A Tracking Clock Recovery Receiver for 4Gb/s Signaling

Extended Abstract

John Poulton
Microelectronic Systems Laboratory
University of North Carolina -
Chapel Hill
jp@cs.unc.edu

William J. Dally
Computer Systems Laboratory
Stanford University
billd@csl.stanford.edu

Steve Tell
Microelectronic Systems Laboratory
University of North Carolina -
Chapel Hill
tell@cs.unc.edu

Abstract

We have previously described a design for a 4Gb/s signaling system that uses transmitter equalization to overcome the frequency-dependent attenuation due to skin effect in transmission lines. We present here experimental results from an implementation of this idea in 0.5μ CMOS, showing the effectiveness of a simple transition-filter equalization technique. Our experimental chips use a tracking clock recovery receiver, in which a 21-phase clock is servoed to center every other clock on the center of the data "eye". This method is contrasted with the oversampling method of clock recovery. Oversampling clock recovery has the advantage that it can reject jitter up to the lesser of the minimum transition frequency or the data clock frequency, but it introduces quantization jitter of $\pm 1/2k$ of the bit cell, where k is the number of samples per cell. Tracking recovery gives better performance when there is little jitter above the cutoff frequency of the tracking control loop, avoids quantization jitter entirely, and allows transmitter encoding with much longer run-lengths. Electrical measurements in very high-speed signaling systems are quite difficult to perform with conventional instrumentation, particularly for on-chip signals; this paper describes our design for simple CMOS analog samplers that enable observation of on-chip signals.

1. Introduction

The performance of many digital systems is limited by the interconnection bandwidth between chips, boards, and cabinets. In current high-performance digital systems, full-swing unterminated signaling methods, unsuitable for data rates over 100MHz, are giving way to current-mode differential signaling, and a number of such systems have been demonstrated at signaling rates up to about 2Gb/s [1,2,3,4]. An IEEE standard (1596) for current-mode differential signaling has been adopted.

Even the best signaling systems are limited to speeds of 2Gb/s or so and distances under 1m by the frequency dependent attenuation of copper lines. However, integrated circuit scaling now makes it possible to build very sophisticated circuitry into each I/O pad on a chip that can compensate for the characteristics of the physical interconnect, cancelling the dominant sources of timing and voltage noise.

At last year's conference, we introduced the technique of transmitter equalization to overcome the frequency-dependent attenuation of copper transmission lines [5], and since then at least one other group has described a similar design [4]. We have implemented this idea in an 0.5μ CMOS chip and will present here some initial experimental results from this device. Specifically, we will show that transmitter equalization, built using a simple transition filter, can cancel much of the amplitude and timing noise for a very lossy PC board interconnect system.

Where [5] described the details of the equalized 4Gb/s transmitter, this paper describes the receiver portion of our experimental system in more detail. In Section 2, we review the elements of the problem we are trying to solve, derive a simple set of formulas for calculating frequency dependent attenuation, and compare these with measured losses. Section 3 briefly reviews the transmitter design. Section 4 describes the tracking clock recovery receiver. Section 5 compares tracking clock recovery with oversampling receivers, showing the complementary benefits and problems of the two approaches. Section 6 describes the test chip and board we have constructed and details a simple

and novel means of providing on-chip instrumentation. Section 7 presents some preliminary experimental results. Section 8 summarizes the work performed so far and explores possible avenues of future research.

2. Frequency-Dependent Attenuation

Skin-effect resistance and dielectric absorption cause the attenuation of a conventional transmission line to increase with frequency. With a broadband signal, as typically used in digital systems, the superposition of unattenuated low-frequency signal components with attenuated high-frequency signal components causes intersymbol interference that degrades noise margins and reduces the maximum frequency at which the system can operate. Intersymbol interference both reduces the signal-to-noise ratio of a signal (amplitude noise) and increases jitter (timing noise), as sketched in Figure 1. The problem here is not the magnitude of the attenuation, but rather the interference caused by the frequency-dependent nature of the attenuation. The high-frequency pulse has sufficient amplitude at the receiver for proper detection. It is the offset of the pulse from the receiver threshold by low-frequency interference that causes the problem.   Equalization reduces the unattenuated low-frequency portion of the signal, reducing the absolute size of the signal eye, but removing most of the amplitude and timing noise.

(a) Eye vertical size reduced by 2A-1.   (b) Trailing-edge jitter reduces   (c) Equalization recovers full eye width,
                                                   horizontal size of eye.   reduces amplitude noise.
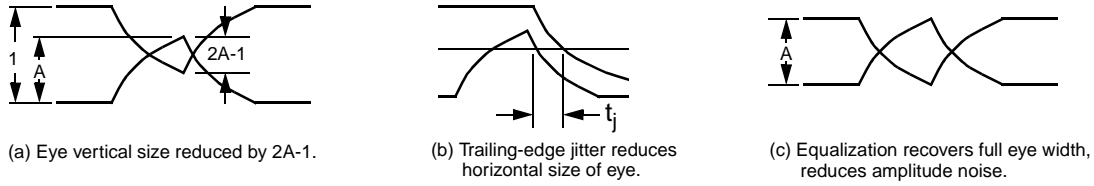
Figure 1: Frequency dependent attenuation causes intersymbol interference that (a) reduces the height of the data eye and (b) introduces trailing-edge jitter. Equalization (c) reduces the height of the eye by the high-frequency attenuation A, but removes the amplitude noise and recovers the full eye width.

Frequency dependent attenuation, and its resultant intersymbol interference, arises from several physical phenomena in transmission lines, mainly loss due to skin effect and to dielectric absorption. At high frequencies (above 100MHz), current is carried primarily on the surface of the conductor, dropping off to a value of e^{-1} at a depth of

$$\delta = (\pi f \mu \sigma)^{-1/2} \tag{1}$$

where σ is the conductivity of the material (5.8×10^7 mhos/m for copper) [6]. A thin strip-guide with width w has a high-frequency resistance per unit length of

$$R(f) = \frac{1}{2w}\left(\frac{\pi \mu f}{\sigma}\right)^{1/2} \tag{2}$$

However, the skin effect does not affect the conductor's resistance unless the frequency is above a frequency f_s, where the skin depth is half the conductor thickness t, given by

$$f_s = \frac{1}{(t/2)^2 \pi \mu \sigma} \tag{3}$$

Above this frequency,

$$R(f) = R_{DC}\left(\frac{f}{f_s}\right)^{1/2} \tag{4}$$

If losses are not too large, then a signal travelling down a transmission line of length L and characteristic impedance Z_0 is attenuated by

$$A(f) = e^{-\alpha_R L} = \exp\left(\frac{-R(f)}{2Z_0}L\right) \tag{5}$$

Dielectric absorption introduces additional attenuation

$$A(f) \; = \; e^{-\alpha_D L} \; = \; \exp\!\left(\frac{-\pi \sqrt{\varepsilon} f \tan\delta}{c} L\right) \tag{6}$$

where c is the speed of light, ε is the dielectric constant and $\tan\delta$ the loss tangent for the dielectric [7].

