

A 0.8- μm CMOS 2.5 Gb/s Oversampling Receiver and Transmitter for Serial Links

Chih-Kong Ken Yang and Mark A. Horowitz

Abstract—A receiver targeting OC-48 (2.488 Gb/s) serial data link has been designed and integrated in a 0.8- μm CMOS process. An experimental receiving front-end circuit demonstrates the viability of using multiple phased clocks to overcome the intrinsic gate-speed limitations in the demultiplexing (receiving) and multiplexing (transmitting) of serial data. To perform clock recovery, data is $3\times$ oversampled so that transitions can be detected to determine bit boundaries. The design of a transmitter for the high-speed serial data is also described. The complete transceiver occupies a die area of $\sim 3 \times 3 \text{ mm}^2$.

I. INTRODUCTION

THE NEED for low-cost high bit-rate communications has generated interest in using modest CMOS technology to achieve Gb/s bit-rates [1]. This paper describes one method of achieving SONET OC-48 (2.488 Gb/s) rates with 0.8- μm CMOS technology.

Because the data-rate is near the unity gain frequency for a fanout of one inverter chain (1.4 GHz for this process), data samplers and output drivers running at the data-rate is not feasible. Instead, a parallel architecture is proposed to demultiplex the data stream at the input and multiplex the data at the output, hence lowering the on-chip clock frequency. N precisely spaced clocks that run at $1/N$ th the data-rate are used to sequence the input sampling and output switching.

For clock recovery, conventional analog phase-locked loop (PLL) techniques would set the loop bandwidth at less than one-tenth the internal clock frequency. The loop must then be designed with very low noise sensitivity to keep the phase noise to one-half the bit time. This work explores the alternative of oversampling each data bit and using digital logic to detect data transitions and select the proper bit-values. Although the sampling positions may not be precisely centered in the data-eye, the digital control loop can operate at a higher bandwidth of roughly the data transition rate. For example, with 8B10B encoding for dc-balancing, the maximum run-length of five guarantees a transition rate of one-fifth the data-rate enabling phase tracking of 500 MHz.

Sections II, III, and IV describe the receiver architecture, the implementation of the front-end sampling and clock generation, and test chip results, respectively. Last, Section V describes the architecture and implementation of the output driver.

Manuscript received June 7, 1996; revised July 29, 1996.
The authors are with the Center for Integrated Systems, Stanford University, Stanford, CA 94305 USA.
Publisher Item Identifier S 0018-9200(96)08078-X.

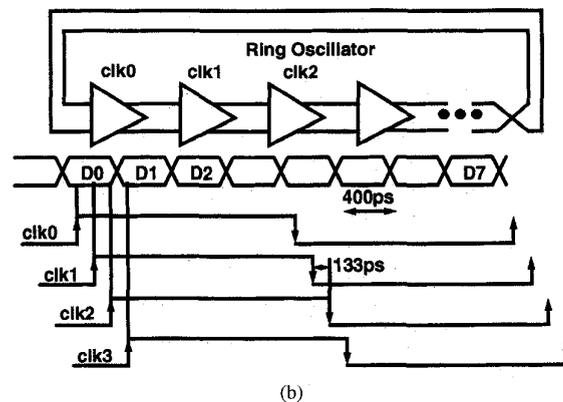
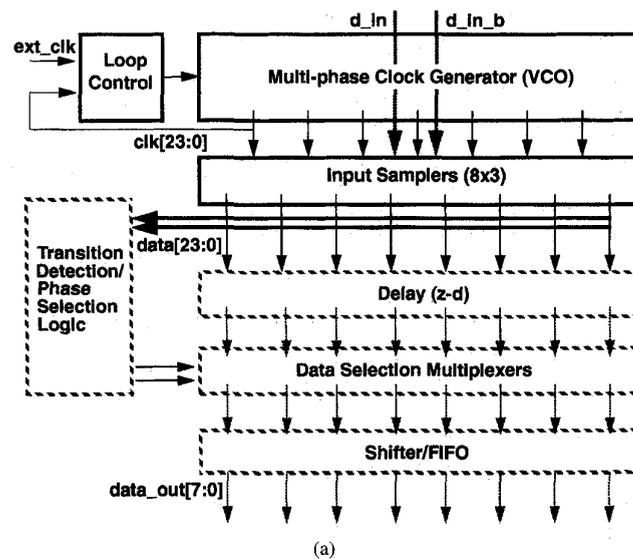


Fig. 1. (a) Architectural floor plan and (b) demultiplexing and oversampling with clock edges.

II. RECEIVER ARCHITECTURE

The architecture of the oversampling receiver is shown in Fig. 1(a). An on-chip PLL locks to an external source that runs at one-eighth the data-rate for the 1:8 demultiplexing. The delay stages of the oscillator are tapped to yield edges, spaced by the sampling interval, that span the period of oscillation. A bank of samplers uses these edges to both demultiplex and oversample the data, Fig. 1(b). Transitions in the data are detected and processed to generate control signals that decimate the oversampled data to actual data bits. The oversampled data are delayed by the number of cycles

needed to generate the control signal for the data selection muxes so that corrections from the control are applied to the corresponding data bits.

The bit error rate and the maximum data-rate for this architecture are limited by the choice of oversampling factor, the input sampler design, and the clock jitter. In the worst scenario, where two samples straddle the center of the data-eye, the deviation from the optimal sampling position is half the bit time divided by the oversampling factor. While this deviation decreases by $O(1/n)$, the number of clock edges, the number of samplers, area, power, and decision logic complexity increase by $O(n)$.

