

A 2.4 Gb/s/pin Simultaneous Bidirectional Parallel Link with Per-Pin Skew Compensation

Evelina Yeung, *Student Member, IEEE*, and Mark A. Horowitz, *Fellow, IEEE*

Abstract—This paper describes voltage and timing margins and design trade-offs in low-cost parallel links. Results from a transceiver prototype demonstrate that per-pin skew compensation improves timing margins in these parallel links and can be implemented with reasonable cost overhead. Single-ended and simultaneous bidirectional links are viable alternatives to the traditional differential and unidirectional systems—these links require fewer pins and wires for the same bandwidth, and the additional noise sources, while significant, can be managed by careful circuit and package design.

Index Terms—Parallel links, simultaneous bidirectional links, single-ended links, skew compensation, timing error, voltage noise.

I. INTRODUCTION

IN multi-chip digital systems, the overall system performance depends on both the on-chip computation speed and the I/O bandwidth. The need for high I/O bandwidth has led to the widespread use of point-to-point parallel links [1]–[4]. For these links, the design goal is to increase the bit rate per I/O while maintaining low cost (in area, power, and complexity) for the I/O circuitry. The cost constraint is very important for systems that have large numbers of high-speed I/Os.

Conventional parallel links are generally source-synchronous, with a clock sent along with the data signals for receiver timing recovery [1]. Technology scaling decreases the cost of transistors faster than it decreases the cost of I/O pins, making signalling setups that reduce pin count attractive alternatives. However, these schemes introduce larger voltage noise and hence require careful design and more complex circuitry for robust operation. One such scheme is single-ended signalling [2], [5]–[8], where the receiver compares each signal to a shared reference. Noise coupled to this shared reference voltage decreases signal margins of the links. Further reduction in pins can be achieved with simultaneous bidirectional signalling [9]–[14], where signals in both directions are superimposed on the same wire. The receiver in each transceiver subtracts its own transmit signal from the line voltage to generate the receive signal. The coupling of the transmit signal to the receive signal creates a number of extra noise sources,

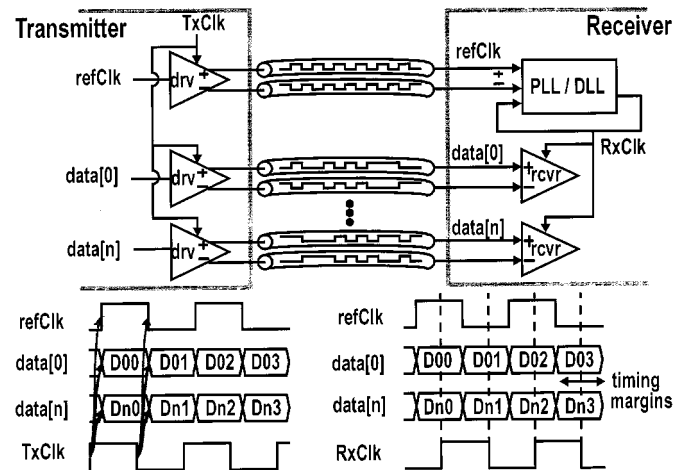


Fig. 1. Conventional source-synchronous unidirectional and differential point-to-point parallel link, where the clock is sent along with data for easier receiver timing recovery.

which also make signal margins dependent on the relative timing of the transmit and receive signals.

Designers have used current-integration [15]–[17] to filter the high-frequency reference noise in single-ended links. This paper investigates the voltage noise sources in single-ended and simultaneous bidirectional links, and extends the use of current-integration to simultaneous bidirectional receivers.

Using low cost electrical components can also help to reduce the cost of a system, but these wires are usually poorer in matching and hence create larger differences in their transmission delay paths. Consequently, larger inter-signal timing skews at the receiver reduce receiver timing margins. This paper demonstrates a per-pin skew compensation architecture and evaluates its benefits and design trade-offs.

Section II reviews source-synchronous point-to-point parallel link design and examines the voltage and timing errors for different signalling schemes. Section III describes a parallel link test chip in detail, explaining the architecture, clock recovery, input receiver design, and link performance. In addition to the core functions, the chip contains some measurement and testing circuits, such as voltage samplers to probe high-speed on-chip signals and a number of receiver clock generation circuits to experiment with jitter tracking. Section IV presents and interprets results from the per-pin skew compensation tests, the jitter tracking experiments and the voltage margin and voltage noise measurements. Finally, in Section V, we summarize our findings and discuss the implications on the design of high-performance and low-cost parallel links.

Manuscript received April 2, 2000; revised June 26, 2000. This work was supported in part by the Defense Advanced Research Projects Agency, and by Grants from LSI Logic and Texas Instruments Incorporated.

The authors are with the Computer Systems Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: evelina@vlsi.stanford.edu).

Publisher Item Identifier S 0018-9200(00)09435-X.




Fig. 2. Inter-signal skew reduces timing margins at receiver. As inter-signal skew increases, overall timing margin of the link decreases.

II. SOURCE-SYNCHRONOUS POINT-TO-POINT PARALLEL LINK

Fig. 1 shows a conventional interface architecture that forms the framework of modern source-synchronous point-to-point parallel link designs. It also shows the timing of the corresponding interface signals. Signalling is unidirectional and differential, and a stream of uncoded non-return-to-zero (NRZ) binary data is sent along each pair of tightly coupled wires. All data signals ($data[0-n]$) and a reference clock signal ($refClk$) are transmitted synchronously (hence the name source-synchronous) on both edges of the transmitter clock ($TxClk$). At the receiver, a phase-locked loop generates a global receiver clock ($RxClk$) by delaying the received reference clock by half of a bit time. This $RxClk$ is then used to sample all incoming data signals in the middle of their transitions to maximize timing margins.

The presence of timing errors, however, shifts the transition edges of the received data signals relative to the transition edges of $refClk$ and narrows timing margins. In parallel links, the phase error of concern is the inter-signal phase error or, more precisely, the deviation in phase of each data signal relative to $refClk$ (and hence to $RxClk$). This phase error can be decomposed into a dc phase offset (skew) and the dynamic phase noise (jitter). The inter-signal skew problem is illustrated in Fig. 2. Any static phase offset in clock recovery shifts the sampling point away from the optimal center and further narrows the timing margins. While designers have different opinions on the magnitude of inter-signal skew resulting from circuit mismatches in careful designs, most agree that the skew coming from interconnect mismatches is becoming a problem. Delay measurements of commercial parts have shown skews as large as 50–60 ps per meter of cable, per meter of printed-circuit board trace, or per connector [18], [19]. The total mismatch as a percentage of bit time obviously gets worse as the bit times continue to scale.

