

MIT Open Access Articles

*An energy-efficient equalized
transceiver for RC-dominant channels*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Kim, Byungsub, and Vladimir Stojanovic. "An Energy-Efficient Equalized Transceiver for RC-Dominant Channels." IEEE Journal of Solid-State Circuits 45.6 (2010): 1186–1197. © Copyright 2010 IEEE

As Published: <http://dx.doi.org/10.1109/JSSC.2010.2047458>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/72662>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



An Energy-Efficient Equalized Transceiver for RC-Dominant Channels

Byungsub Kim, *Student Member, IEEE*, and Vladimir Stojanović, *Member, IEEE*

Abstract—This work describes the architecture and circuit implementation of a high-data-rate, energy-efficient equalized transceiver for high-loss dispersive channels, such as RC-limited on-chip interconnects or silicon-carrier packaging modules. The charge-injection transmitter directly conducts pre-emphasis current from the supply into the channel, eliminating the power overhead of analog current subtraction in conventional transmit pre-emphasis, while significantly relaxing the driver coefficient accuracy requirements. The transmitter utilizes a power efficient non-linear driver by compensating non-linearity with pre-distorted equalization coefficients. A trans-impedance amplifier at the receiver achieves low static power consumption, large signal amplitude, and high bandwidth by mitigating limitations of purely-resistive termination. A test chip is fabricated in 90-nm bulk CMOS technology and tested over a 10-mm, 2- μm pitched on-chip differential wire. The transceiver consumes 0.37–0.63 pJ/b with 4–6 Gb/s/ch.

Index Terms—Equalized on-chip interconnect, RC-dominant wire, charge injection FFE, pre-distortion FFE, trans-impedance receiver, eye sensitivity.

I. INTRODUCTION

NETWORKS-ON-A-CHIP (NoCs) [1]–[3] are increasingly used in multi-core processors creating the need for fast, energy and area efficient global on-chip interconnects. However, the power inefficiency and latency of traditional repeated interconnects [3], [4] limit the performance gains of more advanced NoC architectures that need efficient global interconnections to realize their full potential [2], [3]. To overcome these repeater limitations, several techniques have been explored in the past [5], [6]. However, only recently [7]–[10], equalization at the transmitter (Tx) and receiver (Rx) over RC-dominant wires has been proposed to improve both the latency, energy and area-throughput efficiency.

An equalizing Tx flattens the link transfer function by suppressing the lower frequency portion of the channel response, eliminating the intersymbol interference (ISI). This allows

faster data transfer at lower power, since the suppression of low frequency decreases the voltage swing along the wire. In off-chip links, a feed forward equalization (FFE) Tx is typically implemented as an analog current or voltage summing/subtracting finite impulse response (FIR) filter [11]–[14], which consumes extra energy in addition to the signal energy injected into the wires. In on-chip link designs, pulsewidth pre-emphasis (PWP) [7] and capacitive peaking [8] reduce the complexity of equalization circuits but with limited throughput (bandwidth density) of only 2 Gb/s/ch (1 Gb/s/ μm) [8].

This paper reports a pre-distorted Charge-Injection (CI) FFE Tx and a trans-impedance-amplifier (TIA) Rx to improve the data rate and bandwidth density as well as energy and area efficiency. The CI FFE Tx eliminates the power wasted in analog subtraction of the conventional FFE by injecting digitally pre-computed FFE currents [10]. A full 3-tap FFE enables strong equalization on lossy channels and increases data rate up to 6 Gb/s/ch (3 Gb/s/ μm). The CI FFE also relieves the relative accuracy requirements for FFE coefficients. In addition, digital pre-distortion of CI FFE coefficients utilizes power-efficient nonlinear drivers. At the Rx , a TIA provides small input (termination) impedance while suppressing the static current, providing wide bandwidth and large received current amplitude, which is mapped to a large voltage at TIA output [10].

II. ON-CHIP INTERCONNECTS

Design of on-chip interconnects should be driven by system-level relevant metrics like energy efficiency (pJ/bit) and data rate density (i.e., data rate per wire pitch—Gb/s/ μm) [9]. Previous analysis [7], [9] indicates that RC-dominant, relatively narrow wires maximize these metrics leading to highest network throughput for given power and area constraints.

An RC-dominant channel requires a different signaling strategy than typical off-chip RLC transmission lines. In RC-dominant channels, 50-Ohm impedance matching is neither necessary nor efficient for two reasons: 1) the characteristic impedance of the wire is not 50 Ohm but can actually be co-designed with circuits to maximize relevant system-level metrics [9]; 2) the large channel loss (e.g., 40 dB and 46 dB at 2 GHz and 3 GHz for the 10-mm wire) suppresses the reflected wave from impedance mismatch [7], [9], [15]. The transfer function exponentially depends on wire length and square-root of frequency resulting in a time-response close to an exponentially decaying function, which is a pre-requisite for small-tap FFE implementation.

III. LINK OVERVIEW

Fig. 1 shows the block diagram of the proposed link. The Tx and Rx are connected through a 10 mm long differential wire.

Manuscript received November 23, 2009; revised February 22, 2010; accepted March 22, 2010. Current version published June 09, 2010. This paper was approved by Associate Editor Jafar Savoj. This work was supported by the Interconnect Focus Center, one of five research centers funded under the Focus Center Research Program, a DARPA and Semiconductor Research Corporation program, IBM and Trusted Foundry for chip fabrication, Intel Corporation and Center for Integrated Circuits and Systems at MIT.

B. Kim was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. He is currently with Intel Corporation, Hillsboro, OR 97124 USA (e-mail: byungsub@mit.edu).

V. Stojanović is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: vlada@mit.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2010.2047458




Fig. 1. A link overview.

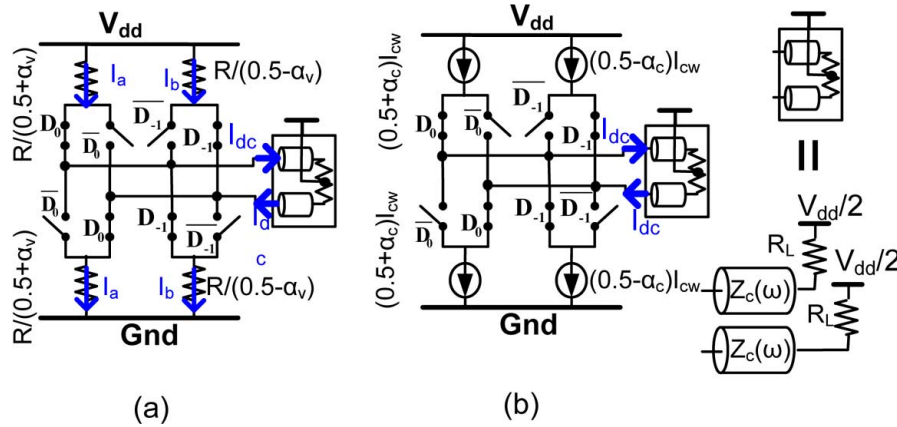


Fig. 2. Comparison between voltage dividing (VD) and current switching (CS) drivers. (a) VD. (b) CS.

The Rx is terminated with a TIA. Two current sources at the Tx provide bias current I_b for the TIA through the wire and set proper common mode voltage levels: V_T^+ , V_T^- , V_S^+ and V_S^- . During data transmission, the Tx computes and injects pre-emphasis currents I_T^+ and I_T^- into the wire. The TIA at the Rx converts the arriving currents I_R^+ and I_R^- into voltages V_S^+ and V_S^- , which are sampled by the decision feedback equalizer (DFE) module. The DFE extends the achievable data rate range by compensating the higher channel loss and mismatches from desired exponential impulse response roll-off.

