



## Nano Scale Disruptive Silicon-Plasmonic Platform for Chip-to-Chip Interconnection

### Chip-to-chip interconnect analysis Requirements and needs

Deliverable no.: D2.1  
Due date: 01/31/2012  
Actual Submission date: 01/31/2012  
Authors: ST  
Work package(s): WP2  
Distribution level: RE<sup>1</sup> (NAVOLCHI Consortium)  
Nature: document, available online in the restricted area of the NAVOLCHI webpage

#### List of Partners concerned

Partner number	Partner name	Partner short name	Country	Date enter project	Date exit project
1	Karlsruher Institut für Technologie	KIT	Germany	M1	M36
2	INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM VZW	IMCV	Belgium	M1	M36
3	TECHNISCHE UNIVERSITEIT EINDHOVEN	TU/e	Netherlands	M1	M36
4	RESEARCH AND EDUCATION LABORATORY IN INFORMATION TECHNOLOGIES	AIT	Greece	M1	M36
5	UNIVERSITAT DE VALENCIA	UVEG	Spain	M1	M36
6	STMICROELECTRONICS SRL	ST	Italy	M1	M36
7	UNIVERSITEIT GENT	UGent	Belgium	M1	M36

<sup>1</sup> **PU** = Public  
**PP** = Restricted to other programme participants (including the Commission Services)  
**RE** = Restricted to a group specified by the consortium (including the Commission Services)  
**CO** = Confidential, only for members of the consortium (including the Commission Services)

***Deliverable Responsible***

Organization: STMicroelectronics – Interconnect Systems Group  
Contact Person: Alberto Scandurra  
Address: STMicroelectronics  
Stradale Primosole, 50  
95121 Catania  
Italy  
Phone: +39 095 – 740 4432  
Fax: +39 095 – 740 4008  
E-mail: alberto.scandurra@st.com

***Executive Summary***

This document describes the main characteristics of both on-chip and off-chip communication systems, highlighting their requirements and needs depending on applications.

***Change Records***

Version	Date	Changes	Author
V1.0	07/12/2011	First draft.	Alberto Scandurra (ST)
V2.0	24/01/2012	Integrated contribution from AIT.	Alberto Scandurra (ST) Christoforos Kachris (AIT) Emmanouil-Panagiotis Fitrakis (AIT)
V3.0	30/01/2012	Final version to be reviewed by the consortium.	Alberto Scandurra (ST) Christoforos Kachris (AIT) Emmanouil-Panagiotis Fitrakis (AIT)
V4.0	31/01/2012	Final version reviewed by AIT.	Alberto Scandurra (ST) Christoforos Kachris (AIT) Emmanouil-Panagiotis Fitrakis (AIT) Ioannis Tomkos (AIT)

## Contents

REFERENCES .....	4
<b>1 INTRODUCTION: SYSTEMS ON CHIP AND SYSTEMS IN PACKAGE .....</b>	<b>7</b>
<b>2 SYSTEMS ON CHIPS INTEGRATION ISSUES.....</b>	<b>10</b>
On-chip interconnect.....	10
Limits of electrical interconnect .....	11
<i>Physical issues</i> .....	13
<i>Performance issues</i> .....	13
From bus to Network on Chip .....	14
New paradigms .....	18
<i>Optical interconnect</i> .....	18
<i>Transmission lines</i> .....	20
<b>3 FROM SYSTEMS ON CHIP TO SYSTEMS IN PACKAGE .....</b>	<b>21</b>
General concepts.....	21
SiP technologies.....	23
<i>Flip-chip</i> .....	23
<i>Through Silicon Via</i> .....	24
Chip-to-chip physical channel requirements .....	26
Limits of chip-to-chip electrical channel .....	26
Novel solutions .....	30
<i>AC-coupling</i> .....	30
<i>Optical interconnect</i> .....	30
Potential benefits of plasmonics .....	32

## References

- [1] A. Scandurra – STBus Communication System: Concepts and Definitions (Public specification, STMicroelectronics web site)
- [2] W.J. Dally, B. Towles – Principles and practices of interconnection networks (Morgan Kaufmann)
- [3] M. Coppola, R. Locatelli, G. Maruccia, L. Pieralisi, A. Scandurra - Spidergon: a novel on-chip communication network (SOC working conference, Tampere, 2004)
- [4] ITRS web site, <http://www.itrs.net>
- [5] System Drivers, International Technology Roadmap for Semiconductors (ITRS), 2011
- [6] Interconnects, International Technology Roadmap for Semiconductors (ITRS), 2011
- [7] William J. Dally and Brian Towles, Route Packets, Not Wires: On-Chip Interconnection Networks, DAC 2001, June 18-22, 2001, Las Vegas, Nevada, USA
- [8] Overcome Copper Limits with Optical Interfaces, Altera White pages, WP01161, April 2011
- [9] A.Carpenter, J.Hu, M.Huang, H.Wu, P.Liu - A Design Space Exploration of Transmission-Line Links for On-Chip Interconnect
- [10] R.Canegallo, L.Perugini, A.Pasini, M.Innocenti, M.Scandiuazzo, R.Guerrieri, P.L.Rolandi – System on Chip with 1.12mW-32Gb/s AC-Coupled 3D Memory Interface
- [11] R.J.Drost, R.D.Hopkins, I.E.Sutherland (Sun Microsystems, Inc.) – Proximity Communication
- [12] A.Fazzi, R.Canegallo, L-Ciccarelli, L.Magagani, F.Natali, E.Jung, P.Rolandi, R.Guerrieri – 3-D Capacitive Interconnections With Mono- and Bi-Directional Capabilities
- [13] E. Ozbay, “Plasmonics: Merging Photonics and Electronics at Nanoscale Dimensions”, Science 311, 189-193 (2006).
- [14] M. Born and E. Wolf, Principles of Optics (Pergamon Press, New York, 1959).
- [15] S. A. Maier, Plasmonics: Fundamentals and Applications (Springer, 2007).
- [16] E. N. Economou, “Surface plasmons in thin films”, Phys. Rev. 182, 539-554 (1969).
- [17] M. P. Nezhad, A. Simic, O. Bondaenko, B. Slutsky, A. Mizrahi, L. Feng, V. Lomakin, and Y. Fainman, “Room temperature subwavelength metallo-dielectric lasers”, Nature Photonics 4, 395-399 (2010).
- [18] R. Perahia, T. P. Mayer Alegre, A. H. Safavi-Naeini, and O. Painter, “Surface-plasmon mode hybridization in subwavelength microdisk lasers”, Appl. Phys. Lett. 95, 201114 (2009).
- [19] R.F. Oulton, V. J. Sorger, T. Zentgraf, R-M. Ma, C. Gladden, L. Dai, G. Bartal, and X. Zhang, “Plasmon lasers at deep subwavelength scale”, Nature 461, 629–632 (2009).
- [20] M. A. Noginov, G. Zhu, A. M. Belgrave, R. Bakker, V. M. Shalaev, E. E. Narimanov, S. Stout, E. Herz, T. Suteewong, and U. Wiesner, “Demonstration of a spaser-based nanolaser”, Nature 460, 1110–1112 (2009).
- [21] M. T. Hill, Y-S. Oei, B Smalbrugge, Y. Zhu, T. de Vries, P. J. van Veldhoven, F. W. M. van Otten, T. J. Eijkemans, J.P. Turkiewicz, H. de Waardt, E. J. Geluk, S-H. Kwon, Y-H. Lee, R. Nötzel, and M. K. Smit., “Lasing in Metallic-Coated Nanocavities”, Nature Photonics 1, 589-594 (2007).
- [22] M. T. Hill, M. Marell, E. S. P. Leong, B. Smalbrugge, Y. Zhu, M. Sun, P. J. van Veldhoven, E. J. Geluk, F. Karouta, YS. Oei, R. Nötzel, C-Z Ning, and M. K. Smit, “Lasing in metal-insulator-metal sub-wavelength plasmonic waveguides”, Opt. Express 17, 11110 (2009).
- [23] V. J. Sorger and Xiang Zhang, “Spotlight on plasmon lasers”, Science Vol. 333, no. 6043, 709-710 (2011).

