

A 22Gb/s, 10mm On-Chip Serial Link over Lossy Transmission Line with Resistive Termination

Hyo Gyuem Rhew, Michael P. Flynn
University of Michigan, Ann Arbor, MI
Email: {hrhew, mpflynn}@umich.edu

JunYoung Park
Samsung Electronics, Suwon, South Korea
Email: jy0401.park@samsung.com

Abstract—An on-chip serial-link scheme incorporating an interleaved voltage-mode driver, interleaved samplers and an optimally resistively terminated transmission line, enables an energy-efficient very-high-speed long-range data communication. The link is more than twice as fast and more than twice as energy efficient as the fastest reported on-chip link, yet has more than triple the communication range. A 10mm prototype link achieves a data rate of 20Gb/s with an energy consumption of 1.36pJ/b and a measured BER better than 10^{-11} . A 10Gb/s prototype achieves an energy efficiency of 680fJ/b with a measured BER of less than 10^{-13} .

I. INTRODUCTION

Global interconnect has suffered from technology scaling in terms of latency, power consumption, bandwidth and the complexity of floorplanning. Several techniques have been presented to improve global signaling, but these suffer from high latency, limited bandwidth [1, 2, 3, 4], design complexity of equalization circuits [3], and the need for placement of negative impedance converters along interconnect lines [5]. An on-chip transmission line with pulsed current-mode signaling offers higher bandwidth (8Gb/s), but achieves much less bandwidth density than densely packed wires employing repeaters [6]. A resistively-terminated transmission line link reported in [7] achieves a data rate of 9Gb/s but consumes 600mW.

We present a global on-chip serial link scheme that achieves high data rates, low latency, and high bandwidth through serial communication over an optimally-terminated lossy transmission line. Optimum resistive termination eliminates dispersion in a lossy link to enable a single long (>10 mm) wire bandwidth of more than 20GHz and eliminate the need for transmit or receive equalization. Improved transmitter, receiver and link design improves the energy efficiency by a factor of 37 compared to the 9Gb/s resistively-terminated link in [7]. An efficient interleaved voltage mode transmitter drives the line. Digital-like, interleaved, voltage mode drivers are used for simplicity, speed and energy efficiency. Interleaved 10GS/s comparators sample the received signal at an optimum phase dictated by a digitally controlled sampling clock. Compared to [7, 8], better link design enables a wire width reduction from $8\mu\text{m}$ to only $2\mu\text{m}$. The prototype 22Gb/s is fully self-contained and self-testing. We also present a 10Gb/s prototype that is optimized for lower speed.

We present measured data from two prototypes implemented in 65nm CMOS; one optimized for high-speed operation and the other optimized for low-power. The high

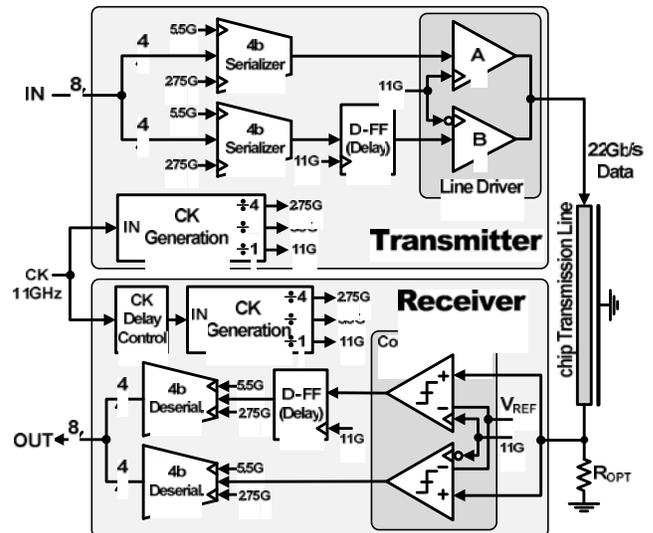


Fig. 1. Transceiver architecture with an on-chip transmission line link terminated with optimal resistance.

speed 10mm link achieves a measured energy-per-bit of 1.36pJ/b at 20Gb/s and an estimated latency of 54ps over a 10mm link. The 10Gb/s 10mm link achieves a measured energy-per-bit of 680fJ/b at 10Gb/s. The total receiver and transmitter circuit area for the 22Gb/s prototype is only $2450\mu\text{m}^2$.

