscillator and the output swing and common-mode voltage of the oscillator and clock distribution driver are programmable. The clock multiplier module contains an "electrical measurement unit" (EMU), an 8-bit analog-to-digital converter (ADC), that provides a way to measure various voltages within the module.

The half-bit-rate clock is distributed over an inductor-loaded network. The inductors are chosen to resonate the network at the clock frequency and thereby increase the impedance of the network and lower the power required for clock distribution.

The transceiver is designed to operate at a single speed of 6.25 Gb/s. Providing a wide range of data rates in a serial link presents a number of technical challenges and always ends up burning more power than a single-rate link of equivalent performance. Principal among these challenges is the difficulty of providing a wide tuning range in an *LC*-PLL (in practice, this is limited to 2:1, since lower rates can be more easily supported by dividing the clock).

Manuscript received April 18, 2007; revised August 29, 2007.

J. Poulton, R. Palmer, A. M. Fuller, and T. Greer are with Rambus, Inc., Chapel Hill, NC 27514 USA (e-mail: jp@rambus.com).

J. Eyles was with Rambus, Inc., Chapel Hill, NC 27514 USA. He is now an independent consultant.

W. J. Dally and M. Horowitz are with Stanford University, Stanford, CA 94305 USA.

Digital Object Identifier 10.1109/JSSC.2007.908692

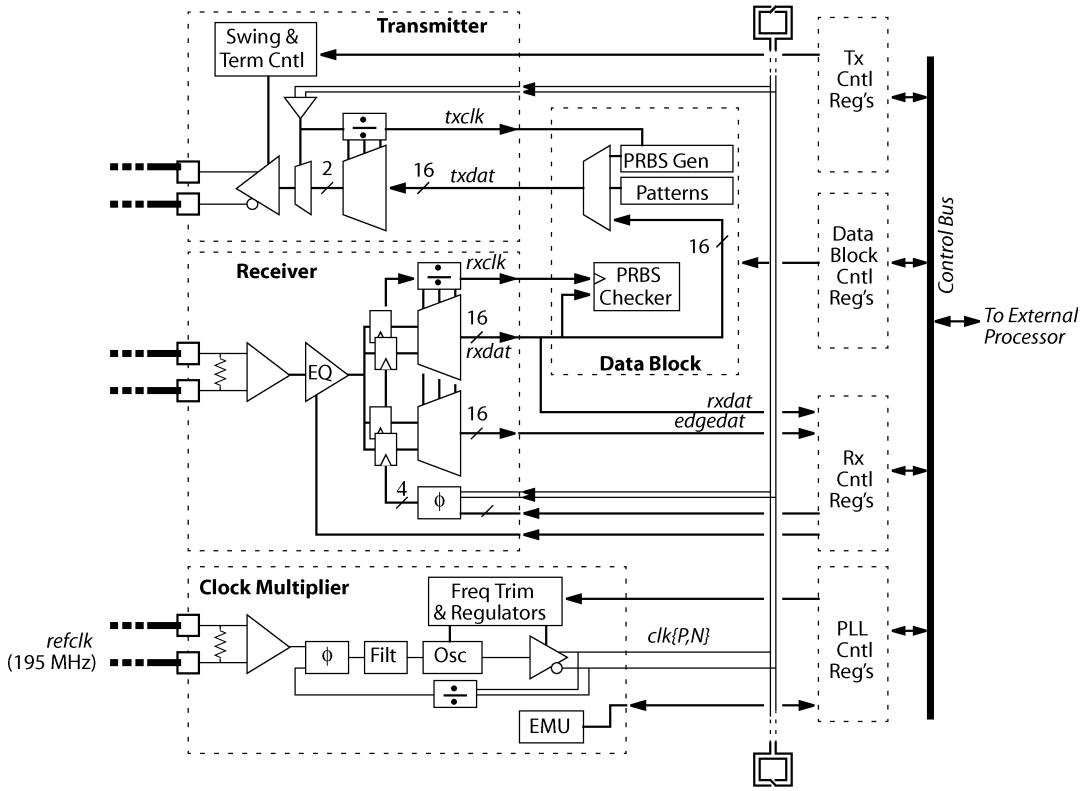

Fig. 2. Transceiver architecture.

The frequency of operation was chosen to best utilize the performance of the target technology, the TSMC 90 n "G" process, a standard 1-V bulk CMOS logic process. At 6.25 Gb/s, 1 UI is about four FO4 delays in this process. At this speed, it is relatively easy to buffer the half-bit-rate clock and to build the critical high-speed structures, including the half-bit-rate 2:1 mux in the transmitter and the edge/data samplers in the receiver. At much higher speeds, these structures either dissipate unreasonable amounts of power or introduce large timing uncertainties. The analyses in [1] and [3] suggest that 1 UI $= 4 \bullet$ FO4 is about the optimum data rate for maximizing bandwidth/power.

The transmitter comprises a 16:1 multiplexer and a "voltage-mode" output stage. The output stage is designed to produce output voltages between 50 and 300 mV peak-to-peak differential (Vppd). The output swing and termination impedance are digitally programmable. Voltage swings are near GND ("low-common-mode" signaling), with the common-mode voltage Vcm$= 0.5 \times$ V(out)ppd.

The receiver contains a two-stage input amplifier. The first stage provides gain and level shift from near-GND to near-Vdd. The second stage is a digitally adjustable source-degenerated differential amplifier that provides equalization for the link. A conventional Alexander phase detector follows, consisting of four offset-trimmed samplers and a digitally adjustable phase rotator that produces quadrature sampling clocks from the distributed clock $clk\{P,N\}$. The two data samplers (upper) and two edge samplers (lower) drive two essentially identical 2:16 demultiplexers. Data samples are delivered to the chip core. Coordinated data and edge samples are also provided for a software clock-and-data-recovery (CDR) algorithm and for a software adaptive equalization algorithm, which are both described in Section IV.

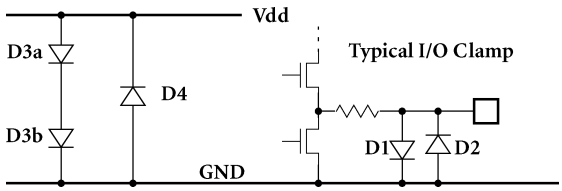

Fig. 3. ESD protection strategy; crossed diodes were used to protect both the transmitter and receiver.

The data block contains a $2^{23} - 1$ PRBS checker and error counter that accepts data from the receiver and a multiplexer that provides data for the transmitter. This mux can select the output of a $2^{23} - 1$ PRBS generator, a programmable pattern memory block, or parallel data loopback from the receiver.

A set of registers control all programmable functions. These registers, associated with transmitters, receivers, data blocks, and the clock multiplier, are connected to a common 8-bit control bus, which is interfaced to an external control processor. In a production version of the link, this processor would be implemented on-chip or replaced by custom control logic. We are not counting the power consumed by the control processor in our power estimates; however, we synthesized a block of logic that directly implements the code running on the control processor and found that this logic block consumes about 400 μW, so processor power is negligible.

Equalization is performed entirely within the receiver in a linear equalizer. The principal advantage of this approach is that

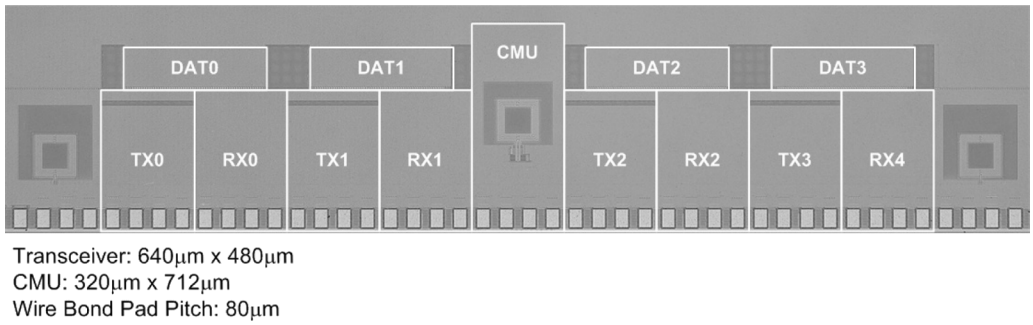

Fig. 4. Die photograph of the test chip.

the equalization setting can easily be set adaptively without the need for a back-channel to the transmitter.