- [24] T. Nikolajsen, K. Leosson, and S. I. Bozhevolnyi, "Surface plasmon polariton based modulators and switches operating at telecom wavelengths", *Appl. Phys. Lett.* 85, 5833–5835 (2004).
- [25] J. A. Dionne, K. Diest, L. A. Sweatlock, and H. A. Atwater, "PlasMOStor: A Metal-Oxide-Si Field Effect Plasmonic Modulator", *Nano Lett.* 9, 897–902 (2009).
- [26] A. V. Krasavin, T. P. Vo, W. Dickson, P. M. Bolger, and A. V. Zayats, "All-Plasmonic Modulation via Stimulated Emission of Copropagating Surface Plasmon Polaritons on a Substrate with Gain", *Nano Lett.* 11, 2231-2235 (2011).
- [27] M. C. Gather, K. Meerholz, N. Danz, K. Leosson, "Net optical gain in a plasmonic waveguide embedded in a fluorescent polymer", *Nature Photonics* 4, 457-461 (2010).
- [28] J. Grandidier, G. C. des Francs, S. Massenot, A. Bouhelier, L. Markey, J. C. Weeber, C. Finot, A. Dereux, "Gain-assisted propagation in a plasmonic waveguide at telecom wavelength", *Nano Lett.* 9, 2935-2939 (2009).
- [29] Y. Li, H. Zhang, N. Zhu, T. Mei, D. H. Zhang, J. Teng, "Short-range surface plasmon propagation supported by stimulated amplification using electrical injection", *Opt. Express* 19 (22), 22107-22112 (2011).
- [30] J. Zhang, S. Zhu, H. Zhang, S. Chen, G.-Q. Lo, D.-L. Kwong, "An ultracompact surface plasmon polariton-effect-based polarization rotator", *IEEE Photon. Techn. Lett.* 23 (21), 5986689, 1606-1608 (2011).
- [31] Q. Li, S. Wang, Y. Chen, M. Yan, L. Tong, M. Qiu, "Experimental demonstration of plasmon propagation, coupling, and splitting in silver", *IEEE J. on Selected Topics in Quant. Electron.* 17 (4), 5594609, 1107-1111 (2011).
- [32] T. Holmgaard, Z. Chen, S. I. Bozhevolnyi, L. Markey, A. Dereux, "Design and characterization of dielectric-loaded plasmonic directional couplers", *J. Lightwave Technol.* 27 (24), 5521-5528 (2009).
- [33] S. Sederberg, V. Van, A. Y. Elezzabi, "Monolithic integration of plasmonic waveguides into a complimentary metal-oxide-semiconductor- and photonic-compatible platform", *Appl. Phys. Lett.* 96 (12), 121101 (2010).
- [34] Jin Tae Kim, Suntak Park, Jung Jin Ju, Sangjun Lee, and Sangin Kim, "Low bending loss characteristics of hybrid plasmonic waveguide for flexible optical interconnect", *Opt. Express* 18, 24213-24220 (2010).
- [35] Yu Nanfang, R. Blanchard, J. Fan, Jie Wang Qi, C. Pflugl, L. Diehl, T. Edamura, S. Furuta, M. Yamanishi, H. Kan, F. Capasso, "Plasmonics for laser beam shaping", *IEEE Trans. on Nanotechn.* 9 (1), 5200530, 11-29 (2010).
- [36] Takuma Aihara, Kyohei Nakagawa, Masashi Fukuhara, Yen Ling Yu, Kenzo Yamaguchi, and Mitsuo Fukuda, "Optical frequency signal detection through surface plasmon polaritons", *Appl. Phys. Lett.* 99, 043111 (2011).
- [37] M. Sandtke and L. Kuipers, "Slow guided surface plasmons at telecom frequencies," *Nat. Phot.* 1, 573–576 (2007).
- [38] S. I. Bozhevolnyi, V. S. Volkov, E. Devaux, J.-Y. Laluet, T. W. Ebbesen, "Channel plasmon subwavelength waveguide components including interferometers and ring resonators", *Nature* 440 (7083), 508-511(2006).
- [39] Jin Tae Kim, Jung Jin Ju, Suntak Park, Min-su Kim, Seung Koo Park, and Myung-Hyun Lee, "Chip-to-chip optical interconnect using gold long-range surface plasmon polariton waveguides," *Opt. Express* 16, 13133-13138 (2008).
- [40] <http://www.ict-platon.eu>
- [41] S. Papaioannou, K. Vyrsokinos, O. Tsilipakos, A. Ptilakis, K. Hassan, J. C. Weeber, L. Markey, A. Dereux, S. I. Bozhevolnyi, A. Miliou, E. E. Kriezis, and N. Pleros, "A 320 Gb/s-

Throughput Capable 2 x2 Silicon-Plasmonic Router Architecture for Optical Interconnects”, J. Lightwave Technol. 29, 3185-3195 (2011).

## 1 Introduction: Systems on Chip and Systems in Package

Systems on Chip (SoCs) are complex systems containing billions of transistors integrated in a unique silicon chip, implementing highly complex functionalities by means of a variety of modules communicating with the system memories and/or between themselves via a proper communication system.

The building-blocks of a SoC (also called Intellectual Properties or IPs) can be distinguished into two classes; **initiators**, which are all the blocks able to generate traffic, i.e. write data into a storage element (SE), typically a memory, and read data from a SE; and **targets**, which are blocks able to manage the traffic generated by the initiators (the SEs themselves).

Typical examples of initiators are:

- **Processors**, which have strict requirements in terms of latency and bandwidth. Their bandwidth must be limited in some way to allow the other initiators to be serviced.
- **Real time initiators**, such as audio/video blocks, are more latency tolerant than the processors, but have strict needs in terms of bandwidth.
- **DMAs**, which have no particular requirements in terms of latency or bandwidth. Normally they can work using the remaining bandwidth, i.e. the part unused by the processors and real time initiators.

Among the targets the following classes can be identified:

- **External fast memories** comprise high performance memories such as SDRAM (Synchronous Dynamic Random Access Memory) and DDR (Dual Data Rate) SDRAM, mainly for real-time applications like video, typically working at around 400 MHz. Their speed is limited by physical constraints imposed by pads.
- **Slow memories** are usually low performance memories like SRAM and Flash, used for storage of huge amounts of data, and whose access is managed by caches, typically working at around 200 MHz. Their speed is limited by the application.
- **Peripherals** are slow memories such as I2C and Smartcard, used where no high-speed performance is required, working at around 50/100 MHz.

The different IPs communicate with each other via the onchip communication system, that can be implemented either as a bus [1], or as a Network on Chip (NoC) [2] according to the novel communication paradigm developed over the last few years [3].

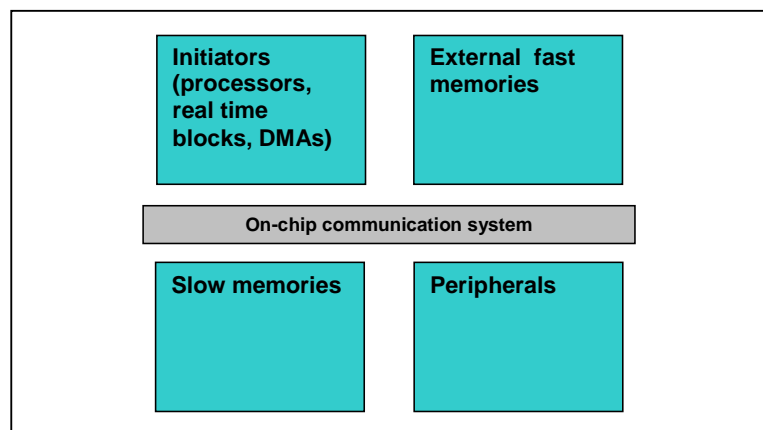


Figure 1-1: Typical SoC architecture

Normally the CPUs run at the highest speed and the memory system represents the SoC bottleneck in terms of performance. This approach, using different islands running at different frequencies in a GALS (Globally Asynchronous Locally Synchronous) system, is widely used today. It is expected that this will eventually become so widespread that no global clock distribution will be required.

While such systems work today, it is expected that the number of IPs in SoCs for consumer stationary applications will rise to above 100, and the aggregated data rate to above 100Tb/s, by 2016 [ITRS 2007 System Drivers, "SOC Consumer Stationary Design Complexity Trends"]. In this context, it is clear that huge strain will be put on the on-chip communication system.

Current systems are affected by at least two meaningful issues from a technological point of view, as highlighted below.

- The ever decreasing feature size in CMOS silicon processes allows digital logic to shrink significantly between subsequent fabrication nodes, for example a shrink of 55% could be expected when comparing a digital IP implemented in 90nm and 65nm. However analogue and IO cells have been unable to match this rate of decrease, leading to increasingly pad limited designs in many complex SoCs. A pad-limited design can be viewed as wasteful since the digital logic is not implemented as densely as it might be were it the only contributing factor to the device area.
- The transition to sub 32nm design introduces a dichotomy between supporting low voltage, high speed IO logic; for example DDR3 1.5V @ 800MHz+; and higher voltage interconnect technologies, for example HDMI, SATA, USB3 etc. The lower voltage DDR3 interface requires a gate oxide thickness of 30Å, while HDMI would require 50Å - incompatible within a standard process.

By splitting a traditional single SoC in to multiple dice these pressures can be alleviated. A system composed of more than one die is usually referred to as System in Package (SiP). An example SiP would be composed of a 32nm die comprising high speed CPUs, DDR3 controller(s) and differentiating IP, connected to a 55nm die comprising analogue PHYs. Thanks to the reduced set of analogue IP the 32nm die gets the maximum benefit from the shrink. SiP technology offers many significant benefits, including:

- **Footprint** – More functionality fits into a small space. This extends [Moore's Law](#) and enables a new generation of tiny but powerful devices.
- **Speed** – The average wire length becomes much shorter. Because [propagation delay](#) is proportional to the square of the wire length, overall performance increases.
- **Power** – Keeping a signal on-chip reduces its [power consumption](#) by ten to a hundred times. Shorter wires also reduce power consumption by producing less [parasitic capacitance](#). Reducing the power budget leads to less heat generation, extended battery life, and lower cost of operation.
- **Design** – The vertical dimension adds a higher order of connectivity and opens a world of new design possibilities.
- **Heterogeneous integration** – Circuit layers can be built with different processes, or even on different types of wafers. This means that components can be optimized to a much greater degree than if they were built together on a single wafer. Even more interesting, components with completely incompatible manufacturing could be combined in a single device (see figure 1.2).






Figure 1-2: Example of heterogeneous integration

- **Circuit security** - The stacked structure hinders attempts to [reverse engineer](#) the circuitry. Sensitive circuits may also be divided among the layers in such a way as to obscure the function of each layer.
- **Bandwidth** - 3D integration allows large numbers of vertical vias between the layers. This allows construction of wide bandwidth buses between functional blocks in different layers. A typical example would be a processor plus memory 3D stack, with the cache memory stacked on top of the processor. This arrangement allows a bus much wider than the typical 128 or 256 bits between the cache and processor. Wide buses in turn alleviate the memory wall problem.




Figure 1-3: Detail of electrical wires between dice

## 2 Systems on Chips integration issues

### *On-chip interconnect*

The function of an **on-chip interconnect** is to distribute the signals to and among the various circuit/system functions on a chip.

The fundamental development requirement for interconnect is to meet the high-speed transmission needs of chips, despite further scaling of feature sizes.

