FP6-2002-IST-1-002131

PICMOS

Photonic Interconnect Layer on CMOS by wafer-scale integration

STReP - Specific Targeted Research Project

IST – Information Society Technologies

D1.1. Definition of applications

Contributors	
LEOM	**I. O'Connor** F. Gaffiot
UGent – PARIS	D. Stroobandt J. Dambre M. De Wilde

Date of preparation : Mar. 30 2005

Project Co-ordinator: D. Van Thourhout *Revision*: 1.1

Dissemination level: PU (public)

Table of Contents

List of Figures

List of Tables

1. Introduction

The semiconductor industry is moving quickly towards applications requiring operating frequencies of the order of tens of GHz (the 10GHz barrier should be reached in 2008 according to the ITRS, Figure 1). This is fuelling the demand for greater device and interconnect bandwidth. The leading semiconductor companies are investing heavily in fabrication technology capable of producing devices with feature sizes in the deep-submicrometer/nanometer range (the printed gate length of devices should be around 25nm in 2007). This increasing level of integration is the main route to achieving ever-higher computing speed.

One of the most urgent needs concerns the development of new interconnect technologies for systems on chip. This report will analyse the origins of this current concern (section 2) and the alternatives to classical interconnect techniques (section 4). We continue with an analysis of advantages and disadvantages of the optical interconnect solution at the root of the PICMOS project (section 5), and detail potential applications in section 6.

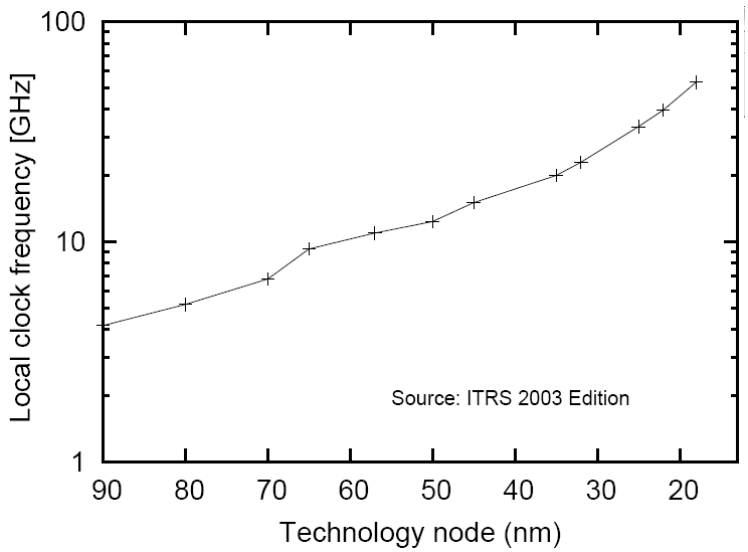

Figure 1 Progression of local clock frequency with technology node

2. Current and future integrated interconnect issues

2.1. The influence of classical scaling strategies on interconnect

Two main factors influence the overall operating frequency of a chip: (i) device switching times and (ii) interconnect bandwidth. Current technologies can achieve unprecedented transistor transition frequencies 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, and has been for some time, the IC design bottleneck from computing capacity to communication

bandwidth and flexibility. Interconnects are thus emerging as the major obstacle to increasing system operating frequencies (Figure 2).

The operating frequency of an integrated circuit was until recently largely dominated by the maximum frequency f_{max} of MOS transistors. Based on single-pole roll-off in transistor characteristics, devices are as a general rule useful for circuit operating frequencies around $f_{max}/10$. There are a few ways to improve device speed, but the main technique that has been employed for several decades is scaling. MOS maximum frequency is inversely proportional to transit time t_d, roughly approximated[1] by

$$t_d = \frac{L^2}{mV_{ds}}$$

where L represents transistor gate length, μ is carrier mobility and V_{ds} is the drain-source voltage. It is therefore clear that reducing L by a factor S^2 increases f_{max} by $1/S^2$. Other quantities must also scale with L: gate oxide (dielectric) thickness t_{ox} reduces, gate dielectric permittivity increases to reduce leakage, junction depth reduces, supply voltages reduce and doping concentration increases to keep the electric field constant over decreased transistor dimensions. Of course, by reducing L, the integration density (number of transistors per cm^2) also improves. Consequently for constant circuit area, complexity rises with technology node (as does power...)

However, device performance is only part of the equation when it comes to considering overall circuit performance. Wire delay becomes preponderant from the 250nm node on, as shown by the ITRS (Figure 2). According to this figure, wire delay already represented 80% of the overall delay of critical paths at the 100nm node.