As an example, we performed measurements on a 1m 50Ω stripguide of width 8mil (w=200μm), 0.7mil thick (t=18μm) in a PC board built with GETEK dielectric. This line has R_{DC} of about 5Ω/m, and f_s is about 54MHz. At high frequencies, return current flow is confined to regions in the planes above and below the trace and roughly the width of the line, so the resistance is (very approximately) doubled. The loss tangent for GETEK is about 0.01at the frequencies between 1 and 10GHz. Both measured (using an Hewlett-Packard 8753D Network Analyzer) and calculated attenuations are plotted in Figure 2. The calculated attenuation from each of the two effects is shown separately in the figure; at 2GHz skin effect attenuation is about 0.56 and dielectric loss is about 0.65.
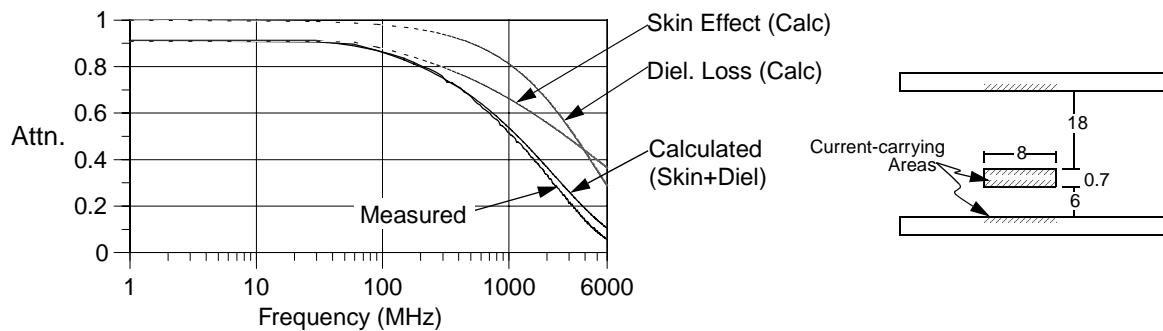

Figure 2: Attenuation in an 8-mil stripguide of dimensions shown at the right.

In this example, our overly simplified model for losses gives quite reasonable agreement with measurement. However, in any critical signaling design, there is no substitute for careful measurement and accurate modeling. In a "real" signaling system package parasitics introduce perturbations in the signal path that may be the dominant source of timing and amplitude jitter.

3. Equalizing Transmitter

Equalization eliminates the problem of frequency-dependent attenuation by filtering the transmitted or received waveform so that the concatenation of the equalizing filter and the transmission line response gives a flat frequency response with linear phase. As described in [5], we equalize the line using a 4GHz 5-tap FIR filter built into the current-mode transmitter. The FIR filter is implemented as a transition filter, as this is much easier to build in CMOS. Equalizing at the transmitter allows us to use a simple receiver that just samples a binary value at 4GHz. Equalizing at the receiver would require an A/D of at least a few bits resolution or a high-speed analog delay line, both difficult circuit design problems. A discrete-time FIR equalizer is preferable to a continuous-time passive or active filter as it is more easily realized in a standard CMOS process and provides linear phase.

A block diagram of the transmitter is shown in Figure 3. A set of retiming and distribution latches accepts 10 bits of data, D_{0-9}, at 400MHz and outputs 10 groups of 5 data bits to each of 10 filters. These bits represent the data bit for one of the transmitter segments, and the 4 preceding data bits. Each filter looks up a strength with which to drive out its bit, based on these five data bits, as described below. The DACs sink current differentially from the loads on the outputs Out+,Out-. DACs are timed by a set of 10 clocks, generated in a DLL. DAC n is turned on by clock phase ϕ_n and turned off by the phase ϕ_{n+1}.
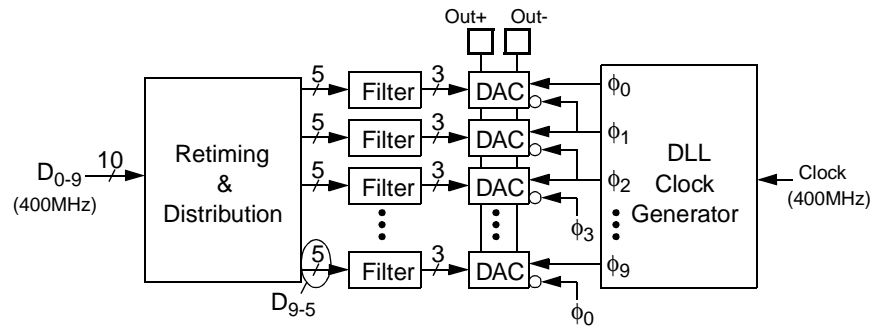
Figure 3: The transmitter is realized using 400MHz circuitry. A 10-phase clock sequences 10 data bits from 10 current-steering DACs, whose output strength is determined by the filter, based on the values of 5 successive bits of data.

Circuit details of the filter and DAC are shown in Figure 4. The filter compares data bit D_n with the four previous bits. A "find first 1" detector determines the most recent data bit that differs from D_n and indexes a small RAM, in which drive strengths for each of the five possibilities are stored. Drive strengths are generally reduced with increased distance (m-n) to the most recent bit for which D_n and D_m differ. The 3-bit strength value fetched from the filter RAM is passed to the DAC in dual-rail form for convenience in implementing the DAC circuitry. The DAC is a binary weighted current-steering circuit. The Off clock ϕ_{n+1} and the strength bits L_k and H_k turn one of the current steering transistors on prior to the rising edge of the On clock ϕ_n, which energizes the current source transistor. The rising edge of the Off clock turns off the current cell; delays for On and Off clocks are closely matched by device sizing.
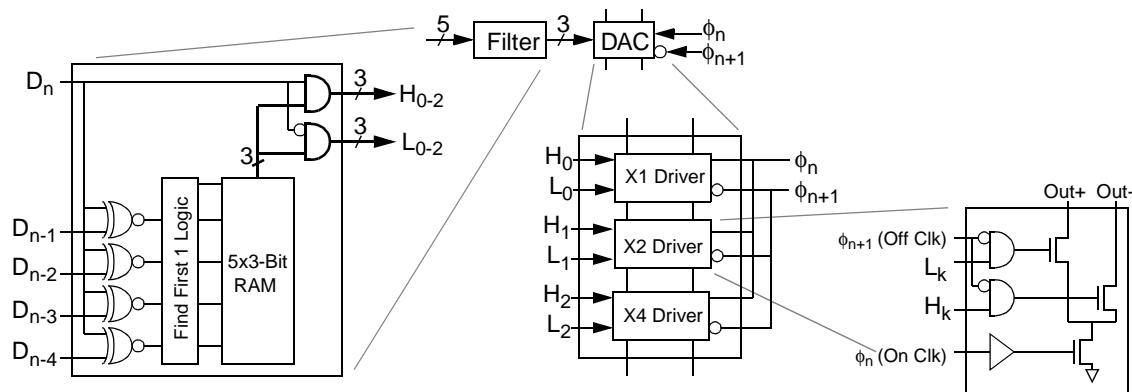

