Distributed Loss-Compensation Techniques for Energy-Efficient Low-Latency On-Chip Communication

Anup P. Jose, *Member, IEEE*, and Kenneth L. Shepard, *Senior Member, IEEE*

***Abstract*—In this paper, we describe the use of distributed loss compensation to provide nearly transmission-line behavior for long on-chip interconnects. Negative impedance converters (NICs) inserted at regular intervals along an on-chip line are shown to reduce losses from more than 1 dB/mm to less than 0.3 dB/mm at 10 GHz. Results are presented for a 14-mm 3-Gb/s on-chip double-data-rate (DDR) link in 0.18-μm CMOS technology, with a measured latency of 12.1 ps/mm and an energy consumption of less than 2 pJ/b with a BER $< 10^{-14}$. This constitutes a factor-of-three improvement in power and a one-and-a-half-times improvement in latency over an optimally repeated *RC* line of the same wire width.**

***Index Terms*—Interconnections, on-chip networks, transmission lines.**

I. Introduction

OVER the past few decades, improvements in integrated circuit density and performance have been achieved by scaling down transistors. Latency (rather than pipelined throughput) of on-chip wires is important for many applications such as buses between cache memories and processors [1]. Although the latencies of local interconnects scale accordingly, the delay per unit length of on-chip wires, as determined by a diffusive *RC*-limited response and as measured relative to gate delays, approximately doubles every technology generation [2], [3] as wire resistances per unit length increase and gate delays decrease with scaling. Furthermore, these wire delays (D) grow quadratically with wire length, $D \propto R_{\rm wire}C_{\rm wire}L^2$. Wire bandwidths, which are inversely proportional to D, also degrade. Buffers (or repeaters) are traditionally added to make the interconnect latency linear with wire length, with a simple relationship guiding an optimal number of repeaters (and their sizing) to minimize interconnect delay [2]. Wide wires can be used to improve overall latency, requiring fewer numbers of repeaters of larger area (to drive the larger wire capacitance) to achieve a delay-optimal solution. Overall energy per bit and routing density, however, degrade with wire widening [2]. Despite continuing tradeoffs, fundamental improvements in interconnect latency, bandwidth density (bits per second per unit routing width), and energy per bit for on-chip wires can be achieved with alternatives to full-rail *RC*-limited interconnect buffered by CMOS inverters.

Reference [4] considers the use of sharp current pulses (driver preemphasis) to emphasize the inductance-dominated region of on-chip interconnects, thereby achieving latencies close to the speed of light. Because there is no loss compensation, wide wires $(> 3\ \mu{\rm m})$ are required for modest lengths ($\cong$3 mm). Although both energy per bit and latency are improved over optimally repeated *RC* lines, this approach is not easily scalable to long interconnect lengths without the use of excessively wide wires.

This study investigates the use of negative impedance converters (NICs) to provide shunt transmission-line loss compensation, allowing transmission-line behavior for long on-chip wires [5]. Distributed loss compensation, which was originally employed for long-distance telephony [6], has found recent application in both distributed amplifiers [7] and distributed oscillators [8] in CMOS technology. Bipolar current-mode signaling, which is similar to that used for low-voltage differential swing (LVDS) links [9], is employed. Offset-compensated sense-amplifier detection at the receiver further reduces the far-end voltage swing required for error-free operation of the link.

Link latency is determined by the phase velocity in the compensated transmission line. The bandwidth improvement (through effective reduction in transmission-line α) comes at the cost of some degradation in wire latency due to the added capacitance of the NICs. In Fig. 1, we compare (through circuit simulation with *RLC* models for the associated wire segments) the energy per bit and link latency for two 14-mm links of different wire widths (the Al interconnect offers a sheet resistance of 0.108 $\Omega/{\rm square}$, typical of a 0.18-μm technology) operating with a throughput of 3 Gb/s. By comparing data at the same wire width, latency and energy per bit can be examined for the two design styles at the same bandwidth density. One link operates as a conventional single-ended optimally repeated full-rail (1.8-V) link, in which both repeater size and number of repeaters are optimized to minimize latency [2].[1] The second link operates as a compensated differential transmission line

Manuscript received August 21, 2006; revised December 23, 2006. This work was supported in part by the MARCO/DARPA C2S2 Center, IBM, and Intel Corporation.

A. P. Jose was with the Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA. He is currently with AMD's Boston Design Center, Boxborough, MA 01719 USA (e-mail: anup.jose@amd.com).

K. L. Shepard is with the Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: shepard@cisl.columbia.edu).

Digital Object Identifier 10.1109/JSSC.2007.897165

[1]Repeater insertion is done under the constraint that repeaters are spaced equally along the wire and are all the same size.

0018-9200/$25.00 © 2007 IEEE

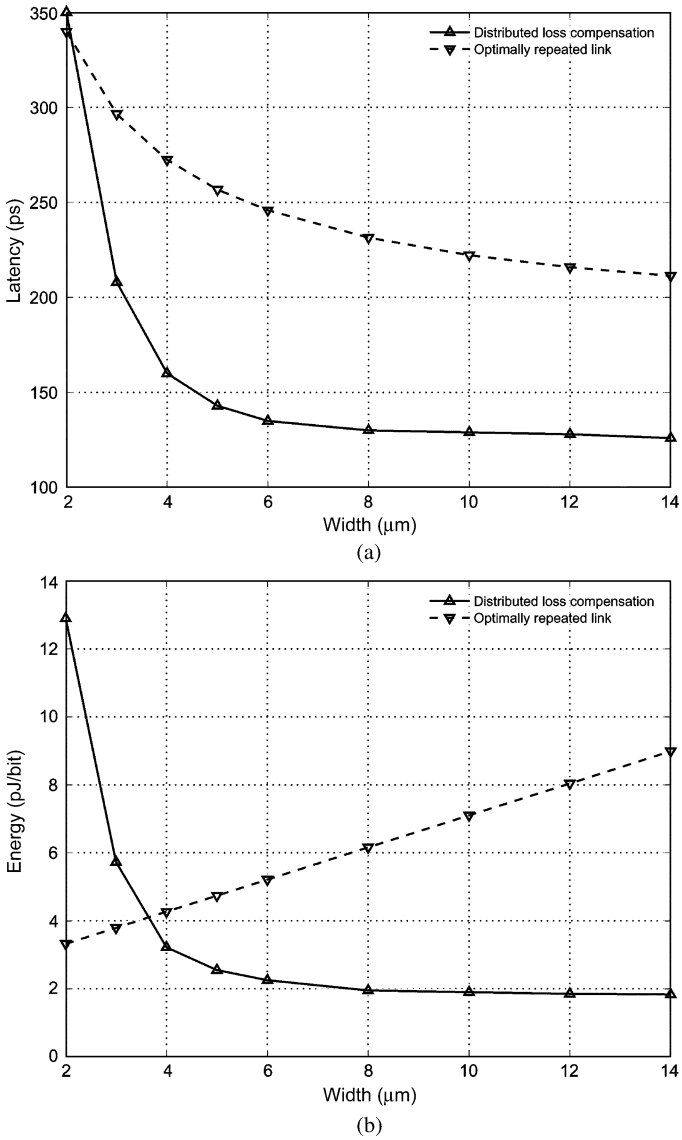

Fig. 1. (a) Wire latency and (b) bit energy (pJ/bit) as a function of wire width for two 14-mm links operating at 3 Gb/s—one with repeaters and one with distributed loss compensation.