IV. TRANSMITTER

A. Voltage Dividing Driver Versus Current Switching Driver

Before we explain the CI FFE driver, we introduce a current switching (CS) driver to more easily compare the CI FFE driver and a conventional voltage dividing (VD) driver [12]–[14], which is known to be more power efficient than a current mode logic (CML) driver [13]. The VD driver implements the FFE function via programmable resistive voltage divider while the CS driver adds/subtracts currents as shown in Fig. 2.

Fig. 3 shows the average supply currents (I_{vd} , I_{cs} and I_{ci} of the VD, CS, and CI drivers) versus the VD driver's output impedance R for 4 Gb/s data transmission. For fair comparison, all three drivers are matched for the same signal strength at Nyquist frequency f_N . The CI driver in this paper has the equivalent driving (channel) current of the VD driver with $R = 625$ Ohm, which is larger than the channel's characteristic




Fig. 3. Supply currents of VD, CS, and CI drivers for the same signal driving ability versus the VD driver's output impedance R . CS and CI drivers are matched for the same signal driving ability to the VD driver of given R in Fig. 2.

impedance $|Z_c(2\pi f_N)| \sim 160$ Ohm. In this $R \gg |Z_c(2\pi f_N)|$ region, I_{vd} , and I_{cs} converge, and VD and CS drivers burn $\approx 2x$ the power of a CI driver.

B. Charge-Injection FFE

Fig. 4 compares the CS FFE (a) and the proposed CI FFE (b) when the consecutive three-bit pattern $D_0D_{-1}D_{-2} = '011'$ is being transmitted, as marked in Table I. The CS driver computes

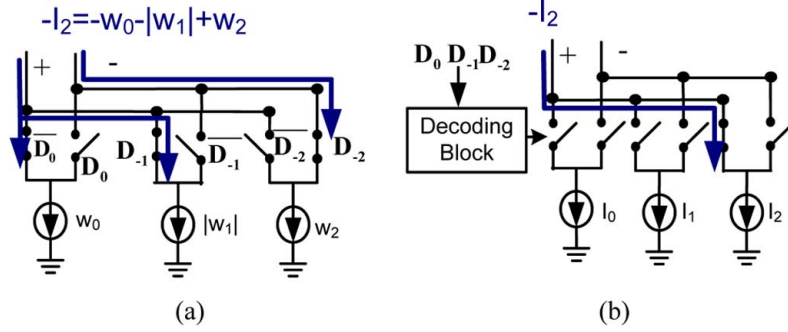


Fig. 4. Comparison between (a) a conventional Current-Switch FFE and (b) a Charge-Injection FFE when data pattern is $D_0 D_{-1} D_{-2} = '011'$.

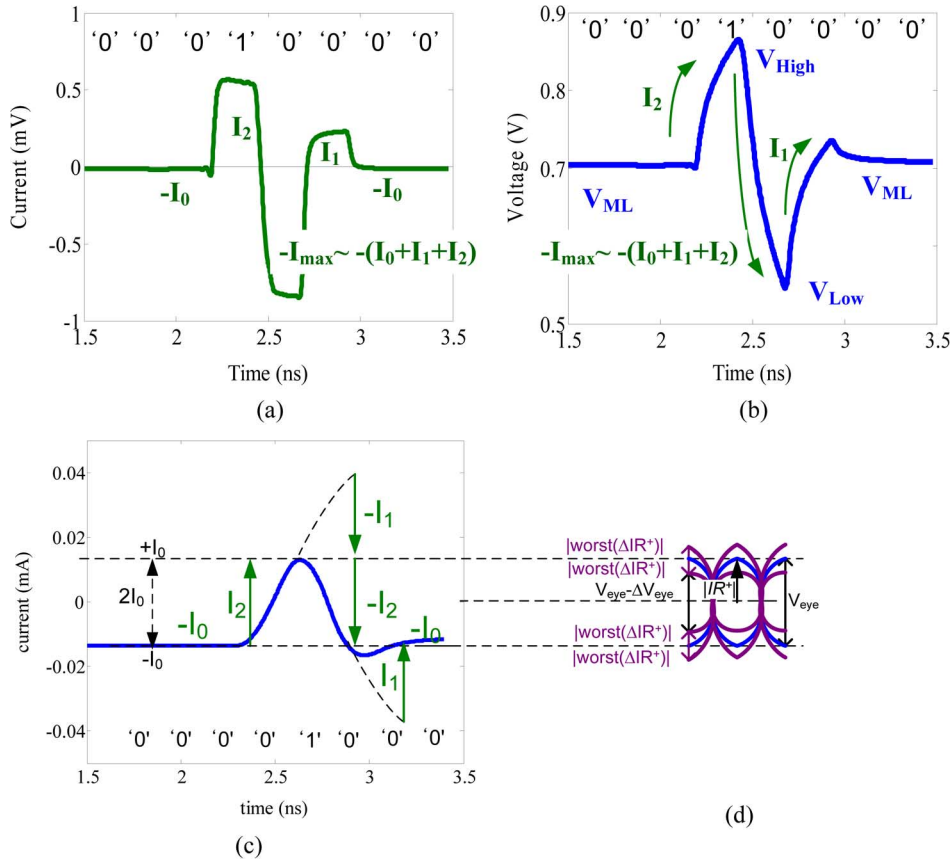


Fig. 5. Simulated (a) I_T^+ current, (b) V_T^+ voltage, and (c) I_R^+ current in Fig. 1 when an isolated '1' pattern is being transmitted at 4 Gb/s, and (d) illustration on eye reduction by I_R^+ current perturbation.

the FFE sum ($-I_2$ in Fig. 5(a)) by addition/subtraction of currents (w_0 , $|w_1|$ ($w_1 < 0$), and w_2) drawing more current from the supply than the current flowing into the channel. Our CI FFE drives the pre-computed $-I_2$ current directly into the channel. Note that this is similar to the unequalizing multi-level modulation drivers, e.g., [16]. Inherently this concept suffers from the exponential growth of driver segments with number of bits (taps) encoded in the output symbol. To prevent this exponential growth in our scheme, we combine the segments through addition only, maintaining linear growth in number of segments with number of taps. As a result, the CI FFE driver draws only the half CS FFE current, with same number of driver segments.

Table I presents this mapping from the 3-tap CS FFE sum I_T^+ to the corresponding 3-tap CI FFE currents for all data patterns

$D_0 D_{-1} D_{-2}$ without exponential complexity growth. Since CS FFE coefficients w_0 , w_1 , and w_2 can span I_T^+ list, another three positive variables I_0 , I_1 , and I_2 are able to span the same list by *addition only*, avoiding the power lost in current subtraction. Note that the list is symmetric with opposite polarities, and therefore, the CI FFE requires only three distinct positive currents (I_0 , I_1 , and I_2) since w_0 , $w_2 > 0$ and $w_1 < 0$ in a typical RC-dominant channel. In hardware implementation, I_0 , I_1 , and I_2 current sources can be connected to the channel independently for $\pm I_0$, $\pm I_1$, or $\pm I_2$ with proper polarity, or together for $\pm(I_0 + I_1 + I_2)$.

