A 2.5Gb/s Bidirectional Signaling Technology

Matthew Haycock and Randy Mooney
Intel Corporation
Microcomputer Research Labs
5200 NE Elam Young Parkway
Hillsboro, OR 97124-6497

Abstract -- **This form of bidirectional signaling was introduced in 1995 [1], and then was implemented in a very large multi-processor machine [2]. That implementation demonstrated the stretch scalability of this bidirectional signaling scheme by using a single design for both short links across a few centimeters of board trace and long links across several meters of coaxial cable. This work describes the results of scaling the bidirectional signaling circuits described in [1] from a 0.6μm to a 0.25μm CMOS process and of adding circuits to dynamically adjust the output buffer impedance to match the interconnect impedance. When the impedances in the system are well matched, the bidirectional link operates at an aggregate transmission rate of 2.5Gb/s per wire (1.25Gb/s per direction simultaneously) over several centimeters of stripline on a printed circuit board. A comparison of the bidirectional and source-parallel terminated unidirectional forms of signaling, on a 'bandwidth per wire' basis, illustrates some advantages of the bidirectional signaling scheme.**

I. Introduction

This paper describes the results of scaling the bidirectional signaling circuits described in [1] from a 0.6μm to a 0.25μm CMOS process and demonstrates this signaling method transmitting data at 2.5Gb/s per wire. Section II contains a brief review of the bidirectional signaling scheme, explains the voltage swings and margins present on the link, and describes the circuits used. Section III describes the circuits used to automatically match the output buffer impedance to the board impedance and to adjust the two voltage reference levels. The results from a test chip are reported in Section IV. Section V summarizes the work and highlights some advantages this type of signaling offers.

II. Bidirectional Signaling

A. Overview

In the bidirectional point-to-point signaling scheme, two drivers are placed at opposite ends of the same wire and data is transmitted simultaneously in two directions. Each driver is impedance matched to the line and functions both as a driver and as a parallel termination for the driver at the opposite end of the wire. This source-parallel termination removes reflections from the line and produces a signal swing between either Vcc and $1/2V_{cc}$ or $1/2V_{cc}$ and V_{ss}, depending on the state of the driver at the opposite end of the line. Two references generated on the chip, at $1/4V_{cc}$ and $3/4V_{cc}$, are connected across the link to couple common mode noise to the inputs of the differential amplifier receivers. This also allows the link to track variations in the power supplies of the communicating components. The point-to-point connections between two bidirectional data ports are illustrated in Figure 1.

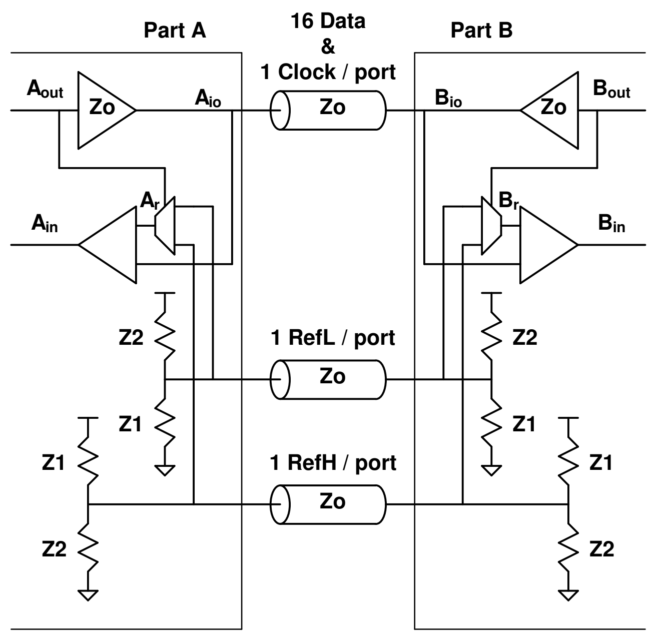

Figure 1. Bidirectional point-to-point connections between two ports.

The two references are connected across the link and are shared by the differential amplifier receivers of the data and clock signals. Sharing the references across several (17 in this case) differential receivers allows rejection of low frequency common mode noise and provides some of the advantages of full differential signaling without the overhead of two wires per signal. Reducing the number of receivers which share the references would increase the effective frequency of common mode noise that could be rejected. Each reference is built using an output driver with both the N and P devices on and set to an impedance such that the Thevenin equivalent circuit is a resistance of Z_o to a power supply equal to the desired reference voltage.

B. Swings and Levels

This bidirectional signaling scheme produces $1/2V_{cc}$ signal swings and $1/4V_{cc}$ margins, which are illustrated in the example in Figure 2.