operating with a fixed far-end swing of 70 mV. The sizing and biasing of the NICs are assumed to be constant (the sizing and biasing matches the implementation of Section III) with enough NICs distributed along the line to provide an attentuation constant (α) of ≈0.2 dB/mm at 1 GHz. The resulting latency, shown in Fig. 1(a), is significantly lower for the loss-compensated transmission line than for the optimally repeated line for wire widths greater than 2 μm. In both cases, the latencies decrease asymptotically with increasing widths. For an optimally repeated line, the delay of each repeated line segment is independent of the wire width, with the number of segments decreasing with increasing width [2]. For very large widths, the line degenerates into a large capacitance driven by a large driver with a negligible wire *RC* delay component. This "silicon" delay determines the asymptote of the solid curve in Fig. 1(a) for wide wires. For the loss-compensated transmission-line case, the resulting number of NICs decreases with increasing width. For very small wire widths, the capacitance introduced by the large number of NICs needed to

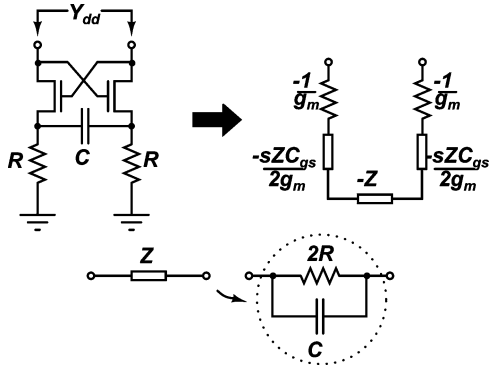

Fig. 2. Schematic of the NIC architecture used in this implementation.

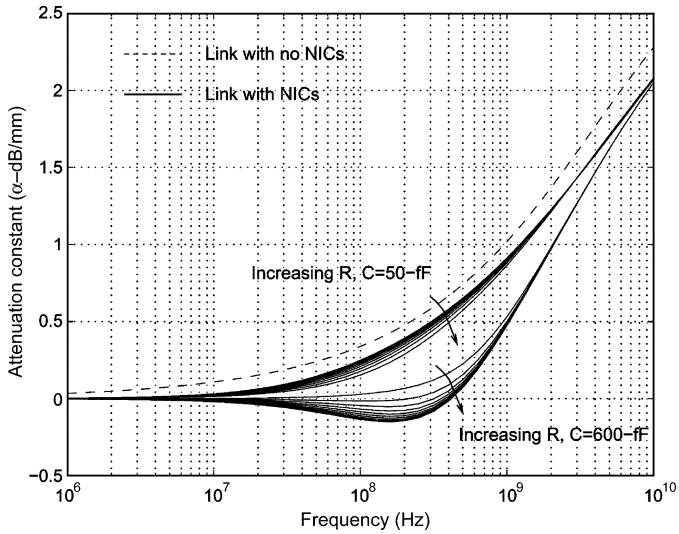

Fig. 3. Attenuation constant (α) of the line.

cancel out interconnect losses leads to increasing wire delays. For large wire widths, the NIC link offers a smaller latency when compared to an optimally repeated link, which is a direct consequence of the (asymptotically reducing number of) NICs. Fig. 1(b) compares the energy per bit for both links. The linear increase in energy per bit for an optimally repeated wire can be attributed to the linear relationship between energy and wire capacitance $(E \propto CV^2)$. This curve assumes a worst case data pattern of alternating ones and zeros with a 3-GHz data rate. The loss-compensated transmission-line case shows a decrease in energy consumption for increasing width as the number of NICs (and their static power consumption) are reduced.

In Section II, we consider the design issues associated with distributed loss compensation, including appropriate sizing and placement of the NICs to ensure stability. Section III considers the overall design of the prototype 3-Gb/s link, and Section IV presents measurement results. Section V concludes.

II. Negative Impedance Converters

The NIC element employed in this design is shown in Fig. 2 and consists of two cross-coupled transistors with a source-degeneration network. Fig. 2 also shows the equivalent small-signal model, in which the transistors are characterized by an input capacitance $C_{\rm gs}$ and transconductance g_m. Ignoring

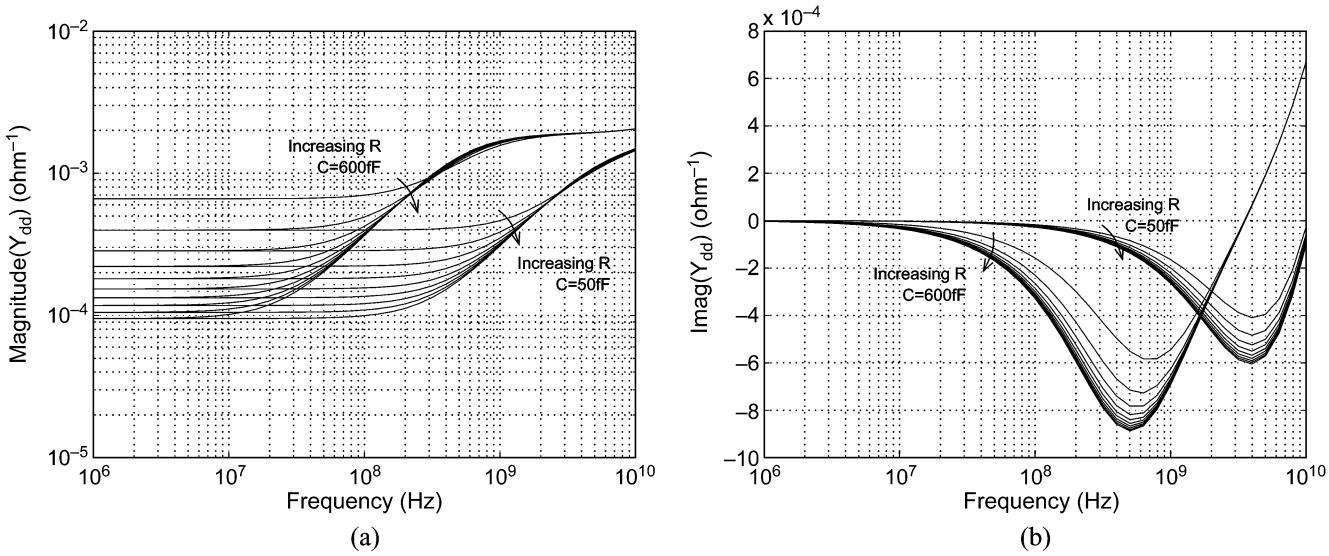

Fig. 4. (a) Magnitude of $Y_{\rm dd}$. (b) Imaginary part of $Y_{\rm dd}$ for various values of R and C.

gate-to-drain overlap capacitance, the differential admittance of the NIC, $Y_{\rm dd}$, is given by

$$Y_{\rm dd} = \frac{-g_m/(2R)}{g_m + 1/R}\left[\frac{1+s2RC}{1+\frac{s(C_{\rm gs}+2C)}{g_m+1/R}}\right]. \tag{1}$$

The resistors degenerate the differential admittance $Y_{\rm dd}$ of the NIC at low frequencies, while, as frequencies increase, the capacitor acts to shunt this degeneration and increase $Y_{\rm dd}$, providing the admittance of a "negative" capacitance [10]. This allows the loss compensation to match the frequency response of α, delivering more compensation at higher frequencies and helping ensure stability for the compensated line (as described below). In the absence of the degeneration network ($R = C = 0$), $Y_{\rm dd}$ reduces to the well-known $-g_m/2$.

