# A Serial-Link Transceiver Based on 8-GSamples/s A/D and D/A Converters in 0.25-$\mu$m CMOS

Chih-Kong Ken Yang, *Member, IEEE*, Vladimir Stojanovic, Siamak Modjtahedi, Mark A. Horowitz, *Fellow, IEEE*, and William F. Ellersick, *Member, IEEE*

***Abstract*—This paper presents a transceiver that uses a 4-bit flash analog-to-digital converter (ADC) for the receiver and an 8-bit current-steering digital-to-analog converter (DAC) for the transmitter. The 8-GSamples/s converters are 8-way time interleaved. Digital compensation reduces the input offset of the ADC comparators to less than 0.6 LSB, improves the accuracy of the interleaved sampling clocks to within 10 ps, and reduces systematic coupling noise to less than 18 mV on the 800-mV signal swing. 1.1-nH bondwire inductors distribute the parasitic capacitances at the transceiver input and output, reducing attenuation by 10 dB at 4 GHz. Equalization algorithms using the converters compensate for the 1.5-GHz transceiver bandwidth to allow 8-GSamples/s multilevel data transmission.**

***Index Terms*—Analog–digital conversion, data transmission, digital–analog conversion, equalization, modulation, offset correction, receiver, serial link, transceiver, transmitter.**

LINKS that communicate multigigabits per second over cables or backplane traces are needed as system bandwidths increase. The 3-dB bandwidth of these copper wires scales with the distance and is approximately 1 GHz for a 10-m coaxial cable or 1-m PCB trace. Thus, the wires filter multigigabit-per-second binary signals causing intersymbol interference (ISI). To achieve very high bit rates, the transceiver compensates for the filtering with equalization and uses multilevel modulation.

On short wires without significant ISI, CMOS links have been reported with bit times as short as the delay of an FO-4 inverter (4 Gb/s in 0.5-$\mu$m CMOS [3], and 10 Gb/s in 0.18-$\mu$m [7]). Higher bit rates using multiple-level (4-PAM) data transmission have been recently demonstrated ([8], [12]) at multigigahertz rates using 2-bit digital-to-analog converters (DACs) and analog-to-digital converters (ADCs) with preprogrammed analog filters for channel equalization. 56Kb and xDSL modems transmit even more bits/symbol, approaching the channel capacity by using high-resolution ADCs and DACs and sophisticated signal processing, but run at much lower data rates. This paper describes two link components, an interleaved 8-bit DAC and an interleaved 4-bit ADC, that achieve a symbol time of a single FO-4 inverter delay, matching the fastest binary transceivers. Converters with this resolution allow the use of more sophisticated communications techniques [6] to continue increasing CMOS link performance.

Relatively simple converter structures achieve the high sample rates. Some of the limitations inherent in these simple architectures are compensated using digital correction. To reduce the low-pass filter caused by the input and output capacitance, the use of inductances to distribute the capacitance is demonstrated in a manner analogous to distributed amplification [5]. The remaining frequency response of the converters and channel are equalized to demonstrate data transmission.

The architecture of the transceiver is discussed in Section I. Section II and Section III describe the circuit design and performance for the receiver and transmitter, respectively. The techniques for digital compensation and equalization are described in Section IV along with the experimental results.

## I. Architecture

Fig. 1 illustrates the transceiver architecture. The receiver consists of eight time-interleaved 4-bit flash ADCs [11], each clocked by a phase-shifted 1-GHz clock. Digitally controllable offset adjustments [4] in each of the comparators compensate for offsets due to device mismatch, reference ladder mismatch, and systematic dc noise. Data received by the interleaved ADCs is demultiplexed and either stored or compared with the contents of a 1024-symbol memory. The transmitter comprises eight time-interleaved 8-bit DACs, also clocked at 1 GHz. The DACs are not expected to achieve eight bits of absolute accuracy, but this resolution allows the limits of the DAC resolution to be explored, and improves overall link performance after digital calibration. A programmable memory stores the sequence to be transmitted, providing flexibility to explore a range of modulation, equalization, and calibration techniques. The ADC and DAC each need Nyquist bandwidth of 4 GHz so that the circuits do not limit link performance.

Capacitance is minimized at the I/O for high data bandwidth. No electrostatic discharge (ESD) structures are included in the design. The input and output each use four pads, with two converters connected to each pad. Long bondwires can be optionally introduced between the pads to add inductance between the capacitance of the interleaved ADCs and DACs, improving the frequency response. These *LC* circuits add signal delay between converters, but these delays can be compensated by phase-shifting the ADC and DAC clocks.

The multiple clock phases used for the receiver and transmitter are generated from two phase-locked loops (PLLs) locked to an external divided-by-4 reference clock running at a nominal

Manuscript received March 20, 2001; revised June 20, 2001. This work was supported by UC Micro Project 99-128 and a Marco grant.

C.-K. K. Yang and S. Modjtahedi are with the University of California, Los Angeles, CA 90095 USA (e-mail: yang@ee.ucla.edu).

V. Stojanovic and M. A. Horowitz are with Stanford University, Stanford, CA 94305 USA.

W. F. Ellersick is with Stanford University, Stanford, CA 94305 USA, and also with Analog Devices Inc., Wilmington, MA 01887 USA

Publisher Item Identifier S 0018-9200(01)08213-0.

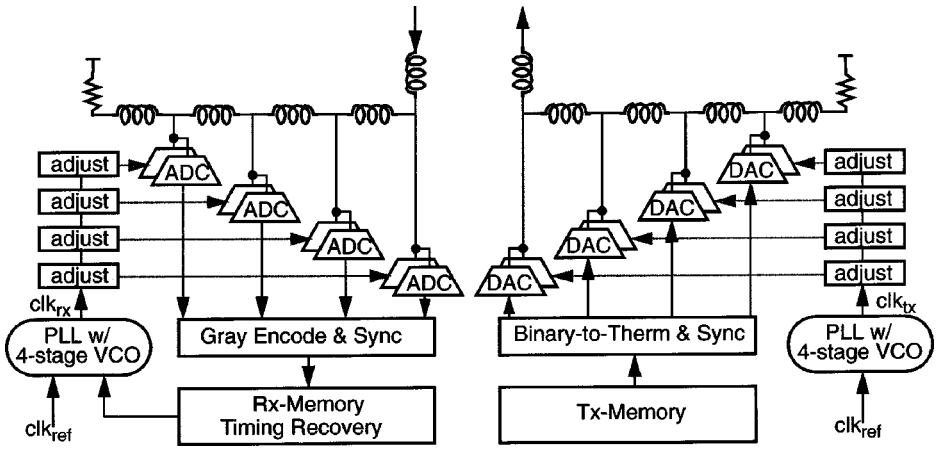


Fig. 1. Transceiver architecture (Rx on the left and Tx on the right).

