2D parallel optical interconnects between CMOS ICs

O. Rits[a], M. De Wilde[b], G. Roelkens[a], R. Bockstaele[a], Richard Annen[c],
Martin Bossard[c], Francois Marion[d], R. Baets[a]

[a]Ghent University – IMEC, INTEC department, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium
[b]Ghent University – IMEC, ELIS department, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium
[c]Helix AG, Seefeldstrasse 45, CH-8008 Zurich, Switzerland
[d] CEA-LETI, optronics department, rue des Martyrs 17, 38054 Grenoble Cedex 9, France

ABSTRACT

The CMOS IC industry thrives on the down-scaling drive for ever smaller transistors, leading to faster, smaller and more complex digital systems. These ICs are interconnected by electrical tracks running on Printed Circuit Boards. Due to different frequency-dependent sources of signal degradation, the performance of these electrical interconnects lags behind the IC performance. As the electrical interconnect bottleneck increasingly impacts overall system performance, the interest in optical interconnects at the inter-chip level is growing. An important question to answer is how and where such optical interconnects should be implemented. Therefore, we first discuss the weaknesses of electrical interconnects and the potential benefits of optical interconnects. From this we then consider the implications on the introduction of optical interconnects and we argue why integration is of key importance for the successful introduction of optical interconnects at this level. Finally we describe how we have implemented optical inter-chip interconnects in a demonstrator system and go into more detail on the different levels of integration in this demonstrator system.

Keywords: parallel optical interconnect, opto-electronics.

1. INTRODUCTION

In the early days of digital electronics, the performance of the system was determined by the active components only. The electrical interconnection between two components did not limit the performance of the system. Things have slightly changed since then, one could say. CMOS Integrated Circuit (IC) technology is pushing for increasingly smaller transistors, increasingly larger chip-size and increasingly faster data rates. All this results in faster and more complex ICs, obeying Moore's Law , which dictates that data density in CMOS IC's shall double every 18 months.
The standard technology for data transport between CMOS IC's is based on metallic tracks on a Printed Circuit Board (PCB). Due to the above described evolution inside digital electronic systems, these electrical interconnects are becoming more and more difficult to design (requiring low-loss dielectrics and transmission lines with well-controlled impedance) and critical for the overall system performance (high-speed serial communication over backplanes requires overhead, such as SERDES circuitry). This so-called interconnect bottleneck is two-fold: firstly the data rate per interconnect channel is increasing and secondly the total number of required high-speed interconnect channels is increasing. In that respect, optical interconnections are regarded as viable alternatives to the copper PCB tracks. Such optical interconnects offer several advantages over their electrical counterparts, specially the possibility for increased bandwidth density, longer link lengths and possibly lower power consumption makes them attractive for introduction inside digital systems.

In section two, we start by analyzing the issues with electrical interconnects. In section three we first summarize the possible benefits of optics, secondly we will argue why integration is an important aspect for optical interconnects inside digital systems and finally we describe the demonstrator we designed and fabricated. In section 4, 5 and 6, we go into more detail on the different levels of integration.

2. THE ELECTRICAL INTERCONNECT BOTTLENECK

At low signal frequency, electrical interconnects can be modeled as a lumped RC (Resistor-Capacitor) network. At higher frequencies however, these electrical interconnects behave like transmission lines, characterized by a distributed

network of inductance, capacitance, resistance and conductance[1-3]. Based on this transmission line model, we can now identify different sources of signal degradation as signal frequencies keep on increasing and signal rise times keep on decreasing.

2.1 Transmission line loss

Ideally, a copper PCB track behaves as a lossless transmission line. However at high clock frequencies, copper tracks display two important sources of loss. First of all, the resistance of the transmission line will increase with the square root of frequency due to the skin effect. Secondly the dielectric conductance results in leakage currents through this dielectric material. Even with constant conductance, this dielectric loss increases proportionally with frequency. The combined effect of the skin effect loss and the dielectric loss results in increasingly higher loss at higher frequencies. The consequence of this frequency-dependent loss is that the interconnection length over which we can transmit a signal decreases as the signal frequency is increased.

2.2 Aspect Ratio Limit

Using the transmission line model and the frequency-dependent loss-effects described above, it can be shown that the maximal bandwidth is given by $B = B_0 * A/L^2$, where A is the cross-sectional area of the interconnect, L is the transmission length and B_0 is a constant[4]. The down-scaling drive of CMOS ICs implies that electrical interconnects need to cope with an increasing number of high speed channels operating at increasingly higher signal frequencies through a decreasing cross-sectional area. The aspect-ratio limit on the bandwidth of electrical interconnects tells us that eventually it will become impossible to transport the required aggregated bandwidth over electrical interconnects. Figure 1 illustrates the implications of the aspect ratio limit. Considering a B_0 constant of 10^{15} (see [4]) and PCB layer thicknesses of 12.6 mil, it shows that the maximum transmission distance decreases with decreasing track pitch.

