Digital Link Pre-emphasis with Dynamic Driver Impedance Modulation

Ranko Sredojević and Vladimir Stojanović
Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Abstract – Digital push-pull impedance-modulating (RM) pre-emphasis driver overcomes the power overhead of equalization in voltage-mode (VM) drivers, improving the output stage efficiency ~2-3x. A compact, fully-digital RAM-DAC implementation with pattern lookup compensates both duty-cycle distortion and driver nonlinearity, while providing a programmable pre-emphasis. A testchip fabricated in 90nm CMOS process shows relatively small signal degradation from dynamic modulation of driver output impedance over a variety of 20" backplanes at 4Gb/s, with energy-efficiency of 2pJ/bit at 100mV of receiver eye.

I. Introduction

Increase in system throughputs and memory bandwidth requirements has put an increased focus on energy-efficiency of high-speed links and I/O circuits in power-limited VLSI systems. In today's large systems-on-a-chip, high-speed links take up a significant portion of power limiting the amount of power left for useful computations [1].

Increase in link energy-efficiency assumes improvements in all its subsystems, of which clocking and signaling transmit-receive chain are the most dominant. To that end, voltage-mode drivers have been introduced instead of current-mode drivers to improve the energy-efficiency of the transmitter [2,3]. However, these voltage-mode drivers suffer from a power penalty when used to implement a transmit pre-emphasis filter [2], which is particularly well suited for asymmetric complexity link channel applications such as memory interfaces [4] and lossy channels with long intersymbol-interference tails such as cables or silicon carriers.

Common transmitters incur significant power penalty in attempt to provide matched, low impedance, termination to the channel over wide range of output voltages necessary for pre-emphasis. However, this matching only notionally improves the performance of the link. In practice, trace fabrication tolerances affect the variation of the channel characteristic impedance, and on-die and packaging parasitic capacitances at transmitter and receiver as well as process variation in the driver output stage can significantly alter the actual termination impedance and matching to the channel. Hence, the blind sacrifice of power efficiency in order to achieve the impedance matching should be re-evaluated in the context of true link performance metrics, to better understand the tradeoff between signal degradation due to transmitter impedance mismatch and power efficiency improvement.

In this work we present a design and experimental verification of a new transmitter pre-emphasis driver that achieves record high energy-efficiency by using a dynamic impedance modulation technique to perform the transmit pre-emphasis without a constant impedance match to the channel. This approach results in a compact, fully digital implementation of the transmit equalizer that favorably compares with VM drivers, achieving ~2-3x better energy efficiency for the same eye opening.

The following sections will first analyze common drivers and the power overhead related to impedance matching, followed by a discussion of how the proposed transmitter balances signal degradation and power efficiency. The remaining sections will present the circuit implementation of this transmitter, measurement results and conclusion.

II. Driver Power Efficiency

Fig. 1 shows power consumption as a function of the output voltage, for common transmit-driver topologies. We assume peak output voltage of $V_{DD}/2$, and a receiver termination of 50Ω. Example constellation points for 2nd order equalization and the corresponding power costs for producing such output level are also marked in Fig. 1 by black dots, for each driver.

Current mode (CM) drivers always consume the same amount of power

$$P_{out}(V_{out}) = 2\left(\frac{V_{DD}^2}{100\Omega}\right). \qquad \text{(CM)}$$

The output is created by steering fixed current through a set of parallel current switches. The low resistance transmit termination (R_{tx}=50Ω) is connected in parallel with the 50Ω channel, effectively utilizing only a half of the driver current.

On the other hand, VM drivers use a resistive divider between V_{DD} and ground, in parallel with the channel, to create various voltages necessary for equalization. As a side effect a wasteful current shunt is created. The amount of current shunted is determined by the transistor sizing, which in turn is determined by the termination impedance. Simple analysis shows that VM drivers are most energy-efficient when transmitting full-swing voltages, since that minimizes the amount of shunted current. When used to create a pre-emphasis driver, VM drivers exhibit significant power penalty as they spend large shunt currents to generate smaller output

