A 4Gb/s/ch 356fJ/b 10mm equalized on-chip interconnect with nonlinear charge-injecting transmit filter and transimpedance receiver in 90nm CMOS

3.6 A 4Gb/s/ch 356fJ/b 10mm Equalized On-chip Interconnect with Nonlinear Charge-Injecting Transmit Filter and Transimpedance Receiver in 90nm CMOS

Byungsub Kim, Vladimir Stojanovic

Massachusetts Institute of Technology, Cambridge MA

This paper presents a transceiver for fast and energy-efficient global on-chip communication, consisting of a nonlinear charge-injecting (CI) 3-tap transmit filter (TX) and a sampling receiver (RX) with transimpedance pre-amplifier (TIA). Recently, pre-emphasis techniques [1-4] have demonstrated significantly better energy-efficiency than repeater interconnects. To further improve energy-efficiency over pre-emphasis techniques that require analog subtraction [3,5], our TX selects a pattern-dependent current to inject into the wire, performing feed-forward equalization (FFE) while mitigating the nonlinearity of the driver. This 3-tap charge-injecting (CI) FFE enables stronger equalization than the capacitively driven TX [1,2] or edge-detection pre-emphasis [4], achieving data-rate of 4Gb/s over a 1cm on-chip wire. The TIA at RX improves bandwidth, signal amplitude, and reduces bias power, breaking the trade-offs in conventional resistor termination [6], and mitigates equalized signal degradation due to impedance changes in dynamic current sensing [7].

Figure 3.6.1 depicts the overall link architecture with waveforms for single bit-flip pattern. Current sources at TX provide a small bias current for the TIA through signaling wires. The TX digitally computes and injects the FFE current (IT+ and IT-) into the wires. At RX, the received current (IR+ and IR-) is converted into a voltage (VS+ and VS-) by the TIA, and the slicers and DFE make a decision about the received bits.

The principle of CI FFE is shown in Fig. 3.6.2 through a conceptual comparison between the pull-down circuit of CI FFE and traditional analog FFE. For FFE coefficients [w_0 w_1 w_2] where only w_1 is negative (typical on a lossy channel) there are only 4 distinct values of IT+ for each direction (inward/outward): I_0, I_1,I_2, and I_0+I_1+I_2. Instead of having current sources with a size ratio of w_0:w_1:w_2, the current sources in CI FFE have a I_0:I_1:I_2 size ratio. Therefore, any value of IT+ can be created by *addition only*. This eliminates the current lost in analog FFE TX on subtraction. For example, when $D_0D_1D_2$='011', a CI FFE consumes only I_2=w_0-w_1-w_2 while an analog FFE consumes I_0+I_1+I_2=|w_0|+|w_1|+|w_2| regardless of the pattern $D_0D_1D_2$. For random data, this reduces the power by half compared to analog FFE [3], with larger savings for lower activity rates.

Figure 3.6.3 shows a detailed block diagram of TX. To improve energy-efficiency, TX is designed in latch-based double-data-rate (DDR) style. TX has a weak driver design of accuracy and and a strong driver designed for efficiency. Since I_0 is much smaller (~30μA) and more sensitive than other current values, a differential current switch is used as the weak driver for accuracy. A strong driver injects large currents (I_1, I_2, and I_1+I_2) on data transitions by driving the eight 5b PMOS or NMOS digital-to-analog converters (DACs) (P_1, P_2, N_1, N_2 for each terminal) rail-to-rail for most efficient charge delivery. Since DAC transistors are changing the region of operation as the output voltage swings from 0.5 to 0.9V (see Fig. 3.6.1), the DAC current and output impedance fluctuate. However, this nonlinearity is compensated by statically tuning each DAC segment (P_1, P_2, N_1, N_2).

Figure 3.6.4 shows a state-transition diagram of output-voltage change and driver control of TX. Since injected current IT+ depends on the transmitted data pattern, TX output voltage VT+ is determined by two consecutive data bits. For example, if $D_0D_1D_2$='011' is being transmitted, the state ('D_1D_2') changes from '11' to '01' (on a transition by D_0) by conducting -I_2 current via P_2^- and N_2^+ segment. At the same time, the drain voltage of P_2^- segment also changes from middle level to high level. Since each state transition has unique voltage-change profile and injected current value, assigning a DAC code for each transition compensates the nonlinear behavior. Except for the transitions between '01' and '10,' coefficients are statically assigned since a single designated DAC segment drives each transition. With small accuracy cost, the transitions between '01' and '10' are done by turning on all strong-driver DACs and a weak driver, since the current errors of P_1/N_1 and P_2/N_2 when all segments are turned-on compared to other individual transitions cancel each other. A decoding block is implemented with small number of simple gates to control the DACs.