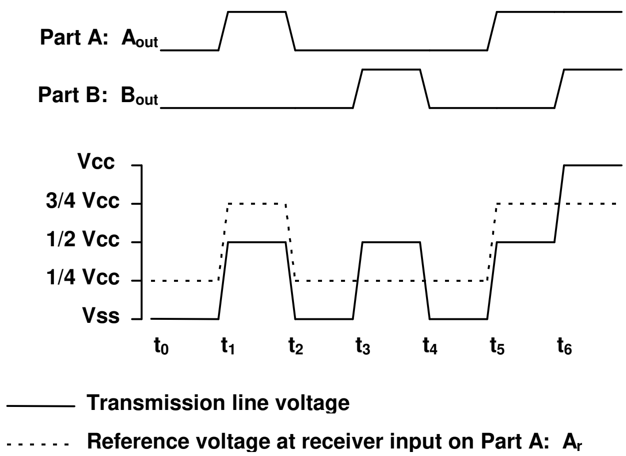

Figure 2: Example bidirectional signal swings and margins

This example refers to the circuit nodes A_{out}, B_{out}, A_{io}, B_{io}, and A_r in Figure 1. Initially at time t_0, A_{out} and B_{out} transmit a '0' which produces a voltage V_{ss} on the line. Additionally, the '0' on A_{out} selects, through the mux, the $1/4V_{cc}$ reference to the amplifier input A_r. At time t_1, A_{out} transmits a '1' while B_{out} continues transmitting a '0'. Since the outputs A_{io} and B_{io} are in opposite states, the line voltage rises to $1/2V_{cc}$. To properly decode the '0' transmitted by B_{out} in the presence of the '1' transmitted by A_{out}, A_{out} also selects the $3/4V_{cc}$ reference to the amplifier input A_r. At time t_2, A_{out} again transmits a '0' and switches the $1/4V_{cc}$ reference to node A_r. At time t_3, B_{out} transmits a '1' which is properly decoded since the $1/4V_{cc}$ reference is selected. At time t_5, A_{out} transmits another '1' and selects the $3/4V_{cc}$ reference to node A_r. Finally, at time t_6, B_{out} transmits a '1' and the line voltage rises to V_{cc}. With the $3/4V_{cc}$ reference selected at A_r, the '1' transmitted by B_{out} is properly decoded by Part A.

For Part A to correctly receive the data transmitted by Part B, A_{out} need only select the correct reference based on the data it transmits. However, this requires that the delays in the data and reference paths to the amplifier inputs be matched as closely as possible.

C. Circuits

The key circuits at the electrical interface are the output buffer, the input amplifier, and the reference multiplexer. The output buffer, shown in Figure 3, is a CMOS buffer with five binary-weighted N and P legs. The impedance is adjusted by enabling or disabling the binary weighted legs. The size of the smallest leg defines the resolution and was chosen to be able to match the buffer impedance to the board impedance within about five percent.

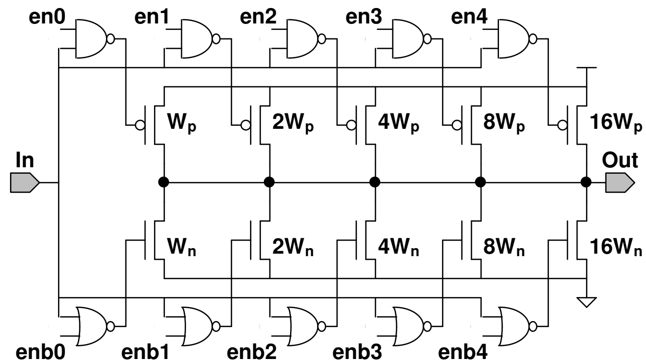

Figure 3: Binary weighted, legged output buffer

The receiver, shown in Figure 4, is a differential self-biased, folded cascode amplifier with rail-to-rail common mode range and is designed to be symmetric and insensitive to process, temperature, and Vcc variations [3].

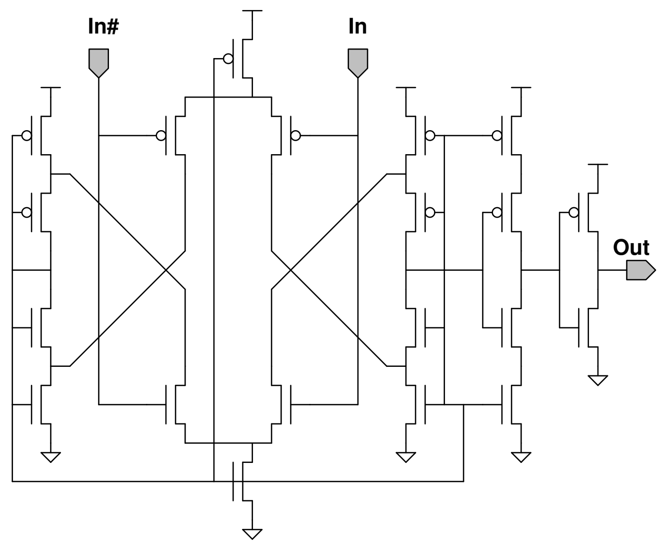

Figure 4: Differential input amplifier

The multiplexer between the references and the differential amplifier is built as a pass gate structure. The parasitic diodes associated with the mux transistors are also used as part of the ESD protection scheme. The select input of the multiplexer is controlled by the predriver of the output buffer in order to match the two paths to the differential amplifier inputs as closely as possible. Matching these two paths minimizes the timing variation in input edges when an output transition occurs. An identical multiplexer is inserted into the data path from A_{io} and B_{io} to the differential amplifiers to match these paths even more closely.