It should be noted that this limit is not yet reached in experiments; the performance of electrical interconnections is more limited by practical considerations, of which the most important ones are described below.
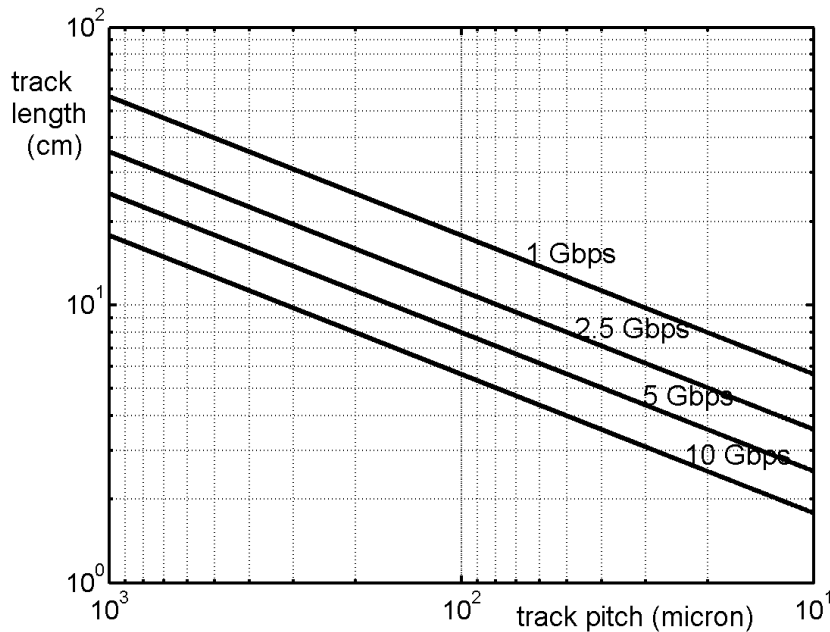

Figure 1: maximal track length vs track pitch at different channel bit rates as defined by the aspect-ratio bandwidth limit

2.3 Reflection Noise

Schematically seen, electrical inter-chip interconnects are simply straight lines connecting the driver and the receiver. Real electrical inter-chip interconnects are rather complex and are routed over the different levels of packaging: chip package, circuit board and subassemblies (Figure 2).
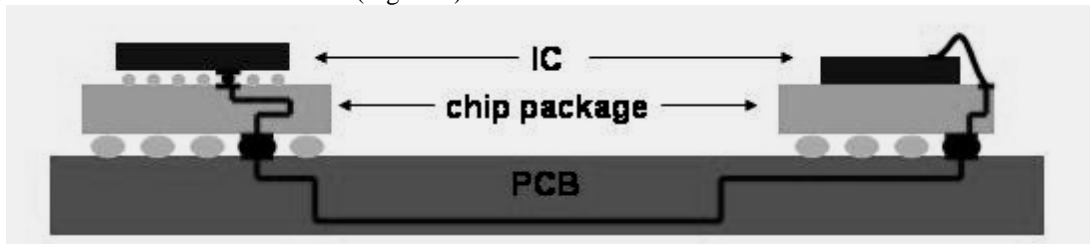

Figure 2: signal routing

More specifically, the signal goes off-chip via a C4 (Controlled Collapse Chip Connection) or wire-bond connection to the chip package. Through the solder bumps or I/O pins of this package, signals travel on to the PCB, towards the other chip package, again through the solder bumps or I/O pins of that chip package. Finally the signal goes back on-chip via a C4 or wire-bond connection. At the PCB level, signal paths encounter other discontinuities like vias, power/ground planes, stubs and backplane connectors. All these discontinuities in impedance added together cause multiple reflections, which in turn increase time delay and can produce overshoot, undershoot and ringing[1].

2.4 Crosstalk

Neighbouring PCB traces exhibit capacitive and inductive coupling. As a result, a signal traveling in one trace can couple to the other trace inducing crosstalk noise into that latter trace. This crosstalk is dependent upon the spacing between the lines, how long the lines run parallel to each other and last but definitely not least the signal rise time. Indeed the crosstalk is proportional to (dv/dt) in case of capacitive coupling and proportional to (di/dt) in case of inductive coupling. As a consequence, decreasing transistor switching times are making this crosstalk, be it capacitive or inductive, only worse. Different solutions exist to overcome this crosstalk (differential signaling, wider track spacing), but they all result in a decreased number of high-speed channels through a given cross-section.

2.5 Packaging and PCB Routing

As stated above, we have witnessed a strong evolution in the chip design and fabrication technology, and with it so has the chip packaging design and fabrication, which evolved from through-hole technology, via surface mount technology to Ball Grid Arrays (BGA) and Chip Scale Packaging (CSP)[5]. This evolution is to serve the continuing search for smaller form factors, smaller pin-pitches and higher pin counts. This evolution however, is also making the PCB routing increasingly complex. Especially the task of having to fan out a high number of signals coming from a tightly packed pin array through some kind of substrate is the most challenging one. Increasing the signal frequency doesn't exactly make this escape problem easier[5].

It is important to note here that, according to [1], the evolution in packaging design and fabrication has lagged considerably behind the evolution in chip design and fabrication. From [1] we quote: "… With current technology, the packaging interconnection delay dominates the system timing budget and becomes the bottleneck of the high-speed system design. It is generally accepted today that package performance is one of the major limiting factors of the overall system performance. … "

Sure electrical interconnect designers are engineering ways to cope with all these effects. Repeaters, equalization, pre-emphasis, differential signaling etc., all enable electrical interconnects to operate at ever higher bit rates. The downside of these clever tricks, is increased cost, increased circuit complexity, increased required silicon real-estate, and increased power dissipation. As such, these tricks are work-around solutions, trying to hide the real nature of electrical interconnects, which dictate that with increasing signal frequency, electrical interconnects will fall short.

3. OPTICAL INTERCONNECTS AT THE RESCUE

3.1 Benefits of optical interconnects

As stated in the introduction, optical interconnects are considered as viable alternatives for electrical interconnects[1,7,8,9]. Miller[7] reviews the physical reasons as to why optical interconnects are favorable over their electrical counterparts. The complete story is beyond the scope of this paper. In stead we will very briefly summarize some important aspects of optical interconnects and then move on to the implementation of optical interconnects inside digital systems.