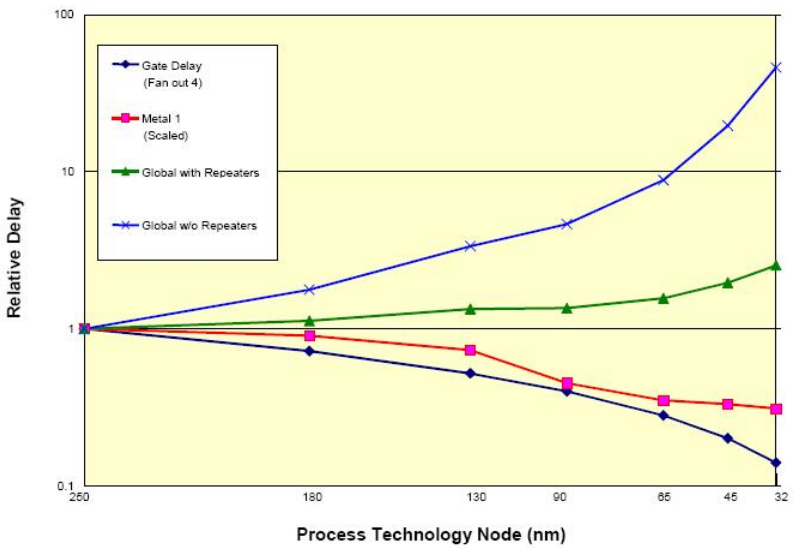

Figure 2 Gate and wire scaling

[1] Note that this gives an overly optimistic view since V_{ds} also scales (though not as much as L)
[2] S represents the scaling factor and is approximately equal to 0.7 per technology generation.

The most essential point to realise is that the physical structure of metallic interconnect cannot scale with transistor size, a point which has many consequences. This can easily be understood through the analysis of basic physical quantities at the root of the most fundamental characteristics of any electrical interconnect, its RC time constant. This quantity defines the maximum signal frequency (proportional to 1/RC) that can be usefully transmitted over an interconnect of a given length, and thus depends upon the unit length resistance and capacitance.

1.1.1. Unit length resistance

For wires, scaling results in the reduction of lateral wire dimensions h and w (wire height and width[3], respectively). The cross-sectional area is therefore reduced by S^2 and so according to

$$R_{wire} = \frac{r}{(h - t_b)(w - t_b)}$$

(where ρ represents the metal resistivity and t_b represents the thickness of the barrier layer[4], shown in Figure 3), the resistance per unit length of wires (for given w) R_{wire} would be increased by a factor slightly less than $1/S^2$.

To counter this problem where the planar footprint w is reduced by S, the cross-sectional area can be maintained constant by using the inverse scaling strategy – by *increasing* the height of conductors, h, by the same factor, S. However, sidewall (and therefore wire-to-wire) capacitance increases with wire height, and this effect is further exacerbated by reduced wire-to-wire spacing (demanded by integration density requirements). Also, conductor height cannot be increased indefinitely for surface roughness reasons. This constraint is reduced in the upper metal layers. It is common practice to route critical links in these layers. Today, in general, we can consider that R_{wire} increases proportionally to $1/S^{\left(2-\frac{n}{N}\right)}$, where N represents the total number of metal layers and n the layer number corresponding to the metal routing layer concerned (from 1 at local routing to N at global routing)

[3] Perpendicular to current flow.

[4] Copper interconnect requires a thin barrier layer on three sides to prevent the metal from diffusing into the surrounding oxide. The barrier thickness t_b for the 0.18μm generation technology is 17nm. Aluminium does not require a barrier layer, explaining why wire resistance did not quite decrease by 50% (although it did decrease significantly) as technologies migrated from aluminium based interconnect to copper.

