

A 27-mW 3.6-Gb/s I/O Transceiver

Koon-Lun Jackie Wong, Hamid Hatamkhani, Mozghan Mansuri, *Member, IEEE*, and
Chih-Kong Ken Yang, *Member, IEEE*

Abstract—This paper describes a 3.6-Gb/s 27-mW transceiver for chip-to-chip applications. A voltage-mode transmitter is proposed that equalizes the channel while maintaining impedance matching. A comparator is proposed that achieves sampling bandwidth control and offset compensation. A novel timing recovery circuit controls the phase by mismatching the current in the charge pump. The architecture maintains high signal integrity while each port consumes only 7.5 mW/Gb/s. The entire design occupies 0.2 mm² in a 0.18- μ m 1.8-V CMOS technology.

Index Terms—I/O, low power, transceiver.

I. INTRODUCTION

TECHNOLOGY scaling has led to increased off-chip data rate. The ITRS roadmap predicts that the aggregate data bandwidth of a chip will exceed several terabits per second (Tb/s) within ten years, as shown in Fig. 1. Widely parallel multi-Gb/s chip-to-chip I/O links are an integral part of these systems. Power consumption of these links is an increasing concern. With higher data rates per I/O port, the design must also provide good signal integrity. To compare power efficiency, this paper uses a normalized power metric of average power per Gb/s (mW/Gb/s). Previously published transceivers have power dissipation on the order of 18–40 mW/Gb/s [1]–[6]. Fig. 2 summarizes their power consumption. This power level would lead to an unaffordable power of 18 W for 1-Tb/s operation. Even with some power reduction from technology scaling, techniques are still needed for further power reduction.

This paper demonstrates a scalable design capable of 7.5 mW/Gb/s in a 0.18- μ m CMOS technology. The design includes features that maintain good signal integrity. The transmitter is source terminated along with slew-rate control and pre-emphasis to equalize the channel. The receiver has digital-offset compensation and sampling-bandwidth control to filter out high-frequency noise. Along with a power-efficient transmitter and receiver, the phase-recovery circuit requires no additional power by introducing static phase offset onto the charge pump. The transceiver operates at 3.6 Gb/s per port.

Section II describes the system and signaling architecture of the transceiver. The details of the transmitter are described in Section III. Sections IV and V describe the design of a low-power receiver and a novel timing recovery technique, respectively. Section VI summarizes the measurement results from an eight-channel test chip.

Manuscript received July 31, 2003; revised November 18, 2003. This work was supported by UCMicro 01–102.

K.-L. J. Wong, H. Hatamkhani, and C.-K. K. Yang are with the University of California, Los Angeles, CA 90095-1594 USA (e-mail: jwong@icsl.ucla.edu).
M. Mansuri is with Intel Corporation, Hillsboro, OR 97124 USA.

Digital Object Identifier 10.1109/JSSC.2004.825259

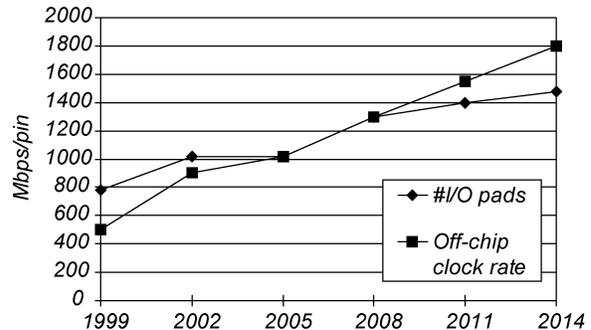


Fig. 1. ITRS roadmap prediction.

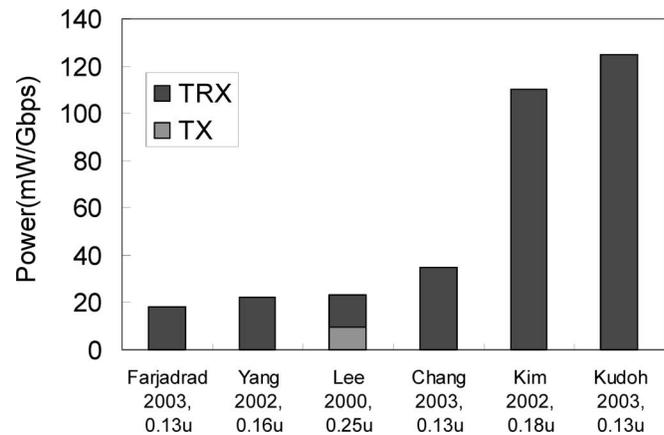


Fig. 2. Power comparison of previously published works.

II. TRANSCIEVER ARCHITECTURE

The transceiver shown in Fig. 3 is designed for widely parallelized half-duplex I/Os where each physical pin is capable of transmitting and receiving but not both simultaneously. The design of each I/O cell targets a bit-time of three fanout-of-four (FO-4) inverter delay which is equivalent to 3.6-Gb/s data rate in the 0.18- μ m technology.

The transmitter maintains 50- Ω impedance matching to the channel to reduce the impact of signal reflections. Fig. 4 shows two common signaling techniques, high common-mode (HCM) and low common-mode (LCM) signaling. In HCM signaling [Fig. 4(a)], the driver transistor operates in saturation. Termination is provided by a 50- Ω resistor. In LCM signaling [Fig. 4(b)], a much lower V_S of 500 mV is used. The driver actively pulls up and pulls down. With high gate voltage, the devices are operating in triode region; hence, NMOSs are used for pull-up. Impedance matching is achieved by precisely controlling the gate voltage. The comparison of the power of the output driver is shown in Table I. The table includes the power needed to switch

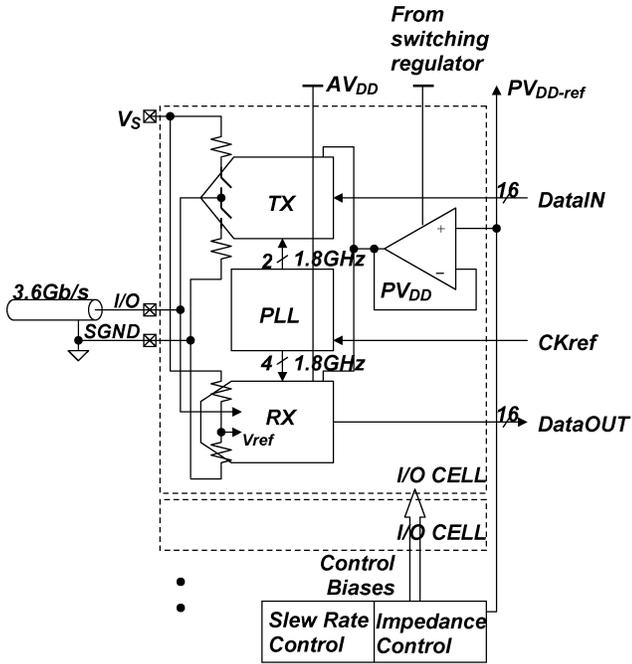


Fig. 3. Transceiver architecture.

the driving transistor. Despite the extra drive transistor for LCM signaling, the total power is significantly lower.

We chose LCM signaling for the design using a dedicated driver supply of $V_S (= 0.5 \text{ V})$. The design is single-ended for minimal power but can be easily extended to differential signaling with only a modest increase in power consumption. In the event that an external V_S is not available, power-efficient switching regulators have been demonstrated to have efficiencies $>80\%$ [7]. A low-dropout linear voltage regulator can then be used to provide a ripple-free V_S . Assuming 70% efficiency of the switching regulator and dropout voltage of 0.2 V, the signaling power would increase by 1.3 mW, which is still smaller than HCM signaling.