Phase tolerant clocking, in which the clock frequency on each component is the same but the phase between them arbitrary, is used at the receiver to capture transmitted data. Parallel data and a source synchronous clock are transmitted in phase across the link. After passing through the differential amplifier receiver, the source synchronous clock is delayed by one half bit cell time, to the center of the data bit cell, using a delay locked loop (DLL). The delayed clock is then used to latch the received data into the phase tolerant fifo. The data is read from this fifo and passed to the core using the receiving components clock. The latency through this fifo can be up to one cycle, depending on the phase difference between the two components [4]. The phase-tolerant clocking is illustrated in Figure 5.

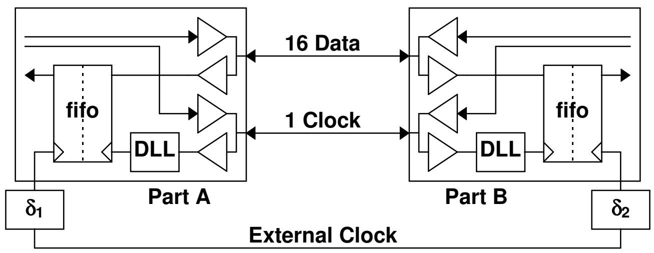

Figure 5: Phase-tolerant clocking between two components

III. Automatic Impedance Control

The technique to adjust the output driver impedance to match a representative board impedance using Time Domain Reflectometry (TDR) principles was described by T. Knight in [5]. Since then, other mechanisms to adjust the impedance of or the current in an output driver have been reported in [6] and [7].

A. Adjusting the Data and Clock Buffers

The automatic impedance control circuits implemented in this bidirectional signaling scheme form a closed loop system in which the board impedance is sampled by driving a voltage step into a series-terminated length of board trace. For the time equal to twice the electrical length of the sampled line, the voltage of the initial step into the line is proportional to a ratio of the buffer and board impedances. In a perfectly matched interface, the initial step into the line will be at $1/2V_{cc}$. This voltage level is compared to an on-chip generated $1/2V_{cc}$ reference, sampled and filtered, and then is used to either increase or decrease the buffer impedance. In the steady state or locked condition, the voltage level of the initial step into this sampled line will alternate above and below $1/2V_{cc}$. When the loop locks in this condition, the impedance of the Data and Clock output buffers is set to the lower of these two impedance values. Since the loop only attempts to correct for process and low frequency temperature and voltage variations, the loop operates at a reduced clock rate to relax internal timing requirements. A block diagram of the closed loop control system is shown in Figure 6.

The functionality of the automatic impedance control circuits are described with the help of the waveforms in Figure 7, where the waveform labels correspond to circuit nodes in Figure 6.

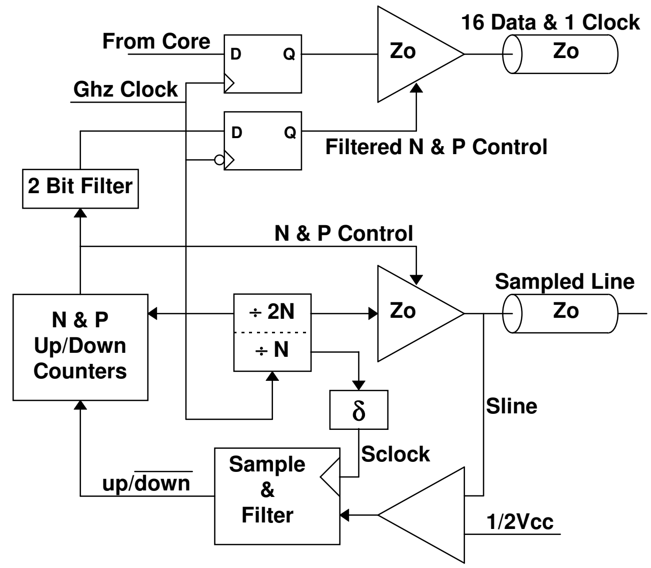

Figure 6: Automatic buffer impedance control block diagram

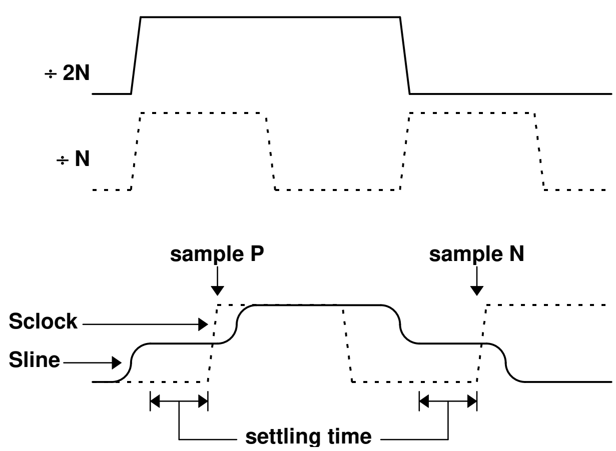

Figure 7: Automatic impedance control sampling waveforms

The core clock is divided by 2N and N to form the two clocks used by the impedance control loop. The 2N clock is driven on to the series-terminated board trace to form the signal Sline where the initial voltage step is proportional to the buffer to board impedance ratio. The step in a rising edge allows the P device to be sampled while the step in the falling edge allows the N device to be sampled. The N clock is delayed in time to form the sampling clock Sclock which samples the voltage level near the end of the step, just before the reflected wave arrives at the sampling latch input. This maximizes the time for any ringing on the line to settle and for the differential amplifier to resolve. The required settling time is the sum of the amplifier delay, the sampling latch setup time, and the time for ringing due to the package-to-board discontinuity to settle, and defines the minimum length of the sampled line. Both edges of the 2N clock are used to clock the state machine counters, while just the rising edge of N clock is used to clock the sampling latch and the digital filtering circuits.

