A 20Gb/s 136fJ/b 12.5Gb/s/μm On-Chip Link in 28nm CMOS

Meisam Honarvar Nazari, Azita Emami-Neyestanak
Electrical Engineering Department, California Institute of Technology, Pasadena, CA 91101
Email: meisam@caltech.edu, azita@caltech.edu

***Abstract*—A high data rate, low power on-chip link in 28nm CMOS is presented. It features a double-sampling receiver with dynamic offset modulation and a capacitively-driven transmitter. The functionality of the link was validated using 4-7mm minimum-pitch on-chip wires. It achieves up to 20Gb/s of data rate (13.9Gb/s/μm) with BER$< 10^{-12}$. It has better than 136fJ/b of power efficiency at 10Gb/s. The total area of the transmitter and receiver is 1110μm^2.**

***Index Terms*—Interconnect, On-chip signaling, Double-sampling, Dynamic offset modulation, Crosstalk.**

I. Introduction

As VLSI technologies and multi-core processor chips continue to scale, long on-chip wires will present increasing performance limitations. While transistors favor from technology scaling, the shrinking cross-sectional area of the on-chip wires increases electrical resistance and hence their latency, which has a quadratic relation with the wire length. Simple inverter-based repeaters can partially mitigate the latency problem, where an optimal design makes the repeated wire delay linear with length instead of quadratic. However, the associated power and area become prohibitive as the technology scales due to the increased number of repeaters per unit length, Fig. 1. Low-swing differential signaling [2]-[6], current-mode signaling [7], [8], equalization [6], [9] and transmission lines [2], [7] have been employed to resolve the energy and latency problem of the repeated links. However, these techniques are becoming less adequate in meeting bandwidth density and power requirements.

Due to the RC nature of the on-chip wires, binary signals suffer from a long train of post-cursor inter-symbol interference (ISI). To eliminate ISI, equalization techniques such as decision feedback equalization (DFE) can be utilized, but, the long post-cursor tail necessitates many DFE taps, which results in significant power overhead. This problem is exacerbated as the technology scales. RC signal emulation in a DFE, is also an attractive solution to eliminate many taps of post-cursors [6], [9]. The main limiting factor in this technique is to meet the timing requirement in the feedback loop, especially at high data rates. In this paper, we present an on-chip link using minimum-pitch wires for high-speed signaling to address the bandwidth requirement of future microprocessors. We propose a double-sampling technique with a feed-forward dynamic offset modulation (DOM) to achieve high data rates over minimum-pitch and long on-chip wires that suffer from excessive loss and latency.

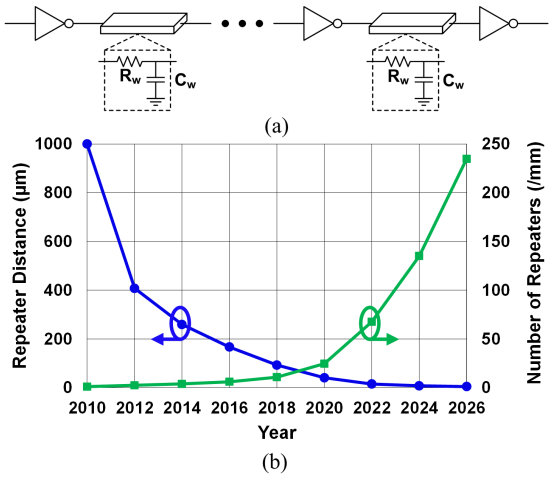

Fig. 1. (a) Repeater-based on-chip link, (b) the distance and number of repeaters for an optimally repeated wire in different technology nodes [1].

In order to further improve data rate and reduce power consumption, a capacitively-driven transmitter is employed [4]. The emphasis of the proposed design is low power consumption, high bandwidth density and scalability to future technology nodes.