Figure 4: Circuit details of transmitter internals.

4. Tracking Clock Recovery Receiver

A block diagram of the tracking clock recovery receiver is shown in Figure 5. Differential inputs are terminated on-chip in an adjustable terminator. A demultiplexing receiver produces 20 samples of the incoming bit stream, timed by a 20-phase DLL-based clock generator. The 20 clock phases are equally spaced on half-bit-cell intervals, and the overall phase of the clock generator is servoed so that the samples represent 10 successive bit values, and the neighboring 10 edge values. The data is re-timed into the chip's internal clock regime in a set of retiming latches, and both the current 10 bits and the previous 10 bits are output by a set of "framing" latches. A contiguous 10 bits of data is selected from these 20 bits by a funnel shifter that frames the data to compensate for arbitrary delay across the interconnect. The funnel shifter is set during initialization by transmitting a known bit pattern. A matcher examines the incoming data and, with the help of a small finite state machine, adjusts the funnel shifter for proper framing.
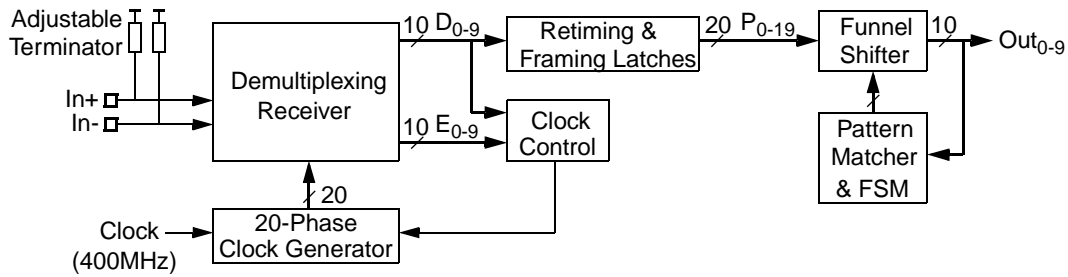
Figure 5: Block diagram of tracking clock recovery receiver.

The details of the demultiplexing receiver are shown in Figure 6. The samplers are gate isolated sense amplifiers, similar to the latches used in [8]. The sense amp outputs precharge high during ϕ low, so the addition of an R-S latch at the output causes these sense amps to behave as edge-triggered flip-flops, a considerable convenience for the remainder of the circuit design.
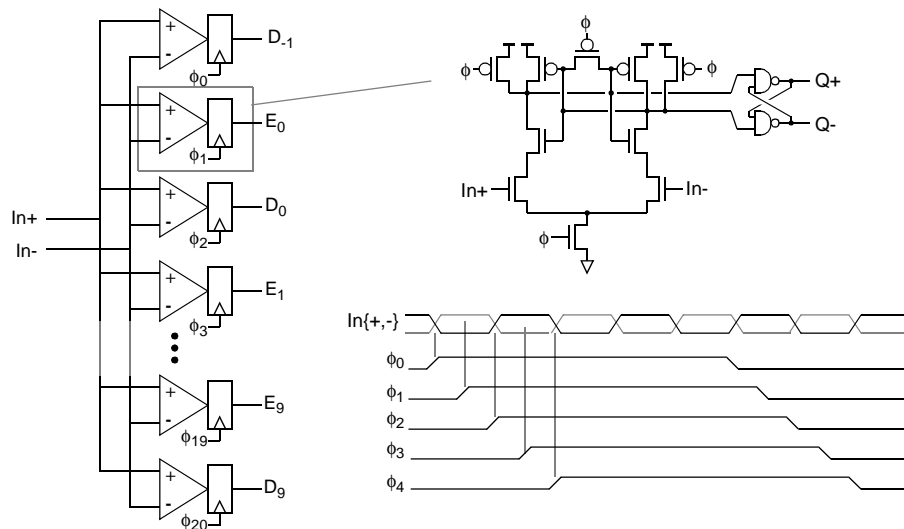

Figure 6: Demultiplexing receiver.