This impedance control scheme can be implemented with just the one pin required for the sampled line. However, adding a second pin for the Sclock input and connecting the far end of the sampled line to the Sclock input would eliminate the need for N clock, but would also require a dual edge triggered sampling latch since both edges of the 2N clock would be used as Sclock.

B. Adjusting the Reference Buffers

The $1/4V_{cc}$ and $3/4V_{cc}$ references are automatically adjusted in a similar manner using another instance of the same closed loop circuit.

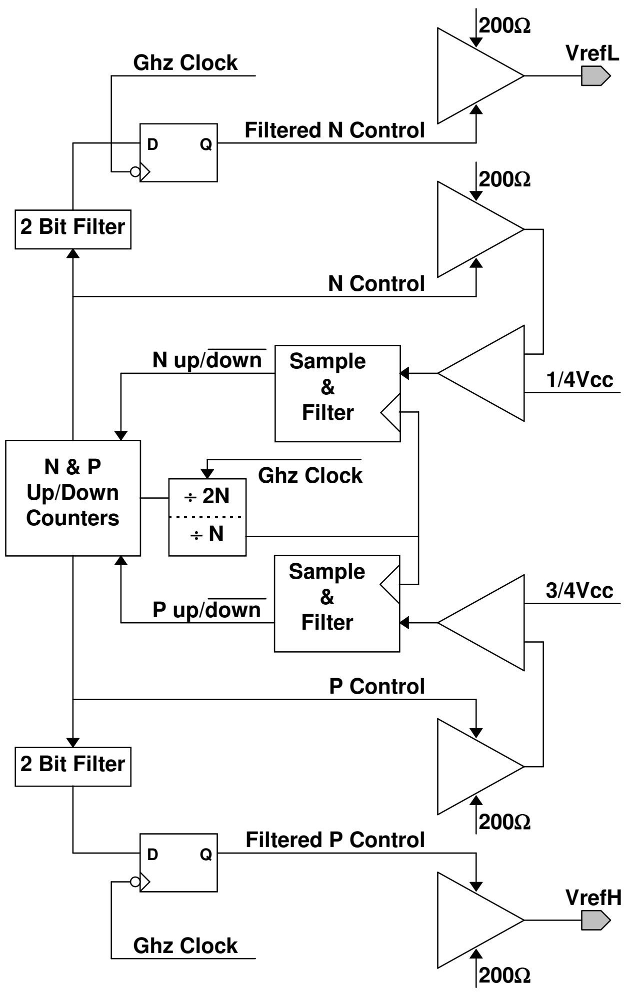

Figure 8: Automatic reference control block diagram

However, instead of matching the references to an external line or resistor, each reference is compared with and matched to an on chip resistive reference. The $1/4V_{cc}$ and $3/4V_{cc}$ references are built from the resistive divider formed by an output buffer with both the N and P devices on. Since one of the two devices is set to a relatively large impedance and is in saturation, and adjustment of this device yields coarse impedance changes, the large impedance device is not adjusted. Thus, only the upper or lower transistors of each reference buffer need be adjusted. For example, the $1/4V_{cc}$ reference is made by fixing the PMOS device to about 200Ω, and then adjusting the NMOS device until the reference matches $1/4V_{cc}$. This will occur when the NMOS device is about 67Ω, and the resulting resistive divider produces a voltage reference with a 50Ω Thevenin impedance to $1/4V_{cc}$. The $3/4V_{cc}$ reference is made in a similar manner by fixing the NMOS device and adjusting the PMOS device. A block diagram showing the way in which the references are adjusted is shown in Figure 8.

The buffers which are sampled and adjusted to match the on chip references are identical to the actual voltage reference buffers, but are not connected to external pins. Since only the upper or lower half of each reference is adjusted, one instance of the impedance control state machine is used to adjust both references.

The on chip $1/4V_{cc}$, $1/2V_{cc}$, and $3/4V_{cc}$ references used in the automatic impedance control circuits are built using n-well resistors and are capacitively filtered.

C. Impedance Update

One important design parameter in any digital automatic impedance control circuit is when to actually update the buffer impedance. On a unidirectional point-to-point bus or on a shared bus, the driver impedance can be updated when the buffer is off. Another option is to update the PMOS impedance when the buffer is in the low state and NMOS impedance when the buffer is in the high state. However, this places a data dependency in the impedance update circuits and adds complexity to the output buffer. To simplify the design of the output buffer, and to maximize the performance of the bidirectional interface, another update mechanism was used.

In the bidirectional signaling scheme the output buffer is always on, either in the high or low state, and new data is driven on to the link at each rising edge of the internal clock. To remove the data dependencies from the impedance update logic, the buffer impedance is updated dynamically on the falling edge of the internal clock. This places the impedance update at the center of the bit cell time to minimize any update induced jitter. The falling edge triggered update latches are illustrated in Figure 6 and Figure 8.