As a compromise we chose $3\times$ oversampling. With the 1:8 demultiplexing, 24 clock edges and samplers are required. In comparison to an ideal acquisition system under white Gaussian noise, to achieve the same bit error rate (BER), the signal power is penalized by 3 dB.

III. FRONT-END IMPLEMENTATION

The critical circuits for this architecture are the input sampler and the clock generation. The primary limitations for maximum data-rate are the input sampler's bandwidth and the setup-hold window. Because the sampling cycle is at one-eighth the data-rate, the sampler's bandwidth requirement is relaxed. However, the sampler's setup-hold window must be less than a sample spacing for the oversampled outputs to have significance.

The signal degradation due to the sampling switch, the input offset of the latch, the intrinsic metastability window of the latch, and the clock jitter all contribute to the setup-hold window. Both the sampling clock slew rate and the sampling switch's RC time-constant act as filters that degrade the data being sampled. If the sampled value is less than the input offset of the following latch, the result of the latch may or may not be correct depending on the previous data. If the sampled value is at the input offset, the latch result has additional uncertainty of the intrinsic metastability of the latch. Lastly, the sampling clock position has the uncertainty of clock jitter. All except the last factor must contribute to a setup-hold window smaller than the sample spacing. The contribution of the last factor depends on the behavior of the jitter. Phase movements slower than the tracking of the digital algorithm will not affect the data reception.

A. Sampler

The circuit diagram for the input sampler [2] is shown in Fig. 2(a). The sampling switch is toggled by one of the sampling edges from the VCO. Data is sampled for a clock half-cycle during which the latch is precharged and reset. The sampling ends on the falling edge of the clock, holding the most recent data that arrived for regenerative amplification by the latch for half-cycle. In order to keep the input data edge rate under one-half the bit time ($2.2 \cdot R \cdot C < 200$ ps), the total input capacitance must be kept under 1.8 pF allotting 75 fF per sampler including wire capacitance.

The zero-jitter setup-hold window is determined by the time taken by the input to slew past the latch offset. With worst case

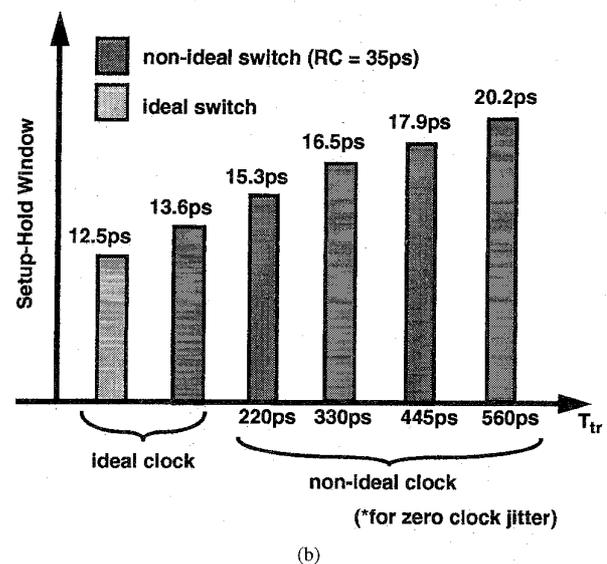
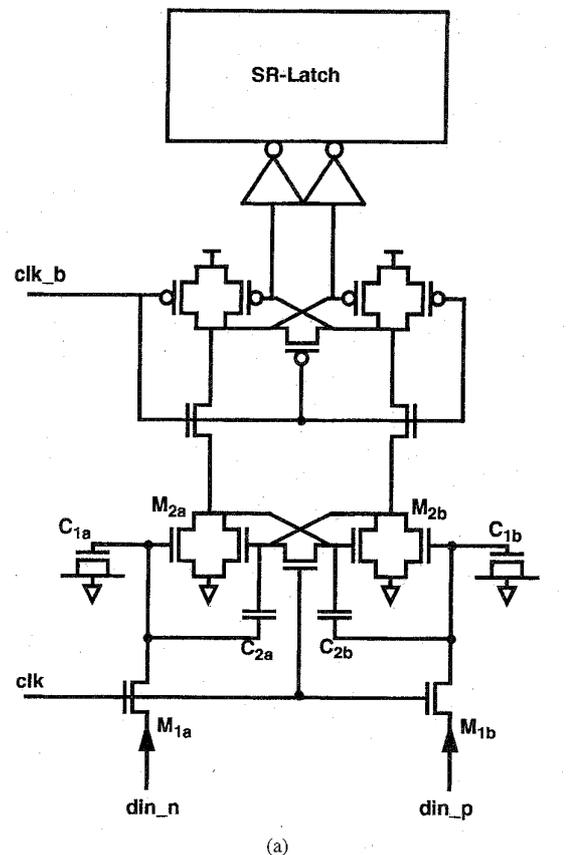


Fig. 2. (a) Input sampler circuit schematic. (b) Setup and hold window dependence on clock switching slope.

mismatches in V_t and K_p , the latch could have an input offset of 50 mV. The offset is dominated by the V_t mismatches of the input transistors ($M_{2a,b}$).

