

Asynchronous Current Mode Serial Communication

Rostislav (Reuven) Dobkin, Michael Moyal, *Senior Member, IEEE*, Avinoam Kolodny, and Ran Ginosar

Abstract—An asynchronous high-speed wave-pipelined bit-serial link for on-chip communication is presented as an alternative to standard bit-parallel links. The link employs the differential level encoded dual-rail (LEDR) two-phase asynchronous protocol, avoiding per-bit handshake and eliminating per-bit synchronization, in contrast with synchronous serial links that rely on complex clock recovery. Novel low-power current signaling driver and receiver circuits are presented, enabling high-speed communication at a very low voltage swing over long wires. In contrast, previous methods employed voltage sensing, resulting in higher swing, higher dynamic power, shorter wires or slower operation. The asynchronous current mode driver is designed to support varying data rates, and it eliminates the need for balanced codes and busy toggling that prevent deep discharge. The data cycle time of the link is equal to a single gate delay, enabling 67 Gb/s throughput in 65-nm technology. Wave-pipelining is employed also by the asynchronous SERDES circuits, to enable such high speed operation. The link was SPICE simulated for 65-nm technology, using wire models obtained by a 3-D EM solver. The link incurs lower power and area relative to synchronous and asynchronous bit-parallel communications, and these relative benefits also scale with technology.

Index Terms—Crosstalk, current signaling, LEDR, SERDES, serial communication.

I. INTRODUCTION

THE performance of VLSI digital logic has increased at an exponential rate over the years thanks to transistor size scaling [1]. While performance of local interconnect follows a similar trend, global wires do not, challenging long range on-chip data communications in terms of latency, throughput and power. High-capacitance of the global interconnect is one of the main sources for losses over the wire, leading to a degraded performance in terms of throughput and power. In addition, as systems-on-chip (SoC) integrate an ever growing number of modules, on-chip inter-modular communications become congested and the modules must turn to serial interfaces, similar to the trend from parallel to serial chip-to-chip interconnects.

Common synchronous on-chip parallel links (multi-wire interconnects) occupy large area, present high capacitive load and incur high dynamic and leakage power and cross-coupling noise, especially when long-range communication is

considered. The problem exacerbates for applications with low-utilization (e.g., network-on-chip [2]) or with high interconnect congestion (e.g., routers, cross-bar switches [3], [4]). The clock frequency of synchronous parallel links is bounded by clock and data uncertainty that worsens as the links get longer. While standard synchronous serial links, employing clocks similar to those of a parallel links, are unattractive due to limited bit-rate, novel high performance serial links may provide an alternative for the parallel links.

Synchronous serial links are typically employed for off-chip communications, where pin-out limitations call for a minimal number of wires per link. Source-synchronous protocols are often used for those applications [5]–[10]. A common timing mechanism for serial interconnects injects a clock into the data stream at the transmitting side and recovers the clock at the receiver. Such clock-data recovery (CDR) circuits often require a power-hungry PLL, which may also take a long while to converge on the proper clock frequency and phase at the beginning of each transmission. If the receiver and transmitter operate in different clock domains, the transaction must also be synchronized at both ends, incurring additional delay and power. Alternatively, an asynchronous data link employs handshake instead of clocks. Traditional asynchronous protocols are relatively slow due to the need to acknowledge transitions [11], [12]. In [13], asynchronous protocols share data lines, but their performance depends on wire delays.

High-speed serial links, having data cycle of a few gate delays (down to single gate-delay cycle), have been recently proposed [14]–[22]. These fast links employ wave-pipelining [23]–[25], low-swing differential signaling, fast clock generators and asynchronous protocols. In addition, these links require channel optimization to support wide-bandwidth data transmission over the link wires. A wave-front train serialization link was presented in [19]. The serializer is based on a chain of MUXes (similar to [26]). The link is single-ended and employs wave-pipelining. The link data cycle is approximately 7 FO4 (fan-out of four) delays (3 Gb/s@180 nm). Wave-pipelined multiplexed (WPM) routing technique was presented in [20], [21]. WPM routing employs source synchronous communication and its performance is limited by the clock skew and delay variations. Employing low-voltage differential pairs for on-chip serial interconnect was discussed in [17] and [18], where data was sampled at the receiver without any attention to synchronization issues. A three-level voltage swing was presented in [27], requiring non-standard amplifiers.

Circuits that had originally been designed for off-chip communications [5], [10] were adopted for on-chip serial link in [16]. An output-multiplexed transmitter is connected to a multiplexed receiver, requiring clock calibration at the receiver side. Both transmitter and receiver use multi-phase DLL circuits. The link employs low-swing differential signaling and transfers

Manuscript received July 06, 2008; revised November 16, 2008; accepted March 15, 2009. First published September 01, 2009; current version published June 25, 2010.

R. (Reuven) Dobkin, A. Kolodny, and R. Ginosar are with the VLSI Systems Research Center, Electrical Engineering Department, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: rostikd@tx.technion.ac.il; kolodny@ee.technion.ac.il; ran@ee.technion.ac.il).

M. Moyal is with the Intel Corporation, IDC, Haifa, Israel (e-mail: miki.moyal@intel.com).

Digital Object Identifier 10.1109/TVLSI.2009.2020859

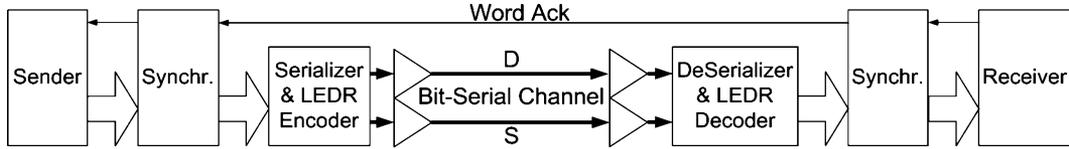


Fig. 1. Serial communication link.

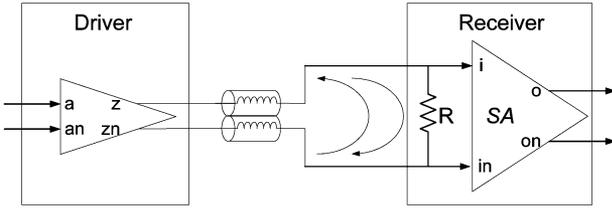


Fig. 2. Typical “current mode” signaling: Current mode driver, differential channel, and voltage mode receiver.

eight-bit words. The output-multiplexed architecture delivers better performance than input-multiplexing (down to $2 \cdot \text{FO4}$ data cycle), but at the expense of much higher output capacitance (that grows linearly with the word-width). A fabricated chip demonstrated an operational 3-mm link.

In this paper, we consider a single FO4 gate delay bit-cycle serial link (Fig. 1). Previously [14], [15], we investigated a high-speed [0]serial link with voltage sensing at the receiver side. The link employed novel high-speed serializer and de-serializer wave-pipelining circuits. Wave-pipelining was also employed over the interconnect wires, together with differential encoding. Link throughput was independent of word width. The link was found to be operational at its highest speed (67 Gb/s at 65 nm) with up to four millimeters of length. The high bandwidth of that novel high-speed asynchronous serial link can be traded off for power and area, reducing the overall cost of inter-modular communication [22], [28].

This paper describes shifting from voltage to current signaling. The high capacitance of the long interconnect is the main contributor to signal degradation over the link. Fast full-swing transitions result in high dynamic currents, dissipating power and causing crosstalk noise. Current signaling offers a way of avoiding high voltage swings. Typical “current-mode” signaling methods, e.g., LVDS, are based on a current driver and a voltage sensor using a termination resistor at the receiver (Fig. 2). Since the receiver measures voltage over the termination resistor using a voltage-mode sense amplifier, similar to voltage-mode signaling, such “current-mode” methods still depend on the voltage swing over the channel, limiting performance by high channel capacitance.

Genuine current-mode sense-amplifiers have been employed in SRAM applications, where speed is critical during read operation and voltage swings are problematic due to sizing limitation of the bit line drivers [29]–[31]. Similar solutions are applied to global interconnect and are reported to dissipate less power and achieve higher throughput than an interconnect buffered with an optimal number of repeaters [32]–[36]. Unfortunately, the bandwidth of previously reported current-mode receivers is insufficient for achieving our target data cycle of a single FO4 gate delay over the link.