With impedance matching of 50Ω to the source and load, the peak-to-peak swing is 250 mV with a common-mode voltage of 250 mV. The channel is terminated to $V_S/2$. The termination resistors are comprised of NMOSs, which are controlled by an impedance-controlled feedback loop. With the target sensitivity of the receiver being $<35 \text{ mV}$, the transceiver can tolerate a channel with 12-dB attenuation (6.5 m of RG58 or 40 cm of FR-4 PCB at 1.8 GHz) by using pre-emphasis.

A divide-by-eight (450-MHz) reference clock (CKref) is distributed to each I/O cell by using a low-jitter clock distribution technique [8]. The clock frequency is multiplied by four with a low-power low-jitter phase-locked loop (PLL) [8]. Multiphased outputs of the PLL are used for data recovery and data transmission. The design targets mesochronous data inputs in which each port has the same data frequency but with variable phase. Since the transmitter (Tx) and receiver (Rx) are not both operating simultaneously, to reduce power and area, they share the same PLL.

The architecture uses separate analog (AV_{DD}) and digital (PV_{DD}) supplies. As described in Section III, the transmitter tracks process, voltage, and temperature (PVT) variations with a

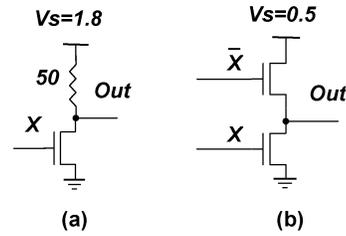


Fig. 4. (a) High common-mode signaling. (b) Low common-mode signaling.

TABLE I
SIMULATED POWER DISSIPATION FOR VARIOUS DRIVER ARCHITECTURES AT 3.6 Gb/s

Power ($V_{Sw}=250\text{mV}$)	Driver Gate	Signaling	Total
HCM (Single-ended) (mW)	0.3	7.874	8.174
HCM (Differential) (mW)	0.3	7.811	8.111
LCM (Single-ended) (mW)	0.824	0.67	1.494
LCM (Differential) (mW)	1.542	0.347	1.889

bias voltage. To further minimize power, the bias voltage is conveniently used to set the supply voltage for the remaining digital logic (PV_{DD}). This supply allows the logic to operate with a constant gate speed regardless of PVT. With the logic designed to operate at maximum V_{DD} in the slow corner, the regulated supply minimizes the power consumption at other corners. To maximize the power savings, a switching regulator controlled by the tracking logic can be used to produce the digital V_{DD} of the chip. A low-dropout local voltage regulator is included in each I/O cell to provide a ripple-free PV_{DD} to the digital part.

In Fig. 5, the two modes of operation, transmit and receive, are shown. In the transmit mode, the Rx block is disabled. To ensure impedance matching with the channel and reduce the signal reflection at the source end of the channel, the driver is source terminated for both pull-up and pull-down. In the receive mode, the transmitter is programmed to turn on both pull-up and pull-down with half the device width (each 100Ω). With the voltage divider formed by two $100\text{-}\Omega$ resistors, the equivalent input impedance of the receiver is 50Ω , biased at 0.25-V dc. V_{ref} is terminated with the same network as the signal pin with the intention of matching the frequency response of the reference and the data inputs. Because the two paths will not be ideally matched, some on-chip common-mode noise will convert into differential noise. Bandwidth control of the receiver is introduced Section IV to further reduce the noise.

III. TRANSMITTER

The transmitter architecture is shown in Fig. 6. The input data is at 225 Mb/s to ease the system testing. A 16:1 multiplexer serializes the input data into 3.6-Gb/s data at the output. The 16:1 multiplexer is a binary tree of 2:1 multiplexers. The clock signals that drive the multiplexers are 1.8 GHz, 900 MHz, 450 MHz, and 225 MHz. Each clock is divided down from a 1.8-GHz PLL clock.

Fig. 7(b) shows the schematic of the output driver (Drv). The source impedance is adjusted by adapting the pre-driver supply, PV_{DD} . However, using the same PV_{DD} for both up and down paths yields different impedances. As a result, a bottom NMOS M_1 is added to match the pull-down impedance across different

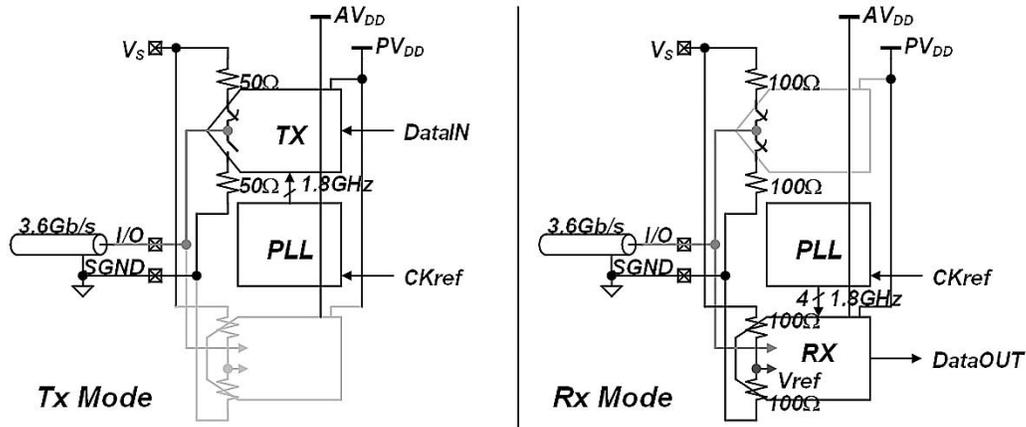


Fig. 5. Operations of transceiver.

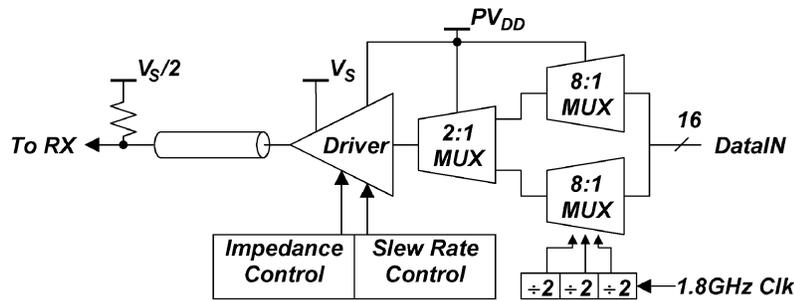


Fig. 6. Transmitter block diagram.