IV. RESULTS

The bidirectional signaling circuits described were implemented on a test chip fabricated on a 0.25µm CMOS process. The test chip contains two bidirectional data ports each with 16 data bits, 1 clock, and 2 references. Each data port contains a high speed logic analyzer unit capable of driving parallel data and comparing received data. The circuits associated with each data port are capable of operating up to 1.5GHz. The internal clock is generated by an on chip phase locked loop (PLL), which multiplies an external clock input by a factor of 16, and is distributed in a delay-matched tree network. The test chip also contains a set of six programmable noise generators to emulate the noise environment of a large microprocessor, and circuits for several additional experiments. The chip contains approximately 2.5 million transistors and is packaged in a plastic land grid array (PLGA) package. The component is operated at ambient case temperatures from a 1.8V supply.

The test chip is depicted in Figure 9 where the major blocks are highlighted. The two bidirectional signaling ports are located at the bottom edge of the chip, between which are located the PLL and the automatic impedance (AZ) control circuits.

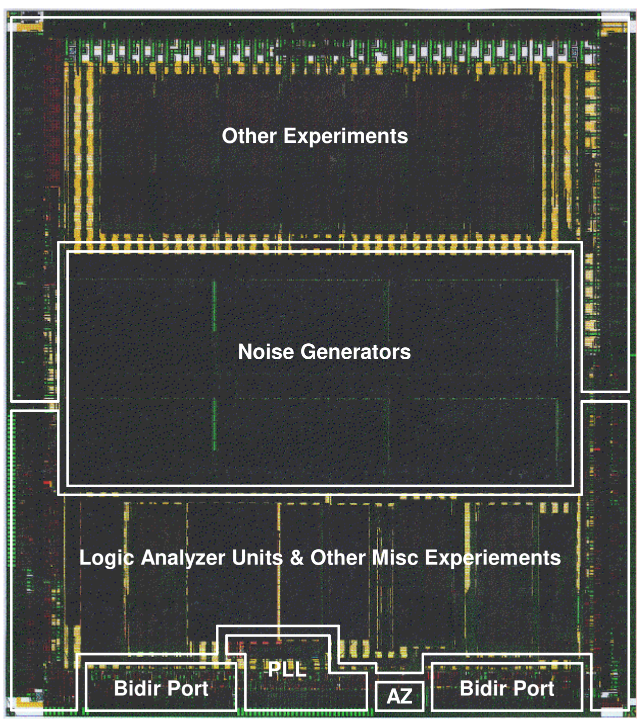

Figure 9: Test chip

A test board to evaluate the performance of the bidirectional signaling experiments over standard printed circuit board interconnects was constructed. The Data, Clock, and Reference signals of the bidirectional ports on two test chip components are connected in a point-to-point configuration across approximately four inches of 50Ω stripline placed on a 5/15 mil pitch. The test board, with two test chip components, is depicted in the photograph in Figure 10.

A scan access port was used to load software-generated random data patterns into the data and compare buffers in each test chip, and provided a way under software control to reset each test chip and to log transmission errors.

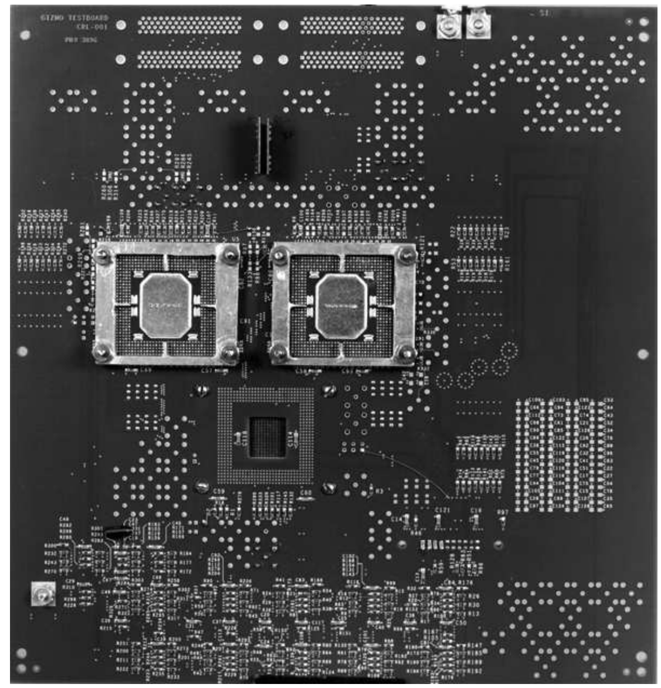

Figure 10: Bidirectional signaling test chips mounted on test board

A. Bidirectional Signaling Performance

The bidirectional ports were operated at 1.25GHz where data was transmitted simultaneously at 1.25Gb/s in each direction, which equates to 2.5Gb/s per wire. The bit error rate (BER or the ratio of the number of bits in error to the number of bits transmitted) measured with the link operating at 2.5Gb/s per wire was < 4.5 x 10^{-13}. An eye diagram showing the link operating at this rate is illustrated in Figure 11.