In this paper, we present a novel genuine current-mode serial link that enables a data cycle of a single FO4 gate delay and is suitable for both constant rate (synchronous) and asynchronous operations. The novel circuits support wires of nearly twice the length of those for voltage-mode, leading to significant cost reduction. These custom circuits are applicable to SoC when implemented and delivered as hard IP cores.

The main contributions of the paper are threefold. First, we present a high-speed LEDR-encoded wave-pipelined bit-serial differential link with a single FO4 gate-delay data cycle, achieving 67-Gb/s simulated throughput when using 65-nm technology over up to 7 mm. Second, we show the scalable superiority of the serial link over standard bit-parallel links. Third, we present novel high-speed low-power current-mode asynchronous signaling circuits, including SPICE simulations for 65 nm.

The rest of the paper is organized as follows. In the next section, we discuss the relative performance of serial and parallel links. In Section III the single gate delay link is presented. Section IV describes in detail the asynchronous current-mode circuits and their performance.

II. PARALLEL VERSUS SERIAL ON-CHIP COMMUNICATION

Parallel links incur a heavy cost in terms of area and power. In addition, the performance of a parallel link performance is bounded by available clock rate and by clock jitter and skew, delay uncertainty incurred by process variations in repeaters [37], in vias [38] and in the interconnect itself [38], [39], crosstalk noise [38], [40], and layout geometries [41]. The link clock frequency should be sufficiently low to enable reliable data sampling at the receiver, when the latest and earliest data arrival worst cases are considered. We adopt the notation of [23] and draw the delay uncertainty for source-synchronous communication in Fig. 3. This leads to following bound for the clock cycle:

$$T_{\text{CLK}} > 2 \cdot (\delta_{\text{MAX}} - \delta_{\text{MIN}}) + 4 \cdot \Delta_{\text{CLK}} + T_{\text{SU}} + T_{\text{H}} \quad (1)$$

where δ_{MAX} and δ_{MIN} are the max and min data delays over the link (which are also the clock uncertainty in source-synchronous communication), Δ_{CLK} is the one side clock skew, and T_{SU} , T_{H} are the setup and hold times of the receiver flip-flop.

The data and clock uncertainty grow with the link length and width [28], [42]. Therefore, the bandwidth of a long-range link is more restricted than a short-range one. In order to support a given bandwidth for a longer range, the cost of the parallel link, e.g., in terms of required shielding, grows at a super-linear rate with length [28].

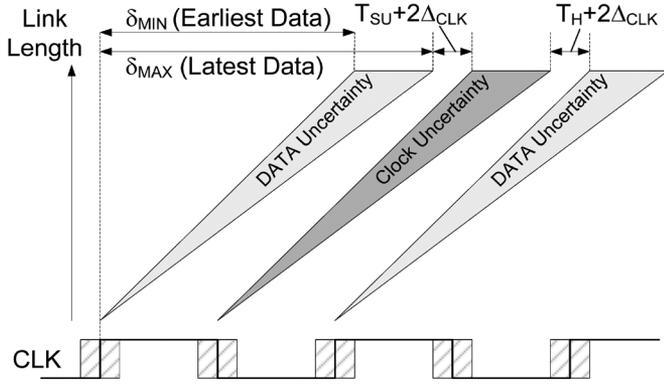


Fig. 3. Parallel link minimal clock cycle is limited by clock jitter and skew and by link-length-dependent delay differences among the parallel wires due to variations and crosstalk.

TABLE I
DATA CYCLES OF SEVERAL HIGH-SPEED SERIAL LINKS

| Name | Data Cycle (# of FO4 Inv. Delays) | Reference |
|-----------|--------------------------------------|-----------|
| WAFT | 7* (3 Gbps at 180nm) | [19] |
| IIP | 3* (5 Gbps at 250nm) | [17] |
| 3-Levels | 3* (1 Gbps at 1200nm) | [27] |
| Pulsed CM | 2 (8 Gbps at 180nm) | [16] |
| This Work | 1 (67 Gbps at 65nm) | |

*Estimated data cycle according to reported maximal throughput and FO4 delay values from [15]

Bit-serial communication links offer an alternative to bit-parallel interconnects, mitigating the issues of area, routability and power, since there are fewer wires, fewer line drivers, and fewer repeaters. To support a throughput similar to bit-parallel links, several asynchronous wide-bandwidth serial link circuits [11], [14]–[21], all operating faster than the system clock, have been proposed. Several of the wide-bandwidth serial links are listed in Table I according to their minimal data cycle. They are insensitive to clock jitter and skew thanks to asynchronous design. Recalling (1), the process variation contribution to the wire skew $\delta_{MAX} - \delta_{MIN}$ of serial links is smaller than in parallel links thanks to the closer placement of repeaters and wires. Cross-talk is also reduced when there are no multiple bits progressing in parallel. Thus, the bandwidth of serial link is less affected by the length than a parallel link [28].

The serial link presented in this work enables the shortest data cycle among the designs in Table I. The high speed is enabled by the following novel features: a wave pipelined shift register, a transmission latch, a transition generator, split architecture, split and merge circuits, LEDR on-the-fly encoding and decoding, wire physical design and layout, and fast current-mode receiver with minimal input voltage swing. All these features are described below.

Comparative analysis of parallel and serial links [21], [28], [43] shows a tradeoff between link length on the one hand and performance parameters such as dynamic and leakage power, active and interconnect area, and latency, on the other hand. For a fixed throughput, the serial link is always preferable in terms of interconnect area and incurs less routing congestion than parallel links. For links longer than a certain length, the serial link

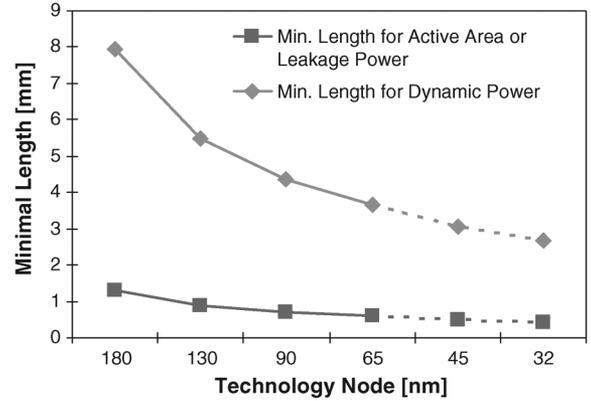


Fig. 4. Minimal length that justifies using a serial link: at longer range, the serial link takes less area and dissipates less leakage or dynamic power than the parallel link [28].

also outperforms the parallel link in terms of active area, leakage and dynamic power [28]. The relative improvement grows with technology scaling, as shown in Fig. 4 for a single gate delay serial link [28]. The figure shows the link length at which a single gate delay cycle serial link becomes superior to a parallel link in terms of leakage and dynamic power (for 8-bit words, equal bit-rate and fully-shielded parallel link with 8·FO4 clock cycle). A detailed comparative analysis is presented in [28].

III. SINGLE GATE DELAY ASYNCHRONOUS SERIAL LINK

The proposed serial link (Fig. 1) employs low-latency synchronizers at the source and sink [44], two-phase NRZ level encoded dual rail (LEDR) data/strobe (DS) encoding [45]–[47] and an asynchronous handshake protocol (allowing non-uniform delay intervals between successive bits), serializer and de-serializer and line drivers and receivers. Acknowledgment is returned only once per word, rather than bit by bit, enabling multiple bits in a wave-pipelined manner over the serial channel. The data over the link wires can be further encoded differentially, inspired by the DS-DE IEEE1355-95 standard [47]. LEDR signaling is preferred over other serial asynchronous protocols for lower power and higher rates [48]. The D and S wires employ (fully shielded) waveguides, enabling multiple traveling signals. On a well-designed waveguide, long wires may carry multiple bits in succession simultaneously.

LEDR encoding is performed on the fly at the very low cost of a XOR and a few transmission gates [14]. LEDR is a systematic code (namely, the original data are included unchanged in the code) and therefore requires no decoder logic at the receiver side, saving power and latency.