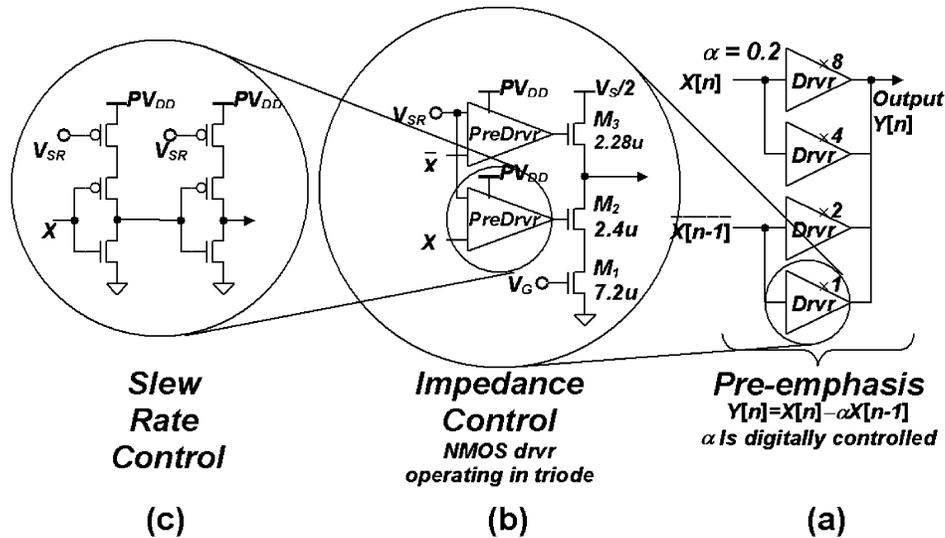


Fig. 7. Schematics of transmitter driver.

processes. The gate of M_1 is connected to a voltage V_G that is independently controlled by a second feedback loop. Fig. 8 shows the impedance controller, which consists of two loops for the up and down impedances. PV_{DD-ref} is distributed to the input of the low-dropout linear regulator. Although an efficient switching regulator [7] is not included in the design, the impedance controller can produce a second voltage, V_X , controlling the regulator. The voltage is higher than the PV_{DD-ref} by the drop-out voltage. The control loop uses an external 50- Ω

resistor as reference. The control voltages are shared among all I/O ports to amortize the power of the control loops.

This paper introduces a novel two-tap pre-emphasis filter for a voltage-mode driver without sacrificing the output impedance matching. The goal is to implement a high-pass filter given by

$$Y[n] = X[n] - \alpha X[n-1]. \quad (1)$$

Increasing the number of taps is feasible in this architecture but has diminishing returns for power dissipation for short channels

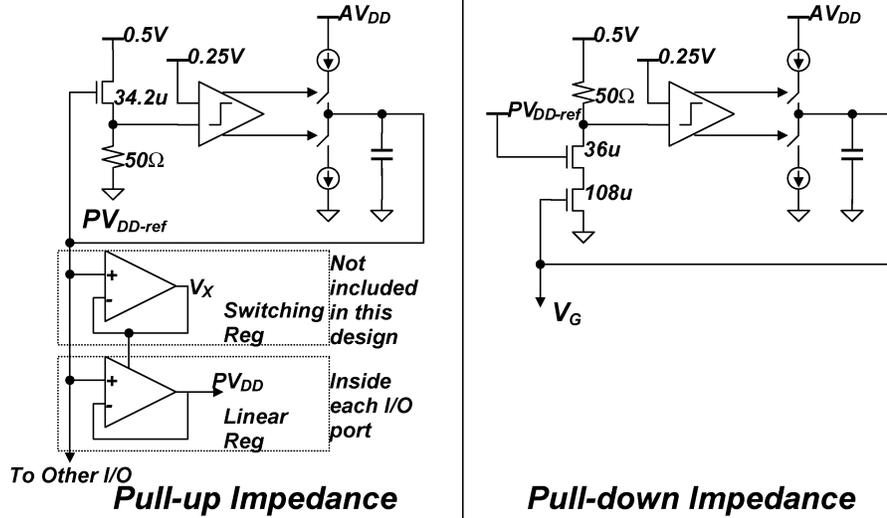


Fig. 8. Impedance controller.

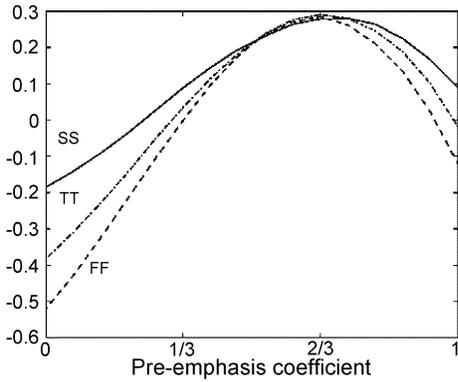


Fig. 9. Simulated DNL of output driver.

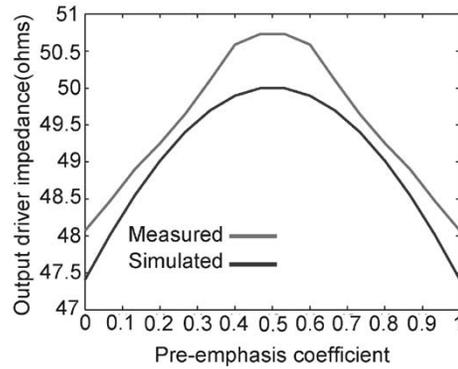


Fig. 10. Simulated and measured output driver impedance.

[1]. To implement $X[n - 1]$, the data is delayed by a half-cycle of the 1.8-GHz clock.

To drive an analog output specified by (1), the entire output driver is divided into four binary-weighted segments, as shown in Fig. 7(a). The output conductance of each segment is directly proportional to the size of the segments. As determined by the digital inputs, each segment either pulls up or down; consequently, the output driver forms a voltage divider, with 16 possible ratios of pull-up and pull-down conductance. These ratios correspond to 16 voltage levels like a digital-to-analog converter. The digital weight determines α of (1). Meanwhile, since all segments are in parallel, the combined output conductance is constant and equals 50Ω regardless of the filter coefficient. Therefore, the driver achieves pre-emphasis while maintaining impedance matching simultaneously. The additional power is the cost of a half-cycle delay and the selection switches.

An example illustrates the output driver operation. Fig. 7 shows the case with $\alpha = 3/15$. Assuming $X[n] = X[n - 1] = 1$, the output voltage is $V_{OUT} = (15R/3)/(15R/3 + 15R/12)V_s = (4 \times V_s)/5 = (1 - \alpha) \times V_s$, where R is the total output impedance (50Ω) of the driver pull-up or pull-down. The output impedance with $\alpha = 3/15$ is $(15R/12) \parallel (15R/3) = R = 50$.

To illustrate the accuracy of generating the 16 voltage levels, Fig. 9 shows a differential nonlinearity (DNL) plot of different simulation corners. The plot shows that the driver linearity is good for the 4 bits. Velocity saturation worsens the driver linearity because the device exits triode region with a lower V_{DS} . In addition, the fast-NMOS, fast-PMOS corner (FF) has worse linearity because V_{GS} is smaller (around 1.2 V), causing the transistor to operate closer to the boundary of the triode and saturation region. In Fig. 10, the simulated output impedance across different pre-emphasis coefficients shows the variation is less than 10%.

To improve signal integrity and minimize simultaneous switching noise, the design limits the output slew rate to roughly 1/3 bit time. A slew rate of 1/3 bit time limits high frequency noise to twice the maximum signal frequency, while introducing only 6% intersymbol interference (ISI) in the signal amplitude, as shown in Fig. 11. In the plot, all axes are normalized to unit bit time and unit output amplitude. The output slew rate is controlled by limiting the pre-driver's slew rate [Fig. 7(c)]. Because the gate capacitance of the driver is relatively process independent, the pre-driver's slew rate is controlled by the drive resistance. An advantage of the architecture is that the NMOS resistance per device width of the pre-driver

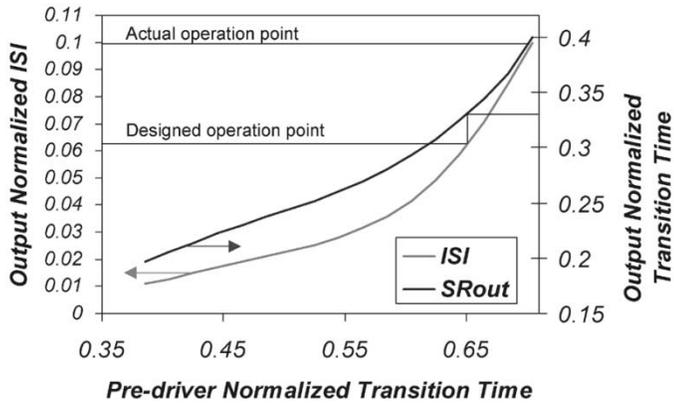


Fig. 11. Normalized ISI due to finite transition time.