II. SYSTEM ARCHITECTURE

Fig. 1 illustrates the overall system. The transmitter, shown on top of Fig. 1, serializes 8b parallel input data and two interleaved drivers launch data over a lossy on-chip transmission line. Termination of the transmission line with optimum resistance permits signal transmission over the high frequency Transverse Electro-Magnetic (TEM) mode. This termination scheme eliminates dispersion and enables high bandwidth communication while suppressing inter-symbol interference (ISI) [8]. A digitally-controlled delay chain sets the optimal sampling time at the receiver. A pair of interleaved comparators in the receiver samples the received signal. The recovered data is converted to CMOS logic levels and deserialized to an 8b parallel data stream.

A. Transmitter

A pair of digital-like interleaved voltage-mode drivers (A and B in Fig. 2) launches serial data on to the transmission line.

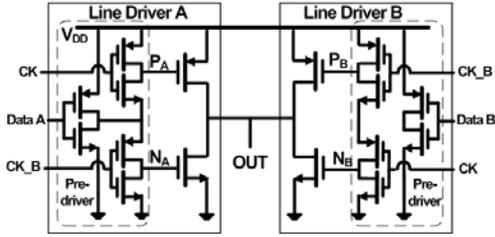


Fig. 2. Interleaved voltage-mode line drivers; A, B.

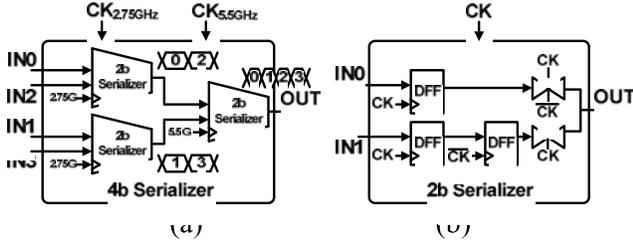


Fig. 3. (a) Block diagram of 4b to 1b serializer. (b) Schematic of 2b to 1b serializer

To simplify the driver circuitry and to allow the use of conventional static CMOS, two interleaved tri-state drivers, operating at 11GHz, feed 22Gb/s data to the lossy transmission line. A voltage-mode scheme is chosen to achieve a large voltage amplitude at the receive-end, enabling better noise immunity and easier decisions by receiver, compared to current mode schemes. The interleaved voltage mode driver requires only a single clock (i.e. CK and its complement CK_B) running at half the total data rate. Driver A drives the line only while the clock is high and driver B drives the line when the clock is low.

The pre-drivers (indicated in the dashed boxes in Fig. 2) enable just one of the four PMOS and NMOS driver devices depending on the polarity of the data.

As shown in the upper half of Fig. 1, the transmitter serializes 8b parallel data and launches the serial data onto the lossy on-chip transmission line. 8b 2.75Gb/s parallel data is first serialized to two interleaved 11Gb/s data streams by two 4b to 1b serializer. The voltage-mode line driver follows and serializes the two 11Gb/s streams to a single 22Gb/s data stream. To facilitate interleaved operation at the line driver, a half clock period delay is introduced in one of the 11Gb/s data streams.

A block diagram of the 4b to 1b serializer and a schematic of the 2b to 1b serializer are shown in Fig. 3. Two 2b to 1b serializers serialize parallel 4b 2.75Gb/s data to 2b 5.5Gb/s data streams with the help of a 2.75GHz clock, and then the next 2b to 1b serializer serialize them to a 11Gb/s bit stream using a 5.5GHz clock.