The data cycle of the serial link cannot be less than one FO4 gate delay [14], [15] due to the digital logic forming the serializer and de-serializer circuits (Figs. 5 and 6, respectively), which consist of fast shift-registers that can deliver and consume one bit every FO4 delay. The fast shift-register (SR) is shown in Fig. 5. It comprises unique transition latches (XL). Each XL is controlled by a differential signal C/CN and consists of a dual-rail inverting control buffer and two separate data paths with XLs. Each XL consists of an inverter and a (weak) keeper that is switched off when the data bit is shifted. The differential

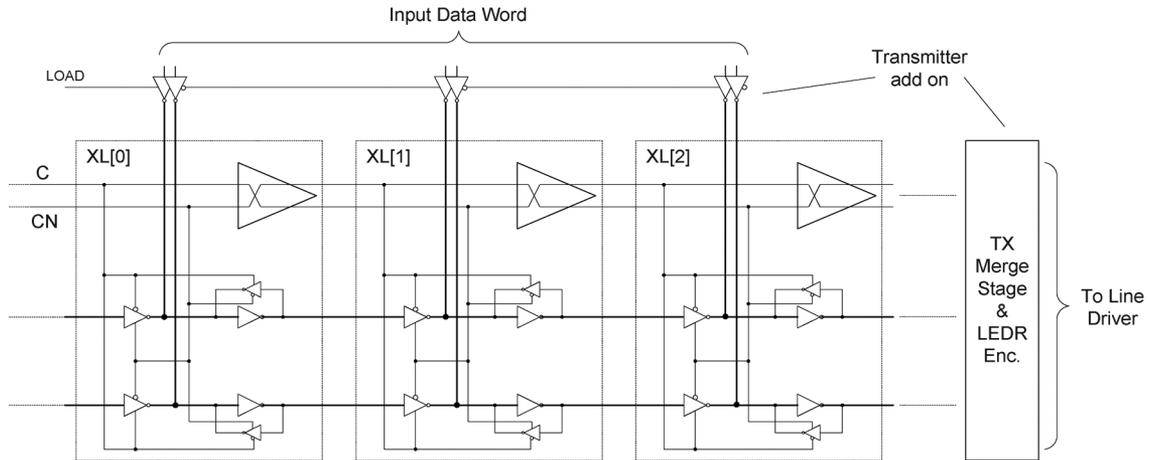


Fig. 5. One gate-delay shift-register architecture with transmitter connections (Serializer) [15].

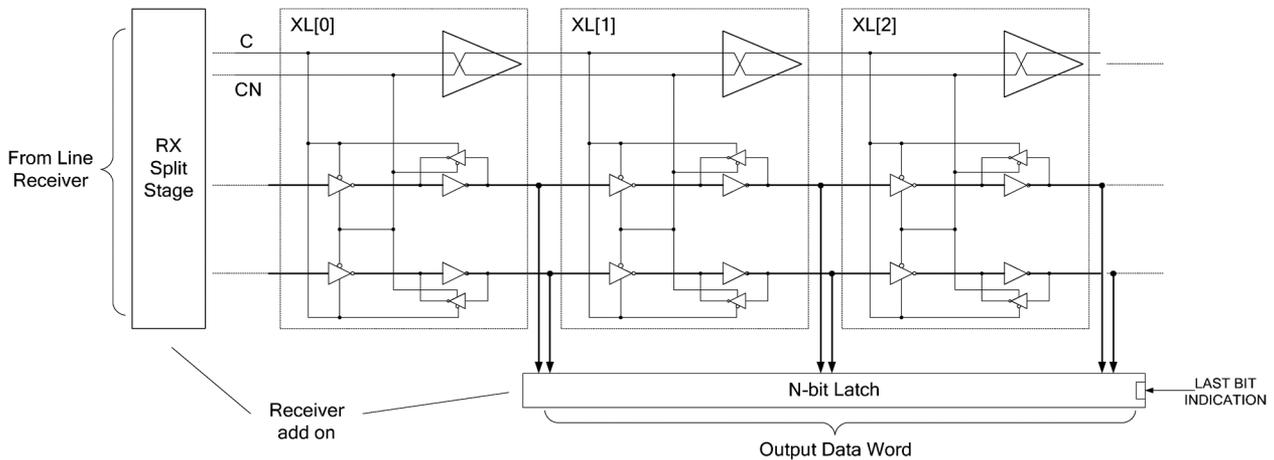


Fig. 6. One gate-delay shift-register architecture with receiver connections (De-Serializer) [15].

control lines C/CN are connected to the switches and keepers of each XL, such that when one switch is open, the other one in the same XL is closed, and the situation is reversed in the neighbor XLs. SR comprises at least two parallel pipes, when even bits are held in the bottom data-path and odd bits are held inside the upper one. At the transmitter (Fig. 5) the data is merged at the SR output. The input data is forked at the receiver (Fig. 6) into the two parallel pipes before the first XL. Control transitions on C/CN propagate without stopping through the control wave-pipeline, shifting data in the pipe. Note the double-data rate operation: data are sampled and shifted by both the rising and falling edges of C/CN. The control signals C/CN are generated using multi-phase clock generator [10]. The clock generator is triggered by the arrival of every newly transmitted word. To enable high-speed operation, the SR components should be properly sized [15].

For wide data words it is more power-efficient to partition the SR into smaller sub-SRs operating in parallel and at a slower rate. This partitioning helps to achieve linear dependence of SR power on its speed [14]. The SR is partitioned down to sub-SRs of eight data bits each.

Thanks to the partitioning (“splitter” architecture [14]), the sub-shift registers are not required to work at the shortest data

cycle of a single gate delay, but rather at data cycles that are at least twice longer. At the transmitter data are loaded by means of three-state gates connected to the XL latches (see Fig. 5). The capacitance presented by the three-state connection does not effect the SR operation since the SR is not required to work at the highest speed. The data coming out of the sub-SRs is merged and encoded by the Merge stage (Fig. 5, [14]). At the receiver size, the data coming out of the line receivers are forked by the Split stage [14] prior to being pushed into sub-SRs working at a reduced speed. The Merge and Split stages employ a few amplification stages (horns) for transistor size matching. The horns should be carefully laid out and contain retainers to reduce the skew.

The maximal data rate of the serial link at its parallel interface (N -bit parallel word pushed by sender, Fig. 1) is bounded by the time required for sending the N bits serially (namely N gate-delays) plus the time for loading the serializer and offloading the de-serializer (several gate delays). If the total required data rate exceeds that of the serial link, the serial link can be duplicated.

Special line driver and receiver circuits must be employed to support the high-rate wave-pipelined communication over the link interconnect. High capacitance interconnect incurs high dynamic currents for voltage mode signaling [14], limiting the link

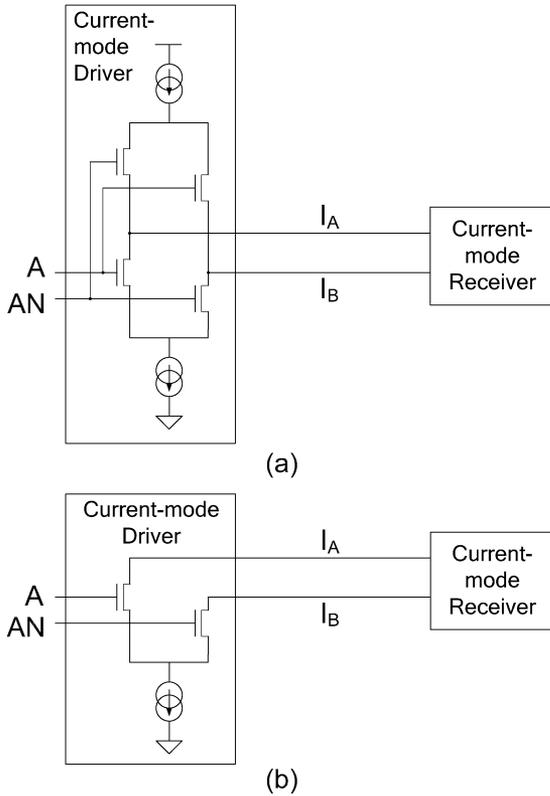


Fig. 7. Driving methods of current-mode signals: (a) bi-directional currents, $I_A = -I_B$, (b) the pull-down driver either blocks one of the two currents or generates two slightly different currents in the same direction.

performance. In the next section we explore novel asynchronous genuine current mode signaling circuits and wires. The circuits significantly reduce the voltage swing over the link interconnect and the wires facilitate current signaling and mitigate cross talks at very high speeds.

IV. HIGH SPEED CURRENT MODE ASYNCHRONOUS SIGNALING

A. Circuit Principles

The proposed novel on-chip communication circuit is comprised of three parts: current mode transmitter, channel wires and current-mode receiver [Fig. 7(a)]. These circuits fit in between the serializer and deserializer of Fig. 1. The differential circuits achieve higher speed over a longer range and better common-mode noise rejection (CMR) than single-ended circuits.