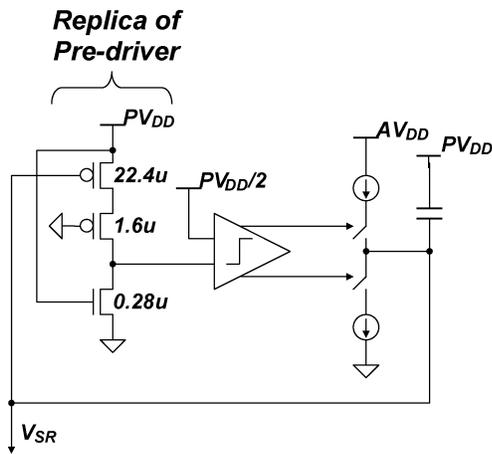


Fig. 12. Slew-rate control loop.

is constant to PVT since PV_{DD} is generated by the impedance control. The ratio of the pre-driver's NMOS size and the driver device capacitance determines the pre-driver's falling slew rate. To control the rising slew rate from the pre-driver, the design uses two PMOSs in series and a control loop. The top device has a voltage V_{SR} from a control loop that maintains a constant total pull-up resistance. The sizing of the top device is significantly larger in order to minimize the capacitance in the signal path (reducing the dynamic power dissipation). The slew-rate control loop (Fig. 12) uses a replica of the pre-driver. The loop turns on both NMOS and PMOS of the dummy pre-driver and adjusts V_{SR} until the output is $PV_{DD}/2$ (equal up and down resistance). The slew-rate control loop is also shared among all transmitters to save power and area.

IV. RECEIVER

The receiver must tolerate noise and amplify the weak input signal to digital levels. To reduce noise without using high bias current, the amplifier bandwidth is limited to reduce total noise power. The minimum bandwidth is the signal bandwidth (1.8 GHz). To reduce the switching power, small device size is used in the front-end samplers. Since small device size leads to significant offsets [9], offset compensation is required to maintain the sampler accuracy. This paper presents a receiver

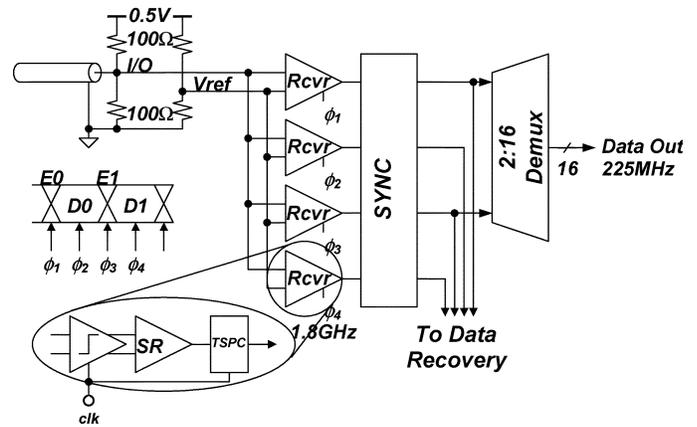


Fig. 13. Receiver block diagram.

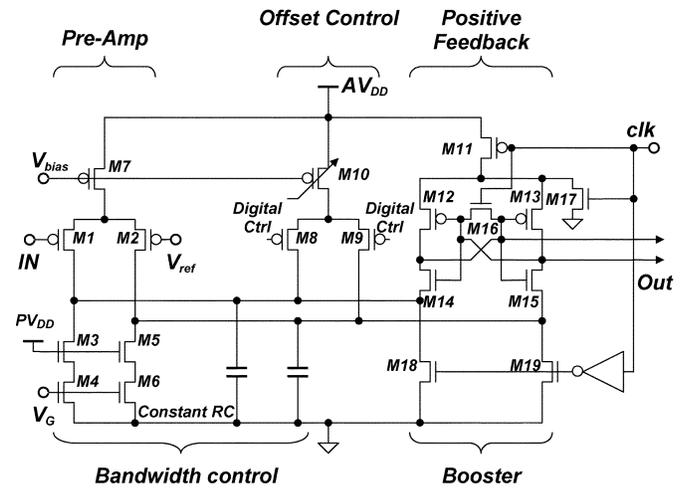


Fig. 14. Schematics of comparator.

that embeds sampling bandwidth control and digital offset compensation with a negligible increase in power consumption.

The block diagram of the receiver is shown in Fig. 13. Four low-power high-speed samplers (*Rcvr*) are used to sample the data with four quadrature phases of the 1.8-GHz clock. Two samples are at the middle of odd and even data eyes, while the other two are at data transitions. Each receiver consists of a comparator, a slew-rate (SR) latch, and a TSPC latch for re-timing. A synchronizer immediately follows to align all the sampled data. The data recovery circuits use all four aligned data to determine the phase of the clock. The recovered data is then passed to the 2:16 demultiplexer to produce 16-bit 225-Mb/s parallel data. Discussion of each major building block follows.

A. Comparator

The target sensitivity of the receiver is <35 mV with input common-mode voltage of 0.25 V. The comparator resolves the sub-35-mV input to digital values at a 1.8-GHz rate (cycle time of six FO-4 inverter delay). The schematic of the comparator is shown in Fig. 14. There are three key components of this comparator: 1) a pre-amplifier that has built-in bandwidth control; 2) an offset-compensation circuit; and 3) a regenerative gain element.

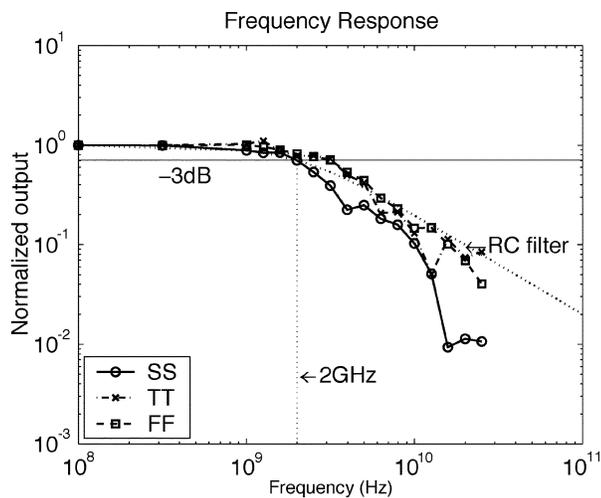


Fig. 15. Frequency response of the comparator.

The pre-amplifier (M1–M7) converts the single-ended input into differential and amplifies the difference of input data (I_N) and the reference voltage (V_{ref}). The PMOS differential pair is used to accommodate the low input common-mode voltage. The tolerable common-mode range of the input devices while maintaining good common-mode rejection is from 0–0.9 V. At high data rate, the pre-amplifier has a gain of less than 2.