The Miller capacitances are roughly cancelled by the cross-coupled capacitors ($C_{2a,b}$) to reduce disturbance on the sampled value. Normally, the overlap capacitance of $M_{2a,b}$ would

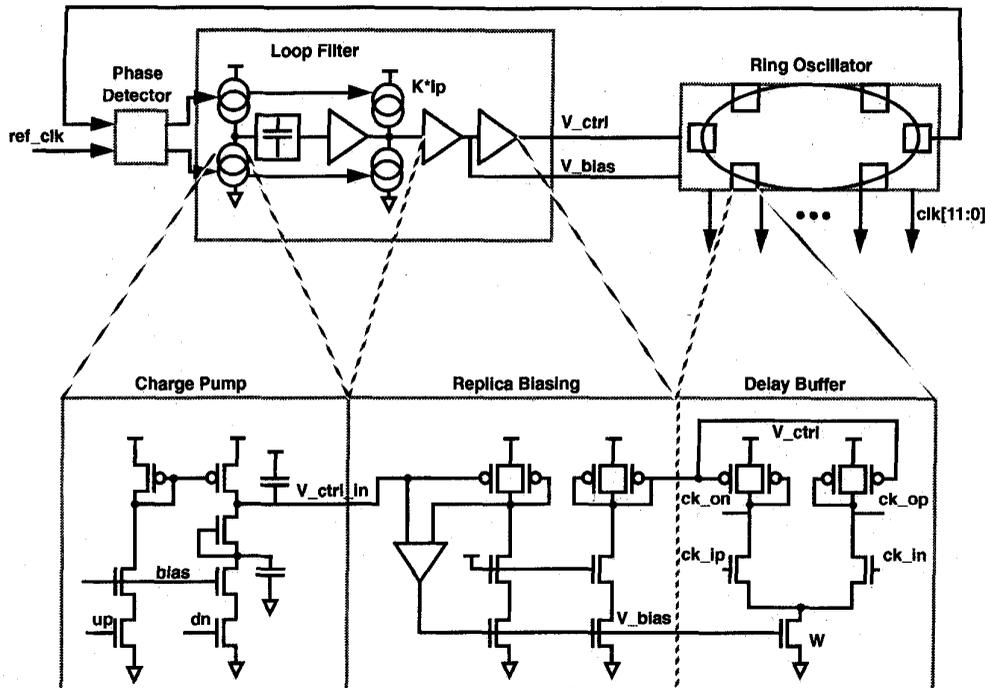


Fig. 3. Clock generation VCO and biasing circuit schematics.

induce charge that slows the latch regeneration, increasing the metastability region.

The input slew rate is reduced by the ON-resistance and capacitance of the sampling switch and the edge rate of the clock transition. The sampling pass-gate ($M_{1a,b}$) is designed to have an ON-resistance of $600\ \Omega$ with $<10\%$ data dependence and sampling capacitance ($C_{1a,b}$) of $60\ \text{fF}$. Part of the sampling capacitance is explicit, reducing its sensitivity to coupled noise. This time constant is significantly smaller than the input transition, to reduce degradation of input stream. The sample to hold transition is toggled by a fast edge rate of $210\ \text{ps}$ driven from a fanout of one inverter. Assuming an input data transition rate of $1/2$ the bit time, this design offers a zero-jitter setup and hold window of $16\ \text{ps}$ ($0.3\ \text{ps/mV}$ of input latch offset). Because the clock edge rate is a considerable fraction of the bit time, slower edge rates would integrate more slowly on the sampling capacitor increasing the setup and hold window. The effect of clock slew rate is shown in Fig. 2(b) where the clock slew rate is changed by incrementing its load by fanout of one.

B. Clock Generation

With $1:8$ demultiplexing and $3\times$ oversampling, 24 evenly spaced clock edges are required. The desired oscillation frequency of $311\ \text{MHz}$ limits the VCO length to six stages. Tapping from each stage generates 12 edges (six true and six complemented). In order to double the number of edges, 50% interpolators are used between every edge.

The PLL and clock generation design are shown in Fig. 3. The sequential phase/frequency detector [3] is designed to have no dead-band to reduce jitter. Because the technology used offers no reliable resistors for a stabilizing zero in the loop

filter, proportional control is summed onto the loop control voltage using a scaled charge pump [4]. The control voltage for a normal zero-stabilized loop filter has the form (1) where I_p is the charge pump current

$$V_c = I_p(s) \times Z(s) = I_p(s) \times \frac{sRC + 1}{sC}. \quad (1)$$

By injecting proportional current, we can get the same equation where K is the proportionality constant and R_o is the output impedance of the replica bias elements

$$V_c = \frac{I_p(s)}{sC} + KR_o I_p(s) = I_p(s) \times \frac{sKR_oC + 1}{sC}. \quad (2)$$

Control voltage ripple at the oscillation frequency can cause duty-cycle error. Since all phases spanning the period are used, the error warps the sample spacing. Offsets in the up and down currents of the charge pump in either magnitude or delay can cause ripple under locked condition. To reduce this error, the phase detector is symmetric with the output loading matched, and a pole is introduced in the down path of the charge pump to roughly match the mirroring pole in the up path (Fig. 3). Furthermore, a third-order pole (from parasitics) is used to suppress the residual ripple.

The loop is designed to have a loop bandwidth at $15\ \text{MHz}$ with phase margin of 70° . A high loop bandwidth is chosen so the loop can track the lower jitter crystal reference frequency. Because the digital logic phase tracking can respond in one cycle, phase drifts less than a sample spacing between decisions will be tracked ($<16.5\ \text{ps}$ per $400\ \text{ps}$ bit) relaxing the design constraints of the VCO elements. In comparison, a conventional PLL scheme with loop bandwidth of one-tenth the VCO frequency would require phase drifts of $<2.5\ \text{ps}$ per $400\ \text{ps}$ bit.

To keep the supply-induced phase drift within the specification, the loop delay elements are designed conservatively for a low supply sensitivity. The delay elements are differential with symmetric loads and replica biasing [3] for supply rejection. Interpolation [Fig. 4(a)] is achieved by two delay elements running at half the current with their outputs tied. For the period of time between the two interpolating clock edges, only one of the two elements is pulling current, thus halving the edge rate and delaying the interpolated edge by approximately half the time between the edges. Also, each of the interpolating clocks from the ring oscillator is buffered to match the delay of the interpolation buffer so that all 24 output clocks have the appropriate phase spacing.

