

Techniques for Fully Integrated Intra-/Inter-chip Optical Communication

Claudio Favi and Edoardo Charbon

Ecole Polytechnique Fédérale Lausanne, Switzerland

Abstract¹ – In this paper we propose to eliminate all data and control pads generally present in conventional chips and to replace them with a new type of ultra-compact, low power optical interconnect implemented almost entirely in CMOS. The proposed scheme enables entirely optical through-chip buses that could service hundreds of thinned stacked dies. Very high throughputs and communication density could be achieved even in tight power budgets. The core of the optical interconnect is a single-photon avalanche diode operating in pulse position modulation. We demonstrate how throughputs of several gigabits per second may be achieved. We also show a systematic analysis and trade-offs of such a system and preliminary results to support its suitability in emerging DSM technologies.

Keywords – Intra-chip & inter-chip communication, low power optical communication, miniaturized optical channel and detector.

1. INTRODUCTION

Chip designers are developing increasingly complex integrated systems that require more and more die space for high throughput I/O pads. As a result, inter- and intra-chip communication is becoming one of the largest sources of noise and power dissipation on chip and also the bottleneck for performance. While transistor count has followed Moore's law; I/O pads have not evolved at the same pace. Moreover, due to bonding inductance, very high bit rates are possible at a cost of prohibitively high currents. Parallelism has often been used, but at a cost of large silicon area.

Traditional alternatives have been flip-chip and chip-level via technology [1]. However, reliability, cost, and flexibility are still open issues, especially when it comes to large inter-chip buses when more than two chips are involved. This is becoming an especially pressing problem with the emergence of multiprocessor and multi-core systems. For this reason, 3D and system-in-package (SiP) techniques have been conceived to enable stacks of inter-bonded dies (see Figure 1, left). The problem with this approach is however the limitation of speed imposed by bonding wires and the power dissipation of drivers.

To overcome these limitations, researchers have turned to wireless solutions based on capacitive, inductive, and optical methods [2],[3],[4]. While capacitive and inductive methods are effective in reducing power and ensuring high speed, they are only appropriate for pairs of chips. Hence, they are ineffective for broadcast and multi-chip systems.

Optical interconnects for inter- and intra-chip communication have been proposed decades ago. Their slow adoption is due mainly to the complexity of receivers whose output needs to be amplified, converted to a digital signal, and synchronized. These functions require area and may dissipate significant power. The lack of compact, low power optical sources has also been an issue.

In this paper a new approach to optical communication is proposed that can be integrated in standard CMOS technologies utilizing a fraction of the area and power of a pad. The proposed approach is amenable to optical broadcasting, optical clock distribution, and optical buses (both vertical and horizontal). The core of the optical interconnect channel is a CMOS single-photon avalanche diode (SPAD) [5]. The device can detect very low photon fluxes, thus ensuring minimal requirements of optical power at the source. Moreover, thanks to its digital output it requires no amplification, no A/D conversion, nor any other type of signal processing.

In SPADs the detection cycle can be as high as a few tens of nanoseconds. Thus, a simple digital modulation scheme must be added to achieve throughputs of several gigabit-per-second (Gbps). We selected pulse position modulation (PPM), a scheme that encodes K bits into 2^K time slots in the total allotted range R . Note that R should be higher than the detection cycle to ensure proper operation of the communication link.

The proposed principle is shown in Figure 1. Optical data signals are generated, for example, in an integrated CPU by a micro LED similar to [7]. A sub-nanosecond optical pulse was recently demonstrated for this device using CMOS drivers that occupy a fraction of the area of a pad.




Figure 1: Example of a conventional SiP (left); High-density inter-chip optical communication scheme.

The detecting section of the channel is represented by a SPAD and integrated PPM decoder. The total area of the receiving system is also a fraction of standard pads. The optical channel may be using integrated micro-optics that can be integrated on chip as a standard issue in most CMOS technologies. Multi-chip vertical buses can be obtained in this way by stacking dies that have been thinned. Optical transmission is ensured by low absorption coefficients of silicon in the visible spectrum.

2. SYSTEM ARCHITECTURE

The proposed optical interconnect architecture comprises a light source and an integrated driver, described in detail in [7], a detector similar to [5], and ultra-fast PPM coder/decoder logic. The PPM decoding process is achieved through a time-to-digital converter (TDC). This component must discriminate the time-of-arrival (TOA) of one or more photons as detected by the SPAD. The total allotted range R comprises a time window for TOA followed by a reset window, or TDC dead time (See Figure 2).

¹ This work was supported by Mobile Information & Communication System (MICS) and Swiss SNF.

We use two methods to interpolate the TOA: coarse and fine. Coarse TOA is measured by means of a counter running at system clock frequency (Fig.2-A). The coarse counter is used as a state machine to control the time window for fine TOA measurement that is achieved with a tapped delay line. When the photon-hit signal enters the delay line, the state of the complete line is latched on the rising edge of the clock. This yields a thermometer representation of the time between hit and the next rising clock edge. The fine controller (Fig 2-B) handles the conversion of the thermometer code onto binary so as to avoid metastability. Note that the delay line is not dynamically adjusted for temperature, voltage, or process variations. To achieve correctness we rely on regular calibration so as to ensure a fix bound on resolution.




Figure 2: TDC Architecture

3. PRELIMINARY RESULTS

We focus the preliminary results on the receiver. The TDC was implemented on a Xilinx XC2VP40 Virtex II Pro FPGA. The TDC design in FPGA is based on [6]. The system clock for our proof-of-concept is 200MHz. The fine chain must hence cover at least 5ns. From experimentation, a chain of 96 elements was sufficient to cover this time window with a maximum of 93 elements used at 20°C.




Figure 3: TDC characteristic differential non-linearity (DNL). We measured both integral (INL) and differential non-linearity (DNL). In Figure 3 the DNL is shown. The INL was below 1LSB. When coupling the TDC with a SPAD, the range must be adapted to the SPAD's dead time so as to keep potential errors due to jitter and afterpulse probability below a certain bound. On the TDC side the shorter the range the higher the throughput. We assume we can control the TDC design through only two parameters N and C . N is the number of fine delay elements and C is the coarse range bits that are used to extend the range by 2^C . Assuming that a single delay element yields to a delay of δ , the total fine range is $R_f = N\delta$. Therefore the measurement windows of the TDC as a function of N and C is

$$MW(N, C) = (2^C + 1)N\delta,$$

where one extra R_f is assumed for TDC reset. The achievable throughput TP as a function of N and C becomes thus

$$TP(N, C) = (\log_2(N) + C) / MW(N, C).$$

The SPAD detection cycle DC was chosen so as to match the range of the TDC.

$$DC(N, C) = (2^C)N\delta.$$

In Figure 4, $TP(N, C)$ in bps is shown by the gray shaded areas, $DC(N, C)$ in seconds by the solid lines.




Figure 4: TDC Throughput and SPAD detection cycle

4. CONCLUSIONS

This paper deals with some of the most important bottlenecks in VLSI design today, namely I/O communication and clock distribution. We propose to use CMOS compatible optical interconnect techniques based on miniaturized optical channels. In our approach we can achieve high throughputs at very low cost in terms of area and power dissipation, so as to represent a real alternative to conventional techniques. We have proposed means to implement such an approach in a systematic way and we have shown preliminary results. Further work in this area is being currently undertaken, including high-speed local clock synchronization, expected to drastically reduce clock distribution power costs with minimal or no area impact.

5. REFERENCES

- [1] K. Brown, "System in Package 'the Rebirth of SIP', *IEEE Custom Integrated Circuits Conference (CICC)*, pp. 681-686, 2004.
- [2] N. Miura, D. Mizoguchi, T. Sakurai, and T. Kuroda, Analysis and Design of Inductive Coupling and Transceiver Circuit for Inductive Inter-chip Wireless Superconnect, *IEEE J. of Solid-State Circuits*, Vol. 40, N. 4, pp. 829 - 837, 2005.
- [3] R. Drost, R. Hopkins, R. Ho, and I. Sutherland, Proximity communication, *IEEE J. of Solid-State Circuits*, Vol. 39, N. 9, pp. 1529-1535, 2004.
- [4] A. Huang, *et al.*, "A 10Gb/s Photonic Modulator and WDM mux/demux Integrated with Electronics in 0.13um soi cmos," *IEEE Intl. Solid-State Circuits Conference (ISSCC)*, pp. 922-929, 2006.
- [5] C. Niclass and E. Charbon, "A CMOS Single Photon Detector Array with 64x64 Resolution and Millimetric Depth Accuracy for 3D Imaging", *IEEE Intl. Solid-State Circuit Conference (ISSCC)*, pp. 364-365, Feb. 2005.
- [6] J. Song, Qi An, and S. Liu, A High-Resolution Time-to-Digital Converter Implemented in Field-Programmable-Gate-Arrays, *IEEE Trans. on Nuclear Science*, Vol. 53, N. 1, 2006.
- [7] H. Zhang, E. Gu, C. Jeon, Z. Gong, M. Dawson, M. Neil, and P. French, "Microstripe-array in GaN Light-emitting Diodes with Individually Addressable Elements," *IEEE Photonics Technology Letters*, Vol. 18, N. 8, pp. 1681-1683, 2007.