First of all, optical waveguides not only provide a very low loss transmission medium, but also do not display any frequency-dependent loss. As a result, optical interconnects do not have any aspect-ratio-limited bit-rate capacity[1,7], offering a far higher bit-rate capacity density over relatively long interconnect lengths than electrical interconnects. Secondly, optical interconnects do not display any frequency-dependent crosstalk. Again, this aspect of optical interconnects makes them ideally suited for a faster and denser interconnection technology. Thirdly, thanks to the availability of low-power output devices (quantum-well modulators or low-threshold lasers) and low-loss transmission media, optical interconnects are capable of transmitting high-speed signals with substantially less power than electrical interconnects.

3.2 Optical Interconnects inside the box?

Optical fiber technology for long-distance telecommunication is a well established technology and has showed to be very successful at replacing its electrical counterpart. Introducing optical interconnects at the inter-chip level however, is a whole different ballgame and simply scaling long-distance optical fiber technology to the inter-chip level won't do the job. More specifically, the problem of introducing optical interconnects at the inter-chip level lies within the so-called last millimeter of the optical interconnects, where no established solutions exist (yet). The most important question to answer here is how and where are we going to integrate the optics with the electronics? Integration turns out to be of key importance for optical interconnects at the inter-chip level. The problem of the last millimeter comprises three different levels of integration:

Electro-optical chips integrated onto the CMOS
Optical interconnects possibly consume less power than electrical interconnects. This however, requires tight integration of the opto-electronic devices with their electronic driver and receiver circuits. Otherwise the total capacity of these devices will become too large, annealing the low-power advantage[10]. This calls for a dense integration of opto-electronic functionality with the electronic functionality. Miller[10] discusses in depth the possibilities for tight integration of optics with electronics.

CMOS integrated analog driver and receiver circuitry
Most optical interconnect technologies use dedicated electro-optical conversion modules, positioned close next to its corresponding data generating IC. The electrical signals from such a module to the data generating chip are routed over a short distance over some kind of substrate (PCB or Multi-Chip-Module). As mentioned earlier however, signal integrity is most critical at the chip packaging level. Furthermore, this approach does not alleviate the complex break-out task. To fully exploit the benefits of optical interconnects and overcome the critical issues of electrical interconnects, we need to provide the optical interconnects directly to the chip. This requires that we integrate the analog driver and receiver circuitry with the digital CMOS functionality. As analog and digital circuits have different boundary conditions and requirements, this integration is not at all an easy and straightforward task.

Package level integration of parallel optical ferrules
Optical interconnects do not suffer from any aspect-ratio-limited bit-rate capacity[4]. This means that optical interconnects will be more relevant and will be introduced more easily in high-aspect-ratio architectures, requiring a large amount of high-speed signals to be routed over a relatively long distance of transmission. These architectures are encountered for example in core IP routers, supercomputers and multi-processor systems.

The implications for optical interconnects are two-fold. First of all, optical interconnects need to be sufficiently parallel to accommodate for a high number of high-speed channels within the same form-factor as electrical interconnects[11]. This implies the need for 2D arrays of optical interconnects. Secondly, the optical signals will need to travel a relatively long distance. At the lengths of interconnects typical in the above mentioned applications, (10 cm to 1 m), guided wave optics are better suited than diffraction limited free-space optics. As a result we will need to extend the chip package with some kind of optical connector to accept a 2-D optical fiber ferrule.

3.3 The IO-demonstrator

In the framework of the IST-IO[12] (Interconnect by Optics) project, funded by the European Union, we have designed and fabricated an optical interconnect demonstrator. We developed a 2-D parallel inter-chip optical interconnect technology with direct on-chip optical access. In this section we will briefly discuss the general architecture of the system demonstrator and some technological choices. In the following sections we will discuss into more detail how we tackled the integration issues. More information on the IO-demonstrator can be found in [12].

General set-up

To accommodate for the 2-D parallel optical channels, we use 2-D arrays of InGaAs-based Vertical Cavity Surface Emitting Lasers (VCSELS) and InP-based photodiodes as the optical source and detectors. To provide on-chip optical access, firstly the digital electronics and the analog optical driver and receiver circuits are integrated on a single silicon die and secondly the 2-D electro-optical arrays are flip-chip mounted onto this die. It is important to note that the electro-optical components exhibit surface normal operation and are through-the-substrate emitting or illuminated devices. These System-on-Chip (SoC) modules are packaged in Ball Grid Array (BGA) packages and mounted onto PCBs, which are in turn assembled on a backboard. A 2-D array of MultiMode (MM) optical fibers serves as the transmission medium between two such packaged SoC modules. Since the electro-optical modules exhibit surface normal operation, these 2-D arrays of fiber ferrules arrive vertically to the chip and need to bent 90 degrees to run along the PCB (see figure3).
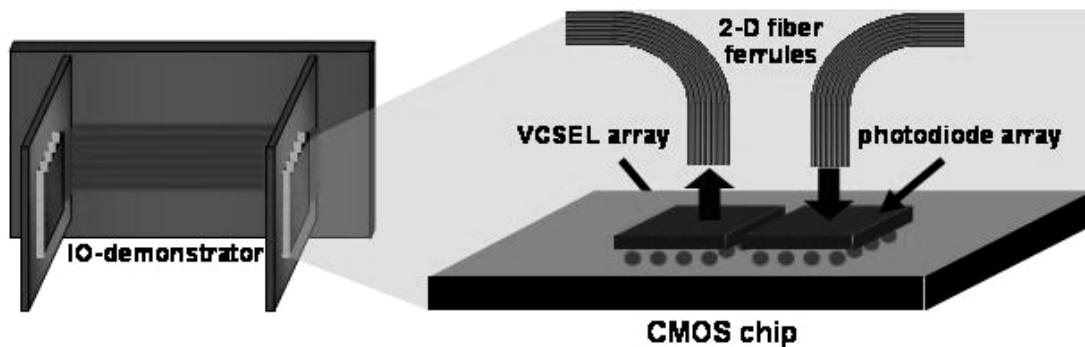