For $R \gg 1/g_m$ and $C \gg C_{\rm gs}$, (1) approximates to

$$Y_{\rm dd} = -\frac{1}{2R}\left[\frac{1+s2RC}{1+s2C/g_m}\right] \tag{2}$$

with a pole and zero at $g_m/2C$ and $(1/2RC)$, respectively. For values that closely match the testchip implementation ($g_m = 4$ mS, $R = 1$ kΩ ($\gg 1/g_m$) and $C = 600$ fF), the zero at $(1/2RC)$ ($\approx$132 MHz) gives a negative admittance that increases with increasing frequency until the pole is reached at approximately $g_m/2C$ ($\approx$660 MHz).

Overcompensation of the transmission-line losses can result in instability, leading to excessive overshoot, oscillations, or "latch-up" of the transmission line. For the doubly terminated transmission lines considered here (in which $S_{11} = S_{22}$ and $S_{12} = S_{21}$), unconditional stability requires that both of the following conditions be satisfied for the S-parameters of the compensated line [11]:

$$k = \frac{1 + \left|S_{11}^2 - S_{21}^2\right|^2 - 2|S_{11}|^2}{2|S_{21}|^2} > 1 \tag{3}$$

$$\left|S_{21}^2\right| < 1 - |S_{11}|^2 \tag{4}$$

where S_{11} is given by $(Z_L - Z_0)/(Z_L + Z_0)$. For matched termination at either ends of the interconnect, $S_{11} \cong 0$ and (4) simplifies to

$$\left|S_{21}^2\right| = e^{-2\alpha l} < 1 \tag{5}$$

for a transmission line of length l. Unconditional stability therefore requires $\alpha > 0$ for all frequencies.

Stability depends on the appropriate choice of g_m, R, and C. The two groups of plots in Fig. 3 shows α for various values of R for two different values of C—50 fF and 600 fF. In all cases, $g_m \approx 4$ mS. For comparison, α for the uncompensated interconnect is also shown. Increasing the value of C clearly enhances the compensation at higher frequencies (20 MHz to $\approx$3 GHz) but also increases the risk of the system becoming unstable ($\alpha < 0$). Higher values of R for $C = 600$ fF make the line unstable. Fig. 4(a) shows $|Y_{\rm dd}|$, which has a zero at $1/(2RC)$ and a pole at $\approx g_m/(2C)$. There is also a right-half-plane zero associated with the device $f_T (= g_m/C_{\rm gs})$ at frequencies >10 GHz. Fig. 4(b) shows the imaginary part of $Y_{\rm dd}$, which is negative for low frequencies (implying a negative capacitance). Increasing C to enhance the compensation leads to a lower crossover (negative to positive) frequency for the imaginary part, making this compensation less effective at high frequencies.

III. System Implementation

A prototype double-data-rate (DDR) 3-Gb/s 14-mm link (on the fifth metal level of a six-level metal process) as designed and fabricated in a 0.18-μm CMOS technology is shown schematically in Fig. 5. The interconnect has a coplanar waveguide (CPW) topology with a line-width and spacing of 8 μm with seven NICs evenly spaced along the line.[2] The resistors inside the NICs are implemented with nonsilicided polysilicon resistors, while the capacitors are implemented with pFETs. Across process corners, these elements have a variance of approximately 10%. Differential operation assures controlled inductance, high common-mode noise rejection, and reduced

[2]The time-of-flight between NICs is approximately 20 ps. With slew rates of more than 75 ps, the NICs appear as a distributed load on the transmission line.

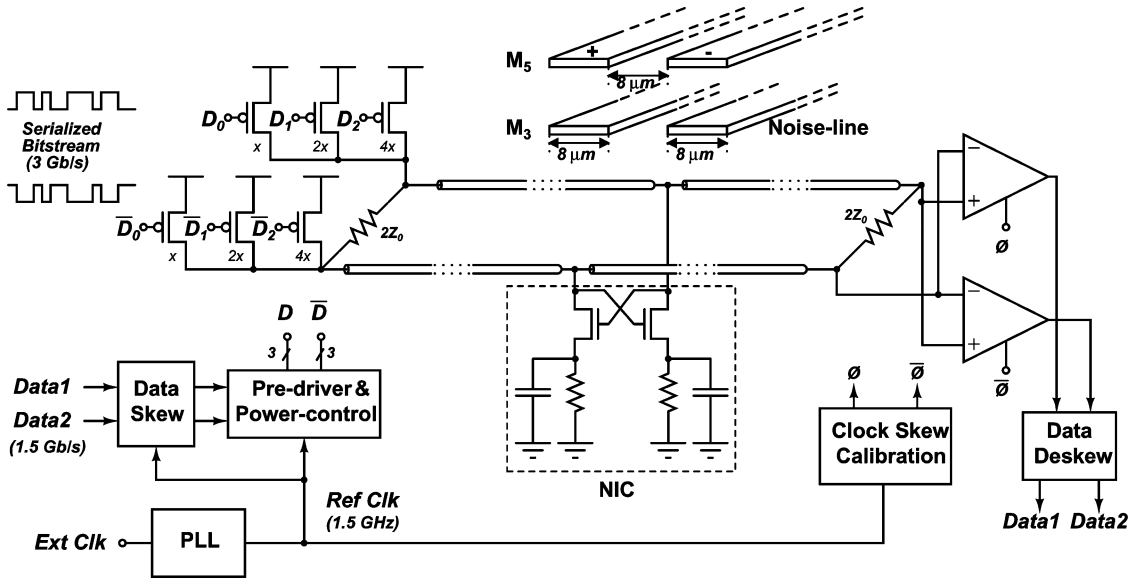

Fig. 5. Overall system architecture of the 3-Gb/s on-chip link prototype.

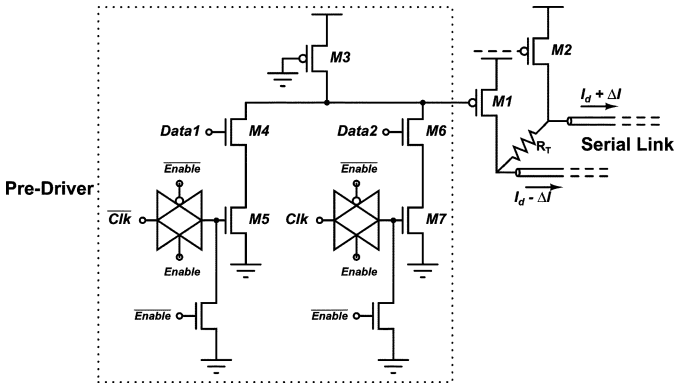

Fig. 6. Driver design.

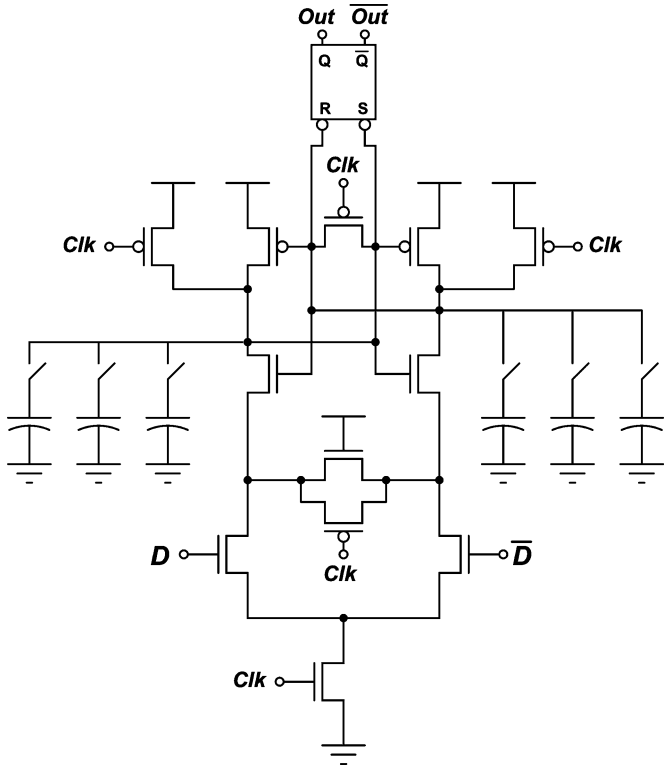

Fig. 7. StrongARM latch with offset calibration capacitors.

shielding requirements.[3] The common-mode voltage at the driver sets the operating point for the NICs. Parallel metal wires run under the main link (on third-level metal) to provide more realistic interconnect loading on the line. The driver and receiver components are assumed to run at the same frequency (mesochronously) although arbitrary skews are accommodated with an automated calibration at startup. The link is doubly terminated with n-type diffusion resistors of value $2Z_0$.