When breaking down any electronic system (e.g. SoC) into its basic components (transistors, diodes, passive circuit elements, etc.) we observe that electronic systems consist of two parts: the basic components and the highly complex interconnect fabric linking them. This interconnect fabric is organized in a **hierarchical way**, from narrow short interconnects between basic elements to longer and larger interconnects for interconnecting circuit blocks. For integrated circuits with well-defined local, intermediate and global interconnect layers, on chip circuit-hierarchy is organized from transistors to logic gates, sub-circuits, circuit-blocks, and finally, bond pad interface circuits.

Table 2-1 presents a structured definition of interconnect technologies based on the interconnect hierarchy. The interconnection hierarchy is depicted in figure 2-1.

<i>Level</i>	<i>Suggested Name</i>	<i>Key Characteristics</i>
Package	3D-Packaging	Traditional packaging of interconnect technologies, e.g., wire-bonded die stacks, package-on-package stacks.
Bond-pad	3D-Wafer-level Package (3D-WLP)	3D interconnects are processed after the IC fabrication, “post IC-passivation” (via last process). Connections on bond-pad level.
Global	3D-Stacked Integrated Circuit/ 3D-System-on-Chip (3D-SIC /3D-SOC)	Stacking of large circuit blocks (tiles, IP-blocks, memory –banks), similar to an SOC approach but having circuits physically on different layers. Unbuffered I/O drivers (Low C, little or no ESD protection on TSVs).
Intermediate	3D-SIC	Stacking of smaller circuit blocks, parts of IP-blocks stacked in vertical dimensions. Mainly wafer-to-wafer stacking.
Local	3D-Integrated Circuit (3D-IC)	Stacking of transistor layers. Requires 3D connections at the density level of local interconnects.

*Table 2.1 : Interconnect technologies based on the the Interconnect Hierarchy, (Source: ITRS Interconnects[6])*

A wide variety of technologies can be used to realize the interconnect technologies described above. Of particular interest here are the so-called “Through-Si-Via” technologies used for 3D-WLP, 3D-SOC, and 3D-SIC interconnect technologies.

A Through Silicon Via (TSV) connection is a galvanic connection between the two sides of a Si wafer that is electrically isolated from the substrate and from other TSV connections. The isolation layer surrounding the TSV conductor is called the TSV liner. The function of this layer is to electrically isolate the TSVs from the substrate and from each other. This layer also determines the TSV parasitic capacitance. In order to avoid diffusion of metal from the TSV into the Si-substrate, a barrier layer is used between the liner and the TSV metal.

Although that these interconnect provide high throughput in the current process technologies, there are several challenges that need to be addressed as discussed in next section, if we want to continue using the silicon-based interconnect in the future process technologies (22nm and below).




Figure 2-1 : Schematic representation of the Interconnects Hierarchy in CMOS wafer

### *Limits of electrical interconnect*

From a technological point of view, interconnects can be classified in the following categories:

- **local interconnect**, used for short-distance communication (typically between individual logic gates), and comprising the majority of on-chip wires; they have the smallest pitch, and a delay of less than one clock cycle;
- **global interconnect**, providing communication between large functional blocks (IPs). Global interconnects have the largest pitch and a delay typically longer than one or two clock cycles.
- **intermediate interconnect**, having dimensions that are between those of local and global interconnects.




Figure 2-2 : On-chip interconnect classification

A key difference between local and global interconnect is that the length of the former scales with technology node, while for the latter the length is approximately constant. From a functional point of view, the two main important and performance-demanding applications of interconnects in SoC are signaling (i.e. the communication of different logic units) and clock distribution. In this context they can be classified as

- **point-to-point links**, used for critical data-intensive links, such as CPU-memory buses in processor architectures;
- **broadcast links**, representing physical channels where the number of receivers (and therefore repeaters) is high and switching activity is also high
- **network links**, targeted at system buses and reconfigurable networks, aiming at serving complete system architectures, whose typical communication load is around several tens of Gb/s.

An *ideal* interconnect should be able to transmit any signal with no delay, no degradation (either inherent or induced by external causes), over any distance without consuming any power, requiring zero physical footprint and without disturbing the surrounding environment. Of course this is not the case, and a number of metrics are used in order to characterize the performance and the quality of real interconnects.

The **propagation delay** (ps) is the time required by a signal to cross a wire. Pure interconnect delay depends on the link length and the speed of propagation of the wavefront (time of flight). Electrical regeneration introduces additional delay through buffers and transistor switching times. Moreover, delay can be induced by crosstalk; this can be reduced by increasing the interconnect width at the expense of bandwidth density. Technology scaling has negligible effect on the delay of interconnect with an optimal number of repeaters.

**Bandwidth density** (Hz/ $\mu\text{m}$ ) is a metric that characterizes information throughput over a unit cross section of an interconnect.

The **power and delay product** (PDP, pJ) is commonly used in the technology design process to evaluate circuit performance.

The **Bit Error Rate (BER,  $\text{s}^{-1}$ )** may be defined as the rate of error occurrences and is the main criterion in evaluating the performance of digital transmission systems. In SoCs, errors come from signal degradation. For an on-chip communication system a BER of  $10^{-15}$  is acceptable; electrical interconnects typically achieve BER figures better than  $10^{-45}$ . This is why BER is not commonly considered in integrated circuits design circles. However, future operation frequencies are likely to change this, since the combination of necessarily faster rise and fall


times, lower supply voltages and higher crosstalk increases the probability of wrongly interpreting the signal that was sent.

Using advanced CMOS technologies, where DSM effects are dominant, the physical design of a SoC is increasingly faced with two types of issue:

- **Physical issues**, related to the difficulties encountered, mainly during the placement of the hard macros and the standard cells, and during the routing of clock nets and communication system wires;
- **Performance issues**, related mainly to the bandwidth requirements of the different IPs, that in order to be fulfilled, would require SoCs to run at very high speeds.

### Physical issues

Figure 2.3 shows the floorplan of an example CMOS chip for consumer applications, where the interconnect, implemented with a NoC solution, is highlighted.



*Figure 2-3: CMOS chip floorplan with NoC placement and routing highlighted*

In this figure the dark squares represent Network Interfaces of the various IPs of the chip (both initiators and targets), the dark rectangles are the nodes, responsible for arbitration and propagation of information, and the dark lines are the physical channels connecting the different NoC devices. All these elements must be physically located in the grey area, which represents the physical space available for interconnect. Because of the shape (quite irregular and with thin regions) and the area size, it is clear that the placement of the interconnect standard cells can be difficult, and that the routing of the wires, that can be also very long, can suffer congestion.

### Performance issues

As far as performance is concerned, two main factors influence the overall operating frequency of a SoC: device switching times and interconnect bandwidth. Current technologies can achieve

unprecedented transistor transition frequencies (GHz) due to short transistor lengths. However, the same is not true for interconnect. Indeed, continually shrinking feature sizes, higher clock frequencies, and growth in complexity are all negative factors as far as switching charges on metallic interconnect are concerned. This situation is shifting the IC design bottleneck from computing capacity to communication.

Feature sizes on integrated circuits and therefore circuit speed have followed Moore's law for over three decades and the CMOS integration capability is still increasing. In this respect, according to the International Technology Roadmap for Semiconductors (ITRS) [4], the RC time constants associated with metallic interconnects will not be able to decrease sufficiently for the high-bandwidth applications destined to appear in the next few years. Internal data rates of processors fabricated in deep submicron CMOS technology have exceeded GHz rates. While processing proceeds at GHz internally, off chip wires have held inter-chip clock rates at hundreds of MHz.

The function of an interconnect is to distribute clock and other signals to and among the various circuits/systems on a chip. The fundamental development requirement for interconnect is to meet the high-speed transmission needs of chips despite further scaling of feature sizes. This scaling down however, has been shown to increase the signal runtime delays in the global interconnect layers severely. Indeed, while the reduction in transistors gate lengths increases the circuit speed, the signal delay time for global wires continues to increase with technology scaling, primarily due to the increasing resistance of the wires and their increasing lengths.

Current trends to decrease the runtime delays, the power consumption and the crosstalk, focus on lowering the RCproduct of the wires, by using metals with lower resistivity (i.e. Copper instead of Aluminum) and by the use of insulators with lower dielectric constant.

### ***From bus to Network on Chip***

The System-on-Chip (SoC) industry has developed rapidly over the last fifteen years from producing VLSI devices that integrated a processor and a few memory and peripheral components onto a single chip to today's high-performance SoCs that incorporate hundreds of IP blocks. This progress is a consequence of Moore's Law (which enables ever-higher levels of integration) and of market economics (where consumers demand ever-more functionality in smaller, lower-cost products with better battery life).

Early SoCs used an interconnect paradigm inspired by the rackbased microprocessor systems of earlier days. In those rack systems, a backplane of parallel connections formed a 'bus' into which all manner of cards could be plugged. A system designer could select cards from a catalogue and simply plug them into the rack to yield a customized system with the processor, memory and interfaces required for any given application. In a similar way, a designer of an early SoC could select IP blocks, place them onto the silicon, and connect them together with a standard **on-chip bus** (see figure 2-4).

However, buses do not scale well. With the rapid rise in the number of blocks to be connected and the increase in performance demands, today's SoCs cannot be built around a single bus. Instead, complex hierarchies of buses are used (as illustrated in figure 2-5), with sophisticated protocols and multiple bridges between them. In this case different busses are used based on the system requirements. For example the CPU bus has high performance and low latency, while other busses can be used for low performance communication such as low speed peripherals. Communication between two remote blocks can go via several buses, and every section of every path must be carefully verified. Timing closure is a growing problem because there is so much that must be checked. Bus-based interconnect is being stretched to its limit, and as the limit is

approached the risk of errors increases rapidly. A new interconnect strategy was required to bring these risks back under control.