Figure 3: (a) IO-demonstartor (b) IO-chip module

Array specification

The array pitch size is 250 μm. This is a standard pitch size for optical interconnections and origins from the specifications of the MT-ferrule[14]. To allow for good coupling, all optical arrays have this pitch size (VCSELs, photodiodes and fibers). To minimize the interconnection length between the optical components and the analog driver and receiver circuits, these analog circuits are integrated directly under the electro-optical components and as such need to fit in this 250 μm pitch size. Table 1 shows the developed array sizes, their corresponding channel bit rate, aggregated bit rate and bit rate density.

Pitch = 250 μm			
Array size	8 x 8	16 x 16	8 x 8
Channel Bit Rate (Gbps)	1.25	1.25	2.5
Aggregate Bit Rate (Gbps)	80	320	160
Data Rate density (Gbps/mm^2)	20	20	40

Table 1: array specifications

Optical wavelength = 970 nm

The wavelength should be chosen such that it allows for good performance of all components of the optical link. This means high InGaAs VCSEL performance, high InP photodiode performance and low glass and plastic optical fiber loss. Another important aspect is the fact that the electro-optical components are through-the-substrate emitting or illuminated devices. The choice of the wavelength should minimize the absorption of energy in the GaAs and InP substrates. A detailed analysis showed that $\lambda = 970$ nm is the best choice for optimal total system performance.

# 4. CMOS INTEGRATED ANALOG DRIVER AND RECEIVER CIRCUITRY

4.1 Driver and receiver circuits

CMOS logic gates cannot directly modulate VCSELs and do not accept photodiode currents; therefore interfacing circuits are required. For each transmitter, a driver circuit converts a CMOS logic input into a corresponding current suitable for direct VCSEL modulation. At the receiver side, each photodiode's photocurrent is converted back to a CMOS logic output by a transimpedance (TIA) and limiting amplifier.

In our 0.35 µm CMOS demonstrator chip, driver and receiver cells are directly integrated in CMOS and located directly underneath the flip-chip mounted optoelectronic devices, arranged in the same 8x8 array format (figure 4 (a)). The driver circuit uses a shunt configuration. It pre-biases the VCSEL with the DC turn-on current and modulates the VCSEL by shunting a constant modulation current parallel to the VCSEL. This circuit concept minimizes power consumption and assures resilient operation for low crosstalk between channels. The driver features up to 1.25 Gbps operation speed and up to 8 mA mean and 8 mA peak-peak drive currents.

The transimpedance amplifier is implemented with two gm/gm gain stages. This concept assures a robust design over all CMOS process corners and, in addition, low crosstalk between channels. The limiting amplifier is simply an inverter, where the decision level is controlled by an active feedback loop. The receiver features speed capability up to 1.25 Gbps and a sensitivity of 40 µA peak-peak photocurrent, which is much higher than the thermal noise limited sensitivity. This way, we could significantly decrease the power consumption and improve the robustness against crosstalk between channels and digital logic.
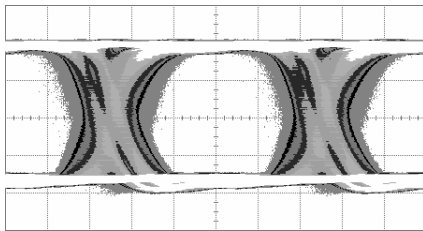
Besides the optical I/O channels located in the center of the chip, our 0.35 µm CMOS ASIC features on-chip pattern generation, bit-error-testing, retiming, and a switch allowing arbitrary connections between all optical interfaces.
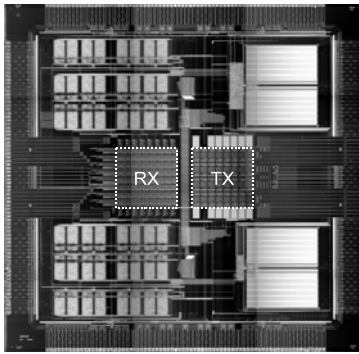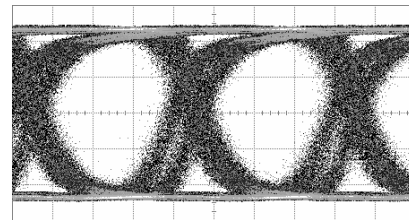

(a) 0.35 µm CMOS ASIC for direct hybridization
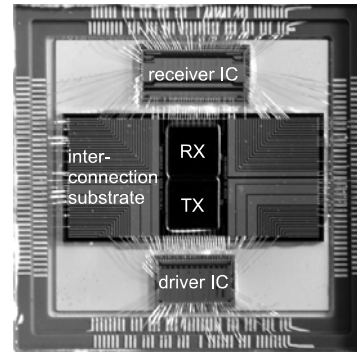
(b) Assembly with 0.18 µm CMOS circuits using wirebonding

(c) Eye diagram of link at 1.25 Gbps consuming 19 mW, using module type (a)

(d) Eye diagram of receiver at 2.5 Gbps and –19.8 dBm optical power, using module type (b)

Figure 4: Photographs of 0.35 µm and 0.18 µm ICs and measured eye diagrams at maximal rated frequency

We have also constructed an optical module using 0.18 µm CMOS driver and receiver circuits. Due to the high cost of this technology, we settled for an approach where separate driver and receiver ICs are wire bonded to an interconnection

substrate carrying hybridized 8x8 VCSEL and photodiode arrays (see figure (b)). In this case, we were limited to pure transceivers with twelve optical inputs and outputs.