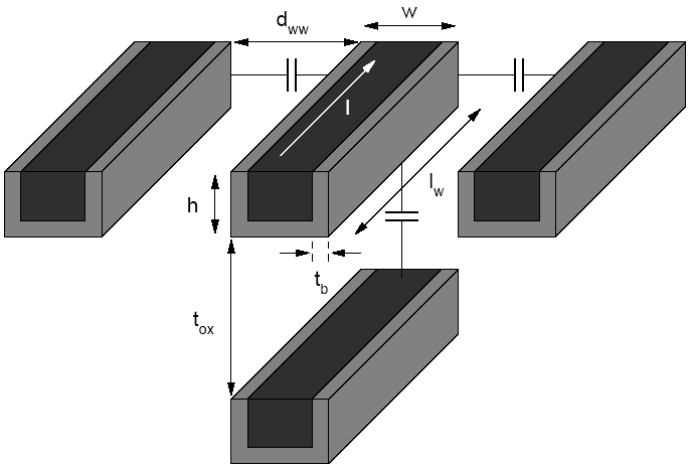

Figure 3 Definition of interconnect dimensions

Thus the only remaining variable to bring down the resistance per unit length of a wire R_{wire} is the metal resistivity. This strategy has already been used in the past by replacing aluminium wires ($\rho = 3.3\mu\Omega\cdot cm$) by copper wires ($\rho = 2.2\mu\Omega\cdot cm$[5] [KAP2002]). Copper also has better electromigration behaviour since it is a denser metal ($8.9mg\cdot cm^{-3}$ for copper against $2.7mg\cdot cm^{-3}$ for aluminium), such that copper interconnect can carry higher currents densities as compared to aluminium interconnect. However, this approach of reducing resistivity through material replacement has without doubt reached its limits. The only metal with resistivity lower than that of copper is silver ($\rho = 1.6\mu\Omega\cdot cm$).

It should be noted here that the skin effect for most on-chip wires is largely negligible at current technology nodes since the skin depth usually exceeds the wire dimensions. The skin depth, δ, is given as

$$d = \sqrt{\frac{r}{pfm}}$$

where f is the frequency of interest, μ is the magnetic permeability of the conductor material. When the skin depth becomes less then the dimensions of the interconnection, the resistance increases and the skin effect must be considered. For example, at a frequency of 1.5GHz the skin depth of copper is 1.7μm, exceeding the vast majority of wire dimensions. The technology node for which the signal frequency is sufficiently high to cause the skin depth of copper to be less than the smallest dimension of the upper metal levels has not yet been reached.

The absolute value of resistance of a wire also depends on its length. Although the length of the local interconnections can decrease with scaling, this is not sufficient to reduce their resistance. Between two technology nodes, the length of an interconnect is reduced by a factor S, but R_{wire} will be increased by a similar factor. Hence the absolute value of

[5] The theoretical value of copper resistivity is $\rho = 1.7\mu\Omega\cdot cm$. The effective value is higher due to several reasons: temperature, barrier layer dimensions, technology, scattering phenomena.

the resistance of this interconnection will be maintained or at best will undergo a slight reduction; insufficient in any event compared to the necessary frequency of operation.

Moreover, interconnect length distribution models (Figure 4) [DAM2004] show that, for roughly constant chip size (in the neighbourhood of 20mm x 20mm), the total number of wires roughly scales as $1/S^2$ with technology scaling (somewhat less due to the increase of high-fanout nets). For global interconnects, where optical interconnect technology can make the most impact, no precise trend can be defined other than that their number increases. In fact, the definition of 'global' interconnect is not clear in itself. They can be defined as being "longer than a constant number of gate pitches", or as being "longer than a constant length". In the first case an initial analysis would show that the total number of global wires scales approximately with $1/S^2$, whereas in the second their number scales almost according to 1/S. However, it should be noted that these trends have not been shown to be valid and should thus only be taken as indications. It is nonetheless clear that the global interconnect problem is not only aggravated by material parameter limitations and scaling dimensions, but also by the increasing number of such global interconnects.

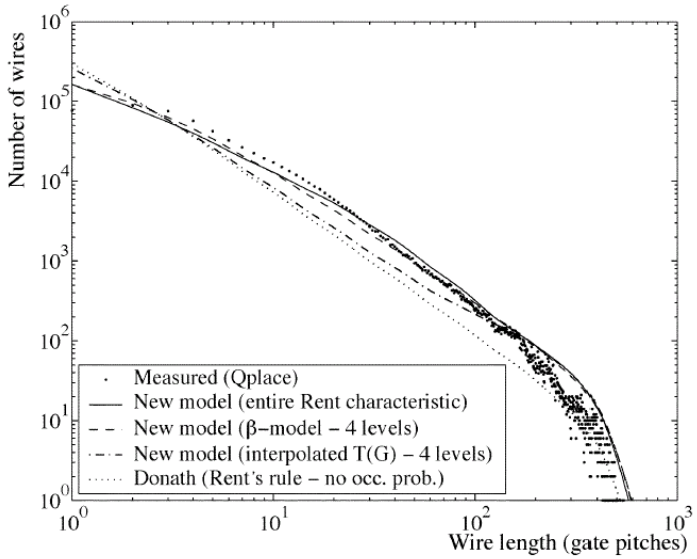

Figure 4 Interconnect length distribution in a typical processor[6]