The control circuitry for the 20-phase clock generator is shown in Figure 7. The clock control circuitry generates an (analog) control that sets the overall phase of the 20-phase clock generator. Its job is to servo this phase shift in such a way that the edge samples E_{0-9} are centered on the data transitions, thereby placing the data samples D_{0-9} at the centers of the data eyes. For each pair of edge/data samples E_n,D_n, the data bit is compared with the data bit in the previous cell D_{n-1} to see if there was a transition. The data sample is also compared with the edge sample; if they are the same, and if a transition was detected, it is inferred that the edge sample is late, indicating that the phase shift should be increased. Up_n and Dn_n signals are generated for each of the 10 edge/data sample pairs, and the results are summed in an analog summer to generate a differential analog phase control voltage pair $P\{+,-\}$. This operation is performed using a switched-capacitor summer. At each stage of this summer, if Up or Dn is asserted, the voltage on a small capacitor is pulled to one of the power rails on one phase of the sample clock for that data cell. On the opposite phase, the charges on these capacitors are dumped onto a much larger pair of capacitors on $P\{+,-\}$. A switched-capacitor circuit similar to the summing stages adjusts the common mode voltage on $P\{+,-\}$ to keep it at the proper value for the phase shifter, described below.
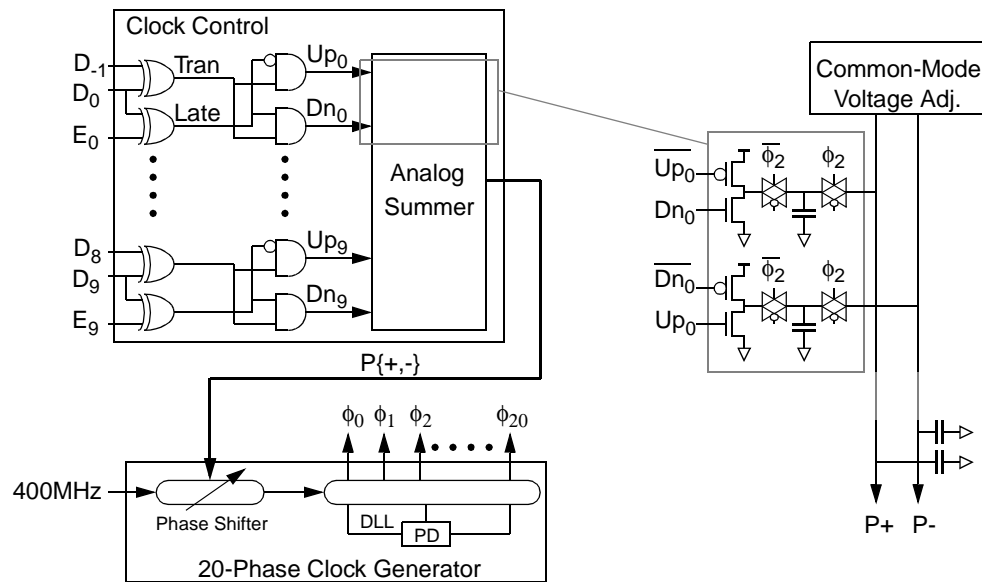
Figure 7: Control circuitry for the clock generator to support tracking recovery.

Figure 8 shows the details of the 20-phase clock generator. At the right of the figure is a delay-locked loop, based on replica-biased differential delay stages. The loop is locked with $180°$ phase shift between ϕ_0 and ϕ_{10}, using a sequential phase detector and charge-pump filter. Circuit techniques are similar to those described in [9,10]. The phase shifter is built using identical delay stages, operating off the same bias and control voltages as the DLL, and thus stabilized by the control loop. The clock input, which arrives on-chip at PECL signal levels, is AC coupled (using linear capacitors built from metal interconnect sandwiches), and DC restored to the mid-swing voltage of the delay line by a delay stage whose inputs and outputs are shorted. This turned out to be an important detail to avoid introducing systematic jitter into the delay elements. A "sacrificial" delay stage ensure that the signal swing into the interpolator proper is amplified to the "native" signal swing within the delay line. The interpolator is a dual input replica-biased stage with two current sources controlled by $P\{+,-\}$, whose common-mode voltage is adjusted, as mentioned above, to keep the total interpolator load current consistent with the normal delay stages. The interpolator adjusts the phase of the clock input to the DLL over a 3-delay range, enough range to compensate for run-time variations in edge timing of the received signal. An initialization circuit, not shown, starts the DLL at minimum delay (to avoid false locking) and initializes the interpolator at mid-range.
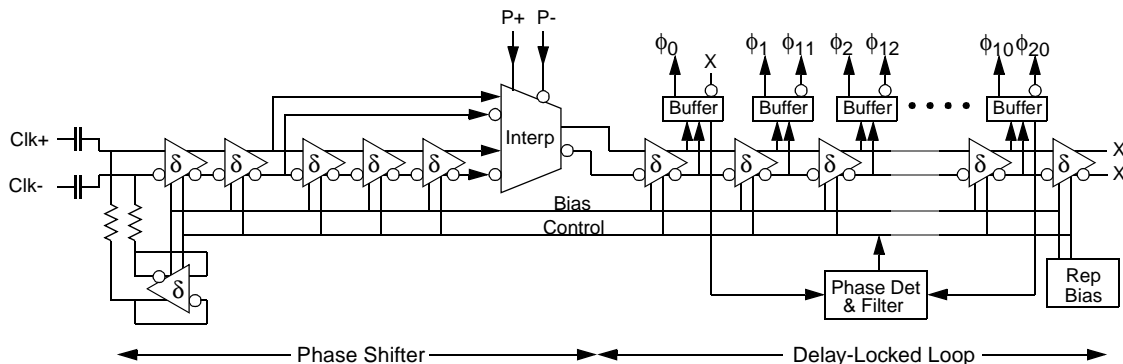

Figure 8: 20-phase clock generator, based on a differential replica-biased delay-locked loop.

5. Tracking vs Oversampling Clock Recover

An oversampling receiver is an alternative to the tracking clock recovery described above. An oversampling receiver operates without adjusting the phase of the receive clock. Instead, the receiver oversamples by taking $k\geq3$ samples[1] each bit period and sorting out the correspondence between bits and samples after sampling.

Oversampling receivers have several advantages compared to the tracking receiver described here. First, they are somewhat simpler to implement. They eliminate the need for the analog summing network (Figure 7) and the interpolating phase shifter (left part of Figure 8). They do, however, require a faster sampling rate, at least three samples per bit, as compared to two samples per bit for the tracking receiver. Also, the logic used to sort out bits from samples, while it operates on the low-frequency side of the demultiplexing receiver, consumes a significant amount of chip area.

Another advantage of an oversampling receiver is the ability to reject very high frequency jitter, up to the lesser of the clock rate or the minimum transition frequency. Oversampling receivers examine all of the samples taken during a window of time, detect the transitions during this window, and then determine the bit pattern that most likely produced the observed transitions. The window is shifted each time a set of samples is taken and the process is repeated. The decision as to whether a given bit is a one or a zero is largely determined by the immediately adjacent transitions giving excellent high-frequency response. Also, since the bits are sorted out after sampling, an oversampling receiver can be non-causal, using samples both before and after a given bit in determining the value of the bit. In contrast, tracking receivers are strictly causal, using only samples before the present bit in determining when to sample, and they have relatively slow clock adjustment loops, passing high-frequency jitter with little attenuation.

Oversampling receivers have two significant disadvantages (in addition to requiring a faster sampling rate). First, to achieve good high-frequency jitter rejection, oversampling receivers require a short run-length reducing the effective signal bandwidth. Second, and more significantly, they introduce quantization jitter with a p-p value equal to $1/k$ of the bit period (*e.g.*, 83ps for k=3 and a 250ps bit cell). This is nearly as much as the 91ps p-p jitter measured on our experimental system (see Figure 12(b) below). This quantization jitter is the uncertainty in the position of each detected transition. A transition is detected when two consecutive samples have different values. The detected transition may have occurred at any time between these two samples. Thus, the uncertainty is equal to the sample period.