In the 0.18 µm assembly, the driver circuits feature additional turn-on and turn-off current peaking in order to improve the VCSEL edge speed[15]. For the receiver circuits, utmost care was taken to avoid crosstalk, supply and substrate noise. The measures taken consist of a fully differential design throughout the circuit, as well as the provisioning of significant capacitive on-chip power decoupling. These measures are costly, and hence usually not taken unless the optical signaling is differential as well. However, this would cut the bandwidth of the optical interface in half. Instead, we use one photodiode per channel, with automatic threshold control via the other TIA input.

Experimental results on a single link indicate that the implemented circuits perform well. Using the 0.35 µm IC, we could demonstrate a complete optical link running at 1.25 Gbps and consuming a power of 19 mW. The 0.18 µm CMOS assembly runs at speeds up to 2.5 Gbps and its sensitivity is better than 10 µA peak-peak photocurrent. The measured electrical receiver output signals are shown in figure 4 (c)-(d).

## 4.2 Robustness to digital switching noise

One of the challenges of direct optoelectronic array hybridization is properly embedding the sensitive receiver circuits into a noisy digital IC environment. If we do not want to spend too much power on laser output and to allow some headroom for interconnection losses, the receiver circuits need to be sensitive. These receiver inputs—where the signals are smallest—are the most susceptible to digital noise.

Careful design—using highly decoupled separate power and ground nets and a layout with minimal coupling between sensitive metal tracks—leaves only coupling through the substrate as remaining noise mechanism. The term *substrate noise* comprises all effects caused by the switching of digital circuit nodes, which change the bulk potential underneath sensitive devices of the analog circuit, or inject current into substrate contacts.
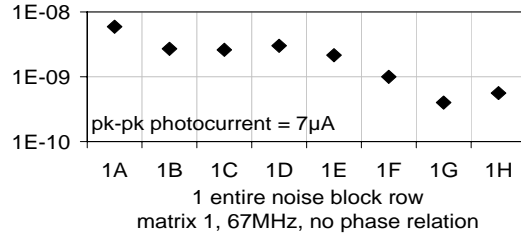
There are several sources of substrate noise. The most important reported *direct* causes[16-18] are capacitive coupling from MOSFET source and drain nodes and impact ionization from the MOSFET channel, which inject current into the substrate. In bulk-type substrates, the effect on the substrate is characterized by short, sharp pulses occurring nearly simultaneously with the switching events that cause them. These pulses affect the threshold voltage of the involved MOSFETs, and capacitive coupling with various circuit nodes is omnipresent. The coupling of this kind of substrate noise with our receiver circuits has been significantly reduced by surrounding all sensitive interface circuit blocks with substrate guard rings.

A major *indirect* cause of substrate disturbances originates from the power supply system and is known as *ringing*. Here, oscillations of digital supply and ground nets—caused by sharp supply and ground current fluctuations—are coupled with sensitive receiver supply and ground nets through a relatively low-impedance connection between the typically numerous substrate contacts of these nets. The ringing effect is normally larger in amplitude and time scale than the direct effect when bondwire-based packages are used. When packaging parasitics are reduced—for instance when the IC is being flip-chipped instead of wirebonded—ringing decreases significantly[19], but the direct effect remains.

We have investigated the effect of direct substrate noise on the receiver circuits of our 0.18 µm CMOS assembly (for an in-depth discussion, see De Wilde et al.[20]). For our measurements we have used dedicated circuits for noise generation and analysis. Our experimental setup was inspired by Nagata et al.[21]. The noise generation circuit has a matrix-based layout with 8 rows of 32 noise cells. The input clock of the noise generation circuit propagates from row to row with an adjustable delay and drives all cells within a row simultaneously. A noise cell switches a capacitance of 15 fF against the substrate. Every cell can be individually programmed to switch its inverters with the clock, with the inverted clock, or not at all. This approach yields a highly customizable noise source in terms of noise location, magnitude, injected current direction and timing.
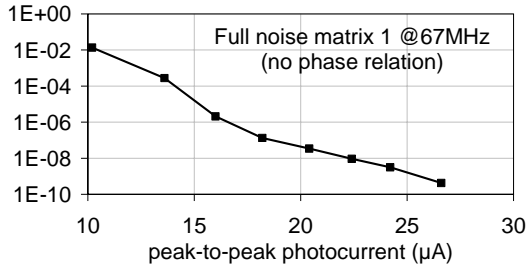
(a) Receiver BER vs. pk-pk photocurrent (at RX12 1GHz PRBS 2^{31}-1)	(b) Receiver BER vs. noise block row

Figure 5: Impact of substrate noise on the 0.18 µm CMOS receiver circuits

We have quantified the impact of substrate noise on the bit error rate (BER) of the receiver circuits. Figure (a)-(b) illustrates that noise generation did affect the observed BER, yet this occurred only at very high noise and low incident optical power levels. Analysis of the measurement data has revealed that bit errors were introduced during the significant power system ringing caused by strong noise generation, but also at much lower noise levels before the onset of the power system response. The latter effect was positively attributed to direct substrate interaction with protection circuitry inside pad cells of a very sensitive analog biasing net. This effect accounts for the noise block row dependency in figure 5 (b). Measurements at the output of the receiver circuits and simulations with our calibrated substrate model have confirmed this coupling mechanism. The observed sensitivity can be eliminated either by removing the protective diodes or by adding an extra ground-connected guard ring. In summary, here, the noise cells have provided a good diagnostic for a very real direct substrate noise coupling threat that otherwise would have remained unnoticed. The substrate noise resilience of the receiver circuits is otherwise excellent.