Figure 2-4: Typical Bus used for interconnection of processors, memory and I/O devices

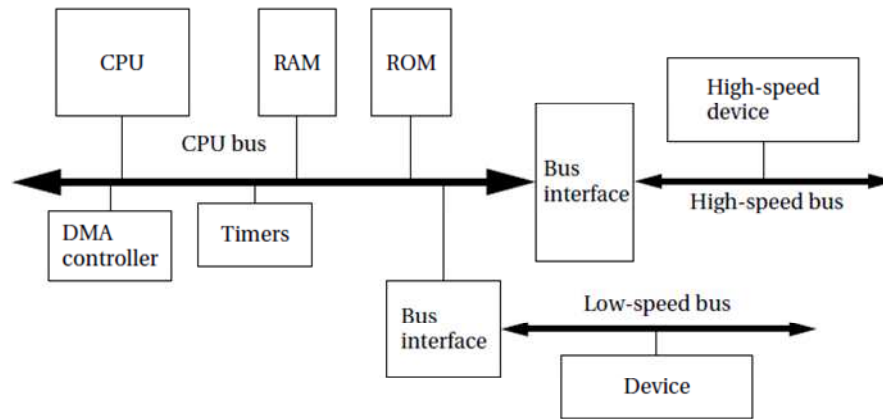


Figure 2-5: Hierarchy of busses in a complex SoC

To overcome these problems of scalability and complexity, **Networks-On-Chip** (NoCs) have been proposed as a promising replacement to eliminate many of the overheads of busses. Instead of connecting these top-level modules (processors, memory, etc) by routing dedicated wires using the busses, they are connected to a network that routes packets between them. This Network-on-Chip framework employs self-timed logic techniques to deliver a robust, correct-by-construction interconnection fabric that allows each client block to operate in its own fully-decoupled timing domain, thereby addressing system-level timing-closure issues.

Such **Networks-on-Chip** have routers at every node, connected to neighbors via short local on-chip wiring, while multiplexing multiple communication flows over these interconnects to provide scalability and high bandwidth. This evolution of interconnection networks as core count increases is clearly illustrated in the choice of a flat crossbar interconnect (usually in a mesh topology) connecting all processors in highly parallel Multi-processor SOCs (MPSoC) as it is shown in figure 2-6. Figure 2-7 depicts the five packet-switched meshes in the 64-core Tiler TILE64 MPSoC. As it is shown in every node of the NoC there connected a processor and a cache using the Tiler NoC Switch.

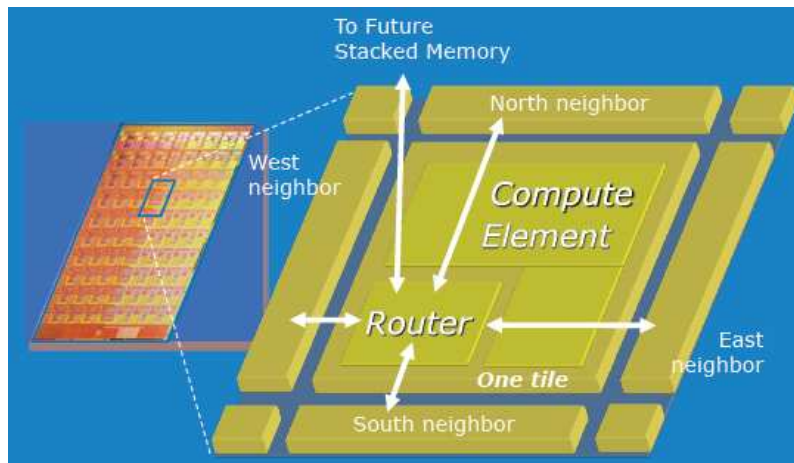


Figure 2-6: Tile design and Mesh topology of NoC

MPSoC design may leverage a wide variety of heterogeneous IP blocks; as a result of the heterogeneity, regular topologies such as a mesh shown above may not be appropriate. With these heterogeneous cores, a customized topology will often be more power efficient and deliver better performance than a standard topology. Often, communication requirements of MPSoCs are known a priori. Based on these structured communication patterns, an application characterization graph can be constructed to capture the point-to-point communication requirements of the IP blocks. To begin constructing the required topology, the number of components, their size and their required connectivity as dictated by the communication patterns must be determined.

An example of a customized topology for a video object plane decoder is shown in Figure 2-9. The MPSoC is composed of 12 heterogeneous IP blocks. In figure 2-8, the design is mapped to a  $3 \times 4$  mesh topology requiring 12 routers (R). When specific application characteristics are taken into account (e.g. not every block needs to communicate directly with every other block), a custom topology is created. This irregular topology reduces the number of switches from 12 to 5; by reducing the number of switches and the links in the topology, significant power and area savings are achieved. Finally, the degree of the switches has changed; the mesh in figure 2-8 requires a switch with 5 input/output ports (although ports can be trimmed on edge nodes). The 5 input/output ports represent the four cardinal directions: north, south, east and west plus an Injection/Ejection port. All of these ports require both input and output connections leading to  $5 \times 5$  crossbars. With a customized topology, not all blocks need both input and output ports; the largest switch in figure 2-9 is a  $3 \times 3$  switch. However in this case the throughput requirements of the links in the custom topology are much higher.






Figure 2-7: The NoC of the Tiler SoC

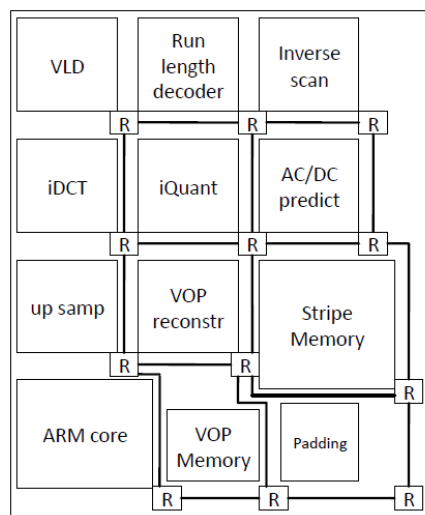


Figure 2-8: Mesh topology of a NoC

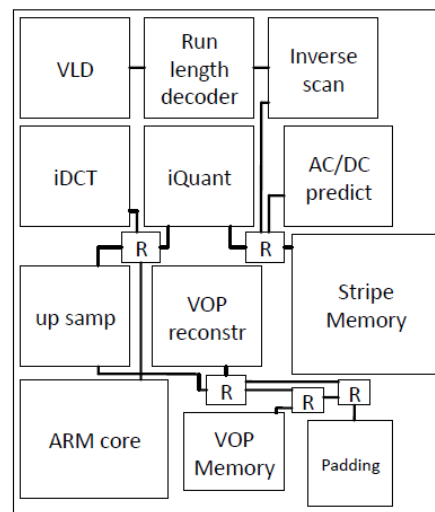


Figure 2-9: Custom topology of a NoC

## *New paradigms*

### **Optical interconnect**

CMOS-compatible optical solutions have been proposed for on-die interconnects (signaling and clock distribution) and I/O. The drivers for on-chip optical interconnects are the utilization of the speed-of-light signal propagation and the large bandwidth of waveguides. For I/O applications, optical solutions focus on increasing the aggregate bandwidth and/or communication distance, while decreasing the power per bit by overcoming the limitations imposed by losses in present package interconnects (metal and dielectric), and by avoiding or minimizing the need for high power equalization and pre-emphasis. Since I/O, signaling and clock distribution require similar optical components, research and production costs are shared.

Because of pitch constraints, as well as delay and power considerations, optical interconnects are not expected to fully replace the lower metal-dielectric interconnect layers in microprocessors. Instead, the focus is on cost-efficient implementations which take advantage of the unique properties of optical architectures to increase overall system performance. For such optical solutions to be viable, the development of CMOS-compatible optical components is of paramount importance. Although significant progress has been made, this area is not yet sufficiently mature to define an intersection with the existing interconnect roadmap.

### **Optical interconnect advantages**

The basic advantages of optical interconnects are speed-of-light signal propagation and large bandwidth, as noted above. However, other potential advantages also exist. Among these are minimum crosstalk between signal transmission paths and multi-wavelength capability. The capability for a single optical path to accommodate multiple wavelengths increases the data-carrying capacity manifold, providing bandwidth densities not achievable by electrical means.

- *Delay*—For the case of on-die signaling, it is possible to define a critical length above which optical interconnects are faster than their metal-dielectric counterpart. The critical length, which depends on the quality of the optical components, has been assessed to be on the order of mm.
- *Signal integrity*—Optical interconnects have the potential for simplifying design and layout constraints arising from undesirable crosstalk in metal-dielectric interconnects.
- *Skew and Jitter*—It has been proposed that the low latency and the absence of crosstalk in optical interconnects can potentially result in low skew and jitter clock distribution. However, advanced clock distribution designs implemented in conventional metal-dielectric systems are expected to meet microprocessor needs.

### **Integration options**

Although a large number of optical architectures have been proposed, most of them fall into one of the following two categories, as follows:

- *Integrated light source architectures*—In this case there are multiple on-die directly modulated light sources (e.g., VCSELS) and on-die detectors. The main disadvantage is the large on-die power consumption/heat dissipation of the sources, and the significant challenges with integrating fast efficient CMOS-compatible light sources.
- *External light source architectures*—These are implementations that utilize one or a few off-die light sources on the package or the board, and on-die modulators and detectors. The main advantage of this family of architectures is that the laser power is off-die (i.e., does not have to be delivered through the die).

The main disadvantage is the coupling losses to bring the light into the chip.

In both cases above, wavelength-specific filters/modulators can be used to implement multiplexing, which enables multiple independent signals transmitted in each channel.