1.1.2. Unit length capacitance

Means are also being investigated by which capacitance, the other part of the RC time constant, can be reduced. For a given width w, (Figure 3), the capacitance per unit length, C_{wire}, can be modelled [HO2001] by

[6] In the current SoC and NoC context, this type of distribution has to be called into question. It is mostly assumed that this type of length distribution is still valid for processors and for large IP-blocks, but for SoC-type of designs, especially with NoC, the shape of the distribution is likely to change, since there is no longer a 1-on-1 relationship between communication signals and physical wires

$$C_{wire} = e_0\left(2Ke_h\frac{h}{d_{ww}} + 2e_v\frac{w}{t_{ox}}\right) + fringe(e_h, e_v)$$

where ε_0, ε_h and ε_v are the absolute, horizontal relative and vertical relative dielectric constants[7]; K represents the switching-dependent effective capacitance factor varying between 0 and 2[8]; d_{ww} represents wire to wire spacing; t_{ox} represents inter-layer dielectric thickness and the fringe term is weakly dependent upon geometry and is about 40fF/μm for the 0.18μm technology for example. The first term in the equation shows that sidewall capacitance could be reduced either by decreasing h (in direct opposition to the resistance requirement, as stated previously) or by increasing wire to wire spacing (in direct opposition to integration density demands, which is not acceptable either). The only remaining quantity to reduce this term is the dielectric constant value $\varepsilon_{h.}$ In the second term, we can see that top and bottom capacitance can be reduced either

- by increasing inter-layer dielectric thickness (but this would increase via length and therefore via resistance), or
- by decreasing w (a common aim with integration density, w is already minimum for a given technology and lithographic accuracy).

So again, the only remaining quantity that can be varied is a material constant, the dielectric constant value ε_v. Hence research efforts have been concentrating on numerous low-k organic and inorganic materials with a wide range of dielectric constants (from air, k=1, to fluorinated oxides, k=3.6). While relative permittivites of less than 2 (at best 1.7-1.8) can be achieved by the use of SiOC-type nanoporous materials or SiLK-type organic materials with an "air gap" approach, the integration complexity of these approaches is higher, the mechanical properties are weaker and moisture or chemical absorption may also increase. In fact, integration and reliability issues that have impeded implementation of low-k dielectrics include thermally or mechanically induced cracking or loss of adhesion, poor mechanical strength, moisture absorption, time-dependent behaviour, chemical interactions (especially during photolithography, etch/clean and dielectric/metal deposition), low electrical breakdown, poor thermal conductivity [HAV2001]. In any event, it will be physically impossible to reduce the dielectric constant by a factor of more than 2-3, since this would imply a material with a dielectric constant lower than 1, that of air.

2.2. Propagation delay

For a given metallic wire of length l_w, the actual resistance and capacitance of the wire will be $R_{wire}\cdot l_w$ and $C_{wire}\cdot l_w$ respectively. An expression for the propagation delay of a signal transmitted from an emitter gate to a receiver gate has been developed in [SAK1983] as:

[7] Dielectrics may be different within layers (ε_h) and between them (ε_v)

[8] If the signals in the left and right neighbouring wires switch in the opposite direction to the signal in the wire under analysis, the effective sidewall capacitances double; while if they switch in the same direction then the effective sidewall capacitance approaches zero. This effect is known as "Miller multiplication". It is supposed in the equation that horizontally adjacent wires carry data while vertically adjacent wires carry supply or ground voltages.

$$t_d = R_{out}(C_{out} + C_L) + R_{out} cl_W + 0{,}4\left[\left(crl_W^2\right)^{1,6} + t_{tof}^{1,6}\right]^{1/1,6} + 0{,}7 rl_w C_L$$

In this expression, R_{out} and C_{out} are, respectively, the output resistance and output capacitance of the emitter gate; C_L is the input capacitance of the receiver gate; and τ_{tof} is the time of flight (i.e. the length of the line divided by the speed of the electromagnetic field).

In the case of global links, this formula shows that the delay in the line becomes predominant. Apart from the strategies described previously to reduce in a limited way the RC time constants of metallic interconnect, gate sizing makes it possible to minimise t_d, and it is possible to show that t_d varies with l_w^2. This increase of delay time with the *square* of line length cannot be avoided. The standard technique to mitigate this problem is to insert an even number of inverter gates along the metal wire to regenerate the signal at distributed points. Such repeater insertion (Figure 5) makes the passive metallic wire into an active element and reduces the influence of length on delay from quadratic to linear when correctly designed. Of course, this comes at the cost of a large number of repeaters, which are thus being diverted from use in data processing to use in interconnect. As the metal used in the wire is usually at a rather high level (to reduce resistance), two via stacks are also required at each repeater point, and can thus become routing obstacles in themselves. As frequencies increase with technology nodes, this technique will require an increasingly high percentage of silicon real estate (possibly as high as 25% at the 45nm node as will be shown in Figure 8) and IC power consumption (clock signal distribution alone can account for up to 50% of total chip power), and will be difficult to sustain.