Fortunately, skew is a static phase error and can be compensated. More and more interface designs have incorporated per-pin deskewing functions [18], [20]–[22]. On start-up, a calibration mode is initiated, where each bit's skew relative to a timing reference is found using some digital control logic. The skew information is stored and is used to control the delay of an adjustable delay chain. Either the local transmitter clock [22] or the local receiver clock [4], [18] is shifted by this amount. The adjustable delay chain can be realized by activating a different number of stages [18], [22], by adjusting the delay per stage or by using phase interpolation [4]. In this paper, a per-pin skew compensation architecture, using phase interpolation to enable full-range compensation, is described in Section III, and

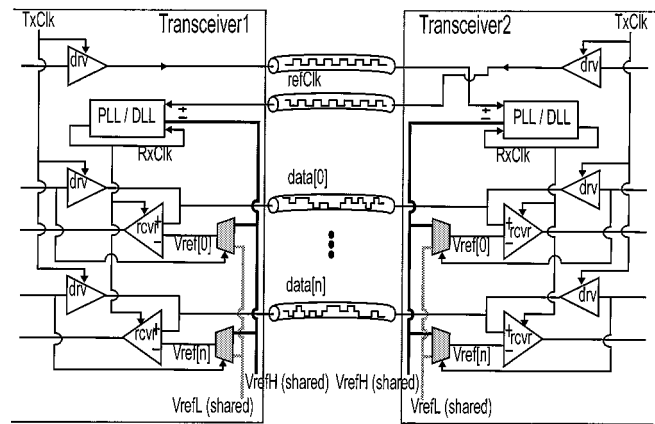


Fig. 3. Single-ended, simultaneous bidirectional parallel link interface. Signals travelling in both directions are superimposed on the same wire. Receiver subtracts out its own transmitted waveform to recover the incoming signal.

the measured inter-signal timing skew results from the test chip are presented in Section IV-A.

Jitter in the received signals also reduces timing margins of the links. Given the balanced nature of the $refClk$ and data signals at the transmitter, the jitter in the data signals may be correlated with the jitter in $refClk$; therefore trying to track the jitter in $refClk$ in receiver timing recovery circuitry may be beneficial. We attempt to answer this question by implementing dynamic phase noise tracking which is described in Section III, and the experimental results of which are presented in Section IV-B.

As mentioned earlier, single-ended signalling reduces the number of pins and wires while delivering the same total bandwidth, but operates with reduced voltage margins because of the presence of larger noise sources. Since each input signal is now compared to a reference voltage V_{ref} , any noise on this reference affects signal margins. The dc component of this noise is usually called reference offset and is caused by mismatch between the reference value and the signal swing. The major source of ac noise is from the coupling of on-chip V_{dd} and Gnd onto the signal wires. V_{ref} and data are coupled to the supplies differently making rejection of power supply noise imperfect—specifically, V_{ref} is more heavily coupled to the power supply at high frequencies than each data signal, so high-frequency power supply noise that is coupled to V_{ref} is not common-mode. Even worse, the magnitude of the power supply noise may also increase in a single-ended system, because the power supply now acts as a shared current return path for the I/O signals.

In addition to the noise on the reference line, another noise source is capacitive and inductive crosstalk coupling between signals. Unidirectional link designers need to worry about far-end crosstalk only, which is often smaller than near-end crosstalk.

Further pin saving is achieved in a single-ended, simultaneous bidirectional parallel interface, illustrated in Fig. 3. Signals travelling in both directions are superimposed on the same wire, giving a tri-level resultant waveform. To recover the incoming signal, the receiver in each transceiver must subtract its own transmitted waveform. This is usually done by multiplexing two shared reference voltages (V_{refH} and V_{refL}) to generate the local

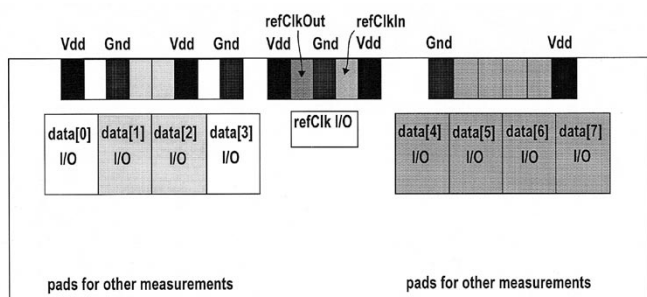


Fig. 4. I/O pad placement. Data pads are laid out with different signal return configurations to study crosstalk.

reference voltage ($V_{\text{ref}}[n]$), which switches to track the transmit signal.

However, noise issues grow even worse for this design due to the extra noise sources induced by the coupling of the transmit signal to the receive signals both on the same wire and on adjacent wires. The extra noise sources caused by the coupling of the transmit signal to the receive signal on the same wire are often termed reverse-channel crosstalk. Clearly, mismatches between the two reference levels reduce margins, but a difference in the timing of the transmitter output and reference also reduces signal margins or can even cause a glitch at the receiver input. Even worse from a noise perspective is that reflections now directly reduce voltage margins. A single reflection of the transmit signal due to impedance discontinuities and termination mismatches will appear as noise to the incoming signal. Reflection noise is less of an issue for double-terminated unidirectional lines since only even reflections reach the receiver. The coupling of the transmit signal to the receive signals on adjacent wires is caused by direct capacitive and inductive crosstalk. In simultaneous bidirectional links, both near-end crosstalk and far-end crosstalk reduce voltage margins.

We present in Section IV a systematic way of measuring the internal voltage and timing margins of links and use these measurement results to study the effects of different signalling setups.

III. PARALLEL LINK TRANSCEIVER TEST CHIP

The parallel link transceiver test chip was fabricated in a $0.35\ \mu\text{m}$ ($0.4\ \mu\text{m}$ drawn) CMOS process. Each test chip has eight single-ended data lines which are capable of simultaneous bidirectional data transmission. Each pin contains high-speed voltage samplers to display on-chip signals and to measure internal voltage margins of the links and inter-signal crosstalk, and per-pin timing adjustment to compensate for inter-signal skew and to measure timing margins. The test chip also has a set of optional unidirectional reference clock (refClk) lines, as shown in Fig. 4, to evaluate dynamic phase noise tracking. The test chip has three operational modes: in the default mode, a refClk signal is unnecessary and a ‘clean’ system clock is used for receiver clock generation; in the second mode, the receiver timing dynamically tracks the phase noise of the source-synchronous refClk signal; and in the third mode, the receiver timing dynamically tracks the phase noise of a filtered

version of the refClk signal using an additional dynamic phase noise tracking loop. The I/O pads are laid out with different signal return configurations, also shown in Fig. 4, to study crosstalk in parallel links. The die occupies a total area of $1.7 \times 3.8\ \text{mm}^2$, and a die photo is shown in Fig. 5.