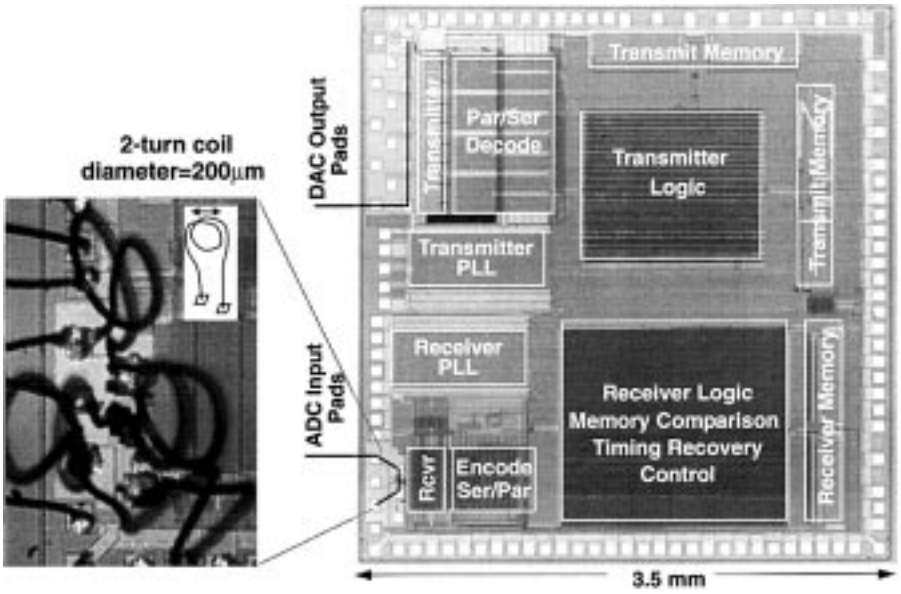


Fig. 2. Chip micrograph with the magnified image of the bonding at the ADC.

250 MHz. Each voltage-controlled oscillator (VCO) [1] uses a ring of four differential buffer stages with the eight internal clock phases tapped and driven to each of the interleaved converters. Phase adjusters are used to compensate for static phase errors from transistor and layout mismatches, and for *LC* delays when inductors are used. Each phase adjuster consists of clock multiplexors and digitally controlled phase interpolators [2]. The interpolators allow 15 interpolation settings (4 bits) between VCO clock phases. At a VCO frequency of 1 GHz (corresponding to 8 GSamples/s), the nominal phase resolution is 8.3 ps ($T_{\rm symbol}/15$). Phase adjustment compensates for transistor and wiring mismatches in the clock path, and for signal delays from the use of inductors in the signal paths.

A die micrograph of the 0.25-$\mu$m CMOS design is shown in Fig. 2. The transceiver system is referenced to $V_{\rm dd}$ with a signal swing of approximately 800 mV, allowing all-nMOS signal paths with maximum headroom. Although the circuits are differential, singled-ended input and output signals are used to allow transmission over inexpensive coaxial cables.

## II. Receiver

The challenge of designing an interleaved receiver with wide bandwidth and 4-bit accuracy is met by using fast sample–hold amplifiers, small transistors to minimize input capacitance, and offset calibration. With a flash architecture, the design of the comparator largely determines ADC performance. As shown in Fig. 3, each comparator contains three stages: a low input capacitance wideband sample–hold amplifier, a clocked latch that regeneratively amplifies the sampled data, and an unclocked amplifier latch. For high bandwidth, minimum-size transistors are used in the sample–hold amplifier. Programmable DACs in the second-stage latch correct for the resulting offset errors with minimal impact on signal bandwidth.

The sample–hold amplifier (Fig. 4) has only unity gain to achieve a high sampling bandwidth of 7 GHz. Because the inputs are applied directly to transistor gates, there is little charge kickback to the input and reference ladder. When clk is high, the circuit acts as a differential amplifier with resistive pMOS

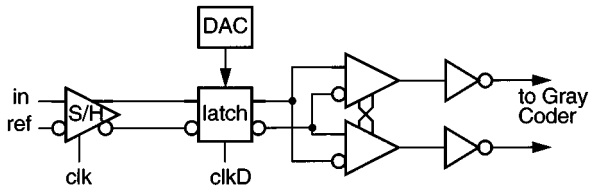


Fig. 3. Comparator block diagram.

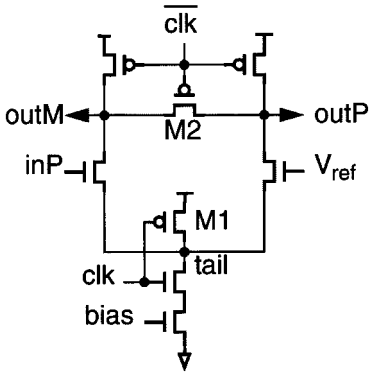


Fig. 4. Sample–hold (first-stage) amplifier design.

loads. M2 serves as a low-resistance differential load with small parasitic capacitance for high bandwidth. When **clk** falls, the current source and load transistors are switched off, holding the sampled value. M1 actively pulls up the tail node to sharply turn off the input transistors and keep the sampling bandwidth close to the raw amplifier bandwidth. The amplifier time constant is low enough to have very little interference from previously sampled values after half a clock cycle (0.5 ns) of settling time, so no reset cycle is needed. The sample time varies little with common-mode input voltage, since the sample–hold amplifier turns off when the **tail** node rises by $V_{gs} - V_{th}$

Both the input and output of the sample–hold amplifier swing from $V_{dd}$ to $V_{dd} - V_{th}$, which maintains the input pair in saturation during amplification. For good linearity during hold, the output must stay in the same range to prevent the input transistors from acting as source followers and keep the pMOS loads fully off. To control the output common mode, a replica bias circuit controls **bias** to set the output swing, and a matched 3/2 inverter chain[10] generates **clk** and $\overline{\mathbf{clk}}$ for each of the time-interleaved ADCs.