This type of equal current interpolation is ideal when the elements are perfect integrators with their outputs as linear ramps. With exponential outputs (RC-load), an offset develops depending on the spacing between the interpolating edge because halving the current is not the same as halving the time constant. Fig. 4(b) shows the offset as a percentage of edge spacing at various edge spacings. Simulations after tape-out showed that equal current interpolation has a systematic offset of 11% of desired sample spacing [Fig. 4(c)]. The sinusoidal amplitude superimposed on the zigzagging is the duty-cycle error due to charge pump mismatch. As shown in the solid line of the figure, by systematically setting the leading interpolating buffer's current $\sim 11\%$ larger than the lagging buffer, the error is reduced to $< 3\%$.

Beside interpolation and duty-cycle error, mismatches in the clock and data sampling paths can also cause nonideal oversampling position, effectively increasing the sample spacing. The layout is carefully done in a bit-slice manner to match each clock path. Furthermore, all clock buffers and samplers are folded and mirrored to reduce process dependent offsets.

Fig. 5(a) shows two bit slices. The six-stage ring oscillator is arranged in a flattened ring to keep interconnect loading the same for each stage. Twelve interpolators are similarly positioned to keep load capacitance the same and minimal. The 24 samplers are arranged in a zigzag so the samplers using the opposite senses of the clocks are located next to each other. The data inputs of the samplers are routed in a balanced tree over the clock generation block in third metal. The data lines are either shielded by V_{DD} plane in second metal or routed over both true and complement signals to reduce noise coupled onto the data lines. The chip micrograph of the array is shown in Fig. 5(b).

IV. FRONT-END RESULTS

A test chip has been implemented in HPCMOS26B MOSIS process to verify the input channel characteristics, the sampling capability, and the quality of the edge positions. The 24 samples are multiplexed onto 12 differential open-drain outputs by using both edges of the 311 MHz clock.

The differential inputs to the test chip are driven from a 3-Gb/s HP BER tester that has programmable patterns. The clocks from the BER tester are divided by eight and used as the trigger input to a 500-MHz pulse generator. The clocks for the test chip are from the pulse generator which can

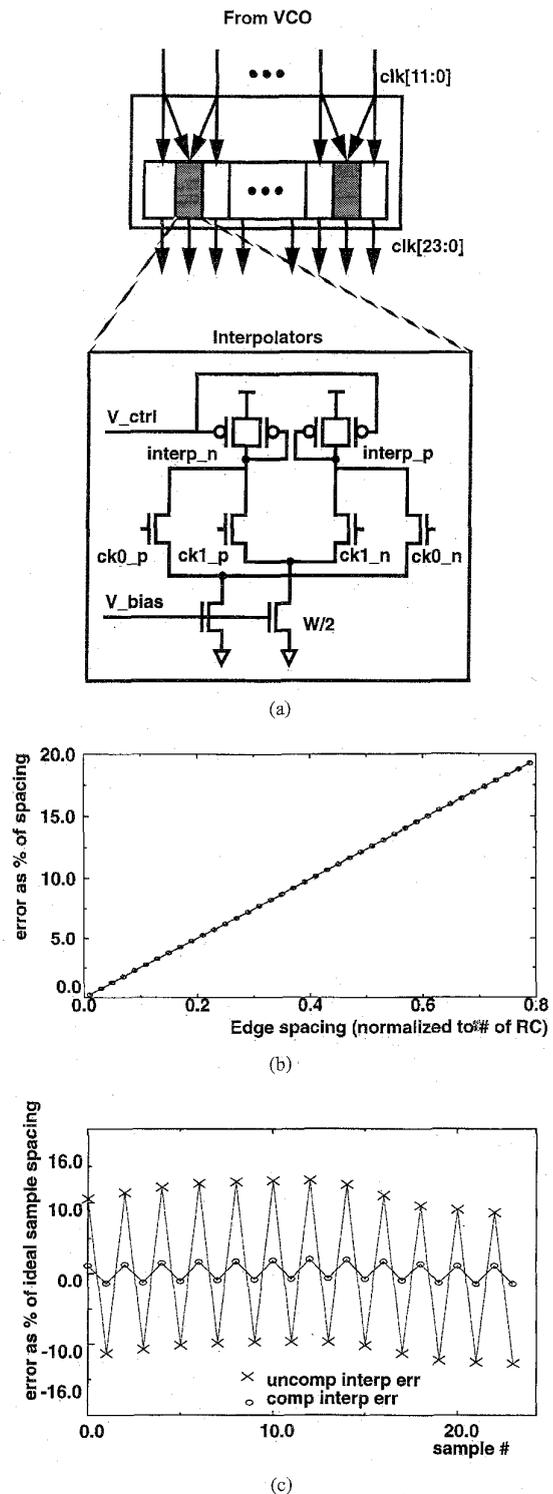
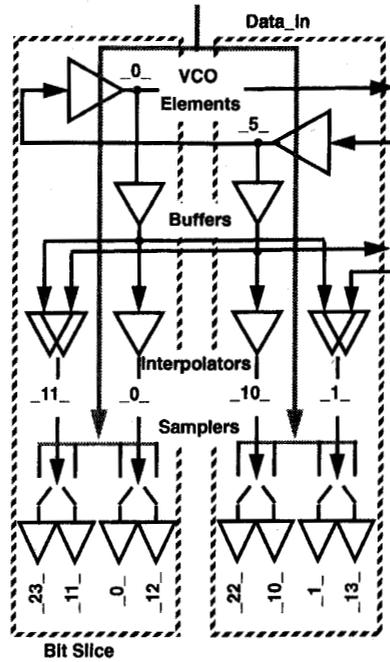
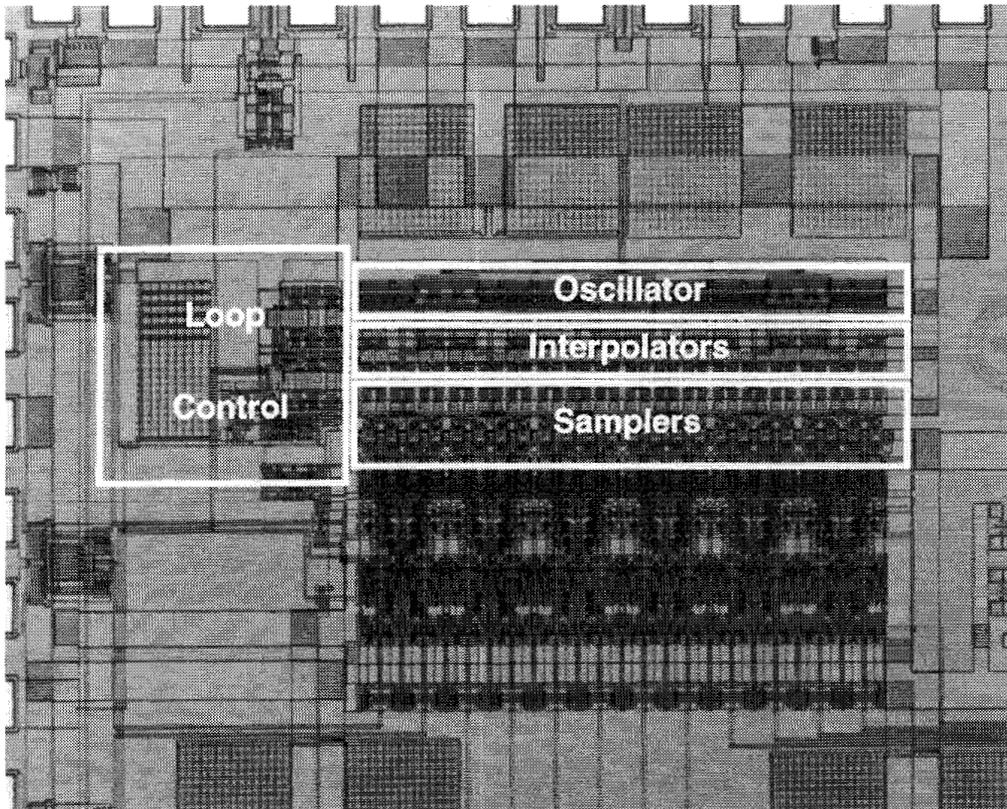