The main components of this on-chip interconnection prototype are described in separate sections below, including the driver and receiver circuits, the link calibration circuits, and the on-chip test circuits for link characterization. The 0.18-μm CMOS technology used in this implementation has a fanout-of-four (FO4) delay (τ_4) of approximately 60 ps.

A. Driver Design

A current-mode driver with pseudo-nMOS predriver, shown in Fig. 6, with input multiplexing to achieve DDR operation is employed. DDR operation, with the modest energy overhead of skewing and de-skewing latches, helps to compensate for the reduction in bandwidth per unit wiring area (bandwidth density) associated with wider wires. The driver/predriver combination multiplexes two bitstreams—*Data1* and *Data2* each having a throughput of 1.5 Gb/s into a single 3-Gb/s bit-stream.

The predriver stage uses ratioed logic to reduce the circuit complexity when compared to static CMOS circuits resulting in a lower area overhead. Transistors M3–M7 form a pseudo-nMOS gate with M3 acting as the pull-up load. Either

[3]In a more realistic interconnect environment with adjacent aggressor lines, these differential lines would be twisted to help render crosstalk noise common-mode.

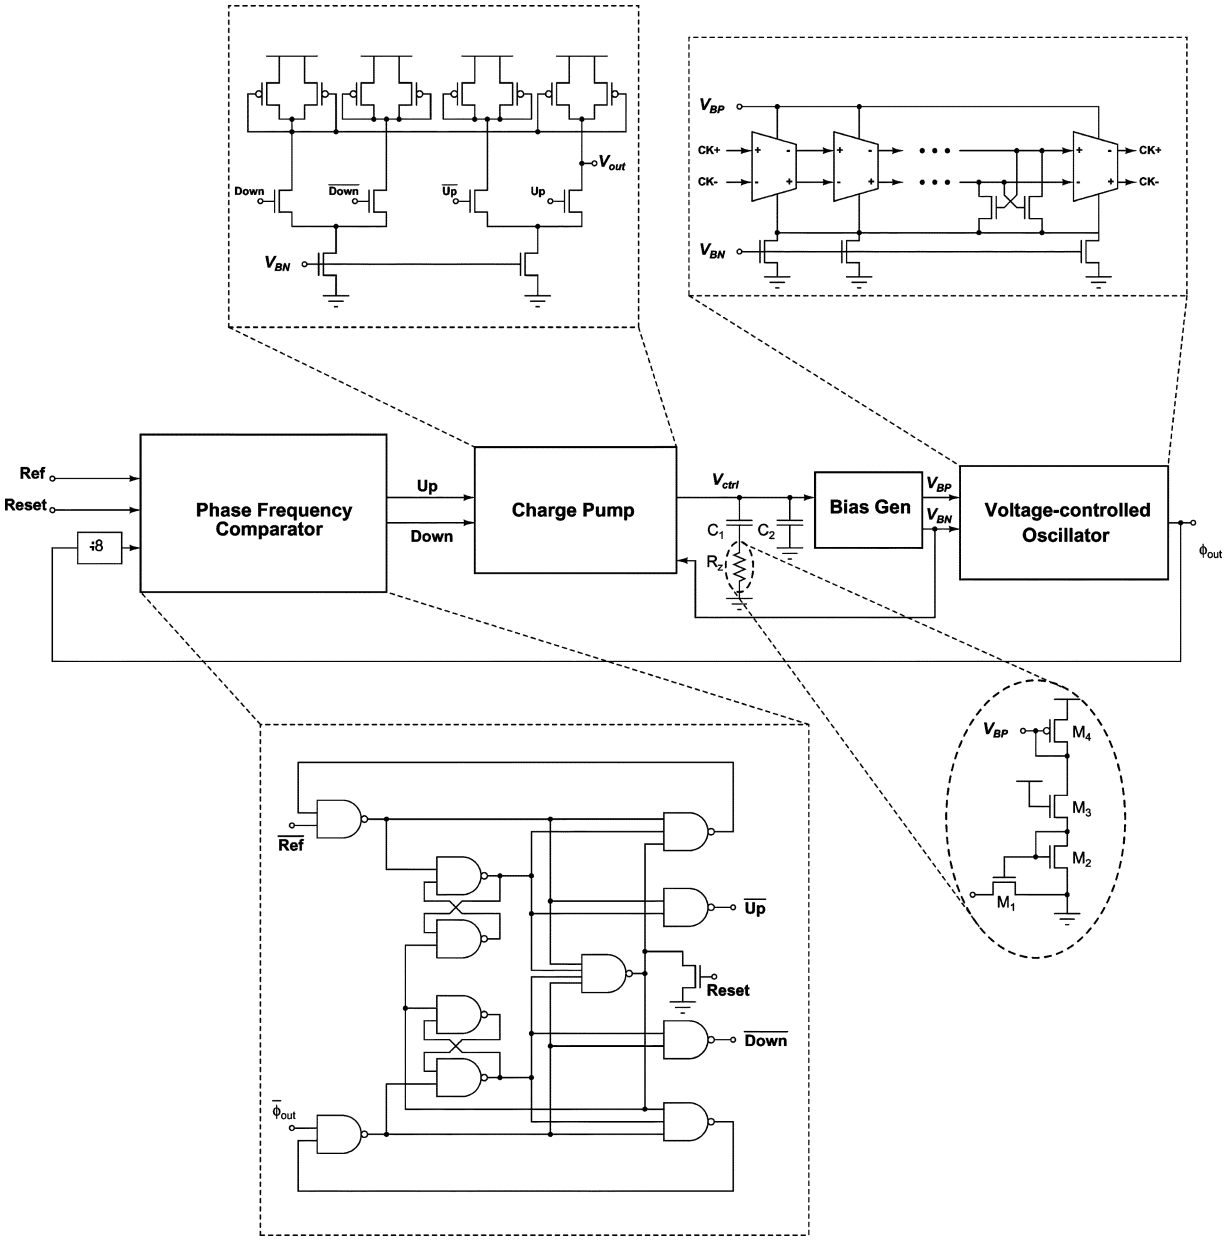

Fig. 8. PLL schematic.