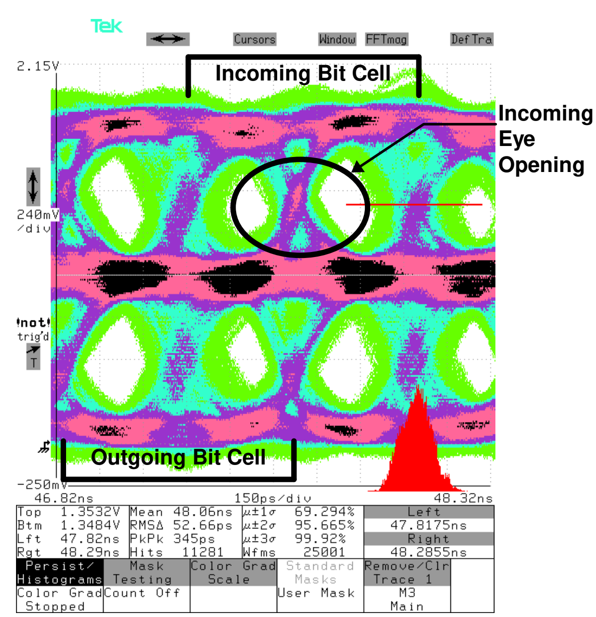

Figure 11: Bidirectional eye diagram at 2.5Gb/s per wire

At this rate, as shown in the eye diagram, a new incoming bit cell and a new outgoing bit cell appear on the wire every 800ps. The phase between the incoming and outgoing bit cells is deliberately set such that the outgoing edges occur in the center of the incoming eye to make observing the separate eyes easier. The eye opening in the incoming bit cell is highlighted. The measured peak-to-peak jitter on the incoming data edge is 345ps and is due primarily to the impedance discontinuity between the PLGA package and the board.

In this source synchronous signaling scheme, the phase between the clocks of the two communicating components is completely arbitrary, and it is important to characterize the link performance as a function of the phase shift between the two components. The timing and voltage noise margins at the receiver, as a function of the phase between the incoming and outgoing bit cells, were measured. For these measurements, the timing window or the voltage noise margin was violated when a single error occurred in transmission of 40GBytes containing approximately 5000 random 32 bit data patterns. Thus, failure was defined to occur when the BER exceeded 2.5 x 10^{-11}.

The timing window was measured by shifting the transmitted clock with respect to the data using impedance-matched external delay lines, and the voltage noise margins were measured by deliberately shifting the $1/4V_{cc}$ and $3/4V_{cc}$ reference levels. Since the resolution of the timing margin measurements was only 100ps, and a performance degradation of 10% at the worst case phase was expected, the timing and noise margins were measured with the link operating at 2.0Gb/s per wire (1.0Gb/s per direction or 1ns bit cells on the wire) to provide a more exact measure of the margins.

With the bidirectional link operating at 2.0Gb/s per wire, a worst case timing margin of 300ps was measured. This means that the transmitted clock can be placed anywhere within a 300ps window to correctly capture the incoming data. Several of the individual terms which erode the 1ns bit cell were characterized and are summarized in Table 1.

Table 1: Measured terms which erode timing margin on bidirectional link operating at 1.0Gb/s per direction or 2.0Gb/s per wire.

1000ps	Total bit cell time
-275ps	Peak-peak pattern dependent jitter
-125ps	Static skew (from package)
-75ps	Duty cycle variation
-50ps	Transmit clock jitter
-175ps	Input jitter & measurement error
300ps	**Total timing margin**

The peak-to-peak pattern dependent jitter of 275ps was measured from an eye diagram similar to that shown in Figure 11, but with the link operating at 2.0Gb/s per wire. The 125ps static skew is due to variations in the signal trace lengths in the PLGA package and to on-die process variations, and could be reduced somewhat by length compensation on the board. The duty cycle variation is due both to the duty cycle of the internal clock and to the nonsymmetric characteristics of the output buffer driving a parallel termination. The transmit clock jitter is the peak-to-peak cycle-to-cycle jitter on the bidirectionally connected transmit clock. The remaining 175ps is attributed to jitter from the input amplifier, which was not measured, and to the uncertainty due to the measurement resolution.

At 2.0Gb/s per wire, the worst case measured voltage noise margins about the $1/4V_{cc}$ and $3/4V_{cc}$ references is approximately 200mV, where the measurement resolution is 10-40mV depending on the reference levels. The noise margins increase to 350mV when only a single output is switched, primarily due to the elimination of the static skew across the port and to a reduction in the cross talk effects.

At 2.0Gb/s per wire, with the transmit clocks and references optimally placed, the measured BER was < 10^{-15}.

B. Comparing Bidirectional and Unidirectional Results

The bidirectional signaling circuits operate in a source-parallel terminated unidirectional manner in the special case when a constant '1' or '0' is transmitted in one direction while random data is transmitted in the other direction. The link operates in this unidirectional manner up to 1.5Gb/s per wire and is limited to this rate by the on-chip PLL. An eye diagram showing the link operating unidirectionally at 1.5Gb/s, with a parallel termination to V_{ss}, is shown in Figure 12. The measured peak-to-peak jitter is 219ps and the BER < 1.5 x 10^{-12}.

