

Future State-of-the-Art Electrical Interconnect

Byungsub Kim, Intel

WINDS 2010 Invited Talk

Workshop on the Interaction between Nanophotonic Devices and Systems

Sunday, December 5th, 2010 • Atlanta, Georgia

Improving parallelism by integrating many cores in a single chip is the key scaling method in modern microprocessor development to continuously improve system performance under a tight power constraint. Since parallelism is maintained by tremendous data transfer through interconnects, the performance-cost efficiency of interconnect has become a critical design consideration. From this design standpoint, nanophotonic on-chip networks are attractive because optical communication potentially provides wider bandwidth at a lower power than conventional repeated electrical interconnects. However, comparing nanophotonics to simplistic electrical interconnect is not enough to determine the winning interconnect because electrical interconnects are also evolving by adopting equalization techniques. In addition, distance, area, and latency are also key design considerations. Therefore, researchers must compare nanophotonic interconnects with the best electrical interconnects, either repeated or equalized, considering all relevant design considerations. In this context, I will intuitively explain the fundamental trade-offs of repeated and equalized electrical interconnects and review the state-of-the-art of electrical interconnects as referencing counterparts of nanophotonic interconnects.

Byungsub Kim received B.S. degree in the department of Electronic and Electrical Engineering at Pohang University of Science and Technology, Pohang, Korea, and M.S. and Ph.D. degrees in the department of Electrical Engineering and Computer Science at Massachusetts Institute of Technology, Cambridge, USA. Dr. Kim was a co-recipient of the Beatrice Winner Award for Editorial Excellence at IEEE International Solid-State Circuits Conference in 2009. He was also awarded the Analog Device Inc. Outstanding Student Designer Award from MIT in 2009. His research interests include high-speed interconnects, computer-aided-design methodologies, and future on-chip networks. He has been with Intel Corporation at Hillsboro, Oregon since March 2010.

Future State-of-the-Art Electrical Interconnect

Byungsub Kim
Intel Corporation

Vladimir Stojanović
Massachusetts Institute of Technology

Abstract

This review intuitively explains the fundamental trade-offs of repeated and equalized electrical interconnects for on-chip networks as referencing counterparts of emerging nanophotonic interconnects in terms of data rate density, energy per bit, and latency. We also review the current state-of-the-art electrical interconnects from literature.

Keywords: *On-chip electrical interconnect, repeated interconnect, equalized interconnect.*

Introduction

Nanophotonics in on-chip network applications has become increasingly important because optical communication potentially provides wider bandwidth at a lower power than conventional repeated electrical interconnects. However, comparing only the bandwidth and power consumption of the two interconnect types is not enough to determine the winning interconnect because electrical interconnects are also evolving by adopting equalization techniques. In addition, distance, area, and latency are also key design considerations. Therefore, researchers must compare nanophotonic interconnects with the best electrical interconnects, either repeated or equalized, considering all relevant design considerations. In this context, based on the previous literature [1]-[11], we intuitively explain the fundamental trade-offs of repeated and equalized electrical interconnects and summarize the current status of electrical interconnects.

Fair Interconnect Metrics

For a fair comparison, we consider data rate density, energy per bit, and latency for a target distance in this review. Due to limited metal area, we want to maximize aggregated data rate for a given cross-sectional metal width. Therefore, data rate density, which is data rate per interconnect divided by the interconnect pitch, is a proper bandwidth metric. Energy per bit is a proper power metric since it is normalized and independent of any other conditions. Latency is another key metric affecting the system's utilization of functional blocks.

Trade-offs of Repeated Interconnects

Early repeated interconnect researchers focused on minimum delay [5], minimum delay-energy product [6], or maximum data rate density [7]. A recent work [2] presented three dimensional trade-offs among data rate density, energy per bit, and latency, demonstrating that we can trade off latency for data rate density while keeping the same energy per bit constant. For example, the two repeated interconnect designs in Fig. 1 burn the same energy per bit, but with different data rate and latency because both designs have the same total capacitance; the same elements (wire segments and buffers) are located differently. Using a simplified RC model, the delays of both interconnects are approximately