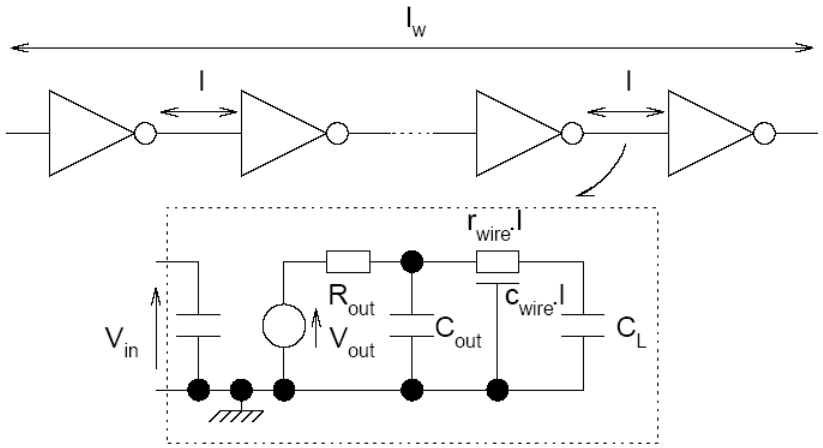

Figure 5 Repeater chain and equivalent circuit

2.3. Predicted "end of the road" for metallic interconnect

This problem has been identified for a long time, but in the 2003 edition of the ITRS[9] roadmap [ITR2003], the interconnect problem was given a new dimension and was

[9] International Technology Roadmap for Semiconductors

summarised thus: "For the long term, material innovation with traditional scaling will no longer satisfy performance requirements. Interconnect innovation with optical, RF, or vertical integration ... will deliver the solution". Table 1, drawn from the 2003 ITRS [ITR2003], shows that by 2010, high performance integrated circuits will count up to two billion transistors per chip and will work with clock frequencies of the order of 10GHz. Coping with electrical interconnects under these conditions will be a formidable task.

	2005	2007	2010	2013	2016	2018
DRAM half-pitch (nm)	80	65	45	32	22	18
Local wiring pitch (nm)	190	152	108	76	54	42
Chip size at production (mm^2)	310	310	310	310	310	310
Functions per chip (MTransistors)	697	1108	2212	4424	8848	14046
Total local interconnect length (m/cm^2)	907	1117	1784	2544	4208	5035
On-chip local clock (GHz)	5.204	9.285	15.079	22.980	29.683	53.207
Max. number of metal levels	11	11	12	12	14	14
Max. power (W)	167	189	218	251	288	300
Max. package pin-count	1760	2140	2782	3616	4702	5426

Table 1 Trends in some integrated circuit parameters

Moreover, this is only part of the picture. It is important to realise that timing is already no longer the sole concern with the physical characteristics of metallic interconnect, although it remains the predominant issue especially when considering propagation time constraints on wires where the length is of the order of the chip size (above 1cm). It is becoming increasingly difficult to ignore several other non-ideal effects, opposed to digital design[10]. Decreasing feature sizes and increased switching activity result in higher capacitive and inductive coupling (interconnect noise). Lower supply voltages render signals more vulnerable to this noise, and also to voltage droop. Two ways of reducing the influence of delay and noise constraints is to increase wire spacing and to use wire-shielding techniques. Both of these approaches cause interconnect resources to be used less efficiently and consequently result in routing congestion or even in non-routability. Finally, the use of metal tracks to transport a signal over a chip has a high cost in terms of power. Power consumption due to capacitive wire charging and repeater insertion contributes to an increase in the silicon real estate and energy required to transfer data even over relatively short distances. Clock distribution for instance requires a significant part (30-50%) of total chip power in high-performance microprocessors[11]. Furthermore,

[10] Digital design is based on abstract descriptions of functions. Taking into account physical phenomena is inherently opposed to this approach.