Figs. 5(a), (b), and (c) show the simulated waveforms of I_T^+ , V_T^+ , and I_R^+ with arrows illustrating impact of I_T^+ current values on I_R^+ , respectively, defined in Fig. 1. Table I also lists the cor-

TABLE I
CI-FFE MAPPING AND VOLTAGE TRANSITIONS

$D_0D_{-1}D_{-2}$	FFE sum	IT^+	VT^+ transition
1 1 1	$w_0+w_1+w_2$	I_0	$V_{MH} \rightarrow V_{MH}$
1 1 0	$w_0+w_1-w_2$	$-I_1$	$V_{High} \rightarrow V_{MH}$
1 0 1	$w_0-w_1+w_2$	$I_0+I_1+I_2$	$V_{Low} \rightarrow V_{High}$
1 0 0	$w_0-w_1-w_2$	I_2	$V_{ML} \rightarrow V_{High}$
0 1 1	$-w_0+w_1+w_2$	$-I_2$	$V_{MH} \rightarrow V_{Low}$
0 1 0	$-w_0+w_1-w_2$	$-(I_0+I_1+I_2)$	$V_{High} \rightarrow V_{Low}$
0 0 1	$-w_0-w_1+w_2$	I_1	$V_{Low} \rightarrow V_{ML}$
0 0 0	$-w_0-w_1-w_2$	$-I_0$	$V_{ML} \rightarrow V_{ML}$

responding V_T^+ transitions for $D_0D_{-1}D_{-2}$. While consecutive '0's are transmitted, the Tx draws I_0 current from the channel ($I_T^+ = -I_0$), and V_T^+ stays in the middle-low level voltage (V_{ML}). Since I_T^+ stays constant at $-I_0$, this current is not attenuated by the channel, and thus, I_R^+ stays at $-I_0$ which corresponds to bit '0' for the Rx . When an isolated '1' is transmitted, the Tx injects I_2 into the channel ($I_T^+ = I_2$) raising V_T^+ from V_{ML} to V_{High} and I_R^+ from $-I_0$ to I_0 . Although I_2 amplitude is much larger than $2I_0$, the impact of the abrupt I_T^+ change on the I_R^+ is attenuated to $2I_0$ by the high-frequency channel loss. On data transition from '1' to '0', I_T^+ changes to $-I_{max}$ (theoretically $-(I_0 + I_1 + I_2)$ but approximately $-(I_1 + I_2)$ since I_0 is much smaller than other currents) decreasing I_T^+ from V_{High} to V_{Low} . The role of $-I_{max}$ injection could be intuitively explained by a superposition of $-I_1$ and $-I_2$ injections. The $-I_1$ portion of $-I_{max}$ suppresses delayed overshoot (depicted as dashed curve in Fig. 5(c)) from previous I_2 during data transition from '0' to '1', keeping the I_R^+ at $+I_0$. The $-I_2$ portion of $-I_{max}$ current further pushes I_R^+ value down to $-I_0$, just as $+I_2$ previously raised I_R^+ from $-I_0$ to I_0 . To finish the transition, I_T^+ becomes I_1 causing V_T^+ transition from V_{Low} to V_{ML} at Tx and correcting the delayed undershoot caused by the previous $-I_2$ portion of $-I_{max}$ at Rx . To send consecutive '0's again after finishing transition, I_T^+ settles to $-I_0$ keeping I_R^+ at $-I_0$. In the next section we compare the power efficiency of CI FFE and CS FFE.

C. CI FFE Power Efficiency

The power of a CS FFE, P_{CS} , is independent of the data pattern and is calculated as

$$P_{CS} = V_{dd}(w_0 + |w_1| + w_2) = V_{dd}(I_0 + I_1 + I_2). \quad (1)$$

The average power of CI FFE for random data P_{CI} is the average of currents in Table I.

$$P_{CI} = \frac{V_{dd}}{2}(I_0 + I_1 + I_2) = \frac{P_{CS}}{2} \quad (2)$$

The CI FFE burns $2x$ less power than the CS FFE for the random data pattern. At lower link utilizations, CI FFE is even more power-efficient since it only draws large current on bit transitions, while the CS FFE always draws its peak current ($=I_0 + I_1 + I_2$). Table II(a) summarizes the linearized values of I_0 , I_1 , I_2 , $I_0 + I_1 + I_2$, and I_{max} (ideally $I_0 + I_1 + I_2$ but slightly different due to non-linearity effect) as well as the corresponding

TABLE II
SUMMARY OF CI FFE CURRENT VALUES AND THE CORRESPONDING FFE COEFFICIENTS

	(a) Current [μA]	(b) Relative ratio [%]	(c) Eye sensitivity	(d) Resolution requirement [%(bits)]	(e) Absolute accuracy [μA]
I_0	14	1.7%	1	10% (3-4b)	1.38
I_1	220	27%	0.79	12.68% (3-4b)	27.9
I_2	558	70%	2	5% (4-5b)	27.9
$I_0+I_1+I_2$	792	100%	N/A	N/A	N/A
I_{max}	836	105%	N/A	N/A	N/A
w_0	286	73.5%	20.7	0.484% (7-8b)	1.38
w_1	-389	-100%	28.2	0.355% (8-9b)	1.38
w_2	117	30.05%	8.47	1.18% (6-7b)	1.38

(a) Current values in μA , (b) relative ratio to the largest coefficient, (c) eye sensitivity, (d) approximate relative accuracy requirement in % (bits) for the design target of eye reduction $\beta = 10\%$, and (e) approximate absolute accuracy requirement in μA for the design target of eye reduction $\beta = 10\%$.

values of w_0 , w_1 , and w_2 used in simulation of Fig. 5. According to Table II(b), I_0 is less than 2% of $I_0 + I_1 + I_2$ due to the large channel attenuation. In an RC-dominant channel, the ratio of I_0 to $(I_0 + I_1 + I_2)$ is proportional to the channel loss at Nyquist frequency f_N as described in (3), which is derived from the first harmonic of the received current (a sinusoidal wave with amplitude I_0) when the Tx current is a square wave with amplitude $(I_0 + I_1 + I_2)$ transmitting the alternating bit pattern '...010101...'.

$$|T(2\pi f_N)| \approx \frac{\pi I_0}{4(I_0 + I_1 + I_2)} \quad (3)$$

The calculated channel loss at 2 GHz from Table II is about 37 dB showing consistency with measured and simulated channel transfer function in Fig. 11.

D. Resolution Requirements

FFE coefficient errors decrease the eye size, degrading the performance of a link. This eye reduction gets worse as the channel loss becomes larger, and thus often limits the performance of the conventional CS FFE. However, in CI FFE, the eye reduction is less dependent on the channel loss, and as a result, the eye is much less ($\sim 10x$) sensitive to the coefficient errors than in the traditional CS FFE. Therefore, at affordable lower coefficient resolutions (4-5 bits) the CI FFE circuits can equalize much higher channel loss than the corresponding CS FFE circuits.

To quantify the robustness of the FFE schemes to the coefficient errors, we define the received eye sensitivity S_x to a FFE coefficient x as the percentage of vertical eye reduction ($|\Delta E_{YE}/E_{YE}|$) divided by the percentage of coefficient perturbation ($|\Delta x/x|$) while other coefficients are fixed. Table II(c) lists the sensitivities to the CS and CI FFE coefficients.

$$S_{w_1} = \left| \frac{\frac{\Delta E_{YE}}{E_{YE}}}{\frac{\Delta w_1}{w_1}} \right|_{\Delta w_0, \Delta w_2=0} \approx \left| \frac{w_1}{I_0} \right| \approx \frac{\pi}{8|T(2\pi f_N)|} \quad (4)$$

Equation (4) is an approximate formula for the eye sensitivity to the critical CS FFE coefficient w_1 . Considering a DC data pattern (all '1's or all '0's), we can derive (4) from the following