### Optical Interconnects in FPGAs

In contrast with electrical interfaces, optical fiber has virtually no loss. A multiple mode fiber (MMF) has a loss of  $\sim 3$  dB/km and  $\sim 1$  dB/km at 850-nm and 1300-nm wavelengths, respectively. A single mode fiber (SMF) has a loss of  $\sim 0.4$  dB/km and 0.25 dB/km at 1300-nm and 1550-nm wavelengths respectively. MMF is less expensive due to its larger core ( $\sim 50$  micron) and has a bandwidth  $\sim 2$  GHz km; while SMF is more expensive due to its smaller core ( $\sim$  a few microns) and has a bandwidth close 100 THz in practice. The laser that drives the optical signal over an MMF is commonly a light emitting diode (LED) or Vertical Cavity Surface Emitting Laser (VCSEL). The MMF is commonly used for reach distances of  $< 1$  km, while SMF is used for reach distances of  $> 1$  km to a few thousand km. At 10 Gbps, the reach distance for a MMF is  $\sim 300$  m. Unlike the copper electrical link, power consumption and penalty of an optical link is relatively independent of reach length. Moreover, unlike an electrical signal, an optical signal is immune to electric-magnetic interference (EMI) and has no amplitude crosstalk, providing better signal integrity resilience. With the wavelength division multiplexing (WDM), multiple channels can be supported with the same optical fiber, enabling channel material savings.

Altera is one of the FPGA vendors that plans to use optical interconnects for the chip-to-chip communication. Figure 2-10 shows an example of an FPGA with optical interfaces. The FPGA in this example is integrated with optics, such as a transmitter optical sub-assembly (TOSA) and receiver optical sub-assembly (ROSA), providing direct optical signal transmitting and receiving without the need for a discrete optical module.

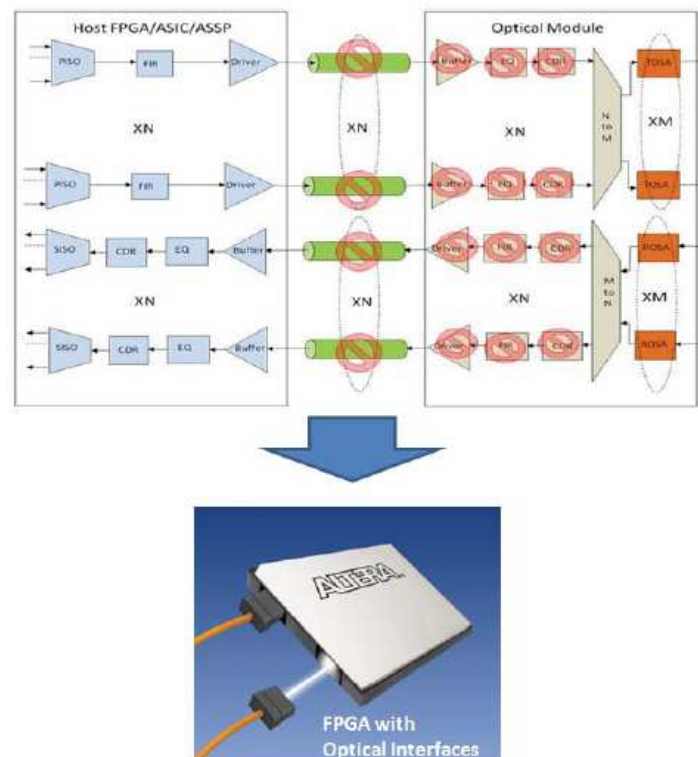


Figure 2-10: Optical chip-to-chip interconnects in FPGAs, source:[8]

## Transmission lines

Transmission lines are common components in RF and microwave circuits. Their characteristics such as impedance, loss, propagation delay, dispersion and crosstalk depend on the structure, size, materials, and fabrication.

A transmission line allows high signaling rate, speed-of-light propagation velocity and can potentially provide sufficient throughput for a range of chip multiprocessors. For all these reasons transmission lines can be considered as good candidates as on-chip interconnect physical medium.

With ever improving transistor performance, a communication system can achieve a data rate of tens of Gb/s per line and an aggregate data rate of Tb/s over on-chip global transmission lines. In medium-sized CMPs, the global network connecting different cores can be entirely based on a multi-drop transmission line system (illustrated in figure 2-11) allowing packet-switching-free communication that is both energy-efficient and low-latency [9].

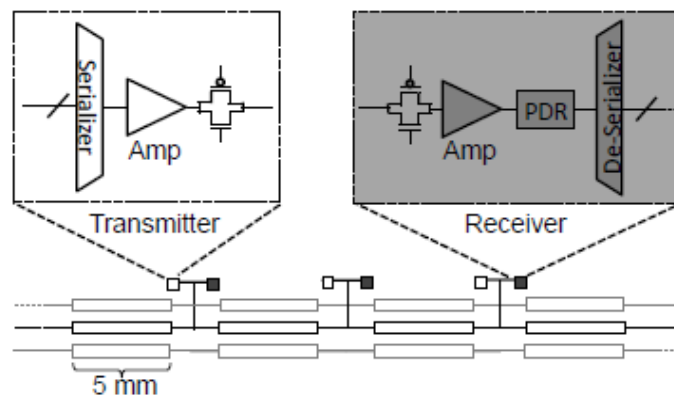


Figure 2-11: Transmission line based interconnect link schematic

In general, the transmission circuit can be as simple as inverter-chain based fully digital circuits and as it becomes more sophisticated, it allows faster data rates at generally reduced per bit energy costs.

Faster transistor speeds in modern and future generation CMOS technologies are an important contributor to the performance of a transmission line link. On-chip transmission lines will operate at many times the core frequency, making serialization and deserialization necessary. Typically, multiple stages of 2:1 MUX/DEMUX are used as serializer/deserializer.

Phase and data recovery (PDR) is another necessary component to ensure the transmitters and receivers can communicate properly, and is independent of transceiver design: After a distance-dependent propagation delay, the transmitted pulses do not align with the receiver's clock. The magnitude of phase delta depends on the sender and can be quickly determined by sending and receiving a short test sequence in an initial calibration step. Data recovery circuits use the clock with the modified phase to ensure correct latching.

### 3 From Systems on Chip to Systems in Package

#### *General concepts*

3D integration technology allows the stacking of different chips and devices in a single package. The maximum benefit is obtained from the use of heterogeneous and highly specialized technologies, and the possibility to make the optimal partitioning early in the design process. The strong increase of the quantitative (i.e. more connectivity for the mobile applications: Bluetooth, WiFi, GPS, FM, UWB, GSM, 3G, 3G+, 4G, WiMax, LTE, DVB-H, WSN, etc., more connectivity for the fix multimedia applications: Internet connection, satellite, DVB-T, home networking, intra-device connections, etc.) and qualitative (higher definition, flexible video format, 3D support, free view point, more and better interactivity, etc.) number of features, required for the current and future multimedia and mobile applications, are exponentially increasing the design complexity. The market success of the related products will be guaranteed if the following challenges are correctly addressed:

- The required performances must be provided in a reasonable power.
- The cost per function must be decreased.
- The development must be done in the right time at reasonable costs (design + manufacturing).

Analyzing the requirements and the potential system architecture solutions of  $M^3$  applications, we can see that:

- The computing performances will be in the order of Tera-operations/s and certainly addressed by Multi-core architectures (from tens to several tens of cores) thanks to the capabilities provided by the nanometer technologies 32 nm and beyond. However, issues remain to be solved on the memory hierarchy and its throughput. The best system partitioning (at chip level) trade-off: performance (computing power, memory throughput, power consumption, size and form factor) vs. cost (on-chip, off-chip) should be done.
- The crucial elements, other than the computing power, contributing to the success of these future applications, will be the ability of the device to manage heterogeneous technologies: digital, analog, RF, discrete devices (resistors, inductances, and capacitors) and software. The analog and RF components don't have the same level of down scaling as the digital logic. Even worse, the size of the analog and RF components could increase with the scaling down of the technology. Again, the best system partitioning (at chip level) trade-off: technology vs. cost (on-chip, off-chip) should be done.
- The I/O and analogue circuitries (about 40 percent of the original device) only scales at about half the rate of the digital logic. We lose an order of magnitude in dice/wafer over several process generations due to the fact that the dice does not scale as quickly as digital transistors. By partitioning the different layers of the silicon it's possible to overrule this limitation allowing designers to get the best of each technology.
- The physical characteristics related to the packaging should be also considered such as: electrical (voltage, signal integrity, throughput, etc.), power consumption, thermal, size and form factor.

The 3D integration (3DI) enables the integration of different types of chips and devices in a single package (figure 3.1) or a compact subsystem providing a maximum benefit from highly specialized and heterogeneous technologies. However, in order to take full advantage of the 3DI, the decision must come upfront in the architecture planning process rather than as a packaging

decision after circuit design is complete. This requires taking 3D design space into account right from the start of the system design in order to distribute its different parts into a new set of chips that will be stacked. So far, huge investments have been made on the fabrication side to provide these technologies but very few results exist on the design methods and tools side. Capability to design, build, validate 3D products is likely to become the next enabler to further silicon integration.