[11] A distinction has to be made here between levels in a hierarchical clock tree (as will be described in section 6.2). A clock tree is split into three parts: global (upper routing levels), secondary (intermediate routing levels) and local (lower routing levels). It seems fair to assume that while architectural modifications such as GALS act on both global and secondary clock trees, technological modifications to

since metallic interconnect based SoC power dissipation is intrinsically linked to switching frequency, tomorrow's architectures will require power over the 100W mark to be able to operate in the 10GHz range and above. At this level, thermal problems will jeopardise system performance if not strictly controlled.

Thus, even with the most optimistic estimates for RC time constants using low-resistance metals such as copper as well as low-k dielectric materials, global interconnect performance required for future generations of SoCs, as predicted by the ITRS, cannot be achieved with metal. The problem is that means of improving one criterion often comes at the expense of another. It is difficult but essential to achieve the simultaneous improvement of all the important criteria for a given application. Identification of important parameters is covered in section 3.

Naturally, for semiconductor companies the ultimate criterion is cost. The budgets at stake in IC fabrication are so high ($3bn for a new fabrication plant in 2002, $100bn predicted in 2015) that a single strategic error can prove fatal to a company. For this reason, the semiconductor industry in general is very conservative and will impose stringent constraints on any new technology:

- compatibility with existing process technologies or requiring minimal process modifications (highly-integrated CMOS or SOI are the preferred technologies)
- high return on investment (economic considerations)

In the following sections, we describe alternative approaches to metallic interconnect. Initially however, we need metrics to compare these alternatives. The following section discusses this point.

3. Important criteria for physical interconnect evaluation

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. Some of these factors can be integrated into a technology figure of merit (outside the scope of this report) which would enable a quick quantification of tradeoffs that can be achieved with a given technology.

Signal quality

Real signals are characterised by their actual frequency content and by their voltage or current value limits. The frequency content will define the necessary channel bandwidth, according to Shannon-Hartley's theorem[12]. Analogue signals are highly sensitive to degradation and the preferred mode of signal transmission over interconnect is digital. Hence the values of the signal are binary, enabling signal regeneration.

---

the clock tree structure would only act on the global distribution part. This, however, only accounts for about 10% of the total clock tree power consumption.

[12] The capacity of a communication channel in bit/s C=B(1+(Ps/Pn)), where B represents the channel bandwidth (Hz) and Ps and Pn represent, respectively, the signal and noise power values (in W).

The signal quality is a function of bw and dr, where bw represents transmission medium bandwidth (Hz), dr represents transmission medium dynamic range = smax/smin, smax is maximum value of signal amplitude (V) (distortion or saturation-limited) and smin represents minimum value of signal amplitude (V) (noise-limited)

Delay[13]
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. Additionally, delay can be induced by crosstalk.

The delay is function of di, dr and dc, where di is pure interconnect delay, dr is regeneration and associated circuitry delay and dc is crosstalk-induced delay.

Signal degradation
There are several types of signal degradation, which can be classed as time-based, inherent and externally induced.

- time-based: non-zero rise-time, overshoot, undershoot, ringing. time-based degradation can be incorporated into the delay term for digital signals. While the whole of these degradations can be assimilated to a quasi-deterministic behaviour which does not exceed the noise margins of the digital circuits, a transformation in temporal space is possible (to contribute to the regeneration delay term). This assumption is however destined to disappear with nanometric technologies, because of a more probabilistic behaviour and especially of weaker noise margins.
- inherent: attenuation (dB/cm), skin effect, reflections (dB)
- externally induced: crosstalk (dB/cm), sensitivity to ambient noise [CAI2001]

Transmission distance
Transmission distance is in practice limited by the allowable tolerance on signal degradation and delay for a given bandwidth and power budget. The maximum interconnect segment length can in fact be calculated, where a segment is defined as a portion of interconnect not requiring regeneration at the required bandwidth and attenuation limits.

Power
Signal transmission always requires power. In the simplest case, it is required to change the charge value on the equivalent capacitor of a metallic wire. In more realistic cases, power will also be required in emitter and receiver circuitry, and in regeneration circuits. A distinction can also be made between static and dynamic power consumption by introducing a factor α representing the switching activity of the interconnect link ($0<\alpha<1$).

[13] It should be noted that latency is linked to delay but is a system-level parameter taking into account transmission protocol requirements. Latency represents the time it takes a bit of data to traverse through a network link (including transmission request/acknowledge exchanges), while delay is the time difference between the transmission of data to its receipt at the other end of the media (only physical delays intervene). Since only criteria for ***physical*** interconnect are taken into account in this section, latency does not appear. It is considered at a higher level in section 6.