The transceiver architecture supports per-pin timing adjustment by adding a variable delay to the global clock in each I/O cell. Fig. 6 illustrates the receiver, which uses current-integrating receivers [17]. In the calibration phase, a clock sequence is sent along each data line, and each variable delay element is adjusted so that RxClk is centered around the transition edges of the calibration clock sequence at the end of the calibration phase. Then a 90° phase shift is added to each delay element so during the data transmission phase the local RxClk is aligned in phase with the incoming data stream. Since the actual receiver is used for timing calibration, this architecture calibrates and compensates for all static inter-signal timing errors at the receiver.

Fig. 7 shows the actual implementation, which uses phase interpolation to realize the variable delay element. Using interpolation allows a 360° unlimited phase adjustment range and hence there is no restriction on the timing of the incoming refClk and data signals relative to the on-chip clocks at the receiver¹.

The core data loop consists of a shared core delay-locked loop (DLL), a shared finite-state-machine controller (FSM), and the eight bidirectional I/O cells. A transmitter delay-locked loop (not shown) generates a transmitter clock (TxClk), and a finite-state-machine clock (FSMClk) at a divided-by-4 frequency. The data source to each I/O transmitter can either be a pseudo-random bit sequence (PRBS) or an externally loaded data pattern.

The core DLL generates six differential clocks at 30° phase spacings [23], [24] that are distributed to all the I/Os using low-swing differential buffers [25], [26]. In the default operation mode, a ‘clean’ system clock (cleanClk) is used for clock generation². As mentioned earlier, on start-up, the chip undergoes a calibration phase during which the transmitter sends a clock stream along each data line. The data pins are calibrated sequentially using the shared FSM. Inside each I/O, the two current-integrating receivers serve as phase detectors that compare the phase of the incoming clock stream to the phase of the local RxClk. In calibrating a data pin, the FSM takes a majority vote of all eight early/late samples collected from its current-integrating receivers in each cycle and decides which direction to adjust the phase controls. When its RxClk is centered around the transition of the incoming clock stream, as shown in Fig. 6, the FSM quadrature-shifts the phase controls and stores them inside its registers. This required quadrature phase shift is performed easily by a change in phase controls—the phase moves by three 30° clock spacings. Then the FSM advances to calibrate the next pin. After all pins are calibrated, the FSM turns off. Data transmission begins, and the stored phase controls inside each I/O keeps its RxClk aligned in phase with the incoming data

¹The interface functions correctly when the maximum inter-signal timing skew between any pair of the refClk and data signals is within one cycle time (or twice the bit time).

²For simplicity, the same ‘clean’ system clock is used in both the transmitter chip and the receiver chip to avoid the control overhead needed to handle any frequency difference between the transmitter and the receiver.

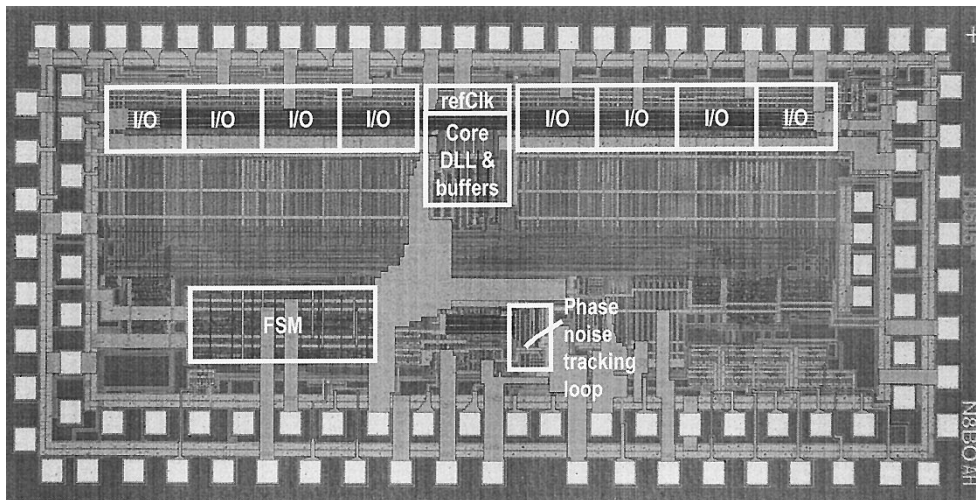


Fig. 5. Die photo.

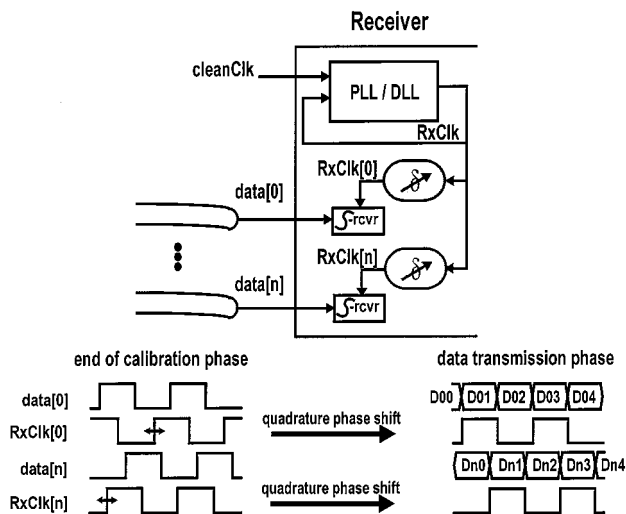


Fig. 6. Receiver section of transceiver architecture. A variable delay is added locally to global receiver clock to support per-pin skew compensation.

stream. As mentioned earlier, a refClk signal is not needed in this operation mode.

Fig. 8 is a schematic of the I/O front-end. The transmitter employs 2:1 multiplexing to transmit data on both clock phases. The open drain output driver is broken down into four segments ratioed 1:2:4:4 to give eleven levels for swing control. The swing control logic is embedded inside the transmitter datapath. The reference-select mux is broken down into four similarly weighted segments to adjust the delay of V_{ref} to match the transmit signal path delay. The two shared reference voltages (V_{refH} and V_{refL}) are externally adjusted to measure internal voltage margins. The signal wire is terminated on each side with a pMOS resistor, whose gate voltage is adjusted externally for impedance control. On-chip voltage samplers are placed at both the data and V_{ref} nodes to probe the internal signals. Finally, two current-integrating receivers [17] are used to integrate the input over the entire bit time, filtering out the high-frequency noise and the potential glitch caused by mismatched V_{ref} and transmit data delays.