In general, oversampling is preferred if the high-frequency jitter that will be rejected by the oversampling receiver is greater than the quantization jitter that oversampling introduces. In our experimental system, described below, these two quantities are comparable. We plan future experiments that include both oversampling and tracking receivers to compare these complementary methods in a controlled environment.

6. Experimental System

We have fabricated the transmitter and receiver, described in the previous section, through the MOSIS service on Hewlett-Packard's CMOS14 process. This process is a 0.5μ epitaxial nwell process designed for 3.3 volt operation. The chip is packaged in a Vitesse 52-pin ceramic leaded chip carrier, a package usually used for GaAs integrated circuits. This section describes the chip, including a simple and novel means for making on-chip measurements, and the test board built to exercise the chip's functions.

The design of the test chip is outlined in Figure 9. Each chip contains a transmitter and a receiver with (almost) completely independent circuitry. The transmitter/receiver pairs operate mesochronously.

---

1. Operating with k=2 samples per bit period is not sufficient. If the sampling clock is aligned with bit transitions, a receiver with k=2 could sample an isolated high bit as 1,2, or 3 consecutive 1s and a pair of high bits as 3, 4, or 5 consecutive 1s. With k=2, it is impossible in many cases to decide if 3 consecutive 1s should be interpreted as a single high bit or a pair of high bits.
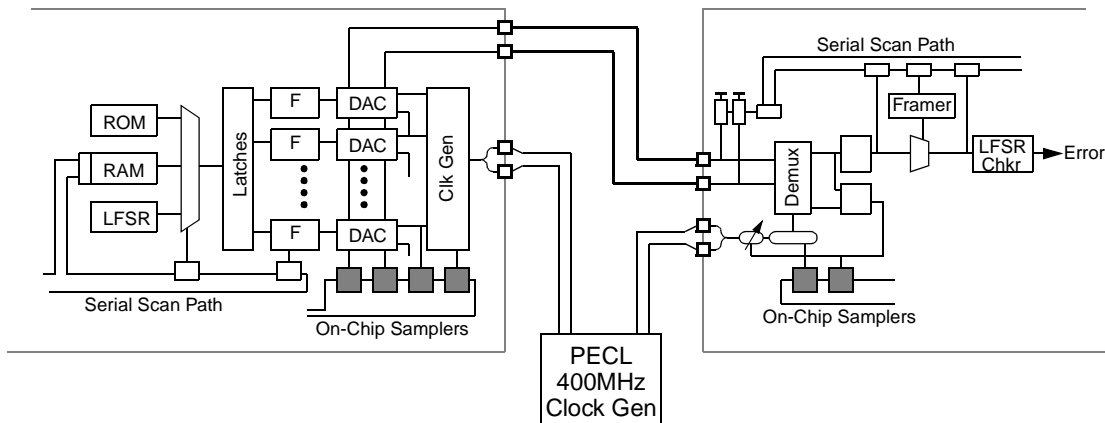
Figure 9: Organization of test chip. Each chip contains one transmitter and one receiver; circuitry for the two functions is almost completely independent.

The transmitter function, shown on the left of the figure, accepts data from one of three data sources: a ROM, which alternately sends the two 10-bit sequences 0101010100 and 1010101011 used to initialize the data framer in the receiver; a RAM, which can be loaded with 8 words of any desired 10-bit data; a linear-feedback shift register, which produces a $(2^{20} - 1)$-bit pseudo-random sequence. The data source for the transmitter, the contents of the pattern RAM, and the strength coefficients for the filters are loaded into registers in the transmitter using a serial scan path.

The receiver's 10 output bits are sent to an LFSR checker, the complement of the LFSR in the transmitter, that analyses the pseudo-random sequence from the transmitter and generates an error signal upon encountering an improper sequence of bits. A serial scan path controls the value of the digitally adjustable terminator, samples the data stream both before and after framing, and monitors or over-rides the framer controller.

Both transmitter and receiver are equipped with another serial data path that links an array of on-chip voltage samplers. Circuit details of the sampler is shown in Figure 10. The sampler is a simple clocked sense amp whose two inputs are the signal to be sampled and a reference voltage, generated off-chip, and brought onto the chip through only an ESD protection network. The reference voltage is low-pass filtered at each sampler with a long poly resistor and a MOS capacitor. The NMOS cap shown in the figure is replaced by a PMOS cap for sampled signals that are referenced to Vdd. The sampler is timed by an off-chip sampling clock, generated at PECL levels and distributed differentially on-chip. Each sampler has a differential buffer amplifier that converts the PECL input clock levels to full-swing CMOS signals. The sampler's output can be loaded into a serial data shifter for collection off-chip. The samplers are quite compact, occupying only about 90 x 180 lambda. Layout was done very carefully to balance capacitances on the two arms of the sense amp. Even with perfect transistor matching, it appears that offsets due to unbalanced charge injection amount to a few 10s of millivolts. We estimate the aperture of the sampler at a few 10s of psec, but plan to measure it on another test chip now in fabrication. The sample clock, generated off-chip in 100E ECL logic, can be adjusted over a 655μsec range with 20psec resolution.
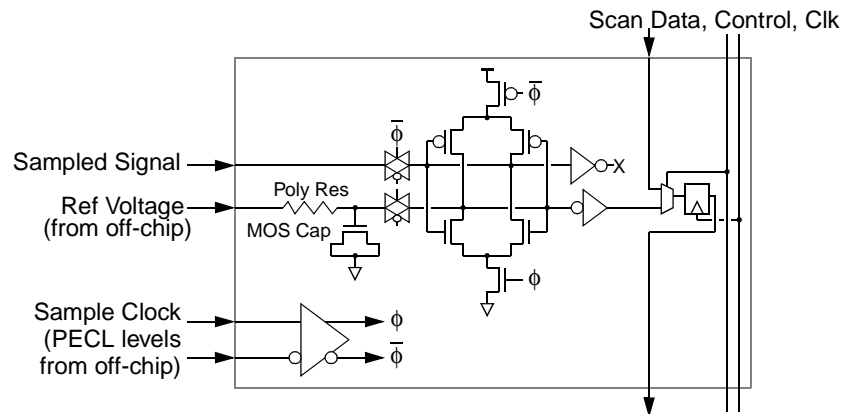
Figure 10: Schematic of on-chip voltage sampler.