Since inputs of the comparator are pseudodifferential, crosstalk or reflections on the signal and common-mode noise from the substrate or V_S , particularly at high frequencies, appear as differential noise at the input. Low-pass filtering has been demonstrated to reduce noise outside the signal bandwidth [10]. In this work, the comparator incorporates a 2-GHz bandwidth filter within the structure. It is set to 10% higher than maximum signal frequency to filter the noise efficiently. The sampling-bandwidth control is built using an RC filter at the output of the pre-amplifier. The bias voltage (PV_{DD} and V_G) that generates the 50- Ω source termination for the transmitter controls NMOS transistors (M3–M6) with scaled-down resistance (≈ 2 k Ω). As a result, the resistance is constant across PVT. With the capacitance relatively constant over different process corners, the RC and, thus, the bandwidth stay relatively constant. The pre-amplifier is active over one half-cycle and is reset over the second half-cycle with transistors M18–M19. The reset phase essentially eliminates any ISI from the previous bits. The simulated frequency responses of different corners are shown in Fig. 15. Because of additional poles from the subsequent comparator, the frequency response rolls off more rapidly than a single-pole filter which in turn provides even better filtering.

Device mismatches of the entire comparator structure are compensated using digital offset compensation. Digital signals control the size of current source (M10) and the steering of the differential pair (M8–M9) to alter the offset of the comparator. The current source M10 is divided into binary weighted segments, which are selected digitally. Since the digitally controlled differential pair M8–M9 fully switches and operates in saturation, the gain and the RC time constant of the pre-amplifier are not disturbed. The offset is calibrated by shorting the differential inputs (shorting devices are not shown), and externally controlling the digital signal until the digital comparator output dithers between zeros and ones.

TABLE II
SIMULATED PERFORMANCE OF RECEIVER

Min. Input Swing	Offset	$< \pm 5$ mV	33 mV
	Supply Noise	$< \pm 8$ mV	
	Thermal Noise	$< \pm 8$ mV	
Power (worst-case estimation)			6.4 mW
Comparator		265 – 301 μ W	
	Pre-Amp	90 μ W	
	Latch+Booster	175 μ W	
	Offset Control	36 μ W	
SR & TSPC latch		185 μ W	
Sync		610 μ W	
Demux		760 μ W	
Clocking		1.84 mW	
Vref generation		1.25 mW	

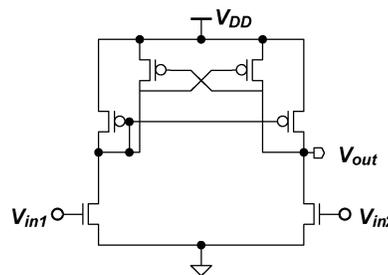


Fig. 16. Schematic of SR latch.

For high signal gain, the pre-amplifier injects differential current into a positive-feedback network (M11–M17). The signal flow is similar to a folded-cascode amplifier. While the pre-amplifier is active (when clk is high), the comparator is reset by disabling the tail current through M11, and equalizing the positive-feedback structure with M16. The switch M17 promotes faster reset by completely turning off positive feedback of M11–M12. When clk is low, M11 provides regeneration current to cross-coupled devices M12–M15. A “booster” (M18–M19) is added to pull the intermediate voltages low. By activating one inverter delay after the positive feedback triggers, it does not impact the signal sensed by the positive feedback circuit. With a low intermediate voltage, the positive feedback regenerates more quickly, and the pre-amplifier is reset, minimizing ISI and data-dependent charge-injection back to the input port. The booster increases the regeneration speed by 10% and reset speed by 15%.

The simulated accuracy is shown in Table II. The estimated device mismatches create ± 75 -mV offset, and it is reduced to ± 5 mV by digital offset compensation. With the bandwidth limited at 2 GHz, the simulated peak thermal noise is ± 8 mV ($\pm 3\sigma$). Because of mixed-mode environments, supply noise can be significant. The differential implementation of the comparator rejects supply noise to the first order. However, large device mismatches imbalances the comparator and degrades the supply-noise rejection. Thus, the supply noise appears as an input-referred noise. Simulation shows that a 10% change in the supply voltage (1.8 ± 0.18 V) can introduce ± 8 mV of input-referred supply noise. Table II also shows the power breakdown of the comparator. Compared with the comparator in [1] with NMOS/PMOS swapped, the new design consumes 30% less power.

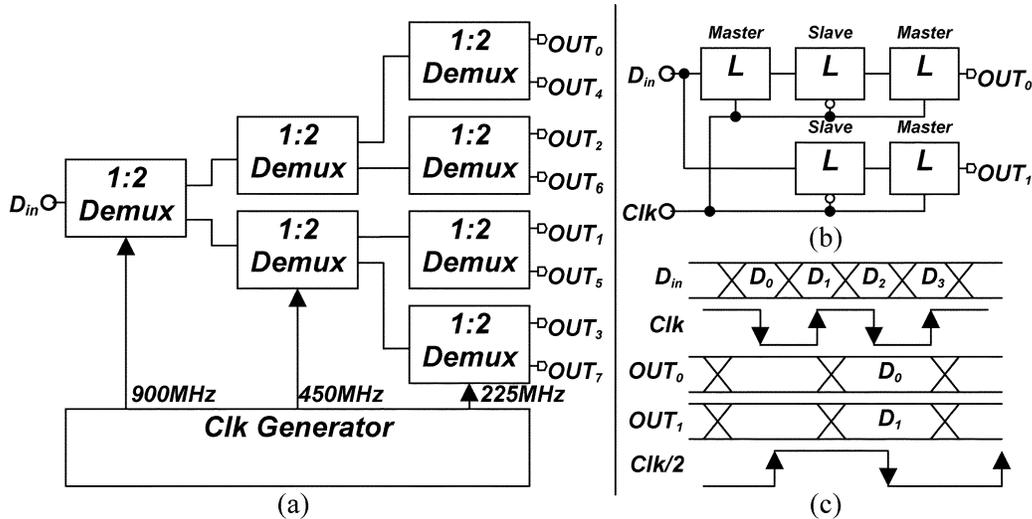


Fig. 17. 2:16 demux.

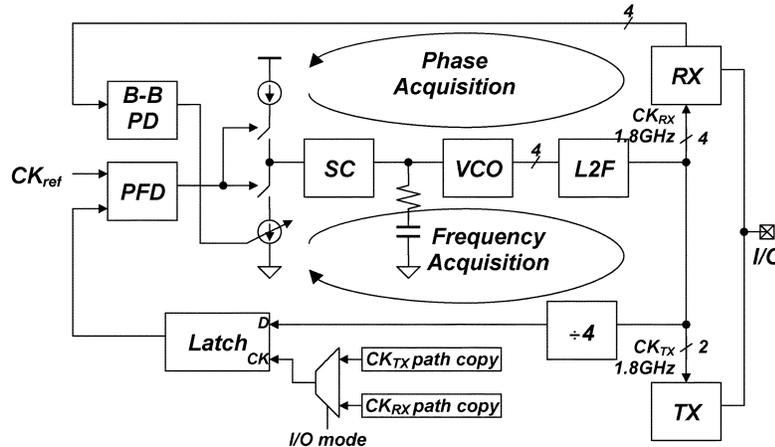


Fig. 18. Data recovery circuits.

B. SR Latch

The SR latch shown in Fig. 16 is used to remove the reset half-cycle from the comparator output and further amplify the signal. The SR latch is designed for proper functionality even at low clock frequency so that the design can be testable. The cross-coupled PMOS load cancels the positive resistance due to the active mirror load and, therefore, keeps the data from discharging even when both inputs are low. After adding the cross-coupled PMOS, the SR latch can hold data from 10 MHz to 2 GHz.