Sampling on-chip signals is a useful technique for the testing and measurements of integrated circuits [27], [28]. A fast on-chip voltage sampler, illustrated in Fig. 9, is placed at every V_{data} and V_{ref} node to display on-chip waveforms and to measure inter-signal crosstalk. A source follower stage between the master and slave prevents charge-sharing between the nodes marked 'hold' and 'sample' which would otherwise impose a bandwidth limitation. Some glue logic allows the enabling or disabling of each sampler and hence different sampler outputs are multiplexed to reduce the total number of pins needed to implement this on-chip probing technique. Using a sampling clock (sampleClk) at a slightly lower frequency f_2 than the on-chip periodic signal frequency f_1 , the sampler output is a replica of the on-chip signal at the beat frequency of the two ($f_1 - f_2$).

The samplers need both time and voltage calibration to ensure the accuracy of resulting waveforms. A changing V_{signal} value causes the sampling pMOS pass gate to turn off at a slightly different point on the rising edge of sampleClk. Therefore, the sampler output exhibits a voltage-dependent time shift that also depends on the slew rate of sampleClk. This time shift is measured to be an almost linear function: 12 ps for every 100 mV that V_{signal} is below the supply. Voltage calibration is also necessary to compensate for nonlinearity of the samplers and up to 100 mV of random offsets.

Another on-chip measurement circuit is the dynamic phase noise tracking loop. As mentioned earlier, given the balanced nature of the refClk and data lines at the transmitter, the phase noise in each received data signal may be correlated with the phase noise in the received refClk; therefore, tracking the dynamic phase variations in refClk at the receiver timing (by moving the local RxClk of each data pin) may be beneficial. The delay in the core data loop clock generation limits the bandwidth of this tracking. Therefore if the phase noise is higher in frequency than this bandwidth, or the phase noise on the inputs is uncorrelated, trying to track the noise will decrease the overall quality of the link.

The delay through most of the circuit stages in the clock generation loop scales with the bit time (T_{bit})—the only excep-

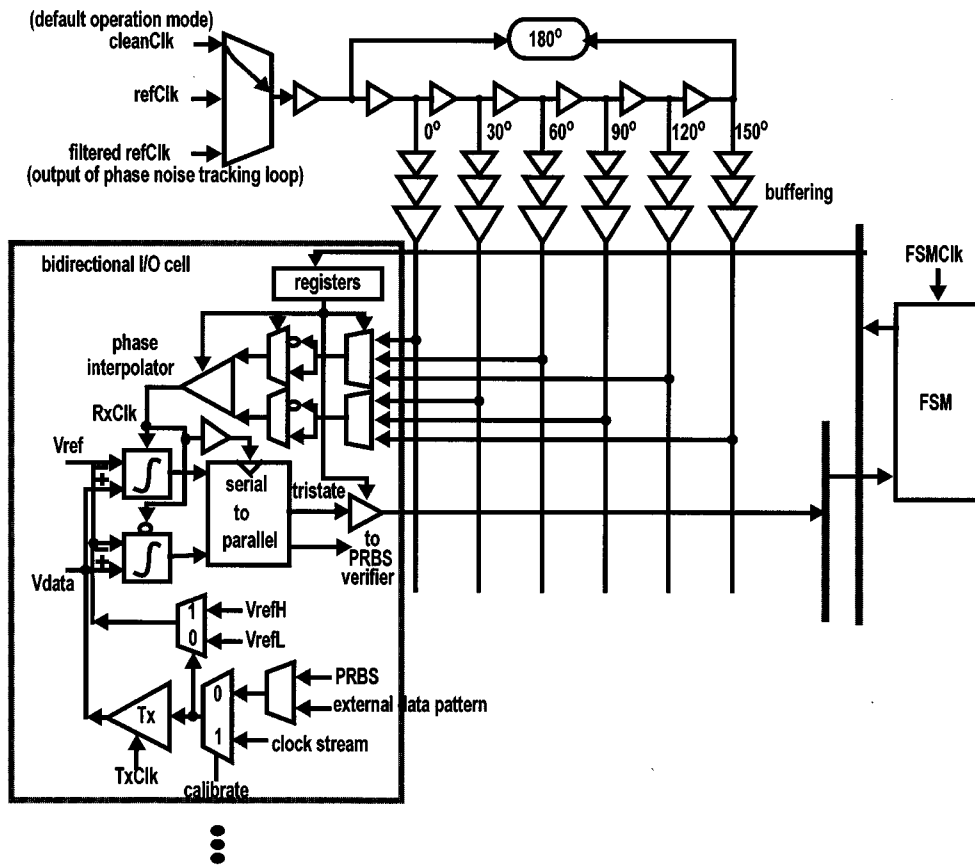


Fig. 7. Transceiver implementation (core data loop). Core DLL generates six differential clocks at 30° spacings that are phase-interpolated to generate a local receiver clock (RxClk) of unlimited phase range in each I/O.




Fig. 8. I/O front-end, using segmented open-drain output driver and current-integrating receivers.

tion is in the differential-to-single-ended converter and its subsequent buffers. The delay also depends on the phase settings inside each I/O cell. The maximum total delay from the received refClk to the local RxClks (T_d) is roughly $3.3 * T_{bit} + 5 * FO4$, where FO4, fanout-of-4 delay, is the delay of an inverter driving a load of four identical inverters. Theoretically, if tracking results in a phase shift of less than 90°, the correction is in the

right direction and hence is beneficial. Using this phase relationship, the maximum ‘track-able’ noise frequency is equal to $1/(4 * T_d)$.

To test whether clock jitter tracking will help in this type of link, the clock for the core DLL has three possible sources as outlined earlier. The default is to use the ‘clean’ clock, an external reference that is driven into the chip. If the main phase noise is below the track-bandwidth allowed by the clock buffer delay, feeding in the received refClk should improve performance. On the other hand, if the phase noise is mostly above the track-bandwidth, but there also is low frequency phase noise, then using a filtered version of the received refClk would perform best. This option is also possible in the test chip by using an phase noise tracking loop, as shown in Fig. 10—the phase noise of the input refClk is filtered by using it to drive the feedback on another DLL. Thus the output of this DLL contains only the low frequency phase noise of the received refClk.