5. ELECTRO-OPTICAL CHIPS INTEGRATED ONTO THE CMOS CHIP

When working with a 2-D array of components, flip-chip technology is ideally suited for dense integration, optimal overall system performance and minimal power consumption. In our system demonstrator, both VCSEL and photodiode arrays are flip-chip mounted on the CMOS using an indium bump technology[22-23]. The indium bumps allow for a high-yield and highly reliable flip-chip. The ductile property of the Indium decreases the mechanical stresses at the ball joints. The CMOS wafers are first post-processed: the solder is deposited on the wafer and reflown. Due to the self-aligning effect of the flip-chip technique, the accuracy on the position of the bumps is better than $\pm$ 0.5 micron. After the flip-chip mounting, a polymer underfill is applied to increase the yield of the flip-chip connections and serves as protection, during further high-temperature processing. Finally the opto-chips are mechanically thinned to 50 mm, and an anti-reflection coating was applied to both the VCSEL and photodiode arrays.

Figure 6 : (a) CMOS bumping, (b) hybridization VCSELs and photodiodes,
(c) underfill, (d) thinning, (e) anti-reflective coating (f) dicing

6. PACKAGE-LEVEL INTEGRATION OF PARALLEL OPTICAL FERRULES

Generally a chip package has three functionalities: it serves as a mechanical protection for the brittle and sensitive CMOS die, it takes care of the thermal management and it provides electrical interfaces to the outside world. For implementing direct on-chip optical access, the chip package will need to provide also optical interfaces to the outside world as well.

As described in section 3.3, the SoC module is assembled in a customized Ball Grid Array package (figure 7.a). It consists of a ceramic package with 5 layers (pads, ground, Vcc and 2 signal layers), 320 I/O pads on a pitch of 1mm, and CuW base plate for dissipation[23]. The BGA has a down-facing cavity. We place a cooler on top of the BGA package for heat removal, meaning that we need to place the optics at the other side. As such, the chip will emit light and will be illuminated through a hole in the PCB (see figure 7.b)
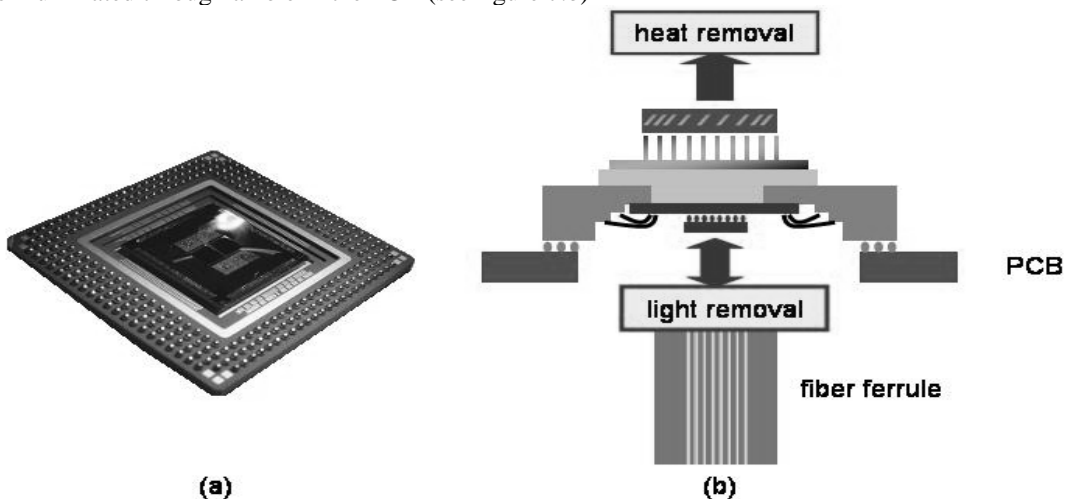

Figure 7 (a) packaged IO-module (b) down facing cavity with heat removal on top and optics through the pcb

