Near Speed-of-Light On-Chip Interconnects Using Pulsed Current-Mode Signalling

Anup P. Jose, George Patounakis, K. L. Shepard
Columbia Integrated Systems Laboratory
Columbia University, New York, New York 10027
Email: {anup,gpat,shepard}@cisl.columbia.edu

Abstract

In this paper, we describe the design of on-chip repeaterless 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 Gbps serial link employing this pulsed current-mode signalling in a 0.18 μm CMOS process is described and measured.

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. However, the RC delay of global wires has been increasing as wire resistance per micron for minimum width wires approximately doubles every process generation [1], [2], leading to issues in both latency and energy for on-chip communications.

- *Latency.* 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. 1 compares the latency of optimally-repeated RC lines of various width; RC latency improves as the wires are widened up until the point at which the "areal" component of the wire capacitance dominates. Dielectric and wires thicknesses here correspond to the 180-nm technology node with aluminum wires. Nonetheless, with increasing clock rates, this optimally-repeated RC interconnection delay is growing to represent a significant number of cycles in deeply-scaled CMOS technologies. This is evident, for example, in the latency to large on-chip L2 caches, which is motivating nonuniform cache architectures in which data accessed most frequently is moved to areas closer to the processor [3].
- *Energy.* By operating full-rail, optimally-repeated RC lines have high bit energies. Fig. 2 shows the bit energy as a function of wire length for optimally-repeated RC lines (the dotted curves) at different wire widths. It is immediately evident that the use of wide wires to minimize latency results in larger bit energies because of larger wire capacitances.

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 interconnection delays are significantly slower than the speed-of-light in silicon dioxide (approximately 5 ps/mm), the true physical limit to information propagation. As shown in Fig. 1, optimally-repeated RC delays are at least a factor of three greater than speed-of-light-limited propagation.

In this paper, we explore the design of on-chip repeaterless interconnect with near speed-of-light latency and low bit energy through low-swing, current-mode operation. 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. In the next section, we consider the properties of lossy on-chip transmission lines and the advantages of pulsed current-mode data transmission. The third section considers detailed design of a prototype on-chip 8 Gbps serial link using this transmission approach. In the fourth section, we present measurement results on the link. The last section concludes and provides a possible context for this work in the design of on-chip networks.

This work was supported by the MARCO/DARPA C2S2 Center and by gifts from IBM and Intel.

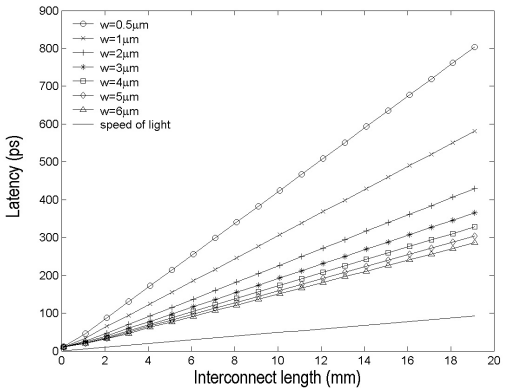

Fig. 1. Link latency as a function of interconnect length, comparing the latency of optimally-repeated RC lines with the speed-of-light latency.

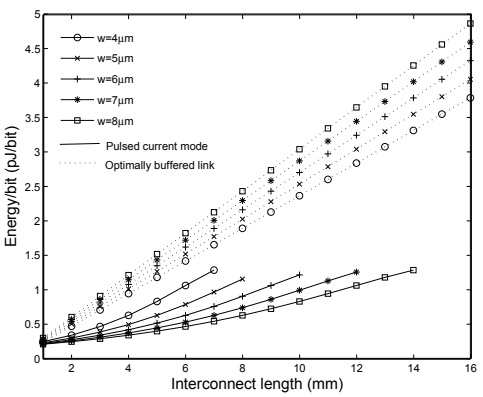

Fig. 2. Energy per bit comparing optimally-repeated RC lines with low-swing current mode links as a function of line length.

Lossy on-chip transmission lines and pulse current-mode signalling

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. 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 conductance of on-chip wires is generally very low except in

cases in which a significant amount of displacement current is collected by the substrate. In the limit of low loss, α is given approximately by $\alpha = R/2Z_0$, where the high-frequency characteristic impedance $Z_0 = \sqrt{L/C}$.

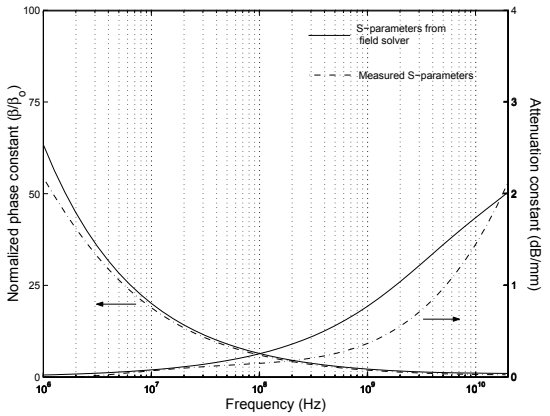

Fig. 3. Real and imaginary part of the propagation constant for the 3-mm-long coplanar transmission line fabricated in this work.

For the on-chip links considered here, we choose a 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 routing directions for adjacent metal layers. The second advantage of a 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. 3, we plot the α and β of the 3-mm-long on-chip transmission line fabricated in our prototype, a coplanar ground-signal-ground configuration. Results are shown for a full-wave integral equation field solution (curve labelled "S-parameter") and as extracted from measured S-parameters of the 3-mm link (details of this measure are described below). 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 modelling by an RLCG representation with with $R = 27\Omega/mm$, $C = 103fF/mm$, and $L = 0.53nH/mm$. The latter works very well because of the well-defined current returns of the shielded coplanar topology. The β is plotted normalized to $\beta_0 = \omega\sqrt{LC}$. Both signal and ground wires are 4 μm wide and 0.53 μm thick, running on fifth-level metal in a six-level metal process. There is 4 μm spacing between the wires. The line shows very little dispersion beyond 10 GHz.

An earlier effort to exploit wire inductance to achieve β_0 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 at a carrier significantly higher in frequency than the transmission bandwidth. We instead choose to mitigate dispersion by means of a RZ signalling scheme in which sharp current pulses are used to transmit data and receiver termination is employed, as shown in Fig. 4. It is well known that the (normalized) power-spectral density for a random RZ code with bit time (time to transmit a single bit) T_b and pulse width T_s is given by [5]:

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

Furthermore, the spectral efficiency, defined as the number of bits per second of data that can be transmitted for each Hertz of

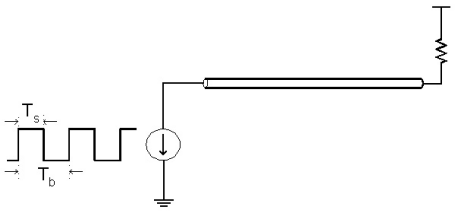

Fig. 4. Pulsed current-mode driver with bit time T_b and pulse width T_s.

bandwidth, is given by $\eta = T_s/T_b$. For the 8 Gbps link considered in this work, $T_b = 125ps$ and $T_s = 96ps$. The low-frequency spectral components fall as the square of T_b/T_s. Additionally, as the pulse width decreases, more of the energy of communication is pushed to higher frequencies. Fig. 5 shows the simulated width (timing margin in ps) of the "margin rectangle" in the data eye (see Fig. 5(a)) 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 there is increased spacing between the pulses. Shorter pulse widths result in lower spectral efficiency, and 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. By increasing spectral content at higher frequencies through narrow pulse widths, Fig. 6 shows that for a fixed-drive current, attenuation increases with decreasing T_s and increasing wire length.

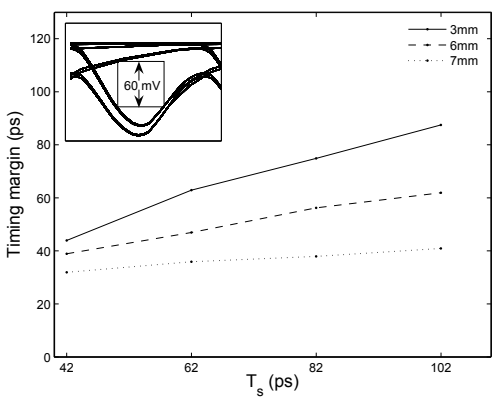

Fig. 5. Far-end timing margin for a constant voltage swing (120 mV) at far-end and a fixed voltage margin of 80 mV. T_b is fixed at 125 ps. *Inset:* "Margin rectangle" filling a data eye.

In Fig. 2, 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 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. 6). 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.

On-chip serial link

To study the use of pulsed on-chip current-mode signalling, we have designed and fabricated a prototype link as shown in

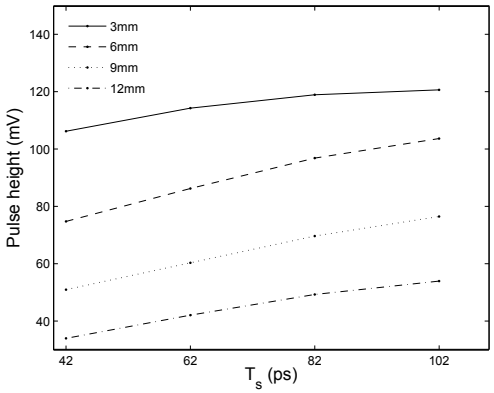

Fig. 6. Far-end voltage swing for a fixed near-end pulse height of 120 mV.

Fig. 7 which achieves 8 Gbps operation with a 1 GHz system clock. A unipolar differential signalling 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 clock skew between driver and and receiver domains to be compensated for with an automated calibration at start-up. The 0.18 μm CMOS technology used in this implementation has a fanout-of-four delay (τ_4) of approximately 60 ps.

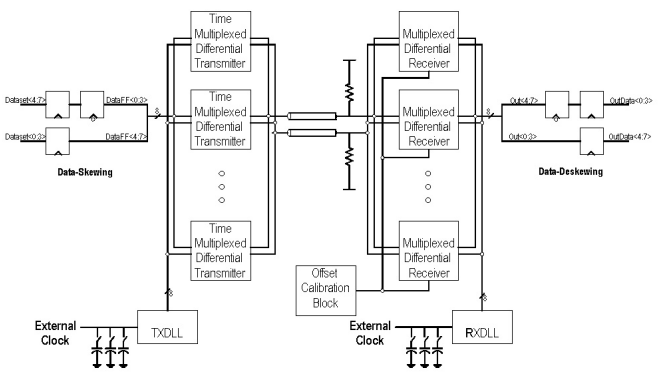

Fig. 7. Overall system architecture of the 8 Gbps on-chip link prototype.

A. Driver design.

Current-mode drivers with output multiplexing are employed in the design. Eight drivers of the form shown in Fig. 8 drive out onto the differential coplanar interconnect. Output multiplexing, despite being costly in terms of power, is chosen over input multiplexing [6] for performance and to allow scalability to 16 Gbps links in the target technology. 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 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.

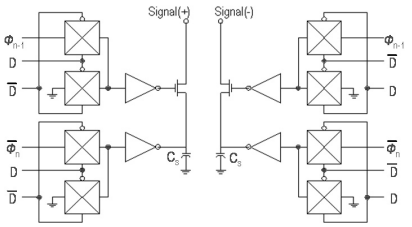

Fig. 8. Driver design.

B. Receiver design.

The receivers used in this design are the StrongARM gate-isolated sense-amplifier latches [7]. 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. Positioning these 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 as in the design of Ref. [6].

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 system clock as shown in Fig. 7. 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) on 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. 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 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 [6]. This offers considerable power savings (approximately 75%) over a DLL using source-coupled delay elements. A dynamic phase-only detector is employed [8].

D. Testing environment.

Other circuits are included in the design to facilitate testing and characterization. Input patterns to the link can be generated from a PRBS generator (consisting of a 17-bit 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.

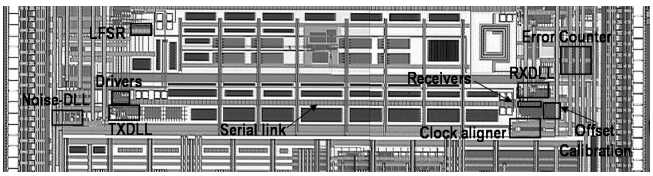

Fig. 9. Die photo of prototype link.