We test the data communication between two chips for link performance. Each data channel consists of bond wires, package wiring, PC board (GETEK) traces totalling >6 inches (3 inches on each board, drawn radially from the package to balance the traces), a coaxial cable ranging from 36 to 42 inches and some SMA connectors. The link speed is limited by clock generation, which was designed for a clock period of 8 fanout-of-4 delays ($FO4 = 193$ ps in this process). At 3.3-V supply, the bidirectional links achieve a data rate of 2.4 Gb/s/pin (1.2 Gb/s in each direction) with no reception error observed for the entire testing

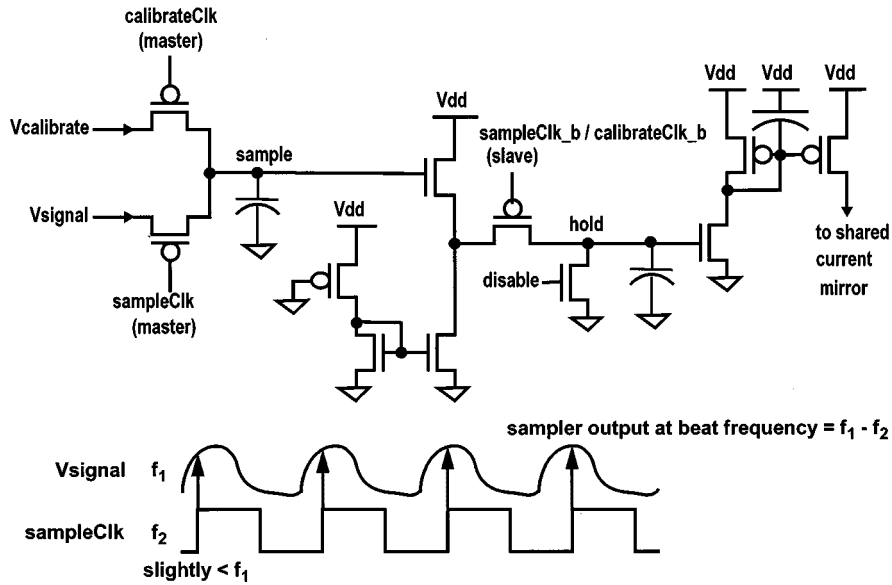


Fig. 9. On-chip voltage sampler (buffered sample and hold).

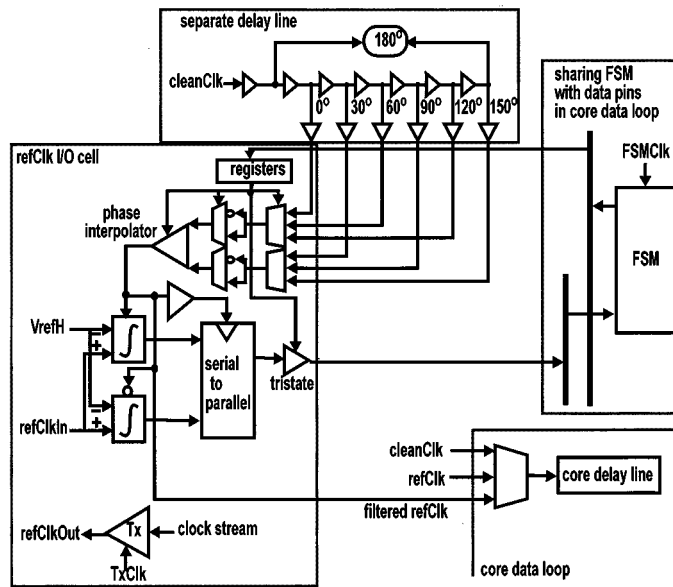


Fig. 10. Dynamic phase noise tracking loop. The received refClk signal is filtered by another dual-loop DLL.

period of more than 15 hours, representing a bit-error rate (BER) $< 8 \times 10^{-15}$. At this data rate, the links require a minimum signal swing of 194 mV on each side in the pins with worst-case crosstalk. The chip dissipates < 1 W total power when all the links are running at 2.4 Gb/s/pin at their largest swings (about 430 mV) and when all on-chip measurement blocks are active.

IV. VOLTAGE AND TIMING NOISE MEASUREMENTS

With the built-in testing and measurement capabilities, we are able to measure the internal voltage and timing margins of the links, illustrated in Fig. 11, in a systematic way and study the voltage and timing noise sources. To measure voltage margins, the links are first calibrated by setting V_{ref} at the middle of the nominal signal swing V_{swing} . Keeping the phase controls

inside all I/O cells (and hence the positions of all local RxClks) fixed, V_{ref} is moved up and down, and the first boundary points at which each link starts to fail are recorded, the difference of which is the voltage margin. This measurement has a 1-mV resolution. To measure timing margins, we set V_{ref} at the middle of the nominal signal swing and calibrate the links. Then, while keeping V_{ref} fixed, we measure the timing margin of each link, by shifting the local RxClk at nominal timing steps equal to 8.7 ps in both directions. The boundary points at which bit errors start to appear are recorded, and the interval between these two points is the timing margin.

The signal margins of bidirectional links are measured in similar steps. Because the transmitter output swing is fairly linear in bidirectional signalling, a fixed $V_{\text{ref}L}$ equal to $1.5 * V_{\text{swing}}$ below the supply is used. The voltage margin of each link is measured by varying $V_{\text{ref}H}$ while keeping RxClk fixed. The timing margin is measured by shifting RxClk while keeping $V_{\text{ref}H}$ fixed. Each passing value in the signal margin measurements has a BER $< 10^{-11}$. Unless otherwise specified, all measurements are taken with all of the circuit blocks turned on to simulate the power supply noise in a real mixed-signal system.

A. Inter-signal Timing Skew

To test the per-pin skew compensation capability, two sets of experiments are carried out using the setup described earlier. In both tests, the unidirectional links run at 1.2 Gb/s at their maximum swings. Calibration results are shown in Fig. 12. The bars³ show receiver timing margins of different signal pins,⁴ their calibrated eye centers, and ideal eye centers.

Initially we used a 36-inch cable in each data channel and carefully match the delays of all paths. The links are calibrated, and the results show a maximum phase difference of 191 ps in the calibrated eye centers between the fastest pin (data[1]) and

³The positions of the centers and the widths of the eyes are scaled appropriately by the bar charts.

⁴Unfortunately, data[6] is mistakenly bonded to a non-I/O pin and its measurement results are ignored.




Fig. 11. Voltage and timing margins of links.

