

Ultra-Short-Reach Interconnects for Package-Level Integration

Anthony Chan Carusone, Behzad Dehlaghi
Department of Electrical and Computer Engineering
University of Toronto, Canada

Rudy Beerkens, Davide Tonietto
Huawei Canada Research Centre, HiSilicon Division
Ottawa, Canada

Abstract

Package-level integration of disparate technologies on either silicon interposers or organic packaging substrates is a key enabler for future optical subsystems. High-performance computing and networking applications will require tremendous interconnection bandwidth between co-packaged dies. To make package-level integration seamless, transceivers for these ultra-short-reach links must fit within vanishingly small area and power footprints.

Enabled in part by silicon photonic modules for coherent optical communication, the data rates processed by individual line cards has pushed beyond 1 Tb/s. Once converted into electrical signals, shuttling the data between chips on the line card is expensive both in terms of energy dissipation, and in terms of the circuit area required for the associated chip-to-chip transceivers. Consider the interface depicted in Fig. 1a between a CMOS die processing data and an optoelectronic transceiver relaying the data either within an equipment rack or beyond. State-of-the-art serializer-deserializer (SerDes) circuits at either end of such links employ sophisticated clocking and equalization techniques. In a next-generation ASIC required to process an aggregate bandwidth of 4-8 Tb/s, a conventional SerDes power consumption of 10 pJ/bit would represent more than half of a typical ASIC's power budget! The ASIC's total power budget is not likely to increase to accommodate the I/O since CMOS chips exceeding 100 W/cm² encounter difficult thermal dissipation challenges [1]. Therefore, assuming a total power budget of 130 W, and allocating 15% of the total power for I/O, an energy efficiency of 2.5 pJ/bit is needed.

An interest in transceivers for so-called "Ultra-Short-Reach" (USR) electrical interconnects has emerged to address this challenge. Two alternative solutions are depicted in Fig. 1b and 1c, both having the CMOS die and optoelectronic transceiver co-packaged offering improved signal integrity and area density. A model of the USR links, with parameters tailored to Fig. 1c, is summarized in Fig. 2. The model's electrical response is presented in Fig. 3 for 20 mm trace lengths. If a constant trace width is maintained (6 μm in Fig. 3a and 3b), differential striplines with ground shields offer the lowest crosstalk, whereas single-ended and unshielded links offer much tighter lane pitch and (arguably) still tolerable crosstalk. If constant lane pitch is maintained (12 μm in Fig. 3c and 3d), single-ended links offer wider traces, hence, lower loss. They also provide more space between neighboring channels than differential links, and therefore offer lower crosstalk. Furthermore, the 30-Gb/s eye diagrams in Fig. 4 demonstrate that single-ended USR links do not necessarily require double-termination; in fact, a low source impedance offers increased eye opening in this case.

Ultra-short electrical interconnect offers many benefits over conventional chip-to-chip interconnect over printed circuit boards. First, relatively low channel losses obviate the need for equalization circuitry, eliminating a significant consumer of circuit power and area. Second, the short distance and parallel interconnect allows circuitry to be shared among many parallel I/O transceivers, amortizing their power and area consumption. Specifically, circuits for clock generation, distribution, and alignment appear in each transceiver in conventional SerDes, but may be shared among many USR lanes [2]. Third, single-ended signaling can be used, reducing the number of pads and, hence, chip area required for electrical I/O, and reducing the number of low-impedance nodes that need be driven, hence, power. Example 20 Gb/s USR transceivers employing single-ended signaling are presented in [3], [4], in both cases resulting in less than 0.6 pJ/bit energy consumption. These benefits result in very compact transceiver circuits, allowing the benefits to compound: the on-chip distances over which high-speed signals must be propagated are reduced, thereby further reducing the area and power of high-speed buffer circuitry.

REFERENCES

- [1] G. Shahidi, "Evolution of cmos technology at 32 nm and beyond," in *IEEE Custom Integrated Circuits Conference (CICC)*, Sept 2007, pp. 413–416.
- [2] T. Dickson, Y. Liu *et al.*, "A 1.4 pJ/bit, Power-Scalable 16x12 Gb/s Source-Synchronous I/O With DFE Receiver in 32 nm SOI CMOS Technology," *IEEE J. Solid-State Circuits*, vol. 50, no. 8, pp. 1917–1931, Aug 2015.
- [3] J. Poulton, W. Dally *et al.*, "A 0.54 pJ/b 20 Gb/s ground-referenced single-ended short-reach serial link in 28 nm CMOS for advanced packaging applications," *IEEE J. Solid-State Circuits*, vol. 48, no. 12, pp. 3206–3218, Dec 2013.
- [4] B. Dehlaghi and A. Chan Carusone, "A 20 Gb/s 0.3 pJ/b single-ended die-to-die transceiver in 28 nm-soi cmos," in *Custom Integrated Circuits Conference (CICC)*, 2015 IEEE, Sept 2015, pp. 1–4.




Fig. 1: (a) Conventional PCB interface to an optical transceiver. (b) Ultra-short reach (USR) links on an organic packaging substrate; (c) USR links on a silicon interposer.

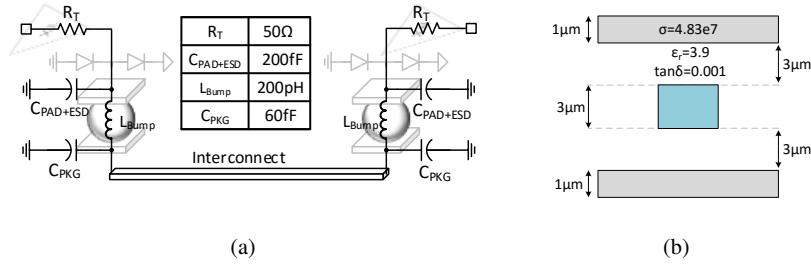


Fig. 2: (a) USR link model, with μ C4 solder bumps, interconnect, and on-die capacitances. (b) Cross-section of copper stripline interconnect.

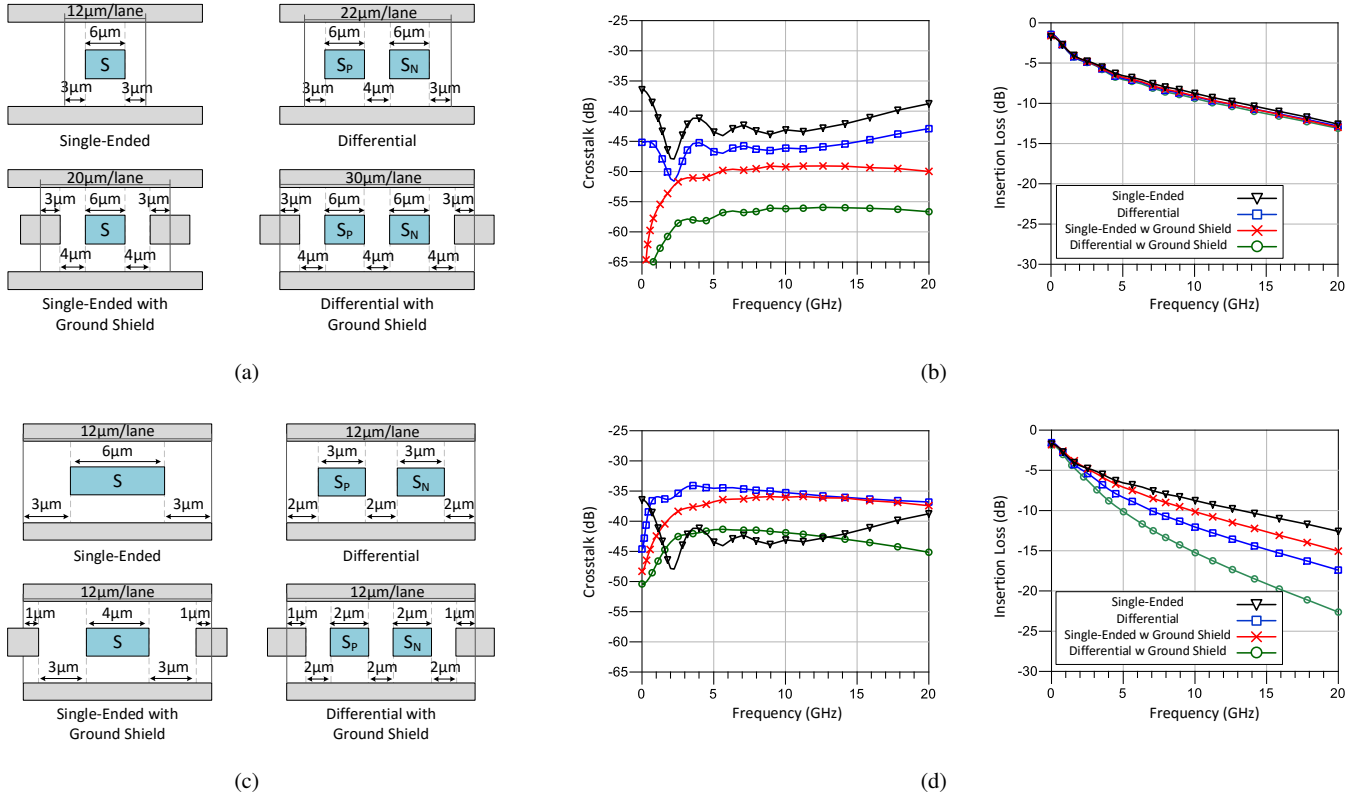


Fig. 3: Silicon interposer striplines 20 mm long maintaining the same trace width: (a) cross-sections; (b) crosstalk and through responses. Striplines 20 mm long maintaining the same lane-to-lane pitch of 12 μ m: (c) cross-sections; (d) crosstalk and through responses.

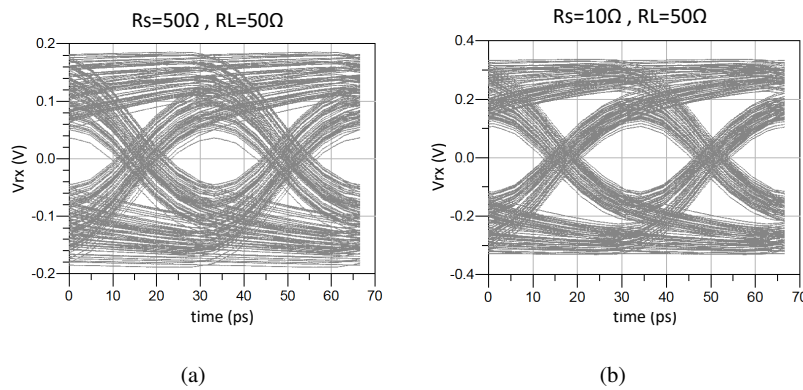


Fig. 4: Simulated 30-Gb/s single-ended eye diagrams with (a) 50-Ohm source and termination impedances; (b) 10-Ohm source and 50-Ohm termination impedances.