B. Receiver

A pair of comparators, at the end of the transmission line, samples the 22Gb/s data and deserialize to two parallel 11Gb/s data streams. Finally, two 1b to 4b deserializers deserialize these two interleaved data streams to a 2.75Gb/s 8b parallel bit stream.

The first stage of the receiver is a pair of comparators

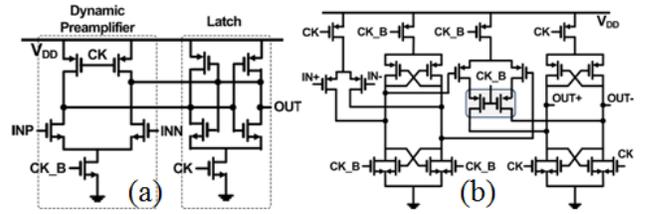


Fig. 4. (a) Schematic of comparator of high-speed serial link. (b) Schematic of comparator of low-power serial link.

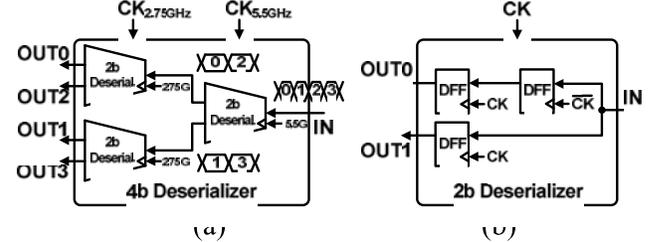


Fig. 5. (a) Block diagram of 1b to 4b deserializer. (b) Schematic of 1b to 2b deserializer.

operating in an interleaved fashion to sample the received signal (see Fig. 1). One comparator of the pair samples the data at the rising edge of the 11GHz clock and the other at the falling edge of the clock. Fig. 4 shows schematics of the comparators used in the high-speed and low-power link prototypes. In the high-speed system, a dynamic comparator is adopted for high speed operation [9]. In the low-power system, a two-stage structure achieves low metastability and low power consumption. Each stage of this two-stage comparator is comprised of a preamplifier followed by a latch. A PMOS pair is added in the second stage preamplifier to reduce kick-back from the output. A 61% reduction in kick-back noise is observed in simulation.

As shown in Fig. 1, since the two output data streams from the comparator pair differ in timing by half a clock period, one of the data streams is delayed by a half clock period to time align the two comparator outputs. Two 1b to 4b deserializers recover the original 8b parallel data. A block diagram of the 1b to 4b deserializer and a schematic of the 1b to 2b deserializer are shown in Fig. 5. In the 1b to 4b deserializer, a 1b to 2b deserializer deserializes 11Gb/s data to two 5.5Gb/s streams. Next two 1b to 2b deserializers deserialize the two 5.5Gb/s data streams to a 4b 2.75Gb/s parallel data stream.

C. Transmission Line Design and Termination of the Line

The characteristics of the prototype, 10mm long, 2 μ m wide microstrip transmission line are summarized in Table I. A 2 μ m line width is chosen to achieve both low driver power consumption and a large voltage amplitude at the receive-end. The estimated output voltage amplitude is 52% of the input voltage amplitude. The large receive amplitude means that comparator offset in the receiver is not a concern in this design.

Optimum resistive termination at receive-side of the lossy transmission line supports signal transmission over the high frequency TEM mode which enables signaling at the speed-of-light. This resistive termination eliminates dispersion

TABLE I
DESIGNED TRANSMISSION LINE CHARACTERISTICS

Distributed Resistance	Distributed Inductance	Distributed Capacitance	Characteristic Impedance (Z_0)
9.9 k Ω /m	476 nH/m	61 pF/m	88 Ω

due to the RC delay of low frequency components. The attenuated output voltage amplitude at the receive-end is [8]:

$$V_{OUT} = V_{IN} \exp\left(-\frac{R_L}{2Z_0}\right), \quad (1)$$

where V_{IN} and V_{OUT} are the input and output voltage amplitudes, and R_L and Z_0 are the total parasitic series resistance and characteristic impedance of the transmission line, respectively.