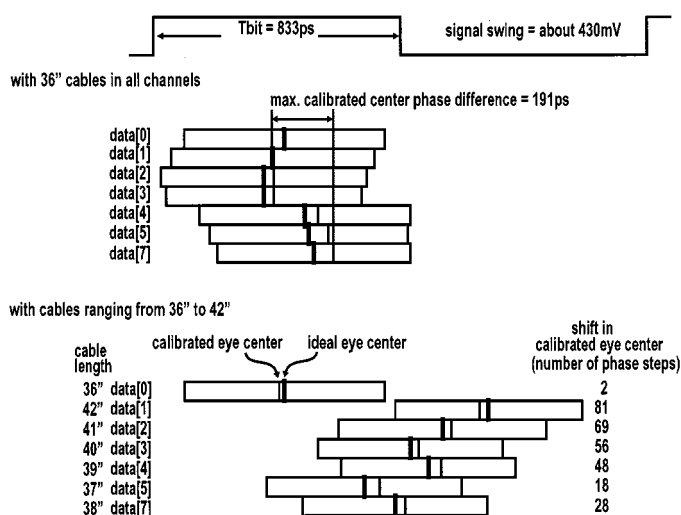


Fig. 12. Receiver timing margins for skew compensation tests. Calibrated eye centers shift as skews increase.

the slowest pin (data[7]). The on-chip data waveforms indicate that approximately 100 ps of this difference is due to inter-signal skew, about half of which can be attributed to differences in the signal traces in the packages used. The calibration results show one possible problem with our calibration scheme. For the signals with significant neighbor coupling (data[4], data[5], and data[7]), the crossing point between the signal and V_{ref} moves by about 90 ps when the neighbor signals transition in sync with the signal compared to when the neighbor signals are idle. These two components account for the observed 191 ps maximum phase difference in the calibrated eye centers. Then cables ranging from 36 to 42 inches in length are used to deliberately introduce more skew. Calibrated eye centers shift as skews increase, showing that the circuit is able to deal with larger skews without reducing timing margins.

The overhead for implementing this per-pin skew compensation scheme is modest. The phase muxes, phase interpolator, and associated registers double the size of the I/O cell, and the total area becomes $192 \times 275 \mu\text{m}^2$. The static currents in the phase muxes, phase interpolator and differential-to-single-ended converter draw an additional 13.7 mW of power per I/O cell.

Receiver timing margins using different clock inputs for receiver clock generation

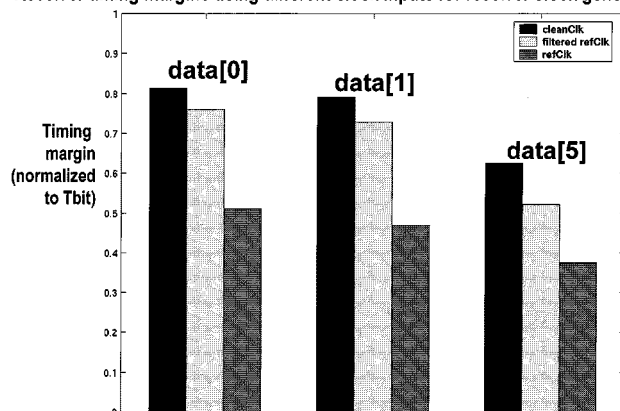


Fig. 13. Receiver timing margins of unidirectional links using different clock inputs to core DLL. Using clean Clk results in largest timing margins, and using received refClk gives smallest margins, showing that the phase noise contains mainly high-frequency components.

B. Dynamic Phase Noise

To evaluate dynamic phase noise characteristics of interface signals, receiver timing margins of unidirectional links are measured using the three different inputs to the core data loop. The results are shown in Fig. 13 for the three pins with different signal return configurations, specifically data[0], data[1], and data[5], running at 900 Mbps unidirectional data rate. Our earlier analytical model predicts a maximum noise tracking frequency of about 52 MHz at this data rate.

The data clearly indicates that for this system the dominant phase noise is high-frequency noise, an expected result for a DLL-based system. Since there are no voltage-controlled oscillators (VCOs) to accumulate jitter near the loop bandwidth, most of the jitter is likely to be cycle-to-cycle jitter. As mentioned earlier, if the refClk signal carries both high-frequency and low-frequency noise, using the filtered refClk will give the best performance among all three options. Therefore, the fact that using the filtered refClk is also worse than using cleanClk indicates that this system experiences very little low-frequency phase drift. The extra jitter in the filtered refClk, resulted from the phase noise tracking loop circuitry, reduces receiver timing margins. One interesting result is that timing margins degrade

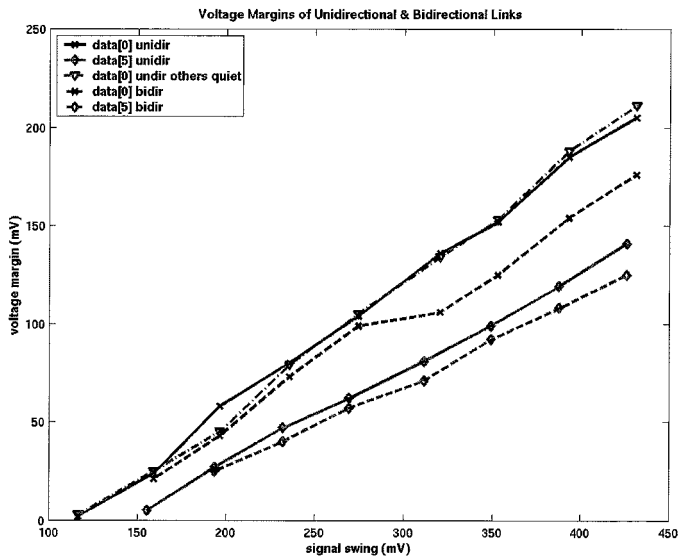


Fig. 14. Voltage margins of unidirectional and bidirectional links as signal swings vary.

from data[0] to data[1] to data[5] in all operation modes as inter-signal crosstalk increases. These voltage margins are further described in the next section.

C. Voltage Noise

To characterize the voltage noise sources, we measure voltage and timing margins of different signal pins in both unidirectional and simultaneous-bidirectional operations under different conditions by varying the transmission signal swings. The data points give a set of roughly straight lines which allow us to analyze the voltage noise sources and extract their values. Fig. 14 illustrates some of the measurement results for data[0] and data[5] transmitting PRBS data at 1.2 Gb/s unidirectional data rate or 2.4 Gb/s bidirectional data rate.

We define voltage margin as the difference between the dc voltage swing and the total noise, and postulate that the voltage noise sources decompose into two groups: noise sources which are fixed in value and noise sources whose values change proportionally to the signal swing. The negative value of the y -intercept is the fixed noise, and the slope of the line corresponds to $(1 - \text{proportional noise})$. Using a linear fit to analyze the data points for unidirectional data[0] when all other data signals are idle, we see about 75 mV of fixed noise and 33% proportional noise.