Physical footprint
The physical footprint is function of the area of the emitter and receiver circuitry, as well as that of the transmission medium. The medium footprint term must take into account the physical area of the interconnect (length and width) as well as the pitch between adjacent interconnect media.

Electromagnetic compatibility
Signal transmission requires the regeneration of a signal at a receiver point spatially distant from its emission point. The energy used to propagate the signal in the transmission medium can escape into the surrounding environment and perturb the operation of elements close to the transmission path. This criterion takes into account the amount of energy (in terms of amplitude and frequency) leaving the transmission path along its entire length.

4. Alternative interconnect strategies

The problem facing us, as described in section 2, is that evolutionary solutions will not be sufficient to meet the performance roadmap developed by the ITRS. To tackle the issues developed above, radically different interconnect approaches displaying a highly improved data-rate to power ratio must be developed. At present, the most prominent ideas are the use of integrated radio frequency or microwave interconnects [CHA2001], 3D (non-planar) integration [BAN2001] and optical interconnects [MIL2000] (the PICMOS project focuses on this last concept). More advanced concepts such as nanowires and integrated superconductors cannot be considered to be sufficiently mature to realistically represent an alternative in the timescale of interest.

4.1. Radiofrequency interconnect

In general, radio-frequency signals can be transmitted through free space or waveguide mediums. At chip scale, the free space transmission is not possible because the optimal dimensions of the antenna are approximately $1mm^2$. On the other hand, guided transmission (by MTL[14] or CPW[15] of widths typically from 10-100µm) can be carried out with attenuation lower than 2dB/cm up to 200GHz. A wireless LAN type interconnect architecture (Figure 6) can exploit this high bandwidth by associating IP blocks, RF transceivers, capacitive couplers as near-field antennas and MTL/CPW as shared RF communication channel [CHA2001]. The access mode, managed by the transceivers, would exploit CDMA[16] or FDMA[17] encoding techniques in order to reduce any inter-channel interference.

[14] Microstrip transmission line
[15] Coplanar waveguide
[16] Code division multiple-access
[17] Frequency division multiple-access

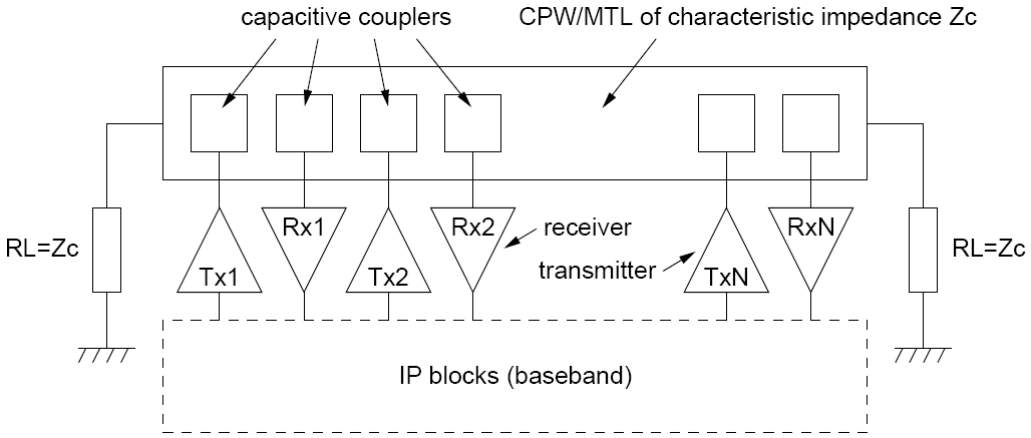

Figure 6 Integrated RF interconnect channel concept

The advantage of this technique lies primarily in the possibility of offering several IP blocks access to a broadband communication channel through a single physical channel, without major process modifications. This solution is thus intrinsically adapted to multiprocessor architectures for parallel computing, and the same approach could also alleviate the constraints on the input/output pin density. Concerning reconfigurable and fault-tolerant architectures, this solution is also relatively attractive. The access algorithm authorises each channel to allot itself a unique address code electronically, thus allowing dynamic reconfiguration of the network. Nevertheless, although the MTL/CPW bandwidth is very high, the fundamental problem with this approach is the inherent incompatibility of the CMOS circuit operating frequencies (GHz) with the maximum carrier frequency (100GHz). The bandwidth could be fully exploited only by using heterogeneous integration for high-speed transceiver circuits. Whatever the technological solution selected, it is clear that the complexity, silicon area and power consumption of RF interconnect are very high.

4.2. Vertical (3D) integration

3D Integration (schematised in Figure 7) allows the creation of multi-layer integrated circuits.