The optimum termination resistance at the receive-end of the transmission line, R_T , is chosen according to [8]:

$$\exp\left(-\frac{R_L}{2Z_0}\right) = \frac{R_T}{R_L + R_T}. \quad (2)$$

Since, according to (1), a larger characteristic impedance leads to larger output voltage amplitude, a microstrip transmission line with the highest metal layer (M9) as a signal line and the lowest metal layer (M1) as a ground plane is chosen to achieve the maximum characteristic impedance for the given technology. The estimated parasitic series resistance of 99 Ω for 10mm line results in an optimum termination resistance of 130 Ω , according to (2).

III. MEASUREMENT RESULTS

Both the 20Gb/s and the 10Gb/s transceivers, each with a 10mm microstrip transmission line terminated with a poly resistor are fabricated in 65nm digital CMOS technology. Die micrographs are shown in Fig. 6. The total transceiver area, including clock network and clock drivers, but excluding the transmission line, is 1580 μm^2 for the low-power prototype and 2450 μm^2 for the high-speed prototype.

Test data patterns are generated by an on-chip 16b PRBS generator. For the low-power prototype, the transmit and receive data streams are compared using a commercial BER tester. Because of limited bandwidth of this BER tester, BER measurements for the high speed prototype are made with an on-chip, custom, error calculation block. The energy consumption of the transmitter, receiver, and entire system (including clock drivers, but excluding the PRBS generator and error calculation block) are measured while maintaining a BER of less than 10^{-13} and 10^{-11} for the low-power and high-speed prototypes, respectively. As shown in Fig. 7, the high-speed type serial link achieves a maximum rate of 22Gb/s data rate, while the entire system consumes 1.36pJ/b at 20Gb/s. The low-power type consumes 680fJ/b at 10Gb/s. Although the latency cannot be measured directly from our test setup, EM simulations show the latency to be 54ps for a 10mm link, which corresponds to the speed of light latency in the SiO₂ medium of this technology.

Eye-diagrams at the receive-end of the transmission line, measured using dedicated probing pads and a 50 Ω GSG probe,

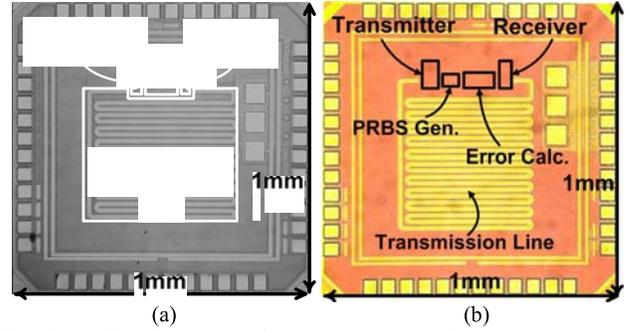


Fig. 6. (a) Die micrograph of low-power type serial link. (b) Die micrograph of high-speed type serial link.

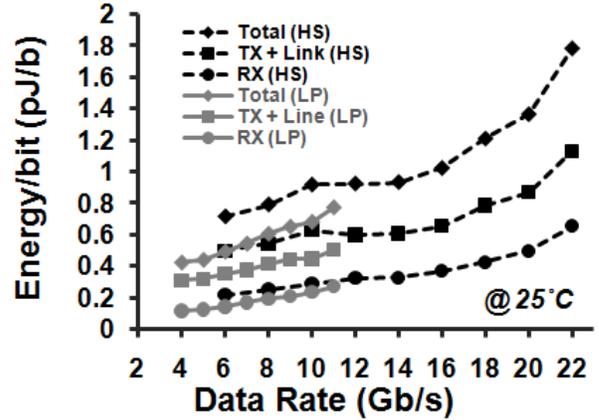


Fig. 7. Measured energy-per-bit and data rates at 25°C. HS: high-speed type serial link, LP: low-power type serial link.