We find that more than half of the proportional voltage loss can be attributed to the way we define the signal swing, which we have defined to be the difference in dc levels when the transmitter outputs a '1' and a '0' permanently. However, when a bit stream with alternating zeros and ones is transmitted, the signal swings to only about 83% of the dc swing at the midpoints of the bit time. This 17% signal loss includes the attenuation in the transmitter board trace, which is measured to be 3%. Measurements show that another 3% is lost in the receiver board trace, and the channel loss in the cable is small enough to be ignored. Using the on-chip voltage samplers, we find that capacitive coupling from the signal onto its local V_{ref} induces a proportional

Fixed noise

data[0]	data[1]	data[5]
75mV	64mV	69mV

Proportional noise

	data[0] others quiet	data[0]	data[1]	data[5]
Unidirectional	33%	34%	37%	51%
Bidirectional	41%	45%	46%	57%

Fig. 15. Extracted fixed noise and proportional noise from voltage margin measurements.

noise of about 3% at the receiver end, and that reflection noise is small in this doubly terminated unidirectional link. Since all other data pins are idle in this set of measurements, there is no crosstalk from them. The refClk pins are always on, and their toggling activities induce a 4% crosstalk on data[0]. The proportional noise components we have identified add up to 27%.

Because we adjust the reference voltage directly in measuring the internal voltage margins, any receiver offset in an individual receiver or any reference offset is eliminated from the voltage margin measurement results. However, any difference in the receiver offsets of the two current-integrating receivers forms another source of fixed noise. Coupling onto the reference line seems to be the dominant source of fixed noise in data[0] and is caused mainly by internal clock coupling. We measure each noise component separately by considering the effect on the differential signal ($V_{\text{data}[0]} - V_{\text{ref}[0]}$). The coupling from the transmitter clock (38 mV) is correlated with data signal transitions; the coupling from the receiver clock (42 mV) depends on the phase control settings; the coupling from the system clock input (28 mV) and the coupling from the sampling clock for the on-chip voltage samplers (close to 0 mV) are uncorrelated to the signal transitions. Therefore, in the worst case, these noise sources add up constructively. The fixed noise extracted from the earlier voltage margin measurements is smaller than the summation of the measured peak-to-peak values of these coupling noise components, showing that the averaging effect of current integrating receivers improves signal margins.

We also measure individual noise components of the other data pins and observe consistent noise behavior when only the measured pin is active and all the others are idle. Then we measure their voltage margins when all data signals transmit PRBS data. A summary of the extracted voltage noise values for data[0], data[1], and data[5] is shown in the tables in Fig. 15. The fixed noise across these three pins spans a range of only 11 mV, providing evidence for consistent fixed noise sources in all pins.

However, the proportional noise differs significantly for pins with different signal return configurations due to the different forms of crosstalk, with data[5] consistently having the largest proportional noise. Interestingly, the proportional noise in data[1] is very close to that in data[0] in both unidirectional and bidirectional operations, showing that using a supply pin for every two I/O signals gives approximately the same performance as having alternating power and signal pins. The near-end crosstalk from the transmit signal to the receive signal on the same wire contributes to the increase in proportional

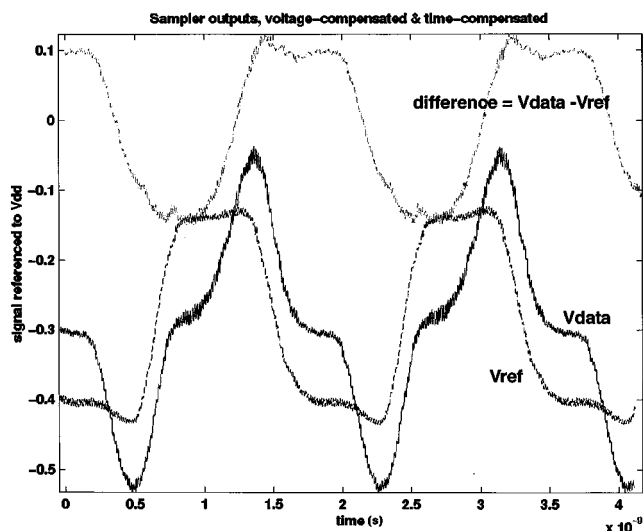


Fig. 16. Bidirectional on-chip signals. Voltage margin falls to a minimum when the transmit and receive signals are in quadrature phase as shown.

noise in all pins as the links go from unidirectional to simultaneous bidirectional signalling.

As noted earlier, the voltage margin of a bidirectional link may change as the phase relationship between the transmit and receive signals varies. Fig. 16 shows 2.2-Gb/s 280-mV-swing bidirectional on-chip signals. When the receive and transmit signals are set up to be in quadrature phase as shown, the voltage margin generally falls to a minimum. In this design, the switching of the reference and the transmitter output is well matched, and the induced glitch is small. The effect of this glitch is further reduced by the current-integrating receivers. In fact, varying the phase relationship across the bit time changes voltage margins by only 20 mV for these signals, which represents a proportional voltage noise of about 7%, and has no appreciable effect on timing margins. This observation is the combined effect of mismatched timing and any reflection of the transmit signal.

As we can see, single-ended signalling and simultaneous bidirectional signalling save pins and wires but create larger voltage noise and hence require larger signal swings to achieve the same data rate when compared to the traditional differential and unidirectional system.

V. CONCLUSION

Voltage and timing error sources limit the performance of a link and affect its robustness. The voltage and timing errors unique in parallel links, such as inter-signal timing skew and inter-signal crosstalk, impose greater challenges as the performance increases. Mass integration of I/Os requires low cost per I/O, and the use of low-cost solutions, such as using cheaper electrical components, single-ended signalling, and simultaneous bidirectional signalling, further increases the voltage and timing errors.

Experimental results from a parallel link transceiver prototype have shown that per-pin skew compensation improves timing margins in high-performance parallel links and can be implemented with a reasonable cost overhead. The phase noise

in high-speed interface signals carries significant high-frequency components, and experimental results have shown that the clock buffer delay makes tracking the jitter of a source-synchronous reference clock in the receiver difficult. Using a 'clean' clock for receiver timing recovery clock generation is the best strategy for jitter. Low-frequency phase drifts in the signals can be compensated by a periodic calibration in a system capable of skew compensation, making the source-synchronous reference clock signal unnecessary.

Measurement results also show that single-ended and simultaneous bidirectional links are viable alternatives to the traditional differential, unidirectional systems. They allow significant pin saving for the same bandwidth. The additional voltage noise sources, while significant, can be managed by careful design in circuits and in packaging.

ACKNOWLEDGMENT

The authors thank Dr. S. Sidiropoulos for providing ideas and guidance for this work, J. Kim for layout assistance, and Prof. C. K. Yang and W. Ellersick for helpful discussions.

REFERENCES

- [1] *Scalable Coherent Interface (SCI)*, IEEE Standard, 1956.
- [2] E. Reese *et al.*, "A phase-tolerant 3.8 GB/s data-communication router for a multiprocessor supercomputer backplane," in *ISSCC 1994 Dig. Tech. Papers*, Feb. 1994, pp. 296–297.
- [3] M. Galles *et al.*, "Spider: A high-speed network interconnect," *IEEE Micro. J.*, vol. 17, pp. 34–39, Jan./Feb. 1997.
- [4] K. Gotoh *et al.*, "A 2B parallel 1.25 Gb/s interconnect I/O interface with self-configurable link and plesiochronous clocking," in *ISSCC 1999 Dig. Tech. Papers*, Feb. 1999, pp. 180–181.
- [5] N. Kushiyama *et al.*, "A 500-megabyte/s data-rate 4.5 M DRAM," *IEEE J. Solid-State Circuits*, vol. 28, pp. 490–498, Apr. 1993.
- [6] K. Donnelly *et al.*, "A 660 MB/s interface megacell portable circuit in 0.3- μ m–0.7 μ m CMOS ASIC," in *ISSCC 1996 Dig. Tech. Papers*, Feb. 1996, p. 290.
- [7] S. Sidiropoulos, C. K. Yang, and M. Horowitz, "A CMOS 500 Mb/s/pin synchronous point to point link interface," in *Proc. 1994 IEEE Symp. VLSI Circuits*, June 1996, pp. 43–44.
- [8] B. Lau *et al.*, "A 2.6-GByte/s multipurpose chip-to-chip interface," *IEEE J. Solid-State Circuits*, vol. 33, pp. 617–626, Nov. 1998.
- [9] K. Lam, L. Dennison, and W. Dally, "Simultaneous bidirectional signalling for IC systems," in *Proc. 1990 IEEE Int. Conf. Computer Design*, Sept. 1990, pp. 430–433.
- [10] L. Dennison, W. Lee, and W. Dally, "High-performance bidirectional signaling in VLSI systems," in *Proc. Symp. Integrated Systems*, Mar. 1993, pp. 300–319.
- [11] R. Mooney, C. Dike, and S. Borkar, "A 900 Mb/s bidirectional signaling scheme," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1538–1543, Dec. 1995.
- [12] T. Takahashi *et al.*, "A CMOS gate array with 600 Mb/s simultaneous bidirectional I/O circuits," *IEEE J. Solid State Circuits*, vol. 30, no. 12, Dec. 1995.
- [13] M. Haycock and R. Mooney, "A 2.5 Gb/s bi-directional signaling technology," in *Hot Interconnects V Symp. Rec.*, Aug. 1997, pp. 149–156.
- [14] T. Takahashi *et al.*, "110 GB/s simultaneous bi-directional transceiver logic synchronized with a system clock," in *ISSCC 1999 Dig. Tech. Papers*, Feb. 1999, pp. 176–177.
- [15] S. Sidiropoulos and M. Horowitz, "Current integrating receivers for high speed system interconnects," in *Proc. 1995 IEEE Custom Integrated Circuits Conf.*, May 1995, pp. 107–110.
- [16] —, "A 700 Mb/s/pin CMOS signalling interface using current integrating receivers," in *Proc. 1996 IEEE Symp. VLSI Circuits*, June 1996, pp. 142–143.
- [17] —, "A 700-Mb/s/pin CMOS signaling interface using current integrating receivers," *IEEE J. Solid-State Circuits*, vol. 32, pp. 681–690, May 1997.
- [18] D. Cecchi, M. Dina, and C. Preuss, "A 1 GB/s SCI link in 0.8- μ m BiCMOS," in *ISSCC 1995 Dig. Tech. Papers*, Feb. 1995, pp. 326–327.

- [19] SCIZZL Documents on IEEE Standards Project P1596.8 (Parallel Links for the Scalable Coherent Interface) (1996). [Online]. Available: <ftp://ftp.SCIZZL.com/P1596.8/960926kbl.pdf>
- [20] *High-Performance Parallel Interface—6400 Mbit/s Physical Layer*, HIPPI-6400 PH, 1998.
- [21] *Gigabit Ethernet*, IEEE Std. 802.3z, 1998.
- [22] T. Sato *et al.*, "5 GByte/s data transfer scheme with bit-to-bit skew control for synchronous DRAM," in *Proc. 1998 Symp. VLSI Circuits*, June 1998.
- [23] S. Sidiropoulos and M. Horowitz, "A semi-digital DLL with unlimited phase shift capability and 0.08–400 MHz operating range," in *ISSCC 1997 Dig. Tech. Papers*, Feb. 1997, pp. 332–333.
- [24] —, "A semi-digital dual delay locked loop," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1683–1692, Nov. 1997.
- [25] J. Maneatis and M. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1273–1282, Dec. 1993.
- [26] J. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1723–1732, Nov. 1996.
- [27] P. Larsson and C. Svensson, "Measuring high-bandwidth signals in CMOS circuits," *Electron. Lett.*, vol. 29, no. 20, pp. 1761–1762, Sept. 1993.
- [28] R. Ho *et al.*, "Applications of on-chip samplers for test and measurement of integrated circuits," in *Proc. 1998 IEEE Symp. VLSI Circuits*, June 1998, pp. 138–139.



Evelina Yeung (S'96) received the B.S. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 1994, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1996. She is currently working toward the Ph.D. degree in the same department. Her research interest is in high-performance and low-cost parallel link design. She has been with Marvell Semiconductor, Inc., Sunnyvale, CA, since June, 2000. She has previously held internship positions at Hewlett-Packard Laboratories, MIPS Technologies (Silicon Graphics, Inc.), and Lawrence Berkeley National Laboratory.

Ms. Yeung has received honors including the U.S. Department of Energy Undergraduate Research Fellowship, Senior Women's Scholarship, and Student Speaker at the College of Engineering Commencement at the University of California, Berkeley. She is a member of Eta Kappa Nu and Tau Beta Pi, and an Asia/Pacific Scholar at Stanford University.



Mark A. Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, and the Ph.D. degree from Stanford University, Stanford, CA.

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

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