of the two branches of the pull-down network (M4–M7) are turned on in each half cycle depending on the input data. The driver consists of two pFETs M1-2 along with a termination resistor R_T to minimize the effect of reflections as well as crosstalk noise. The driver transistors are pFETs, allowing nFETs to be used for the cross-coupled transistor pair in the NICs. A larger driver transistor size is required than would be necessary if nFETs were used, but this is more than compensated by the smaller NIC devices for a given g_m that are enabled by using nFETs rather than pFETs. Transistor M3 is sized relative to the pull-down transistors to keep M1–M2 in saturation. One of transistors M1–M2 is sourcing current at any instant, resulting in a steady-state common-mode current I_d, upon which a bipolar differential signal current (ΔI) is superimposed. The total current of $2I_d$, drawn from the power supply during normal operation, is obtained when either M6–M7 or M4–M5 of either pre-driver is switched on—M1 and M2 cannot be both turned on at any given instant, except in the offset calibration mode discussed in Section III-D. The ratio $\Delta I/I_d$ is given by $R_T/(R_T + 2Z_l)$, where Z_l is the impedance looking into each half of the line. The value of R_T has to be chosen carefully so as to achieve the right compromise between reflection and far-end voltage swing—maximizing R_T maximizes the near-end voltage-swing, but this comes at the expense of increased reflection at higher frequencies. Because of resistive losses in the line, the common-mode voltage on the wire and the associated bias currents of the NICs decrease toward the far-end of the line, the devices of which are sized larger to provide uniform g_m. There are multiple copies of the predriver/driver with varying sizes to dynamically control the drive current $(2I_d)$ from 3.0 to 6.0 mA in steps of 0.35 mA. Larger driver currents boost signaling levels, increasing the

magnitude of the far-end swing, allowing crosstalk noise to be overcome dynamically as necessary. Larger driver currents also increase the g_m of the devices in the NIC, increasing NIC bandwidth. Adjustment of drive current also allows the NICs to be trimmed to compensate for variability of process parameters and ensure stability of the compensated interconnect.

B. Receiver Design

The receivers, shown in Fig. 7, are StrongARM gate-isolated sense-amplifier latches [12]. For a clock slew time of 75 ps, this latch provides a typical aperture time of 15 ps. A digitally trimmed capacitive load is used for input offset cancellation which is typically on the order of a few tens of millivolts. Increasing the size of the transistors to lower this offset voltage significantly degrades the overall performance of the receiver and increases the loading at the far end of the interconnect. This required scaling factor can be as large as 50, as shown in [13]. Positioning the trimming capacitors at the output of the latch offers better offset control for smaller capacitance (and switch) sizing over adding these at the drains of the differential input pair [13].

A 320-Ω n-type diffusion resistor is used for line termination at the receiver.[4] This is slightly larger than the (high-frequency) $2Z_0$ of the line ($Z_0 \approx 150\ \Omega$) to boost far-end voltage swing while not creating an impedance discontinuity large enough to produce significant reflection at the far end. With the multiplexed detection, far-end capacitive loading is approximately 9 fF on each leg of the differential link. The resulting Z_0C_L time constant (2 ps) is sufficiently low that it does not introduce significant ISI.

C. Phase-Locked Loop (PLL) Design

As shown in Fig. 8, the design includes a PLL for on-chip clock multiplication. The design largely follows that of Maneatis [14] and is shown in Fig. 8. The VCO used in this implementation [15] is a modified version of the one used in [14]. The significant difference between these two schemes is the shorting of the current source drains ($V_{\rm tail}$) to reduce the variation in the tail current (due to varying $V_{\rm DS}$), further reducing power-supply-induced jitter. With the tail nodes tied together, the VCO becomes two single-ended rings. To rectify this, a cross-coupled transistor pair is used prior to the last stage as shown in the figure.

The loop filter R_z is implemented with an FET biased in the triode region, as shown in the inset of Fig. 8 [16]. For transistor M1, $R_{\rm on}^{-1} = \mu C_{\rm ox}(W/L)(V_{\rm GS} - V_{\rm TH})$, which is equal to the transconductance of M2, if both transistors have the same geometry. This resistor value along with the capacitor C1 sets the zero of the PLL transfer function. The value of $V_{\rm GS} - V_{\rm TH}$ and hence the transconductance is set by transistors M2 and M3 along with the bias voltage $V_{\rm BP}$. Since this triode resistance value is proportional to $1/\sqrt{I_D}$ (through $V_{\rm BP}$), where I_D is the buffer bias current, the expressions derived in [14] still hold well. A second capacitor C2 is added to reduce the variation in $V_{\rm ctrl}$ by introducing a higher order pole. The value of C2 is typically set to one-tenth the value of C1 for an appropriate compromise between input-jitter rejection and stability of the feedback loop—a higher frequency for the second pole implies greater phase margin (and hence increased stability), but it also allows more input jitter to get through to the output of the PLL.

[4]These resistors have a variance of approximately 4% in this technology, resulting in little effect on the resulting far-end response.

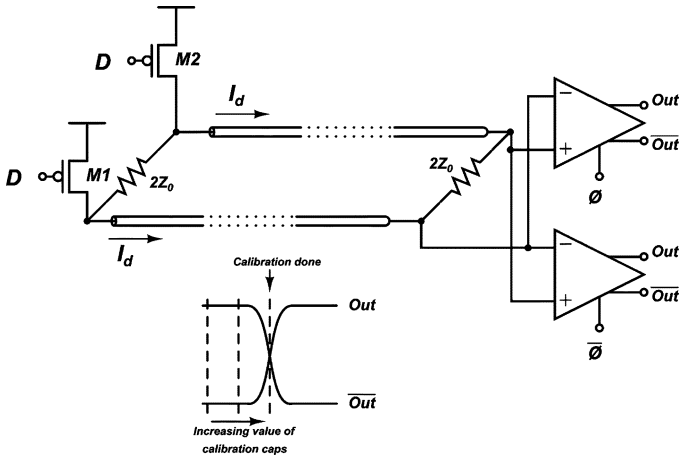

Fig. 9. Receiver offset calibration.

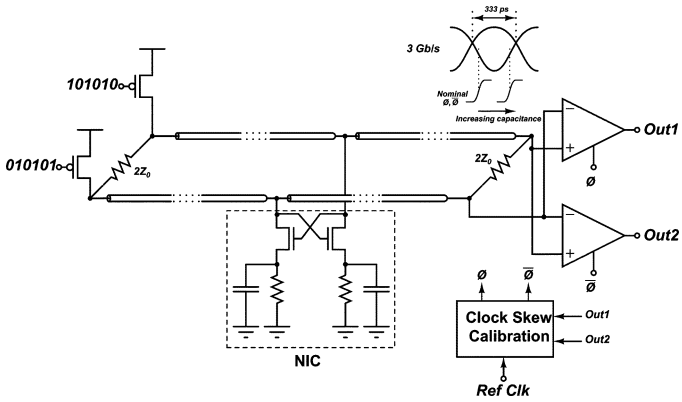

Fig. 10. Clock-skew calibration at the receiver end.