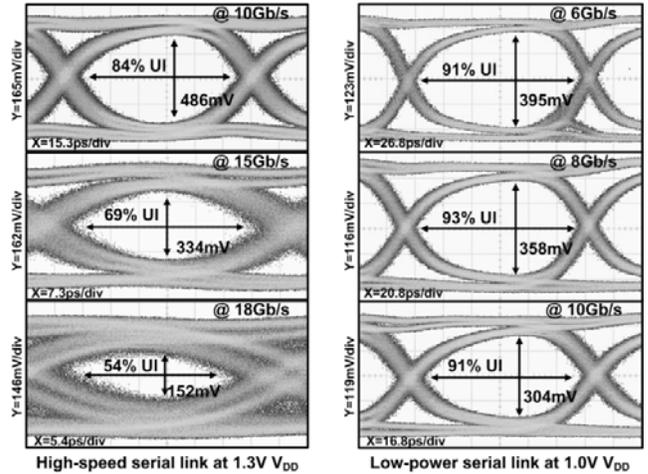


Fig. 8. Measured eye-diagram of the receive end of the 10mm serial links with a GSG probe at 25°C.

are shown in Fig. 8. Direct eye measurements were done for data rates up to 18Gb/s, which is the limit of the 15GHz GSG probes used. For the high-speed type serial link, the horizontal eye width is 54% UI and the vertical eye opening is 152mV at 18Gb/s. For the low-power type serial link, the horizontal eye width is 91% UI and the vertical eye opening is 304mV at 10Gb/s. The large eye opening proves the noise immunity of

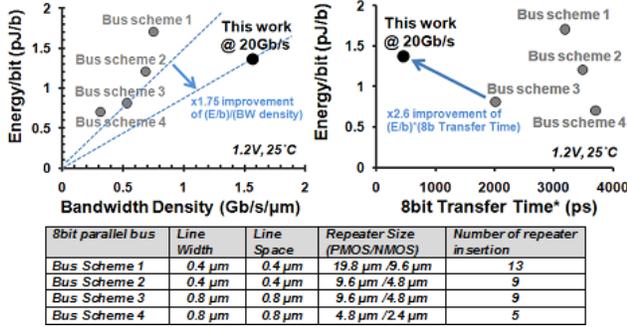


Fig. 9. Performance comparison with a conventional bus schemes that utilizes the same metal level (M9) at 1.2V VDD and 25°C. A crosstalk coefficient of 15% is assumed for the transmission line.

*8bit transfer time = (link latency) + (clock cycle)

the proposed scheme, and also demonstrates the effective elimination of dispersion and the suppression of ISI.

The serial link achieves better energy efficiency than conventional parallel busses with the same bandwidth density and data transfer time. Fig. 9 compares the measured performance of the prototype to the simulated performance of conventional parallel busses. For a crosstalk coefficient of 15%, the spacing between the 2μm wide M9 microstrip transmission lines is 10.8μm. We compare the 20Gb/s serial link with several 16bit and 8bit parallel bus schemes each with this same overall width and spacing, and each optimized in terms of energy-per-bit and bandwidth density. The transfer time of 8bit data over link is considered for fairer comparison instead of directly using link latency. The serial link shows a 75% improvement of (energy-per-bit)/(bandwidth density) and 2.6 times improvement of (energy-per-bit)/(8bit data transfer time) compared to the most optimized parallel bus.

IV. CONCLUSION

An on-chip serial transceiver with an optimally terminated transmission line for global signaling is implemented in 65nm CMOS technology. Table II summarizes the measured performance of the prototype and compares it to recently published on-chip global link systems. The prototype achieves the fastest data rate and the lowest link latency.

The combination of voltage-mode line driver, optimum

resistive termination, and SerDes techniques enables very low latency, high data rate and low ISI. The high-speed prototype device achieves a data rate of 22Gb/s with a BER lower than 10^{-11} . The 10Gb/s prototype device consumes 680fJ/b at a data rate of 10Gb/s with a BER lower than 10^{-13} . These links demonstrate the highest combinations of speed, energy efficiency and link distance of any on-chip serial link. Furthermore energy efficiency and bandwidth density are better than that of an optimized parallel bus in the same technology.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the assistance of Mohammad Mahdi Ghahramani.

REFERENCES

- [1] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-chip interconnects," *ISSCC Dig. Tech. Papers*, pp. 414-415, Feb. 2007
- [2] R. Ho, I. Ono, F. Liu, R. Hopkins, A. Chow, J. Schauer, and R. Drost, "High-speed and low-energy capacitively-driven on-chip wires," *ISSCC Dig. Tech. Papers*, pp. 412-413, Feb. 2007.
- [3] B. Kim and V. Stojanović, "A 4Gb/s/ch 356fJ/b 10mm equalized on-chip interconnect with nonlinear charge-injecting transmit filter and transimpedance receiver in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 66-67, Feb. 2009.
- [4] J. Seo, R. Ho, J. Lexau, M. Dayringer, D. Sylvester and D. Blaauw, "High-Bandwidth and Low-Energy On-Chip Signaling with Adaptive Pre-Emphasis in 90nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 182-183, Feb. 2010.
- [5] A. P. Jose and K. L. Shepard, "Distributed loss-compensation techniques for energy-efficient low-latency on-chip communication," *IEEE Journal of Solid-State Circuits*, vol. 42, pp. 1415-1424, Jun. 2007.
- [6] A. P. Jose and G. Patounakis, and K. Shepard, "Pulsed Current-Mode Signaling for Nearly Speed-of-Light Intrachip Communication," *IEEE Journal of Solid-State Circuits*, vol. 41, pp. 772-780, Apr. 2006.
- [7] J. Park, J. Kang, S. Park, and M. P. Flynn, "A 9Gbit/s serial transceiver for on-chip global signaling over lossy transmission lines," *IEEE Trans. Circuits Syst. I*, vol. 56, pp. 1807-1817, Aug. 2009.
- [8] M. P. Flynn and J. J. Kang, "Global signaling over lossy transmission lines," *Proc. IEEE/ACM ICCAD*, 2005, pp. 985-992.
- [9] C. C. Lee and M. P. Flynn, "A 12b 50MS/s 3.5mW SAR assisted 2-stage pipeline ADC," *IEEE Symp. VLSI Circuits, Dig. Tech. Papers*, pp. 239-240, Jun. 2010.

TABLE II
PERFORMANCE COMPARISON WITH RECENT WORK. HS: HIGH-SPEED TYPE SERIAL LINK, LP: LOW-POWER TYPE SERIAL LINK.

	Data Rate	BER	Eye Width (% UI)	Energy/bit (pJ/b)	Energy/bit/link length (pJ/b/10mm)	Transceiver Area	Link Latency (ps/mm)	Link Length
This work: HS	20Gb/s	<10⁻¹¹	54 at 18Gb/s	1.36	1.36	2450μm²	5.4 (from sim.)	10mm
This work: LP	10Gb/s	<10⁻¹³	91	0.68	0.68	1580μm²	5.4 (from sim.)	10mm
[3]	4Gb/s	<10 ⁻⁶	50	0.36	0.36	1760μm ²	N/A	10mm
[4]	4.9Gb/s	<10 ⁻¹⁰	N/A	0.34	0.68	N/A	N/A	5mm
[5]	3Gb/s	<10 ⁻¹⁴	N/A	2.0	1.43	N/A	12.1	14mm
[6]	8Gb/s	<10 ⁻¹⁴	N/A	> 3.4	> 11.2	N/A	6.4	3mm