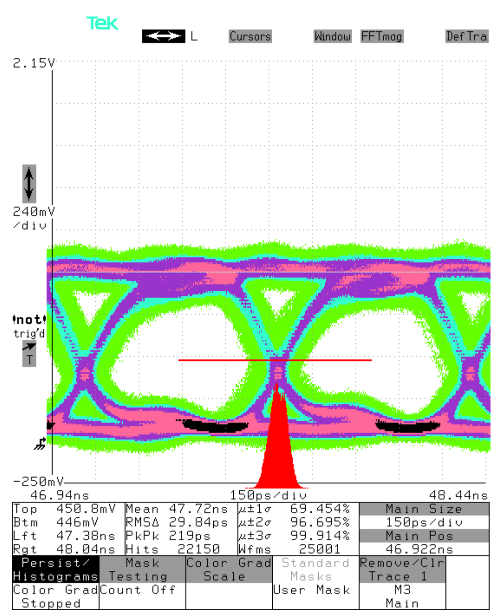

Figure 12: Unidirectional eye diagram at 1.5Gb/s.

When the bidirectional and unidirectional links are compared on an equivalent *bandwidth per wire* basis at 1.5Gb/s per wire, substantially larger timing and voltage noise margins exist on a bidirectional link operating at 0.75Gb/s per direction (1.5Gb/s per wire) than on a unidirectional link operating at 1.5Gb/s. Operating the link bidirectionally also relaxes the bandwidth requirements on the input and output circuits since the bit cells are 2x larger. There are also system protocol advantages of a full duplex link since no turn around cycles are required. Transmission occurs simultaneously and continuously in each direction.

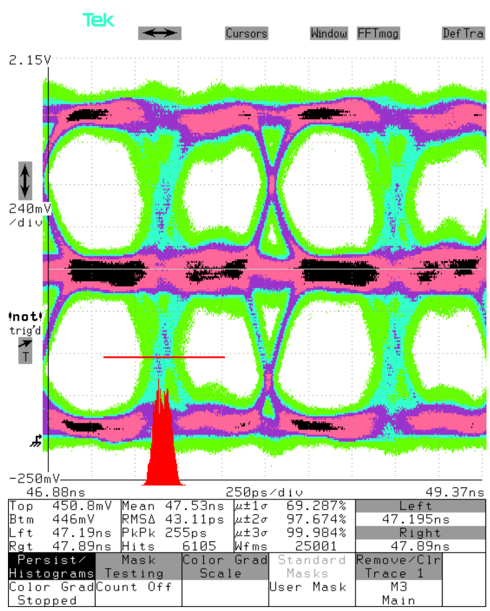

Figure 13: Bidirectional eye diagram at 0.75Gb/s per direction or 1.5Gb/s per wire.

To compare the bidirectional and unidirectional results at 1.5Gb/s per wire, an eye diagram of the link operating bidirectionally at 0.75Gb/s per direction (1.5Gb/s per wire) is shown in Figure 13. When comparing Figure 13 with Figure 12, note that the horizontal time base (ps/div) in each figure is different, but that the vertical deflections (mV/div) are the same.

While the unidirectional link certainly operates at 1.5Gb/s, substantially more timing and voltage noise margins exist on the bidirectional link operating at the equivalent bandwidth per wire of 0.75Gb/s per direction. The measured timing and voltage margins under these two conditions are summarized in Table 2.

Table 2: Comparison of bidirectional and unidirectional margins at the equivalent bandwidth per wire of 1.5Gb/s.

Mode and Rate/Dir	Timing Margin	Noise Margin
Bidir at 0.75Gb/s	700ps	425mV
Unidir at 1.5Gb/s	200ps	125mV

When the two signaling methods are compared on a *bandwidth per direction* basis, the margins on the bidirectional link are lower than those on the unidirectional link. However, when compared on a *bandwidth per wire* basis, the margins on the bidirectional link are more than 3x those on the unidirectional link since the bidirectional bit cells are twice the length of the unidirectional bit cells.

C. Automatic Impedance Control

Correct functionality of the automatic impedance control loop is illustrated in Figure 14. The top waveform, made with the automatic impedance control circuits disabled, shows the buffer undersized with respect to the board impedance. In this case, two full round trips of the initial voltage step into the sampled line are required for the voltage to reach the V_{cc} or V_{ss} rail. The bottom waveform, made with the automatic impedance control circuits enabled, shows the buffer impedance has automatically been adjusted to match the board impedance. In this case, the initial voltage step into the line occurs at $1/2V_{cc}$ and after one round trip delay, the voltage reaches either the V_{cc} or V_{ss} rail.

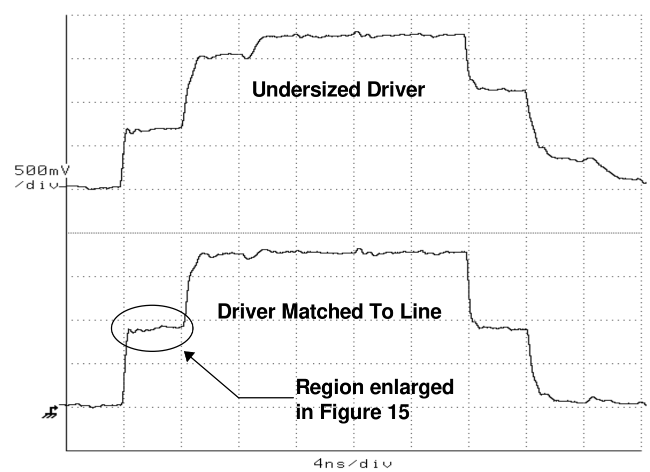

Figure 14: Undersized driver impedance automatically adjusted to match the line impedance.