D. Link Calibration

When the link is turned on, receiver offets are first calibrated as shown in Fig. 9. During this receiver calibration mode, the transmitter is configured to source the common-mode current I_d on both lines. This is done by setting the data inputs to each of the two drivers appropriately, such that both driver transistors (M1 and M2) in Fig. 6 are never both off at any instant. Following this, a calibration sequence, implemented with an on-chip finite-state machine (FSM), is performed to tune the position of the receiver clock edge to optimally sample the data as shown in Fig. 10. For a throughput of 3 Gb/s, the eye-opening at the far-end allows an aperature window of $\approx$250 ps for correct detection. The transmitter is configured to send a bitstream consisting of alternating 0's and 1's. This data pattern results in the smallest possible eye-opening at the transmit end due to the finite rise time at the output of the driver. In addition, this data pattern results in maximum energy in the high-frequency region, which also implies maximum signal attenuation. These two factors together account for the smallest possible eye-opening at

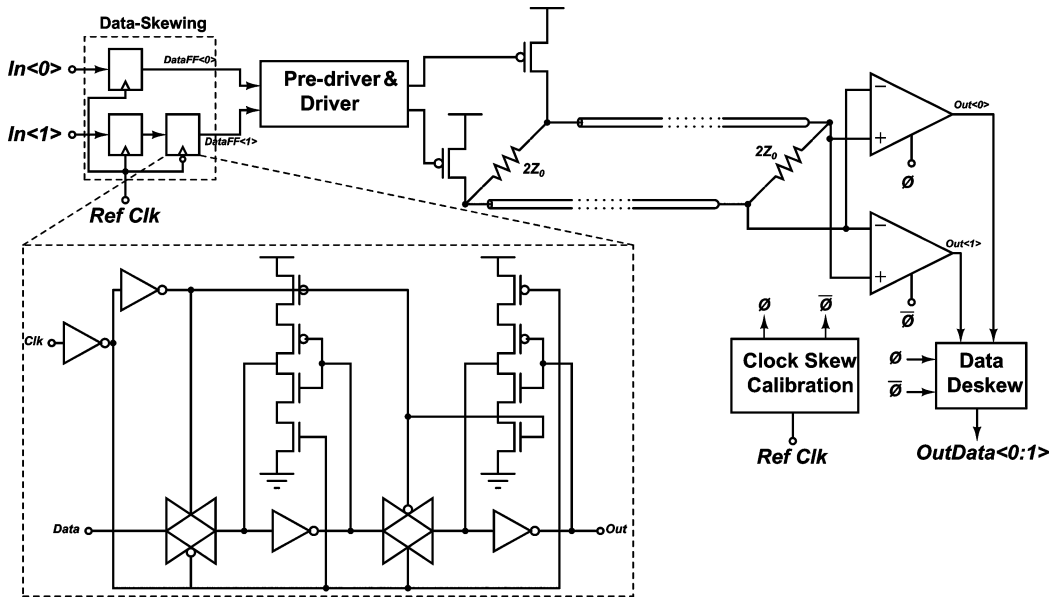

Fig. 11. Skewing/de-skewing circuits.

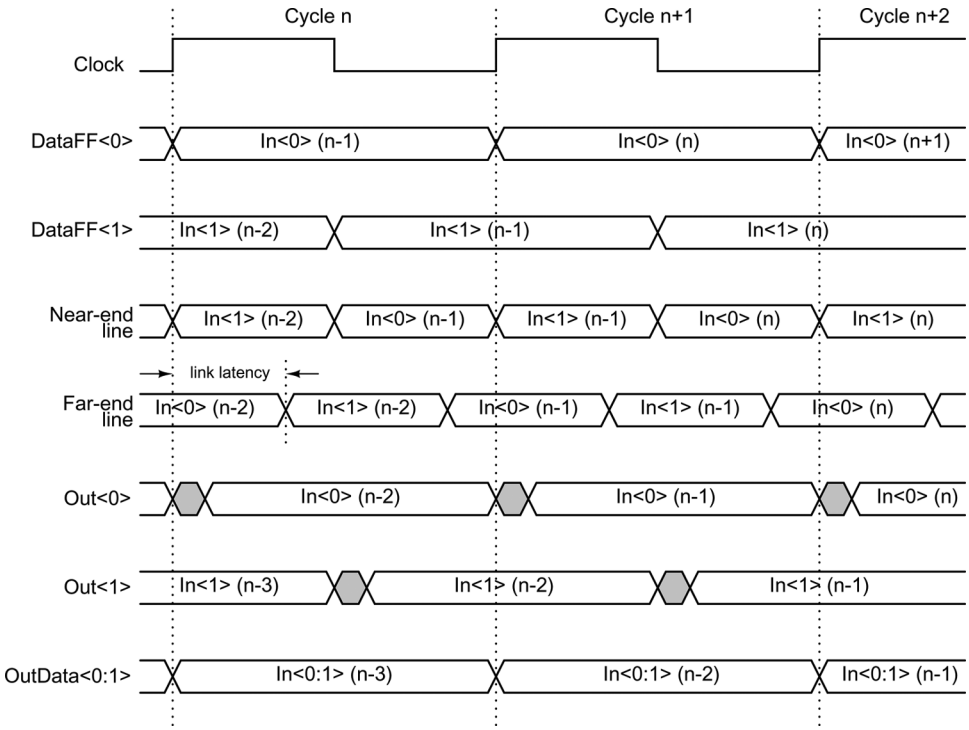

Fig. 12. Timing diagram.

the receiver end. The clock delay elements consist of inverter stages with digitally trimmed capacitive loading, providing a maximum delay of 250 ps in steps of ≈8 ps. The calibration controller varies this clock delay from minimum to maximum and vice versa to position the clock edge at the optimal location.

E. Data Skewing and Deskewing

The system is designed to operate with two cycles of latency, including data skewing and de-skewing, with a cycle time of 667 ps (1.5 GHz). These skewing (de-skewing) latches are shown in Fig. 11 with the timing diagram shown in Fig. 12. Bit 0 of the input $In_{(n-1)}$ is latched by the skewing latch at the rising clock edge (cycle n) followed by bit 1, which is latched on at the next falling clock edge. Both the bits of the input $In_{(n-1)}$ are available at the output of the de-skewing latches after two clock cycles (cycle $n+2$), as shown in Fig. 12. The delay of the receiver and transmit clock phases is determined in the calibration sequence described above. The link latency is constrained by the synchronization requirement that the sum of the transmitter delay, link latency, and receiver latch delay must be less than a half a cycle (333 ps). For the system considered here, link latencies up to 270 ps can be accommodated. The 14-mm link considered here has a latency of approximately 170 ps.

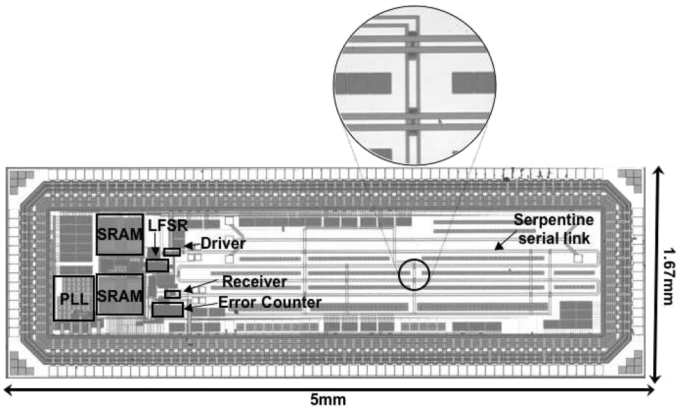

Fig. 13. Die photograph of a prototype link.

F. Test Setup

The design consists of various other blocks to facilitate testing and characterization. Input patterns to the link can be generated either from a PRBS generator (consisting of a 17-b LFSR) or from a 512 $\times$ 16-b SRAM. Demultiplexed results at the other end of the link can be stored into another 512 $\times$ 16-b SRAM or, in the case of the PRBS pattern, can be compared with the transmitted pattern with a 32-b error counter used for measuring BER. Picoprobe pads are included at both the driver and receiver side to allow direct probing of the waveforms on the link as well as for S-parameter measurement for extraction of transmission-line parameters.

In addition, the testing environment also features an identical link running on the third-level metal to add more realistic loading on the main link (running on the fifth-level metal).