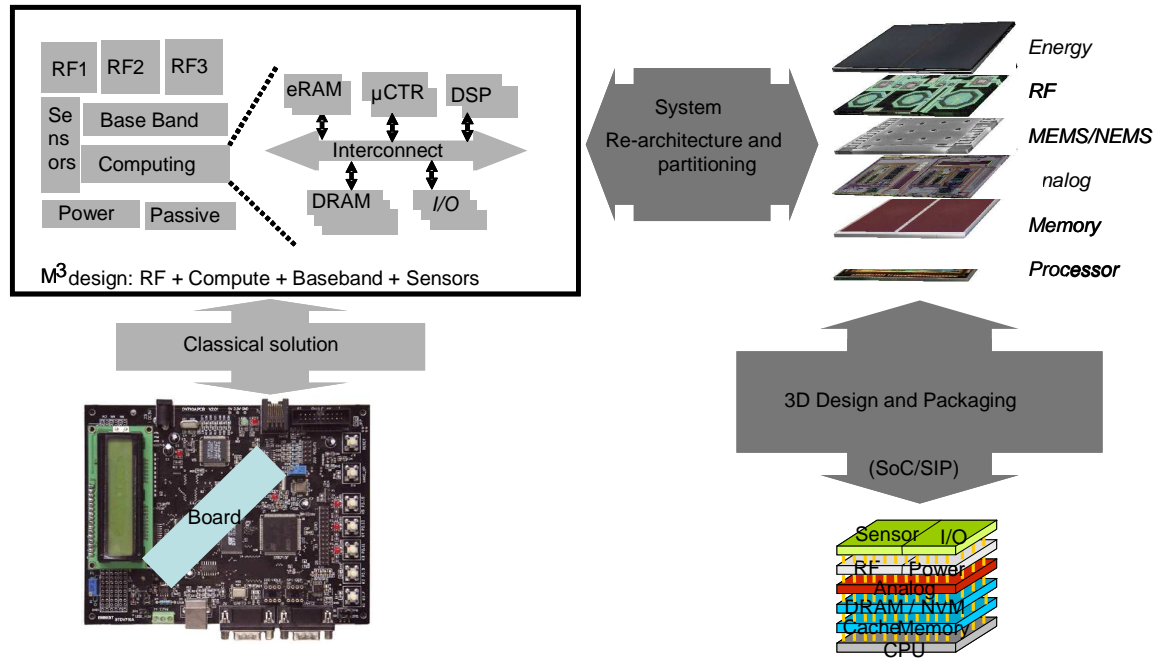


Figure 3-1: 3D integration concept

3D integrated technology is a very attractive option for many advanced consumer products meeting specifications of the next generation of market key drivers such as mobile phones, set-top-boxes and HDTV. By replacing single chip packages with 3D devices, higher transistor density and low power saving are achieved, data travel distances shortens, the manufacturing cost decreases through die reuse generalization. Ultimately, 3D integration helps meeting Tera-scale computing challenges, by allowing increasing memory bandwidth while pushing forward the miniaturization required by the consumer wireless mobile applications while mastering the pin counts.

The first step in achieving 3D integration was to incorporate memories into a 3D memory package, followed by the introduction of memories into a CPU + Memory Multichip Package (MCP), wire bonded. By progressively exploring 2D Planar MCP concept, the next step is to address Substrate Embedded Die MCP and 3D Stacked Die MCP, to achieve more memory bandwidth and introduce the heterogeneous dice assembled into the same package (see figure 3.2).

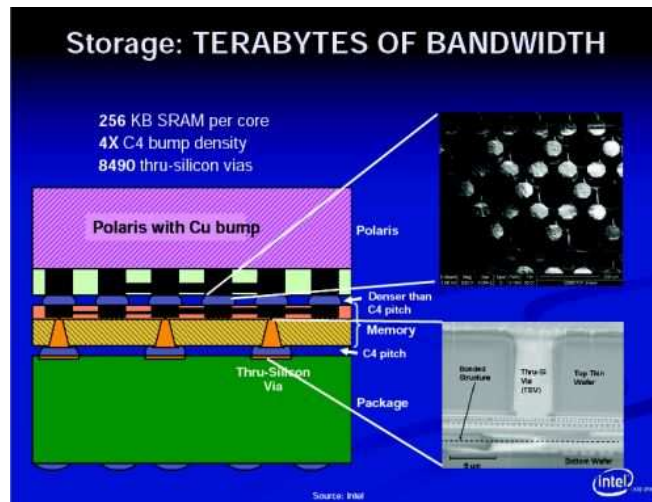


Figure 3-2: Intel's Freya prototype uses stacked dice with through-silicon vias to connect the Polaris cores to memory (Source, Intel)

### SiP technologies

Advanced high speed interfaces have been presented to handle the transfer of large amounts of data between embedded processor cores and main off-chip memories in digital multimedia applications. These approaches support hundreds of Gigabits per second of aggregate I/O bandwidth but they require high power consumption and large chip area occupation. Recently, stacking technologies have been proposed for connecting computing elements and the memory system through micro-bumps to achieve better performance at a lower cost. These solutions address both cost and complexity issues, but they require thousands of bump connections to ensure a high data rate, thus increasing both reliability issues in the flip chip manufacturing process and power consumption due to the large amount of buffers needed to match performance requirements [10].

There are many methods for inter-chip connection, such as wire-bonding, edge connect, capacitive or inductive coupling.

### Flip-chip

Flip Chip describes the method of electrically connecting the die to the package carrier. The package carrier, either substrate or lead frame, then provides the connection from the die to the exterior of the package. In "standard" packaging, the interconnection between the die and the carrier is made using wire. The die is attached to the carrier face up, then a wire is bonded first to the die, then looped and bonded to the carrier. Wires are typically 1-5 mm in length, and 25-35  $\mu\text{m}$  in diameter. In contrast, the interconnection between the die and carrier in flip chip packaging is made through a conductive "bump" that is placed directly on the die surface. The bumped die is then "flipped over" and placed face down, with the bumps connecting to the carrier directly. A bump is typically 70-100  $\mu\text{m}$  high, and 90-125  $\mu\text{m}$  in diameter.

The flip chip connection is generally formed one of two ways: using solder or using conductive adhesive. By far, the most common packaging interconnect is solder. Current solder options are: eutectic (63%Sn, 37%Pb) or high lead (95%Pb, 5%Sn) or lead-free (97.5%Sn, 2.5%Ag) compositions. The solder bumped die is attached to a substrate by a solder reflow process, very similar to the process used to attach BGA balls to the package exterior. After the die is soldered,

underfill is added between the die and the substrate. Underfill is a specially engineered epoxy that fills the area between the die and the carrier, surrounding the solder bumps. It is designed to control the stress in the solder joints caused by the difference in thermal expansion between the silicon die and the carrier. Once cured, the underfill absorbs the stress, reducing the strain on the solder bumps, greatly increasing the life of the finished package. The chip attach and underfill steps are the basics of flip chip interconnect. Beyond this, the remainder of package construction surrounding the die can take many forms and can generally utilize existing manufacturing processes and package formats.




Figure 3-3: Flip-chip concept

Using flip chip interconnect offers a number of possible advantages to the user:

- Reduced signal inductance - because the interconnect is much shorter in length (0.1 mm vs 1-5 mm), the inductance of the signal path is greatly reduced. This is a key factor in high speed communication and switching devices.
- Reduced power/ground inductance - by using flip chip interconnect, power can be brought directly into the core of the die, rather than having to be routed to the edges. This greatly decreases the noise of the core power, improving performance of the silicon.
- Higher signal density - the entire surface of the die can be used for interconnect, rather than just the edges. This is similar to the comparison between QFP and BGA packages. Because flip chip can connect over the surface of the die, it can support vastly larger numbers of interconnects on the same die size.
- Die shrink - for pad limited die (die where size is determined by the edge space required for bond pads), the size of the die can be reduced, saving silicon cost.
- Reduced package footprint - in some cases, the total package size can be reduced using flip chip. This can be achieved either by reducing the die to package edge requirements, since no extra space is required for wires, or by utilizing higher density substrate technology, which allows for reduced package pitch.

### Through Silicon Via

Through Silicon Via (TSV) technology allows stacked silicon chips to be interconnected through direct contact in order to provide high-speed signal processing.

Two main techniques exist: *TSV via last* and *Post Back-End Of Line (BEOL) via first*.

- With the *TSV via last* technique the target is to make a hole through the wafer on the bottom side (top side supposed to have the upper metal layer) reaching the first metal layer; then the hole is metalized and connection assume with a redistribution layer to another connection type (i.e. solder bump).



- With the *Post Back-End Of Line (BEOL) via first* technique the process is applied to a finally processed wafer by etching and metalizing Silicon vias into predefined regions of the chip. The metalized TSVs are connected to the metal layer of the chip by a redistribution layer.

Assembly challenge of this technique is to reach the first metal on a reliable manner without a big hole in order to avoid affecting too much the die size. Holes are done at wafer level on the back side and their size depends on the techniques used to make them and the wafer thickness. 3D wafer level integrated circuits, where interconnections are made at global or intermediate levels of the chip by TSV allows better integration, bandwidth and electrical performances by reducing signal delays, smaller form factor and pin counts, and overall a better packing density. In this type of technology, the most common solutions use **stacked memory devices**, but also **sensors** including CMOS imaging sensors **bonded with DSP**. The future applications aim at packing **multicore processors, caches and memory hierarchy** and other incompatible technologies in a **heterogeneous integration** solution. Figure 3.4 shows the roadmap of the 3D-IC TSV.

Even with the advantages of 3D-IC, there are several major challengers to the adoption of 3D architectures. These challenges need to be overcome for the technology to see widespread adoption:

- Commercial availability CAD tools are required to allow a flexible floor plan, better vertical and horizontal place and route steps, including better thermal modeling and specific DFY constraints.
- System architecture design methodologies based on new hierarchical design flow, system partitioning by using standardized data interfaces.
- Thermal concerns by increased power densities (avoid hotspots overlap).
- Test and reliability issues for repartitioned logic targeting new multiple defects generated by TSV and interferences.

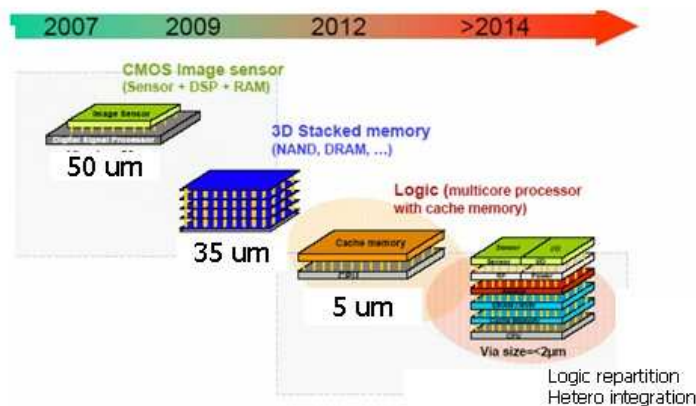


Figure 3-4: 3D-IC TSV Roadmap (courtesy of CEA-Leti)

### ***Chip-to-chip physical channel requirements***

Current multi-die systems developed by STMicroelectronics for consumer applications (Set Top Box, HDTV, etc.) use a 16-bits electrical PHY at the frequency of 450MHz, for an available bandwidth of 7.2Gb/s (monodirectional).