C. 2:16 Demultiplexer and Synchronizer

The 2:16 demultiplexer uses a tree-type architecture [Fig. 17(a)] [11]. Each 1:2 demultiplexer [Fig. 17(b)] consists of two paths. The top path uses an extra latch so that the output data, D_0 and D_1 , are synchronized to the same clock edge. In the synchronizer, the edge samples, E_0 and E_1 , are similarly synchronized to adjust for the quarter-cycle difference. The simulated power summary of the entire receiver is also shown in Table II.

V. DATA-RECOVERY CIRCUIT

The data-recovery circuit is based on a nontraditional dual-loop PLL architecture (Fig. 18). The first loop multiplies the input clock frequency by 4. It generates the 1.8-GHz clock for the transmitter when configured as a driver. The low-power PLL [8] consists of a phase frequency detector (PFD), a charge pump, a switched-capacitor (SC) loop filter, a voltage-controlled oscillator (VCO), low-to-full swing amplifiers (L2F), a frequency divider-by-4, and a retiming latch. The clock paths that drive the clock to the transmitter and receiver are duplicated in the loop's feedback to minimize jitter. Depending on the mode of operation, the appropriate path is selected as the feedback. The secondary loop is active and acquires phase only during receive mode. The secondary loop operates in conjunction with the primary loop. It uses a bang-bang phase detector (B-B PD in the figure) and varies the phase of the output clock by introducing a difference in the up and down currents of the charge pump. A mismatch in the charge pump appears as static phase "error." In this architecture, the error is the desired phase shift. Detailed description of the operation of the bang-bang phase detector and digitally controllable charge pump follows in the next two sections.

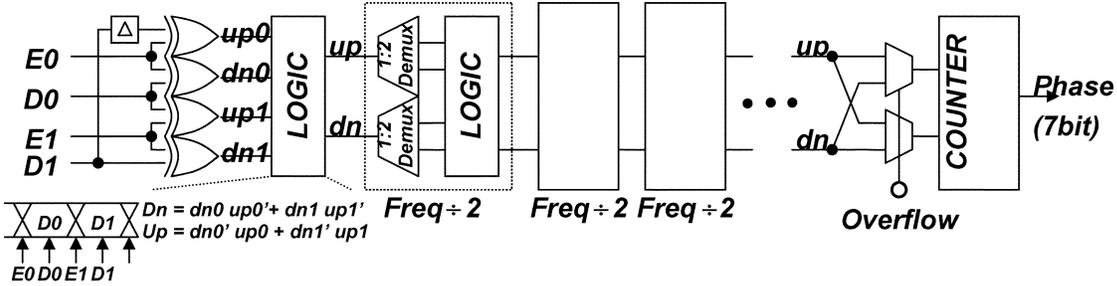


Fig. 19. Bang-bang phase detector.

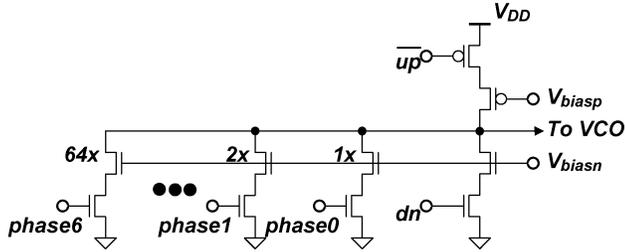


Fig. 20. Schematic of the charge pump.

A. Bang-Bang Phase Detector

A block diagram of the bang-bang phase detector is shown in Fig. 19. The receiver produces two data samples ($D0$, $D1$) and two transition samples ($E0$, $E1$). A bank of four XORs in the phase detector uses the four samples and the delayed version of a data sample to generate two pairs of up/down signals (up/dn). Instead of directly driving a high-speed counter circuit, the up/down signals are first downsampled to a lower rate with a simple logic. This architecture reduces the power consumption of the high-speed up/down signals. The logic, $dn = dn0 \cdot \overline{up0} + dn1 \cdot \overline{up1}$ and $up = dn0' \cdot up0 + dn1' \cdot up1$, combines two pairs of up/down signals into a single pair. The up/down pair is downsampled to half the rate with a pair of 1:2 demultiplexers. The “bandwidth” (or, more precisely, the amount of accumulation) of the bang-bang phase-acquisition loop depends on the update rate of the up/down signals. The proposed architecture uses several stages of the downsampling followed by a low-power counter. A multiplexer is included to program the bandwidth. The nominal bandwidth is set at 1/16 of the data rate. Considering the encoding on the data input, the bandwidth does not degrade the jitter tolerance.

The counter output digitally controls the weight of the charge-pump mismatch. In the case of overflow or underflow by the counter, we avoid a jump from MSB to LSB by reversing the up/down direction so that the counter counts backward.

B. Charge Pump

By controlling charge-pump mismatch, a phase shift of the output clock is introduced. As shown in Fig. 20, the design of the charge pump replaces roughly 1/4 of the pull-down current source with binary-weighted current sources. Since the input clock is 450 MHz, a 25% current mismatch would result in a phase shift of 278 ps (the bit time). The design uses a mismatch slightly larger than 25% to provide a phase adjustment range greater than the bit time. Since the mismatch is a percentage

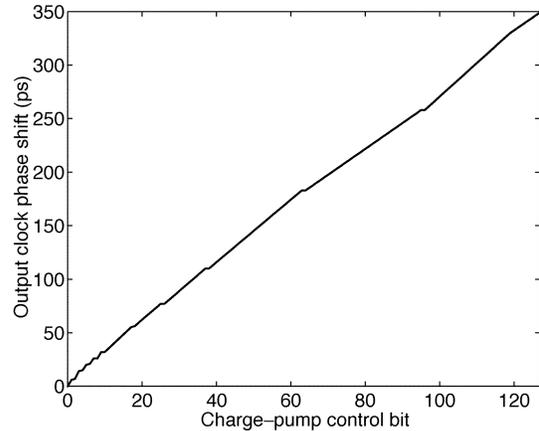


Fig. 21. Linearity of phase shift due to charge-pump control.

of the charge-pump current, the phase range is constant across process corners. Seven bits of phase adjustments are used to guarantee a step size (LSB) of 6 ps. The seven bits are controlled with the outputs of the up/down PD and counter described previously. Fig. 21 shows the measured curve for linearity. The maximum phase shift is 350 ps. The DNL equals one LSB and it occurs at the transition of the MSB. The DNL is mainly due to the device mismatch and the use of binary-weighted current sources. The phase-recovery technique has low power overhead and small area because it does not require additional phase adjustment components such as phase interpolators.

The mismatched charge pump introduces significant ripple at the loop filter. The ripple at 450 MHz would modulate the clock phase in a repeating pattern across four cycles. To reduce the errors, the design uses a switched-capacitor filter after the charge pump (as shown in Fig. 18) to filter the high-frequency modulation. The switching capacitor samples at 450 MHz, creating a notch at 450 MHz. The bandwidth of the switching-capacitor filter does not perturb the PLL response because the PLL bandwidth is <10% of the reference frequency. Fig. 22 shows the simulated eye diagram of the output clock when the charge-pump control bit is set to its maximum, i.e., maximum phase shift. As seen, adding a switch-capacitor filter considerably reduces the phase error.