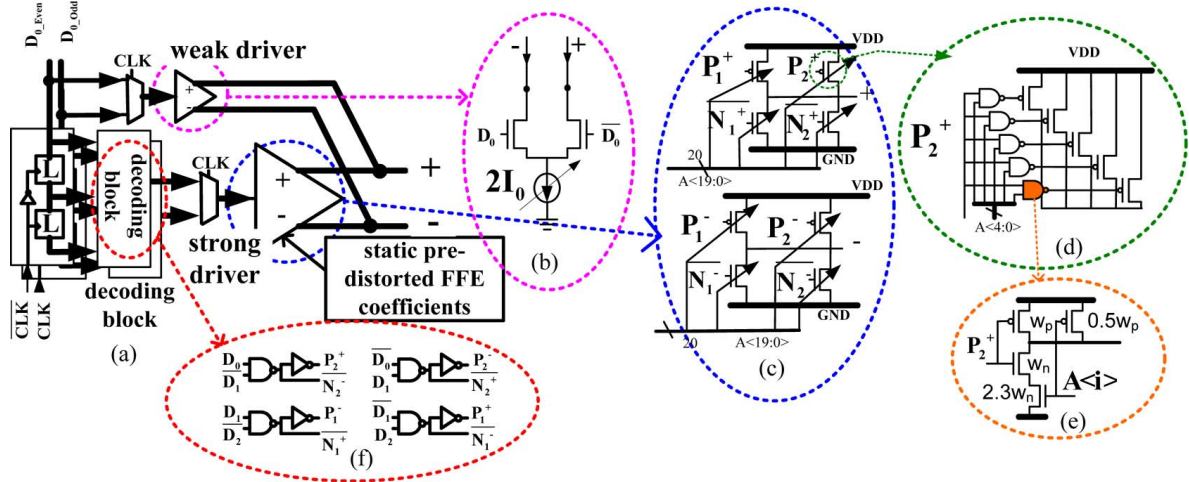


Fig. 6. CI FFE implementation: (a) architecture, (b) weak driver circuit, (c) strong driver circuit, (d) P_2 DAC transistor, (e) skewed NAND gate, and (f) decoding block.

rationales: 1) a good equalizer achieves $E_{YE} \approx 2I_0$ as depicted in Fig. 5(d); 2) the Rx current error is equal to the Tx current error Δw_1 since the channel does not attenuate DC signal; 3) $|\Delta E_{YE}|$ is approximately twice of the received current error as depicted in Fig. 5(d); 4) $w_1 \approx -(I_1 + I_2)/2 \approx -(I_0 + I_1 + I_2)/2$ from Table I using $I_0 \ll I_1 + I_2$; 5) channel transfer function in (3).

As shown in (4) and Table II, the eye sensitivity of the CS FFE is proportional to the channel loss $1/|T(2\pi f_N)|$ (40 dB at 2 GHz), and thus the eye of the CS FFE is highly sensitive to the coefficient error, requiring expensive high-resolution circuits. Since the small received I_0 (or $-I_0$) current for the DC pattern ‘...111111...’ (or ‘...000000...’) is generated by linear addition or subtraction of three large current values: w_0 , $|w_1|$, and w_2 , a small percentage error of w_1 is significantly large for the small received signal height $|I_0|$, greatly reducing the eye size. Therefore, the high eye sensitivity to the coefficient error is a limiting factor in a high-data-rate (i.e., large channel loss) CS FFE design.

In CI FFE, on the other hand, the eye sensitivities are much smaller and less affected by the channel loss since small I_0 is generated by a designated current source while other large current segments (I_1 and I_2) are turned off, instead of being generated from summation/subtraction of large current taps: w_0 , w_1 , and w_2 . This relaxes the eye sensitivity to I_0 as shown in (5), which is derived from the fact that the Tx current error is small ΔI_0 instead of large Δw_1 and the same rationales used to derive (4).

$$S_{I_0} = \left| \frac{\frac{\Delta E_{YE}}{E_{YE}}}{\frac{\Delta I_0}{I_0}} \right|_{\Delta I_1, \Delta I_2=0} \approx \frac{\frac{2|\Delta I_0|}{2I_0}}{\frac{\Delta I_0}{I_0}} = 1 \quad (5)$$

The CI FFE Tx also generates large current errors when large current taps I_1 and I_2 are active. However, these current sources turn on only for a bit-time at data transitions: ‘1’ \rightarrow ‘0’ or ‘0’ \rightarrow ‘1.’ Therefore, the large current errors ΔI_1 and ΔI_2 are modulated and attenuated by the channel by the factor of h_{peak} , the

peak value of the sampled channel response to a unit square pulse, relaxing the eye sensitivity as shown in (6).

$$S_{I_2} = \left| \frac{\frac{\Delta E_{YE}}{E_{YE}}}{\frac{\Delta I_2}{I_2}} \right|_{\Delta I_0, \Delta I_1=0} \approx \frac{I_2 h_{\text{peak}}}{I_0} \approx 2 \quad (6)$$

The high eye sensitivities require high resolution for the CS FFE. Table II(d) summarizes the required relative accuracy (resolution) of each current source to restrict the eye perturbation within a given design target β ($=10\%$). Note that, although the worst absolute eye accuracy requirement is the same for CS and CI FFE, the much higher resolution requirement for CS FFE makes the hardware cost of CS FFE significantly more expensive than the CI FFE. The most stringent accuracy constraints are 0.35% ($>8\text{b}$) for w_1 in CS FFE compared to 5% ($>4\text{b}$) for I_2 in CI FFE, indicating that the CI FFE relaxes the current source accuracy by more than $10\times$ compared to the CS FFE.

E. CI FFE Circuit

Fig. 6(a) describes the architecture of the CI FFE Tx . A latch-pipelined, double-data-rate (DDR) digital decoding block generates switching signals for driver segments. The driver consists of weak and strong segments to appropriately pull up and down non-transient (I_0) and transient (I_1 , I_2) currents, respectively.

The weak segment conducts small current I_0 . Although only 10% of the relative accuracy is required to bound the eye change by I_0 error within 10%, the absolute accuracy is high because the nominal value of I_0 is small ($\sim 60 \mu\text{A}$). Therefore, we use a current switch for the weak segment as shown in Fig. 6(b). The tail current source is $2I_0$ instead of I_0 to make $2I_0$ swing because the weak segment only pulls down and is unable to inject I_0 into the channel.

The design of the strong segment focuses on power-efficient current delivery since the transient currents (I_1 and I_2) have large amplitude and more relaxed accuracy constraints. The strong segment consists of four 5-bit digital-to-analog converter

(DAC) transistors (P_1 , N_1 , P_2 , and N_2) for each differential terminal as illustrated in Fig. 6(c). P_1 and N_1 generate I_1 , and P_2 and N_2 generate I_2 . The 5-bit accuracy on I_1 and I_2 theoretically bounds the received eye error to less than 10% because of the channel attenuation. The strong segment pull-up is necessary to keep a proper common mode voltage level for the driver because of the pull-down bias ($\sim 80 \mu\text{A}$) current for the TIA at the Rx . Without the strong driver pull-up, the common mode voltage becomes too low to keep the driver on. Each DAC transistor is an array of binary weighted transistors with enable signal as shown in Fig. 6(d). With DAC transistor gate nodes driven rail-to-rail, this topology delivers the maximum current for a given parasitic capacitance, achieving good power efficiency. For example, for the same transistor area, the topology like the weak segment delivers about $5\text{--}8\times$ smaller current than the strong driver to keep the tail transistor in saturation. The enabling NAND gate in Fig. 6(e) is skewed for fast response to signal input (P_2), improving the pre-driver energy efficiency. The enable pMOS is only half the size of the signal pMOS minimizing the loading on the signal path while being strong enough to keep the output voltage at V_{dd} when disabled.

However, the strong CI FFE driver behaves nonlinearly due to output impedance change. The sources of the impedance change are: 1) data-dependent switching of driver segments with different impedances (from bit-time to bit-time) and 2) segment output impedance fluctuation due to output voltage change in a strong driver segment (within a bit-time). For example, the output impedance is set by the current source of the weak segment while Tx current is I_0 , but the output impedance becomes small and a function of the drain voltage of P_2 when the P_2 transistor is conducting I_2 current. In our simulation, this nonlinear behavior reduces the current by 13% on average, by 27% at maximum, introducing additional degradation in signal quality. According to Table II, the eye reduction is more than 300% in CS FFE due to the high eye sensitivity (> 28), completely closing the eye. Therefore, we cannot use this nonlinear driver in CS FFE. However, in CI FFE, the eye is weakly sensitive to the coefficient errors (~ 2.5), resulting in 32.5% eye degradation. In CI FFE, this eye reduction can be further relieved by compensating the nonlinearity with static pre-distortion.