$$T_{da}=T_{dsa}=(R_{inv}/2+2R_{wire})(2C_{inv}+2C_{wire})=R_{inv}C_{inv}+R_{inv}C_{wire}+4R_{wire}C_{wire}+4R_{wire}C_{inv} \quad (1)$$

$$T_{db}=2T_{dsb}=2(R_{inv}+R_{wire})(C_{inv}+C_{wire})=2R_{inv}C_{inv}+2R_{inv}C_{wire}+2R_{wire}C_{inv}+2R_{wire}C_{wire} \quad (2)$$

where T_{dsa} and T_{dsb} are the corresponding segment delays. In reasonable designs, $R_{inv}C_{wire}$ terms are the most significant since the inverter resistance R_{inv} and the wire capacitance C_{wire} dominate. Although T_{dsb} is smaller than T_{dsa} , the total delay T_{db} is larger than T_{da} . Therefore, distributing repeaters as in Fig. 1 (b) is worse off for delay.

The minimum toggling period (reciprocal of data rate) can be approximated as 3 or 4 times the segment delay. Using 3, the minimum toggling periods are approximately

$$T_{sa}=3T_{dsa}=3(R_{inv}/2+2R_{wire})(2C_{inv}+2C_{wire})=3R_{inv}C_{inv}+3R_{inv}C_{wire}+12R_{wire}C_{wire}+12R_{wire}C_{inv} \quad (3)$$

$$T_{sb}=3T_{dsb}=3(R_{inv}+R_{wire})(C_{inv}+C_{wire})=3R_{inv}C_{inv}+3R_{inv}C_{wire}+3R_{wire}C_{inv}+3R_{wire}C_{wire} \quad (4).$$

The minimum toggling period is reduced by distributing repeaters as in Fig. 1 (b) compared to Fig. 1 (a) because the segment delay is reduced and the increased number of segments does not affect the minimum toggling period. Therefore, we can trade off delay for data rate by distributing repeaters differently while burning the same energy per bit.

Using more sophisticated models than the simplified and conceptual equations (1)-(4), Kim [2] presented trade-off contours of repeated and equalized interconnects in Fig. 2 when circuit and wire sizes are jointly optimized for the fair interconnect metrics using an aggressively scaled 32-nm technology model. The plot shows that we can trade off latency for data rate density under the same energy per bit constraint. If the energy budget increases, then the contour lines move toward higher data rate density and lower latency.

Trade-offs of Equalized Interconnects

Equalized interconnects can offer higher performance at a smaller power cost than repeated interconnects [1]-[4]. In equalization, the transmitter (Tx) and receiver (Rx) are high-pass-filter circuits (Fig. 3(a)), canceling the low-pass-filter characteristic of the wire so that the overall frequency response becomes flat. Without transmit equalization, a transmitted square pulse (dashed blue in Fig. 3 (b)) will be dispersed at the Rx (solid blue in Fig. 3 (b)), limiting the maximum toggling rate. An equalizing Tx sends a more sophisticated pulse (dashed red in Fig. 3 (b)) which is attenuated by the channel to a shorter and narrower pulse (solid red in Fig. 3 (b)) at the Rx . The equalizing Rx further reduces the pulse width by eliminating the DFE-tap in Fig. 3 (b). The narrower equalized pulse allows higher data rate.

Unlike repeated interconnects, we rarely trade off latency for data rate in equalization since equalization relies on flattening the channel's low pass filter response. Once flattened, the latency is the flight time of the electromagnetic wave, which is roughly determined by the wire distance.

Power consumption of the equalization strongly depends on the wire distance and the data rate. As shown in Fig. 3 (a), typically, the channel attenuation increases exponentially with

the wire distance and the square root of the data rate. To maintain the R_x pulse amplitude above a margin, the T_x controls the T_x amplitude. Therefore, the signaling energy exponentially increases with the wire distance and the square root of the data rate. In general, an equalized interconnect burns less power than repeated one because the effective voltage swing along the wire is typically smaller than V_{dd} . However, equalized interconnects are not always superior to repeated ones, especially for too short or too long distances, because the power overhead of the sophisticated high-pass-filter circuits used in equalization may exceed the power savings by voltage swing reduction or too large attenuation might be beyond the equalization ability. Recently, repeating equalized interconnects has been suggested to mitigate the impact of too large channel attenuation [11] for too long wires.

Fig. 2 shows the trade-off of equalized interconnects over a 10 mm distance in 32-nm technology when circuits and wires in Fig. 3 (a) are jointly optimized [2]. Unlike in repeated interconnects, in equalization, we cannot trade off latency for data rate density. The energy cost exponentially increases from 34fJ/b to 240fJ/b as the target data rate density increases. Due to exponential increment of channel attenuation, equalization can barely achieve 2Gb/s/um while repeated ones can achieve higher than 5Gb/s/um with additional energy and latency costs. Therefore, the best choice of the electrical interconnect type depend on many design considerations including energy budget and target data rate density.

State-of-the-art Electrical Interconnects

Fig. 4 shows the metrics of electrical interconnects in literature [3], [4], [9]-[11]. These are all equalized (and plus RF [10]) interconnects. Fig. 4 contains only measured silicon data. Over a 10 mm distance, equalized interconnects operate up to 2-3Gb/s/um with only 400-600fJ/b in 90-nm CMOS ASIC technology [3], [4]. An equalized interconnect over 5mm [11] in similar technology operates about twice the data rate density (~4Gb/s/um) for a similar energy cost (~400fJ/b) compared to the 10-mm ones because halving the distance exponentially reduces the channel attenuation. With more advanced technologies such as 22-nm technology, we can expect higher data rate density with a smaller energy cost.

References

- [1] B. Kim and V. Stojanović, "Equalized Interconnect for On-chip Network: Modeling and Optimization Framework," *ICCAD*, Nov. 2007.
- [2] B. Kim and V. Stojanović, "Characterization of Equalized and Repeated Interconnects for NoC Applications," *D&T of Computers*, Oct. 2008.
- [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 Technology," *ISSCC* 2009.
- [4] B. Kim and V. Stojanović, "An Energy-efficient Equalized Transceiver for RC-dominant Channels," *JSSC*, June 2010.
- [5] H. B. Bakoglu and J. D. Meindl, "Optimal Interconnection Circuits for VLSI," *IEEE Trans. Electron Devices*, May 1985.
- [6] R. Ho, "On-Chip Wires: Scaling and Efficiency," doctoral dissertation, Dept. of Electrical Engineering, Stanford, Univ., 2003.
- [7] V. Deodhar and J.A. Davis, "Designing for Signal Integrity in Wave-Pipelined SoC Global Interconnects," *IEEE Int. SoC Conf.*, 2005.
- [8] E. Mensink et. al., "A 0.28pJ/b 2Gb/s/ch Transceiver in 90nm CMOS for 10mm On-chip interconnects," *ISSCC* 2007
- [9] S. Tam et. al., "Simultaneous tri-band on-chip RF-interconnect for future network-on-chip," *VLSI* 2009.
- [10] J. Seo et. al., "High-bandwidth and low-energy on-chip signaling with adaptive pre-emphasis in 90nm CMOS," *ISSCC* 2010.
- [11] A. Joshi et. al., "Designing Energy-efficient Low-diameter On-chip Networks with Equalized Interconnects," *IEEE Hot Interconnects*, 2009.




Fig. 1. Two repeaters burning the same energy per bit and having different data rate and latency.

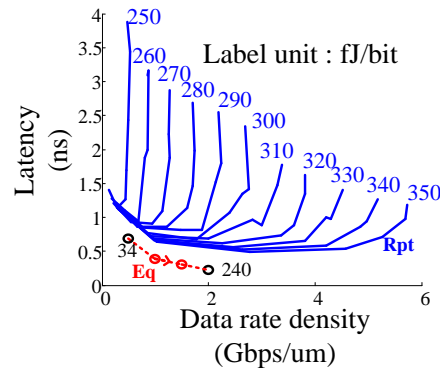


Fig. 2. Projected trade-offs for repeated (Rpt) and equalized (Eq) interconnects over 10mm-long wires assuming aggressive scaling [2]. The labels are energy per bit. Energy per bit cost is constant along an Rpt contour line and increases toward arrow direction along the Eq trade-off line.





Fig. 3. An example of equalized interconnect [1], [2].




Fig. 4. Current state-of-the-art electrical interconnects over 10-mm (a) and 5-mm (b) wires in 90-nm CMOS ASIC technology [3], [4], [8]-[10].