A TIA pre-amplifier is used in RX to combine the benefits of voltage-mode (VM) and current-mode (CM) signaling. In VM, as the termination impedance decreases, the received voltage decreases and static current increases, but the bandwidth also increases [6], while in CM, both the signal current and bandwidth increase. The TIA termination at RX provides small input impedance without large static current. The TIA has 860Ω input impedance for high bandwidth (~350MHz) and high current amplitude (~7.5μA) at 2GHz, and the received current is converted into voltage V_{TIA} (~19mV) by a 2.7kΩ TIA gain in simulation. The RX with TIA roughly doubles the received voltage at half the static current compared to VM signaling with direct 860Ω termination. After the TIA, the received voltage is fed into a 1-tap loop-unrolled DDR DFE [8]. The DFE is designed based on threshold-controllable Strong-Arm slicers and latches to minimize power and area cost.

A test link over 1cm on-chip differential M8 wires (0.6μm wide, 0.4μm spaced) is implemented in 90nm 1.2V CMOS process. By sweeping threshold voltage and clock delay of RX slicers, the BER and signal probability distributions are measured *in situ* [8] at room temperature. Figure 3.6.5 shows the *in situ* measured eye diagram at 4Gb/s with 98mV vertical and 50% UI horizontal eye. A BER test showed no errors within the test-setup-limited window of ~10^6 bits. The large eye size indicates good supply-noise immunity. Without equalization, the eye is completely closed. Only 3b out of the possible 5b DACs are used when equalized, implying potential tuning cost reduction. Although the DFE is implemented in the receiver, it is not used at 4Gb/s since the FFE provides sufficient equalization. Figure 3.6.6 shows the energy breakdown and comparison with most recent related work [1]. The energy breakdown is extracted by multiplying the measured current with energy breakdown ratio from postlayout simulation. The total transceiver energy cost is 356fJ/b. The pre-driver energy of 85fJ/b could have been smaller (~30fJ/b) since the pre-driver was oversized to handle higher data-rates. Since the DFE is not used, the DFE overhead of 30fJ/b is not included. The test-chip results show that this link design can operate about 2× faster at similar pitch for ~30% higher energy cost compared to [1]. The measured eye is larger than the eye at 1.75Gb/s reported in [1], increasing link robustness. Figure 3.6.7 is a die micrograph of the transceiver. The TX area is 1120μm^2 and RX area is 640μm^2.

Acknowledgement:
We acknowledge the support of the Interconnect Focus Center, a Semiconductor Research Corporation Program and Trusted Foundry for chip fabrication, Intel Corporation and Center for Integrated Circuits and Systems at MIT.

References:
[1] E. Mensink, D. Schinkel, E. Klumperinck, et al., "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-chip interconnects," *ISSCC Dig. Tech. Papers, pp. 314-315, Dec.* 2007.
[2] R. Ho, I. Ono, F. Liu, et al., "High-Speed and Low-Energy Capacitive-Driven On-Chip Wires," *ISSCC Dig. Tech. Papers, pp. 412-413, Dec.* 2007.
[3] B. Kim and V. Stojanovi "Equalized Interconnect for On-Chip Networks: Modeling and Optimization Framework," *ICCAD*, pp. 552-559, Nov. 2007.
[4] L. Zhang, J. Wilson, R. Bashirullah, et al, "A 32Gb/s On-chip Bus with Driver Pre-emphasis Signaling," *CICC*, pp. 265-268, Sep. 2006.
[5] J.F. Bulzacchelli, M. Meghelli, S.V. Rylov, et al., "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology," *J. Solid-State Circuits, vol. 41, no. 12, pp. 2885-2900, Dec.* 2006.
[6] R. Bashirullah, W. Lieu, R. Cavi, D. Edwards, "A 16Gb/s Adaptive Bandwidth On-chip Bus based on Hybrid Current/voltage Mode Signaling," *Symp. VLSI Circuits,* pp. 392-393, Jun. 2004.
[7] N. Tzartzanis, W.W. Walker, "Differential Current-Mode Sensing for Efficient On-Chip Global Signaling," *J. Solid-State Circuits,* vol. 40, no. 11, pp. 2141-2147, Nov. 2005.
[8] V. Stojanovic, A. Ho, B. Garlepp, et al. "Adaptive Equalization and Data Recovery in Dual-Mode (PAM2/4) Serial Link Tranceiver," *Symp. VLSI Circuits,* pp. 348-351, Jun. 2004.

 978-1-4244-3457-2/09/$25.00

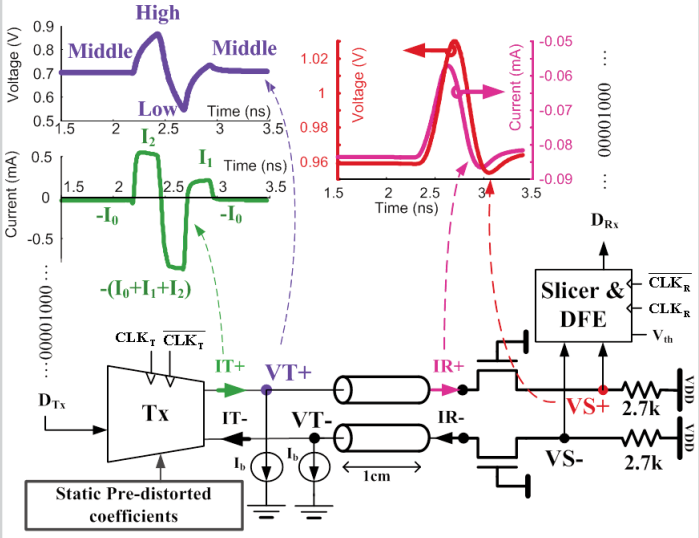

Figure 3.6.1: Overall link architecture with waveforms when '0001000' pattern is transmitted.

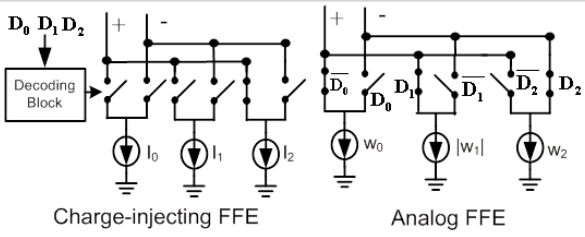

$D_0D_1D_2$	FFE sum	IT+
1 1 1	$w_0+w_1+w_2$	I_0
1 1 0	$w_0+w_1-w_2$	$-I_1$
1 0 1	$w_0-w_1+w_2$	$I_0+I_1+I_2$
1 0 0	$w_0-w_1-w_2$	I_2
0 1 1	$-w_0+w_1+w_2$	$-I_2$
0 1 0	$-w_0+w_1-w_2$	$-(I_0+I_1+I_2)$
0 0 1	$-w_0-w_1+w_2$	I_1
0 0 0	$-w_0-w_1-w_2$	$-I_0$

Figure 3.6.2: Conceptual comparison between the pull-down circuits of charge injection FFE and analog subtraction FFE when $D_0D_1D_2$='100'.

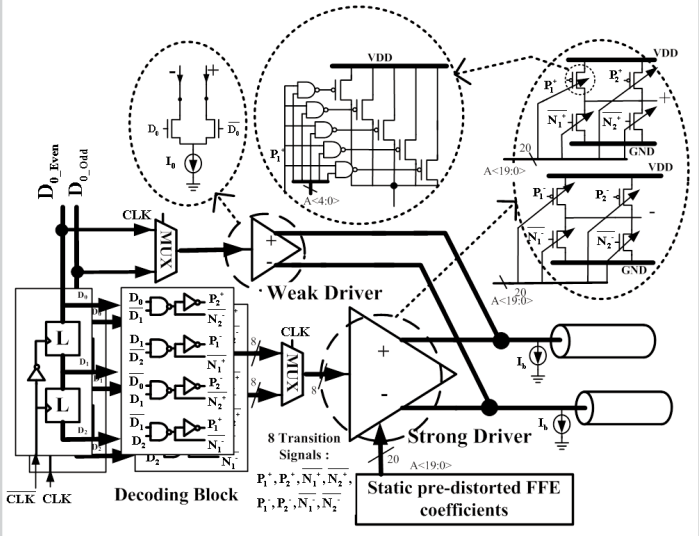

Figure 3.6.3: Block diagram of TX.

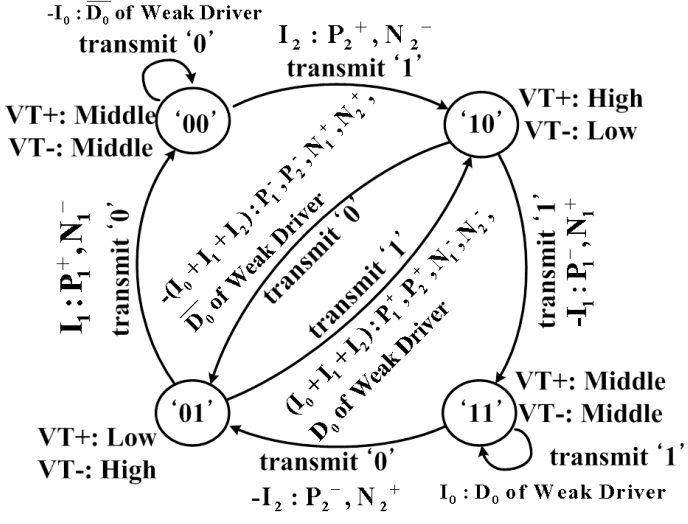

Figure 3.6.4: State transition diagram for TX control – 'D_1D_2' state with 'D_0' transition.

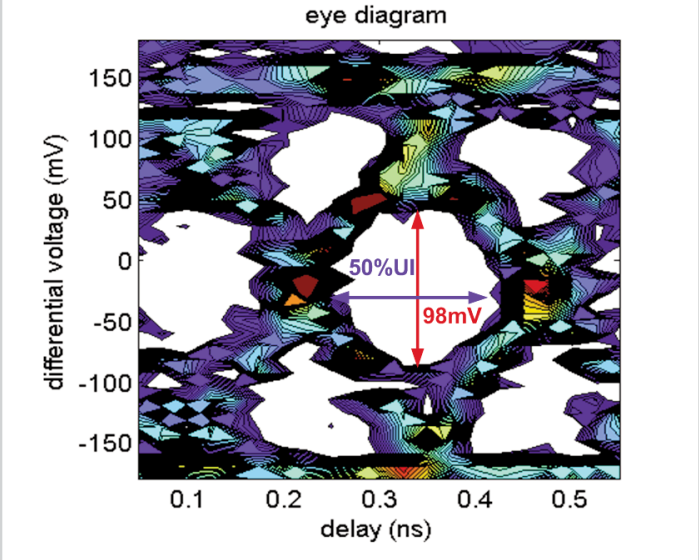

Figure 3.6.5: Measured eye diagram with FFE and without DFE equalization.

		This work		[1]
Data rate		4Gb/s		2Gb/s
Data rate density		2Gb/s/μm		1.1628Gb/s/μm
Channel length		1cm		1cm
Vertical Eye		98mV		N/A
Horizontal Eye		50%UI@(4Gb/s) measured on-chip		44%UI @ (1.75Gb/s) after 50Ohm buffer
Measured Energy Cost (fJ/b) at room temperature	total	356		280
	Tx	total	277	
		Strong driver	51	
		Weak driver + TIA bias	48	
		Decoder	24	
		Clock	67	
		Pre-driver	87	
	Rx	total	79	
		Slicer/Latch	30	
		Clock	49	
Technology		90nm		90nm

Figure 3.6.6: Comparison with recent work.

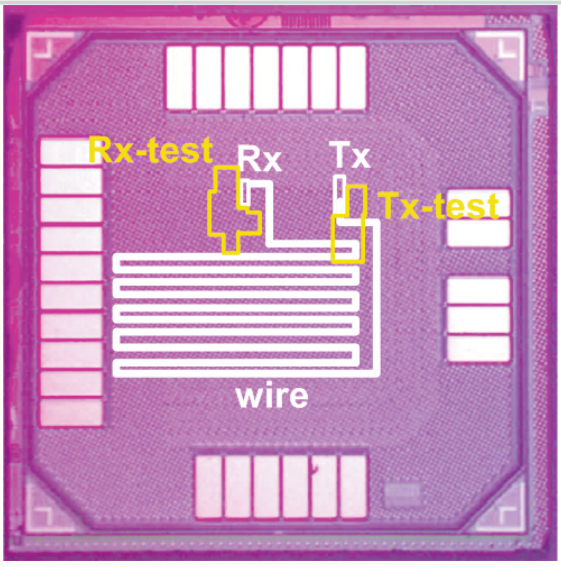

Figure 3.6.7: Micrograph of 1×1mm² die.

978-1-4244-3457-2/09/$25.00