The second-stage latch amplifies the sampled value, and compensates for offsets throughout the comparator. As will be shown, even reference ladder errors, clock coupling, and offsets in the subsequent digital flip-flops are corrected. The latch is based on that used in the StrongArm processors [9]. When **clkD** is low, the latch resets and its outputs are pulled high, eliminating interference from previously sampled data. When **clkD** rises, the input difference causes a current difference that feeds into the cross-coupled inverters, causing both outputs to fall, one falling more rapidly. The inverters then regeneratively amplify to full swing. To compensate for offsets, current-mode DACs introduce an offset current into each latch, as shown in Fig. 5. Each DAC consists of two sets of nMOS current sources to provide four bits of compensation resolution. Since the DACs do not load the input, they only slightly affect the input capacitance and regeneration rate. To set the step size of the offset correction DACs, a bias generator outputs $V_{biasDAC}$ of approximately 2 $V_{th}$. $V_{biasDAC}$ increases slightly with $V_{dd}$ to help track the dependence of the latch offsets on $V_{dd}$.

The gain of the second stage is exponentially related to the amplification time, so its output should be sampled just before it is reset. An RS latch [3] or clocked circuits sample about two FO-4 delays before reset due to long sampling times and clock skew. This leaves only two FO-4 delays for amplification, reducing the gain of the second stage enough that the hysteresis and offset of the next stage(s) are significant. The unclocked latch shown in Fig. 6 integrates the second-stage latch output during the entire regeneration time, then holds it during reset and provides additional gain. The latch is based on simple op-amps with pMOS inputs which hold their value when the second-stage latch outputs goes high during reset.

The operation of the unclocked latch can be understood with **inP** high and **inM** low, as denoted with "+" and "−" in Fig. 6. The upper pMOS transistors pull $\overline{\mathbf{oM}}$ and **mirM** high, and **mirM** high pulls $\overline{\mathbf{oP}}$ low and **mirP** low. Then **inM** and **inP** go high when the second-stage latch is reset, turning off the pMOS input transistors and holding the values at $\overline{\mathbf{oM}}$ and $\overline{\mathbf{oP}}$. The cross-coupled nMOS transistors M1 and M2 are needed to avoid degraded output levels. Since the second-stage latch outputs equalize before they are fully reset high, all the pMOS inputs are partially ON during reset. M1 and M2 prevent **mirP** from rising unless **inP** is lower than **inM**, and thus prevent **oM** from being pulled down. M1 and M2 introduces little hysteresis because **mirM** and **mirP** drop below $V_{th}$ during reset of the second-stage latch.

Complementary CMOS drivers drive the long wires out of the comparator array to the Gray coder so there is no net charge transfer out of the comparator, and thus avoid local supply bounce. The receive data is then retimed and encoded with high-speed CMOS logic. A resistor ladder between $V_{dd}$ (the input signal return) and $V_{ref}$ (a supplied reference) generates the 15 reference voltages for the comparators. The ladder uses metal 4 shielded by metal 3. Each step is 66 Ω for a total ladder resistance of 1 kΩ. An active bias circuit (shown in Fig. 7) controls the current for the resistor ladder so that the internal $V_{ref}$ matches that of an external signal. The digital offset correction in the comparator compensates for small mismatches in the ladder resistance.

In the layout, both the input and reference ladder are tightly coupled to the signal return so that noise appears as common mode to the differential comparators. To minimize size (and thus wiring capacitance), the comparator array only contains the analog circuitry of the ADC, with digital encoding performed in a separate block. Comparators with complementary clocks are interleaved, sharing the same input, reference, and supply wires, so that charge kickback and noise on the isolated analog supply are cancelled to first order.

Despite the near-minimum-size devices used, the input capacitance of 1.9 pF results in a pole at 3.3 GHz. This capacitance can be distributed using inductors between every pair of ADCs. By inserting four 1.1-nH inductors, a lumped 50-Ω transmission line is constructed. This distributed ADC technique trades bandwidth for delay—about 100 ps down the entire *LC* line.

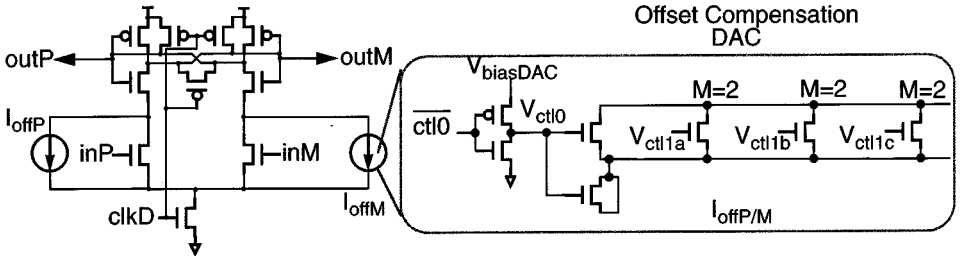


Fig. 5. Clocked offset-compensation (second-stage) latch with design of offset-current DAC.

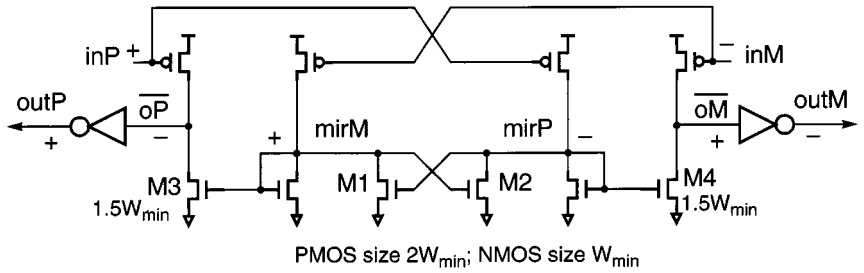


Fig. 6. Output unclocked (third-stage) latch.

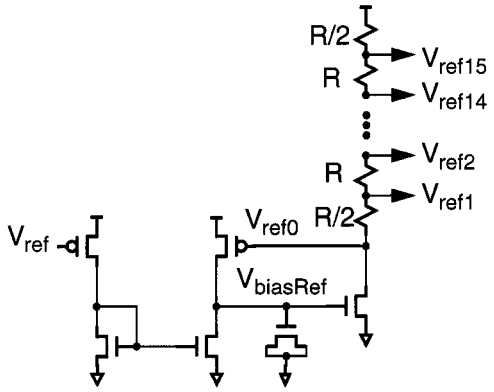


Fig. 7. Reference resistor ladder.

The phase interpolator driving the sampling clock of each ADC can compensate for the delay and optimize the sampling time. In measurements, the phase adjusters reduce static timing errors from 47 to 10 $ps_{pp}$

Fig. 2 shows the ADC input pads on a chip with inductors formed by two-turn coils of bondwire. Implementing these inductors in a flip-chip package is preferable but is beyond the scope of this work. Although precise inductors cannot be bonded, simulation shown in Fig. 8(b) indicates that a 30% variation in inductance is acceptable. Furthermore, because the inductance per unit length depends on the log of physical dimensions, moderate changes in shape and size only result in small changes in the inductance value. For modest control of the 200-$\mu$m diameter of the coils, we wind the bondwires around a 30-gauge wire. Although coupling between the inductors changes the frequency response, it can be compensated with equalization.

The measured frequency response with and without inductors is shown in Fig. 8(a). A bandwidth of 6 GHz with inductors indicates that the bandwidth is primarily limited by the sample–hold amplifier bandwidth. The response is measured by driving asynchronous sine waves into the ADC and fitting ideal sine waves to the data captured in the on-chip memory. The frequency response matches the simulation of a detailed model of the ADC input path, shown on the right, and is flatter with inductors because reflections are reduced.

To achieve four bits of resolution, the offsets in each comparator are compensated by digital calibration. The calibration procedure sweeps a known input voltage at the ADC input to the switching point of each comparator in sequence. The digital value of the 4-bit offset correction DAC in each comparator is adjusted until the comparator outputs toggle. The primary advantage of this calibration method is that no switches or capacitors are needed, allowing simple high-bandwidth circuits to be used. The correction reduces errors from 3 to 0.6 $LSB_{pp}$. Half an LSB of p–p random noise is measured.

## III. Transmitter

The transmitter uses eight interleaved DACs. Each 8-bit DAC uses a current-steering architecture with an array of current sources (Fig. 9). To improve linearity without significant area overhead, the eight bits are divided into three bits (MSB) of thermometer encoded data and five bits (LSB) of binary data. The output current of all the current sources sums at a 50-$\Omega$ resistor which also serves as the termination resistor. Each current source consists of two nMOS transistors in series. The top transistor stays in saturation as long as the output voltage stays above $V_{DSAT}$ of the transistor (800-mV swing). To time-interleave the converters, two clock signals with different clock phases drive the gates. The ON period of the current (the symbol) is the overlap of the two clocks. Symbols are multiplexed sequentially onto the channel by using the appropriate clock phases for each DAC. The predriver of the bottom switch ANDs the data value with its clock to conditionally switch the

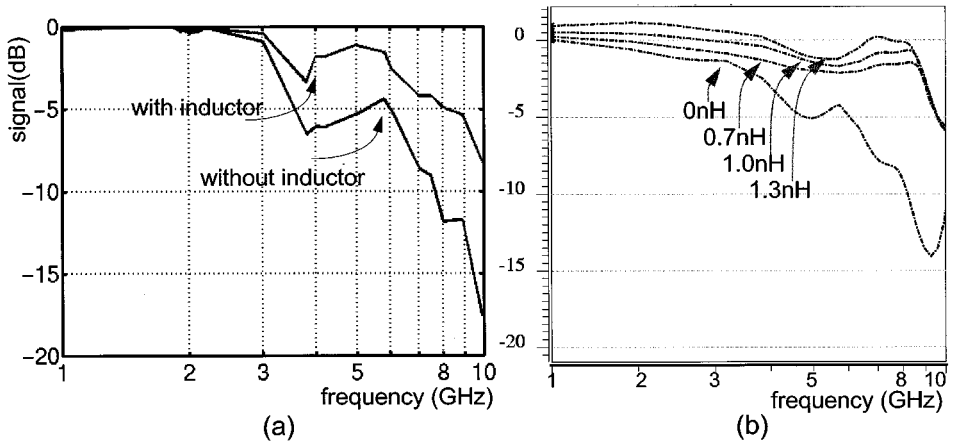


Fig. 8. Receiver bandwidth (a) measured with and without inductors and (b) simulated with various inductances.

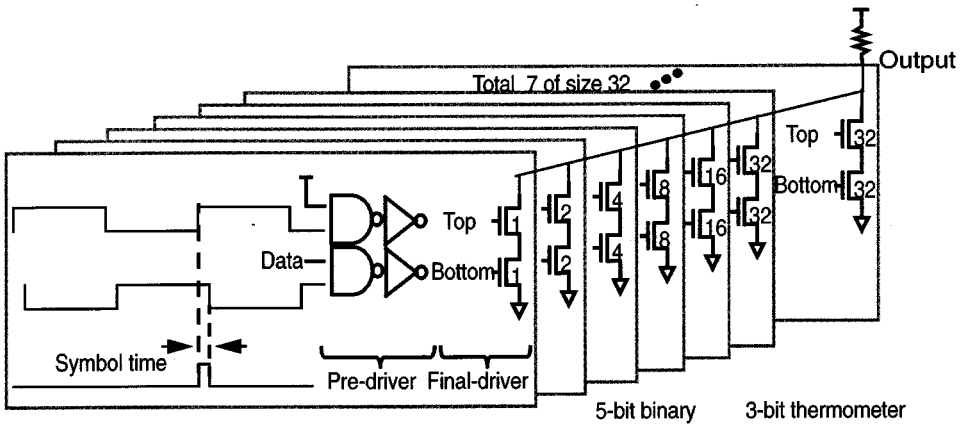


Fig. 9. D/A converter current steering architecture.

current source. A dummy NAND gate drives the top switch to match the bottom predriver delay.

Charge injection and glitches are unavoidable in this simple design. However, the design has reasonable linearity because many of the systematic error sources result in either a dc offset or an error proportional to the output signal. The injection appears as a periodic pattern. As will be shown in Section IV, digitally compensating for this injection is critical. A similar charge injection occurs when the top clock turns ON while the bottom clock is still OFF. The intermediate node capacitance is initially low and then charges to $V_{pdrv} - V_{th}$, where $V_{pdrv}$ is the supply of the predriver (roughly $V_{dd}$). Since $V_{pdrv} - V_{th}$ is constant, the injection is a dc offset. Lastly, the interleaving architecture eliminates nonproportional glitches that occur when switching from one DAC value to another. Handing off between two DACs will cause a glitch due to the overlap (or nonoverlap) of current, but the glitch is linearly proportional to the switched current and hence can be compensated.Section IV describes measurements of the pulse response of each DAC, and compensation with a time-interleaved equalization algorithm.

The clock phases for each DAC that determine the start and stop time of each output pulse are generated with separate programmable phase adjusters. Both the width and phase position of the pulse from each DAC are independently adjustable. As in the ADC, this allows compensation for systematic phase errors such as differences in on-chip clock routing or phase errors from the multiphase VCO. If inductors are inserted, the differences in delay between DACs are also compensated.

The DAC output swing is controlled by the gate voltage of the output current sources. An adjustable on-chip voltage regulator controls the supply of the predrivers. Since the switching speed of the predrivers may limit output bandwidth, the predrivers are inverters that drive a fan-out of less than two.

We take similar layout considerations with the transmitter as the receiver. Clocks are routed to inject noise equally onto all nodes. Dummy devices are used at array boundaries to minimize edge effects. We minimize the output capacitance using minimum size devices. Due to the large currents of 32 mA, we use 40-$\mu$m-wide metal output traces to reduce resistance and handle electromigration. However, the wide traces introduce significant interconnect capacitance.

The primary bandwidth limitation is the *RC* time constant at the output node. Because of the interleaving, the total output capacitance is 4 pF of which 2.5 pF is from the interconnect. As shown in Fig. 10, the output 3-dB bandwidth of 1.6 GHz comprises the dominant pole of the transceiver system. At 4 GHz, the rolloff can be as large as $-11$ dB. By introducing 1-nH inductors, the frequency response is improved above 2.5 GHz, for an attenuation of roughly $-5$ dB at 4 GHz. The 1-nH inductors are too small because the DAC output capacitance was under-

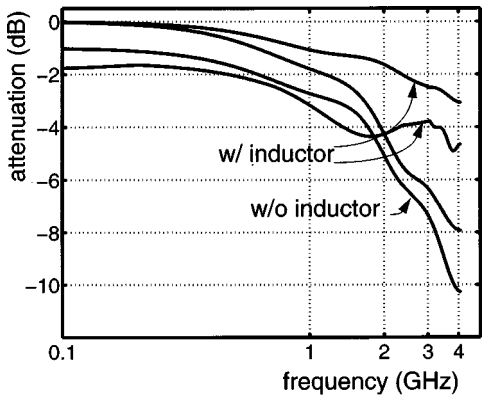


Fig. 10. Measured transmitter frequency response (envelope of all eight DACs) with and without inductor.

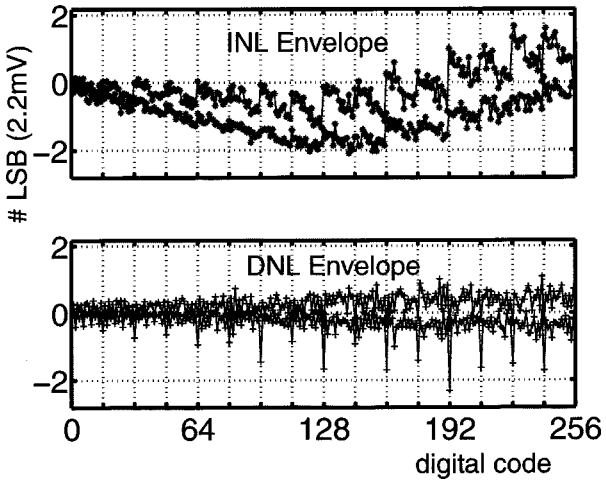


Fig. 11. Measured INL and DNL of all eight D/A converters shown as an envelope (LSB $= 2.2$ mV).

estimated; simulation shows that lower attenuation is achieved with larger inductors.

The nonlinearity of the DACs is characterized by measuring the integral nonlinearity (INL) and differential nonlinearity (DNL). Since each converter has a different characteristic, the nonlinearity is measured by transmitting a pulse of varying amplitudes from each converter. Fig. 11 illustrates measured results from all eight transmitters by plotting the upper and lower bounds of the INL and DNL. The increase in the envelope reflects an increase in noise at the output when transmitting higher digital counts. The proportional increase in noise is due to noise on the power supply of the predriver causing noise on each of the current source in the DAC. The INL exhibits a parabolic shape due to the finite output impedance of the current source. The DNL shows a systematic device mismatch from layout that repeats every 16 counts (fourth bit). This mismatch fortuitously corresponds to overlapping codes and hence only modestly decreases the resolution. The combined INL and DNL indicates a low-frequency resolution of approximately six bits after calibration.

## IV. Experimental Method and Results

Transceiver calibration involves the correction of four main noise sources: receiver offset (discussed previously in Section II), static noise coupling onto the signal, static timing errors in the clock phases, and intersymbol interference (ISI).

