

Pulsed Current-Mode Signaling for Nearly Speed-of-Light Intrachip Communication

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

Abstract—In this paper, we describe the design of on-chip repeater-less interconnects with nearly speed-of-light latency. Sharp current-pulse data transmission is used to modulate transmitter energy to higher frequencies, where the effect of wire inductance is maximized, allowing the on-chip wires to function as transmission lines with considerably reduced dispersion. A prototype 8-Gb/s serial link employing this pulsed current-mode signaling in a 0.18- μm CMOS process is described and measured.

Index Terms—High-speed integrated circuits, integrated circuit interconnections, on-chip communication, serial links, transmission lines.

I. INTRODUCTION

OVER THE PAST few decades, improvements in integrated circuit density and performance have been achieved by scaling down transistors. Local interconnects that span a few gate pitches also scale accordingly. However, the RC delay of global wires has been increasing, especially when measured relative to gate delay, as wire resistance per micron for minimum width wires approximately doubles with every process generation [1], [2]. In this paper, we show that relying on full-rail RC -limited interconnect buffered by CMOS inverters for on-chip communications is slow (high latency) and energy inefficient when compared with techniques that exploit wire inductance.

As the latency of an RC -line grows quadratically with line length, buffers (or repeaters) are traditionally added to make the interconnect latency linear with wire length, with simple relationships guiding an optimal number of repeaters to minimize interconnect delay [1]. Fig. 2 compares the latency of optimally repeated RC lines of various widths;¹ RC latency improves as the wires are widened but this improved latency comes at the cost of higher power, more silicon area and greater use of

Manuscript received September 5, 2005; revised December 19, 2005. This work was supported by the MARCO/DARPA C2S2 Center and by gifts from IBM and Intel.

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

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

Digital Object Identifier 10.1109/JSSC.2006.870922

¹Dielectric and wire thicknesses here correspond to the 0.18- μm technology node with aluminum wires.




Fig. 1. Maximum distance that can be covered for various technology nodes in both the cases.

routing resources as more repeaters are required to drive the additional capacitance. The latency reduces with the addition of repeaters until the point at which latency is determined primarily by “silicon” delay—the delay of the repeaters driving gate capacitance and areal-dominated wire capacitance.

Nonetheless, with increasing clock rates, this optimally repeated RC interconnect delay is growing to represent a significant number of cycles in deeply scaled CMOS technologies. The right-hand part of Fig. 1 shows the distance a signal can travel with an optimally repeated RC delay from the center of the chip if a cycle is assumed to be approximately seven fanout-of-four (FO4) delays and wires are assumed to be of minimum width for the technology node. The chip itself is shown as 20 mm \times 20 mm. The number of repeaters required under these scaling assumptions is rapidly increasing. The (projected) number of repeaters in a typical microprocessor block as a function of total cell area reaches 35% at the 45-nm node and 70% at the 32-nm node [2].

In many cases, the additional latency of on-chip interconnects can have a significant impact on system performance. This is evident, for example, in access times to large on-chip L2 caches, in which on-chip wire latency results in higher miss penalties. This is motivating nonuniform cache architectures in which data accessed most frequently is moved to areas of the cache array closer to the processor [3]. Large on-chip caches with single, discrete hit latencies lead to wasted cycles.

By operating full-rail and without taking advantage of inductance, optimally repeated RC lines also have high bit energies. Fig. 3 shows the energy per bit as a function of wire length for optimally repeated RC lines (the dotted curves) at different wire widths in 0.18- μm CMOS technology with a wire thickness of 0.53 μm . It is immediately evident that the use of wide wires to minimize latency results in larger bit energies because of larger wire capacitances.

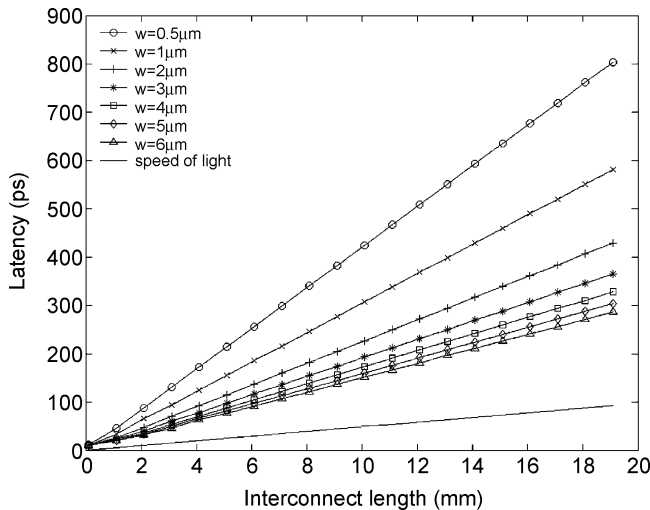


Fig. 2. Link latency as a function of interconnect length, comparing the latency of optimally repeated RC lines with the speed-of-light latency.

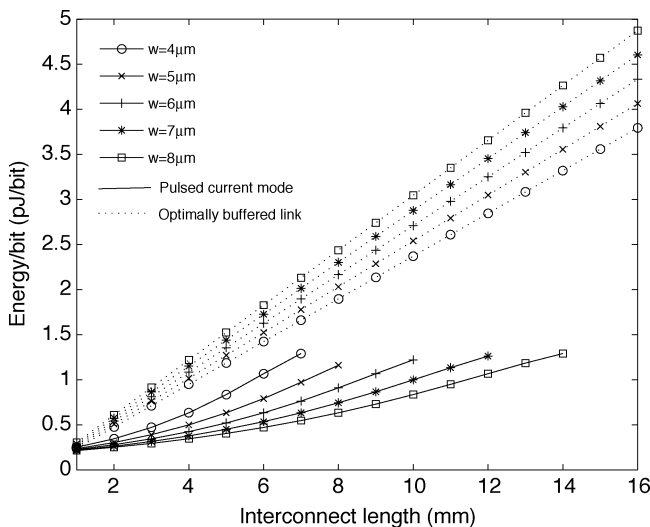


Fig. 3. Energy per bit comparing optimally repeated RC lines with low-swing current-mode links as a function of line length.

Most of these challenges come about because of the constraints imposed by full-rail RC -limited on-chip data transmission. The fact remains that these optimally repeated RC interconnect delays are significantly higher than the speed of light in silicon dioxide (approximately 5 ps/mm), the true physical limit to information propagation. As shown in Fig. 2, optimally repeated RC delays (in 0.18- μm technology) are at least a factor of three greater than speed-of-light-limited propagation. This fact is also represented by the left-hand part of Fig. 1 in which the maximum radius of propagation of a speed-of-light link is contrasted with the optimally repeated RC case.

In this paper, we explore the design of on-chip repeater-less interconnects which use inductance to achieve nearly speed-of-light latency and low bit energies through reduced-swing, current-mode operation. Earlier efforts to exploit wire inductance to achieve nearly speed-of-light propagation characteristics resulted in the development of a 1-GHz link operating with phase-shift keying on a 7.5-GHz sinusoidal

carrier [4]. This type of modulation results in a relatively large power dissipation and poor spectral efficiency due to the need to modulate using a carrier significantly higher in frequency than the transmission bandwidth. Instead, we explore the use of sharp current-pulse data transmission to modulate transmitted energy to higher frequencies, where the effect of wire inductance is maximized, allowing the on-chip wires to function as relatively dispersionless (albeit lossy) transmission lines. To reduce the complexity of the link and to test the limits of this relatively simple driver pre-emphasis, more complex equalization schemes [5] are not employed. In Section II, we consider the properties of lossy on-chip transmission lines and the advantages of pulsed current-mode data transmission. Section III considers detailed design of a prototype on-chip 8-Gb/s serial link using this transmission approach. In Section IV, we present measurement results on the link. Section V concludes and provides a possible context for this work in the design of on-chip networks.

II. LOSSY ON-CHIP TRANSMISSION LINES AND PULSED CURRENT-MODE SIGNALING

As frequency increases, the inductive part of the impedance of on-chip wires increases (relative to the resistive part) and this can be exploited to achieve low wire latency. A digital signal is broadband, containing significant energy content up to the *knee frequency*, $f_{\text{knee}} \cong 0.5/t_r$, where t_r is the slew rate [6]. Beyond f_{knee} , the power spectral density of a digital signal typically falls off with greater than -20 -dB/decade slope. Phase velocity and attenuation increase with frequency [4], with both having the effect of “smearing out” edges of switching waveforms and adding delay (as measured, for example, to the 50% point of a full-rail signal). Repeater insertion improves effective latency and bandwidth for long wires by amplifying the high-frequency content of full-rail switching signals but it comes at the cost of significant power dissipation; the repeaters themselves also add delay. Lower latencies, latencies that approach speed-of-light propagation, can instead be achieved by relying on only the high-frequency content in the digital signal to carry information, exploiting the LC -dominated high-frequency transmission-line response of the interconnect.

For a transmission line characterized by a resistance (R), inductance (L), capacitance (C), and conductance (G) per unit length, the propagation constant is given by $\gamma = \alpha + j\beta = \sqrt{(R + j\omega L)(G + j\omega C)}$. The phase velocity at frequency ω is given by $v_p = \omega/\beta$. The conductance of on-chip wires is generally very low except in cases in which a significant amount of displacement current is collected by a lossy substrate or by other signal lines with a lossy connection to power or ground.² In the limit of low loss, α is given approximately by $\alpha = R/2Z_0$, where the high-frequency characteristic impedance $Z_0 = \sqrt{L/C}$.

For the on-chip links considered here, we choose a ground-signal-ground (GSG) coplanar waveguide topology, which offers two principle advantages over a microstrip configuration. First, a microstrip configuration is not consistent with layout image, which generally enforces orthogonal preferred

²Capacitive coupling to the substrate is rare in the metal-dense environment of typical digital integrated circuits.

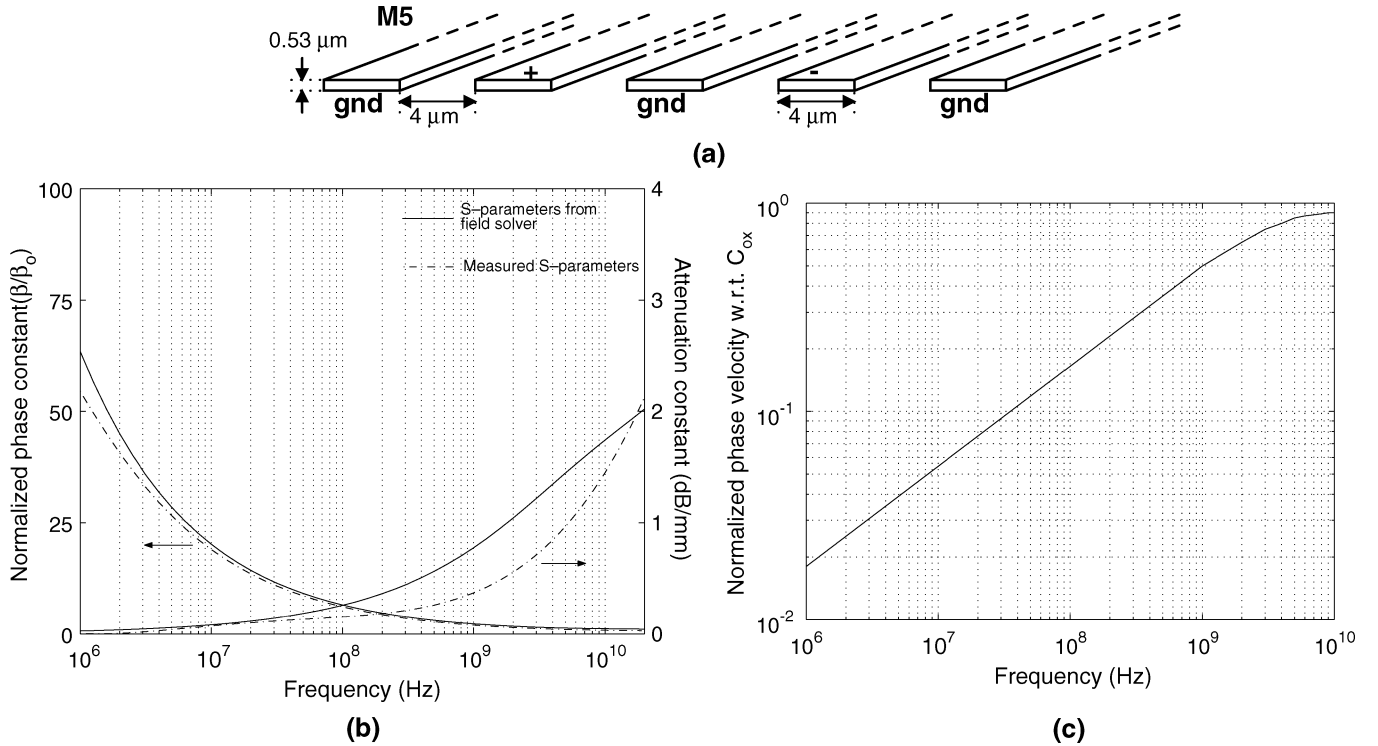


Fig. 4. Real and imaginary part of the propagation constant for the 3-mm-long coplanar transmission line fabricated in this work.

routing directions for adjacent metal layers. The second advantage of this coplanar topology is that there is more flexibility to provide limited engineering of Z_0 ; in particular, the characteristic impedance can be increased slightly by engineering the distance between the signal and ground lines, helping to reduce attenuation in the link. In Fig. 4(b), we plot the α and β of the 3-mm-long on-chip transmission line fabricated in our prototype, a coplanar GSG configuration. Both signal and ground wires are $4\ \mu\text{m}$ wide and $0.53\ \mu\text{m}$ thick, running on fifth-level metal in a six-level-metal aluminum process as shown in Fig. 4(a). There is $4\text{-}\mu\text{m}$ spacing between the wires. Results are shown for a full-wave integral equation field solution (curve labeled “S-parameters from field solver”) and as extracted from measured S-parameters of the 3-mm link (details of this measurement are described in Section IV). The calculated full-wave S-parameters overestimate the loss primarily because they do not include some loss-reducing aspects of the measured links, including additional “strapping” of the ground return lines to each other. We also find that this link is well modeled by an *RLCG* representation with $R = 27\ \Omega/\text{mm}$, $C = 103\ \text{fF}/\text{mm}$ and $L = 0.53\ \text{nH}/\text{mm}$. The latter works very well because of the well-defined current returns of the shielded coplanar topology. The β is plotted normalized to $\beta_0 = \omega\sqrt{LC}$. The line shows very little dispersion beyond 10 GHz. Fig. 4(c) shows the corresponding phase velocity; beyond 10 GHz, nearly speed-of-light-limited propagation is achieved.

In this work, we mitigate dispersion by means of a return-to-zero (RZ) signaling scheme in which sharp current pulses are used to transmit data and receiver termination is employed, as shown in Fig. 5(a). It is well known that the power-spectral den-

sity for a random RZ code with bit time (time to transmit a single bit) T_b and pulsewidth T_s is given by [7]

$$\mathcal{P}(f) = \frac{T_s^3}{2T_b^2} \text{sinc}^2(\pi f T_s) \left[1 + \frac{1}{T_b} \sum_{n=-\infty}^{n=\infty} \delta\left(f - \frac{n}{T_b}\right) \right].$$

Furthermore, the spectral efficiency, defined as the number of bits per second of data that can be transmitted for each hertz of bandwidth, is given by $\eta = T_s/T_b$. Fig. 5(b) compares the normalized power spectral density of an RZ code with that of a non-return-to-zero (NRZ) code and clearly shows the increase in high-frequency content when using an RZ signaling scheme. For the 8-Gb/s link considered in this work, $T_b = 125\ \text{ps}$ and $T_s = 96\ \text{ps}$. The low-frequency spectral components fall as the square of T_b/T_s . As the pulsewidth decreases, more of the energy of communication is pushed to higher frequencies, but a decreasing T_s also means lower spectral efficiency. Fig. 6 shows the simulated width (timing margin in ps) of the “margin rectangle” in the data eye (see Fig. 6 inset) for a fixed voltage margin (rectangle height) of 60 mV. The drive current is adjusted for a constant 120 mV pulse height at the far-end of the line; T_b is fixed at 125 ps. For long wires, larger pulse widths show proportionally larger eye “collapse” due to increased intersymbol interference (ISI), both because there is more dispersion in the pulses themselves and because of reduced spacing between the pulses. Because shorter pulse widths result in lower spectral efficiency, there is a direct trade-off between spectral efficiency and ISI mitigation. Smaller values of T_s will be necessary for longer links or links with intrinsically higher loss. The

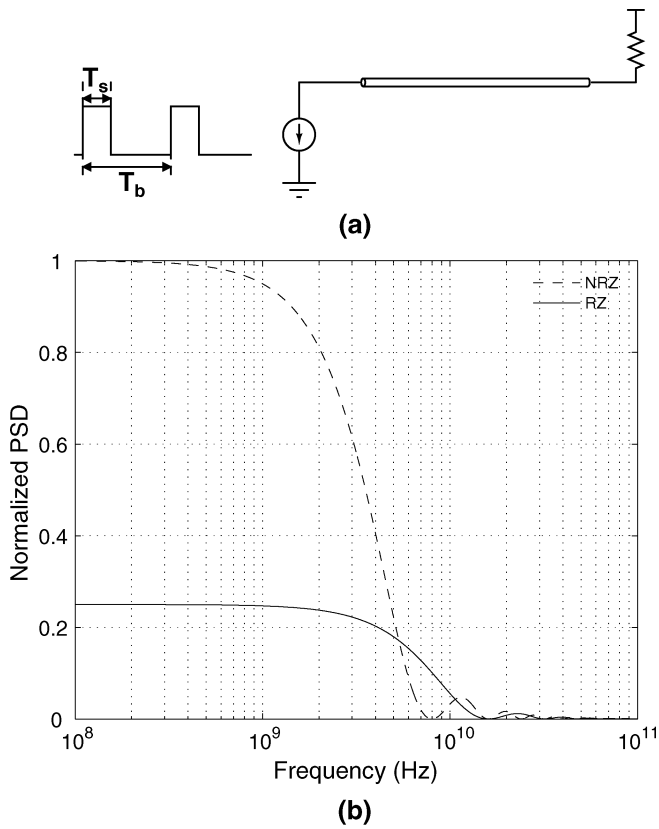


Fig. 5. Pulsed current-mode driver with bit time T_b and pulsewidth T_s .




Fig. 6. Far-end timing margin for a constant voltage swing (120 mV) at far-end and a fixed voltage margin of 60 mV. T_b is fixed at 125 ps. Inset: “Margin rectangle” filling a data eye.

improvement in ISI performance for such long links more than offsets the loss in spectral efficiency and also gets around the need for using more complex equalization techniques. By increasing spectral content at higher frequencies through narrow pulse widths, Fig. 7 shows that for a fixed-drive current, attenuation increases with decreasing T_s and increasing wire length.




Fig. 7. Far-end voltage swing for a fixed near-end pulse height of 240 mV.




Fig. 8. Far-end voltage swing for various driver strengths in the case of a 6 mm interconnect.

In Fig. 3, in addition to the energy-per-bit for the optimally repeated RC line, we also show the energy-per-bit for a pulsed current-mode link (with a 96-ps pulse width) for a fixed 120-mV pulse height at the far end of the line. Bit energy increases with increasing wire length because the drivers must provide more current to maintain the far-end pulse height in the presence of increased attenuation (see Fig. 7). Fig. 8 shows the far-end pulse height for various driver strengths in the case of a 6 mm long interconnect. This gives an idea of the minimum pulsewidth (and hence maximum throughput) possible for a given power budget and interconnect length. Bit energy increases with decreasing wire width for the same reason. Transmission length is limited in each case; for wires that are too narrow or too long, the wire degenerates into an RC line and significant “transmission line” pulse propagation is not possible. Decreasing wire width increases the transition frequency between the RC and LC domains making transmission-line propagation increasingly difficult at lower frequencies [4].

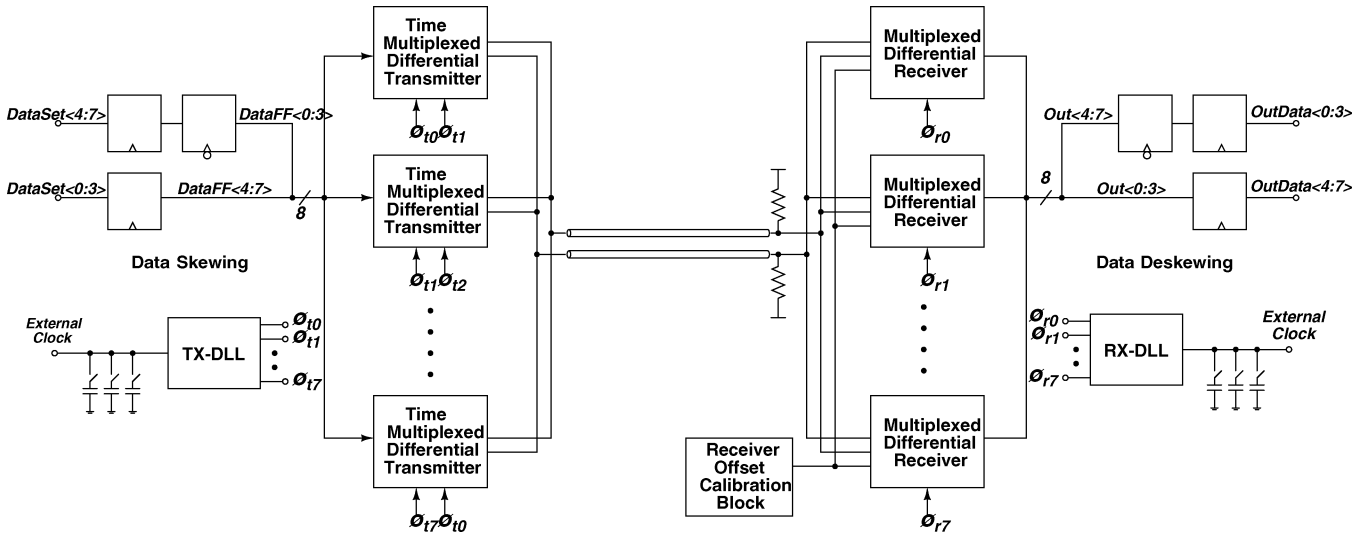


Fig. 9. Overall system architecture of the 8-Gb/s on-chip link prototype.




Fig. 10. Driver design.

III. ON-CHIP SERIAL LINK

To study the use of pulsed on-chip current-mode signaling, we have designed and fabricated a prototype link as shown in Fig. 9 which achieves 8-Gb/s operation with a 1-GHz system clock. A unipolar differential signaling scheme is employed which offers good immunity to power supply noise but is more prodigal in its use of routing resources (the coplanar waveguide structure has a G-S-G-S-G configuration). Driver and receiver components are assumed to be clocked by the same global clock, although the system does allow for clock skew between driver and receiver domains to be compensated with an automated calibration at start-up. The main components of this on-chip interconnection prototype are described in separate sections below, including the driver and receiver circuits, the link calibration, the data skewing and deskewing timing, and the on-chip test circuits for link characterization. The 0.18- μm CMOS technology used in this implementation has an FO4 delay (τ_4) of approximately 60 ps.

A. Driver Design

Current-mode drivers with output multiplexing are employed in the design. An output multiplexed architecture, with multiple copies of the wire drivers, achieves higher performance than input multiplexed designs at the expense of higher power dissipation [8]. An input multiplexed scheme using full-swing

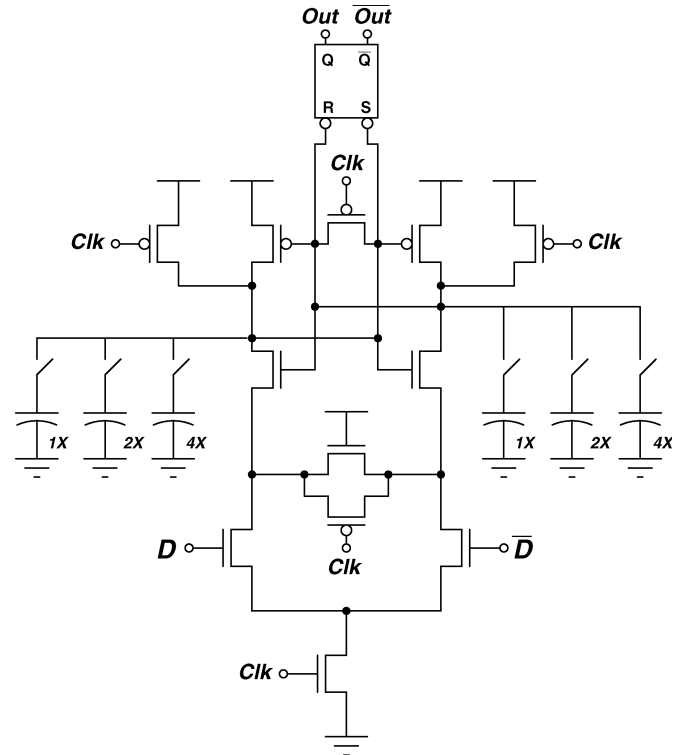


Fig. 11. StrongARM latch with offset calibration capacitors.

static CMOS gates as pre-drivers cannot achieve a bit duration smaller than $4\tau_4$ while the output multiplexed approach can easily achieve a bit duration of 125 ps ($\approx 2\tau_4$ in 0.18- μm technology). This driver design also allows scalability to 16-Gb/s links in this target technology (bit duration $\approx \tau_4$).

Eight drivers of the form shown in Fig. 10 drive out onto the differential coplanar interconnect. The high output impedance of these drivers improve power-supply-noise immunity over voltage-mode drivers. The current pulses are of magnitude 4 mA and width 96 ps with approximately fanout-of-one sizing of the predriver inverters. This driver design allows both edges

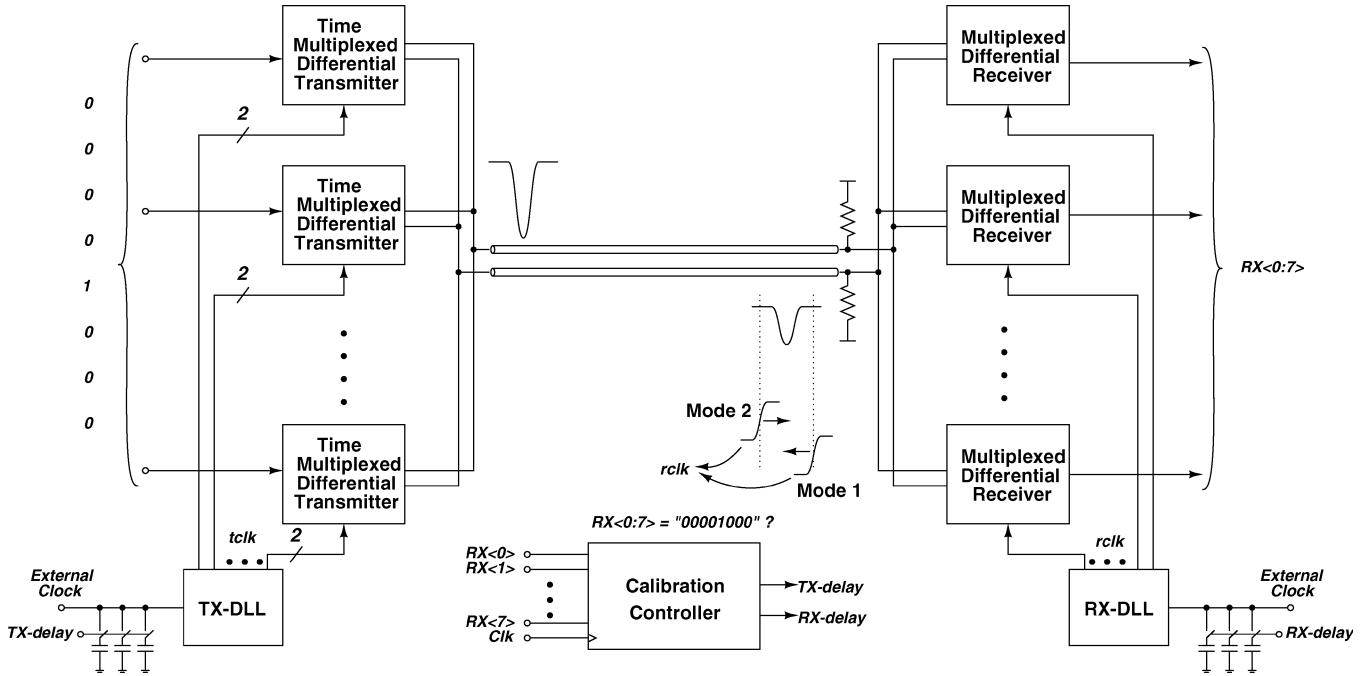


Fig. 12. Schematic of the skew-calibration control loop.

of the current pulse to be controlled by the rising output of the predrivers, permitting significant predriver skewing and improving predriver logical effort. Also, since all the critical transitions are controlled by pFETs, the design is not sensitive to nFET–pFET tracking issues. The capacitor C_s (acting as a dynamic current source) sinks a significant amount of the high-frequency current pulse reducing the need to size up the nFETs of the predriver.

B. Receiver Design

The receivers, shown in Fig. 11, are the StrongARM gate-isolated sense-amplifier latches [9]. 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 [8]. 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 [8].

A (silicided) 90- Ω polysilicon resistor is used for line termination at the receiver. This is slightly larger than the (high-frequency) Z_0 of the line (80 Ω) 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 48 fF on each leg of the differential link. The resulting $Z_0 C_L$ time constant (4 ps) is sufficiently low that it does not introduce significant ISI.

C. Link Calibration

When the link is powered up, the receiver latch offsets are automatically calibrated out. Following this, a calibration sequence is performed to tune the position of the driver and receiver clocks to sample the incoming data. The driver and receiver clocks can both be delayed from the external clock as shown in Fig. 12. The calibration sequence, determined by on-chip digital control, consists of sending the bit sequence “00001000.” To begin the calibration, the receiver clock ($rclk$) is delayed its maximum amount and the driver clock ($tclk$) is not delayed to provide the maximum relative delay in the “positive” direction ($rclk$ lags $tclk$ by 149 ps). This delay is then reduced (both by reducing the receiver delay and increasing the driver delay) in steps of approximately 3 ps until the correct pattern is captured in the receiver latches. This procedure is repeated beginning with the maximum driver delay and minimum receiver delay (for the maximum relative delay in the “negative” direction; $rclk$ leads $tclk$ by 56 ps), reducing the magnitude of this delay in steps of approximately 3 ps until the correct pattern is captured. These two delay settings are averaged to set the correct transmit and receiver clock delays. This approach compensates for clock skew and driver and link latency in the positioning of the sample clocks. It also allows this system to operate in a mesochronous environment where different parts of the chip are clocked at the same frequency but different phases [10].

The delay-locked loops (DLLs) used in the design are based on a regulated inverter delay line, consisting of two inverter chains connected at each stage by cross-coupled inverters [8]. This offers considerable power savings (approximately 75%) over a DLL using source-coupled delay elements. A dynamic phase-only detector is employed [11].




Fig. 13. System timing diagram.

D. Data Skewing and Deskewing

The system is designed to operate with two cycles of latency, including data skewing and deskewing, with a cycle time of 1 ns. These skewing and deskewing latches are shown in Fig. 9 with the timing diagram shown in Fig. 13. Bits 0–3 of the input $DataSet(n-1)$ are latched by the skewing latches at the rising clock edge (cycle n) followed by bits 4–7 which are latched on at the next falling clock edge. All eight bits of the input $DataSet(n-1)$ are available at the output of the deskewing latches after two clock cycles (cycle $n+2$) as shown in Fig. 13. The delay of the receiver and transmit clock phases is determined in the calibration sequence described above. The link latency is constrained by synchronization requirements such that the sum of the transmitter delay, link latency, and receiver latch delay must be less than a half a cycle (in this case, 500 ps). For the system here, this allows link latencies up to 280 ps to be accommodated. The 3-mm links considered here have a link latency of less than 20 ps.

E. Testing Environment

Other circuits are included in the design to facilitate testing and characterization. Input patterns to the link can be generated from a pseudo-random bit sequence (PRBS) generator (consisting of a 17-bit linear feedback shift register, LFSR) or from a 512×16 -bit SRAM. Demultiplexed results at the other end of the link can be stored into another 512×16 -bit SRAM, or in the case of the PRBS pattern, can be checked with a duplicate LFSR at the receiver. Picoprobe pads are included at both the driver and receiver side to allow direct probing of the waveforms on the link and to allow for direct network analysis and TDR/TDT characterization of the line. and to allow for direct network analysis and TDR/TDT characterization of the line.




Fig. 14. Die photo of prototype link.

In addition, crosstalk noise can be injected into the link from an on-chip noise generator that consists of an aggressor parallel link running beneath the victim link. This link is controlled with a 16-stage DLL, allowing aggressor current pulses of varying widths to be positioned with 62.5-ps resolution. The minimum width of this pulse is also 62.5 ps. The small resolution of both the width and position of the noise pulse calls for a low supply voltage sensitivity for precise noise injection. Because of the more aggressive performance requirements, this DLL consists of a source-coupled delay line with symmetric loads [12]. Three different current drives can be used on the drivers of the aggressor link: 1.44 mA, 2.08 mA, and 2.69 mA.

IV. MEASUREMENT RESULTS

The die photo of the prototype link is shown in Fig. 14 with a 3-mm interconnect length in a TSMC $0.18\text{-}\mu\text{m}$ process. Fig. 15 shows the *measured* eye diagram at the receiver for the link operating at 8 Gb/s; a 120-mV swing and well-defined eye are evident.

In Fig. 16, the latency obtained from simulations of the measured S-parameters is shown and contrasted with the speed-of-light latency in SiO_2 . The propagation velocity for the line is slightly lower than the speed-of-light velocity because of “slow-

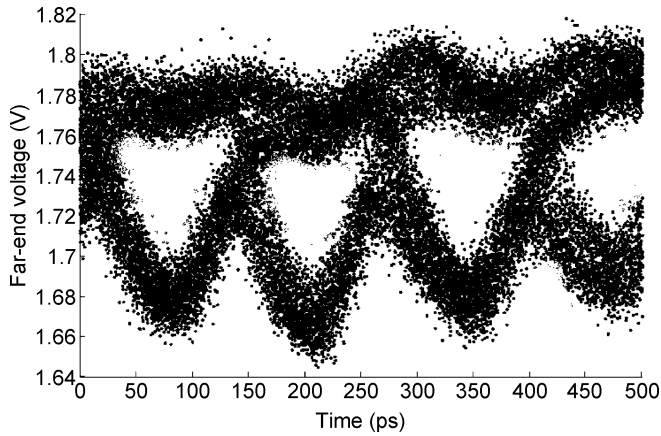


Fig. 15. Measured eye diagram for the link operating without injected noise.




Fig. 17. Measured eye diagram for the link operating with injected noise.

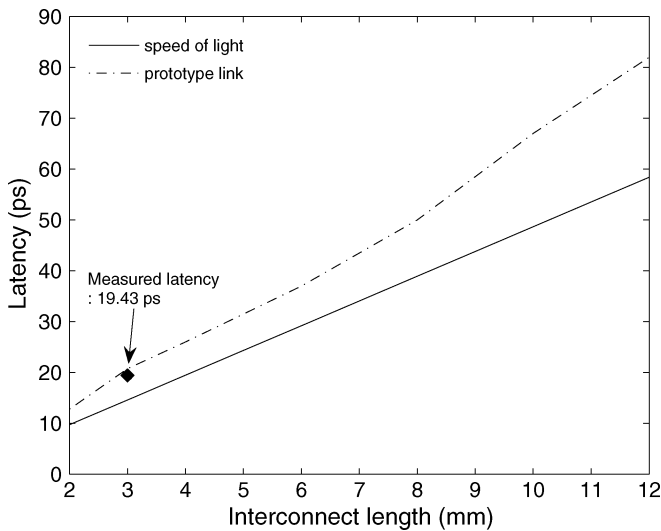


Fig. 16. Measured and simulated interconnect latency.

wave” effects in which some of the displacement current associated with the wire’s capacitance does not find a easy high-frequency return path. Displacement currents to the substrate, for example, can contribute to this slow-wave response. The *measured* latency of the 3-mm link, determined through probing, is noted in Fig. 16.

In Fig. 17, we show the same receiver eye diagram for the case that the aggressor link is engaged with 2.69 mA (the peak drive strength) of drive current and is positioned in time in the first half of the second bit position. Some closing of the eye is observed; although it is not significant enough to upset the link.

The link itself, through the action of the drivers, consumes approximately 0.29 pJ, consistent with the predictions of Fig. 3. Nonetheless, a very significant 3.1 pJ per bit is dissipated in the DLLs and skewing and deskewing latches to achieve the aggressive serialization implemented here. Table I summarizes the measured results and performance characteristics. Less aggressive serialization (such as DDR) will allow the energy benefits of pulsed current-mode links to be realized without this overhead.

TABLE I
SUMMARY OF CHARACTERISTICS AND MEASURED PERFORMANCE

Technology	0.18 μ m CMOS, 6-level Al
Nominal V_{DD}	1.8 V
Throughput	8 Gbps
Link length	3 mm
Link latency	19.4 ps (\approx 55 ps for a conventional link)
Link structure	GSGSG configuration on M5 layer
M5 thickness	0.53 μ m
Line width	4 μ m
Line spacing	4 μ m
Power consumed (Drivers)	0.29 pJ/bit (0.7 pJ/bit for a conventional link)
Far-end differential swing	120 mV
Bit error rate at 8 Gbps	$< 10^{-14}$
Transmission-line loss	6 dB

V. CONCLUSIONS AND FUTURE WORK

We have demonstrated the use of pulsed current-mode signaling to achieve (nearly) speed-of-light latency across lossy on-chip transmission lines. This implies the use of carefully designed circuit and interconnect structures for global interconnect and may make the most sense within the context of network-on-chip [13], [14] architectures. The repeater-less techniques presented here could easily be extended to links as long as 8 mm in the current technology. More advanced copper metallization would allow even longer links. Future work will consider the design of truly differential bipolar current-mode links, which would offer the advantage of a simpler interconnect fabric (no ground returns required).

REFERENCES

- [1] R. Ho, K. W. Mai, and M. A. Horowitz, “The future of wires,” *Proc. IEEE*, vol. 89, no. 4, pp. 490–504, Apr. 2001.
- [2] P. Saxena, N. Menezes, P. Cocchini, and D. A. Kirkpatrick, “The scaling challenge: can correct-by-construction design help?,” in *Proc. Int. Symp. Physical Design*, 2003, pp. 51–57.
- [3] C. Kim, D. Burger, and S. W. Keckler, “Nonuniform cache architectures for wire-delay dominated on-chip caches,” *IEEE Micro*, vol. 23, no. 6, pp. 99–107, Nov–Dec. 2003.

- [4] R. T. Chang, N. Talwalkar, C. P. Yue, and S. S. Wong, "Near speed-of-light signaling over on-chip electrical interconnects," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 834–838, May 2003.
- [5] R. Farjad-Rad, C.-K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4- μ m CMOS 10-Gb/s 4-PAM pre-emphasis serial link transmitter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 580–585, May 1999.
- [6] H. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [7] I. L. W. Couch, *Digital and Analog Communication Systems*. Englewood Cliffs, NJ: Prentice-Hall, 2001.
- [8] 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.
- [9] 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.
- [10] D. G. Messerschmitt, "Synchronization in digital system design," *IEEE J. Sel. Areas Commun.*, vol. 8, no. 8, pp. 1404–1419, Oct. 1990.
- [11] S. Sidiropoulos, D. Liu, J. Kim, G. Wei, and M. A. Horowitz, "Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2000, pp. 124–127.
- [12] 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, Nov. 1996.
- [13] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Proc. Design Automation Conf.*, 2001, pp. 684–689.
- [14] H.-S. Wang, L.-S. Peh, and S. Malik, "Power-driven design of router microarchitectures in on-chip networks," in *Proc. Int. Symp. Microarchitecture (MICRO)*, Nov. 2003, pp. 105–116.



Anup P. Jose (S'01–M'06) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, India, in 2001, and the M.S. and Ph.D. degrees in electrical engineering from Columbia University, New York, NY, in 2003 and 2006, respectively. 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 worked on 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 working as a Senior Design Engineer with the high-speed I/O group at AMD's Boston Design Center.

Mr. Jose received the 2005 Best Paper Award at the European Solid-State Circuits Conference.



George Patounakis (S'00) received the B.S. degree in electrical engineering from Rutgers University, New Brunswick, NJ, in 2000, and the M.S. degree in electrical engineering from Columbia University, New York, NY, in 2001. He is currently working toward the Ph.D. degree at Columbia University in the Columbia Integrated Systems Laboratory.

His research interests include interfacing biological molecules with silicon microelectronics, on-chip power management, and high-speed intrachip interconnects.

Mr. Patounakis is the recipient of the Columbia SEAS Presidential Fellowship and the Intel Ph.D. Fellowship.



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 in the VLSI Design Department at 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, New York, NY, where he is now an Associate Professor. He 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 received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992. At IBM, he received Research Division Awards in 1995 and 1997. He was also the recipient of an NSF 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 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 ICCAD, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD.