

# A 0.3- $\mu\text{m}$ CMOS 8-Gb/s 4-PAM Serial Link Transceiver

Ramin Farjad-Rad, *Student Member, IEEE*, Chih-Kong Ken Yang, *Member, IEEE*, Mark A. Horowitz, and Thomas H. Lee, *Member, IEEE*

**Abstract**—An 8-Gb/s 0.3- $\mu\text{m}$  CMOS transceiver uses multilevel signaling (4-PAM) and transmit preshaping in combination with receive equalization to reduce intersymbol interference due to channel low-pass effects. High on-chip frequencies are avoided by multiplexing and demultiplexing the data directly at the pads. Timing recovery takes advantage of a novel frequency acquisition scheme and a linear phase-locked loop that achieves a loop bandwidth of 35 MHz, phase margin of 50°, and capture range of 20 MHz without a frequency acquisition aid. The transmitted 8-Gb/s data are successfully detected by the receiver after a 10-m coaxial cable. The  $2 \times 2 \text{ mm}^2$  chip consumes 1.1 W at 8 Gb/s with a 3-V supply.

**Index Terms**—Clock recovery, multi-level signaling, receiver equalizer networks, serial links.

## I. INTRODUCTION

AS THE demand for higher data-rate communication increases, low-cost high-speed serial links using copper cables become more attractive for distances of 1–10 m [1], [2]. For multi-gigabit/second (Gb/s) applications, the data rate is limited by the cable skin-effect loss and the process technology. The 10-m coaxial cable (PE-142LL) used in this work has a  $-3\text{-dB}$  bandwidth of 1.2 GHz. This design differs from existing Gb/s links [1], [2] in its use of a receiver equalizer in combination with a transmitter filter to compensate for the cable characteristics. High on-chip frequencies are avoided by multiplexing and demultiplexing the data directly at the pads. To reduce the symbol rate, a four-level pulse amplitude modulation (4-PAM) is used. A new proportional phase detector for data recovery is proposed, which does not suffer from the stability and bandwidth limitations of traditional bang-bang loops. A novel frequency acquisition architecture enables the receive phase-locked loop (PLL) to lock to the input stream under all process variations. The focus of this paper is the design and implementation of the high-speed link receiver. Details of the transmitter architecture are discussed in [5].

## II. SYSTEM ARCHITECTURE

Implementing truly optimal detection methods (in the information theoretical sense) for multi-Gb/s rates demands high complexity and large area [3]. Instead, square pulses, which can be generated and detected with modest complexity, are




Fig. 1. Two 4-PAM eye diagrams: (a) slow transition and (b) sharp transition.

used here as the basis communication symbols [4]. At rates well above the channel bandwidth, however, square pulses result in severe intersymbol interference (ISI), which reduces the data-eye openings. For a given data rate, the 4-PAM scheme reduces the symbol rate to half compared to a conventional 2-PAM system. This symbol rate reduction lowers not only the ISI in the channel but also the maximum required on-chip clock frequency.

To invert the channel, a pre-emphasis filter at the transmitter and an equalizer at the receiver are used. The pre-emphasis transmitter has a two-tap symbol-spaced finite-impulse response (FIR) filter that is used to cancel the tail of the cable pulse response for two subsequent symbol intervals [5]. The receiver equalizer is a one-tap half-symbol-spaced FIR filter, which is described by the following equation:

$$V_{eq}(n \cdot t_S) = V_i(n \cdot t_S) - \alpha \cdot V_i\left(n \cdot t_S - \frac{t_S}{2}\right) \quad (1)$$

where  $t_S$  is the symbol period or sampling interval. This equalizer, using half-symbol-spaced sample values, can equalize the signal over a frequency range that is double that of the transmitter filter without an aliasing effect. Thus the high-frequency components of the signal that were not compensated by the transmitter filter can be equalized in the receiver. In the time domain, the receiver equalizer sharpens the transition edges of the signal. Sharper transition edges result in a larger timing margin for signal detection, especially in multilevel signaling systems. Fig. 1 shows two eye diagrams for a 4-PAM system with different slew rates; clearly the eye diagram with sharper transition results in a larger eye opening [Fig. 1(b)].

The effects of the receiver and transmitter filters for a 0.2-ns pulse (5 Gsym/s) at the near and far ends of the 10-m channel are shown in Fig. 2. The unfiltered pulse response remains at

Manuscript received August 1, 1999; revised November 29, 1999. This work was supported by the Powell Foundation.

The authors are with the Center for Integrated Systems, Stanford University, Stanford, CA 94305 USA.

Publisher Item Identifier S 0018-9200(00)02989-9.




Fig. 2. Pulse shape with and without filtering.

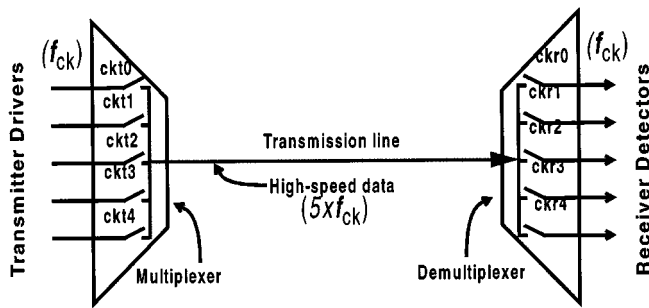


Fig. 3. Multiplexing and demultiplexing the high-speed signal onto the transmission line.

a large value 0.2 ns after its peak (next symbol sample point), while the preshaped equalized signal is zero at that point. All the filter tap weights can be programmed to accommodate different channel characteristics.

The on-chip frequency requirement is further reduced to the symbol rate (1/10 bit rate) by performing 5 : 1 multiplexing and 1 : 5 demultiplexing directly at the chip pads, allowing five symbols to be transmitted every clock cycle [7]. The abstract view of this architecture is shown in Fig. 3. The five symbols correspond to 10 bits that include four data symbols and one symbol for line coding. In this design, coding is performed on-chip to guarantee a high enough transition density for clock recovery.

### III. CIRCUIT IMPLEMENTATION

The block diagram of the complete transceiver chip is depicted in Fig. 4. The transmitter, comprising five identical drivers, uses different clock phases from a five-stage differential ring oscillator (TX-VCO) to multiplex the data stream onto the 50- $\Omega$  line. The detailed transmitter design is described in [5].

The receiver performs 1 : 5 demultiplexing at its input pads by sampling the signal with five out of ten clock phases from a five-stage differential ring oscillator (RX-VCO). The five additional alternate clock phases allow 2 $\times$  oversampling to recover timing and provide required samples for the input equalizer with half-symbol-spaced tap spacings. After equalization, the recovered data samples are converted to bits (binary data) by a bank

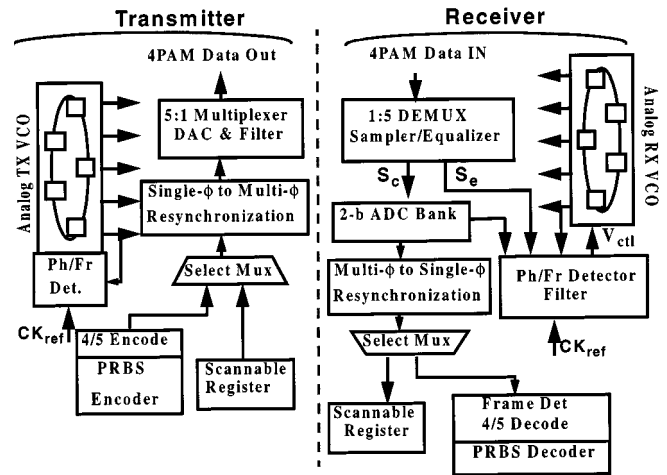


Fig. 4. Transceiver general architecture.

of five 2-bit analog-to-digital converters (ADC's). Finally, the bits in each pack of 10-bit data are pipelined properly and synchronized to a global clock.

A  $2^7 - 1$  pseudorandom bit sequence (PRBS) encoder and decoder, as well as a scannable transmit/receive data register, are provided on-chip for bit error rate (BER) testing. The 5/4 sym decoder removes the extra line-code symbol Fig. 4. As the important functions of the receiver are timing recovery, equalization, and 4-PAM data detection, each of these topics is discussed separately in the following sections.

#### A. Timing Recovery

Timing recovery uses data transitions to adjust the phase of the sampling receiver clocks. There are two main approaches for timing recovery from a serial data: oversampling data recovery and tracking phase detection. In the oversampling technique, each transmitted symbol is sampled  $N$  times ( $N \geq 3$ ), and the sample that is closest to the symbol center is selected by logic as the data [7]. This approach allows very fast timing recovery but suffers from large input loading (due to the large number of samplers) and phase quantization error. Furthermore, it requires complex logic to process many samples at high frequency.

In the tracking phase detection technique, a data phase detector measures the phase difference between the transition edge of the transmitted symbol and the sampling clock. This error value is used to align the sampling point at the symbol center. Traditional proportional tracking data PLL's offer good loop stability and bandwidth, but most suffer from a systematic phase offset. Sampling transitions by the same mechanism as the symbol centers reduces the systematic phase offset in data recovery. However, conventional sampling digital loops use bang-bang control, resulting in limited bandwidth and stability [6]. In this work, we have designed a novel proportional tracking phase detector to overcome these problems.

Fig. 5 shows the receiver 2 $\times$  oversampling front end that is part of the phase detection scheme. When the receive PLL is locked properly to the input data, half of the ten samples represent symbol values at the center of the symbols ( $S_c$ ), and half are samples at the data transitions ( $S_e$ ). The  $S_c$  samples are digitized by 2-bit flash ADC's and result in the received data bits that are

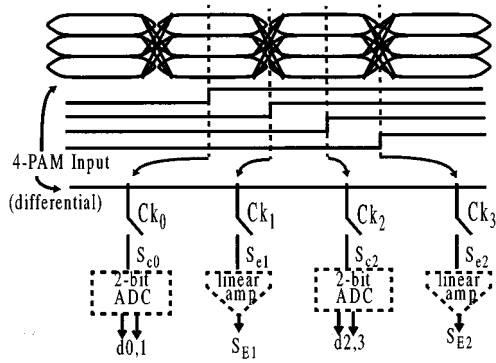


 Fig. 5. Receiver 2 $\times$  oversampling front end.


Fig. 6. Proportional tracking phase detection method (sampling clocks lags the data).

next resynchronized to a global clock. The  $S_e$  samples are amplified by linear amplifiers and kept as analog values ( $S_E$ ) that are used as part of a linear phase detector for timing recovery.

Fig. 6 illustrates the proposed phase detection method for a special case of two-level data and lagging sampling clock. Arrows in the figure show the clock sampling points only at symbol boundaries. When the loop is not in lock and a transition occurs, the edge samples are nonzero and a monotonic function of the phase difference ( $\Delta\phi$ ) between sampling clock edge and data zero crossing. This function can be approximated by a linear function, when the sampling edge occurs within the data transition interval and the loop is near its locking point. Thus, for edge samples within this interval of interest, we have

$$S_e \approx k \cdot \Delta\Phi \quad (2)$$

where  $k$  is the slope of the transition edge. The  $S_e$  values are added together with correct polarity, determined by the direction of each transition, and used to adjust the loop control voltage to correct for the phase error. As the correction on the loop control voltage is proportional to the phase error, this method results in a proportional loop control. Therefore, this PLL combines the advantages of both a linear and a sampling loop. Also, the analog edge samples ( $S_e$ ) at transitions are zero when in lock, resulting in zero sum voltage (no ripple) on the loop control line. In bang-bang control, fixed-amplitude correcting pulses are always applied to the control line that result in ripple and, hence, timing error.

Note that in a differential 4-PAM stream, there are three distinct transition types (Fig. 7). Of these three types, only *type1* makes a transition to the same magnitude but opposite polarity, which results in a zero crossing that occurs exactly at the midpoint between two symbols and that therefore can be used for

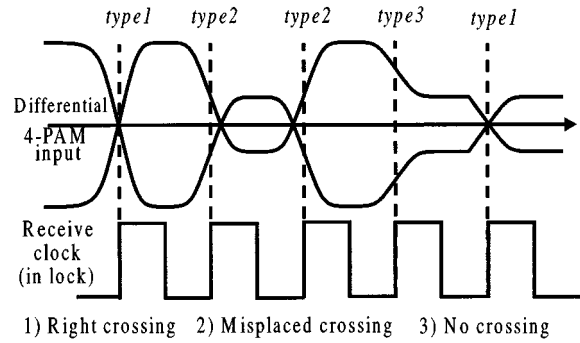


Fig. 7. Three types of transitions in a 4-PAM symbol stream.

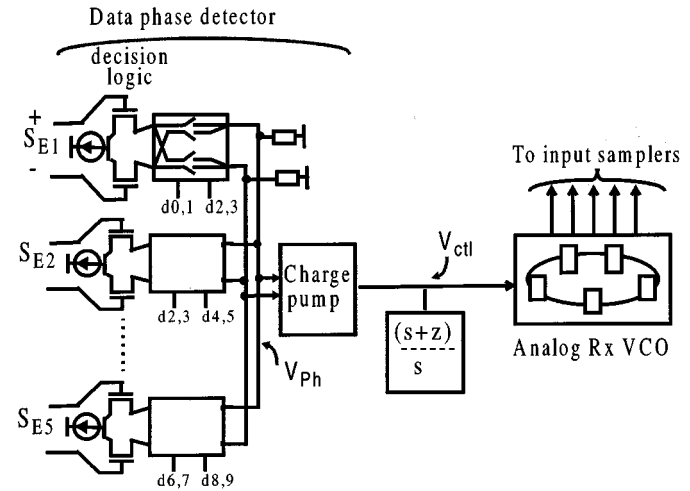


Fig. 8. Proportional data phase detector architecture.

clock recovery. The two other types are ignored as they convey wrong phase information. In every cycle (five symbols), one *type1* transition is guaranteed by the transmitter's 4/5 sym encoder.

Fig. 8 shows the block diagram of the data phase detector that performs the proposed phase detection technique. The five amplified analog edge samples  $S_E$  are each fed into a decision logic block of the phase detector. Based on the two symbol values before and after the transition (2-bit data from ADC's, e.g., d0d1), the phase detector adds the  $S_E$  values of *type1* with correct polarity to the control voltage of the loop and ignores the other two types of transitions by turning off all the switches of that stage. The add/subtract function is done by current summing the differential analog samples with correct polarity at the output of the phase detector ( $V_{ph}$ ).

A charge pump (Fig. 8) converts  $V_{ph}$  to a proportional current using a differential voltage-to-current converter ( $V-I$ ), as shown in Fig. 9. The voltage offsets in the charge pump and phase detector stages directly translate into a phase offset between the sampling clocks and input data. Random offset due to transistor mismatches is reduced by increasing the device sizes and careful layout. The systematic offset of the charge pump ( $V-I$ ) is cancelled using an offset calibration loop that forces the charge pump to inject zero net charge (current) into the loop filter when differential  $V_{ph} = 0$ , as shown in Fig. 9. The calibration circuit has an exact replica of the main  $V-I$ , whose inputs are tied together and set equal to  $V_{ph}$  common-mode voltage




Fig. 9. Charge-pump offset calibration replica circuit.

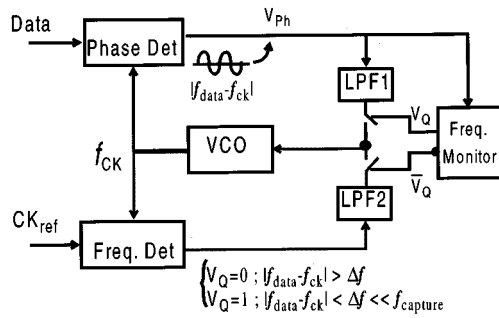


Fig. 10. Frequency acquisition loop for data phase detector.

( $V_{phcom}$ ). The replica  $V-I$  that has a capacitor at output acts as a charge integrator. Thus, the source ( $I_{up}$ ) and sink ( $I_{dn}$ ) currents should be exactly equal to avoid charging the replica circuit output to either of the supply rails. The two  $I_{up}$  and  $I_{dn}$  currents are forced equal by a differential comparator and a current trimming circuit combination (Fig. 9) that compares the replica output to the loop control voltage ( $V_{ctl}$ ) and makes the replica output equal to  $V_{ctl}$  by trimming  $I_{up}$  and  $I_{dn}$ . A replica of the trim currents is applied to the main  $V-I$ . As a result, the loop charge pump generates equal  $I_{up}$  and  $I_{dn}$  when its differential inputs are equal, i.e., when  $V_{ph} = 0$ , or  $V_{ph+} = V_{ph-} = V_{phcom}$ .

To make loop dynamics (gain, bandwidth, phase margin) track process variations and frequency of operation, the loop filter design proposed in [8] is used. As there is no ripple on the loop control voltage when in lock, owing to the phase detector architecture, the loop filter does not require a third-order pole capacitor to damp the control voltage ripple. Therefore, the loop theoretically has only two poles and one zero, and is stable for an infinite range of bandwidths ( $BW > f_{zero}$ ) and loop gains. However, the capacitive loading of the VCO stages on the loop control line introduces a third-order pole that can make the loop unstable for very large gains. The loop gain, and consequently the bandwidth, increases with the number of useful (*type I*) transitions per cycle, and the slew rate of input data signal [ $k$  in (2)]. Using the 4/5 sym encoder, the density of *type I* transition varies from a minimum of one to a maximum of five transitions per clock cycle. The input slew rate ( $k$ ) is determined by the signal amplitude and transition time, which is limited by channel bandwidth. Hence, the loop parameters are

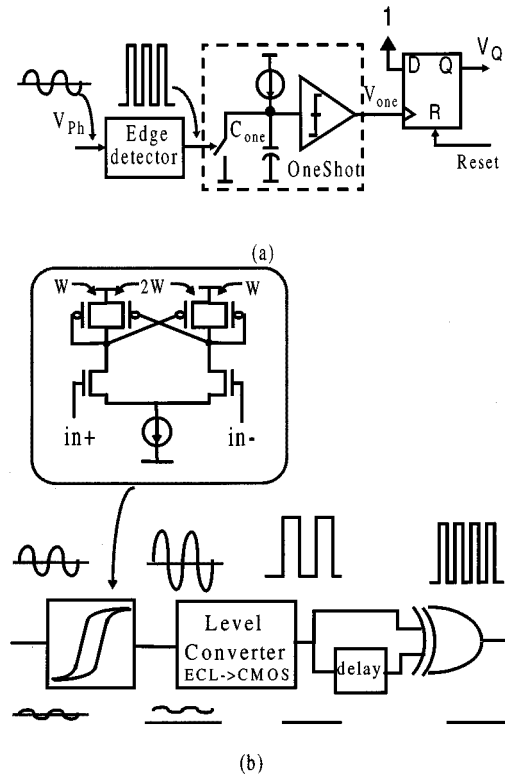


Fig. 11. Frequency monitor: (a) top view and (b) edge detector.

chosen carefully to guarantee a loop bandwidth of  $>20$  MHz and a phase margin  $>45^\circ$  at the worst operating conditions (lowest and highest loop gains). The loop is optimized for a random data sequence with an average *type I* transition density of two per cycle, a differential input amplitude of 1 V (500 mV single-ended), and a risetime of 200 ps. In this condition, the loop has a bandwidth of 35 MHz ( $BW/f_{ref} > 0.07$ ) and phase margin of  $50^\circ$ .

As the phase detector has a limited frequency capture range, a frequency acquisition aid is employed to help acquire lock to a local reference clock at startup (Fig. 10). When the Rx-VCO frequency is different from that of the incoming data, cycle slipping occurs. During cycle slipping, sweeping of the clock phase causes the phase detector output ( $V_{Ph}$ ) to oscillate between *early* and *late* signals. The frequency of this oscillation (sweep speed) is equal to the frequency difference between the receive clocks and the incoming data. A frequency monitor circuit activates the frequency acquisition loop if the frequency difference is large, and activates the data recovery loop (deactivates frequency acquisition) when this difference is smaller than the capture range of the PLL. Fig. 11(a) shows the top view of the frequency monitor circuit. If there is a considerable frequency difference, the oscillations at the phase detector output ( $V_{Ph}$ ) cause the edge detector to produce pulses that continuously discharge  $C_{one}$  and keep the one-shot circuit output ( $V_{one}$ ) at zero. Once the VCO frequency is close enough to the incoming data frequency (within the data PLL capture range), the pulse rate of the edge detector decreases such that  $C_{one}$  can charge high enough to switch  $V_{one}$  to one. At the rising edge of  $V_{one}$ ,  $V_Q$ , which is reset to zero at startup, is asserted and hands loop control over to the data phase detector. The edge detector is de-

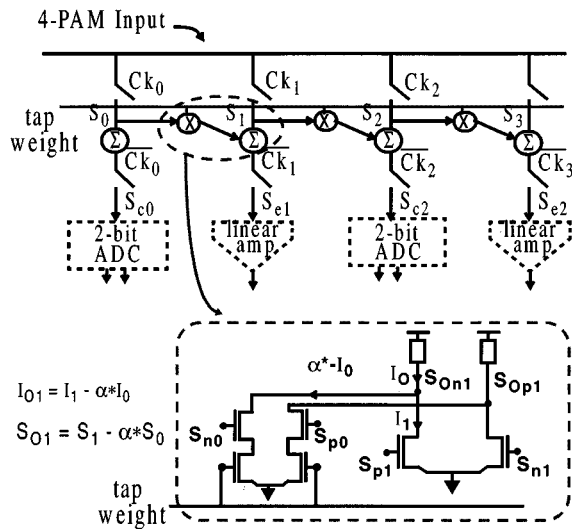


Fig. 12. 1: Five demux samplers and equalizers.






Fig. 14. Simulated eye diagrams at 10 Gb/s for 10-m coaxial cable: (a) no filtering, (b) transmit emphasis, and (c) receiver equalization and transmit preemphasis.

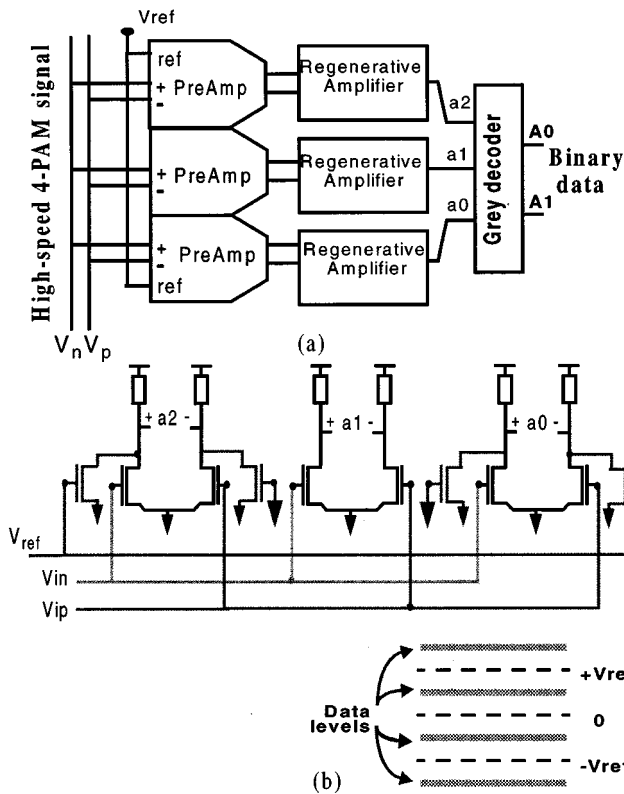


Fig. 13. (a) 2-bit differential flash ADC and (b) differential preamplifier with one reference voltage.

signed to have hysteresis [Fig. 11(b)], using positive feedback in its first stage amplifier. Thus, it reacts only to oscillation amplitudes larger than a certain threshold level, which helps prevent erroneous transitions due to noise.

**B. Equalization**

Since equalization has to be performed at a very high frequency (symbol rate) on each data sample, speed limitations of the process make it impractical to implement this equalizer as a digital FIR filter. Thus, equalization is performed in the analog domain directly on the sampled values before they are used

by other blocks. Fig. 12 shows the architecture of the one-tap half-symbol-space equalizer, where receiver  $2\times$  oversampling provides the required samples.

Having the present and former differential samples, the equalizer subtracts the weighted value of the former sample from the

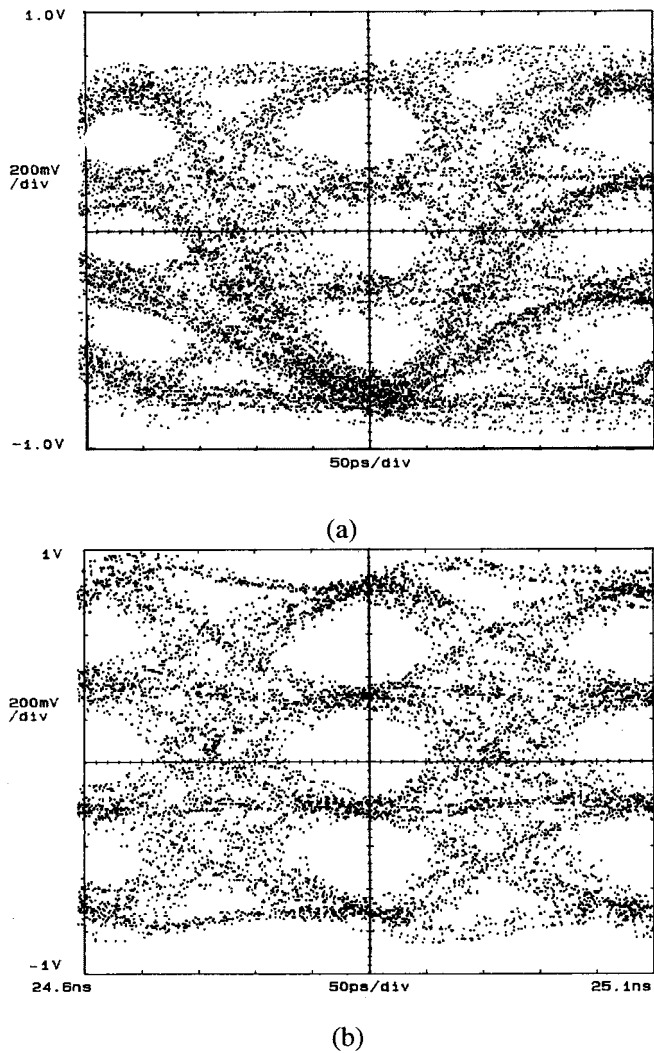


Fig. 15. Differential data-eye over 10-m cable with preemphasis (a) at 10 Gb/s and (b) at 8 Gb/s.

present sample. This operation is done by current summing two differential values with opposite polarity, as shown in Fig. 12. The weighted current is controlled by two tail NMOS transistors that act as a resistor and should therefore operate in the triode region. The equalizer further improves the data eye area (height  $\times$  width) up to 40% by sharpening the signal transitions.

### C. Data Detection

To convert the four-level analog symbols into digital bits, five 2-bit flash ADC's are implemented after the receiver 1:5 demultiplexer (Fig. 13). Each ADC consists of three preamplifiers and regenerative latches, followed by a gray coder that converts thermometer code into binary code, as shown in Fig. 13(a). Using gray coding in the 4-PAM data makes thermometer-to-binary conversion easier as well.

Comparison versus the reference voltage is performed in the preamplification stage. Since a differential signaling scheme is used, only one reference voltage value is required to differentiate among the four input levels, as shown in Fig. 13(b). Also shown is how this reference voltage is applied to the three




Fig. 16. Test setup for BER measurements.

preamplifiers. The middle stage is a balanced differential comparator, and the two other stages are unbalanced by two transistors, which are controlled by  $V_{ref}$ . To balance the output capacitive loading of the preamplifiers, identical dummy transistors, with grounded gates, are placed on the opposite branch of the differential pair.

## IV. MEASUREMENTS

Fig. 14 shows three different simulated eye diagram at 10 Gb/s after the 10-m coaxial cable (PE-142LL). Fig. 14(a) is without transmit preemphasis and receive equalization, Fig. 14(b) is with preemphasis alone, and Fig. 14(c) is both preemphasis and equalization applied. The improvement in the eye diagram in these three conditions shows the necessity of the two filters.

The actual transmitter achieves a symbol rate of 5 Gsym/s (10 Gb/s) with an eye opening of 200 mV and 90 ps, and 4 Gsym/s (8 Gb/s) with an eye opening of 350 mV and 110 ps over 10 m of coaxial cable, using preemphasis (Fig. 15). Symbols without preemphasis after the 10-m cable show an eye opening with 60-mV height and 50-ps width at 4 Gsym/s. The transmitter output has an adjustable amplitude with a maximum of 1.2 V and a jitter of 11 ps (p-p) and 2 ps (rms).

The BER measurements are performed using the test setup shown in Fig. 16. The PRBS encoder in the transmitter generates a 10-Gb/s pseudorandom sequence that is sent over the line. The receiver detects the serial signal from the line and, after proper framing, sends it to the PRBS decoder. Whenever there is a bit error in the received sequence, the PRBS decoder generates an error pulse. The number of these pulses per second is the system BER. The valid data window is measured by connecting the receive and transmit PLL's to two clock sources, as shown in Fig. 16, and varying the delay of one clock source versus the other until a rapid increase occurs in BER. To set the reference voltage ( $V_{ref}$ ) for the receiver 2-bit ADC's (Fig. 13), a differential dc voltage equal to the reference voltage level of the 4-PAM data is applied to the receiver and  $V_{ref}$  is adjusted manually to the point where the ADC's outputs toggle. The receiver successfully detects an 8-Gb/s, 4-PAM data stream after 10 m with a 3-V supply. At data rates higher than 8 Gb/s, the receive PLL fails due to increased high-frequency noise in the loop. Raising the supply to 3.3 V allows the receiver to perform up to 9 Gb/s. The decision logic of the data phase detector injects undesired charge onto the VCO control line, causing error in sampling clock phases and data detection. At 8 Gb/s over 10 m, the receiver had a BER of  $10^{-7}$  for a time window of 50 ps,

TABLE I  
PERFORMANCE SUMMARY

<b>Transmitter performance</b>	
Maximum transmitter rate	10Gb/s @3.3V, 8Gbps @ 3V
Output jitter @ 8Gbps	11ps (p-p), 2ps (rms)
Max. eye opening @8Gbps	350mV, 110ps (10-m cable)
Max. eye opening @10Gbps	200mV, 90ps (10-m cable)
<b>Receiver performance</b>	
Maximum receive rate	9Gbps @3.3V, 8Gbps @3V
Data PLL jitter @ 8Gbps	28ps (p-p), 4ps (rms)
Data PLL capture range	>20MHz
Min. swing to capture lock	$\pm 400\text{mV}$
Min. swing to maintain lock	$\pm 300\text{mV}$
Data PLL dynamics	BW $\sim 35\text{MHz}$ , Ph.m. $\sim 50^\circ$
<b>Power dissipation @ 8Gbps, 3V</b>	
Output driver	220mW
Analog (2 PLLs)	750mW
Input samplers and logic	130mW
Total	1100mW

whereas at 6 Gb/s, the BER decreased to  $10^{-15}$  for a window of 150 ps.

Receiver equalization helps reduce the required transmitter preemphasis for the 10-m cable, effectively allowing the use of longer cables for the link. The receiver equalizer is adjusted manually, as is the transmitter preemphasis filter. However, as opposed to the transmitter output, the equalized waveform in the receiver cannot be viewed and used to set the optimized tap weight value. Therefore, the equalizer tap is adjusted to minimized the measured BER.

The receiver data-recovery PLL requires that the input symbols have a minimum peak-to-peak swing of 800-mV differential (400-mV swing on each line) to acquire lock and 600-mV differential swing to maintain lock. This PLL has a capture range of >20 MHz for a symbol stream with one transition per cycle (five symbols). The frequency acquisition circuit switches the loop control to the data phase detector when there is less than 100-kHz frequency difference between the transmitter and receiver reference clocks. The receive PLL has a jitter of 28 ps (p-p) and 4 ps (rms) when locked to the incoming data signal.

The chip occupies  $2\text{ mm} \times 2\text{ mm}$  of die area. The transceiver die photo is shown in Fig. 17. Table I summarizes the transceiver chip performance.




Fig. 17. Transceiver die photo.

## V. CONCLUSIONS

Using parallelism, 4-PAM modulation, and analog transmit and receive FIR filters, data rates of over 8 Gbps are achievable in conventional CMOS technology over long copper cables. Performance is further enhanced by a novel high-bandwidth linear data-recovery PLL with zero systematic offset that reduces the bit error rate due to random phase errors. A new frequency detector design guarantees frequency acquisition of the data-recovery PLL under all process variations.


## ACKNOWLEDGMENT

The authors would like to thank L. Sampson and S. Krishnan for fabrication assistance, J. Namkoong, K. Yu, A. Hajimiri, K. Falakshahi, and B. Ellersick for helpful discussions, and LSI Logic for fabrication.

## REFERENCES

- [1] J. Poulton and W. J. Dally, "A tracking clock recovery receiver for 4Gb/s signaling," in *Proc. Hot Interconnects Symp.*, Aug. 1997, pp. 157–170.
- [2] A. Fiedler *et al.*, "A 1.0625 Gb/s transceiver with  $2\times$  oversampling and transmit pre-emphasis," in *ISSCC Dig. Tech. Papers*, Feb. 1997, pp. 238–239.
- [3] P. J. Black and T. Meng, "A 1-Gb/s, four-state, sliding block viterbi decoder," *IEEE J. Solid-State Circuits*, vol. 32, pp. 797–805, June 1994.
- [4] R. Farjad-rad *et al.*, "An equalization scheme for 4-PAM signaling over long cables," in *Proc. IEEE CAS Mixed Signal Conf.*, July 1997, pp. 19–22.
- [5] R. Farjad-Rad, C.-K. Yang, M. Horowitz, and T. Lee, "A 0.4- $\mu\text{m}$  CMOS 10-Gb/s 4 PAM serial link pre-emphasis transmitter," *IEEE J. Solid-State Circuits*, vol. 34, pp. 580–585, May 1999.
- [6] T. Hu and P. Gray, "A monolithic 480 Mb/s AGC/decision/clock-recovery circuit in 1.2- $\mu\text{m}$  CMOS," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1312–1318, Dec. 1993.

- [7] C.-K. Yang, R. Farjad, and M. Horowitz, "A 0.5- $\mu$ m CMOS 4-Gb/s serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 33, pp. 713–722, May 1998.
- [8] J. Maneatis and M. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1273–1282, Dec. 1993.



**Ramin Farjad-Rad** (S'95) was born in Tehran, Iran, in 1971. He received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1993 and the M.Sc. degree in electrical engineering from Stanford University, Stanford, CA, in 1995, where he is currently working toward the Ph.D. degree.

He worked at SUN Microsystems Laboratories, Mountain View, CA, on a 1.25-Gb/s serial transceiver for the fiber channel standard during the summer of 1995. During summer 1996, he was with LSI Logic, Milpitas, CA, where he examined different multi-gigabit/s serial transceiver architectures. He has received four U.S. patents.

Mr. Farjad-Rad received the Bronze Medal of the 20th International Physics Olympiad, Warsaw, Poland.



**Chih-Kong Ken Yang** (S'93–M'98) received the B.S. and M.S. degrees in electrical engineering and the Ph.D. degree from Stanford University, Stanford, CA, in 1992 and 1998, respectively.

He is an Assistant Professor at the University of California, Los Angeles. His research interests are in the area of VLSI circuit design with emphasis on high-speed interfaces.

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



**Mark A. Horowitz** 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 Yahoo Founder's Professor of Electrical Engineering and Computer Science at Stanford University. His research area is in digital system design, and he has led a number of processor designs including MIPS-X, one of the first processors to include an on-chip instruction cache; TORCH, a statically scheduled, superscalar processor that supported speculative execution; and FLASH, a flexible DSM machine. He has also worked in a number of other chip design areas, including high-speed and low-power memory design, high-bandwidth interfaces, and fast floating point. In 1990, he took leave from Stanford to help start Rambus Inc., a company designing high-bandwidth memory interface technology. His current research includes multiprocessor design, low-power circuits, memory design, and high-speed links.

Dr. Horowitz received a 1985 Presidential Young Investigator Award and an IBM Faculty Development Award, as well as the 1993 Best Paper Award at the International Solid State Circuits Conference.



**Thomas H. Lee** (M'96) received the S.B., S.M. and Sc.D. degrees in electrical engineering, all from the Massachusetts Institute of Technology (MIT), Cambridge, in 1983, 1985, and 1990, respectively.

He joined Analog Devices in 1990, where he was primarily engaged in the design of high-speed clock recovery devices. In 1992, he joined Rambus, Inc., Mountain View, CA, where he developed high-speed analog circuitry for 500 megabyte/s CMOS DRAM's. He has also contributed to the development of PLL's in the StrongARM, Alpha, and K6/K7 microprocessors. Since 1994, he has been an Assistant Professor of electrical engineering at Stanford University, where his research focus has been on gigahertz-speed wireline and wireless integrated circuits built in conventional silicon technologies, particularly CMOS. He holds 12 U.S. patents and is the author of a textbook, *The Design of CMOS Radio-Frequency Integrated Circuits* (Cambridge, MA: Cambridge Press, 1998), and is a coauthor of two additional books on RF circuit design. He is also a cofounder of Matrix Semiconductor.

Dr. Lee has twice received the "Best Paper" award at the International Solid-State Circuits Conference, was coauthor of a "Best Student Paper" at ISSCC, and recently won a Packard Foundation Fellowship. He is a Distinguished Lecturer of the IEEE Solid-State Circuits Society, and was recently named a Distinguished Microwave Lecturer.