IV. Measurement Results

The die photograph of the prototype link is shown in Fig. 13 with a 14-mm interconnect length in a TSMC 0.18-μm process. One NIC consumes an area of approximately 200 μm^2, dominated by the area consumed by pFETs used in implementing the capacitors. Transistors that are 22 μm wide (0.18 μm long) are employed in the NIC. Fig. 14(a) shows the *measured* eye diagram (using picoprobes) at the far end for the link operating at 3 Gb/s; a 70-mV swing (for the lowest power setting) and well-defined eye are evident. The measured BER for this power setting was $< 10^{-14}$. In Fig. 14(b), we show the same receiver eye diagram for the maximum power setting with 6 mA of drive current.[5] The overall shape of the eye diagram remains the same, indicating a linear relationship between the eye-opening width and the drive current. Fig. 15 shows the transmit-end eye diagram at the maximum power setting. A comparison between Figs. 14 and 15 shows the low-pass filtering effect of the interconnect which is minimized by the effect of the NICs in boosting the frequency response of the interconnect up to $\approx$8 GHz.

The loss compensation characteristic of the interconnect is illustrated in Fig. 16, which shows the attenuation constant α

[5]The picoprobes add approximately 100 fF of loading to the receivers, degrading the $Z_o C_L$ time constant from 2 ps to approximately 20 ps, but still insignificant when compared with the bit period of 333 ps.

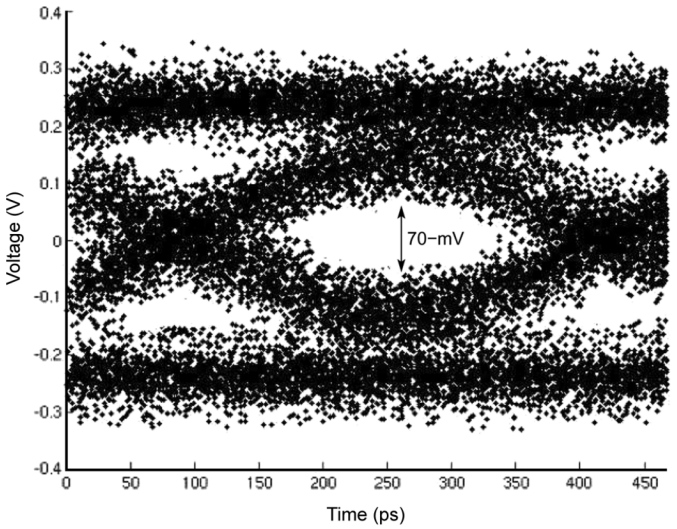

(a)

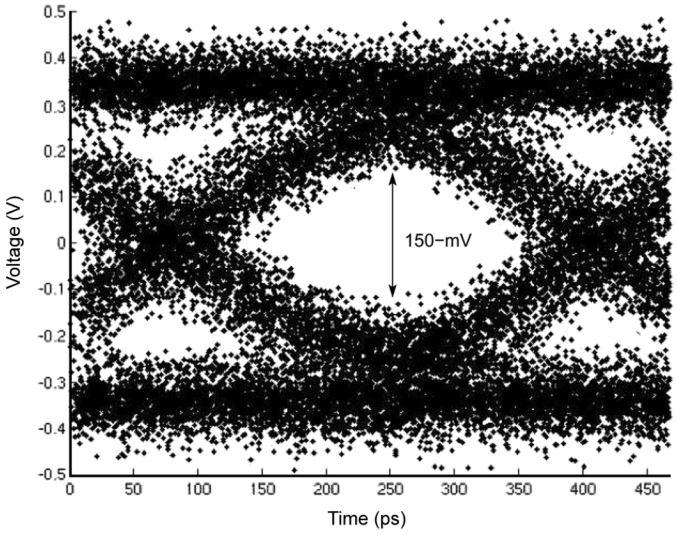

(b)

Fig. 14. Measured far-end eye diagram for (a) minimum power setting and (b) maximum power setting (β).

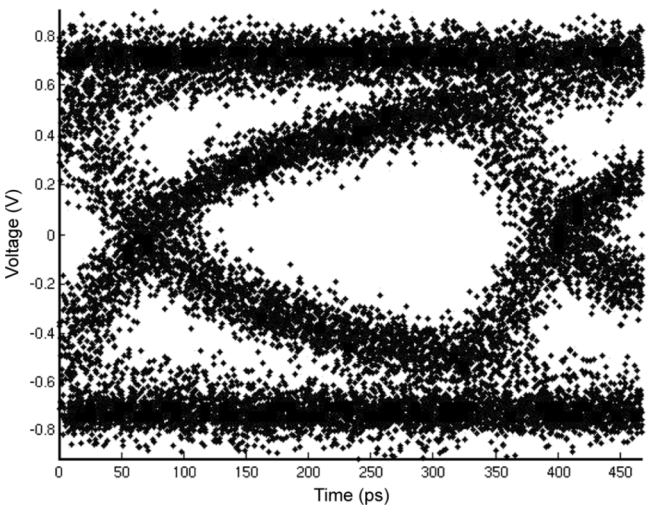

Fig. 15. Measured eye diagram at the near-end of the link operating at maximum power.

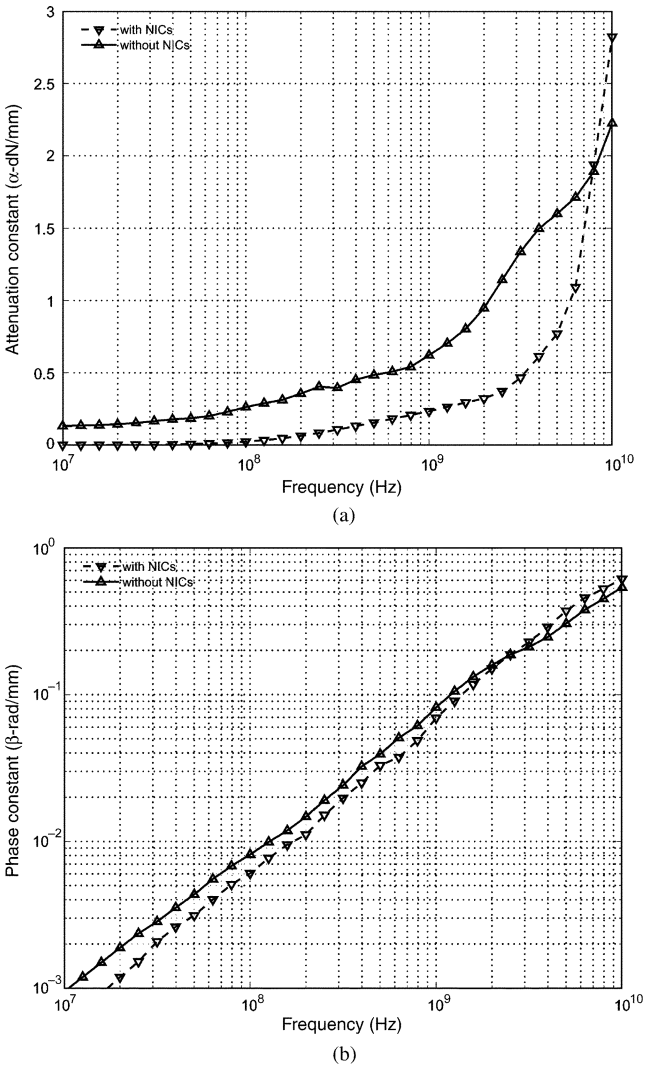

Fig. 16. (a) Measured attenuation constant (α) and (b) measured phase constant (β) as a function of frequency.

for the 14-mm interconnect as extracted from the measured S-parameters with and without enabled NICs; the NICs, when enabled, reduce α by approximately 0.9 dB/mm at 3 GHz. Inactive NICs add approximately 40 fF of loading (primarily gate capacitance) to the interconnect every two millimeters; this additional loading only increases α by approximately 0.1 dB/mm (at 3 GHz) compared with a wire with no NICs at all. Also shown in the same figure is the phase velocity of the interconnect, which clearly shows the tradeoff between loss compensation and latency—lower attenuation implies higher latency. The *measured* latency of the 14-mm link, determined through probing, is 10.7 ps/mm with the NICs present but unbiased and 12.1 ps/mm with the NICs biased on. The latency of the link in the absence of any NIC loading would be approximately 8 ps/mm. This compares with a latency of more than 18.6 ps/mm for an identical optimally buffered link.