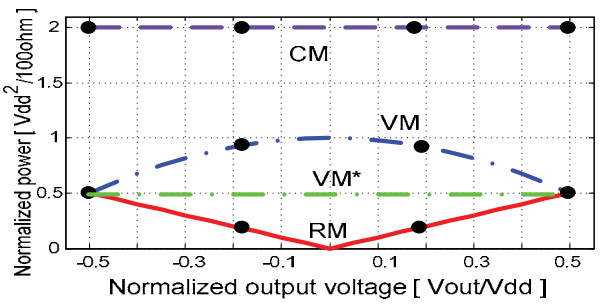

Figure 1. Normalized driver power vs. normalized output voltage

978-1-4244-5760-1/10/$26.00 ©2010 IEEE

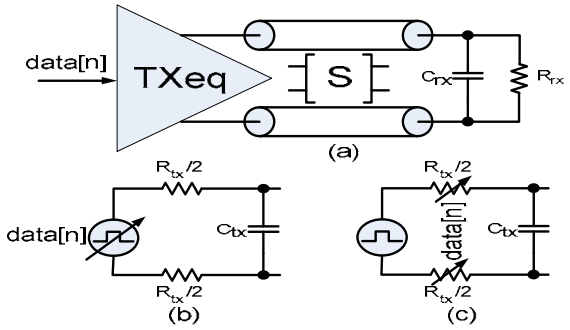

Figure 2. (a) Schematic for study of impedance modulation (b) Equivalent schematic of VM driver, (c) Equivalent schematic of RM driver

voltages, as shown in Fig. 1. The power of the VM driver is given in

$$P_{out}(V_{out}) = \left(\frac{V_{DD}^2}{100\Omega}\right)\left(1 - 2\left(\frac{V_{out}}{V_{DD}}\right)^2\right), \qquad \text{(VM)}$$

explaining why some designs use VM only in series-terminated mode without any equalization [3].

An interesting redesign of VM [5] reduces the amount of crowbar current by introducing additional termination switches between the channel input ports achieving

$$P_{out}(V_{out}) \approx \frac{1}{2}\left(\frac{V_{DD}^2}{100\Omega}\right). \qquad \text{(VM*)}$$

It is obvious that when zero output voltage is transmitted, there is no current delivered into the channel, but VM* consumes the same amount of power as for the maximum signal being sent, thus shunting still exists in this driver. This is only necessary in the absence of supply regulator, in order to achieve a desirable constant supply current.

If we completely eliminate the shunt current, we obtain a driver that always delivers all the current to the channel. In this case, as can be seen in Fig. 2, the power efficiency achieved is

$$P_{out}(V_{out}) = \left(\frac{V_{DD}^2}{100\Omega}\right)\left|\frac{V_{out}}{V_{DD}}\right|. \qquad \text{(RM)}$$

From Fig. 2 we see that, to achieve equalization, such an equalizer would either have to modulate the supply (Thevenin equivalent source, Fig. 2b), or modulate the impedance (i.e. Thevenin impedance, Fig. 2c), and thus dynamically change the voltage divider between the source and channel/receiver. Since fast and energy-efficient dynamic modulation of the supply is very difficult (would require a voltage regulator with multiple adjustable output voltages), we focus on the potentially more energy-efficient implementation of pre-emphasis driver with impedance modulation. This driver changes the impedance to change the voltage divider ratio, while keeping the Thevenin voltage source constant, always delivering all the current through the channel and eliminating the shunt current. As we modulate the resistance seen from the channel into the transmitter, we call this driver the R-mode (RM). In the next section we explore the effects of impedance modulation on link performance over lossy channels.

III. Impedance Modulation

Although, not providing a standard impedance match at the transmitter may sound as an unorthodox idea, in link design practice, an ideal impedance match is impossible to achieve over a typically large link bandwidth due to parasitic capacitances (wiring, pads, ESD) and package inductances. For example, parasitic capacitance at the output of the transmitter is at least ~400fF in well-optimized designs [4, 5]. At 7Gb/s data rate, this parasitic capacitance leads to termination impedance (50 + j 127)Ω at Nyquist frequency, which is significantly different than the nominal 50Ω termination.

Impedance matching is an intuitive but qualitative proxy for describing signal quality, not a fundamental constraint for proper operation. In view of practical obstacles for achieving this perfect matching and possible energy-savings, we proceed to explore the performance of the proposed unmatched transmit driver by evaluating these trade-offs through the ultimate link performance metrics such as energy-efficiency (pJ/bit) for a given bit-error rate (BER) – or equivalent eye opening.

Assuming the model in Fig. 2c, we first explore the impulse response dependence on transmit and receive termination resistance, assuming that they are fixed. The first step is obtaining the transfer function, which is easiest to calculate from Z or Y parameters. To do this, we use the measured S-parameters of a channel and convert them to Z-parameters [7]

$$Z(\omega) = Z_0 \frac{I + S(\omega)}{I - S(\omega)} \qquad (1)$$

where Z_0 is the termination matrix during measurement, ω is frequency and **I** is the identity matrix. In this form it is easy to add impedances in series or in parallel with the channel, on either end, and by sampling in frequency it is possible to obtain the impulse response through inverse discrete Fourier transform (IDFT).

For simulations we first use B3, 3'' Peters' channel [8], as a conservative estimate, because it has low loss and shows more degradation due to impedance mismatch and reflections. In Fig. 3a we can see dependence of 6.25Gb/s pulse response on R_{tx} and R_{rx}, assuming C_{tx}=C_{rx}=300fF (or 600fF single ended parasitic capacitance). To provide a fair comparison of the signal quality, Fig. 3b shows normalization of the responses, where the scaling due to voltage divider was removed. It can be observed that the dominant influence on reflections is from the receiver termination (Fig. 3b), while the transmitter side

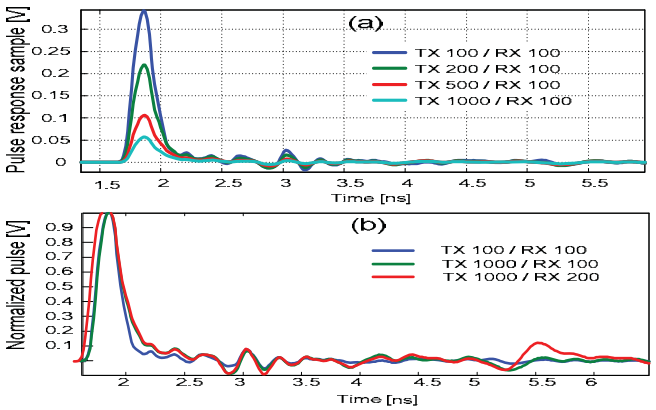

Figure 3. (a) Pulse response of Peters' 3'' channel vs. Rtx/Rrx interface resistance, (b) Normalized pulse response for Peter's 3'' channel to eliminate the scaling vs Rtx/Rrx

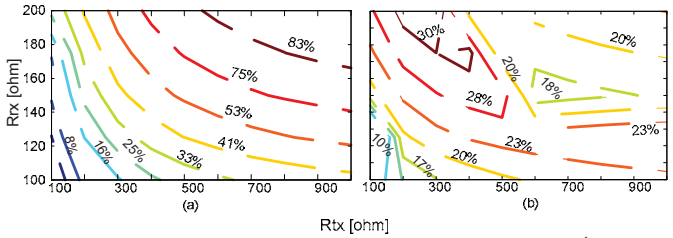

Figure 4. Relative degradation of the worst-case eye equalized by 3rd order TX filter as a function of transmit and receiver interface resistance

increases dispersion and scales the response due to the resistive divider, from Fig. 2.

To visualize signal degradation from additional reflection and dispersion in Fig. 4a and 4b we show relative degradation of the worst-case eye-opening as a function of transmit and receive termination resistance. In Fig. 4a we can see Peters' B3 channel, while Fig. 4b shows results for more lossy 32" Peters' B32 channel [8]. To remove the scaling, introduced by the transmit termination, we use the normalized pulse responses. Lossy B32 channel shows significantly less degradation as the possible reflections from the receiver experience twice the attenuation of the channel until they appear at the output. In this scenario the receiver impedance can be increased to a certain extent to increase the gain of the system. Finally we note that in practice, when using dynamic impedance modulation, we expect less degradation than what is indicated in Fig. 4, which is an estimation of the worst case scenario. This is because RM driver is not a good termination only when it is producing small voltages, which is, due to equalization, only when there are no transitions in the bit-stream, and then there are no reflections. When the transitions are occurring, the RM mode driver switches into the low impedance mode to drive the channel as strongly as possible, and in our design this mode is reasonably well matched.

As the channel is low-pass, we can expect that time-varying transmit impedance would be 'interpolating' between the impulse responses shown in Fig. 3. In this mode, time-invariant transfer function of the system does not exist since poles of the system change due to changes in the transmitter impedance. However, by designing the range of these impedances such that the lowest position of the moving poles is always well beyond the 3dB bandwidth of the channel's dominant pole, the moving poles can be neglected.

V. Transmitter Architecture

The RM driver is implemented as a 7-bit segmented push-pull DAC, as shown in Fig. 5. The LSB and MSB segments of

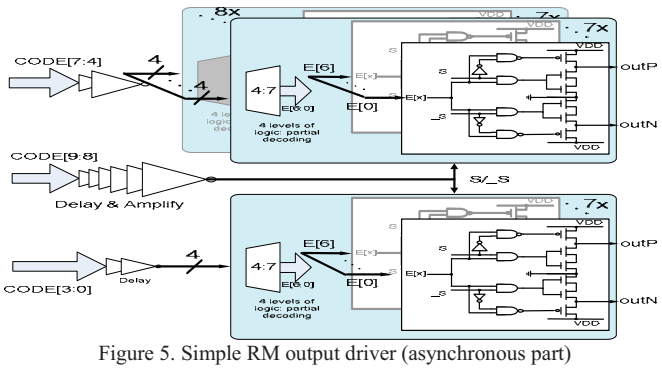

Figure 5. Simple RM output driver (asynchronous part)

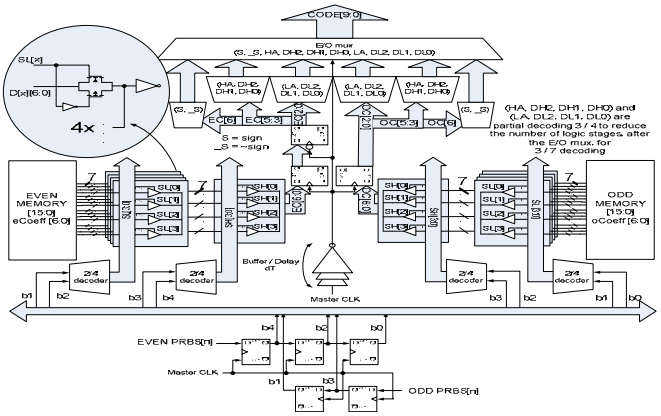

Figure 6. Decoder path of the transmitter (synchronous part)

the DAC are thermometer coded for energy-efficiency as well as to allow regular sub-segment layout and pre-driver/decoder load balancing. The total driver output impedance ranges from 50Ω - 3kΩ. Each sub-segment is designed as a push-pull pseudo-differential CMOS bridge.

The decoder data-path is shown in Fig. 6. Incoming bit-stream is generated by two independent $2^{31}-1$ PRBS generators. To achieve equalization we use a 4-bit sliding window to address into memories. Each memory contains 16 7-bit words (signed 6-bit value) representing a static fully programmable equalization map. We used a dual-data rate (DDR) configuration, with decoding logic partitioned by the DDR multiplexer, to balance the power dissipation with number of unclocked stages in the DAC predriver. The clock in the first pipeline stage is delayed for latching and E/O multiplexing to provide more time for memory readout. Although these common design practices were used to balance the power dissipation and jitter, we rely on nonlinear maps to compensate the dominant effects of driver nonlinearity as well as segment skew and duty-cycle distortion. Small static and dynamic timing errors are mapped to voltage domain by the bandlimited channel, and can be corrected by changing the magnitude of transmitted samples in the lookup table [7].

VI. Measurement Results

Our setup assumes capacitive coupling of the channel. Equivalently receiver side channel interface should have high impedance to both ground and supply, and 100Ω differential termination with no more than ~300fF of equivalent differential capacitance. The DAC transfer curve in Fig. 7 shows significant static nonlinearity. It is compensated with transmitter look-up tables similar to [6]. The address into these static 7-bit tables is a 4-bit long input data pattern, which allows simultaneous pre-emphasis and driver pre-distortion.

In Fig. 8, we show the eye diagrams of the equalized and unequalized AC-coupled link running at 4Gb/s over 26" of FR4 backplane. At this speed, the eye is equalized with

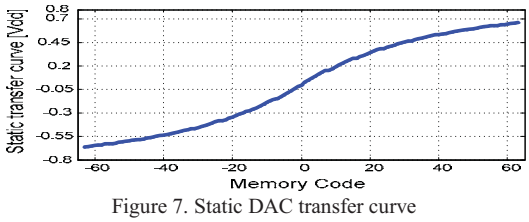

Figure 7. Static DAC transfer curve

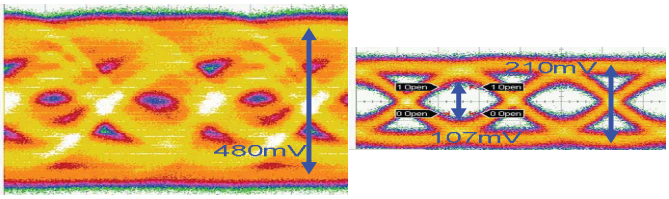

Figure 8. Equalization of FR4 backplane at 4Gb/s

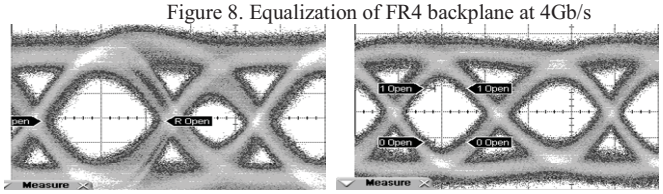

Figure 9. Joint pre-emphasis and duty-cycle correction

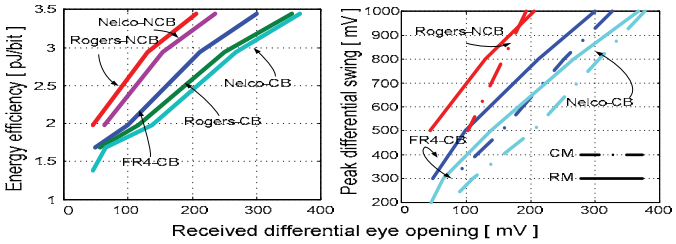

Figure 10. Energy-efficiency over various backplane channels

effectively a two-tap pre-emphasis filter, pre-distorted to correct for static DAC linearity and mapped into the lookup tables. To illustrate the effect of duty-cycle compensation, we distort the duty cycle of the external link clock so that even/odd eyes at the receiver are 292mV/226ps and 128mV/113ps, respectively. Decreasing the strength of symbols with longer duration the eye diagrams are balanced with 207mV/171ps and 198mV/178ps, in Fig. 9.

Fig. 10 shows the energy-efficiency of the driver equalization over backplane channels with 8-10dB loss at Nyquist frequency of 2GHz, from low-loss Nelco and Rogers channels, with and w/o counterboring to lossy FR4 channel. As expected, the equalizer performance is the best on dispersive channels and the worst on low-loss channels without counterboring, where transmit pre-emphasis cannot equalize distant reflections from long backplane via stubs. The achieved eye openings compared to those of a commercial CM driver for same peak output swing, show a relatively small degradation. The common mode fluctuations were spread-out evenly through the symbol period in these measurements and scaled linearly from 30mV to 75mV for eye openings in Fig. 10.

The transmitter power is illustrated in Fig. 11, showing the decoder path overhead of 5mW at 4Gb/s, and linear scaling of driver power with output swing, as predicted in Fig.1. Normalized by performance of the received eye opening, the

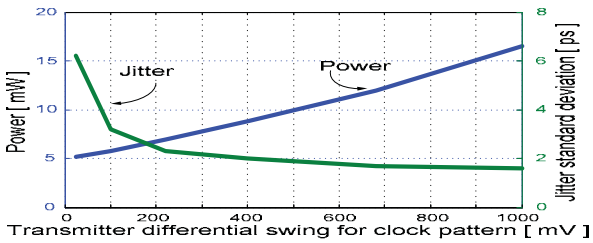

Figure 11. Power and jitter vs. output swing for a clock output pattern

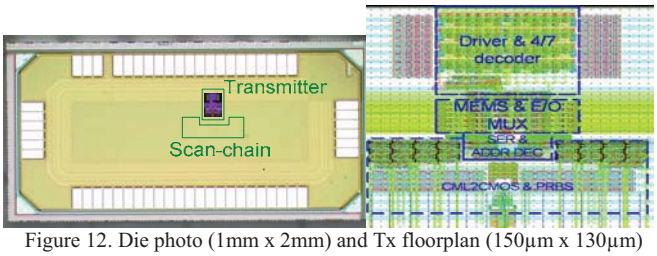

Figure 12. Die photo (1mm x 2mm) and Tx floorplan (150µm x 130µm)

pre-emphasis RM driver stage achieves 2-3x better energy-efficiency than VM and 5-8x that of CM.

Die photo and transmitter layout are shown in Fig. 12. Due to design error in transition from serializer to the memory address decoder the lookup table address was reduced to 2-bit instead of 4-bit address. Decoding speed was also reduced due to large load the address decoder presents to the serializer's quasi-minimal output buffers. As the 2nd order lookup table would have been ~4x smaller, we can expect that the decoder overhead, Fig. 11, would have been 1.5 – 2mW in this case, bringing us to ~1.25pJ/bit, for 100mV eye opening.

VII. Conclusions

This work presents simple and efficient transmit driver for high-speed links. Novelty of the approach is in re-evaluation of the costs and benefits of the impedance matching in common transmit drivers. While perfect termination results in pulse response that is easier to equalize, it is practically impossible to achieve in GHz-wide baseband systems due to parasitic and process variations. By careful analysis and design, the power overhead of common pre-emphasis drivers can be eliminated, while signal degradation due to reflections and additional dispersion can be controlled. This prediction is confirmed on a prototype driver with modulated series termination shown to effectively equalize wide a variety of 20'' backplane channels. The resulting RAM-DAC based, circuit implementation is fully digital and very compact as efficient inverter-style drivers can be used. Proper optimization and co-design of channel parameters along with exploration of power efficiency vs. signal degradation tradeoff in transmitter and receiver can provide a way for further improvements in link energy-efficiency.

Acknowledgements

This work is supported by C2S2 and IFC FCRP Centers.

References

[1] J.L Shin, *et al.* "A 40nm 16-core 128-thread CMT SPARC SoC processor," *IEEE ISSCC,* pp. 98-99, Feb. 2010.

[2] H. Hatamkhani, J. K-L. Wong; R. Drost, and C-K. K. Yang "A 10-mW 3.6-Gbps I/O transmitter," *VLSI Circuits Symposium*, pp. 97- 98, June 2003

[3] J. Poulton, et al., "A 14-mW 6.25-Gb/s Transceiver in 90-nm CMOS," *IEEE JSSC, Vol. 42, No. 12, December 2007, pp. 2745 -2756.*

[4] H. Lee, *et al.* "A 16 Gb/s/Link, 64 GB/s Bidirectional Asymmetric Memory Interface," *IEEE JSSC,* , vol.44, no.4, pp. 1235-1247, April 2009.

[5] W. Dettloff, et al., "A 32mW 7.4Gb/s Protocol-Agile Source-Series-Terminated Transmitter in 45nM CMOS SOI", *ISSCC, February 2010, pp. 370-371*

[6] C. Hayun, *et al.* "A 12.5-Gbps, 7-bit transmit DAC with 4-tap LUT-based equalization in 0.13µm CMOS," *IEEE CICC*, pp. 563-566, 2008.

[7] V. Stojanovic and M. Horowitz, "Modeling and Analysis of High-Speed Links," *IEEE CICC, September 2003*

[8] William Peters, et al., "ATCA Channel Data for Backplane Ethernet Task Force"