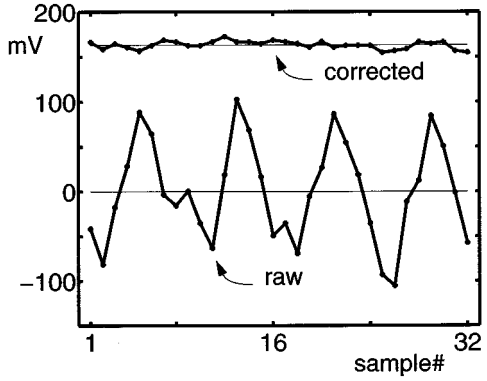


Fig. 12. DC offsets of transceiver with inductors before and after calibration.

Calibrating the static coupling and timing errors is complex because the coupling noise varies with the phase position of the sampling/drive clocks. This requires an iterative procedure of phase tuning and coupling cancellation. Initially, the clocks are adjusted to create the proper pulsewidth and location. The coupling noise is then compensated with a decision-directed adaptation by digitally adding a small signal that is equal to a constant minus the clock coupling. Since the coupling noise correction may cause a small error in timing, i.e., the position of the pulse and width could change, the output pulse timing is then readjusted. The process repeats for each DAC sample position for 32 symbols, which is the duration of the divide-by-four PLL reference clock and the period of the coupling noise. This process generally takes only one iteration.

A sequence of all zeroes before and after calibration is shown in Fig. 12 for the transceiver with inductor bonding. In the noninductor case, noise of roughly 36 $\mathrm{mV_{p-p}}$ can be cancelled to 6 $\mathrm{mV_{p-p}}$. Coupling noise increases significantly in the inductor case since long bondwires have much bigger mutual inductance, thereby effectively acting as antennas. In this case, the calibration procedure is critical in reducing the coupling from 208 to 18 $\mathrm{mV_{p-p}}$. Moreover, the higher dc resolution of the transmitter and the programmability over a long sequence allow us to program the transmitter to compensate for residual periodic offset errors in the receiver. By sweeping the interpolator setting while sampling a pulse, receiver clock phases are tuned to sample at the middle of the symbol.

Once the static offsets are cancelled, the pulse response of each time-interleaved channel is characterized to determine the equalizer weights for predistortion of each time-interleaved converter. Parasitic filtering spreads the pulse response over several symbol times causing ISI. Since the pulse amplitude also depends on the width of the pulse, phase-interpolation quantization errors cause pulse gain mismatches. Also, with inductors, each time-interleaved channel has a different frequency response depending on the position of the DACs and ADCs in the chains of inductors. By finding the pulse response from each time-interleaved DAC to all the ADCs, an interleaved channel matrix is obtained. This not only allows the use of standard zero-force equalizer (ZFE) or minimum-mean-square-error equalizer (MMSE) algorithms to formulate the equalizer matrix

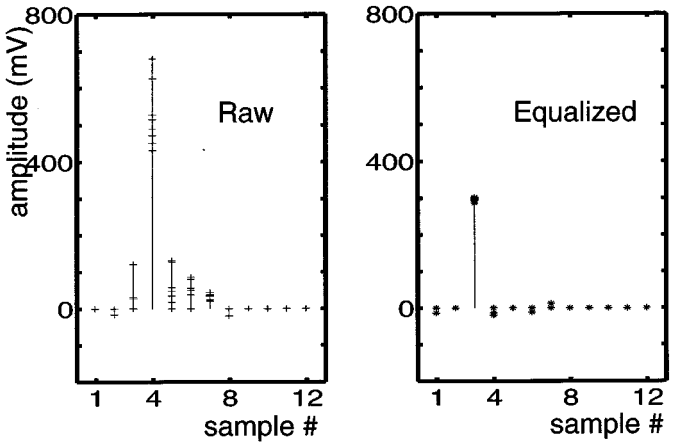


Fig. 13. Measured time-domain pulse response of the transceiver channel with inductor as sampled by the receiver both equalized and unequalized.

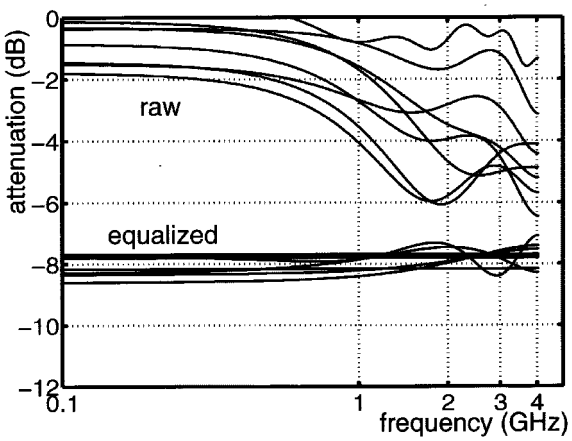


Fig. 14. Measured frequency response of the transceiver channel with inductor both equalized and unequalized.

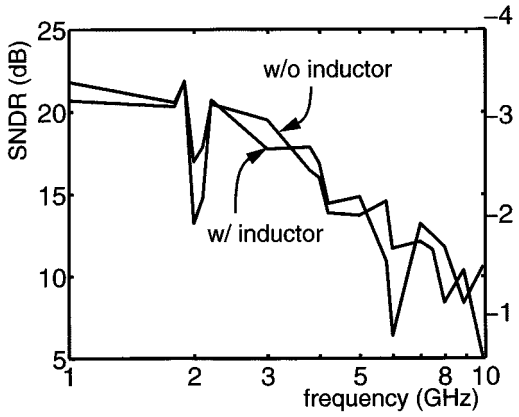


Fig. 15. Measured receiver SNDR with and without inductor.