The link, through the action of the drivers, consumes approximately 2 pJ/b, which is consistent with the predictions of Fig. 1(a). This includes the power consumed in the NICs as well as the predrivers.[6] The use of DDR allows serialization to be performed with the limited power overhead of skewing and de-skewing latches. Table I summarizes the measured results and performance characteristics. Simulated comparison number for the energy per bit and latency of the optimally repeated link are indicated, which are consistent with the values of Fig. 1.

TABLE I
SUMMARY OF CHARACTERISTICS AND MEASURED PERFORMANCE

Technology	0.18-μm CMOS, 6-level Al
Nominal V_{DD}	1.8 V
Throughput	3 Gb/s
Link length	14 mm
Link latency	12.1 ps/mm ($\approx$ 18.6 ps/mm for an identical conventional link)
Link structure	Differential configuration on M5 layer
M5 thickness	0.53 μm
Line width	8 μm
Line spacing	8 μm
Power consumed (Drivers and pre-drivers)	2 pJ/bit (6 pJ/bit for an identical conventional link)
Far-end differential swing	70 mV at minimum power 150 mV at maximum power
Bit error rate at 3 Gb/s	$< 10^{-14}$ (minimum power)

V. CONCLUSION

We have demonstrated the use of NICs to compensate for transmission-line losses in on-chip global interconnects. In contrast to repeaterless implementations, the use of NICs (as shunt "repeaters") improves bandwidth at a latency cost. The resulting latencies are still significantly lower than conventional optimally buffered links. Bit energies are also substantially less (factor-of-three) for line-widths $>$ 8 μm as compared with optimally buffered *RC* lines. Bit energies can be further improved by running the link at higher bit rates as most of the power dissipation is static. Higher bandwidth copper wires in deep-submicrometer technologies will allow the advantages of distributed loss compensation to be achieved with narrower wires, resulting in a 33% increase [2] in the bandwidth density.

REFERENCES

[1] C. Kim, D. Burger, and S. W. Keckler, "Nonuniform cache architectures for wire-delay dominated on-chip caches," *IEEE Micro*, no. 6, pp. 99–107, Nov.–Dec. 2003.
[2] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proc. IEEE*, vol. 89, no. 4, pp. 490–504, Apr. 2001.
[3] P. Saxena, N. Menezes, P. Cocchini, and D. A. Kirkpatrick, "Repeater scaling and its impact on CAD," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 23, no. 4, pp. 451–463, Apr. 2004.
[4] A. P. Jose, G. Patounakis, and K. L. Shepard, "Pulsed current-mode signaling for nearly speed-of-light intrachip communication," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 772–780, Apr. 2006.

[6]The power consumed in the PLL is not included in this energy-per-bit figure as this is not an overhead of the link itself.

[5] A. Jose and K. L. Shepard, "Distributed loss compensation for low-latency on-chip interconnects," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2006, pp. 19–20.
[6] J. G. Linvill, "Transistor negative impedance converters," *Proc. IRE*, vol. 41, pp. 725–729, Jun. 1953.
[7] H.-T. Ahn and D. J. Allstot, "A 0.5–8.5 GHz differential CMOS distributed amplifier," *IEEE J. Solid-State Circuits*, vol. 37, no. 8, pp. 985–993, Aug. 2002.
[8] F. O'Mahony, C. P. Yue, M. A. Horowitz, and S. S. Wong, "A 10-GHz global clock distribution using coupled standing-wave oscillators," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1813–1820, Nov. 2003.
[9] *Low Voltage Differential Signals (LVDS) for Scalable Coherent Interface (SCI)*, IEEE Std. 1596.3-1196, 1994.
[10] S. Galal and B. Razavi, "10-Gb/s limiting amplifier and laser/modulator driver in 0.18- μm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2138–2146, Dec. 2003.
[11] D. M. Pozar, *Microwave Engineering*, 2nd ed. New York: Wiley, 1999.
[12] J. Montanaro *et al.*, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1712, Nov. 1996.
[13] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1591–1599, Nov. 2000.
[14] J. G. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1723–1732, 1996.
[15] J. G. Maneatis, J. Kim, I. McClatchie, J. Maxey, and M. Shankaradas, "Self-biased high-bandwidth low-jitter 1-to-4096 multiplier clock generator PLL," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1795–1803, Nov. 2003.
[16] B. Razavi, "A 2.4-GHz CMOS receiver for IEEE 802.11 wireless LANs," *IEEE J. Solid-State Circuits*, vol. 34, no. 10, pp. 1382–1385, Oct. 1999.

Anup P. Jose (S'01–M'06) received the B.Tech. degree from the Indian Institute of Technology, Madras, India, in 2001, and the M.S. and Ph.D. degrees from Columbia University, New York, NY, in 2003 and 2006, respectively, all in electrical engineering. His research focused on low-latency low-power interconnects for on-chip networks.

He was with the IBM T. J. Watson Research Center, Yorktown Heights, NY, for the summers of 2002–2004, where he was involved with on-chip jitter measurement circuits and an on-chip spectrum analyzer. During the summer of 2005, he was with the IBM Austin Research Laboratory, Austin, TX, where he was responsible for the design of on-chip sampling circuits in 65-nm technology. He is currently a Senior Design Engineer with the high-speed I/O group at AMD's Boston Design Center, Boxborough, MA.

Dr. Jose was the recipient of the 2005 Best Paper Award at the European Solid-State Circuits Conference.

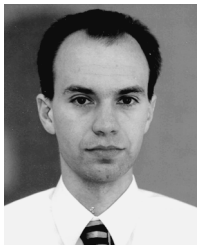

Kenneth L. Shepard (S'85–M'92–SM'03) received the B.S.E. degree from Princeton University, Princeton, NJ, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

From 1992 to 1997, he was a Research Staff Member and Manager with the VLSI Design Department, the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors. Since 1997, he has been with Columbia University, where he is now Associate Professor. He also served as Chief Technology Officer of CadMOS Design Technology, San Jose, CA, until its acquisition by Cadence Design Systems in 2001. His current research interests include design tools for advanced CMOS technology, on-chip test and measurement circuitry, low-power design techniques for digital signal processing, low-power intrachip communications, and CMOS imaging applied to biological applications.

Dr. Shepard was the recipient of the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992. At IBM, he was the recipient of Research Division Awards in 1995 and 1997. He was also the recipient of a National Science Foundation CAREER Award in 1998 and IBM University Partnership Awards from 1998 through 2002. He was also awarded the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS and was the Technical Program Chair and General Chair for the 2002 and 2003 International Conference on Computer Design, respectively. He has served on the program committees for ISSCC, VLSI Symposium, ICCAD, DAC, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD.