Low-power Data Communication Circuits for Advanced Integrated Systems

by Azita Emami

Advanced electronic systems require complex architectures that consist of many integrated circuit (IC) chips or modules. Examples of such systems are high-performance multiprocessors, servers, backplanes of internet routers, consumer electronics, and biomedical devices. There is also a growing demand for multiple processing and memory units integrated on a single chip. The high-bandwidth communication between these ICs and modules is a critical design issue. Large numbers of high-speed inputs and outputs (IOs) are required to create efficient interfaces between different processors, memory units and other modules located at varying distances from each other. With the continuous scaling of feature sizes in chip manufacturing technology, the speed of on-chip data processing as well as the level of integration will continue to scale. In order to enhance the overall performance of the system, the bandwidth of the interconnections needs to follow the same trend. To achieve this goal, scaling of the data rate per IO as well as the number of IOs per chip is necessary. While the increased switching speed of transistors allows faster transceiver electronics, the scaling of interconnect bandwidth has proven to be very difficult, and new approaches are necessary. The main challenges are: power and area limitations, properties of highly-scaled technologies, and characteristics of the signaling media.

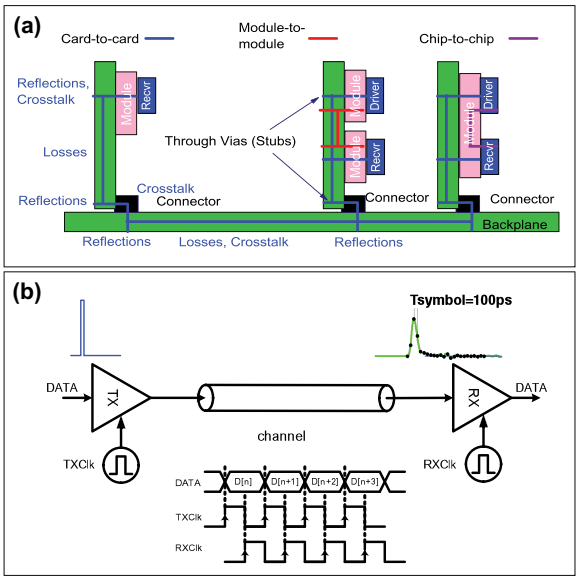

Figure 1: (a) Electrical links at different levels; (b) Synchronous data transmission.

Synchronous data transmission, shown in Figure 1, is the most common interconnection scheme in integrated systems. A precise clock signal and a driver are used to launch the data into the channel media. The channel is defined here as a complete signal path from the transmitter to the receiver, which includes the packages, Printed-Circuit-Board (PCB) traces, vias, and connectors (Figure 1a). Different modulation and coding schemes can be used for transmission of the information. The most basic and widely used scheme is binary data transmission by current or voltage amplitude modulation. The goal is to recover the transmitted bits at the receiver with a very low Bit-Error-Rate (BER) at very high speed (data rate). The information is recovered by the receiver, which samples the incoming signal using a synchronous clock.

Under support from the Lee center, we have studied new solutions for high-bandwidth chip-to-chip, module-to-module and intra-chip interconnections in highly-scaled technologies. We envision that a promising approach for data communication in advanced integrated systems will be an optimum combination of electrical, optical, and wireless signaling techniques. Accurate link modeling and understanding the fundamental limitations of different signaling methods have been among major goals of this effort.

Novel All-digital Synchronization Loop

One of the key building blocks of a synchronous data transmission system is the timing recovery loop at the receiver. Driven by power dissipation issues in high density circuits, timing and voltage margins are continually being reduced. It is therefore essential to precisely sample the signal at the time that the Signal-to-Noise-Ratio (SNR) is maximum. We designed a novel fully digital architecture that removes the need for an analog Phase-Locked-Loop (PLL). In this approach, an

extra clock phase is used to continuously monitor the incoming signal and to find the location of the transitions. After the completion of each monitoring period, the functionality of the data and monitor clocks are switched. The ping-pong action between the two clock phases is the key to the effectiveness of this architecture. A 10Gb/s data receiver based on this method has been implemented in a 90nm CMOS technology, and is currently under test and measurement. The results confirm that this receiver is capable of recovering an arbitrary amount of phase shift between clock and data without use of an analog PLL[1].

Low-power Optical Transceiver

Instead of furthering the complexity of electrical links, a different approach is to change the signaling media entirely. The possibility of using optics for interconnection at short distances has been a subject of considerable research and analysis. By providing a high-capacity channel, optical signaling can potentially close the gap between the interconnect speed and on-chip data processing speed. Although optical signaling involves electrical-to-optical (EO) conversion and vice versa (OE), its large channel bandwidth can simplify the design of the transceiver electronics. The maximum data rate of an optical link is in fact limited by the performance of the optical devices and the speed of on-chip electronics. The number of optical IOs per chip is usually limited by the power consumption and area of the electronics. For parallel optical signaling at short distances, one can either use fiber-bundles, integrated waveguides or free space to send collimated beams in parallel. In all cases the cross-talk among the beams is relatively low, avoiding another problem with large numbers of electrical IOs.

A three dimensional configuration for parallel optical interconnection in free space is shown in Figure 2.

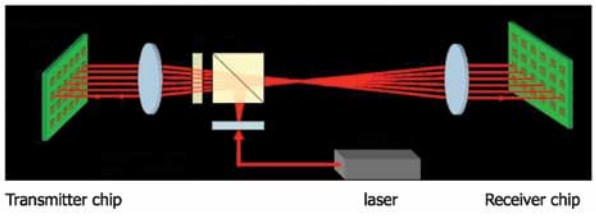

Figure 2: Parallel optical chip-to-chip interconnection over free space, with optical devices flip-chip bonded to CMOS chips.

The integration of dense two-dimensional arrays of optical devices with standard CMOS ICs has been demonstrated, and allows a huge chip-to-chip interconnection bandwidth[2,3]. In prior work, to achieve low-power and low-area optical IOs, a novel receiver front-end was designed using a double sampling/integrating technique. This approach facilitated a number of interesting scalable solutions such as parallelism, de-multiplexing, and an efficient baud-rate clock and data recovery. The transceiver operated at data rates as high as 16Gb/s and consumed less than 130mW[3]. These promising results proved the merits of optical signaling and the importance of continuing of this approach. Under support from the Lee center, we identified some of the shortcomings of the double sampling/integrating receiver and have found novel solutions for them. The integrating front-end suffers from headroom problem and needs DC-balanced signaling. In the new front-end, we are applying the double sampling technique to a resistive front-end. The added resistor converts the optically generated current signal to a voltage signal and automatically limits the voltage swing at the input node. The double sampling extends the bandwidth of the system beyond the time constant of the input node.

Efficient electrical, optical and hybrid signaling for on-chip and chip-chip networks are among the most important requirements of future computing and communication systems. We will continue to focus on understanding the fundamental limitations and developing of ultra-low-power solutions. Our effort will include optimum pulse shaping at the transmitter, novel cross-talk cancellation and equalization techniques, and further use of optics to enhance performance. LEE

Azita Emami is Assistant Professor of Electrical Engineering.

Read more at: http://www.mics.caltech.edu

References

[1] Matthew Loh, Azita Emami-Neyestanak, "All-Digital CDR for High-Speed High-Density I/O," submitted to the *2010 IEEE Symposium on VLSI Circuits*.

[2] A. Emami-Neyestanak et al. "CMOS Transceiver with Baud Rate Clock Recovery for Optical Interconnects," Proc. *IEEE VLSI Symp on Circuit*, June 2004.

[3] S. Palermo, A. Emami-Neyestanak, M. Horowitz, "A 90nm CMOS 16Gb/sTransceiver for Optical Interconnects," *IEEE Journal of Solid-State Circuits*, **43**, No. 1235-1246, pp. 889-896, May 2008.