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 to be positioned with 62.5 ps resolution. Because of the more aggressive performance requirements, this DLL consists of a source-coupled delay line with symmetric loads [9]. Three different current drives can be used on the drivers of the aggressor link: 1.44 mA, 2.08 mA, and 2.69 mA.

Measurement results.

The die photo of the prototype link is shown in Fig. 9 with a 3 mm interconnect length in a TSMC 0.18 μm process. Fig. 10 shows the *measured* eye diagram at the receiver for the link operating at 8 Gbps; a 120 mV swing and well-defined eye are evident.

In Fig. 11, 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 lightly lower than the speed-of-light velocity because of "slow-wave" effects in which some of the displacement current associated with the wire's capacitance does not find a easy high-frequency return path. Displacements 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. 11.

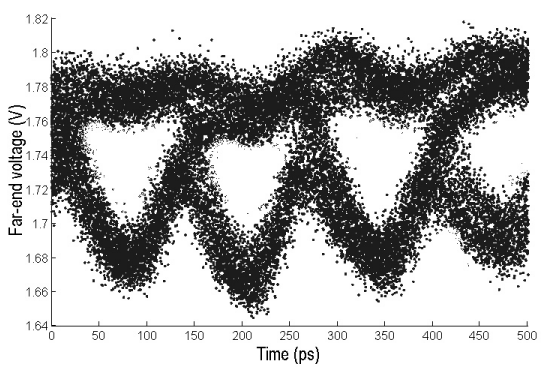

Fig. 10. Measured eye diagram for the link operating without injected noise.

The link itself, through the action of the drivers, consumes approximately 0.29 pJ, consistent with the predictions of Fig. 2. 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. Less aggressive serialization (such as DDR) will allow the energy benefits of pulsed current-mode links to be realized without this overhead.

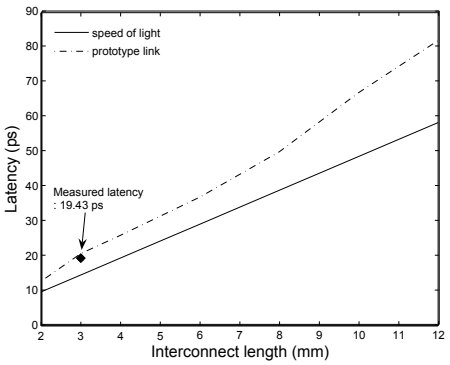

Fig. 11. Measured and simulated interconnect latency.

Conclusions and future work

We have demonstrated the use of pulsed current-mode signalling 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 [10], [11] architectures. We would like to expand this work to consider bipolar differential current-mode links, which would offer the advantage of a simpler interconnect fabric (no ground returns required) and reduced energy around dc, which should improve ISI performance.

References

[1] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proceedings of the IEEE*, pp. 490–504, April 2001.

[2] P. Saxena, N. Menezes, P. Cocchini, and D. A. Kirkpatrick, "The scaling challenge: can correct-by-construction design help?" in *Proceedings of the International Symposium on 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*, pp. 99 – 107, November - December 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 Journal Solid-State Circuits*, vol. 38, pp. 834 – 838, May 2003.

[5] I. L. W. Couch, *Digital and analog communication systems*. Prentice-Hall, 2001.

[6] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed I/O circuit techniques," *IEEE Journal Solid-State Circuits*, pp. 1591 – 1599, November 2000.

[7] J. Montanaro and et al, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE Journal Solid-State Circuits*, vol. 31, pp. 1703–1712, November 1996.

[8] S. Sidiropoulos, D. Liu, J. Kim, G. Wei, and M. Horowitz, "Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers," in *Symposium on VLSI Circuits Digest of Technical Papers*, 2000, pp. 124 – 127.

[9] J. G. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 11, pp. 1723 – 1732, 1996.

[10] W. J. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Proceedings of the Design Automation Conference*, 2001, pp. 684 – 689.

[11] H.-S. Wang, L.-S. Peh, and S. Malik, "Power-driven design of router micrarchitectures in on-chip networks," in *Proceedings of the 36th International Symposium on Microarchitecture (MICRO)*, November 2003.