Fig. 4. (a) Interpolation schematics. (b) Equal-current interpolation error at various edge spacing. (c) Simulated systematic error and correction.

operate either synchronous to the pattern source or at its own frequency. The 12 outputs are sampled by a 2 Gsamples/s logic analyzer that has a 64 Kb buffer per channel. The sampled data are downloaded onto a workstation to test various decision algorithms.



(a)



(b)

Fig. 5. (a) Layout bit slice. (b) Sampler/clock generation die micrograph.

Table I summarizes the performance of the test chip. The 40 ps setup-hold time is due to internal clock jitter and sampling window. The window was measured by applying

a transition at the input and using the tester to move the input by 5 ps steps. The samplers' outputs would be fixed before and after the transition. Within the setup-hold win-

TABLE I
PERFORMANCE SUMMARY

Supply range @ 2.5Gbps	4.4 - 5.5V
PLL frequency range	80 - 385MHz
Data rate range	320Mbps - 2.7Gbps
Min. input amplitude	< 100mV
Input time constant	~150ps
Input setup and hold	40ps +/- 5ps
Supply sensitivity of VCO (closed loop)	0.4ps/mV
Power - (Samplers + PLL.)	1W
Die Size	3mm x 3mm
(Samplers + PLL)	(0.5mm x 1.5mm)
Technology	0.8- μ m MOSIS

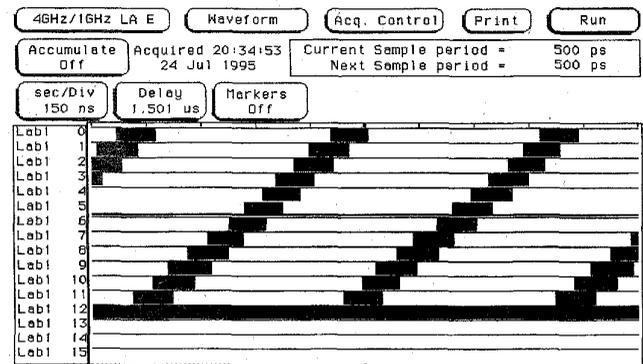
dow, the sampler output would vary with the jitter on the clock.

The minimum input amplitude receivable by the input sampler was measured at lower frequency to eliminate input channel degradation; the pulse generator limits the measurement. The input attenuation is observed using a single-ended pulse of HIGH in a sea of LOW with respect of an adjustable reference. With varying pulse widths, the reference measures the on-chip amplitude to estimate the RC constant.

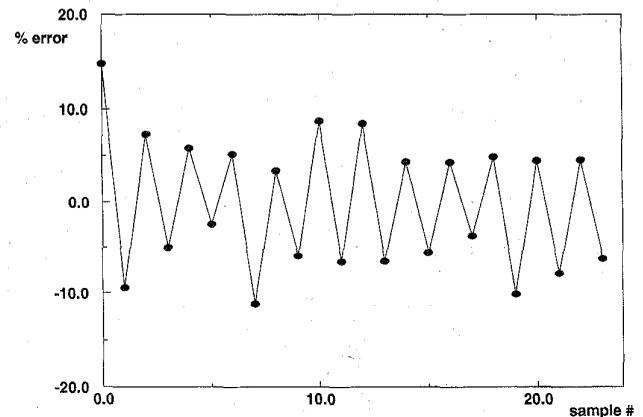
The quiescent jitter numbers are small because the phase-shifted clocks stagger the current pulses onto the supply spreading the noise across the period. To test performance with noise, externally controllable shorting transistors on-chip induce supply bumps. A resistor outside the chip isolates the external supply and the internal supply to allow bump. The magnitude of a supply bump is detected by sensing the drain current of a transistor whose gate-to-source is the supply voltage. A low frequency square wave is applied to the shorting transistor to induce a sharp supply step to model the worst case scenario where the VCO accumulates phase error continually until the loop compensates. The noise sensitivity is measured by observing the shift of the staggered clocks with respect to the data at the oversampled output upon a supply bump. The 0.4 ps/mV sensitivity is a measure of the peak jitter per millivolt of supply noise.

Fig. 6(a) shows output data received by the logic analyzer for a repeating 10000000 input pattern at a data frequency that is not synchronous with the clock frequency. The pulse of ONE walks linearly across all sampling edges at the difference frequency demonstrating the position of each sampling edge and the duration that each sampling edge samples the ONE. At least two samplers sample the data bit during the entire walk so that the oversampling spacing is sufficient for the bit to be determined using a simple algorithm under this relatively quiet environment. Below the logic analyzer output shows the sampling position error as a percentage of sample spacing.

Pseudorandom input patterns are driven to the test chip under various operating conditions such as on-chip supply bumps and input SNR degradation. The supply bumps simulate noise from a digital back-end that would induce phase noise to the sampling clocks. The reduced SNR models the attenuation within a link such as an optical network. The data collected



(a)



(b)

Fig. 6. (a) Logic analyzer output with walking pulse of 00000001. (b) Sampling position error.

from the logic analyzer are used to test various decision algorithms. From a C-language version of the algorithm, data is successfully received even with induced jitter larger than 200 ps (one-half bit time) of peak-to-peak jitter. For the data sets available, no errors were found for all cases with SNR > 17 dB.

V. OUTPUT DRIVER

After demonstrating a working design of the receiver, the problem of transmitting the bit rate needs to be addressed. To properly test the bit error rate, we designed a 4 Gb/s transceiver using 0.6- μ m CMOS technology operating at 3.3 V.

A. Architecture

The process technology limitation on generating pulses at the data-rate is addressed in the receiver by using parallel samplers. The same limitation occurs in the transmitter where a single output device and its predriver would not be switched quickly enough to drive the data-rate without significant intersymbol interference.

A similar but reversed architecture as the receiver can be used with multiple drivers that multiplex onto the same output pad. Each driver contains two switches sequenced by clock phases tapped from a ring oscillator. The bottom switch is

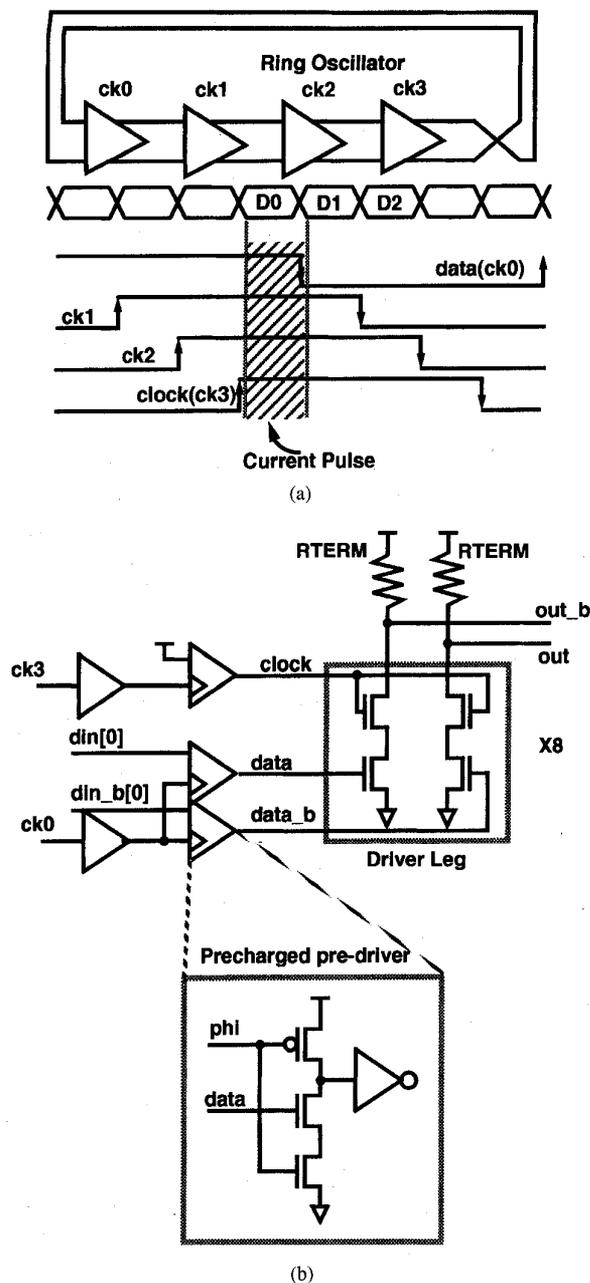


Fig. 7. (a) Output driver timing. (b) Circuit schematic.

driven by data qualified by a clock phase, $CK0$. The top switch is driven by a buffered version of clock phase, $CK3$, one ring oscillator tap earlier than the complement of $CK0$ [Fig. 7(a)]. A current pulse of the bit width would flow during the period the two clock phases overlap. The driver is duplicated with the switches driven by the appropriate clock phase allowing the current pulses to occur in sequence.