The idea here is to extend the package with some kind of mechanical feature, which we will call the spacer plate from now on, in order to align the 2-D fiber ferrule to the 2-D arrays of VCSELs and photodiodes. The spacer plate consists of a thin plate with 4 standard MT pins to accept two 8x8 MT-like connectors (one for the VCSELS and one for the photodiodes) and a central hole for the protruded fiber connector. The critical step is the mounting of the spacer plate on the BGA package. This should be done in such away, that it assures good coupling between the optical components. Therefore, we need to accurately align the spacer plate to the opto-electrical chips. As discussed in the previous chapter, the flip-chip process has an accuracy of better than 1 micron. Note also that we are dealing with MultiMode fibers, requiring alignment accuracies of around 5 to 10 micron. As a result, correct alignment of the spacer plate to the CMOS will assure good coupling between the optical components. Note that the spacer plate serves a vertical and lateral reference for the connector and as such, tight control of the lateral and vertical alignment of the spacer plate to the CMOS chip is needed. We have developed two different approaches to tackle this package-level integration.
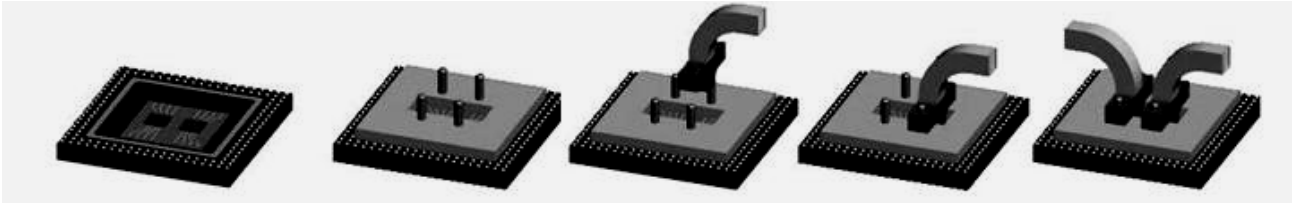

Figure 8: spacer plate mounted on BGA package for correct fiber ferrule alignment

6.1 Alignment pins on CMOS: silicon bench approach
In this approach, the guiding MT pins are mounted in two precise-machined silicon structures (benches), which are in turn flip-chipped onto the CMOS. The high accuracy of the flip-chipping process ensures good overall alignment. Here the spacer has oversized pin holes.

The advantage of this approach is the high alignment accuracy, the self-alignment character of the flip-chip process, and the low cost of the silicon benches. The down-side is that it requires space on the CMOS chip and it actually also requires a large CMOS since the pins are 4.6 mm apart in one direction and 2.5 mm apart in the other direction. Finally, the vertical alignment remains a difficult aspect.
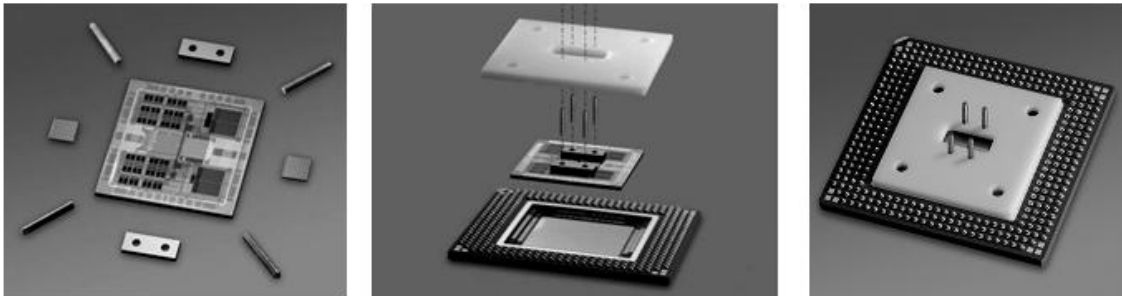

Figure 9: view on the assembly of the silicon benches, MT pins and opto-electronic components (photos by CEA-LETI)

6.2 Alignment pins in spacer plate: 3-D index alignment
In this approach, the spacer plate contains precisely machined pin holes for holding the pins in the correct position. On top of that, high-accuracy alignment features are patterned on the top surface. Such alignment features are also patterned on the CMOS chip. In combination with a 3-D Coordinate Measurement Microscope (3-D CMM), these features enable us to measure the 3-D position (x,y,z) and orientation (α,β,γ) of both the spacer plate and the CMOS chip. Using this information, we can translate and rotate the spacer plate into the desired position and orientation. Once the spacer is correctly aligned to the CMOS, we fix the spacer plate to the packaged module using UV-curable adhesives. The 3-D CMM is the key to extending conventional 2-D index alignment to the full 3-D space. More information on this 3-D index alignment can be found in O.Rits et al.[24].

This approach offers greater flexibility than the silicon bench approach, allows to avoid contact with the delicate and sensitive CMOS chip and to achieve the full 3-D alignment in one alignment step. The downside is the expected lower alignment accuracy, and the requirement of a highly accurate spacer plate.
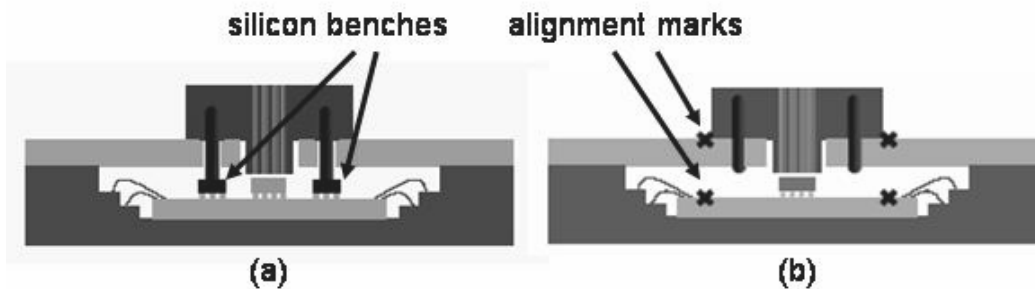

Figure 10 : (a) pins in silicon benches on CMOS (b) pins contained in spacer plate

7. CONCLUSIONS

We have discussed the issues with using electrical inter-chip interconnects for high-speed applications, where the practical implementation inside digital electronic systems is becoming increasingly difficult. We have also discussed optical inter-chip interconnects as alternatives. After a review of the major drawbacks of electrical interconnects and the potential benefits of optical interconnects, we argued integration to be a key-factor for the successful introduction of such optical interconnects at the inter-chip level: CMOS integrated driver and receiver circuitry, electro-optical chips integrated on the CMOS chip and package level integration of parallel optical ferrules. We have described the system demonstrator we designed and fabricated in the framework of the IST project IO and described the three levels of integration in that demonstrator into more detail.

