9.5 High-Bandwidth and Low-Energy On-Chip Signaling with Adaptive Pre-Emphasis in 90nm CMOS

Jae-sun Seo¹, Ron Ho², Jon Lexau², Michael Dayringer², Dennis Sylvester¹, David Blaauw¹

¹University of Michigan, Ann Arbor, MI, ²Sun Microsystems, Menlo Park, CA

Long on-chip wires pose well-known latency, bandwidth, and energy challenges to the designers of high-performance VLSI systems. Repeaters effectively mitigate wire RC effects but do little to improve their energy costs. Moreover, proliferating repeater farms add significant complexity to full-chip integration, motivating circuits to improve wire performance and energy while reducing the number of repeaters. Such methods include capacitive-mode signaling, which combines a capacitive driver with a capacitive load [1,2]; and current-mode signaling, which pairs a resistive driver with a resistive load [3,4]. While both can significantly improve wire performance, capacitive drivers offer added benefits of reduced voltage swing on the wire and intrinsic driver pre-emphasis. As wires scale, slow slew rates on highly resistive interconnects will still limit wire performance due to inter-symbol interference (ISI) [5]. Further improvements can come from equalization circuits on receivers [2] and transmitters [4] that trade off power for bandwidth. In this paper, we extend these ideas to a capacitively driven pulse-mode wire using a transmit-side adaptive FIR filter and a clockless receiver, and show bandwidth densities of 2.2-4.4 Gb/s/µm over 90nm 5mm links, with corresponding energies of 0.24-0.34 pJ/bit on random data.

Figure 9.5.1 shows the basic concept, which we refer to as a single-data-rate (SDR) scheme to distinguish it from a double-data-rate (DDR) scheme described later. Sending data using differential return-to-zero (RZ) pulses improves wire latency, as the wires start each transition already equalized, but requires two (albeit smaller) transitions for each bit. Each differential signal expands in the opposite direction in RZ signaling, resulting in an equivalent eye height with NRZ signaling (see Fig. 9.5.1). A three-tap transmit FIR filter creates sharp, criticallydamped pulses on the wire, allowing these pulses to complete as fast or faster than a long-tailed NRZ transition, in exchange for higher transmit energy. An inverse-FFT Matlab analysis using extracted wire models gave optimal tap values of 1, -1.5, and 0.5. Note that the taps sum to zero for this RZ signal, unlike the taps for NRZ signaling [1]. A simple clockless hysteresis receiver, often used in off-chip communication [6], detects the RZ pulse at the end of the wire (Fig. 9.5.2). To detect sharp RZ pulses, hysteresis receivers are simpler than clocked sense-amplifiers [1] or decision feedback equalization (DFE) receivers [2,4], because they do not need a clock edge carefully positioned on the pulse, a requirement made difficult by link and process variations. Hysteresis receivers consume no energy with idle inputs, unlike clocked receivers with precharge and evaluate [1,2], and also reduce clock load and simplify timing verification. In exchange, they are less efficient in evaluating switching inputs and need careful noise margin checks. The hysteresis receiver drives a subsequent flip-flop. The hysteresis circuit was designed to support bandwidths up to 6 Gbps, well above the target for the rest of the link. The transmitter and receiver, including series capacitors, consume 650µm² and 80µm² per bit, respectively.

A second series capacitor at the end of the wire separates the receiver and wire biases [7]. The hysteresis receiver needs a Vdd/2 input bias, because hysteresis uses contention between NMOS inputs and cross-coupled PMOS pull-ups. The series capacitor minimizes the load for this Vdd/2 bias, which is generated by an on-chip capacitive divider. In contrast to the receiver inputs, the wire requires a higher bias voltage because the series capacitors are all implemented with compact NMOS native devices that exhibit their largest capacitance at high gate voltages. Because the wires send RZ pulses, both differential lines remain at the same DC voltage when inactive; this allows a simple biasing of the differential wires through leaky PMOS transistors. Bias circuits for NRZ signaling are more complicated or impose DC-balanced data restrictions [1,2].

In the 3-tap TX filter, a rising transition generates a positive pulse followed by a negative pulse on the wire. If another data transition (in this case, a falling edge) immediately occurs, the trailing negative pulse of the current rising data bit will be adjacent to the leading negative pulse of the successive falling data bit (see Fig. 9.5.3). This suggests that if the two bits were partially overlapped, such that the two negative pulses coincided, the pre-emphasis would effectively double, causing the wire voltage slew rate to also double. Therefore, this circuit supports sending data on the wire at double the bandwidth, as in DDR links, with a new bit each clock phase: the increased pre-emphasis allows the wire to keep up with

the circuits by generating suitably sharper pulses. This doubling of the TX preemphasis, and hence wire performance, happens only when two transitions occur back-to-back. A data transition followed by constant data would not double the pre-emphasis, as the hysteresis receiver allows the wire response to be slower. In other words, the transmitter adaptively employs higher pre-emphasis only when needed, without any special encoding. This DDR scheme uses the same circuits as the SDR scheme, but adds a differential amplifier in front of the receiver to improve its performance.

We fabricated a 90nm CMOS testchip that included 3-bit 5mm links using conventional repeaters, a previous capacitive driver [1,2], and the proposed SDR and DDR schemes. The proposed and previous schemes both employ optimally twisted differential M5 wires with 0.28 μ m width and 0.28 μ m spacing (2× minimum pitch), while the conventional scheme is optimally repeated on single-ended M5 wires with 0.56 μ m width and 0.56 μ m spacing for the same footprint (representing an 11% delay increase with 21% lower energy over the optimal delay point in the width/spacing design space). M4 and M6 layers are filled with densely packed orthogonal interconnects. The wires were chosen to be 5mm to match "long" wire lengths most commonly found in high-performance systems running at 2.5 GHz. Longer wires would require a larger wire pitch to overcome series losses and are typically flopped in the system architecture. All capacitive driver schemes employed a 100mV differential swing at the end of the wire.

Pseudo-random binary sequence (PRBS) data is generated off-chip and directly sent to the on-chip test structures. The measured BER for both SDR and DDR schemes are less than 10⁻¹⁰. Energy versus performance characteristics are shown in Fig. 9.5.4. The proposed schemes improved the performance over prior approaches to 2.5 Gb/s (SDR) or 4.9 Gb/s (DDR), while achieving energy consumption of 0.24 pJ/b (SDR) or 0.34 pJ/b (DDR). The performance was limited by smaller-than-expected capacitance of the NMOS native devices, which resulted in reduced signal swing. As shown in Fig. 9.5.4, adaptive transmitter pre-emphasis in the DDR scheme provides ~2X bandwidth density improvement with 38% increase in energy consumption, due to additional clock energy and the differential amplifier. Measured energy per bit scales with data activity factors as seen in Fig. 9.5.5. Figure 9.5.6 shows the probed waveforms of the transmitter output and receiver input signals through on-chip samplers [8] for a 000010000 data pattern in the DDR scheme.

The proposed transceiver design for repeaterless on-chip communication demonstrates high bandwidth density, low latency, and low energy consumption. A 3-tap transmitter significantly reduces ISI over narrow wires, and a simple hysteresis receiver recovers the resulting low-swing RZ pulse. With DDR signaling the transmitter pre-emphasis is adaptively controlled, enabling a data rate of 4.9 Gb/s/ch and bandwidth density of 4.4 Gb/s/µm over 5mm on-chip links with 0.34 pJ/b energy consumption.

References:

[1] R. Ho, T. Ono, F. Liu, et al., "High-Speed and Low-Energy Capacitively-Driven On-Chip Wires," *ISSCC Dig. Tech. Papers*, pp. 412-413, Feb. 2007.

[2] E. Mensink, D. Schinkel, E. Klumperink, et al., "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-chip interconnects," *ISSCC Dig. Tech. Papers*, pp. 414-415, Feb. 2007.

[3] N. Tzartzanis, and W. Walker, "Differential Current-Mode Sensing for Efficient On-Chip Global Signaling," *IEEE J. Solid-State Circuits*, vol. 40, no. 11, pp. 2141-2147, Nov. 2005.

[4] B. Kim, and V. Stojanovic, "A 4Gb/s/ch 356fJ/b 10mm Equalized On-chip Interconnect with Nonlinear Charge-Injecting Transmit Filter and Transimpedance Receiver in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 66-67, Feb. 2009.

[5] D. Schinkel, E. Mensink, E. Klumperink, et al., "A 3-Gb/s/ch Transceiver for 10-mm Uninterrupted RC-limited Global On-Chip Interconnects," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 297-306, Jan. 2006.

[6] N. Miura, H. Ishikuro, T. Sakurai, and T. Kuroda, "A 0.14pJ/b Inductive-Coupling Inter-Chip Data Transceiver with Digitally-Controlled Precise Pulse Shaping," *ISSCC Dig. Tech. Papers*, pp. 358-359, Feb. 2007.

[7] J. Bae, J.-Y. Kim, and H.-J. Yoo, "0.6pJ/b 3Gb/s/ch Transceiver in 0.18µm CMOS for 10mm On-Chip Interconnects," *IEEE International Symposium on Circuits and Systems*, pp. 2861-2864, May 2008.

[8] R. Ho, B. Amrutur, K. Mai, et al., "Applications for On-Chip Samplers for Test and Measurement of Integrated Circuits," *Symp. VLSI Circuits*, pp. 138-139, Jun. 1998.





Figure 9.5.1: Conceptual bandwidth and latency benefits of proposed RZ signaling are shown using simulated waveforms. Such RZ signaling requires a multi-tap transmitter, and hence more transmit energy to achieve the improvements.



Figure 9.5.3: Adaptive pre-emphasis through DDR signaling, including simulation waveforms of intermediate points along the on-chip link.



Figure 9.5.2: Proposed transmitter and receiver circuits with waveforms when '001100' pattern is sent over on-chip links. Note that only 01 (rising) or 10 (falling) patterns generate pulses on the wire. Transceiver remains idle with consecutive 0s or 1s.







Figure 9.5.6: Measured waveforms of transmitter output and receiver input for a 000010000 data pattern in the proposed DDR scheme.