Several methods for differential current-mode drivers are available, producing different types of symbols. First, one wire may push current towards the receiver while the second wire of the differential pair pulls current backwards from the receiver [Fig. 7(a)], i.e., $I_A = -I_B$. Second, when using only pull-down drivers, one wire may be pulling current from the receiver while the second wire is disconnected, passing no current [Fig. 7(b)], e.g., $I_A > 0$, $I_B = 0$. Third, both differential wires may be pulling currents of two slightly different values. The driving circuit is similar to Fig. 7(b), but the voltage swings of A/AN are reduced, affecting transistor conductivity rather than

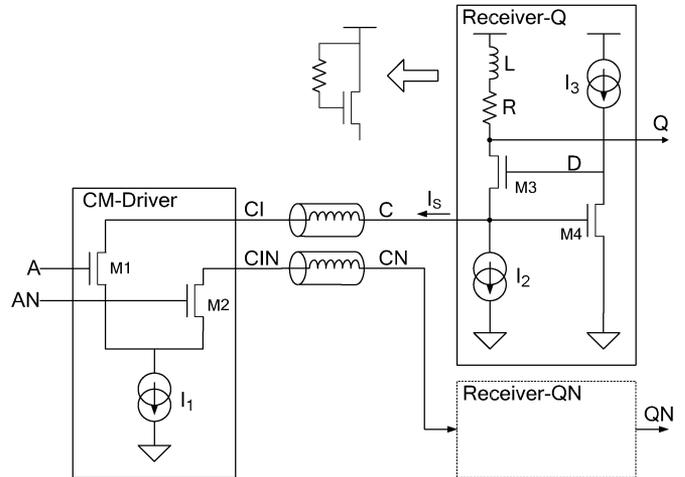


Fig. 8. Current-mode driver, channel and receiver.

turning them on and off. Either $I_A > I_B > 0$ or $I_B > I_A > 0$. The current swing is largest at the first method, reduced at the second method, and minimized at the third one. As is explained below, this current swing affects both link distance and operating frequency. The differential current-mode driver (CM-driver) in Fig. 8 draws current I_1 through either C or CN line according to the input A/AN.

The channel wires are designed as waveguides, enabling multiple traveling signals. At signal propagation velocity of at least $c/10$ ($30 \mu\text{m/ps}$) on a well-designed waveguide, and at the desired data rate of one bit per 15 ps (the expected FO4 inverter delay at 65 nm), a 1-mm wire may carry at least two successive bits simultaneously. At the planned data rate and wire dimensions, the link operates at the *RLC* region [49], requiring fewer repeaters than parallel wires which operate at the *RC* region [50]. The lines should be placed in the same metal layer to avoid inter-layer conductivity differences and speed degradation of vias. No termination is employed at the line ends. To facilitate the required throughput, high-metal (e.g., M5 and higher) wide lines are used. Using high metal layers increases the total interconnect area. However, the additional area is small relative to the size of multi-line parallel links. Skin effect is mitigated by line partitioning. Current return paths are placed in the same metal layer as the wires. We use the interconnect layout of Figs. 13 or 14 for crosstalk mitigation.

At the receiver side there are two identical but separate current mode receivers. The main concept of current sensing is to measure input current while minimizing the voltage swing on the channel and at the receiver's input port (C/CN), in contrast with typical "current mode" receivers that measure the voltage at the receiver's input.

The current mode receiver (Fig. 8) converts the input current into a low voltage swing signal on output Q. Inductance L in series with resistor R compensates for the input and output capacitance. Since a real inductor is difficult to achieve in silicon, active inductance can be employed (transistor and resistor in Fig. 8) [51]. The size of the active inductance is determined empirically to maximize output voltage swing. The receiver operates as follows. When no current is drawn by the channel and

the CM-driver, the current through R is I_2 (defined by the receiver current source) and the output voltage at Q is maximal:

$$V_Q^{\text{MAX}} = V_{CC} - I_2 \cdot R. \quad (2)$$

Once either A or AN switch on, current $I_S \leq I_1$ is drawn from the receiver through the channel, and the output voltage at Q drops to its minimal level:

$$V_Q^{\text{MIN}} = V_{CC} - (I_2 + I_S) \cdot R. \quad (3)$$

I_S is the channel current at the receiver input. When using the third current driving method of Fig. 7(b), I_S is not switched completely off, in order to enable higher speed. On the other side, $I_S \leq I_1$ due to losses over the channel, mostly due to channel length, wire aspect ratio, signal frequency and wire parasitic capacitances. Additional losses may be incurred by inductance, reverse waves (no termination) and noise.

Combining (2) and (3) we obtain the following expression for the maximal Q/QN output swing:

$$V_Q^{\text{MAX}} - V_Q^{\text{MIN}} = I_S \cdot R. \quad (4)$$

The voltage swing at output Q depends solely on I_S , namely the input current swing. The voltage swing grows when the transmitted current swing grows and channel attenuation is reduced. In other words, we should expect a linear dependence of performance on the invested power. Note that V_Q^{MIN} is bounded by the total voltage envelope over the driver transistors, the channel and V_{DS} of $M3$. Note also that the output Q/QN is designed to traverse a low voltage swing (positively biased close to the high rail voltage). All circuit transistors operate in their saturation region. I_2 should be reduced as much as possible to mitigate the standby power ($I_2 \ll I_1$).

The voltage swing at the input to the receiver is minimized by means of a very fast feedback loop consisting of $M3$ and $M4$. Receiver input C is a virtual ground point that is ideally kept at a constant voltage (as shown below, the achieved swing is as low as 50 mV). The input impedance of the receiver is very low. For fast operation the capacitance of node D should be minimized, facilitating fast charge/discharge following voltage fluctuations of input C . The entire feedback path from node C through $M4$, node D and $M3$ should be faster than the bit cycle time; the circuit simulated in this work achieved very short loop delay by using carefully tuned nMOS transistors, as demonstrated by the fast C and Q/QN waveforms of Fig. 17, discussed below.

The current-mode receiver circuit employs only nMOS transistors, contributing to its high-speed operation.

B. Adaptive Drive Control Circuit

The channel characteristic impedance and effective resistance depend on signal frequency as shown in Fig. 9 and, (5), where δ is the skin effect depth, R , L , and C are the interconnect resistance, inductance and capacitance per unit length, respectively and σ is shunt conductance [52], [53]. Note the 50% change in characteristic impedance value in Fig. 9. The impedance change

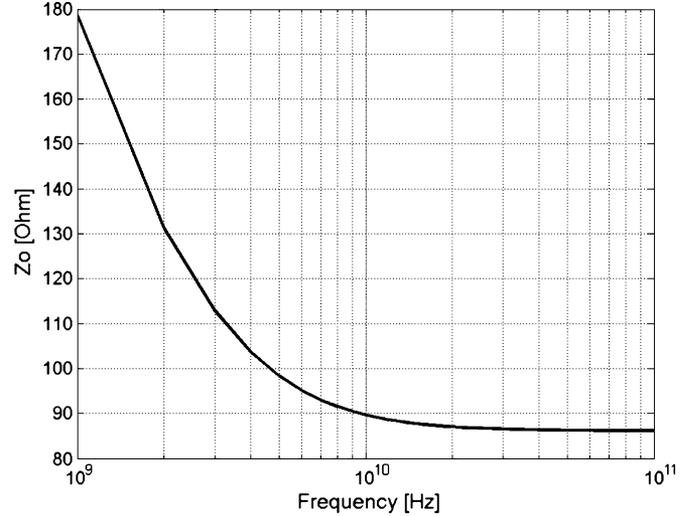


Fig. 9. Characteristic channel impedance frequency dependency. Based on data from PTM [54] and (5). Drawn for constant R , L , C .

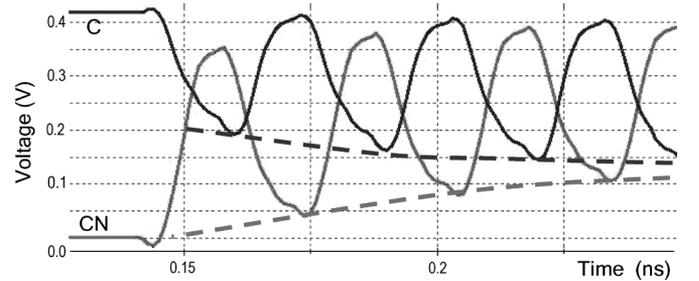


Fig. 10. Channel voltage/current drift due to frequency change.

between DC and high frequency values depends on the wire characteristics (process, width and length). Fig. 9 shows a typical case, while for real implementation the impedance change between DC and high frequency values may be different.

$$Z_0 = \sqrt{\frac{R + j\omega L}{\sigma + j\omega C}} \quad R = R_{DC} + \delta \cdot (1 + j) \cdot \sqrt{\omega}. \quad (5)$$

When signal toggling slows down or stops, the effective frequency is reduced, the effective resistance decreases towards its DC value but the characteristic impedance is increased. Conversely, when fast toggling resumes, the effective frequency is increased, the effective resistance grows due to the skin effect but the channel impedance decreases, calling for a stronger drive. These impedance changes lead to an undesired distortion of the channel differential signal. A typical example in voltage mode is shown in Fig. 10; similar results are observable in current mode.

To solve the frequency-dependent degradation problem, the transmitter circuit is amended by an adaptive control circuit, designed to compensate for changes in the effective channel impedance (Fig. 11). The inverters and AND gates constitute inertial delays and control a variable load on the driver output. The inverter chain delay is similar to the shortest data cycle (d_4). When the input is stable (or switches slowly), the drive

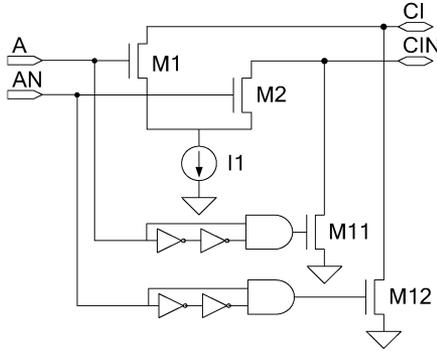


Fig. 11. Transmitter differential driver with adaptive control M11, M12 pull-down transistors provide variable load on the channel.

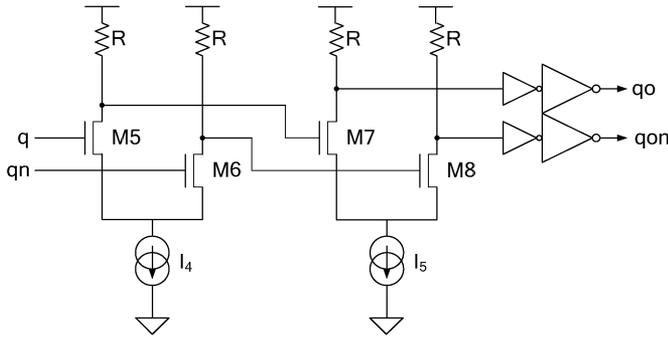


Fig. 12. Output amplification stage for full-swing reconstruction (here inductance L can also be placed in series to R for performance enhancement).

strength is reduced. When the input toggles fast, the AND gate never turns on and the drive strength is increased. The driver circuit is completely symmetric ($M1 = M2$, $M11 = M12$). The driver and receiver transistors are sized empirically to maximize receiver output swing, and should be adapted to link length and expected current in the link wires. For instance, in the 7 mm, $I_1 = 3$ mA link, the transistor widths are $M1 = M2 = 11$ μm , $M11 = M12 = 3$ μm , $M3 = 18$ μm , and $M4 = 8$ μm .

The adaptive control is capable of handling fast data transients. During a long period of no toggling, the characteristic channel impedance is high. Hence, once a new transition arrives, the first transmitted toggle of the channel may be distorted. The adaptive control mitigates this effect because it presents a reduced impedance to the drive at the time of this first toggle. Shortly afterwards, the AND gate turns off and the extra load is removed (see also Fig. 16). By that time, the channel impedance decreases, and normal transmission continues.

C. Receiver Output Stage

As mentioned above, the Q/QN output of the first stage of the receiver is a biased low-swing signal. Full swing must be restored before it can drive standard digital logic (receiver SR). The swing restoration amplifier output stage is shown in Fig. 12. It consists of two differential amplifiers. The first amplifier unbiases the signals and the second one creates the full swing differential output. Similarly to the circuit in Fig. 8, small inductors may be connected in series with resistors R in Fig. 12 for better performance.

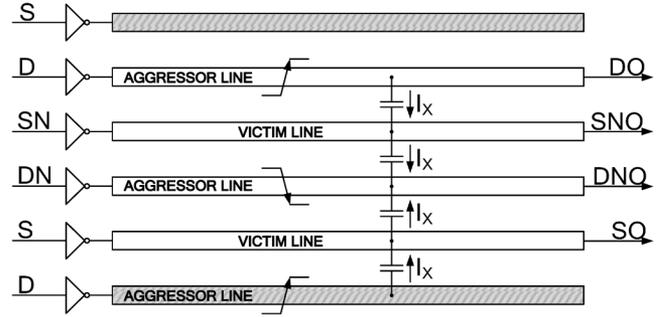


Fig. 13. Link wire arranged for active shielding: D, S are LEDR code bits, DN, SN are their inverted values. When D/DN lines toggle and S/SN are quiet, cross talk of D/DN on S/SN is cancelled out and S/SN lines serve as shields for D/DN lines.

When the link is not utilized, the driver and the receiver enter a standby mode, in which all current sources are switched off. This operation is controlled by an additional line from the transmitter. In standby mode the circuit consumes minor leakage current as described in the next section.

D. Wire Layout LEDR CM and VM Links

To reduce crosstalk over the interconnect we take advantage of the LEDR encoding characteristics: only one signal of D and S toggles per each transmitted bit. When LEDR symbols are signaled differentially, there are always two concurrent opposite voltage transitions per every transmitted symbol. In [14], a special version of active shielding [55] is employed for voltage mode signaling (Fig. 13). Two dummy S, D wires are added on the sides to actively shield the D, S signals, respectively. As noted above, all wires are laid out on the same metal layer. This arrangement minimizes cross talk and provides shielding as follows: Each toggling wire is surrounded by two quiet wires, and each quiet wire is surrounded by two wires that toggle in opposite directions.

Note that the structure of Fig. 13 not only mitigates the Miller effect (crosstalk caused by capacitive cross-coupling of adjacent wires), but it also reduces the proximity effect [52], since differential wires carrying opposing currents are separated by twice longer distance than if they were adjacent. If the wire width and separation are 1 μm , two neighboring wires that switch in opposite directions observe a 15% effective increase of resistance [52]. In the layout of Fig. 13, the separation of the differential pair for LEDR interconnect is 3 μm (instead of 1 μm), resulting in only 3% effective increase in resistance, namely five times smaller.

Most of the return current in Fig. 13 passes through the differential partners. However, the dummy wires create a mismatch and hence some of the return current passes outside this structure, constraining the maximal possible signaling speed. A balanced differential signaling layout has been adopted instead for current mode operation (Fig. 14). The two complementary wires of each pair are spaced wider than the minimum distance to minimize the proximity effect [52], and the two pairs are sufficiently separated from each other to avoid crosstalk. Recall that only one pair transfers a toggling transition at any one time in the same vicinity. SPICE simulations based on electromagnetic

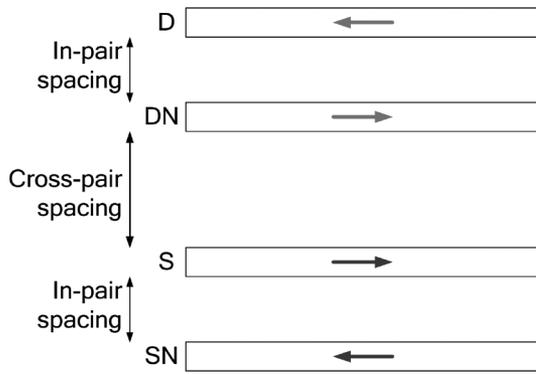


Fig. 14. Differential pair layout for LEDR code bits D,S. Return paths are well defined. In-pair spacing mitigates proximity effect; cross-pair spacing minimizes crosstalk.

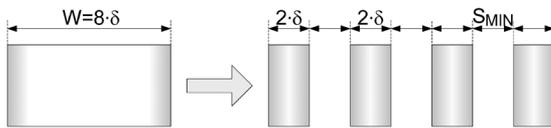


Fig. 15. Skin effect mitigation.

solvers indicate certain advantage to this latter method. Optionally, a shielding wire may be inserted between the two pairs, as well as on the outside of the entire structure.

Skin effects must also be addressed in the link interconnect design. There is a tradeoff between high-data rate and the length of interconnect: to communicate over a longer range we need to employ wider interconnect at higher metal layers. However, high-frequency switching causes skin effects, limiting the current through wide lines and thus constraining performance. Commonly, this effect is mitigated by partitioning wide lines into wires no wider than twice the depth δ of the skin effect (Fig. 15). Minimal spacing is employed between the wire partitions, and the interconnect contains merging stages along its length at constant intervals. Note that the frequency bandwidth considered in skin effects is related to the rise and fall times of the signals, about 10 ps (100 GHz).

Process variations in wires and in vias may significantly affect wire delays by increasing wire resistance [38]. In order to reduce the skew caused by process variations, the serial link wires should be placed close to each other and should not be interleaved with other wires. In addition, multiple via connections are essential, and minimal size vias should be avoided when possible. Increased wire resistance directly affects the performance of the serial link and may require speed reduction. However, since the serial link employs fewer wires and occupies smaller area than an equivalent-bandwidth parallel link, the serial link is less affected by variations [28]. A practical approach should employ speed adjustment and possibly link trimming after fabrication.

E. Performance

The link was SPICE simulated and shown operational at the highest speed of 67 Gb/s up to a range of 7 mm. Custom layout in 65-nm technology and link wires 2 μm wide were

used. Full shielding on the same metal layer was used, as described in Section IV-D. Signal and shield wires were partitioned down to 0.5- μm segments and there was 5 μm spacing between the signal and shield lines (layout of Fig. 14 with shielding was employed in the simulations). *RLC* parameters, obtained by Raphael-like three-dimensional field solver,¹ were employed for channel modeling. Various data patterns were simulated, covering fast and slow data rates as well as various bit combinations.

Fig. 16 shows circuit operation for an asynchronous input. Slow to fast transients are enabled by the adaptive control circuit; pull-down currents through M11, M12 can be observed in Fig. 16(c). For higher frequencies the current through the pull-down transistors is zero.

Voltage swing at the receiver input is lower than 50 mV [Fig. 17(a)]. The total voltage drop over the 7-mm interconnect is less than 100 mV and is even smaller for shorter links. These very low voltage fluctuations over the link result in very low currents consumed for charging and discharging the parasitic capacitances of the interconnect.

The current swing through the M3 transistor (Fig. 8) is about 1.2 mA [Fig. 17(b)]. This results in 200-mV voltage swing at the Q/QN output [Fig. 17(c)]; a 150- Ω resistor R was used in Fig. 8). The 200-mV output is unbiased and amplified by the output stage (Fig. 12). Note that the S and D paths should be placed very close to each other in order to reduce the skew as a result of process variations. In addition, trimming of I_4 and I_5 sources (Fig. 12) can help to optimize the performance after fabrication.

Power dissipated by the 7-mm link was measured using random pattern data signals. Note that the driver circuit has no power sources; rather, it draws its current from the receiver circuit through the channel. The total current drawn from the power supply was shown to be constant, and on average it was equal to the standby current. This happens thanks to differential operation of the circuit. The total average current that flows through the channel is similar to I_1 , 3 mA, and is about 30% of the total link current. Peak currents are observed only at the digital inverters at the output stage of the receiver; all other circuits consume a relatively constant current.

Leakage power was measured by turning off all current sources of the circuit (one in the driver, Fig. 11, four in the two receivers, Fig. 8, and two in the amplifiers of the receiver output stage, Fig. 12). Power results are summarized in Table II. The receiver consumes 10-mA current. Its dynamic power is 12 mW. Leakage power is negligible relative to dynamic power. Assuming 20% utilization of the link, the total power of a single link is 4.8 mW. A complete LEDR communication system comprises two current-mode links and its expected total power is about 10 mW. The power computation for SERDES SRs depends linearly on splitter architecture depth and word size [14]. For 16-bit SR with two sub-SRs under 20% utilization, the combined power of the serializer (in the transmitter) and deserializer (in the receiver) is about 25 mW. The synchronization circuits (Fig. 1) work at a lower rate than the serializer, driver, receiver and deserializer, since they operate on

¹From Synopsys Inc., available [online] at <http://www.synopsys.com>

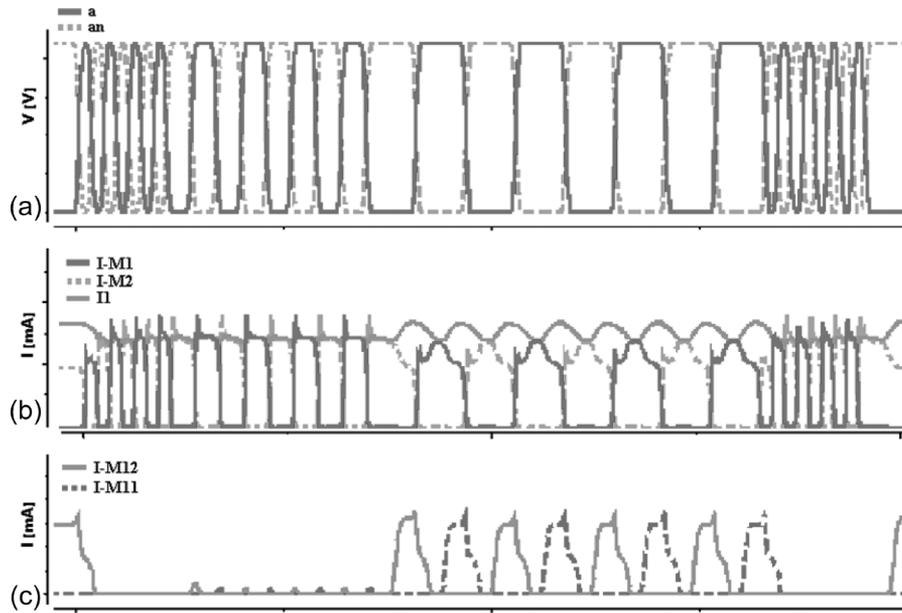


Fig. 16. (a) Input sequence (fast, slow, slower, fast); (b) transmitter current in transistors M1, M2; (c) adaptive control currents in pull-down transistors M11, M12: during no data or slow data current is drawn from channel, but during fast data there is no pull-down current.

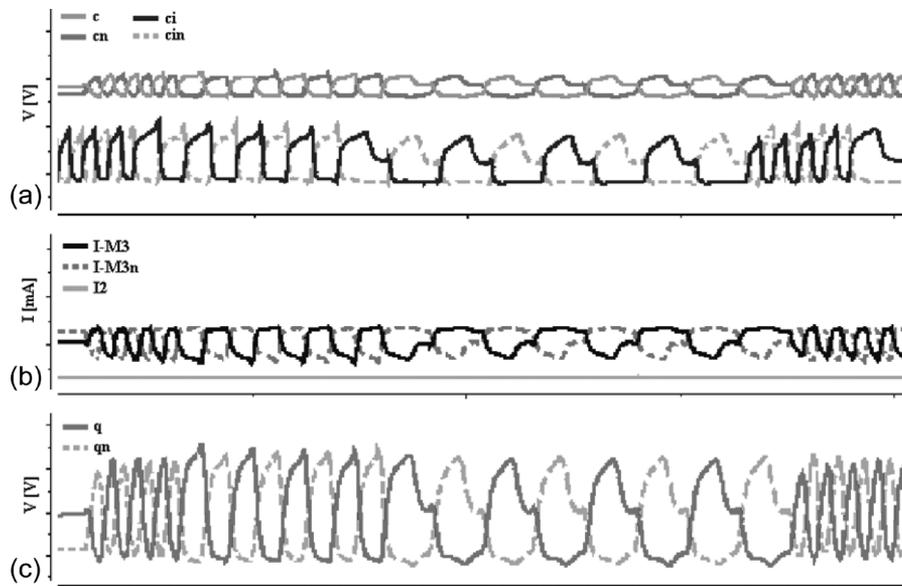


Fig. 17. (a) Channel swing: input swing of 150 mV, output swing: 80 mV; (b) receiver current through resistor (M3); (c) receiver output voltage (q, qn) of 200 mV.

parallel words, and therefore their power is insignificant. Thus, the total power of the link, when transferring 16 bit words at 20% utilization, is about $25 + 2 \times 4.8 = 35$ mW.

We have performed Monte Carlo simulations varying the threshold voltage V_T and the channel length L_{EFF} . These simulated in-die process variations, as well as low voltage corners, were found to affect mostly the first output stage amplifier of Fig. 12. This sensitivity may be mitigated by common adaptive trimming circuits that can compensate for process variations and for some voltage and temperature variations.

The obtained 7-mm range is nearly twice longer than the one achieved with the voltage mode circuits of [14]. This result matches the expectations of relative voltage and current mode performance from [33]. Therefore, the current mode commu-

TABLE II
POWER RESULTS

| Power (mW) | Receiver | Output Amplifiers | Total |
|----------------------------------|----------|-------------------|-------|
| Dynamic Power | 11.7 | 12.3 | 24.0 |
| Leakage Power | 0.07 | 0.08 | 0.15 |
| Peak Power | 11.93 | 16.50 | 28.4 |
| Total Power with 20% Utilization | 2.4 | 2.5 | 4.8 |

nication presented in this paper is preferred over voltage mode channels.

Current mode communication is more efficient than voltage mode in terms of dynamic power. The dynamic power of voltage mode driver and receiver of the 4-mm 67-Gb/s serial link [14] was 18 mW (under 100% utilization). The dynamic power of the 7-mm 67-Gb/s current mode link is 24 mW (Table II). Thus, only 33% increase in power is required for 75% longer link.

The power and frequency of this circuit can be traded off for link length. For example, by reducing frequency the link length can be extended. For shorter links, power dissipation can be reduced by readjusting the circuit current sources.

Future technologies, with scaled supply voltage, may limit the performance of the current mode link. Either speed or length may have to be traded off. To support high bandwidth, current mode repeaters, constructed of our line receivers and drivers, may be required.

V. CONCLUSION

This paper has shown for the first time that on-chip serial interconnect can achieve a data cycle as short as a single gate delay: the asynchronous serial link achieves throughput of 67 Gb/s. This achievement facilitates efficient on-chip communications in fast and large digital chips.

Novel high-speed on-chip serial links outperform parallel links for long range communication. The serial links occupy significantly smaller area, require less power and reduce routing congestion and noise. The relative improvement over parallel links grows with technology scaling. High-speed serial link with data cycle of a single FO4 gate delay was presented, enabling repeater-free communication over 7-mm distance at 67-Gb/s for 65-nm technology. The single gate delay serial link employs two-phase transition based LEDR encoding and differential signaling, and exploits wave-pipelining inside its SERDES circuits and over the link wires. Novel current mode signaling is explored for speed and range improvement. The novel current mode circuits support fast data rate transients, enabling both synchronous and asynchronous operation. Current mode signaling was found to be more efficient than voltage mode, enabling almost twice longer links at the same high speed. Novel low-crosstalk layout structures, especially suited for LEDR encoding, were presented. We advocate current mode asynchronous link as a viable alternative to common long-range wide parallel links.

ACKNOWLEDGMENT

The authors are grateful to Intel Research for supporting parts of this research and for providing feedback and advice. They also would like to thank their students D. Nahmanny, J. Moukabel, and R. Salach for performing the Cadence simulations of the current and voltage mode circuits.

REFERENCES

- [1] International Technology Roadmap for Semiconductors 2005 [Online]. Available: www.itrs.net/Links/2005itrs/Home2005.htm
- [2] W. J. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *DAC*, 2001, pp. 684–689.
- [3] A. Lines, "Asynchronous interconnect for synchronous SoC design," *MICRO*, vol. 24, no. 1, pp. 32–41, 2004.
- [4] G. De Micheli and L. Benini, *Networks on Chip*. San Mateo, CA: Morgan Kaufmann, 2006.
- [5] C. K. K. Yang, "Design of high-speed serial links in CMOS," Ph.D. dissertation, Stanford Univ., Stanford, CA, 1998.
- [6] S. Sidiropoulos, "High performance inter-chip signaling," Stanford Univ., Stanford, CA, Tech. Rep. CSL-TR-98-760, 1998.
- [7] W. F. Ellersick, "Data converters for high speed CMOS links," Ph.D. dissertation, Stanford Univ., Stanford, CA, 2001.
- [8] H. O. Johansson, J. Yuan, and C. Svensson, "A 4 Gsamples/s line-receiver in 0.8 μm CMOS," *VLSI Circuits*, pp. 116–117, 1996.
- [9] C. Svensson and J. Yuan, "High speed CMOS chip to chip communication circuit," in *Proc. IEEE Int. Symp. Circuits Systems (ISCAS)*, 1991, pp. 2228–2231.
- [10] M. J. E. Lee, "An Efficient I/O and Clock Recovery for TERABIT Integrated Circuits Design," Ph.D. dissertation, Stanford Univ., Stanford, CA, 2001.
- [11] J. Teifel and R. Manohar, "A high-speed clockless serial link transceiver," in *ASYNC*, 2003, pp. 151–161.
- [12] W. Bainbridge and S. Furber, "Delay insensitive system-on-chip interconnect using 1-of-4 encoding," in *ASYNC*, 2000, pp. 118–126.
- [13] R. Ho, J. Gainsley, and R. Drost, "Long wires and asynchronous control," in *ASYNC*, 2004, pp. 240–249.
- [14] R. Dobkin, Y. Perelman, T. Liran, R. Ginosar, and A. Kolodny, "High rate wave-pipelined asynchronous on-chip bit-serial data link," in *ASYNC*, 2007, pp. 3–14.
- [15] R. Dobkin, R. Ginosar, and A. Kolodny, "Fast asynchronous shift register for bit-serial communication," in *ASYNC*, 2006, pp. 117–126.
- [16] A. P. Jose, G. Patounakis, and K. L. Shepard, "Pulsed current-mode signaling for nearly speed-of-light intrachip communication," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 772–780, 2006.
- [17] I. Saastamoinen, T. Suutari, J. Isoaho, and J. Nurmi, "Interconnect IP for gigascale system-on-chip," in *Proc. Eur. Conf. Circuit Theory and Design (ECCTD)*, 2001, pp. 116–120.
- [18] T. Suutari, J. Isoaho, and H. Tenhunen, "High-speed serial communication with error correction in 0.25 μm CMOS technology," in *Proc. IEEE Int. Symp. Circuits Systems (ISCAS)*, 2001, pp. 618–621.
- [19] S. J. Lee, K. Kim, H. Kim, N. Cho, and H. J. Yoo, "Adaptive network-on-chip with wave-front train serialization scheme," *VLSI Circuits*, pp. 104–107, 2005.
- [20] A. J. Joshi and J. A. Davis, "Wave-pipelined multiplexed (WPM) routing for gigascale integration (GSI)," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 8, pp. 889–910, 2005.
- [21] A. J. Joshi, G. G. Lopez, and J. A. Davis, "Design and optimization of on-chip interconnects using wave-pipelined multiplexed routing," *TVLSI*, vol. 15, no. 9, pp. 990–1002, 2007.
- [22] S. Ogg, E. Valli, B. Al-Hashimi, A. Yakovlev, C. D' Alessandro, and L. Benini, "Serialized asynchronous links for NoC," in *DATE*, 2008, pp. 1003–1008.
- [23] W. P. Burleson, M. Ciesielski, F. Klass, and W. Liu, "Wave-pipelining: A tutorial and research survey," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 6, no. 3, pp. 464–474, 1998.
- [24] J. Xu and W. Wolf, "A wave-pipelined on-chip interconnect structure for networks-on-chips," in *Proc. High Perf. Interconnects Symp.*, 2003, pp. 10–14.
- [25] J. Xu and W. Wolf, "Wave pipelining for application-specific networks-on-chip," in *CASES*, 2003, pp. 10–14.
- [26] G. Lakshminarayanan and B. Venkataramani, "Optimization techniques for FPGA-based wave-pipelined DSP blocks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 13, no. 7, pp. 783–793, 2005.
- [27] C. Svensson and J. Yuan, "A 3-level asynchronous protocol for a differential two-wire communication link," *IEEE J. Solid-State Circuits*, vol. 29, no. 9, pp. 1129–1132, 1994.
- [28] R. Dobkin, A. Morgenshtein, A. Kolodny, and R. Ginosar, "Parallel versus serial on-chip communication," in *SLIP*, 2008, pp. 43–50.
- [29] E. Seevinck, P. J. Beers, and H. Ontrop, "Current-mode techniques for high-speed VLSI circuits with application to current sense amplifier for CMOS SRAM," *IEEE J. Solid-State Circuits*, vol. 26, no. 4, pp. 525–535, 1991.
- [30] M. Sinha, S. Hsu, A. Alvandpour, W. Burleson, R. Krishnamurthy, and S. Borkar, "High-performance and low-voltage sense-amplifier techniques for sub-90 nm SRAM," in *SOC*, 2003, pp. 113–116.
- [31] M. Sinha, S. Hsu, A. Alvandpour, W. Burleson, R. Krishnamurthy, and S. Borkar, "Low voltage sensing techniques and secondary design issues for sub-90 nm caches," in *Proc. ESSCIRC*, 2003, pp. 413–416.
- [32] A. Maheshwari and W. Burleson, "Current-sensing techniques for global interconnects in very deep sub-micron (VDSM) CMOS," in *IEEE Comp. Soc. Workshop VLSI*, 2001, pp. 66–70.

- [33] R. Bashirullah, W. Liu, and R. K. Cavin, III, "Current-mode signaling in deep submicrometer global interconnects," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 11, no. 3, pp. 406–417, 2003.
- [34] V. Venkatraman and W. Burleson, "Robust multi-level current-mode on-chip interconnect signaling in the presence of process variations," in *Proc. ISQED*, 2005, pp. 522–527.
- [35] S. Xu, V. Venkatraman, and W. Burleson, "Energy-aware differential current sensing for global on-chip interconnects," in *Proc. MWSCAS*, 2006, pp. 718–722.
- [36] V. Venkatraman and W. Burleson, "An energy-efficient multi-bit quaternary current-mode signaling for on-chip interconnects," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2007, pp. 301–304.
- [37] B. D. Winters and M. R. Greenstreet, "A negative-overhead, self-timed pipeline," in *Proc. ASYNC*, 2002, pp. 37–46.
- [38] L. Scheffer, "An overview of on-chip interconnect variation," in *SLIP*, 2006, pp. 27–28.
- [39] A. B. Kahng and R. O. Topaloglu, "Generation of design guarantees for interconnect matching," in *SLIP*, 2006, pp. 29–34.
- [40] R. Weerasekera, D. Pamunuwa, L. R. Zheng, and H. Tenhunen, "Minimal-power, delay-balanced smart repeaters for interconnects in the nanometer regime," in *SLIP*, 2006, pp. 113–120.
- [41] K. Moiseev, S. Wimer, and A. Kolodny, "Timing optimization of interconnect by simultaneous net-ordering, wire sizing and spacing," in *Proc. IEEE Int. Symp. Circuits Systems (ISCAS)*, 2006, pp. 329–332.
- [42] Y. Suwen, B. D. Winters, and M. R. Greenstreet, "Energy efficient surfing," in *Proc. ASYNC*, 2005, pp. 2–11.
- [43] A. Morgenshtein, I. Cidon, A. Kolodny, and R. Ginosar, "Low-leakage repeaters for NoC interconnects," in *Proc. IEEE Int. Symp. Circuits Systems (ISCAS)*, 2005, pp. 600–603.
- [44] R. Dobkin, R. Ginosar, and C. P. Sotiriou, "High rate data synchronization in GALS SoCs," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 14, no. 10, pp. 1063–1074, 2006.
- [45] M. E. Dean, T. E. Williams, and D. L. Dill, "Efficient self-timing with level-encoded 2-phase dual-rail (LEDRL)," in *ARVLSI*, 1991, pp. 55–70.
- [46] D. H. Linder and J. H. Harden, "Phased logic: Supporting the synchronous design paradigm with delay-insensitive circuitry," *IEEE Trans. Comput.*, vol. 45, no. 9, pp. 1031–1044, 1996.
- [47] *DS-DE*, IEEE 1355-1955 [Online]. Available: <http://grouper.ieee.org/groups/1355/index.html>
- [48] P. McGee, M. Agyekum, M. Mohamed, and S. Nowick, "A level-encoded transition signaling protocol for high-throughput asynchronous global communication," in *Proc. ASYNC*, 2008, pp. 116–127.
- [49] C. Svensson, "Electrical interconnects revitalized," *TVLSI*, vol. 10, no. 6, pp. 777–788, 2002.
- [50] Y. I. Ismail *et al.*, "Repeater insertion in RLC lines for minimum propagation delay," in *Proc. IEEE Int. Symp. Circuits Systems (ISCAS)*, 1999, pp. 404–407.
- [51] M. Green, *High-Speed circuits design techniques 2008*, private communication.
- [52] H. Johnson and M. Graham, *High-Speed Digital Design. A Handbook of Black Magic*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [53] —, "High bandwidth, low latency global interconnect," in *Proc. SPIE Conf.*, 2003, vol. 5117, pp. 126–134.
- [54] Predictive Technology Model (PTM), Interconnect Models 2005 [Online]. Available: <http://www.eas.asu.edu/~ptm>
- [55] H. Kaul, D. Sylvester, and D. Blaauw, "Active shields: A new approach to shielding global wires," in *Great Lakes Symp. VLSI*, 2002, pp. 112–117.



Rostislav (Reuven) Dobkin received the B.Sc., M.Sc., and Ph.D. degrees in electrical engineering from the Technion—Israel Institute of Technology, Haifa, in 1999, 2003, and 2009, respectively. His Ph.D. research focused on high-speed asynchronous communication for SoC.

Through the years 1997 to 2000, he worked within RAFAEL ASIC Experts design group, and between 2001 and 2002 he led a VLSI design group in IC4IC Ltd., developing a family of chips for communications. He is VLSI Group Manager with Ramon Chips

Ltd. He also leads a number of research projects in VLSI Systems Research Center at the Technion. His research interests are VLSI architecture, parallel architectures, asynchronous logic, high-speed interconnect, synchronization, GALS systems, SoC, NoC.



Michael Moyal (SM'03) was born in Tel Aviv, Israel. He received the B.S. degree in mathematics and physics from the University of Oregon, Eugene, the M.S.E.E. degree from Oregon State University, Corvallis, in 1982, and the Ph.D. degree from the Universität Der Bundeswehr Munich.

From 1982 to 1985 and 1988 to 1995, he was with the Communication Produce division of Advanced Micro Devices (AMD) as a Deputy Manager of SLIC/SLAC analog Line cards. From 1985 to 1988, he was with IXYS designing smart DMOS and Motion Control ICs. From 1995 to 1999, he was with Siemens AG (later Infineon Technologies) as Senior Manager responsible for the department of ISDN and xDSL transceivers. He is a Founder of Xignal and served on its Board and as CTO from 1999 to 2003 and as Chief Scientist from 2003 to 2005. He is currently with Intel Corporation, Israel, as a CTO analog Lan Access Division and a Lecturer at Tel Aviv University, Israel. His current interest are wide BW data converters and HF-CMOS, and transceivers architectures. He has published six papers in international journals and 18 patents.

Dr. Moyal serves on the ISSCC Technical program committee.



Avinoam Kolodny received the Ph.D. degree in microelectronics from the Technion—Israel Institute of Technology, Haifa, in 1980.

He joined Intel Corporation, Israel, where he was engaged in research and development in the areas of device physics, VLSI circuits, electronic design automation, and organizational development. He has been a member of the Faculty of Electrical Engineering at the Technion since 2000. His current research focuses primarily on interconnects in VLSI systems, at both physical and architectural levels.



Ran Ginosar received the B.Sc degree from the Technion—Israel Institute of Technology, Haifa, in 1978, and the Ph.D. degree from Princeton University, Princeton, NJ, in 1982, both in electrical and computer engineering. His Ph.D. research focused on shared-memory multiprocessors.

He worked at the AT&T Bell Laboratories on multiprocessor architectures from 1982 to 1983 and joined the Technion in 1983. He was a visiting Associate Professor with the University of Utah in 1989–1990, and a visiting faculty member with Intel Research Labs from 1997 to 1999. He co-founded several companies in various application areas of VLSI systems. He is an Associate Professor with the Departments of Electrical Engineering and Computer Science, and he serves as head of the VLSI Systems Research Center at the Technion. His research interests include VLSI architecture, many-core computers, asynchronous logic and synchronization, and networks on chip.