ACKNOWLEDGEMENT

REFERENCES

1. "Signal Integrity", http://www.csee.umbc.edu/vlsi/reports/si_chapter.pdf
2. "Modeling and Simulation of High Speed Digital Circuits and Interconnects", Mohammad S. Sharawi, Computer and Software Engineering, Philadelphia University, Amman 19392 - Jordan.
3. S.K. Tewksbury, "Interconnections within Microelectronic Systems", West-Virginia University, Dept. of Electrical and Computer Engineering, PO Box 6101, Morgantown, WV 26506,
1. http://stewks.ece.stevens-tech.edu/sktpersonal.dir/sktwireless/microinterx.pdf
4. D. Miller and H. Ozaktas, "Limit to the bit-rate capacity of electrical interconnects from the aspect ratio of the system architecture," Journal of Parallel and Distributed Computing, vol. 41, no. 1, pp. 42–52, Feb. 1997.
5. D. Wiens, "Printed Circuit Board Routing at the Threshold, Advanced Technology for the New Millenium", Mentor Graphics, May 2000.
6. D. Maliniak, "Wringing More Performance form FR-4", Electronic Design, September 30, 2002
7. D.A.B. Miller, "Physical Reasons for Optical Interconnection", Int. J. Optoelectronics 11, 155-168 (1997)
8. D.A.B. Miller, "Rationale and challenges for optical interconnects to electronic chips," Proc. of IEEE, vol. 88, no. 6, pp. 728–749, June 2000.
9. N. Savage, "Linking with light," IEEE Spectrum, vol. 39, no. 8, pp.32–36, Aug. 2002.

10. D.A.B. Miller, "Dense Two-Dimensional Integration of Optoelectronics and Electronics for Interconnections", Heterogeneous Integration: Systems on a Chip, Anis Husain, and Mahmoud Fallahi, Eds. SPIE Critical Reviews of Optical Engineering, Vol. CR70 (SPIE, Bellingham, 1998), 80 – 109.
11. Huang, D., Sze, T., Davidson H., and Esener S., "Can Optical Interconnects be Sufficiently Parallel to Support the Needs of In-Box Digital Systems", Trends in Optics and Photonics Series Vol. 90
12. The European Union IST research program IO IST-2000-28358. For more information on the IO project, see http://io.intec.ugent.be
13. More information on parallel optical interconnect modules can be found at http://www.paralleloptics.org
14. Satake, T.; Arikawa, T.; Blubaugh, P.W.; Parsons, C.; Uchida, T.K. : "MT multifiber connectors and new applications", Proceedings 44th Electronic Components and Technology Conference, 1994, pages:994 - 999
15. R. Annen and H. Melchior, "Fully integrated 0.25 µm CMOS VCSEL driver with current peaking," *Electronics Letters*, vol. 38, no. 4, pp. 174-175, Feb. 2002.
16. S. Donnay and G. Gielen, Eds., *Substrate Noise Coupling in Mixed-Signal ASICs*. Dordrecht, The Netherlands:Kluwer Academic Publishers, 2003.
17. M. Badaroglu, G. Van der Plas, P. Wambacq, L. Balasubramanian, K. Tiri, I. Verbauwhede, S. Donnay, G. G. E. Gielen, and H. J. De Man, "Digital circuit capacitance and switching analysis for ground bounce in ICs with a high-ohmic substrate," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 7, pp. 1119–1130, July 2004.
18. D. K. Su, M. J. Loinaz, S. Masui, and B. A. Wooley, "Experimental results and modeling techniques for substrate noise in mixed-signal integrated circuits," *IEEE Journal of Solid-State Circuits*, vol. 28, no. 4, pp. 420–430, Apr. 1993.
19. B. E. Owens, S. Adluri, P. Birrer, R. Shreeve, S. K. Arunachalam, K. Mayaram, and T. S. Fiez, "Simulation and measurement of supply and substrate noise in mixed-signal ICs," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 2, pp. 382–391, Feb. 2005.
20. M. De Wilde, W. Meeus, P. Rombouts and J. Van Campenhout, "A simple on-chip repetitive sampling setup for the quantification of substrate noise," *IEEE Journal of Solid-State Circuits*, accepted for publication.
21. M. Nagata, J. Nagai, T. Morie, and A. Iwata, "Measurements and analyses of substrate noise waveform in mixed-signal IC environment," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 6, pp. 671–678, June 2000.
22. F. Marion, J. Routin, R. Bockstaele, O.Rits, "Packaging for Optical Rx/Tx Arrays: From BGA to CSP packaging", 37th International Symposium on Microelectronics, Long Beach, California, USA, Nov. 2004
23. Tissot J.L., Marion F., "Collective flip-chip technologyfro infrared focal plane arrays", SENSOR MATTER 12 (7), 389 – 396, 2000
24. O. Rits, R. Bockstaele, K. Naessens, R. Baets, "Packaginf Solution with a 3-Dimensional Alignment Approach for 2-Dimensional Parallel Direct Optical Interconnects to the Chip", 15th European Microelectronics and Packaging Conference & Exhibition, Brugge, Belgium, June 12- 15, 2005