II. Architecture

Minimum-pitch wires can have a slow exponential response to a fast transition, with a time-constant (τ) much larger than the bit time (T). Instead of conventional equalization techniques, in this work, we propose to employ a mostly-digital double-sampling technique to break the trade-off between the data-rate and the on-chip wire time-constant [10]. Fig. 2 shows the top-level architecture of the proposed receiver. As shown in Fig. 2(b), the input voltage is sampled at the end of two consecutive bit times (V[n-1], V[n]) and these samples are compared to resolve each bit: $\Delta V[n] = V[n] - V[n-1] > 0$ results in 1, and $\Delta V[n] < 0$ results in 0. Note that the overall sampling rate is still equal to the data rate. This double-sampled voltage ($\Delta V[n]$) will be input-dependent due to the channel transfer function, Fig. 3(b). To resolve this problem, the offset of the next stage comparator can be dynamically modulated to provide a constant voltage at its input regardless of the data sequence. In this receiver, an offset proportional to the previous sample, V[n-1], is applied to the comparator. For instance, in case of a large $\Delta V[n] = \Delta V_{max}$ (e.g. a one after many zeros), or a very small $\Delta V[n]$ (e.g. a zero

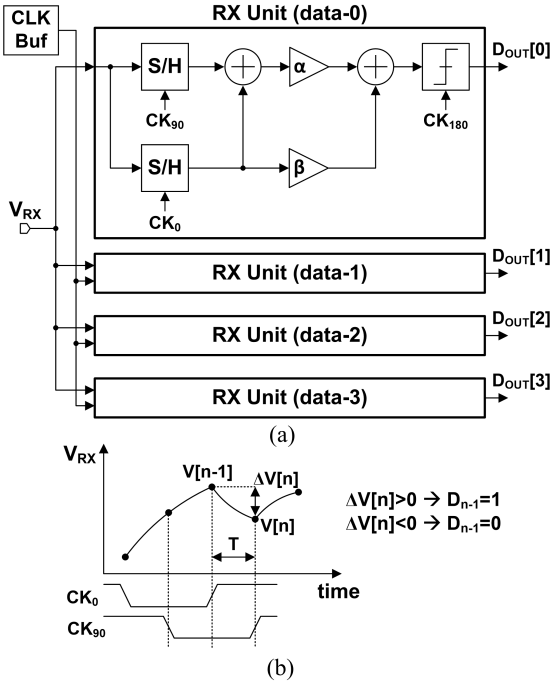

Fig. 2. Receiver top-level architecture, double-sampling technique and DOM.

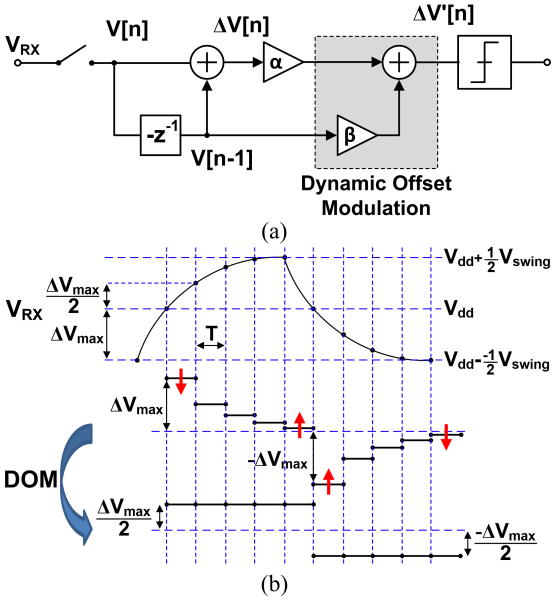

Fig. 3. (a) Z-domain representation of the double-sampler and the dynamic offset modulation. (b) Operation of the dynamic offset modulation.

after many zeros), an offset equal to $\frac{\Delta V_{max}}{2}$ is subtracted resulting in $\Delta V'[n] = \frac{\Delta V_{max}}{2}$, and $\frac{-\Delta V_{max}}{2}$, respectively as shown in Fig. 3(b). The same scenario is true for the opposite case. Fig. 3(a) shows the z-domain representation of the double-sampling and the dynamic offset modulation technique. Assuming a dominant pole of $\omega_p = \frac{1}{\tau}$, it can be shown that for an exponential signal, dynamic offset modulation can eliminate the input dependency of the double-sampled voltage if DOM gain, β, is equal to

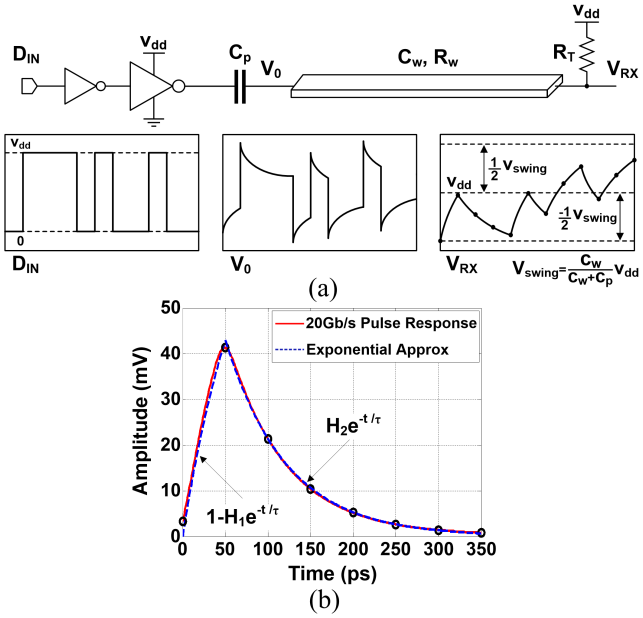

Fig. 4. (a) Capacitively-driven transmitter, (b) double-sampling technique to resolve the received data.

$$\beta = \alpha(1 - e^{\frac{-T}{\tau}}), \tag{1}$$

where α is the main path gain. This results in a constant double-sampled voltage, $\Delta V'[n]$, equal to

$$\Delta V'[n] = \frac{\alpha}{2}(1 - e^{\frac{-T}{\tau}}) = \frac{\alpha}{2}\Delta V_{max}. \tag{2}$$

Another advantage of this technique is the capability to perform immediate demultiplexing at the front-end. A quarter-rate architecture (multiplexing factor of 4) is employed in this design. As a result, the comparators will operate in a fraction of the data rate.

Utilizing low-swing signaling also reduces the power consumption in an on-chip interconnect, where most of the power is associated with the dynamic charging and discharging of the wire capacitance (C_w). A separate supply can be employed for an inverter-based transmitter to reduce the signal swing and hence improve power efficiency. However, it is not desirable to have multiple supplies on chip, as it makes the power distribution complicated. An alternative approach to achieve low swing is to drive the wire through a capacitor, C_p. This helps reducing the signal swing on the wire through a capacitive voltage divider. Ignoring the parasitic capacitance associated with the driver and the receiver, the resulting signal swing at the receiver side will be equal to $\frac{C_w}{C_w+C_p} \times V_{dd}$. This capacitor also pre-emphasizes transitions and reduces the driver's load. Because it acts as a high-pass filter, the capacitor increases the bandwidth of the wire by almost a factor of $\frac{C_w}{C_p}$ and decreases latency. Fig. 4 shows the capacitive driver and the resulting signals at the input and the output of the wire. As shown in Fig 4(b), the receiver input has an exponential behavior.

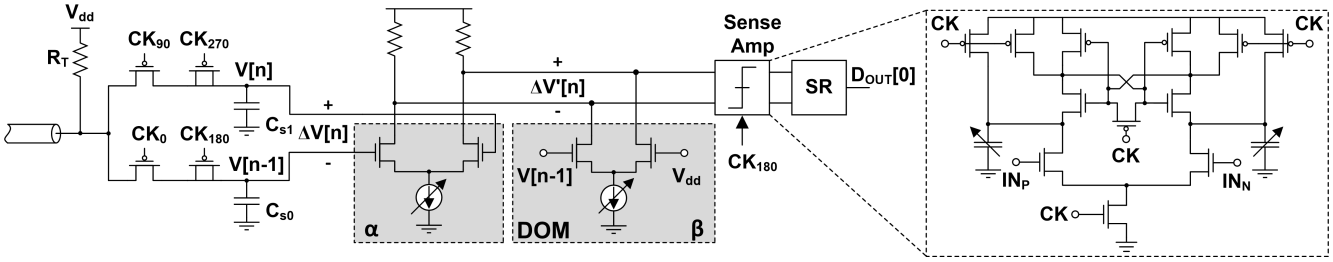

Fig. 5. Transistor level schematic of the receiver front-end and the StrongArm sense amplifier with capacitive offset cancellation.

III. Implementation

As mentioned in previous section, in the proposed design a capacitive driver is employed to achieve small voltage swing and reduce power consumption. It should be noted that the coupling capacitor limits the maximum number of consecutive ones/zeros due to the voltage drift associated with the high-pass behavior of the link. As a result, the coupling capacitor, C_p is optimized to reduce the time constant of the drift process while providing reasonable voltage swing and bandwidth enhancement. In this design a PMOS transistor realizes a 400fF capacitor for the driver. This results in about 140mV voltage swing over a 7mm wire and less than about 1mV drift in voltage after more than 40 consecutive ones/zeros. The termination resistor sets the receiver's DC voltage to V_{dd}. The high DC voltage at the input of the receiver guarantees best operation of the PMOS samplers as shown in Fig. 5. It also biases the PMOS coupling capacitor in the accumulation regime to ensure maximum capacitance and hence least area. The sampler utilizes dummy switches to reduce charge injection. The residual error due to charge sharing and clock feed-through is removed by the double sampling technique, which performs the single-ended to differential conversion immediately at the receiver input. The input voltage is sampled by a bank of four sample/holds (S/H) driven by quarter-rate clock phases). The sampling capacitor C_s is chosen such that enough SNR for BER$< 10^{-12}$ is achieved. An amplifier with about 6dB gain provides isolation between the sensitive sampling node and the sense amplifier. It also creates a constant common-mode voltage and prevents input dependent offset. A StrongARM sense amplifier is employed to achieve high speed and low power. The sense amplifier has a separate offset cancellation for mismatch compensation through the variable capacitors shown in Fig. 5. The wires are implemented using minimum-pitch ($0.36\mu m$ width, $0.36\mu m$ spacing) M_7 layer in the 9-metal process where M_6 and M_8 layers are densely populated to mimic orthogonal interconnects in a microprocessor chip. A shielded wiring structure is employed to minimize coupling noise from adjacent lines, as shown in Fig. 6. This provides noise immunity while the double-sampling technique eliminates sensitivity to common-mode interferences at the receiver. Fig. 6 shows the simulated and measured characteristics of the 5mm and 7mm on-chip wires.

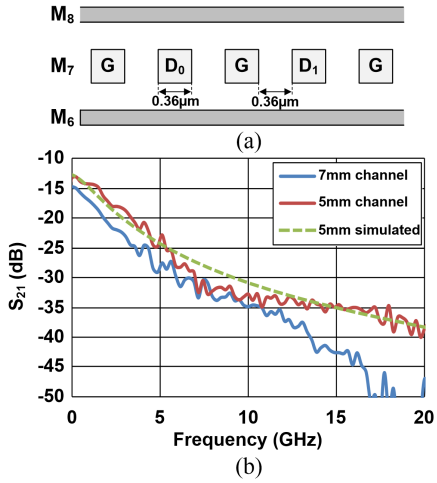

Fig. 6. (a) Shielded single-ended on-chip wire. (b) Simulated and measured characteristics of the on-chip wires.

IV. Measurement Results

The link prototype was fabricated in 28nm LP CMOS technology with the receiver and transmitter occupying less than $950\mu m^2$ and $160\mu m^2$, respectively. The functionality of the transceiver was validated using single-ended on-chip wires with different lengths (4-7mm). PRBS-7 to 31 data was generated off-chip and sent to the on-chip transmitters. Fig. 7 shows how power consumption (including all clock buffers) changes with increasing the data rate. The 5mm link operates up to 18Gb/s while achieving better than 164fJ/b power efficiency. As the link length and thus wire capacitance increases, signal swing at the receiver degrades. This result in the maximum measured data rate of 15Gb/s for the 7mm link. At this data rate the power efficiency is about 180fJ/b. For the 1.44μm wire pitch, the bandwidth density for the 5mm and 7mm links is 12.5 and 10.42Gb/s/μm, respectively. As this design employs mainly digital blocks, the power consumption almost linearly scales with data rate, Fig 7(a). An optimally repeated version of the link with the same geometry was also simulated for comparison purpose. The proposed scheme offers over 4x improvement in energy efficiency and about 40% lower latency compared to the repeated link. The receiver offers a peak energy efficiency of 136fJ/b at 10Gb/s data rate for 7mm wires. The transceiver was also tested using a 4mm wire. This link is comprised of two adjacent

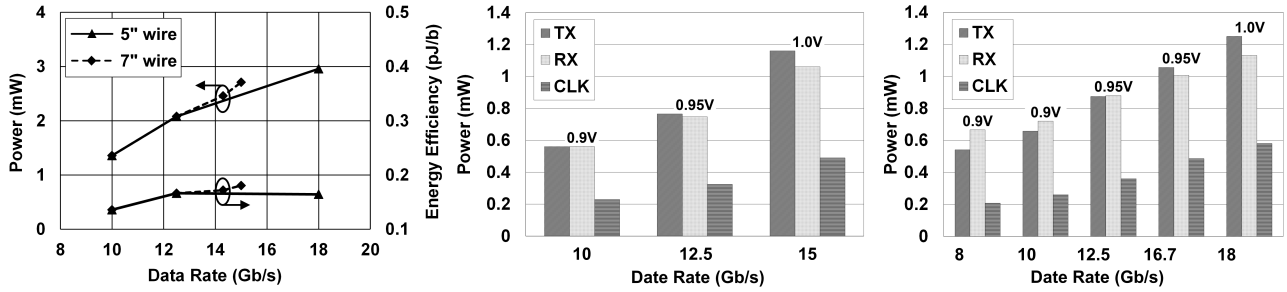

Fig. 7. Total power consumption (a) and power breakdown for the 5mm (b) and 7mm (c) links.

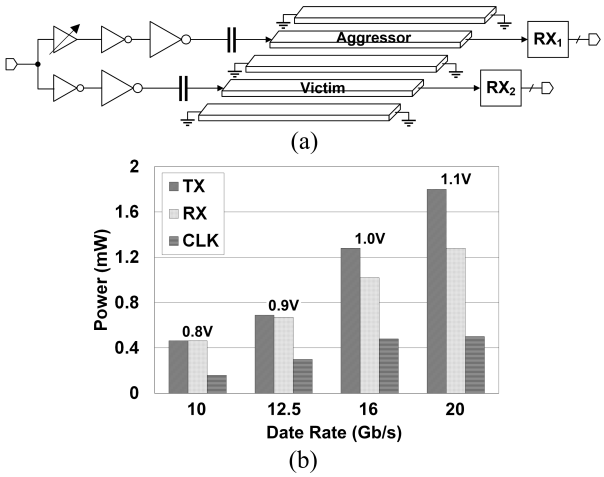

Fig. 8. (a) Crosstalk measurement setup, (b) power consumption of the 4mm link in the presence of an aggressor.

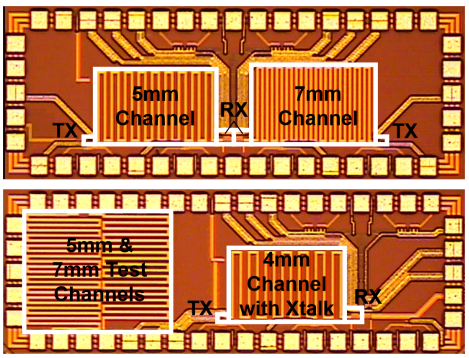

Fig. 9. Die Micrograph.

wires to investigate the effect of crosstalk, Fig. 8. Once the aggressor is activated while V_{dd}=0.95V the SNR drops due to the crosstalk noise, which causes an increase in the BER. By increasing the supply to 1V we could restore the BER, which translates into about 5% degradation in the SNR due to the crosstalk. This level of crosstalk noise is comparable to a twisted differential architecture. The advantage of the shielded structure is that it eliminates the Miller capacitance that exist in a differential pair and hence offers a better power efficiency as well as area efficiency. The immediate single-ended to differential conversion provided by the double-sampling technique also minimizes the sensitivity to common-mode noise. Maximum data rate of 20Gb/s with BER$< 10^{-12}$ and 180fJ/b of energy efficiency was achieved over this link. Table I summarizes the performance of the proposed link and compares it with prior art.

V. Conclusion

The proposed transceiver for repeater-less on-chip communication demonstrates high bandwidth density, low latency, and low power consumption. The mainly digital architecture is well-suited for highly-scaled technologies. Experimental results validate the functionality of the link in 28nm CMOS. It offers up to 20Gb/s/ch data rate and 13.9Gb/s/μm bandwidth density with better than 180fJ/b energy efficiency.

TABLE I
PERFORMANCE SUMMARY AND COMPARISON

	[2]	[3]	[10]	[5]	[6]	[8]	This Work
Technology	CMOS 90nm	CMOS 90nm	CMOS 65nm	CMOS 65nm	CMOS 90nm	CMOS 90nm	CMOS 28nm
Data Rate	10Gb/s	4.9Gb/s	20Gb/s	10Gb/s	2.0Gb/s	4.0Gb/s	15Gb/s-20Gb/s
Length	5mm	5mm	10mm	6mm	10mm	10mm	4-7mm
Energy Efficiency	0.27pJ/b	0.34pJ/b	1.36pJ/b	1.0pJ/b	0.28pJ/b	0.36pJ/b	0.136pJ/b
Bandwidth Density	0.5 Gb/s/µm	4.38 Gb/s/µm	1.5 Gb/s/µm	2.56 Gb/s/µm	1.16 Gb/s/µm	2.0 Gb/s/µm	10.42-13.9 Gb/s/µm
Bit Energy Efficiency	54 fJ/mm	68 fJ/mm	136 fJ/mm	174 fJ/mm	28 fJ/mm	36 fJ/mm	19.5 fJ/mm

References

[1] International Technology Roadmap for Semiconductors (ITRS) 2011 Update.

[2] H. Ito *et al.*, *IEEE Int. Interconnect Tech. Conf.*, pp. 193-195, 2007.

[3] J. Seo *et al.*, *IEEE ISSCC Dig. of Tech. Papers*, pp. 182-183, Feb. 2010.

[4] R. Ho *et al.*, *IEEE J. Solid-State Circuits*, vol. 43, No. 1, pp. 52-60, 2008.

[5] D. Walter *et al.*, *IEEE ISSCC Dig. of Tech. Papers*, pp. 180-181, Feb. 2012.

[6] E. Mensink *et al.*, *IEEE J. Solid-State Circuits*, vol. 45, No. 2, pp. 447-457, 2010.

[7] A. P. Jose *et al.*, *IEEE J. Solid-State Circuits*, vol. 42, No. 6, pp. 1415-1424, 2007.

[8] B. Kim *et al.*, *IEEE ISSCC Dig. of Tech. Papers*, pp. 66-67, Feb. 2009.

[9] Y. Liu *et. al.*, *IEEE ISSCC Dig. of Tech. Papers*, pp. 182-183, Feb. 2009.

[10] M. H. Nazari *et al.*, *IEEE ISSCC Dig. of Tech. Papers*, pp. 130-131, Feb. 2012.