In the steady state condition when the impedance control loop is locked, the state machine will set the buffer impedance such that the initial voltage step into the line will alternate between two binary steps above and below $1/2V_{cc}$. This is indicated in Figure 14 and enlarged in Figure 15, which is a persistence diagram showing the voltage levels that occurred due to a rising edge at the near end of the sampled line, with the control loop locked, over some elapsed time. The ringing observed just after the initial voltage step is caused by the impedance discontinuity between the PLGA package trace and the sampled line on the board. (Note that another locked condition can occur when a binary step produces an initial voltage step right at $1/2V_{cc}$. In this case, the differential amplifier fails to resolve before the reflected wave arrives at the sampling latch input, causing the state machine to alternate between three binary steps.)

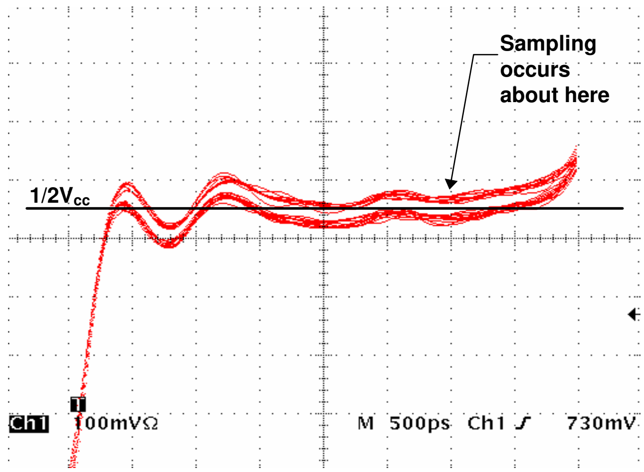

Figure 15: Persistence diagram showing impedance changes to produce an initial voltage level one binary step above and below 1/2Vcc.

To characterize the effects of updating the buffer impedance in the center of the bit cell, provision was made to vary the update rate with respect to the transmitted data rate and to measure the pattern dependent jitter as a function of the update rate. When operating the link bidirectionally at 2.0Gb/s per wire, the pattern dependent jitter was measured with update intervals of 16ns, 16µs, 16ms, and 16s. No measurable differences in the peak-to-peak jitter, measured at these impedance update intervals, was detected.

V. Summary And Conclusions

A bidirectional point-to-point signaling scheme, in which data is transmitted simultaneously in each direction over one wire, has been scaled from a 0.6µm CMOS process to a 0.25µm CMOS process and is capable of transmitting and receiving data at 2.5Gb/s per wire. The link also operates in a source-parallel terminated unidirectional scheme up to 1.5Gb/s. A method, based on TDR techniques, to automatically sample the board impedance and adjust the output buffer impedance to match has been devised and successfully used. The performance of the link operating bidirectionally is compared with the performance of the link operating unidirectionally. On a *bandwidth per wire* basis at 1.5Gb/s, the bidirectional timing and voltage noise margins are about three times the margins observed on a unidirectional link. The bidirectional scheme also offers several circuit and system level advantages.

Point-to-point buses allow inter-chip communication rates to approach the maximum rate a given process technology will deliver. The bidirectional point-to-point signaling scheme described here offers the possibility of extending communication bandwidths, on a per wire basis, beyond what the process technology can deliver to a unidirectional bus.

VI. Acknowledgments

The authors would like to thank Greg Dermer for designing the test board used to characterize the bidirectional signaling method, and for many helpful discussions about high frequency test and characterization. Brad Bloechel provided much help in the lab, without which the testing would not have been possible. We also acknowledge the Portland Technology Development and Design Technology groups at Intel for their help throughout this program.

VII. References

[1] R. Mooney et al., “A 900 Mb/s Bidirectional Signaling Scheme,” IEEE Journal of Solid-State Circuits, vol 30, no 12, pp.1538-1543, December 1995.

[2] J. Carbonaro et al., “Cavallino: The Teraflops Router and NIC,” Hot Interconnects IV Symposium Record, pp. 157-160, August 1996.

[3] M. Bazes, “Two Novel Fully Complementary Self Biased CMOS Differential Amplifiers,” IEEE Journal of Solid-State Circuits, vol. 26, pp. 165-168, February 1991.

[4] E. Reese et al., “A Phase-Tolerant 3.8 GB/s Data-Communication Router For A Supercomputer Backplane,” ISSCC Digest of Technical Papers, Paper FA18.4, pp. 296-297, February 1994.

[5] T. Knight et al., “A Self-Terminating Low-Voltage Swing CMOS Output Driver,” IEEE Journal of Solid-State Circuits, vol 23, no 2, pp. 457-464, April 1988.

[6] A DeHon et al., “Automatic Impedance Control,” ISSCC Digest of Technical Papers, Paper 10.7, pp. 164-165, February 1993.

[7] K. Donnelly et al., “A 660MB/s Interface Megacell Portable Circuit in 0.3um = 0.7um CMOS ASIC,” ISSCC Digest of Technical Papers, Paper 17.4, pp. 290-291, February 1996.