### ***Limits of chip-to-chip electrical channel***

On-chip performance has been increasing much more rapidly than off-chip communication bandwidth because both on-chip transistor density and clock frequency are increasing faster than off-chip input/output density and frequency. This difference occurs because off-chip bonding and wiring are about two orders of magnitude larger than on-chip wiring: on-chip wiring pitch is on the order of 1 micron, while off-chip wiring and ball-bond pitches are on the order of 100 microns, The performance gap between on-chip and off-chip bandwidth makes off-chip bandwidth a performance bottleneck [11].

Advanced high speed interfaces have been presented to handle the transfer of large amounts of data between embedded processor cores and main off-chip memories in digital multimedia applications. These approaches support hundreds of Gigabits per second of aggregate I/O bandwidth but they require high power consumption and large chip area occupation.

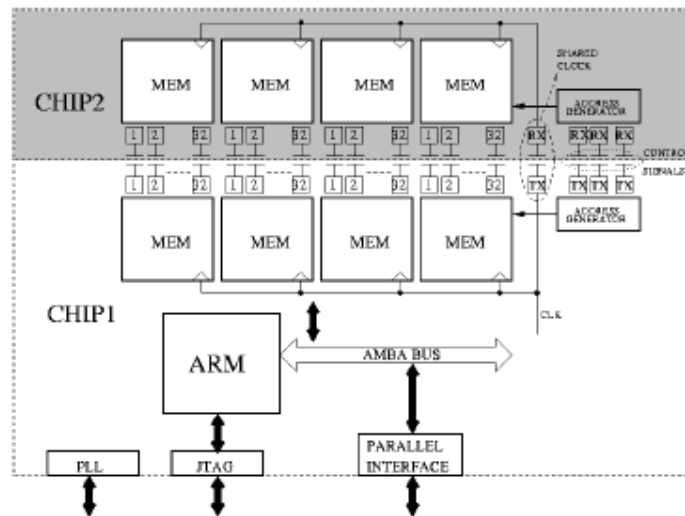


Figure 3-5: System architecture with off-chip memory interface

Recent stacking technologies proposed for connecting computing elements and the memory system through micro address both cost and complexity issues, but they require thousands of bump connections to ensure a high data rate, thus increasing both reliability issues in the flip chip manufacturing process and power consumption due to the large amount of buffers needed to match performance requirements.

Figure 3-6 shows the I/O data rate trend from International Technology Roadmap for Semiconductors. As it is shown in this figure, future systems will require high bit rates for the communication between the various systems.




Figure 3-6: ITRS projected high speed I/O data rates

In the domain of multi-core high performance processors the limitations of the electrical chip-to-chip interconnects has a direct impact on the overall performance of the system. As the number of on-chip processor cores increases, the system performance can keep increasing assuming that the processor cores are fed with instructions and data. Eventually, this is not possible due to the limits of the current electrical interconnects between the processors core and the memory chip (this problem is also known as the memory wall, figure 3-7).

The **memory wall** is defined as a situation where the much faster improvement of processor speed and the increased number of cores in a chip as compared with dynamic random access memory (DRAM) speed will eventually result in processor speed improvements being masked by the relatively slow improvements to DRAM speed and the limited interconnection bandwidth. However, the problem of relative slow improvement in DRAM speed has been avoided by traditional techniques: making caches faster and reducing the miss rate from caches (by increasing the size or associativity, or both). Hence, the main problem remains the slow interconnection bandwidth between the processor and the DRAM memory.




Figure 3-7: Processor-Memory Wall

Currently, the electrical chip-to-chip interconnects are based on high performance transceivers that can reach up to 11Gbps while it is estimated that 28Gbps transceivers will be shipped in 2012. For example, figure 3-8 depicts the block diagram of a typical transceiver that is used to connect FPGA devices with processors, memory or network modules. The transceivers are composed of the Physical Coding Sublayers (PCS) that is used for the digital encoding, the serializer, etc., and Physical Medium Attachment sublayer (PMA) that is used for the physical transmission/reception. Due to the high complexity for the signal integrity, these transceivers consume increased amount of power. Figure 3-9 depicts the current power consumption of the FPGA transceivers that are used for chip-to-chip interconnects for different data rates. As we move to lower process technology (e.g. 28nm supporting 28Gbps) it is estimated that the power consumption of the electrical transceivers will consume a significant portion of the overall power consumption of the systems.

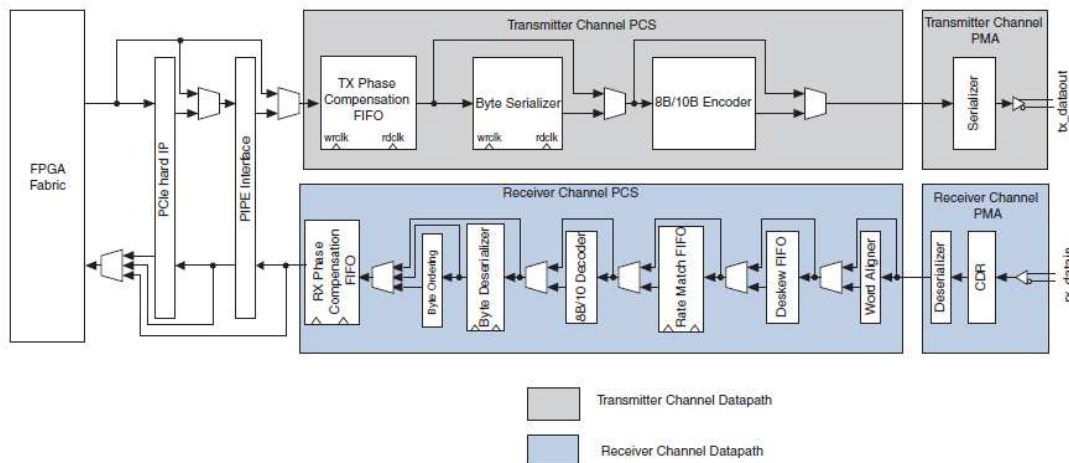


Figure 3-8: High performance chip-to-chip electrical transceivers




Figure 3-9: Power consumption of transceiver

Designers widely use electrical interconnect for chip-to-chip and chip-to-module interfaces over traces on a printed circuit board (PCB), in chip-to-chip over backplane, and in chip-to-chip over copper cable assemblies. At 10 Gbps, the reach distances are approximately 0.3 m for chip-to-chip and chip-to-module interfaces, 1 m for chip-to-chip over backplane, and 7 m for chip-to-chip over copper cable assemblies.

The challenge for electrical based interconnect is that it does not scale with the data rate because of the frequency dependent loss. For example, in the widely used FR-4 copper trace material the loss is ~ 0.5-1.5 dB/in at 5 GHz (Nyquist for 10 Gbps rate), and the loss increases to ~ 2.0-3.0 dB/in at 12.5 GHz (Nyquist for 25 Gbps rate). Return loss and crosstalk can also increase with frequency.

In these copper-based systems, designers typically must compensate for insertion loss signal impairments, such as inter-symbol interference (ISI) or data-dependent jitter (DDJ), return-loss, and crosstalk. Designers adjust for these impairments by using various equalizers, such as a feed-forward equalizer (FFE), continuous time linear equalizer (CTLE), or decision feedback equalizer (DFE), implemented on the transmitter or receiver at the copper channel to ensure that the link performance (that is, bit error rate (BER) <math>< 10^{-12}</math>) is met. However, equalizers consume power and add penalties, especially the DFE. As the data rate increases, insertion loss, return loss, and crosstalk also increase and require even stronger equalizers (that is, more taps or larger DC/AC gains) to compensate for the resulting impairments, and to insure the same performance. This technique in turn, adds more power.

In the case of the telecommunication network and data-center networks, the high power and low performance electrical interconnects have been replaced by optical interconnects. For example in the case of the data center networks, the older copper-based links have been replaced by optical links (e.g. SFP, SFP+ links) that can provide high throughput, low latency and reduced power consumption. At the same time, optical interconnects occupy limited physical space which is used for space saving and better cooling inside the racks. Figure 3-10 depicts the transition of different interconnection schemes from the electrical to the optical domain. Although that for the case of chip-to-chip communications, the electrical technology is still the preferred solution, it is obvious that new technology must be adopted that will meet the future requirements in terms of bandwidth, latency and power consumption.