The target application for the link is chip-to-chip interconnect, usually involving two packages, a relatively short run of PCB trace, and no connectors. At 6.25 Gb/s, such a channel is fairly benign, and -10 dB is a fairly conservative estimate for the attenuation. Our initial estimates of proportional, fixed, and random noise sources predicted that an amplitude of 200 mVppd at the transmitter output would achieve a BER of 10^{-18} for a -10 dB channel. We therefore designed our transmitter to have adjustable output between 300 mVppd (max), which is sufficiently large to cover any gross misestimates of link and channel parameters, and about 50 mVppd (min), small enough to generate significant errors even in fairly good channels. For short-haul links, dc coupling is preferred, since it avoids the need for bulky and expensive coupling capacitors. However, the link will also operate ac-coupled with appropriate encoding.

To provide high-quality termination, the capacitance shunting the termination must be small. In both the transmitter and receiver, shunt capacitance is dominated by wiring and ESD clamps. Since the signaling in this link is near-GND and the signal swing is quite small, we employed an unusual but very simple ESD protection strategy that uses both reverse- and forward-biased junction diodes. The overall ESD protection strategy is outlined in Fig. 3.

The forward-biased device has very small conduction when biased below 0.5 V and therefore has a negligible effect on the termination impedance. There is a small additional contribution to random noise due to the shot noise in the forward-biased junction. The diodes are very small $(1.44 \times 35.84\ \mu\text{m})$, contributing less than 130 fF to the I/O capacitance. This set of clamps was sufficient to protect the test chip to 4 kV (HBM), 200 V (MM), and 300 V (CDM) in ESD stress tests (applied signal–signal, signal–supply, supply–supply) performed on four chips. Some I/O's on some of these chips "failed" (exhibited 20% change in I/V characteristics) at levels higher than these, but we observed no outright failures in any I/O's subjected to ESD overstress, though "failed" devices exhibited some degradation in output amplitude or BER.

A die photograph of the test chip is shown in Fig. 4. The *LC*-PLL reference clock multiplier is centrally located, with a pair of transceivers on each side. The load inductors for the clock distribution are located at each end. A single transceiver occupies 0.307 mm^2, of which 25% is the Vdd bypass capacitor implemented in the 2.5 V thick-ox "native" nMOS device. The *LC*-PLL is 0.228 mm^2, of which 6% is the Vdd bypass capacitor.

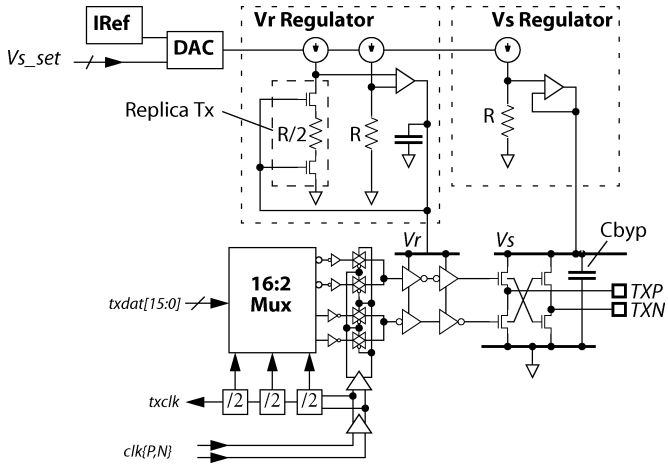

Fig. 5. Transmitter block diagram.

III. Transmitter

Fig. 5 shows a block diagram of the transmitter. The output terminals are driven by an N-over-N stage operating between GND and the regulated supply voltage Vs, which is generated by an on-chip regulator. The voltage Vs, and thus the output swing, are adjusted digitally via a DAC that sets the currents within the regulators. Each transmitter has its own regulator, so that the output swing can be adjusted on a per-transmitter basis between 50 and 300 mVppd. The gates of the output stage are driven by a set of inverters powered from a second regulated voltage Vr. This voltage is generated in a replica-bias circuit that forces a replica transmitter to have the same impedance as a scaled resistor. This method sets the sum of the pullup and pulldown impedance of the transmitter approximately equal to the line impedance. The relative size of pullup and pulldown is fixed at design time to give equal impedance at an assumed operating point.

The differential impedance looking into the transmitter output terminals is fairly independent of common-mode voltage because the pullup is operating as a source follower while the pulldown is common-source. As the common-mode voltage increases, the pulldown's small-signal impedance increases, while

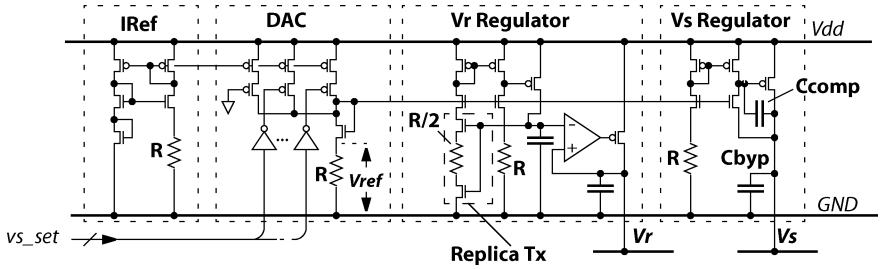

Fig. 6. Regulator details.

the pullup's impedance decreases. During a data transition, the output impedance depends on the details of the trajectories of the drive voltages at the gates of the output transistors. Simulation indicates that the small-signal differential output impedance varies by less than 15% during a transition.

The capacitance Cbyp not only rejects noise on Vs but is also a critical part of the termination. Line currents into $TX\{P,N\}$ flow through the pullup/pulldown impedances in series with the regulator output impedance, which is dominated by $\mathrm{X_{Cbyp}}$ at high frequencies. Cbyp, which is about 36 pF, is implemented in the 2.5-V "native" nMOS device and occupies about 8400 μm^2. We used the thick-oxide device because of oxide reliability concerns, but this capacitor is charged to less than 300 mV, so the oxide is not heavily stressed in any case. Use of the 1-V native device would have provided 2.5$\times$ the capacitance in the same area.

The predriver stage is powered from the termination-control voltage Vr. A pair of inverters is required here so that transition times for both data-edge polarities will be equal. This stage provides fanout, allowing the 2:1 multiplexer to be drawn quite small, thereby minimizing the load on the half-bit-rate distributed clock. Since this fanout stage is powered from a regulated supply, it is fairly immune to power supply noise and introduces very little timing jitter. The variation of Vr across cases tends to make the edge rate of the gate control signals driving the output stage nearly constant across PVT variation, and the edge rate of the transmitter output is also nearly constant.

Fig. 6 is a more detailed view of the regulators. Since the voltages that must be compared within the regulator circuitry are near GND, common-gate (CG) nMOS amplifiers are a good choice for the gain blocks.

A Vt-referenced current source produces a bias voltage for a pMOS DAC. The current in the reference is proportional to $\mathrm{V_{TN}/R}$, so the voltage across the resistor load in the DAC is $V_{\mathrm{ref}} = \mathrm{K} \bullet \mathrm{V_{TN}}$, independent of R, where K can be adjusted digitally via vs_set. The diode-connected nMOS in the DAC is drawn the same size as the nMOS common-gate devices in the two regulator amplifiers, and, assuming high loop gain, the voltage $Vref$ appears at all four of the input terminals of the CG amplifiers.

The Vs regulator is a simple series regulator that forces $Vs = Vref$, thereby allowing Vs to be set directly by vs_set. This regulator has two poles in its loop transfer function: one set by the output impedance of the CG amplifier and the input capacitance of the pMOS series regulator, and the second by the load impedance and the bypass capacitor Cbyp. We chose to add a compensating capacitor to make the amplifier pole the dominant one. This choice saved power in the CG error amp and avoided an extremely large bypass capacitor, but compromised the power supply rejection for the regulator, allowing half of the supply noise to get through the regulator between 10–100 MHz. In a production version of this link, we would want to make the output pole dominant, requiring larger bypass capacitance and higher power consumption in the CG amplifier. For example, with Cbyp implemented in the thin-oxide native nMOS in the same 8200 μm^2 area, it appears possible to build a greatly improved regulator for about 1 mW of additional power dissipation.

TABLE I
TRANSMITTER SIMULATED POWER SUMMARY

Function	Power
Serializer	2.0 mW
Predriver	0.4 mW
Output driver	1.1 mW
Regulator	0.76 mW
Clock buffer	1.1 mW
Total power	**5.4 mW**

The Vr regulator is a two-stage design. The first stage generates a "master" copy of the Vr control voltage. Since the load current for the pMOS stage is near-0, it is easy to make the output pole for this two-pole regulator the dominant one, and power supply rejection is quite good. The second stage is a simple series regulator with a gain of one, and it serves to isolate the "master" Vr from the time-varying load of the transmitter's predriver inverters. The transmitter replica used to set the impedance is drawn very small (1/16th scale); mismatch between the replica devices and the main transmitter contributes about a 5% variation in output impedance.

All of the P+/poly de-salicided resistors R are digitally trimmable by $\pm$20% to account for process variation. Trim is performed using a bench measurement; a production link would use a resistor trim cell and an external reference resistor.

Overall, the transmitter's regulator system consumes about 0.76 mW.

The half-bit-rate distributed clock $dk\{P,N\}$ is received and buffered in two CMOS inverter stages to drive the final 2:1

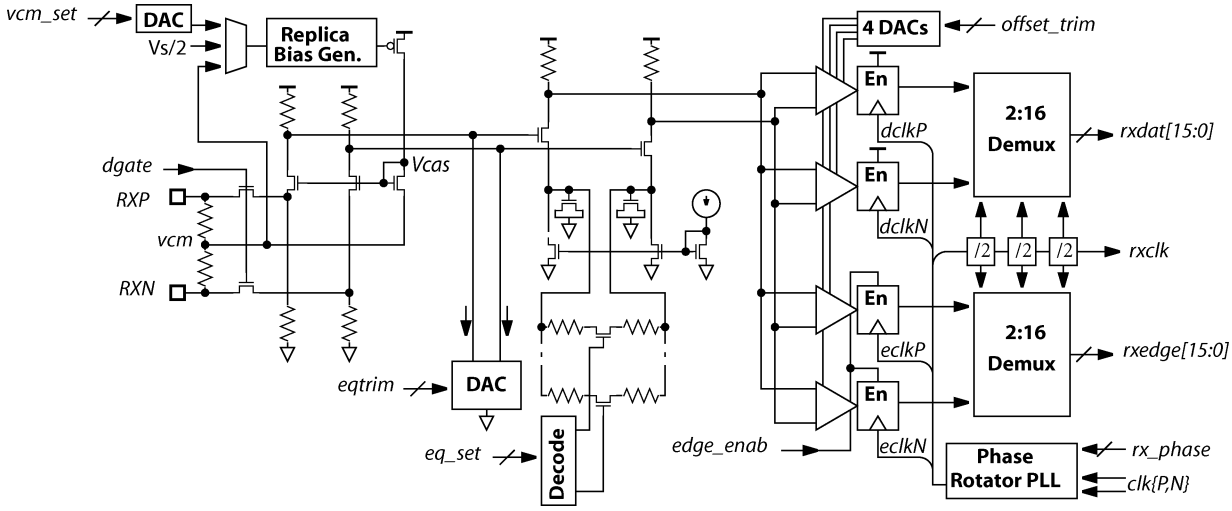

Fig. 7. Receiver overview.

mux stage and a divide-by-2 stage that generates the quarter-rate clock for the 4:2 mux stage. Since this load is relatively small, we kept the fanout of the two-stage buffer to $\mathrm{F} = 2$ to avoid introducing jitter from power supply noise, while still dissipating relatively little power in the clock buffer.

Power breakdown for the transmitter is shown in Table I. Simulated total power at nominal PVT, transmitting a $2^{23}-1$ PRBS, is 5.4 mW.

IV. Receiver

Fig. 7 shows an overview of the receiver. The receiver inputs drive a CG amplifier that provides level conversion for the second-stage amplifier/equalizer as well as gain. The input is terminated for the differential mode only, and about 25% of the conductance of the terminator is represented by the amplifier's input. Correctly terminating the common mode in the CG amplifier would have required burning 5× the power to maintain the proper input common-mode bias. We opted instead to rely on the transmitter's common-mode termination along with the receiver's common-mode rejection.

Large nFET switches enabled by $dgate$ allow the receiver to be disconnected from the line for offset trim. The bias voltage $Vcas$, developed in a replica bias arrangement, sets the input common-mode voltage of the receiver to the nominal output common-mode voltage of the transmitter; alternatively, $Vcas$ can be programmed externally.

The second-stage amplifier is a fairly conventional source-degenerated differential amplifier. The amount of degeneration is controlled digitally via the eq_set input. At a maximum EQ setting, this amplifier provides about 8.7 dB of peaking at the Nyquist frequency relative to dc gain, with a slope of 3 dB/octave, as shown in Fig. 8. The offset of the two input amplifier stages can be trimmed, separately from the samplers, by adjusting the $eqtrim$ DAC. The EQ amplifier drives a bank of four samplers, two data and two edge, derived from StrongARM flip-flops [5]. Each sampler has an associated 8-bit DAC to cancel input offset to within about 1 mV. Each sampler pair (data and edge) drives identical 2:16 demultiplexers composed

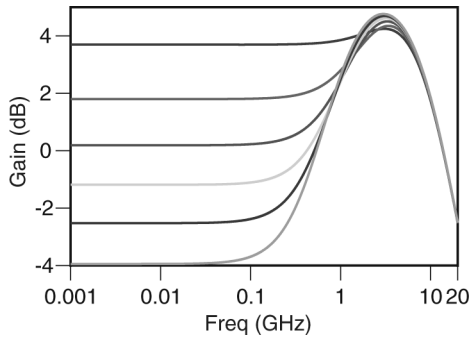

Fig. 8. Second-stage gain versus eq_set setting.

of a binary tree of 1:2 stages. Samplers are equipped with an *enable* input, which, when de-asserted, prevents the samplers from toggling. Edge samplers are enabled infrequently by the software CDR and adaptive EQ algorithm. When disabled, the edge samplers and their associated 2:16 demux dissipate essentially no power. A 7-bit digitally programmable phase rotator, implemented as a PLL, accepts the distributed half-bit-rate clock as input and generates four quadrature half-bit-rate clocks to drive the four samplers.

Input offset is trimmed iteratively. First, both input amplifiers are powered down, allowing each sampler to be trimmed with its inputs at Vdd; next the input amplifiers are powered on, with inputs disconnected from the line by de-asserting $dgate$, and the resulting offset, averaged over the samplers, is trimmed using the $eqtrim$ DAC. Finally, any residual offset is removed at each sampler. The maximum required offset trim across 32 receivers (128 samplers) was about 45 mV for the samplers and 25 mV for the input amplifiers.

The phase-rotator PLL's overall block diagram is shown in Fig. 9. It operates at the same frequency as its reference input and is built around a two-stage differential CMOS oscillator that generates the four sampler clocks, which are first level converted and then buffered. The oscillator uses very small devices in order to save power, and mismatch between these devices introduces fairly large uncertainty in phase positioning, as noted

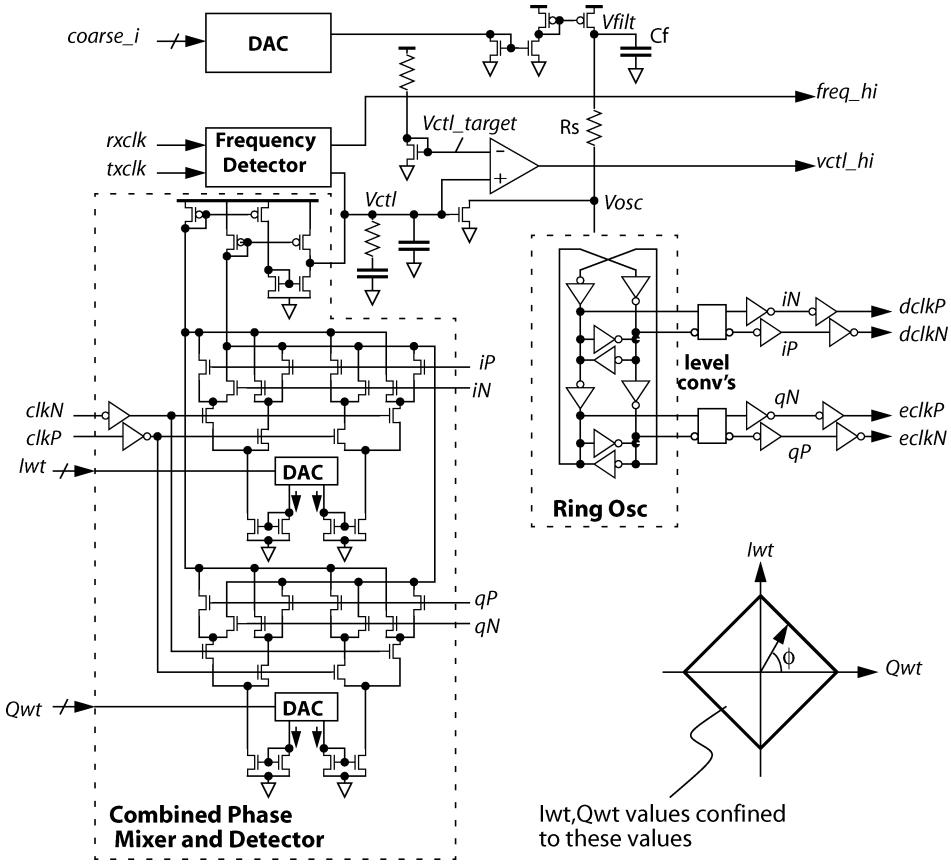

Fig. 9. Phase-rotator PLL.

in the measurements section of this paper. We included trim hardware to cancel these phase offsets, but, thanks to a design blunder, this hardware was ineffective. It appears straightforward to include a workable phase trim function, however, and we have devised and partially tested a trim algorithm that uses a training sequence to adjust the four sample times to near-optimal quadrature positions.

The phase rotator's control element is a combined phase mixer/detector. The phase shift is introduced in this element in the feedback path of the PLL and is controlled digitally by the inputs Iwt and Qwt. It was inspired by the design in [6], though it differs in detail. The phase detector is fundamentally an XOR (type-I) phase detector, and, since it has limited capture range, a separate asynchronous state-machine frequency detector is provided to achieve initial frequency lock. The control voltage $Vctl$ generated by the phase and frequency detectors feeds a shunt regulator that controls the output phase/frequency via the VCO's positive supply rail $Vosc$. The shunt nMOS pulls current through the series resistor Rs to change the voltage on $Vosc$. The current supplied to the VCO through Rs is controlled by a DAC and current mirror, whose input $coarse\_i$ allows the voltage on $Vctl$ at the PLL lock to be varied. The initial frequency lock is obtained by varying $coarse\_i$ until frequency lock is obtained and $Vctl$ is close to $Vctl\_target$, and then control is handed off to the phase detector.

The phase detector consists of four current-weighted XOR's, whose differential outputs are summed in a current mirror that drives $Vctl$. When the PLL is locked, the currents summed onto the two branches of the current mirror are equal. By varying the Iwt and Qwt values such that $|Iwt| + |Qwt| = 1$, as shown in the inset in the lower right-hand side of Fig. 9, a phase shift $-\phi$ that varies linearly is introduced into the feedback path of the PLL, and thus the output phase is rotated by $+\phi$. A similar phase rotator is used in [7], though that design places the rotator in the forward clock path. Since our phase mixer is in the feedback path and its output is current mode, it can be drawn very small and consumes little power relative to forward-path phase mixers.

The PLL's output frequency is the same as its input frequency, so it can operate at very high bandwidth. The shunt regulator's gain is set quite low, so the effective VCO gain seen from $Vctl$ is very low. Therefore, the loop filter components R1, R2, and C1 are small and occupy very little area. An AC model of the PLL control loop is shown in Fig. 10. The series resistor Rs and the filter capacitor Cf are in the high-gain part of the loop, where they introduce a pole-zero pair at $1/[(\mathrm{Rvco}+\mathrm{Rs})\mathrm{Cf}]$ and $1/\mathrm{RsCf}$, respectively (we have ignored the regulator impedance $1/\mathrm{gm2} \gg 1/\mathrm{sCf}$). We suppressed the effect of this pair by making Cf very large ($>$200 pF) and verified that the loop is stable and over-damped across PVT cases with a MATLAB continuous-time simulation. The presence of this pole-zero pair, however, made it impossible to further increase the loop gain and, thus, limited the loop bandwidth to 100 MHz. The small impedance of Cf together with the large impedance of the $coarse\_i$ current source $1/\mathrm{gm2}$ provides excellent power-supply noise rejection for the PLL. Any noise

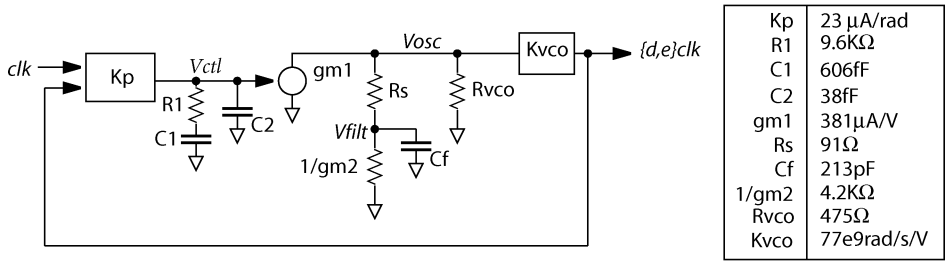

Fig. 10. Phase-rotator PLL loop model.

TABLE II
PHASE-ROTATOR POWER SUMMARY (FROM SIMULATION)

Function	Power
VCO	0.50 mW
Level conv's & buffers	1.75 mW
Phase mixer/detector	0.75 mW
Frequency detector	0.13 mW
Loop control	1.1 mW
Total:	**4.2 mW**

TABLE III
RECEIVER SIMULATED POWER SUMMARY

Function	Power
Input amplifiers	2.3 mW
Samplers	0.5 mW
Deserializers	1.6 mW
Phase rotator:	4.2 mW
Total power	**8.6 mW**

that survives this filter is low frequency and is well inside the PLL loop bandwidth, where it is cancelled by the loop itself. In designs with constrained area, requiring such a large filter capacitor may, of course, be an unattractive alternative, and, in those cases, some other type of regulator may be needed. Transport delay through the oscillator's output buffers and the phase detector is of the order of 0.75 output periods, and, while short, this delay has a nonnegligible affect on loop dynamics. With the component values shown in Fig. 10, a continuous-time model implemented in MATLAB predicts a 54° phase margin. A discrete-time model with 0.75 cycles of feedback delay predicts a phase margin of about 51°.

The phase rotator consumes half of the power in the receiver. Power consumption is summarized in Table II.

The CDR in this link uses a conventional Alexander-type phase detector operating at half bit rate. The CDR logic is implemented in software running on the control processor. The processor occasionally enables the edge samplers and de-serializer, then a set of 16 edge samples, and their associated 16 data samples are processed to compute "early" and "late" indications. As usual

$$\text{Early} = (\text{Edge}_i \oplus \text{Data}_i) \bullet (\text{Data}_{i-1} \oplus \text{Data}_i)$$
$$\text{Late} = \overline{(\text{Edge}_i \oplus \text{Data}_i)} \bullet (\text{Data}_{i-1} \oplus \text{Data}_i).$$

Because the external processor is operated at low clock speed (5 MIPS), performs a number of different run-time functions, and is shared by four links, the overall CDR bandwidth is quite low (about 128 Hz for a jitter amplitude of 0.25 UI). This is sufficient to track timing variations due to temperature and supply-voltage drift in a mesochronous link, but cannot track timing jitter. Extending CDR bandwidth to several megahertz for jitter tracking is fairly expensive in terms of power. It requires running the edge samplers continuously, along with additional logic that we estimate would dissipate about 5 mW, based on a recent in-house 90-nm design.

The early/late indications extracted from the edge sampling operation can be used to adapt the equalization setting in the receiver to existing channel attenuation. This is done as in [8] by correlating early/late indications with the data bits that produced them. We look for certain data patterns, for example ..00010.. and its complement ..11101.. embedded in the received data stream. When these data patterns are found, we examine the early/late indications for the isolated 1(0). Assuming the CDR is locked (the number of early's and late's is equal, averaged over a large number of data transitions), if the leading edge of the isolated bit is always late and the trailing edge is always early, we infer that the link is under-equalized and advance the setting of the linear EQ in the receiver by a step. If the leading edge is always early and trailing edge is always late, we infer that the link is over-equalized and reduce the linear EQ setting by one step.

The power breakdown for the receiver, simulated under nominal PVT, is shown in Table III.

V. CLOCK MULTIPLIER AND CLOCK DISTRIBUTION

The clock multiplication and distribution scheme is shown in Fig. 11. The reference clock is multiplied up to half the bit rate in a conventional CMOS *LC*-PLL with 5-MHz loop bandwidth. This bandwidth is sufficient to reject the PLL's internally generated low-frequency noise, instead tracking in-band noise on the reference clock at both ends of the link. VCO frequency is digitally trimmed by switching in some number of metal–metal capacitors as in [9] and phase-locked using varactors. A conventional divider, phase-frequency detector, charge pump, and loop filter complete the PLL. The EMU (not shown) allows the

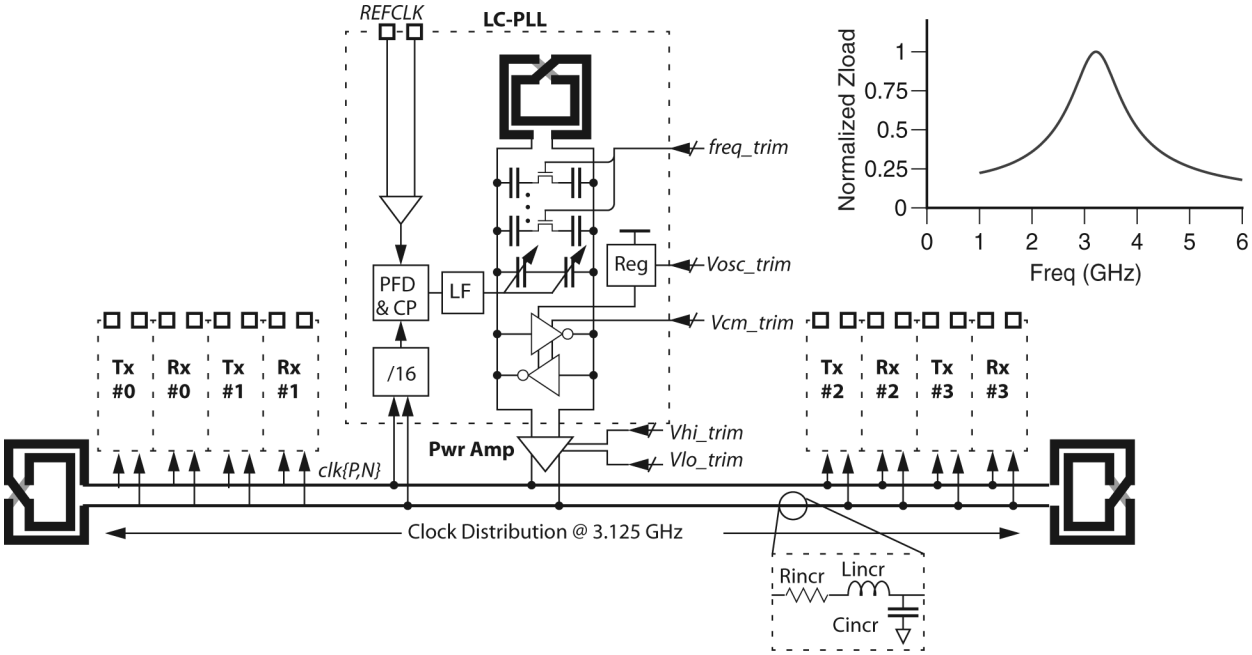

Fig. 11. Clock multiplier and distribution overview. The inset shows the frequency response of distribution network.

control processor to measure the varactor bias, oscillator output voltages, and clock distribution voltages. Software on the processor can thereby set the frequency trim, oscillator swing, and common-mode voltages to optimize power and jitter. A power amplifier buffers the oscillator output onto the clock distribution network, and processor software likewise trims the swing and common-mode voltages on this network so as to optimize power consumption across PVT variations. A pair of inductors at each end of the network resonate the network to minimize power consumption. The network has fairly low Q, so no trim capacitors are required. All inductors are flat differential spiral coils on the top two levels of metal.

The *LC*-PLL's tank inductor is a five-turn, 5.5 nH coil, 142 μm on a side, with a Q of 6.5, limited mainly by the eddy currents in the highly doped bulk of the target process. Fixed tuning capacitors are finger-type metal–metal capacitors. The tuning voltage from the loop drives an accumulation-mode 1-V (thin-oxide) varactor. The tuning range is about $\pm$ 10% to cover process variations in the tuning elements. Losses in the tank are restored by a pair of cross-coupled CMOS inverters whose positive supply voltage is supplied from an on-chip regulator. The regulator output voltage is set digitally by the control processor, nominally to about $0.85 \times \mathrm{Vdd}$. The swing and common-mode output voltage from the oscillator are set to provide optimal drive to the power amplifier that drives the distribution network.

The centrally located clock multiplier and buffer drive 1.2 Vppd into the distribution wiring and out to the four transceivers. The clock is restored at each transmitter and receiver with CMOS inverters. Inductors are connected across the differential clock distribution wires at each end of the network to form a tank circuit with the distribution capacitance. The circuit resonates at 3.125 GHz with $Q = 3.5$, giving a 3.5$\times$ increase in impedance relative to the *RC* network presented by the wiring and loads alone. Resonating the clock load has the significant side benefit of reducing duty-factor distortion in the distributed clock and more generally rejecting phase modulation due to noise sources in the clock multiplier and buffer.

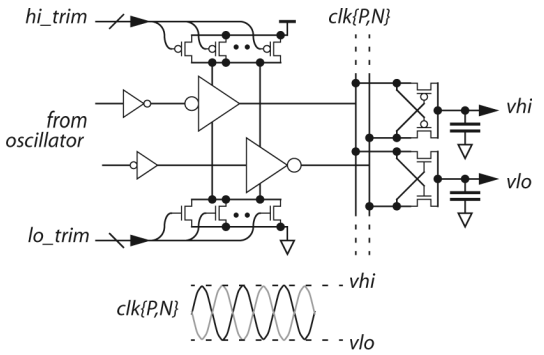

Fig. 12. Clock buffer circuitry.

The clock distribution wires are routed in metal-8, 1.6 mm in each direction from the central clock multiplier. The total tank capacitance is about 0.5 pF, mostly due to the wiring itself. The wires have an incremental inductance of 0.32 nH/mm. Each end of the distribution wiring is loaded with an 8.6 nH differential square spiral inductor of six turns and 140 μm on a side. The network was simulated as a distributed-element model, and the resulting impedance versus frequency curve is shown in the inset of Fig. 11. Process variations introduce approximately a $\pm$5% variation in resonant frequency, and this variation produces negligible power losses in the relatively low-Q network. For a single-frequency link, no tuning is needed; for links that must accommodate a wide range of frequencies, it appears feasible to tune a resonant clock distribution system, but a transmission-line-based system may be preferable.

Fig. 12 is a schematic of the clock buffer The first stage is a CMOS inverter and the second stage is a pair of inverters with tail transistors, digitally trimmed to set the swing and common-

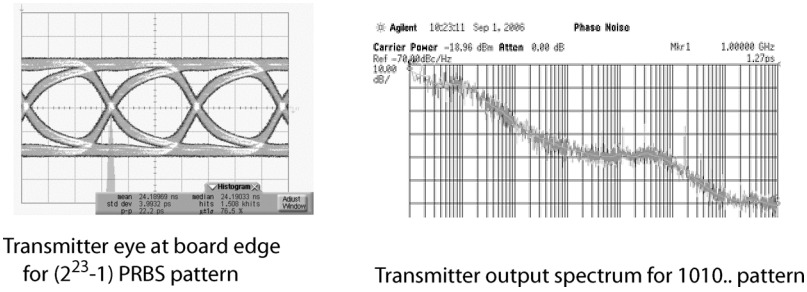

Fig. 13. Output eye and jitter spectrum.

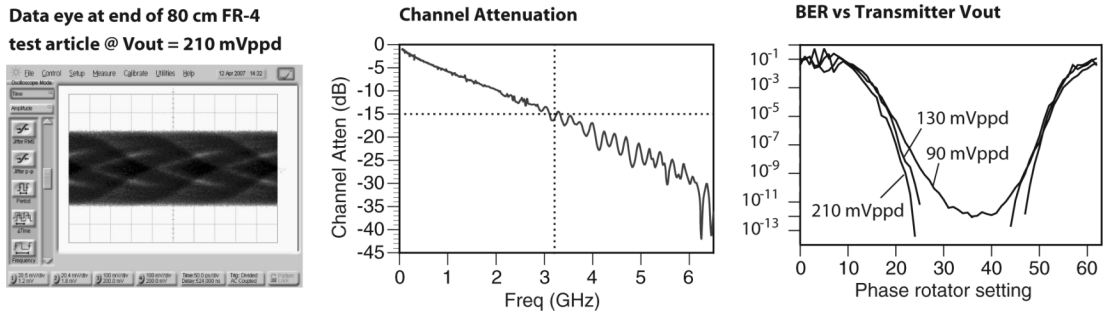

Fig. 14. Test channel attenuation and BER measurements.

TABLE IV
MEASURED VERSUS SIMULATED POWER CONSUMPTION BY MODULE

Module	Measured	Simulated
Transmitter	4.9 mW	5.4 mW
Receiver	8.0 mW	8.6 mW
Clock mult & distr	3.6 mW	4.6 mW

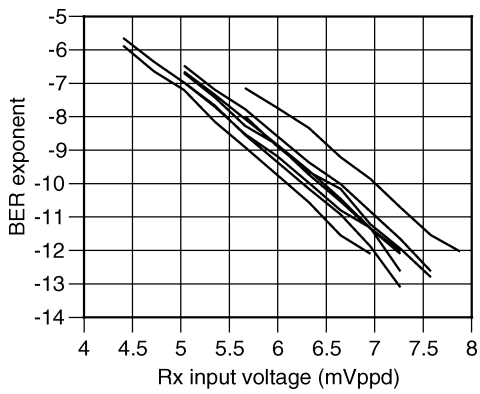

Fig. 15. Receiver sensitivity curves.

mode voltage at the output of the buffer. At startup, vhi and vlo are measured by the EMU, using the "synchronous rectifier" peak detectors shown in the figure. The control processor then trims the pMOS and nMOS tail devices in the buffer to achieve the desired swing, centered on the inverter threshold of a typical clock buffer inverter. An exemplar inverter, with input shorted to output, is provided for measurement by the EMU as well, so that PVT can be tracked by the initialization algorithm.

At nominal PVT, simulated power is 2.24 mW for the PLL and 2.09 mW for the clock buffer. This power is amortized across four transceivers, 1.1 mW per transceiver.

VI. MEASUREMENTS

The experimental four-link test chip was socketed on a test board that provides clocks, power, an off-chip control processor implemented in an FPGA, and fixturing for various measurements. The chip is mounted in a wire-bonded ball-grid-array package, and each high-speed connection is wired via the pogo-pin device socket and about 4″ of FR-4 microstrip to an SMA edge launcher. Unless otherwise noted, measurements were performed at ambient temperature and at the nominal data rate of 6.25 Gb/s, at Vdd $=$ 1 V, and with transmitter output level set to about 200 mVppd. For most tests, and during normal operation, the receiver's edge detector samplers and deserializer operate only about 1% of the time, so their power consumption is negligible.

Power Consumption: Measured power consumption for each component, passing $2^{23} - 1$ PRBS, is shown in Table IV. Total measured power per link is 13.8 mW for a specific power of 2.21 mW/Gb/s. Generally measured power is somewhat smaller than simulated power at nominal PVT, suggesting that test-chip wafers are slightly faster than nominal.

We also measured link power consumption as a function of transmitter output level and found a 14 μW/mVppd slope. As the signal level increases, power consumption in the transmitter output stage increases linearly; the receiver input amplifier also consumes more power as its input bias tracks the increasing common-mode voltage.

Output Jitter: Fig. 13 shows the transmitter output eye, transmitting a $2^{15} - 1$ PRBS pattern, measured at the edge of the test board (left) and the jitter spectrum for a 1010...

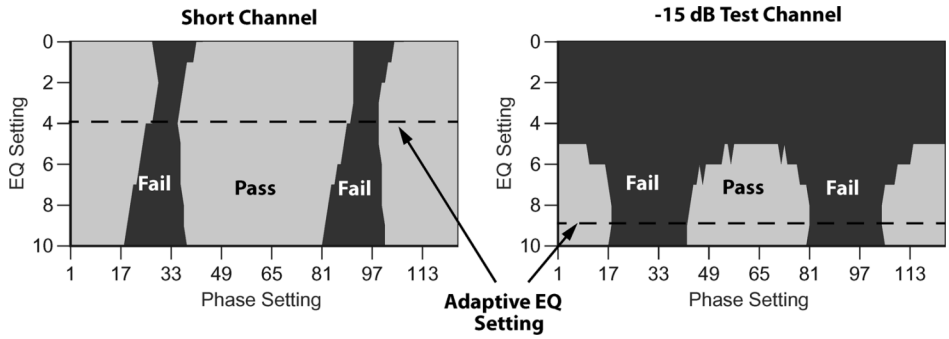

Fig. 16. Eye opening versus EQ setting; the dashed line indicates the EQ setting for edge-based adaptation.

pattern (right). Measured dc swing for the transmitter eye is 210 mVppd, and the vertical eye opening is 125 mVppd. Deterministic jitter, measured using an Agilent 86100C DCA-J, is 14.4 ps, primarily from attenuation in the socket, board traces, and launchers. We measured no increase in jitter when transmitting independent PRBS sequences on neighboring transceivers.

Random jitter was integrated over a bandwidth of 100 Hz to 1 GHz, using an Agilent E4404B spectrum analyzer, and found to be 1.27 ps rms. This result agrees well with the random jitter analysis performed on the DCA-J.

BER Measurements: Fig. 14 shows the channel response for one of several channels used to test the link. This channel, which includes an external test article with 80 cm of 8-mil FR-4 microstrip, exhibits about -15 dB of attenuation at 3.125 GHz. The channel response (center) was constructed in HSPICE from a combination of network analyzer measurements and vendor models. The PCB traces, pogo-pin socket, package, and pad capacitances are all included in the model. The eye at the far end of the test article is shown on the left-hand side of Fig. 14. The signal is about 200 mVppd (dc attenuation is very small), and the eye is completely closed. This test link was not intended to be representative of all possible channels, but is useful in that it stresses SNR by operating the link at an extreme of the attenuation likely to be encountered in many 6.25 Gb/s chip-to-chip links.

BER on this channel was measured using a $2^{23}-1$ PRBS sequence. The right-hand portion of the figure shows the error rate measured at a number of points across the eye, found by varying the phase rotator setting in the receiver, until the BER fell below 10^{-12}. The test was run at transmitter output swings of 90, 130, and 210 mVppd. The 90-mV swing gave a measured BER of 10^{-12} at the center of the eye. Data fitting was used to extrapolate performance at the 130- and 210-mV swings. At the center of the eye, respective BERs of 10^{-23} and 10^{-30} are predicted. For both swings, the results indicate that the link will operate with margin >100 mUI at a BER of better than 10^{-15}.

Receiver Sensitivity and Random Noise: The sensitivity curves for eight receivers, operated with all neighboring transceivers shut down, are shown in Fig. 15. For the worst of these, sensitivity is about 7.9 mV for BER $= 10^{-12}$, and (extrapolated) about 8.9 mV at 10^{-15}. These curves were generated by applying the output of an Agilent J-BERT N4903A through an attenuator into a receiver input on our test board. The BERT

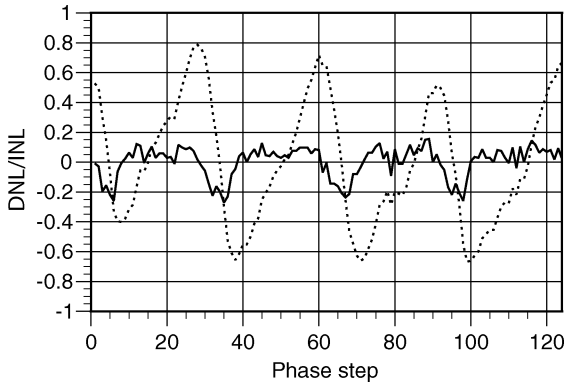

Fig. 17. DNL (solid line) and INL (dashed line) measurements for one receiver.

was programmed to generate a $2^{23}-1$ PRBS pattern. When all four transceivers were operating, sensitivity was about 10 mV for BER $= 10^{-12}$ and 11.4 mV for BER $= 10^{-15}$.

From the receiver sensitivity curves, we can make an estimate of the input-referred unbounded random noise (V_R) using

$$\mathrm{BER} < \exp\left(\frac{-\mathrm{VSNR}^2}{2}\right)$$

$$\mathrm{VSNR} = \frac{V_{\mathrm{IN}}(1-K)-V_N}{V_R}$$

from [10], where K represents the proportional noise, V_N represents the fixed sources of noise, and V_R represents the random noise. If we make the conservative assumption that there are no fixed or proportional noise sources, then the slope of VSNR versus the single-ended input voltage gives us an upper bound on the random noise. Accounting for the random noise appearing in the signal source, we calculate a mean value of 570 μVrms averaged across the eight receivers detailed in Fig. 15.

Sampler Clock Phase Offset: As discussed above, the ring oscillator that generates the four sampler clocks uses very small devices, so we expect large variation in timing placement of the clocks. Measurements indicate 1-σ variation of about 8°, and maximum observed error of about 20° out of 90. A Monte Carlo simulation with 100 iterations at nominal PVT of the oscillator produced nearly identical results. Data sampler clock misplacement reduces the voltage available to the sampler by about 15%.

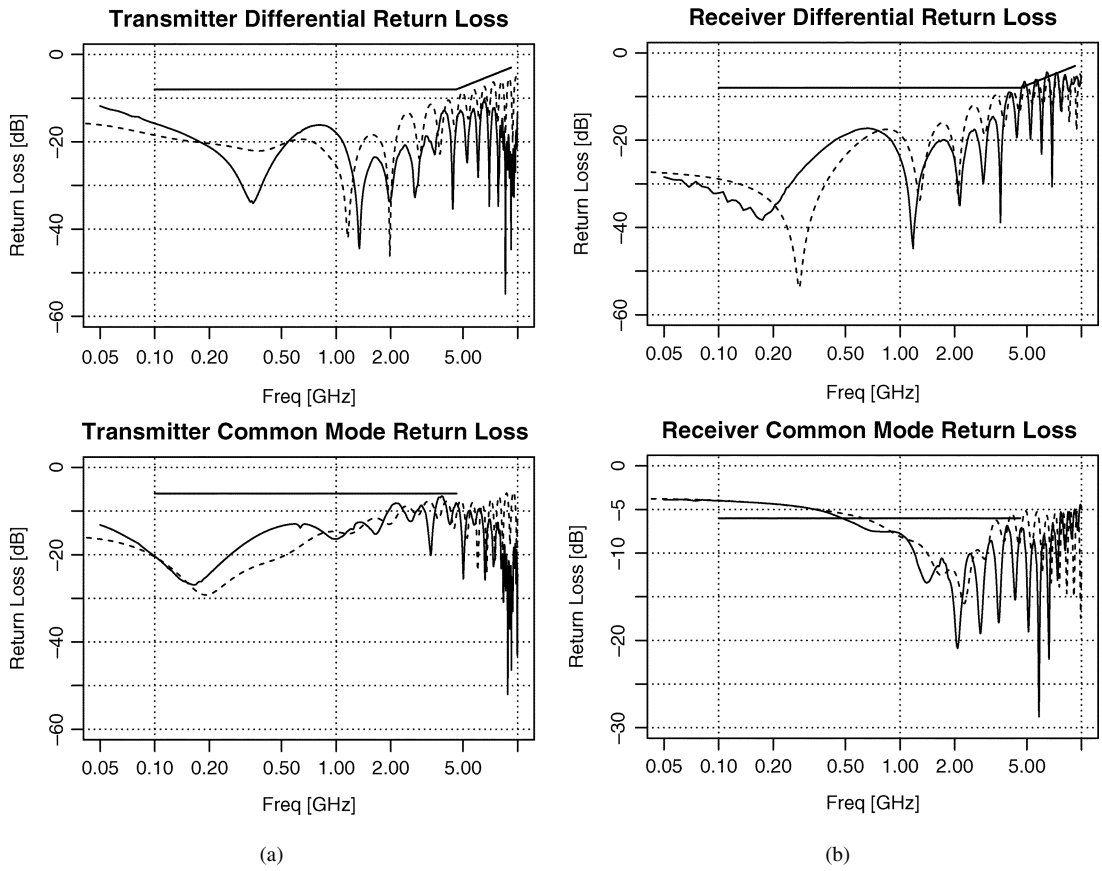

Fig. 18. Simulated (dashed line) and measured (solid line) return loss for (a) transmitter and (b) receiver. Differential S11 (top) and common mode (bottom). The CEI-6 return loss limits are shown for reference.

Adaptive EQ Algorithm: Fig. 16 shows schmoo plots for the receiver's EQ setting for two channels: a short length of coaxial cable (left) and the -15 dB channel used for the BER measurements of Fig. 14 The dashed line shows the setting to which the adaptive EQ algorithm converges. "Pass" in this context means no errors detected over a very short (12 μs) interval of operation, corresponding to a BER of better than 10^{-5}.

CDR Phase-Rotator Linearity: Linearity of the phase rotator was measured for 32 receivers. Fig. 17 shows a plot of the DNL and INL for one of these receivers. Across the 32 measured receivers, DNL < 0.5 LSB and INL < 2.8 LSB. Phase-rotator INL is not critical in the test chip, which operates mesochronously, but is an important parameter for extending this design to plesiochronous operation. DNL and INL compare favorably with other published results [6], [7].

Return Loss: Fig. 18 shows measured and simulated return loss measurements for the transmitter and receiver. This measurement was performed with the transmitter sending a static value, and its output level was set to about 200 mVppd.

The CEI-6SR return loss limits are shown on the plots for reference. Since it is extremely difficult to build a fixture for a direct measurement of the return loss at the device pins, we chose instead to perform the measurements at the test-board edge and include a model for the launchers, board, socket, and package in the simulation. For all four S11's, measurement show a somewhat better match than simulation. The transmitter provides a reasonably good back-match for the common mode, while the receiver provides no common-mode termination (by design), and the match below 0.5 GHz is poor, as expected.

VII. Conclusion

We have described a low-power 6.25 Gb/s serial link suitable for many chip-to-chip applications with measured specific power of 2.2 mW/Gb/s, which is about an order of magnitude lower than most contemporary links. A combination of techniques were used to reduce power, including near-GND voltage-mode signaling with on-chip regulated output swing, a shared *LC*-PLL clock multiplier, resonant half-bit-rate clock distribution, a novel PLL-based phase rotator for the receiver's sampler clocks, and software CDR and adaptive equalization. A set of measurements was presented that demonstrate that the link can operate with BERs below 10^{-15} in a fairly difficult channel under conservative assumptions.

Acknowledgment

The authors would like to gratefully acknowledge the assistance of M. Kellam, F. Quan, and F. Zarkeshvari, Rambus, Inc., who were instrumental in producing the test chip.