F. Pre-Distortion in CI FFE

In a 3-tap CI FFE, a three consecutive bit pattern $D_0D_{-1}D_{-2}$ determines the proper current source, I_T^+ value, and V_T^+ transition as listed in Table I. Since $D_0D_{-1}D_{-2}$ determines I_T^+ and the two sources of the nonlinear error, it also determines the magnitude of the nonlinear error. Therefore, we can correct all nonlinear errors by assigning the digitally compensated currents for all eight cases generated by the three binary combinations. However, in this design, by allowing a small nonlinearity error in $\pm(I_0 + I_1 + I_2)$ case, we reduce the cost of the compensation by statically tuning the three current sources (I_0 , I_1 , and I_2) to cover only six cases.

Except the two cases when $D_0D_{-1}D_{-2} = '101'$ and $'010'$, each segment associated with a CI FFE current (I_0 , I_1 , or I_2) has a unique voltage profile and thus a unique amount of nonlinear error. Therefore, we can compensate the error by statically

tuning each segment. For example, except in two cases, the weak segment, which conducts $+I_0$ or $-I_0$, only turns on when the output voltage is V_{ML} or V_{MH} to transmit $D_0D_{-1}D_{-2} = '111'$ or $'000'$. The difference between V_{ML} and V_{MH} is very small since it is set by the I_0 current. Since the voltage level is in the middle of the supply level, the current source of the weak segment operates in saturation and thus the current error is negligible. When $D_0D_{-1}D_{-2} = '100'$, P_2^+ turns on and conducts I_2 causing V_T^+ voltage transition from V_{ML} to V_{High} . The V_T^+ voltage change weakens the P_2^+ DAC's current and causing a current error in I_2 . Therefore, the static adjustment on P_2^+ strength is enough to compensate this nonlinear error.

The nonlinear errors of the two special cases are small. When $D_0D_{-1}D_{-2} = '010'$, the voltage changes from V_{High} to V_{Low} by conducting a current through the weak driver, N_1^+ , and N_2^+ (P_1^- and P_2^- for the other differential terminal). N_1^+ transistor for this pattern is weaker than for $D_0D_{-1}D_{-2} = '110'$ because the drain voltage at the end of transition is V_{Low} , which is lower than V_{MH} for $D_0D_{-1}D_{-2} = '110'$. N_2^+ transistor, on the other hand, is stronger for this pattern than for $D_0D_{-1}D_{-2} = '011'$ because the start-transition voltage V_{High} is higher than V_{MH} for $D_0D_{-1}D_{-2} = '011'$. In this case, the errors on I_1 and I_2 have opposite polarities, mostly canceling each other. Note that I_0 current is too small to add significant error in this case. The nonlinear error of the other case $D_0D_{-1}D_{-2} = '101'$ is mitigated in the same manner.

Fig. 6(f) shows a simplified circuit implementation of the decoding block to select the driver segment as listed in Table I. Since the Tx compensates nonlinear error by statically tuning the CI FFE coefficients (i.e., strengths of segments), the high speed decoder logic does not carry any coefficient information and can be implemented with very simple logic gates.

The partitioning of driver segments in CI FFE allows compact and hardware-efficient static pre-distortion, which is not possible in CS and VD FFE. In CS FFE, three current sources work together to generate all eight current values listed in Table I. As a result, each current source is turned on all the time and experiences all cases of voltage change. Since the nonlinear error is associated with voltage change, each current source has more than two distinct values of nonlinear error, preventing static compensation. For example, the pull-down w_2 current source in Table I is connected to $'+$ ' node when $D_0D_{-1}D_{-2}$ is $'110'$, $'100'$, $'010'$, or $'000'$. The four patterns cause four distinct V_T^+ transitions, respectively: $V_{High} \rightarrow V_{MH}$, $V_{ML} \rightarrow V_{High}$, $V_{High} \rightarrow V_{Low}$, and $V_{ML} \rightarrow V_{ML}$. Therefore, four different coefficients are necessary for w_2 since the non-linear error can have four distinct values for each case. The hardware implementation is much more difficult since the pre-distortion requires memory to store four different values of each coefficient, and a decoding block must select and assign the right coefficient value within a bit time.

Fig. 7 shows simulated eye diagrams with and without pre-distortion in CI FFE. We calculated coefficients in Fig. 7(a) assuming a fixed drain voltage of DAC transistors while we pre-distorted the coefficient values in Fig. 7(b). The non-pre-distorted eye is about 36% smaller than the pre-distorted one, confirming the analysis in Section IV-E.

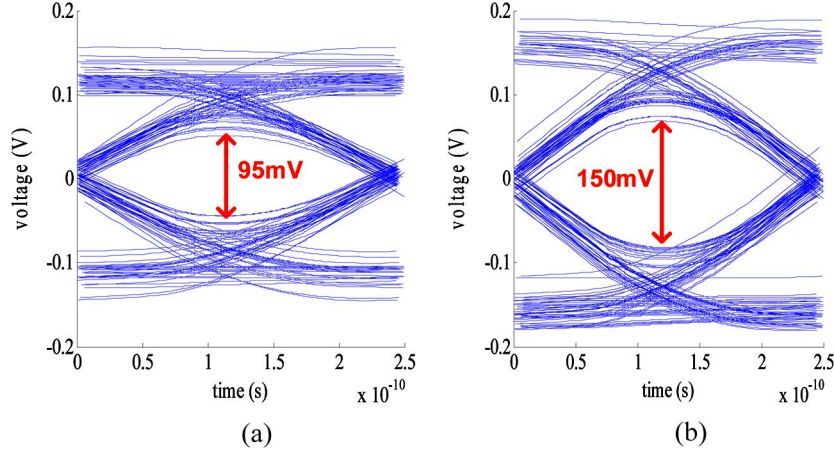


Fig. 7. Eye simulations at 4 Gb/s with 1 V supply: (a) before pre-distortion, (b) after pre-distortion.




Fig. 8. Various metrics versus receiver's termination impedance. Circle markers represent TIA termination. Square markers represent resistive termination for the same bandwidth of TIA.

V. RECEIVER

A. Receiver Termination

The R_x impedance affects the channel transfer function in an RC-dominant channel differently than in an LC-dominant channel. Fig. 8 shows various metrics versus the termination impedance R_L while keeping sufficient voltage swing at T_x . The voltage-mode receiver is usually terminated with a large resistor R_L to achieve large amplitude with small static current. As R_L increases, the voltage amplitude of the received signal increases, while static current decreases reducing static power. However, the 3 dB bandwidth of the channel also decreases requiring increased equalization effort, and hence increased transmit power. By terminating the R_x with a small resistor (current-mode R_x), the received signal current as well as the 3 dB bandwidth increase as shown in Fig. 8. However, the cost of smaller input impedance is the larger static current indicating a trade-off between bandwidth, amplitude, and static power.

We propose adding a TIA to the R_x to change this fundamental trade-off between voltage-mode and current-mode signaling in RC-dominant channels by mitigating the dependence of the small signal gain on the input impedance (and static bias current). While the common-gate TIA is used in [17] to match the 50-Ohm transmission line in an off-chip link, we utilize the TIA in on-chip RC-dominant channel to adjust the termination impedance for best power-efficiency (not impedance matching) while maintaining the link bandwidth. The TIA in Fig. 1 provides small signal input resistance (≈ 860 Ohm) to the channel but requires $2x$ smaller static current (160 μ A) than the resistive termination, while providing the same bandwidth (54 MHz). After current-to-voltage conversion by the TIA, the converted voltage amplitude V_{TIA} is about $3x$ higher than the received voltage with the same resistive termination. Therefore, the TIA can achieve $3x$ higher signal amplitude (which decreases transmit dynamic power) and $2x$ smaller static power for the same bandwidth compared to the resistive termination. This benefit scales up with decrease in the TIA input impedance and increase in the TIA gain.

B. Receiver Circuit

Fig. 9 describes the R_x circuits. The TIA amplifies and converts the input current into voltage on which the following DFE decides the received bit. We implemented the loop-unrolling DFE as a latch-based design to further save power and area [18]. In this design, the selection signals of the MUXs are delayed by one additional differential-input latch stage to further relax the latency requirement by ensuring that the MUXs always take rail-to-rail selection signals. In a regular loop-unrolled DFE, if the input of the sense-amplifier becomes small due to noise, the output signal may not be fully regenerated within a bit time, failing the MUX feedback. In this design, the additional latch helps the partially-regenerated output of the sense amplifier to fully regenerate.

Modified StrongArm sense amplifiers with additional offset compensation ports are used to add or subtract the post cursor to the TIA output by setting the offset/threshold current. A tail current source attached to the output node of the sense amplifier controls the threshold voltage.

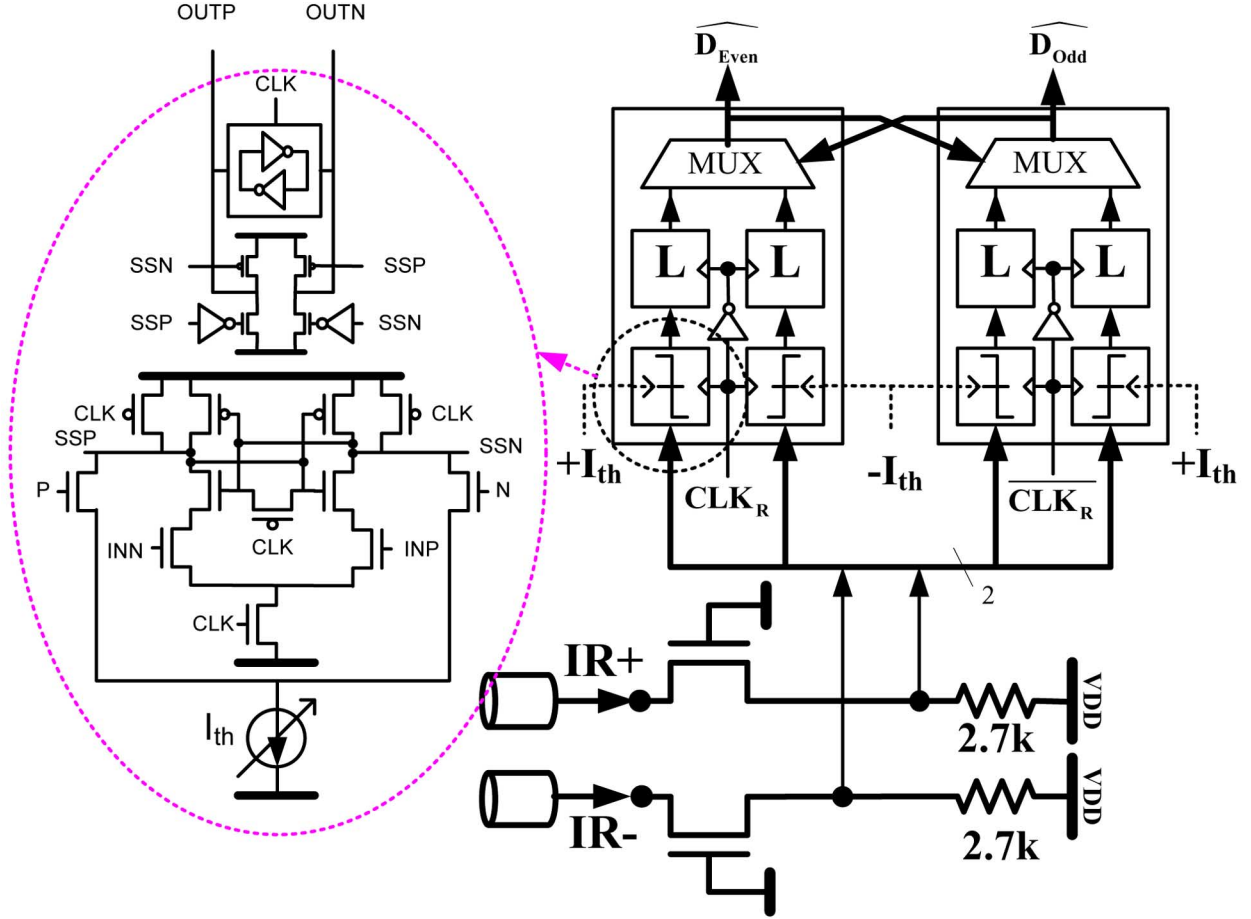


Fig. 9. Receiver circuit.


VI. EXPERIMENT

A. Chip Fabrication

A proof-of-concept chip in 90 nm bulk CMOS process has been fabricated and tested with on-chip test support blocks to measure the link *in-situ*. Fig. 10 shows the die photo overlaid with layout to outline the transceiver and test-support blocks. The channel is a 10-mm-long serpentine differential wire in M8. The wire width and space are $0.6 \mu\text{m}$ and $0.4 \mu\text{m}$, respectively. M7 and M9 layers are filled with supply grid and dummy metal. The transmitter and the receiver areas are $16 \mu\text{m} \times 70 \mu\text{m}$ and $16 \mu\text{m} \times 40 \mu\text{m}$, respectively.

B. On-Chip Test Support Block

Fig. 11 illustrates the block diagram of the on-chip test-support circuits. During test, two pattern generators feed the Tx with a test bit sequence: two pseudo random bit sequences (PRBS) with 31 bit seeds; 64-bit fixed pattern. The two 36-bit snapshot units monitor the received bit sequence sent by Tx . By comparing the transmitted and received patterns at different Tx and Rx clock phases and Rx threshold voltages, we generate the *in-situ* statistical eye diagram and channel pulse response as seen by the Rx . Except the Tx and Rx clocks, all digital control/monitoring is done by the scan chain through slow I/Os

Fig. 10. Test-chip die photo: $1 \text{ mm} \times 1 \text{ mm}$ (Tx : $70 \mu\text{m} \times 16 \mu\text{m}$, Rx : $40 \mu\text{m} \times 16 \mu\text{m}$).

to reduce cost of high-speed I/Os. External analog DC reference current input configures the bias currents of the weak segment and the TIA at Tx , and the thresholds of slicers at the Rx .

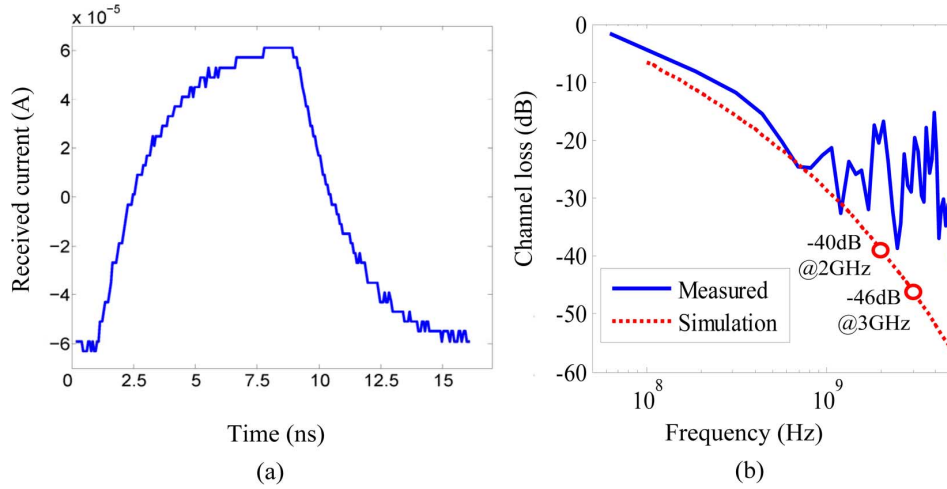


Fig. 12. Measured (a) step response and (b) transfer function. Solid: measured. Dotted: SPICE simulation using an *RLGC* model extracted from 2-D field solver.

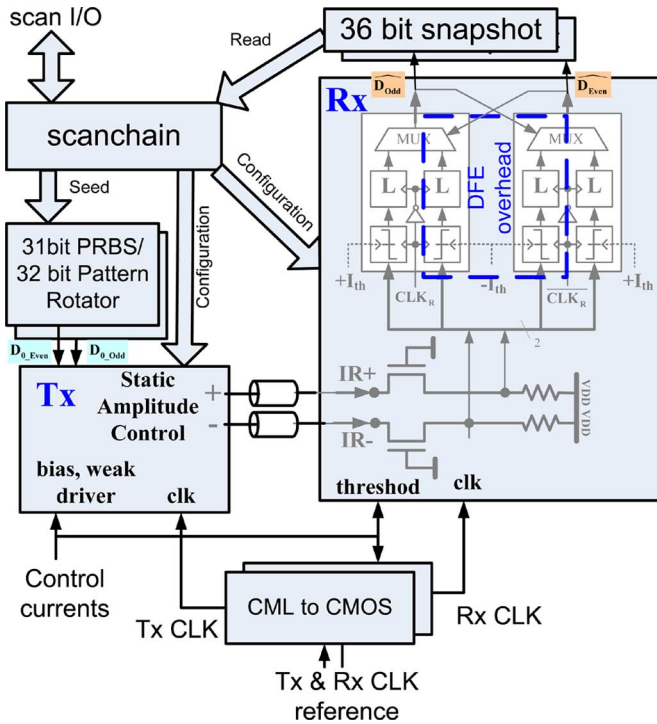


Fig. 11. Test support blocks.

C. Channel Measurement

Fig. 12 shows the on-chip measured step response and current-driving and current-receiving transfer function while only the weak segment drives the wire with a step pattern at 4 Gb/s: 32 consecutive '1's followed by 32 consecutive '0's. The received current is measured *in-situ* by finding the threshold of the *Rx* slicers to get 50% of received '1' at each 62.5 ps-spaced time point. The transfer function is calculated from the measured step response. The high-frequency noise and sampler dither during step estimation are amplified in dividing the spectrum of the step response by a *sinc* function (the spectrum of the square pulse) because the *sinc* function has small amplitude at high frequencies.

The long tail of the step response in Fig. 12(a) reveals significant ISI for an unequalized channel. The measured 50% delay and 90% settling times are about 1.4 ns and 5.5 ns, respectively, or 8.6 and 33 UI at 6 Gb/s/ch. Fig. 12(b) show the measured and simulated transfer functions. Although the noise dominates the signal at high frequencies, due to large signal attenuation and the noise amplification in the conversion of the step response to the transfer function, the noise is still 15 dB smaller than the DC level. The measurement and simulation shows reasonably good match within measured frequency range. The measured channel loss at 690 MHz is about 25 dB, implying much higher losses at 2 GHz and 3 GHz, which are 40 dB and 46 dB, respectively, in simulation. These high losses show the good robustness of the CI FFE compared to other on-chip links [6]–[8]. Additionally, in off-chip links with conventional FFE [11], [12], [15] losses up to 30 dB are typically considered equalizable.

D. Eye Diagram

Fig. 13 presents the measured eye diagrams to achieve close to 100 mV vertical eye to acquire power-performance trade-off. The CI FFE coefficients are calibrated and pre-distorted by monitoring the isolated pulse response similar to the simulation in Fig. 5(c). The eye diagram in Fig. 13(a) is measured at 6 Gb/s with 1.2 V supply voltage. During the measurement, the DFE was fully functional. At this data rate, the eye was closed without DFE since the channel response at this speed requires more than 3 taps to equalize. The measured eye height and width are 87 mV and 60% UI, respectively. Due to the scan-interface speed limitations, the probability of each of the voltage-time points on the eye diagram was collected from 10^4 transmissions. The good quality of horizontal eye opening (60% UI) in comparison with a typical bathtub curve in other works [11] implies that the link would have much lower bit error rate. The eye in Fig. 13(b) is measured at 4 Gb/s with 1.1 V supply, with disabled DFE. The DFE could improve the eye, but we are better off saving the DFE power overhead instead. The measured vertical eye was 109 mV, and the horizontal eye was 80% UI (BER < 10^{-4}). The eye in Fig. 13(c) is measured at 2 Gb/s with 1.1 V supply and disabled DFE. At 2 Gb/s, we also disabled one CI FFE tap




Fig. 13. Measured eye diagram. (a) 6 Gb/s DFE enabled 1.2 V. (b) 4 Gbps DFE disabled 1.1 V. (c) 2 Gbps DFE disabled 1.1 V.




Fig. 14. Vertical eye versus strong driver coefficient change. (a) 6 Gb/s DFE enabled 1.2 V Str(9, 10, 10, 10)*. (b) 4 Gbps DFE disabled 1.1 V Str(9, 5, 9, 5)*. (c) 2 Gbps DFE disabled 1.1 V Str(2, 2, 2, 2)*. *Nominal value of 5-bit strong driver coefficients, Str(P_1, P_2, N_1, N_2).

(N_1 and P_1 strong segments) to demonstrate that further hardware cost reduction is possible at low data rate operation. The vertical eye is 120 mV and horizontal eye is 60% UI.

E. Eye Sensitivity

Fig. 14 shows the vertical eye versus perturbation of the strong segment coefficient measured at 6 Gb/s, 4 Gb/s, and 2 Gb/s with 1.2 V, 1.1 V, and 1.1 V supply, respectively. To capture the eye sensitivity to each coefficient, the eye measurements were taken for perturbations in each coefficient. To shorten the test time, the eye-measurement statistics were taken down to 10^{-3} probability, which is sufficient to capture the eye sensitivities.

Fig. 14(a) confirms the small eye sensitivity to strong segment coefficients. At 6 Gb/s, the vertical eye changes by about 30% at most for 10% coefficient change. Fig. 14(c) illustrates higher sensitivity to I_2 (generated from P_2 and N_2) than I_1 (generated from P_1 and N_1), as discussed in Section IV-C, while Figs. 14(a) and (b) do not show clear evidence of higher sensitivities to I_2 due to approximation errors and non-linearity effects. At 4 Gb/s, the peak sensitivity is about 16% for 20% coefficient change. At 2 Gb/s, the eye is not very sensitive to P_1 and N_1 showing that the channel can be equalized only with 2-taps. Since P_2 and N_2 is the main tap, the eye still shows stronger dependency on P_2 and N_2 .

F. Power Consumption

Fig. 15 presents the measured link energy/bit breakdown at different conditions. The measurements show that operation at 4 Gb/s with 1 V supply is the most energy efficient. The energy cost is relatively flat up to 4 Gb/s because DC energy, switching energy, and channel related energy change differently as data rate increases. For example, the TIA draws DC current so its energy per bit decreases as the data rate increases. Rx energy-cost stays relatively flat up to 4 Gb/s following the CV^2 rule and doubles at 6 Gb/s due to the additional DFE overhead. Tx energy cost, especially the strong driver energy, increases with data rate increase. Since the channel loss becomes larger at higher data rate, the driver must be configured stronger to inject more energy into the channel for loss compensation. For high performance, the result shows that an increase in data rate from 4 Gb/s to 6 Gb/s requires approximately 70% more energy. For lower and fixed data-rate target, the link might be further resized for lower energy cost.

Fig. 16 shows the energy cost versus data rate density plot compared to the most efficient previously reported works [6], [8]. The closest performance and efficiency is reported in [8]. Compared to [8], the maximum achievable data-rate is improved by 3x to 3 Gb/s/ μm (6 Gb/s/ch) with up to 2x energy cost. This is the only on-chip link design to date that asserts the eye




Fig. 16. Comparison to the most relevant works done over a 10-mm link in 90-nm CMOS process. * 5 mm link.




Fig. 15. Measured energy breakdown at different data rates*. *TxOther*: Tx decoder and clock energy. *TxStr*: strong driver energy. *Rx*: Sense amplifier and DFE logic energy. *TIA*: TIA bias energy. *Misc*: energy not included in above list. The decoupling capacitor's leakage currents were found by simulation (72 μ A, 120 μ A, and 192 μ A at 1 V, 1.1 V, and 1.2 V, respectively) and subtracted from the measurements.

quality *in-situ*. The eye quality stays above 60% UI for all operations (maximum 80% UI), which is larger than 44% UI of [8]. Compared to the optimized repeater power consumption [9], the equalized interconnects burn about 4 \times less energy.

VII. CONCLUSION

We report a pre-distorted charge-injection FFE transmitter and a TIA-terminated receiver for RC-dominant on-chip interconnects. The charge-injection FFE consumes 2 \times less power and relaxes the coefficient accuracy requirement by 10 \times compared to the conventional voltage divider FFE architecture. The static pre-distortion technique utilizes a power-efficient nonlinear driver for equalization to further improve the power efficiency. At the receiver end, a TIA-termination is implemented, simultaneously achieving wide bandwidth, high amplitude, and small static power, by decoupling the input small signal resistance from the output transimpedance gain.

Measurements indicate operation up to 6 Gb/s (3 Gb/s/ μ m) data rates at channel losses up to 46 dB with energy cost around 0.63 pJ/b, and 0.37 pJ/b at 4 Gb/s. The eye is measured *in situ*. The eye sensitivity tests illustrate significantly relaxed coefficient

accuracy requirements when compared to the traditional analog FFE by leveraging channel attenuation to minimize the eye reduction due to driver inaccuracy.

The proposed link architecture and circuit techniques are not only applicable to on-chip interconnects [9], [10] but also to other RC-dominant channels such as narrow PCB wires and emerging silicon-carrier based packages [19], [20].


ACKNOWLEDGMENT

The authors thank Fred Chen and Sanquan Song for valuable help, and also thank Ian Young and Alexandra Kern for discussion.

REFERENCES

- [1] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Borkar, "An 80 Tile 1.28TFLOPS network-on-chip in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, vol. 589, pp. 98–99.
- [2] J. Kim, J. Balfour, and W. J. Dally, "Flattened butterfly topology for on-chip networks," *IEEE Computer Architecture Lett.*, vol. 6, no. 2, pp. 37–40, Jul. 2007.
- [3] A. Joshi, B. Kim, and V. Stojanović, "Designing energy-efficient low-diameter on-chip networks with equalized interconnects," in *Proc. 17th IEEE Symp. High-Performance Interconnects*, Aug. 2009, pp. 3–12.
- [4] H. B. Bakoglu and J. D. Meindl, "Optimal interconnection circuits for VLSI," *IEEE Trans. Electron Devices*, vol. ED-32, no. 5, pp. 903–909, May 1985.
- [5] A. P. Jose and K. L. Shepard, "Distributed loss-compensation techniques for energy-efficient low-latency on-chip communication," *IEEE J. Solid-State Circuits*, vol. 42, no. 6, pp. 1415–1424, Jun. 2007.
- [6] S. Tam, E. Socher, A. Wong, and M. F. Chang, "Simultaneous tri-band on-chip RF-interconnect for future network-on-chip," in *IEEE Int. Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2009, pp. 90–91.
- [7] D. Schinkel, E. Mensink, E. A. M. Klumperink, E. Tuijl, and B. Nauta, "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.
- [8] E. Mensink, D. Schinkel, E. Klumperink, E. Tuijl, and B. Nauta, "A 0.28 pJ/b 2 Gb/s/ch transceiver in 90 nm CMOS for 10 mm on-chip interconnects," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, vol. 612, pp. 414–415.
- [9] B. Kim and V. Stojanović, "Equalized interconnect for on-chip networks: Modeling and optimization framework," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design*, 2007, pp. 552–559.
- [10] B. Kim and V. Stojanović, "A 4 Gb/s/ch 356 fJ/b 10 mm equalized on-chip interconnect with nonlinear charge-injecting transmitter filter and transimpedance receiver in 90 nm CMOS technology," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2009, vol. 978, pp. 66–67.

- [11] J. F. Bulzacchelli, M. Meghelli, S. V. Rylov, W. Rhee, A. V. Rylyakov, H. A. Ainspan, B. D. Parker, M. P. Beakes, A. Chung, T. J. Beukema, P. K. Pepljugoski, L. Shan, Y. H. Kwark, S. Gowda, and D. J. Friedman, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, Dec. 2006.
- [12] H. Hatamkhani, K. J. Wong, R. Drost, and C. K. Yang, "A 10-mW 3.6-Gbps I/O transmitter," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, 2003, pp. 97–98.
- [13] H. Hatamkhani and C. K. Yang, "Power analysis for high-speed I/O transmitters," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 142–145.
- [14] C. Menolfi, T. Toifl, P. Buchmann, M. Kossel, T. Morf, J. Weiss, and M. Schmatz, "A 16 Gb/s source-series terminated transmitter in 65 nm CMOS SOI," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2007, vol. 614, pp. 446–447.
- [15] Y. Liu, B. Kim, T. Dickson, J. Bulzacchelli, and D. Friedman, "A 10 Gb/s compact low-power serial I/O with DFE-IIR equalization in 65 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2009, pp. 182–183.
- [16] K. Farzan and D. A. Johns, "A CMOS 10-Gb/s power-efficient 4-PAM transmitter," *IEEE J. Solid-State Circuits*, vol. 39, no. 3, pp. 529–532, Mar. 2004.
- [17] S. I. Long and J. Q. Zhang, "Low power GaAs current-mode 1.2 Gb/s interchip interconnections," *IEEE J. Solid-State Circuits*, vol. 32, no. 6, pp. 529–532, Jun. 1997.
- [18] S. Kasturia and J. H. Winters, "Techniques for high-speed implementation of nonlinear cancellation," *IEEE J. Sel. Areas Commun.*, vol. 9, no. 6, pp. 711–717, Jun. 1991.
- [19] C. S. Patal, "Silicon carrier for computer systems," in *Proc. IEEE/ACM Design Automation Conf.*, Jul. 2006, pp. 857–862.
- [20] J. U. Knickerbocker, P. S. Andry, B. Dang, R. R. Horton, C. S. Patel, R. J. Polastre, K. Sakuma, E. S. Sprogis, C. K. Tsang, B. C. Webb, and S. L. Wright, "3D silicon integration," in *Proc. 58th IEEE Electronic Components and Technology Conf.*, May 2008, pp. 538–543.



Byungsub Kim (S'06) was born in Busan, Korea. He received the B.S. degree in electronic and electrical engineering from Pohang University of Science and Technology, Korea, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA.

In the summers of 2006 and 2007, he was an intern at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he developed DFE-IIR architectures and circuits for a compact I/O. He has been with Intel Corporation, Hillsboro, OR, since March 2010. His research interests include high-speed interconnects, computer-aided-design methods, and future network-on-chips.

Dr. Kim was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 IEEE International Solid-State Circuits Conference. He also received the Analog Device Inc. Outstanding Student Designer Award from MIT in 2009.



Vladimir Stojanović (S'96–M'05) received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2005, respectively, and the Dipl. Ing. degree from the University of Belgrade, Belgrade, Serbia, in 1998.

He is currently an Associate Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge. He was with Rambus, Inc., Los Altos, CA, from 2001 to 2004. He was a Visiting Scholar with the Advanced Computer Systems Engineering Laboratory, Department of Electrical and Computer Engineering, University of California, Davis, during 1997–1998. His current research interests include design, modeling, and optimization of integrated systems, from standard mixed-signal and VLSI blocks to CMOS-based electrical and optical interfaces and networks.