The architecture can drive higher bit-rates because the overlap of two pulses can be made at least one-third smaller than a propagable pulse in a given technology. In order to maximize the output data-rate, the predriver outputs are driven as sharply as possible. The slew rates of the output are

determined by the top switch's rising edge and the bottom switch's falling edge.

Because the design is pushing the maximum data rate, this implementation did not include any slew rate control. However, typically, the slew rate is designed to be approximately one-third the bit time keeping the switching noise (Ldi/dt) low and reducing the high frequency components on the line. In many output drivers with a single predriver output, slowing down the slew rates to acceptable levels causes intersymbol interference at the output. An additional advantage of this architecture for use in slower data-rate communication is that the slew rate control is applied to predriver outputs running at a divided frequency of the data rate. The edges forming each of the data bits are no longer dependent upon the previous bit removing intersymbol interference. Furthermore, using slower clocks than the output rate allows power savings.

There are several concerns regarding this architecture. First, similar to the receiver (Section II), using a divided clock to drive a higher frequency signal imposes a more severe phase noise constraint on the PLL. The phase noise of the VCO and any buffering directly reduce the data-eye. Another concern is intersymbol interference. The large drain capacitance of the parallel legs could potentially increase the RC-time constant, degrading the channel characteristics. The output time constant with the termination resistor should be kept to well under one-third bit time. Data dependent parasitic capacitances at the output could also cause delay variation in the switching. Lastly, the switching noise due to the pseudodifferential nature could induce supply noise back into the clock generation.

B. Implementation

A design of the output driver has been implemented in the MOSIS HPCMOS14TB process (0.6 μm channel length) aiming for 4 Gb/s transmission rate. To match the demultiplexing of the receiver, an output performs 8:1 multiplexing. A four-stage ring oscillator with the same low noise stages is used to generate the eight phases. The switches are two series-stacked NMOS transistors [Fig. 7(b)] that operate as open drain current mode switches. To reduce susceptibility to common mode noise, the output is pseudodifferential.

Fully differential drivers are not used because the transistor sizes required to fully switch the current would be much larger because the $V_{gs} - V_t$ is lower in comparison, hence increasing the output capacitance. Furthermore, large tail (common) node capacitance in a differential pair causes different rising and falling edges. Compensation for this would further increase output capacitance. Using a grounded source driver, careful layout, and sharing the source/drain diffusion, the total on-chip output capacitance is kept under 1 pF to keep the time constant below 50 ps.

The data input is driven by a domino buffer so that the data is precharged low. The clock input is driven by the same buffer to maintain the phase alignment coming from the oscillator. In operation, either DATA or DATA_B input will rise as the precharge is released, discharging the internal node

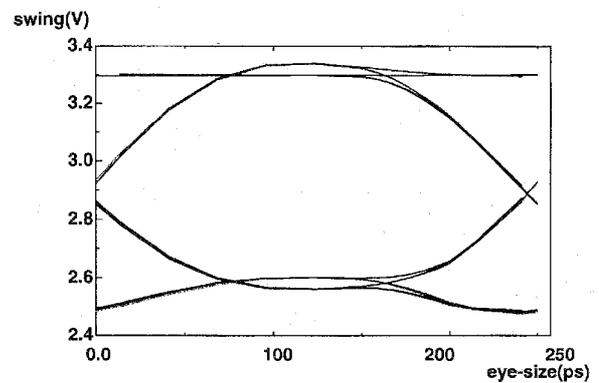
of the series stack. When clock input rises, current for the bit flows through the series stack until the data input falls with the precharge. This process repeats eight times with each of the eight driver legs.

The same clock generation architecture as the receiver is used to generate a quiet clock. If the output is driven to the phase tracking receiver (Section II), only the rate of jitter accumulation is of concern. However, to maintain good tracking of the clean reference frequency and a low peak-to-peak jitter, the loop bandwidth is chosen at $\sim 1/20$ th the input clock frequency. The clock and data buffering that drives the output is matched in the feedback of the PLL to partially track the phase noise these elements introduce due to their higher supply sensitivity. However, a small enough delay is maintained so that the phase margin of the PLL will not be severely degraded. The buffering elements include the low swing to high swing conversion, two inverters to increase the drive strength of the clock and the precharged elements of the predriver. The additional inverter buffering is necessary because the slew rate at the predriver output is required to be as high as possible. The design used fanout of 1.5 inverter at the last stage (10%–90% of 200 ps in 0.6 μm process). The total delay of the path from the oscillator outputs to the driver legs is kept to ~ 700 ps. The worst-case additional jitter (untracked by the PLL) can be approximated using the supply sensitivity of inverters ($\sim 0.02\%/mV$).

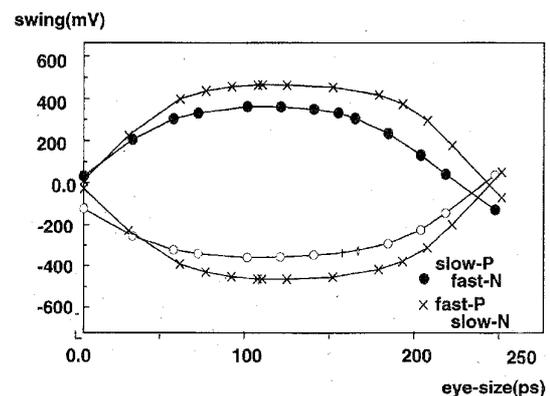
Being driven to both true and complement legs, the clock signal is not sensitive to data dependent loading. By using it to drive the upper output transistor, the internal node is always charged to $V_{DD} - V_t$ ($V_{sb} = 2.5$ V) before the data's falling transition. This shields the data's falling transition from any dependence on the output value removing this as a source of intersymbol interference. Furthermore, because the data input precharges the internal node, the driver output high-low transition is sped up due to charge sharing, relaxing the PMOS size requirements in the predriver. Fig. 8(a) demonstrates the output driving pseudorandom sequence demonstrating low intersymbol interference. In the figure, both V_{OL} and V_{OH} show the overshoot from the 2 nH inductor at the output. The low value dips slightly between bits due to overlapping of consecutive current pulses.

The output swing depends on the current pulled by one of the series stack. Because the swing is designed to be 600–800 mV, the stack behaves as a current source even with the gate pulled to V_{DD} . The output current can be controlled through programming the number of current source legs activated by the predriver [6]. To simplify the current design for testing, the programmability is not implemented. Instead, the predriver's supply is adjustable to control the output swing.

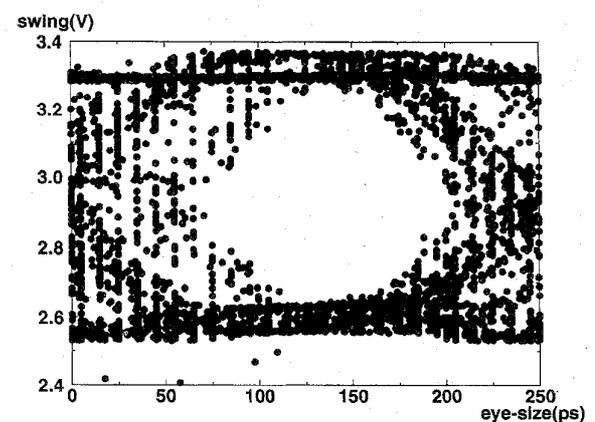
Since the pseudodifferential output does not guarantee constant current, switching noise still exists. To keep the switching noise and output ripple to a minimum, the falling edge of each current pulse should overlap the next current pulse at half the current. Because the data predriver drives the bottom transistor, the Miller capacitance is much greater creating additional delay in the data buffering by ~ 50 ps. Further compensation is necessary because the currents are not half



(a)



(b)



(c)

Fig. 8. (a) Noise free data-eye showing very little intersymbol interference. (b) Eye crossing shifts across process, indicating crossover current pulses. (c) Simulated data-eye with 300 mV noise induced on supply.

when data's (DATA0) falling edge and clock's (CK1) rising cross at half V_{DD} . DC analysis shows an optimal cross point of ~ 2 V. The compensation is achieved by decreasing the fanout of two of the inverters in the buffer chain. By using two inverters, the resulting change in delay is dependent mainly on the fanout difference and less sensitive to $P:N$ process skews [Fig. 8(b)].

Fig. 8(c) shows the simulated data-eye with the output driving a 2-nH inductor and 50- Ω load. With a 400 mV supply bump to the internal V_{DD} , the eye opening is still ~ 125 ps.

VI. CONCLUSION

This work has demonstrated the possibility of receiving and transmitting data at rates near the unity gain bandwidth of the underlying technology, extending the range of usability for CMOS in high bandwidth communication. A test chip of the entire transceiver with the digital phase selection logic is in fabrication.

ACKNOWLEDGMENT

The authors wish to thank Prof. L. Kazovsky and his research group for the use of test equipment and technical support. They are also indebted to Vitesse Semiconductors for their packages, Rambus Inc. for the use of test equipment and technical suggestions, Prof. T. Lee, S. Sidiropoulos, and B. Amrutur for numerous discussions.

REFERENCES

- [1] S. Kim *et al.*, "An 800 Mbps multi-channel CMOS serial link with $3\times$ oversampling," in *IEEE 1995 Custom Integrated Circuits Conf. Proc.*, Feb. 1995, p. 451.
- [2] A. Yukawa, "A CMOS 8-bit high-speed A/D converter IC," *IEEE J. Solid-State Circuits*, vol. 20, no. 3, pp. 775-779, June 1985.
- [3] Motorola MECL Device Data, 3rd ed. Motorola, Inc., pp. 6-81-6-84, 1979.
- [4] I. Novof *et al.*, "Fully integrated CMOS phase-locked loop with 15-240 MHz locking range and ± 50 ps jitter," in *ISSCC Dig. Tech. Papers*, Feb. 1995, pp. 112-113.
- [5] J. Maneatis and M. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1273-1282, Dec. 1993.
- [6] N. Kushiyama *et al.*, "A 500-Megabytes/s data-rate 4.5M DRAM," *IEEE J. Solid-State Circuits*, vol. 28, pp. 490-498, Apr. 1993.



Chih-Kong Ken Yang received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 1992. He is currently pursuing the Ph.D. degree at Stanford University in the area of circuit design for high-speed interfaces.

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

Mark A. Horowitz, for a photograph and biography, see p. 1284 of the September 1996 issue of this JOURNAL.