A printed circuit board, shown in Figure 11, was built to evaluate the operation of the experimental signaling chips. Up to three chips can be installed on the board. One chip has its transmitter and receiver connected by a minimum-length transmission line; the others can be interconnected by 1, 2, or 3 meter PCB traces or by external coaxial or twisted-pair cables. All transmitter outputs and receiver inputs have a resistive probe brought out to SMA connectors for connection to an oscilloscope. A 400 MHz clock from an external signal generator is gated, buffered and distributed to all chips with discrete ECLinPS logic. ECL counters and delay lines allow positioning the trigger signal for the on-chip digital and analog samplers. Oscilloscope trigger outputs at the 400MHz clock rate and at fclk/8 are provided; the former is useful for generating eye diagrams; the latter corresponds to the period of the transmitter's pattern RAM. A 12-bit DAC generates a reference voltage for the analog voltage samplers. Error indications from the receiver on each chip are accumulated in 10 bit counters. The chips' serial scan chains and control signals for the test board logic are connected to a host computer through an optically-isolated standard PC parallel port.
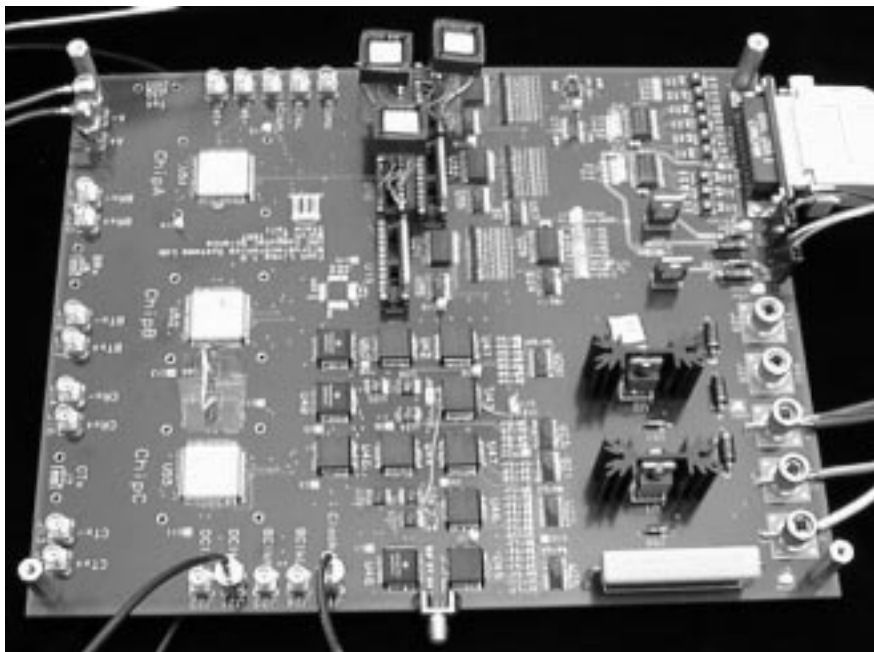

Figure 11: Photo of test board; three of the experimental 4Gb/s signaling chips are mounted on the left side of the board.

7. Experimental Results

7.1 Transmitter Timing Jitter

Measurements of the transmitter's timing jitter are shown in Figure 12. These (and all of the other measurements in this section) were performed using a Tektronix 11801A digital sampling oscilloscope. For transmitter jitter measurements, the transmitter was connected through a short PC-board transmission line to a discrete resistor termination.
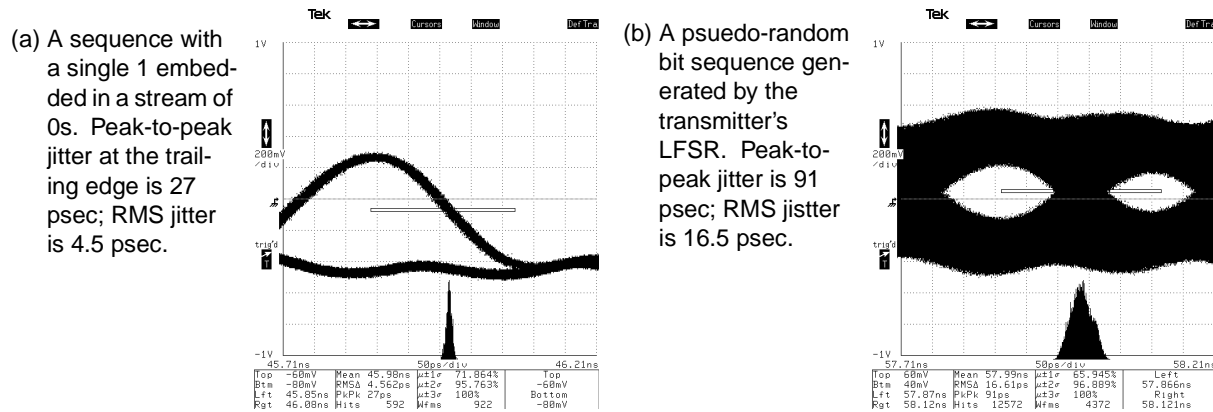
(a) A sequence with a single 1 embedded in a stream of 0s. Peak-to-peak jitter at the trailing edge is 27 psec; RMS jitter is 4.5 psec.

(b) A psuedo-random bit sequence generated by the transmitter's LFSR. Peak-to-peak jitter is 91 psec; RMS jistter is 16.5 psec.

Figure 12: Transmitter timing jitter measured with a transmitter driving a short transmission line terminated with discrete resistors.

Figure 12a shows jitter at the trailing edge of an isolated 1 in a string of 0s. This measurement is indicative of the stability of any one output of the transmitter's clock generator. Single-edge jitter is only about 4.5psec RMS, which is quite low compared to other reported jitter measurements for replica-biased DLLs. However, these measurements are on a chip that has no other active circuitry generating substrate and power supply noise. In an actual application we would expect the single-edge timing jitter to be considerably larger.

Figure 12b shows a histogram of jitter at one particular data cell boundary for a pseudo-random sequence of bits generated by the transmitter's LFSR. Jitter is about three times larger than for the single-edge case. Note that the data eye for the right-hand bit is smaller than the previous data cell. The timing generator in our transmitter exhibits a systematic timing error from cell to cell of about ±10-15psec. We believe this is a result of two effects. First, the buffer that converts the small-swing signals within the DLL to full-swing clocks does not quite preserve the symmetry of the DLL internal signals, and since we are using both clock polarities from each tap of a 180º DLL to operate the transmitter, these asymmetries introduce offsets. Second, the layout of the timing generator and transmitter places slightly different loads on various clock outputs. The first problem can be mostly eliminated with an improved buffer design that we are using in current projects, but which was not available when this transmitter was designed.   The second problem can be avoided by much more careful layout that integrates timing circuitry and output drivers into an array of absolutely identical units.

7.2 Transmitter Equalization

Figure 13 shows the results of an experiment on the effectiveness of transmitter equalization. In this experiment, the transmitter drives a short PC-board transmission line, 1m of #30 twisted-pair line, another short PC-board line, and the terminator. Each of the traces in the figure is the voltage across the differential pair at the terminator. Figure 13a shows an isolated 1 in a string of 0s with equalization turned off. Attenuation is about 0.7 with a corresponding reduction in eye height, and the width of the eye is reduced from the nominal 250psec to about 150psec. Figure 13b shows the same bit pattern with equalization turned on, and weights in the filter RAM set to 4 (out of 7) for unequal successive bits; eye width is now much closer to the nominal bit-cell width. Figures 13c and 13d show a pseudo-random bit pattern with equalization turned off (c) and on (d). Equalization recovers a recognizable eye, but attenuation in this 1m line is high enough that this probably represents the limit of the technique at its present state of development. For these experiments we set all of the weights (except the first) in the filters to the same value; with only observation of eye patterns, there was no improvement in the pseudo-random sequence eye patterns with further

adjustment of weights for non-adjacent bits. We speculate that a 2-tap transition filter may be sufficient to obtain most of the benefits of this equalization technique, allowing much simpler transmitter circuitry than we used in the test chip.
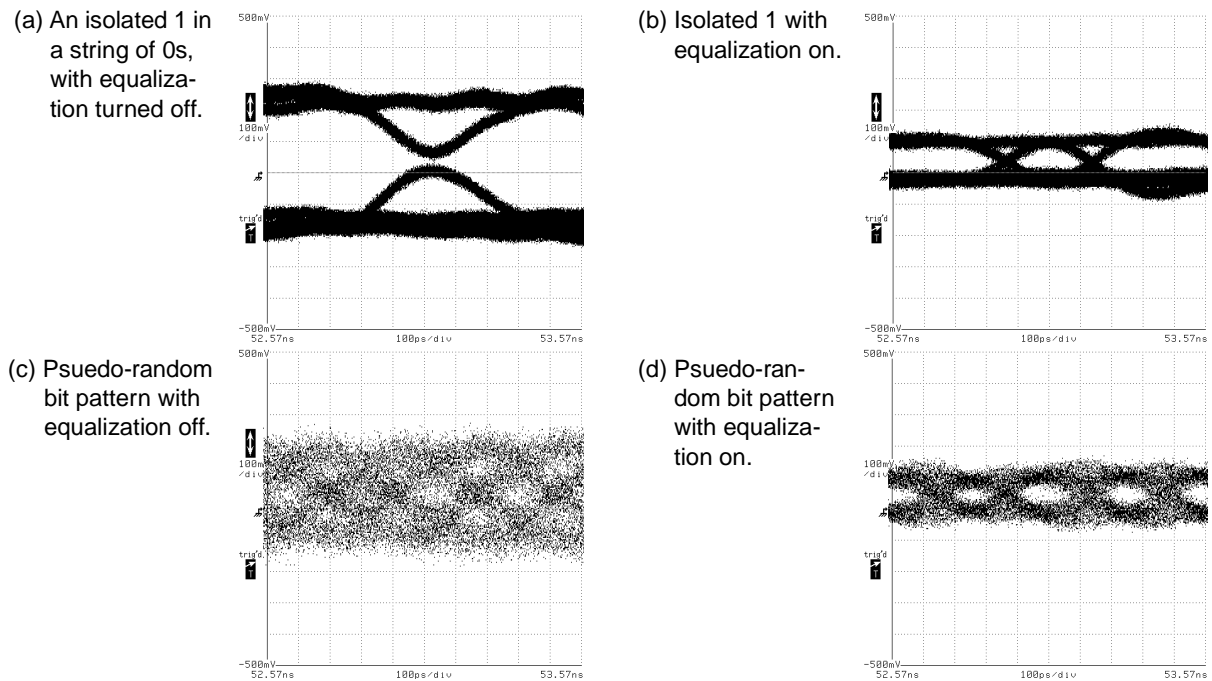
(a) An isolated 1 in a string of 0s, with equalization turned off.

(b) Isolated 1 with equalization on.

(c) Psuedo-random bit pattern with equalization off.

(d) Psuedo-random bit pattern with equalization on.

Figure 13: Transmitter timing jitter measured for a 1m twisted-pair transmission line with and without equalization.

7.3 Transmitter Problems

In the isolated-bit waveforms of the previous figures, various effects that arise from transmitter imperfections and package parasitics can be seen. The "ripple" in an otherwise quiet series of 0s, for example, is due to a slight mis-adjustment of On and Off clock overlap. Some overshoot is evident following an isolated bit (this overshoot contributes much of the amplitude noise in the pseudo-random bit patterns); this is caused by the LC tank circuits formed by the transmitter's output capacitance and bond-wire inductance, as confirmed by SPICE simulations. This effect is probably unavoidable in conventional packaging, but can be overcome almost entirely by terminating the transmitter as well as the receiver. Our future signaling chips will have both ends of the transmission line terminated. A second phenomenon that contributes significant noise is reflection off the capacitance at the receiver. These reflections are reflected again at the (high-impedance) transmitter, causing intersymbol interference at the receiver. Simulations suggest that on-chip transmitter termination will very effectively remove these reflections.

Two design "features" of the transmitter also contributed to signaling problems. First, the transmitter has a clocked (On-clock) current tail transistor (this circuit structure was the only one we could find at the time that would deliver the necessary switching speed). When the Off-clock turns on, well in advance of the time when the DAC cell is supposed to transmit its data bit, charge is injected into the output from the common source node. In our signalling system, this effects introduces an unintentional signal onto the line at about 15% of the amplitude of the actual signal. Second, and more seriously, the gating function within the filter changes the strength by completely turning off one or two of the three drivers in a DAC. The result is that, when equalization is turned on, the output signal is not truly differential. This has the unfortunate effect of sending a common-mode signal down the line, where various non-linearities transform the common-mode signal into a differential (noise) signal. A new transmitter design will take care of both of these problems.

7.4 On-Chip Samplers

Figure 14a shows a graphical display of data from a sampler run in which two of the transmitter clocks and the output differential pair are measured. For comparison, an HSPICE simulation of the same waveforms are shown in

Figure 14b. A quantitative comparison of the simulation run and sampler data must await a full characterization of the sampler's behavior. For example, it is not yet clear whether the substantial differences in clock rise time and output pair amplitude is an artifact of the sampler's slew rate or an indication that we have not properly modeled circuit parasitics in the simulation. The samplers have been used so far only to determine whether internal circuits are behaving qualitatively correctly; in this role, they were very useful in verifying correct operation of the DLLs, for example.

A set of samplers is currently being fabricated on a second test chip with provisions for more careful characterization. We hope that in future chips, fairly precise measurements may be possible.
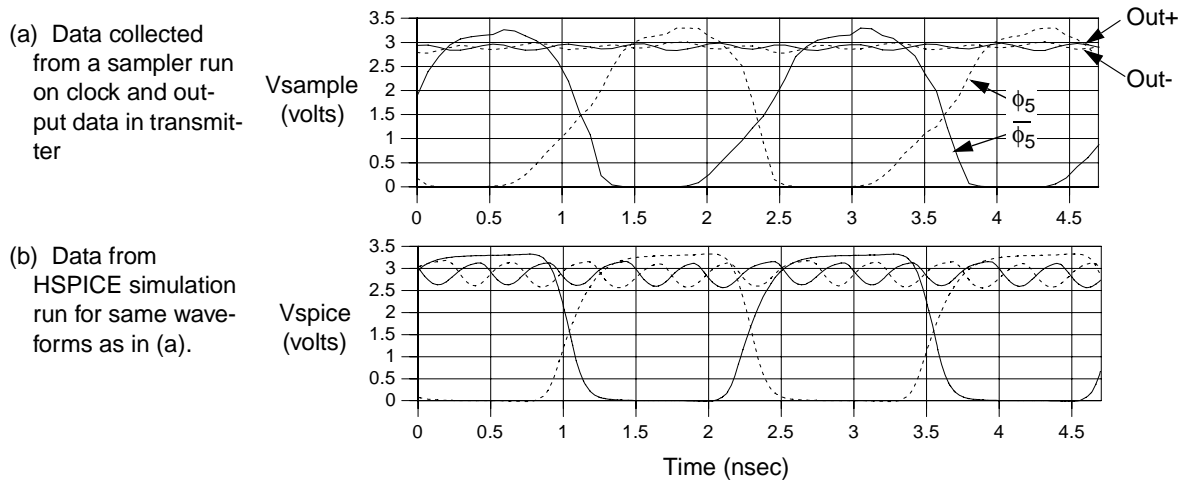

(a) Data collected from a sampler run on clock and output data in transmitter

(b) Data from HSPICE simulation run for same waveforms as in (a).

Figure 14: Sampled and simulated waveforms from the transmitter (one of the clock pairs and the output pair).

7.5 Tracking Receiver Operation

Because of the previously described problems with package-parasitic-induced noise, we were unable to test the receiver at 4Gb/s. We were, however, able to fully test its ability to acquire the framing sequence and to operate the LFSR checker at 2.5Gb/s. Because of various problems with the data transmitter, described above, bit-error rates were unacceptably high (about 10^{-4}). Simulation suggests that the BER can be brought to acceptable levels with a redesigned transmitter.

8. Conclusions and Future Work

We have described an initial experiment to verify the usefulness of transmitter equalization to remove the amplitude and timing noise introduced by frequency-dependent attenuation on a transmission line. Though only partly successful, the test chip and system verified that our simple equalization technique, a FIR filter implemented as a transition filter, is effective for this purpose. It appears, in fact, that only a simple 2-step transition filter is necessary to reap most of the benefits of the approach. This type of filter can be implemented with very simple circuitry. We are able to report only correct operation of the tracking clock recovery receiver; operation at full speed and meaningful measurements of bit error rates will have to wait for the next iteration of the design.

We are commencing the design of a new test chip with a number of improvements. The transmitter will be a true current-steering design with an adjustable on-chip terminator. We plan to include only a 2-step transition filter equalizer in this design, and therefore expect it to be much more compact and lower power. An improved DLL, clock buffer, and layout will solve the systematic timing problems encountered in the first chip. We also plan to build two receiver systems, one with tracking clock recovery and one with oversampling clock recovery. Implemented on the same die, in the same noise environment, and using many of the same circuit components, these should provide a useful experimental comparison of the two clock recovery techniques.

Our longer term plan is to improve both the speed and efficiency of high-performance signaling. In addition to increasing our transmission rate we are currently investigating methods of sending more than one bit per baud. These

include both multi-level signaling and phase modulation using a weighted sum of quadrature square waves. On the efficiency front we are looking at techniques that reduce the area and power dissipation of our prototype signaling systems. In particular we are investigating alternative timing circuits based on regulated inverter delay lines and regenerative circuits that avoid the static power dissipation of the source-coupled delay lines we currently employ.

9. Acknowledgement

10. References

[1] Yang, Ken C-K, and Horowitz, M., "A 0.8-μm CMOS 2.5 Gb/s Oversampling Receiver and Transmitter for Serial Links," *IEEE JSSC*, Vol. 31, No. 12, pp 2015-2023.

[2] Chen, D-L, and Baker, M., "A 1.25Gb/s, 460mW CMOS Transceiver for Serial Data Communication," *Proc. of ISSCC97*, pp 242-243.

[3] Bull Serial Link Technology, B.P. 68, PC F1-1D-16, rue Jean-Jaures, 78340 Les Clayes-sous-Bois, France.

[4] Fiedler, A., Mactaggart, R., Welch, J., and Krishnan, S., "A 1.0625Gbps Transceiver with 2x-Oversampling and Transmit Signal Pre-Emphasis," *Proc. of ISSCC97*, pp 238-239.

[5] Dally, W. , and Poulton, J., "Transmitter Equalization for 4Gb/s Signalling," *Proc. of Hot Interconnects '96*, pp 29-39, Stanford University, Aug 15-17, 1996. Also in *IEEE Micro*, pp 48-56, Jan/Feb, 1997.

[6] Matick, R. E. *Transmission Lines for Digital and Communication Networks*, McGraw-Hill, 1969.

[7] Bakoglu, H.B. *Circuits, Interconnections, and Packaging for VLSI*, Addison-Wesley, 1990.

[8] Montaro, J., et. al., "A 160MHz, 32-b, 0.5-W CMOS RISC Microprocessor," *IEEE JSSC*, Vol. 31, No. 11, pp 1703-1714.

[9] Manetis, J., and Horowitz, M., "Precise Delay Generation Using Coupled Oscillators," *IEEE JSSC*, Vol. 28, No. 12, pp 1273-1282.

[10] Manetis, J., "Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques," *IEEE JSSC*, Vol. 31, No. 11, pp 1723-1732.