References

[1] M.-J. E. Lee, W. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1591–1599, Nov. 2000.

[2] G.-Y. Wei, J. Kim, D. Liu, S. Sidiropoulos, and M. Horowitz, "A variable-frequency parallel I/O interface with adaptive power-supply regulator," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1600–1610, Nov. 2000.
[3] K.-L. Wong, H. Hatamkhani, M. Mansuri, and C.-K. Yang, "A 27-mW 3.6-Gb/s I/O transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 602–612, Apr. 2004.
[4] H. Hatamkhani and C.-K. Yang, "Power analysis for high-speed I/O transmitters," in *Proc. Symp. VLSI Circuits*, 2004, pp. 142–145.
[5] J. Montanaro *et al.*, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.
[6] T. Toifl, C. Menolfi, P. Buchmann, M. Kossel, T. Morf, R. Reutmann, M. Ruegg, M. Schmatz, and J. Weiss, "0.94 ps-rms-jitter 0.016 mm^2 2.5 GHz multi-phase generator PLL with 360 ° digitally programmable phase shift for 10 Gb/s serial links," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, pp. 410–411.
[7] J. Bulzacchelli *et al.*, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.
[8] Y. Hidaka, W. Gai, A. Hattori, T. Horie, J. Jiang, L. Kand, Y. Koyanagi, S. Matsubara, and H. Osone, "A 4-channel 3.1/10-Gb/s trasnceiver macro with a pattern-tolerant adaptive equalizer," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, vol. 50, pp. 442–443.
[9] H. Sjöland, "Improved switched tuning of differential CMOS VCOs," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 49, no. 5, pp. 352–355, May 2002.
[10] W. Dally and J. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998, pp. 297–298.

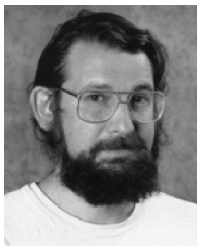

John Poulton (M'85–SM'90) received the B.S. degree from the Virginia Polytechnic Institute and State University, Blacksburg, in 1967, the M.S. degree from the State University of New York, Stony Brook, in 1969, and the Ph.D. degree from the University of North Carolina, Chapel Hill (UNCCH), in 1980, all in physics.

From 1981 to 1999, he was a Researcher with the Department of Computer Science, UNCCH, where from 1995 he held the rank of Research Professor. He performed research on VLSI-based architectures for graphics and imaging and was a principal contributor to the design and construction of several experimental high-performance graphics systems. From 1999 to 2003, he was a Chief Engineer with Velio Communications, where he was engaged in the development of gigabit signaling systems. Since 2003, he has been a Technical Director with Rambus, Inc., Chapel Hill, where he continues to work in the area of high-speed wireline communications.

Robert Palmer was born in Charlotte, NC, in 1975. He received the B.S. degree in physics from the University of North Carolina, Chapel Hill, in 1997.

He joined Velio Communications in 1999, where he worked on signal integrity and custom circuit design for several generations of high-speed serial backplane transceivers. In 2004, he joined Rambus, Inc., Chapel Hill, where he is currently a Principal Engineer working on the development of low-power, high-speed serial links.

Andrew M. Fuller received dual B.S. degrees in physics and computer engineering from the University of North Carolina, Chapel Hill, in 2006.

During the summer of 2006, he was an intern at Rambus Inc., Chapel Hill, NC. Since then, he has been a full employee for Rambus, specializing in testchip characterization, firmware programming, algorithms, and logic design.

Trey Greer (M'00) received the B.S. degree in math and physics from the University of the South, Sewanee, TN, in 1984, and the M.S. degree in computer science from the University of North Carolina at Chapel Hill in 1988.

Since 1988, he has been involved in various startups in the Chapel Hill and Bristol, U.K., areas. His research interests include clocking and equalization for high-speed data links. He is currently with Rambus, Inc., Los Altos, CA.

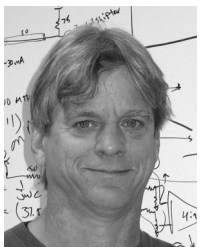

John Eyles was educated at (received the B.S. degree in mathematics from Davidson College in 1974, and the Ph.D. degree in biomedical engineering and mathematics from the University of North Carolina at Chapel Hill in 1982.

From 1983 to 1996, he served on the research faculty in UNC's Computer Science Department as a Senior Investigator in a series of projects building experimental computer graphics hardware, as well as head-motion tracking systems (another enabling technology for virtual reality). In 1996, he joined Hewlett-Packard's Chapel Hill Graphics Laboratory to commercialize UNC's latest experimental graphics system, PixelFlow. From 1999 to 2003, he worked at Velio Communications as a Senior IC Designer for high-speed chip-to-chip signaling systems, then joined Rambus when they acquired Velio's Chapel Hill team. He is currently an independent consultant specializing in high-performance CMOS design and graphics architectures.

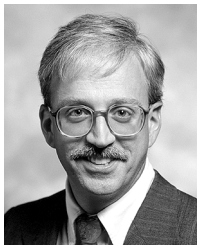

William J. Dally (M'80–SM'01–F'02) received the B.S. degree in electrical engineering from the Virginia Polytechnic Institute, Blacksburg, the M.S. degree in electrical engineering from Stanford University, Stanford, CA, and the Ph.D. degree in computer science from the California Institute of Technology (Caltech), Pasadena.

He is the Willard R. and Inez Kerr Bell Professor of Engineering and the Chairman of the Department of Computer Science at Stanford University. He is also co-founder, Chairman, and Chief Scientist of Stream Processors, Inc. He and his group have developed system architecture, network architecture, signaling, routing, and synchronization technology that can be found in most large parallel computers today. While at Bell Labs, he contributed to the BELLMAC32 microprocessor and designed the MARS hardware accelerator. At Caltech, he designed the MOSSIM Simulation Engine and the Torus Routing Chip which pioneered wormhole routing and virtual-channel flow control. While a Professor of electrical engineering and computer science at the Massachusetts Institute of Technology, his group built the J-Machine and the M-Machine, experimental parallel computer systems that pioneered the separation of mechanisms from programming models and demonstrated very low overhead synchronization and communication mechanisms. At Stanford University, his group has developed the Imagine processor, which introduced the concepts of stream processing and partitioned register organizations. He has worked with Cray Research and Intel to incorporate many of these innovations in commercial parallel computers, with Avici Systems to incorporate this technology into Internet routers, co-founded Velio Communications to commercialize high-speed signaling technology, and co-founded Stream Processors, Inc. to commercialize stream processor technology.

Dr. Dally is a Fellow of the IEEE, a Fellow of the ACM, and a Fellow of the American Academy of Arts and Sciences. He has received numerous honors including the IEEE Seymour Cray Award and the ACM Maurice Wilkes Award. He currently leads projects on computer architecture, network architecture, and programming systems. He has published over 200 papers in these areas, holds over 50 issued patents, and is an author of the textbooks *Digital Systems Engineering* and *Principles and Practices of Interconnection Networks*.

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

He is the Associate Vice Provost for Graduate Education working on Special Programs and the Yahoo! Founders Professor of the School of Engineering at Stanford University. In addition, he is Chief Scientist at Rambus Inc. His research interests are quite broad and span using EE and CS analysis methods to problems in molecular biology to creating new design methodologies for analog and digital VLSI circuits. He has worked on many processor designs, from early RISC chips, to creating some of the first distributed shared memory multiprocessors, and is currently working on on-chip multiprocessor designs. Recently he has worked on a number of problems in computational photography. In 1990, he took leave from Stanford to help start Rambus Inc, a company designing high-bandwidth memory interface technology, and has continued work in high-speed I/O at Stanford. His current research includes multiprocessor design, low power circuits, high-speed links, computational photography, and applying engineering to biology.

Dr. Horowitz has received many awards including a 1985 Presidential Young Investigator Award, the 1993 ISSCC Best Paper Award, the ISCA 2004 Most Influential Paper of 1989, and the 2006 Don Pederson IEEE Technical Field Award. He is a Fellow of IEEE and ACM and is a member of the National Academy of Engineering.