VI. MEASUREMENT

The transceiver is fabricated in a 0.18- μm CMOS technology. Fig. 23 is the die photo. The entire transceiver occupies an area of 440 μm \times 450 μm . The test chip is packaged with 120-pin

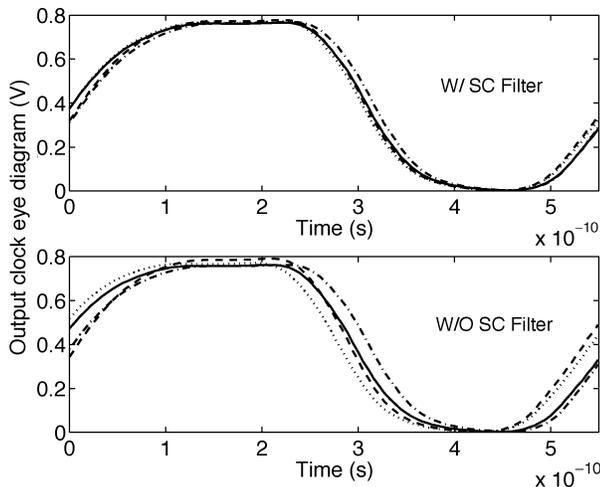


Fig. 22. Simulated output clock when the charge-pump control bit is set to its maximum.

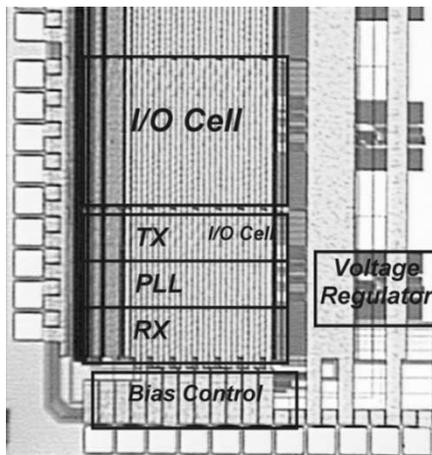


Fig. 23. Die photo.

TQFP. The maximum data rate of the transceiver is 3.6 Gb/s with $AV_{DD} = 1.80$ V and $PV_{DD} = 1.37$ V.

The core transmitter power consumption is 9.66 mW. The power breakdown is shown in Table III. Fig. 24(a) and (b) show the eye diagram after passing through an 8-cm FR4 PCB trace, at 1.6 Gb/s and 3.6 Gb/s respectively. The larger than anticipated ISI at 3.6 Gb/s is primarily due to extra parasitic capacitance of the pre-driver. With the rise and fall times given in Table III, the data from Fig. 11 indicates that ISI is as large as 10% of maximum amplitude. Fig. 25 shows the eye diagram with a channel of an 8-cm FR4 PCB trace followed by 6.5 meters of RG58 cable (total 12 dB loss); the eye is completely closed before equalization. After proper equalization, the eye opening is enlarged to 37 mV (height) and 189 ps (width). The two available taps limit the pre-emphasis performance. To measure the pre-emphasis resolution, a DNL plot with respect to the filter coefficient is measured. The maximum DNL of 0.8 LSB (shown in Fig. 26) indicates that any further increase in number of binary segments may not improve the pre-emphasis accuracy. Fig. 10 illustrates the measured output impedance along with simulated results. The figure shows that the output impedance maintains within 10% variation.

TABLE III
MEASUREMENT RESULTS OF TRANSMITTER AT 3.6 Gb/s

Signaling Power(mW)	0.75
Impedance controller(per I/O)	1.1
Slew rate control(per I/O)	0.1
8:1 Multiplexers	0.4
2:1 MUX and Pre-driver	5.91
Frequency dividers and on-chip PRBS generator	1.4
Total(mW)	9.66
Output data rise time(ps)	112
Output data fall time(ps)	107

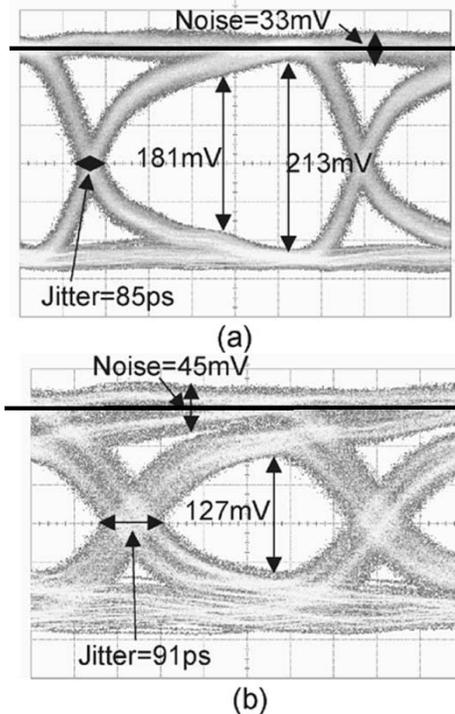


Fig. 24. Eye diagram after 8 cm FR4 PCB trace. (a) 1.6 Gb/s; (b) 3.6 Gb/s.

The timing and voltage margin of the receiver, as shown in Fig. 27, is measured by sweeping the voltage offset and the static phase offset of a clean data input. The plot shows that the minimum input swing is about 35 mV and timing margin is approximately 205 ps. The required input swing is larger than expected due to low-frequency noise on the reference voltage.

To verify switch-capacitor performance, the peak-to-peak jitter at the PLL output clock is measured with the digital scope triggered by the output clock. The measurement results indicate that the p-p jitter is increased by less than 3 ps for entire range of desired phase offset.

The measurement setup cannot collect sufficient data for a bit-error rate. Instead, the error rate is estimated based on noise measurements. Near the center of the data eye, the transmitter has voltage noise of 3.6 mVrms and ISI-induced voltage errors of 63 mV. The overall receiver has an ± 8 -mV offset (± 3 mV from the errors of reference voltage) and a 4-mV rms input-referred noise. The PLL output clock has 6.8-ps rms of jitter. The data-recovery circuit produces 22-ps peak-to-peak dithering when locked, corresponding to 3-LSB steps in the B-B phase detector. Using the eye shape and various noise

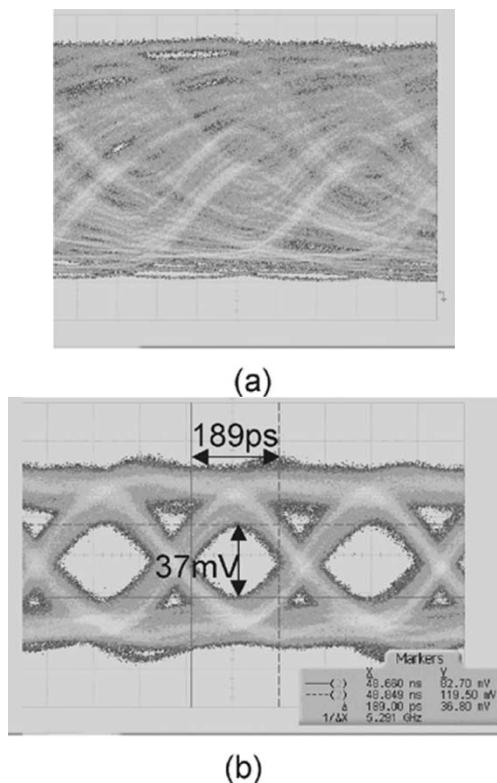


Fig. 25. Eye diagram before and after equalization. (a) Without equalization. (b) With equalization $\alpha = 0.3$.

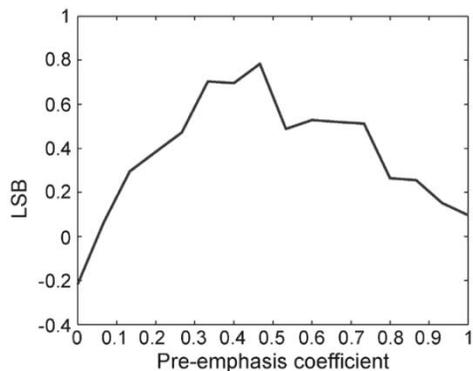


Fig. 26. Pre-emphasis resolution.

sources, the estimated bit error rate is $< 10^{-13}$ for $2^7 - 1$ PRBS data. Measured power consumption of the transceiver components is shown in Table IV. Assuming one transceiver is transmitting and one transceiver is receiving, the complete link dissipates a total active power of 7.5 mW/Gb/s.

VII. CONCLUSION

A 27-mW 3.6-Gb/s parallel I/O transceiver for chip-to-chip applications has been implemented in a 0.18- μm 1.8-V CMOS technology. By comparing the average power per Gb/s operations, this architecture consumed 62.5% less power than the lowest reported so far. The transceiver design demonstrates several circuit and voltage tuning methods of improving signal integrity without excessive power dissipation. For the transmitter, power is significantly reduced using a low common-mode sig-

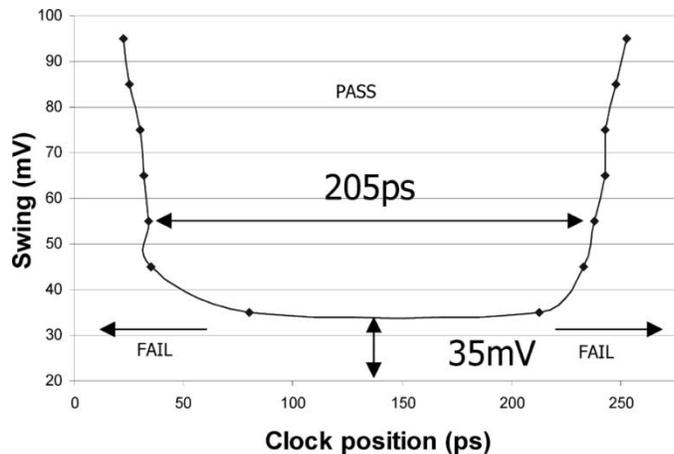


Fig. 27. Receiver timing margin.

TABLE IV
MEASURED POWER OF THE TRANSCIEVER

		Power
Transceiver 1 (TX mode)	TX	9.66mW
	PLL	5.6mW
Transceiver 2 (RX mode)	RX	4.67mW
	PLL	5.6mW
	Phase detector & counter	1.59mW
Total		27.12mW

naling. The dominant remaining power is from pre-driver due to the cost of implementing slew-rate control and pre-emphasis. Similarly, the dominant power consumption of the receiver and timing-recovery design is due to CV^2f power of the comparators, and the phase detectors. The breakdown indicates that over 70% of the power is scalable with technology to the first order.

ACKNOWLEDGMENT

The authors thank National Semiconductor for fabrication.

REFERENCES

- [1] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power, area efficient, high speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591–1599, Nov. 2000.
- [2] F. Yang, J. H. O'Neill, D. Inglis, and J. Othmer, "A CMOS low-power multiple 2.5–3.125-Gb/s serial link macrocell for high IO bandwidth network ICs," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1813–1821, Dec. 2002.
- [3] K.-Y. K. Chang *et al.*, "A 0.4–4-Gb/s CMOS quad transceiver cell using on-chip regulated dual-loop PLLs," *IEEE J. Solid-State Circuits*, vol. 38, pp. 747–754, May 2003.
- [4] J. Kim and M. A. Horowitz, "Adaptive supply serial links with sub-1-V operation and per-pin clock recovery," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1403–1413, Nov. 2002.
- [5] Y. Kudoh, M. Fukaiishi, and M. Mizuno, "A 0.13- μm CMOS 5-Gb/s 10-m 28 AWG cable transceiver with no-feedback-loop continuous-time post-equalization," *IEEE J. Solid-State Circuits*, vol. 38, pp. 741–746, May 2003.
- [6] R. Farjad-Rad *et al.*, "0.622–8.0 Gb/s 150 mW serial IO macrocell with fully flexible preemphasis and equalization," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2003, pp. 63–66.
- [7] A. J. Stratakos, S. R. Sanders, and R. W. Brodersen, "A low-voltage CMOS DC-DC converter for a portable battery-operated system," in *Proc. Power Electronics Specialists Conf.*, vol. 1, 1994, pp. 619–626.
- [8] M. Mansuri and C.-K. K. Yang, "A low-power low-jitter adaptive-bandwidth PLL and clock buffer," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2003, pp. 430–431.

- [9] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, pp. 1433–1440, Oct. 1989.
- [10] S. Sidiropoulos and M. Horowitz, "A 700-Mb/s/pin CMOS signaling interface using current integrating receivers," *IEEE J. Solid-State Circuits*, vol. 32, pp. 681–690, May 1997.
- [11] M. Fukaishi *et al.*, "A 20-Gb/s CMOS multichannel transmitter and receiver chip set for ultra-high-resolution digital displays," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1611–1618, Nov. 2000.



Koon-Lun Jackie Wong was born in Hong Kong. He received the B.S. and M.S. degrees in electrical engineering from the University of California, Los Angeles (UCLA), in 1999 and 2001, respectively. He is currently working toward the Ph.D. degree at UCLA.

He was an intern working on voltage regulators at Broadcom Corporation in summer 1999. In summer 2002, he was with National Semiconductor Corporation working on clock and data recovery for OC-3 applications. He designed high-speed frequency dividers and samplers at IBM in summer 2003.



Hamid Hatamkhani received the B.Sc. and M.Sc. (with highest honors) degrees from Tehran Polytechnic, Tehran, Iran, in 1998 and 2000, respectively. Since January 2001, he has been with the Department of Electrical Engineering, University of California, Los Angeles (UCLA), where he is currently working toward the Ph.D. degree.

His main research interests are high-performance digital and analog integrated circuits design, especially high-speed signaling. In the summer of 2003, he was with Jaalaa Inc., San Diego, CA, working on

the design of power amplifiers for wireless LAN chips.

Mr. Hatamkhani received the Outstanding Student Award from Tehran Polytechnic in 1998. He also received a fellowship from the Department of Electrical Engineering, UCLA, for the first year of graduate study. He served as a scientific committee member of the 1999 Iranian Student Conference on Electrical Engineering.



Mozghan Mansuri (S'97–M'04) received the B.S. and M.S. degrees in electronics engineering from Sharif University of Technology, Tehran, Iran, in 1995 and 1997, respectively, and the Ph.D degree in electrical engineering from the University of California, Los Angeles, in 2003.

She was a Design Engineer with Kavoshgaran Company, Tehran, where she worked on the design of 46–49-MHz cordless and 900-MHz cellular phones from 1997 to 1999. In 2003, she joined Intel Corporation, Hillsboro, OR. Her research interests

include low-power low-jitter clock synthesis/recovery circuits (PLL and DLL) and low-power high-speed I/O links.



Chih-Kong Ken Yang (S'94–M'98) was born in Taipei, Taiwan. He received the B.S. and M.S. degrees in 1992 and the Ph.D. degree in electrical engineering in 1998 from Stanford University, Stanford, CA.

He joined the University of California, Los Angeles, as an Assistant Professor in January, 1999. His current research areas are high-performance mixed-mode circuit design such as clock generation, high-performance I/O, low-power digital design, analog–digital conversion, and low-power high-precision MEMS interface design.

Dr. Yang received the Northrup-Grumman Teaching Award in 2003 and is a IBM Faculty Fellow. He is currently an Associate Editor of the *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II*. He is also a member of Tau Beta Pi and Phi Beta Kappa.