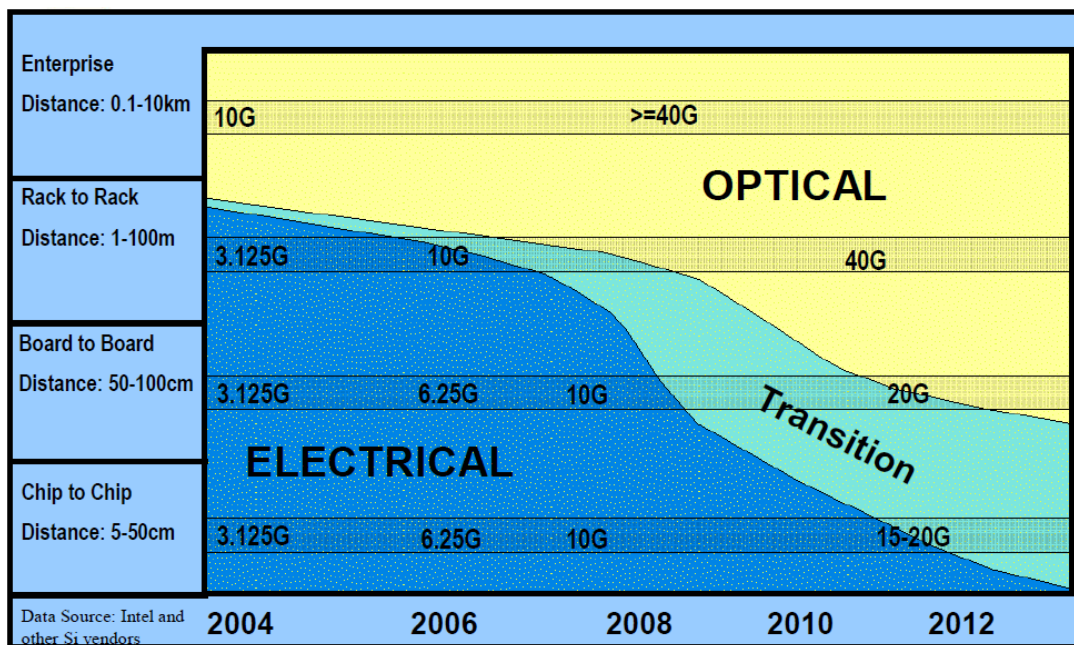


Figure 3-10: Transition to optical domain for different interconnection links

## Novel solutions

### AC-coupling

With this approach dice are stacked on each other and aligned face-to-face; communication electrodes are realized in the upper metal layer of each die and they are connected to dedicated communication circuits; receivers and transmitters exploit the inter-electrode capacitance in order to provide a reliable signal propagation. [12].

This technique has demonstrated the ability to provide a cost-effective integration paradigm: standard packaging procedures provide the assembly accuracy required by contactless connections based on capacitive coupling or on inductive coupling. Works in these fields have proved a throughput of more than 1 Gb/s/pin. The main difference between interconnection approaches based on capacitive or inductive coupling lays in the fact that the inductive signaling is based on the magnetic field flowing through coupled inductors, and the intensity and effectiveness of this field can be more easily increased (by increasing the current flowing in the inductor or the number of turns) with respect to the strength of the electric signal used for the capacitive approach (that is voltage driven).

This leads to an even lower cost for the assembly of systems based on inductive interconnections: chips can be assembled in a face-up configuration thanks to the larger transmission power that is enabled by the current driven approach. The face-up assembly proves to be simple to implement and cost-effective but, on the other hand, the increased distance between the communication structures requires larger size for the vertical interconnections as well as larger power consumption; so, power and area result in being larger than the ones related to solutions with face-to-face assembly and based on capacitive coupling. For these reasons, 3-D communication based on capacitive coupling seems more attractive.




Figure 3-11: Principle of capacitive interconnections

Communication electrodes are realized in the upper metal layer of each die and they are connected to dedicated communication circuits; receivers and transmitters exploit the inter-electrode capacitance in order to provide a reliable signal propagation.

### Optical interconnect

Optical chip-to-chip connections are an active area of investigation within both the industrial and academic sectors of the optical communications market. The attractiveness of the concept for semiconductor suppliers and their customers is the possibility of realizing optical links across high-speed electronic backplanes and motherboards.

The desire for an optical chip-to-chip solution is driven by the I/O needs of future communication systems and the increasingly complex ASICs, microprocessors and digital signal processors that support the system architectures. An optical chip-to-chip communication scheme is an attractive solution to the power, density and signal isolation issues in high-throughput, compact systems.

When considering such schemes, it is important to understand that electrons will continue to power the data-processing engine while photons will be the data path conduit. This means the optical solution must be compatible with the electronics components and vice versa. Further, the optical solution will compete to replace a low-cost, optimized solution with strong support across standards bodies and manufacturers, so it must offer compelling benefits to succeed.

Electronic I/O management schemes are typically serializer/deserializer devices and are monolithically integrated in the ICs. Using ser/des to multiplex the data processed by each IC reduces the resulting number of I/O ports. The appeal of an optical solution lies in its ability to further maximize the data rate and distance of data transmission between mixed optoelectronic-VLSI integrated circuits (OE-VLSI ICs). An optical solution can also further reduce the number of connections using wavelength-division multiplexing (WDM) or optical time-division multiplexing (TDM).

The design community is approaching consensus on a hybrid solution in which the optical components are manufactured independently and attached to the electronic ones via flip-chip, multichip-module or system-on-package schemes. Most of the solutions investigated to date have focused on a vertical approach, typically utilizing vertical-cavity surface-emitting lasers or vertical modulators and detectors. A single additional chip containing all the optical functionality would likewise be an attractive solution. Another possibility is to embed thinned optical devices in polymer waveguides and connect the electronics to the substrate or package containing the polymer waveguides.




Figure 3-12: Chip-to-chip optical interconnect concept

A number of critical issues remain, including multimode vs. single mode, on-chip vs. off-chip optical source, WDM vs. TDM or no multiplexing, and polymer waveguides vs. free-space connections. For data rates of 10 Gbits/second and above (the speed at which an optical chip-to-chip solution begins to be attractive), multimode solutions are challenging because of modal dispersion.

On-chip sources will be difficult to manage, especially thermally, for high-I/O-port ICs such as microprocessors; off-chip sources have the advantage of being independently controlled and monitored, yet coupling the off-chip light source to the optical chip or modulator adds complexity. WDM solutions require a minimal number of input/output paths. Manipulation of

WDM is typically space-consuming, however, and requires either multiple-wavelength, precise lasers or a single, costly, mode-locked laser.

The polymer quality or free-space architecture is particularly important when considering reasonable distances between the communicating chips-and more significantly so for single-mode than multimode solutions. Sophisticated packaging schemes will also be required to address thermal management of dissimilar materials and temperature-sensitive devices, coupling and the need for electronic packaging pick-and-place alignment tolerance.

At the moment, both optical and electronic suppliers are evaluating these paths. The ultimate implementation will lie in a joint solution that fits well into the economic model of all parties.

### ***Potential benefits of plasmonics***

A communication technology based on plasmonics should allow to overcome the bandwidth, foot-print and power consumption limitations of today electrical and optical interconnect solutions.

Such a technology would exploit the ultra-compact dimensions and fast electronic interaction times offered by surface plasmon polaritons to build plasmonic transceivers with a few square-micron footprints and speeds only limited by the RC constants.

The transceivers will be interconnected by free space and fiber connect schemes. The plasmonic transceiver concept aims at overcoming the challenges posed by the need for massive parallel interchip communications. Yet, it is more fundamental as the availability of cheap miniaturized transmitters and detectors on a single chip will enable new applications in sensing, biomedical testing and many other fields where masses of lasers and detectors are need to e.g. analyze samples.

Economically, the suggested technology would be a viable approach for a massive monolithic integration of optoelectronic functions on Si substrates as it relies to the most part on the standardized processes offered by the silicon industry. In addition, the design and production cost of plasmonic devices are extremely low and with the dimension 100 times smaller over conventional devices they will require much lower energy to transfer data over short ranges of multi-processor cluster systems.




Figure 3-13: Principle of plasmonic LASER

With respect to optics, electronics is limited in operation speed. In silicon-photonics optical interconnects, the bandwidth capability of light is utilized to overcome the electronics data rate limitations (light is about three orders of magnitude faster [13]). On the other hand, photonics is limited in miniaturization capability (more than an order of magnitude larger than electronics), because the spatial cross section of conventional light signals is of the order of the operating



wavelength (e.g., a few hundred nanometers for telecom wavelengths) [14]; hence, it is problematic to produce hybrid electronic/photonic chips and achieve acceptable scales of integration. Ideally, one would like to combine the advantages of each technology; electronics-level miniaturization and photonics-level data rates. Plasmonics is a technology that promises to bridge the gap between electronics and photonics, and combine the best of both worlds. Plasmonic interconnects are optical interconnects where surface plasmon polaritons (SPPs) are utilized as information carriers. SPPs constitute of charge density oscillations at the interface between a material of positive permittivity (i.e., a dielectric or a semiconductor) and a material of negative permittivity (traditionally a metal, but semiconductors can also be used at special frequency ranges) [15]. Accompanying the charge oscillations, there is an electromagnetic field propagating along the interface. This electromagnetic field can be used as the optical signal to carry information at optical bandwidths. From the solution of Maxwell's equations for plasmonic structures, it turns out that SPPs can feature subwavelength cross sections; thus, they allow for the design of exceptionally compact optical devices, beyond the diffraction limit that restricts conventional photonics. In addition, the presence of metallic parts in plasmonic geometries allows for the design of structures where both electric currents and optical waves can propagate as signals.

The theory behind plasmonics has attracted attention since the late 60s [16]; but it is only recently that practical application for plasmonic devices has attracted focus, due to fabrication and material challenges. In the last few years, several plasmonic devices have been experimentally demonstrated, like plasmonic nanolasers (optically-pumped [17-20] as well as electrically pumped [21-22]; see also [23] for a review), modulators [24-26], amplifiers [27-29], polarizers [30], couplers [31-32], waveguides [31, 33-34], beam shapers [35], photodetectors [36], slow light devices [37], resonators and interferometers [38].

In 2008, a plasmonic on-board chip-to-chip optical interconnect was demonstrated by researchers in Korea [39]. VCSELs were used as transmitters at 1.3  $\mu\text{m}$  and the transmitted light was received with a photodiode. An array of 4 plasmonic waveguides with a length of 2.5 cm was used for signal propagation, for a total bit rate of  $4 \times 2.5 \text{ Gb/s} = 10 \text{ Gb/s}$ . Note that the Korean researchers used mature technologies for the transceivers, whereas NAVOLCHI aims to implement plasmonic technology for the transceivers.

In 2011, the EU-funded project PLATON consortium demonstrated a  $2 \times 2$  silicon-plasmonic router architecture with 320 Gb/s throughput capabilities for back-plane or Blade-Server optical interconnect applications, supporting  $2 \times 2$  thermo-optic switch operation [40, 41]. In comparison to PLATON, NAVOLCHI includes a focus on chip-to-chip interconnection, as well as on-chip plasmonic transceivers.