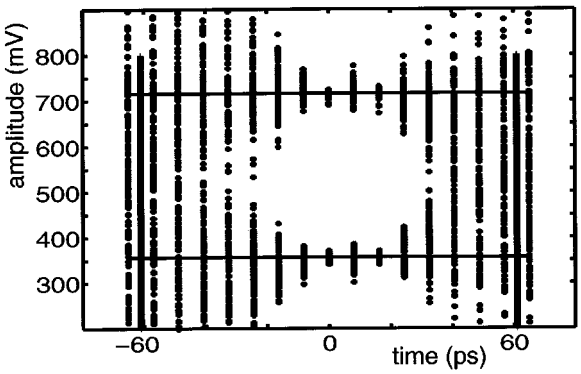


Fig. 16. Schmoo plot of binary data eye at 8 GSymbol/s.

and minimize the ISI, but also allows DAC gain mismatches to be corrected. To avoid quantization error at the receiver, the pulse responses are measured by adding a varying dc voltage to the signal to find the switching point of the MSB comparator in each time-interleaved ADC.

Time-domain pulse responses and the corresponding frequency responses with inductors are shown in Figs. 13 and 14, respectively. The downside of transmitter pre-emphasis-based equalization is that the dc portion of the signal is attenuated, decreasing the swing at receiver (high frequencies cannot be amplified since swing of the transmitter is limited). As the signal amplitude reduces, the clock coupling compensation described earlier becomes even more important in maintaining a clean signal. The residual error is due to uncompensated dc noise, transmitter nonlinearity, pulse-response measurement error, and the finite-length equalizer.

The resolution and jitter of the multiphase VCO significantly impacts performance. The interpolator digital control can adjust the output phase with nominal steps of 8.3 ps and maximum steps of 10 ps. This results in maximum phase error of $\pm 5$ ps, which is about $\pm 5\%$ of a symbol time. The experimental measurements were hampered by an error in layout that caused the PLL to be severely underdamped. This led to larger than expected jitter numbers and a limited operating range. The jitter varies with operating frequency and is at best 28.3 $\text{ps}_{\text{p-p}}$ ($\sigma = 3.1$ ps) for the receiver and 21 $\text{ps}_{\text{p-p}}$ ($\sigma = 2.5$ ps) for the transmitter, with statistically combined jitter of 40 $\text{ps}_{\text{p-p}}$ ($\sigma = 4.0$ ps) for the transceiver, or 32% of a symbol time. The effect of jitter increases proportionally with signal bandwidth and limits the signal-to-noise ratio (SNR) to 17.5 dB at the Nyquist rate of 4 GHz in the transceiver.

The signal-to-noise-plus-distortion ratio (SNDR) of the converters is limited by the jitter. Fig. 15 shows a resolution bandwidth of only roughly 3 GHz for the ADC. The SNDR degrades by only 2 dB over 10% supply change and 0 °C–70 °C temperature variations, indicating that the calibration is robust. Inductors do not improve the SNDR because an increase in the distortion from jitter accompanies the larger received signal. Furthermore, the SNDR drops at 1 and 2 GHz are believed to be caused by the ADC input coupling onto the reference clock, being subsampled by the PLL and exciting its large sensitivity to noise at the loop bandwidth. Similarly, the transmitter SNDR is limited to roughly 16 dB at 4 GHz near the jitter limit.

Despite the large jitter, transceiver operation at 8 GSymbol/s is demonstrated by a schmoo plot that varies both the reference voltage and the sampling phase of the receiver to trace the received data eye. Fig. 16 shows the schmoo plot for binary transmission with inductors with an eye opening of 60-ps width (48% of the bit time) and a 300-mV eye height. Without inductors,

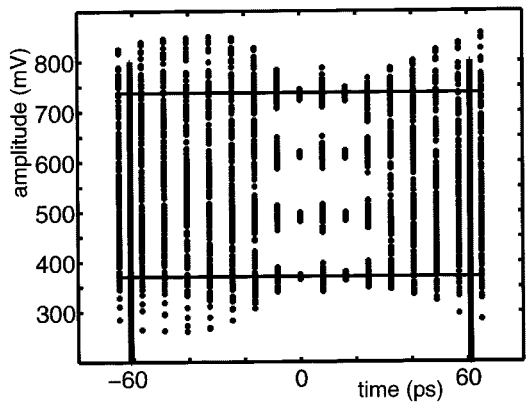


Fig. 17. Schmoo plot of 4-level (4-PAM) data eye at 8 GSymbol/s.

TABLE I
MEASURED TRANSCEIVER RESULTS

| | |
|---|---|
| Technology | NSC 0.25μm CMOS |
| Supply Voltage | 2.5V |
| Jitter Rx with 20kHits | $28ps_{p\text{-}p}$ (σ=3.1ps) |
| Jitter Tx with 20kHits | $21ps_{p\text{-}p}$ (σ=2.5ps) |
| Static Phase Error | $10ps_{p\text{-}p}$ (47 raw) |
| Input (Rx) Bandwidth | 3.2GHz (w/o L)<br>6GHz (w/ L) |
| Rx Input Offset | $0.6LSB_{p\text{-}p}$ (3 raw) |
| Rx Power | 1.1W |
| Rx Area - ADC<br>Logic + PLL | 0.3 x 0.2 mm<br>1.7 x 3.5 mm |
| Output (Tx) Bandwidth,<br>Attenuation at 4GHz | 1.5GHz<br>10dB (w/o L), 5dB (w/ L) |
| Tx DC Compensation | $6LSB_{p\text{-}p}$ (63 raw) |
| Tx Power | 1.2W |
| Tx Area - DAC<br>Logic + PLL | 0.9 x 0.1 mm<br>1.8 x 3.5 mm |

the transmitter bandwidth of 2 GHz caused severe attenuation at 4-GHz Nyquist frequency and no eye was obtained. The schmoo plot for the 4-PAM sequence, inductor case, is illustrated in Fig. 17, showing an eye opening of 45-ps width and 100-mV height. These figures show the worst-case eye by overlaying the eyes from all eight DACs. Due to high jitter, we were unable to verify the 4-PAM eye for a bit-error rate (BER) less than $10^{-4}$, but binary operation has been verified for BER $< 10^{-10}$. Table I summarizes the measurement results.

## V. CONCLUSION

This work demonstrates the feasibility of transmitting and receiving data with interleaved ADCs and DACs that operate at 8 GSamples/s. This opens the door to more complex signaling and equalization to increase link bit rates in band-limited transmission channels. Distributing parasitic capacitances with inductors reduces link attenuation by 10 dB at the Nyquist frequency. Offsets due to processing variations and systematic noise are corrected by designing each transceiver component with digital programmability. The receiver comparator uses four bits of correction in each comparator without switches or capacitors, maintaining high bandwidth. Digital phase interpolators allow timing adjustments for the driving and sampling clocks to compensate for signal delays due to the inductors and clock driver and wiring mismatches. The transmitter uses extra resolution to transmit precompensated sequences and to cancel dc noise patterns. The combined jitter of 40 ps from both underdamped PLLs is left as the predominant cause of signal degradation. The results indicate that with less jitter, interleaved data converters can enable sophisticated communications techniques to further increase CMOS link data rates.

## ACKNOWLEDGMENT

The authors wish to thank K. Chang, A. Emami, D. Liu, and J. Wong for assistance in developing the chip. They also acknowledge National Semiconductor for fabrication. They especially thank P. Prather for the beautiful bondwire inductors.

## REFERENCES

[1] J. Maneatis and M. A. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1273–1282, Dec. 1993.
[2] D. Weinlader *et al.*, "An eight channel 36-GS/s CMOS timing analyzer," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2000, pp. 170–171.
[3] C. K. Yang *et al.*, "A 0.5-$\mu$m CMOS 4.0-Gb/s serial link transceiver with data recovery using oversampling," *IEEE J. Solid-State Circuits*, vol. 33, pp. 713–722, May 1998.
[4] W. Ellersick *et al.*, "GAD: A 12-GS/s CMOS 4-bit A/D converter for an equalized multilevel link," in *Proc. IEEE Symp. VLSI Circuits, Dig. Tech. Papers*, Kyoto, Japan, June 1999, pp. 49–52.
[5] E. Ginzton and W. Hewlett *et al.*, "Distributed amplification," *Proc. IRE*, pp. 956–969, Aug. 1948.
[6] J. Proakis and M. Salehi, *Communication Systems Engineering*. Englewood Cliffs, NJ: Prentice Hall, 1994.
[7] J. Savoj and B. Razavi, "A 10-Gb/s CMOS clock and data recovery circuit with frequency detection," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2001, pp. 78–79.
[8] R. Farjad-Rad *et al.*, "A 0.3-mm CMOS 8-GS/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 757–764, May 2000.
[9] J. Montanaro *et al.*, "A 160-MHz 32-b 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1703–1714, Nov. 1996.
[10] M. Shoji *et al.*, "Elimination of process-dependent clock skew in CMOS VLSI," *IEEE J. Solid-State Circuits*, vol. SC-21, pp. 875–880, Oct. 1986.
[11] W. C. Black and D. A. Hodges, "Time-interleaved converter arrays," *IEEE J. Solid-State Circuits*, vol. SC-15, pp. 1022–1029, Dec. 1980.
[12] J. Zerbe *et al.*, "A 2-Gb/s/pin 4-PAM parallel bus interfaces with transmit crosstalk cancellation equalization and ingtegrating receivers," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2001, pp. 66–67.
[13] W. Ellersick, "Data converters for high-speed CMOS links," Ph.D. dissertation, Stanford Univ., Stanford, CA.

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

He joined the University of California at Los Angeles as an Assistant Professor in January of 1999. His current research areas include high-speed data and clock-recovery circuits for large digital systems (6–10 Gb/s), low-power digital design, and low-power high-precision MEMS interface design.

Dr. Yang is a member of Tau Beta Pi and Phi Beta Kappa.

**Vladimir Stojanovic** was born in Kragujevac, Serbia, Yugoslavia. He received the Dipl. Ing. degree in 1998 from the University of Belgrade, Yugoslavia, and the M.S. degree in 2000 in electrical engineering from Stanford University, Stanford, CA. He is currently working toward the Ph.D. degree at Stanford University, where he is a member of the VLSI Research Group.

He was a Visiting Scholar with the Advanced Computer Systems Engineering Laboratory, Department of Electrical and Computer Engineering, University of California, Davis, during 1997–1998. His current research interests include design and modeling of CMOS-based electrical and optical interfaces, application of digital communication techniques to high-speed links (equalization, noise cancellation), and high-speed mixed-signal IC design.

**Siamak Modjtahedi** received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley, in 1996. Since 1998, he has been working toward the M.S. degree in electrical engineering at the University of California, Los Angeles.

He then joined the Memory Test Division, Teradyne Inc., Agoura Hills, CA, as a Hardware Design Engineer working on the design of memory test systems. He was with the Analog/Mixed Signal Group at Broadcom Inc., Irvine, CA, from June 2000 until February 2001, working on 0.18-$\mu$m CMOS designs.

**Mark A. 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, and the Ph.D. degree from Stanford University, Stanford, CA.

He is Yahoo Founder's Professor of Electrical Engineering and Computer Sciences and Director of the Computer Systems Laboratory at Stanford University. He is well known for his research in integrated circuit design and VLSI systems. His current research includes multiprocessor design, low-power circuits, memory design, and high-speed links. He is also co-founder of Rambus, Inc., Mountain View, CA.

Dr. Horowitz received the Presidential Young Investigator Award and an IBM Faculty Development Award in 1985. In 1993, he was awarded Best Paper at the International Solid State Circuits Conference.

**William F. Ellersick** (M'84–S'97–M'00) received the B.S.E.E. degree from Rensselaer Polytechnic Institute, Troy, NY, in 1983. He received the M.S.E.E. degree in 1994 from Stanford University, Stanford, CA, where he is currently working toward the Ph.D. degree with research on high-speed CMOS links.

He is currently designing high-speed switch and backplane interface circuits at Analog Devices Inc., Wilmington, MA. He worked on shared-optical networks at Ericcson/Raynet, and developed a metal-cutting robot at Cybermation, Inc. His interests are in circuit design, mixed-signal verification, and communications.