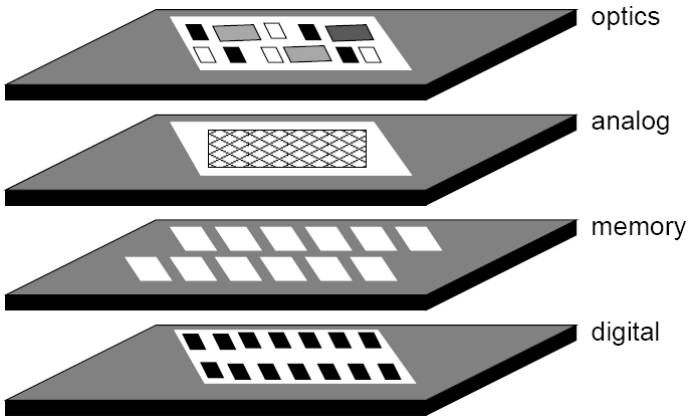

Figure 7 Principle of 3D heterogeneous integration

The essential idea is that a complete planar (2D) chip is divided into a number of blocks, and each block is integrated in an individual silicon layer. The layers are then stacked vertically. Each layer has several levels of interconnection, and dedicated interlayer "vias" ensure the electric interconnection between the layers. The 3D architecture approach offers more flexibility in terms of architectural design, placement and routing. Moreover, *global* interconnections would be all but eliminated by using short interlayer vias to route vertically and not in the plane.

However, difficult technological barriers are opposed to the use of this technique. Heterogeneous integration is still a long way from reaching the status of mature technology. It also appears rather clear that physical problems must be solved, such as the evacuation of thermal energy and the necessarily higher reliability of the components, which must survive the longer and more intricate fabrication process. Even if these issues are resolved, it is also necessary to question the long-term validity of the approach. Indeed, when communication frequency increases, the communication distance decreases at constant power budget. The surface of each layer must thus decrease, which results in a quadratic increase in the number of layers and in the vertical communication distance. Thus it is not obvious that all technological problems will be solved before this technique becomes obsolete compared to bandwidth requirements.

The active interconnect layer is a variant on the general vertical integration technique, specific to the interconnect problem. This solution will not solve the problems inherent to metallic interconnect, but it offsets them for a while. The idea is to graft a layer of active buffers above the last metal layer and use these buffers (which can be arranged in a matrix formation or as a "sea" of buffers) solely as repeaters for long interconnect. This approach frees up active silicon area for computing functions rather than for interconnects. It also reduces congestion by replacing the via stacks (necessary between the last metal layer and the active silicon layer), by single vias going vertically upwards from the last metal layer to the vertically adjacent buffer layer. Power may be marginally reduced (as the via stack contribution to line load is removed). Crosstalk or any disturbance on victim lines should be reduced. As with 3D, doubts can be expressed as to how long this approach could be viable.

4.3. Optical interconnect

A promising approach to the interconnect problem is the use of an optical interconnect layer. Optical data communication technology has revolutionised long-haul communications, local networks, fibre to the home, parallel machine backplanes. The most advanced existing approach, inter-chip optical interconnect, is now beginning to emerge as a relatively mature and commercially viable technology [BER2003] [LEM2004]. For example, NEC[18] has prototyped an inter-chip optical interconnect system using optical fibre, where a switching IC (16 x 16 I/O ports) on a ball-grid array (BGA) package is surrounded by four optoelectronic conversion modules for optical signal communications with other IC. Further, Luxtera[19] recently announced a 10Gbit/s

[18] http://www.nec.com
[19] http://www.luxtera.com

CMOS photonics platform based on optical modulation and enabling practical high-speed optical fibre communication. Luxtera chips are built in a standard Freescale CMOS process typically used for mass production of microprocessors. Such technologies can enable greater connectivity for parallel computing (for example) through the use of optical I/O pads and wavelength division multiplexing.

The PICMOS project focuses on the further downsizing of optical communication. The integration of optical interconnect inside the chip could arguably empower an increase in the ratio between data rate and power dissipation. At the same time it would enable synchronous operation within the circuit and with other circuits, relax constraints on thermal dissipation and sensitivity, signal interference and distortion, and also free up routing resources for complex systems. However, this comes at a price. Firstly, high-speed ***and*** low-power interface circuits are required, design of which is not easy and has a direct influence on the overall performance of optical interconnect. Device-level difficulties expected are obtaining a large enough optical-electrical conversion efficiency, reducing the optical transmission losses while allowing for a sufficient density of photonic waveguides on the circuit and reduction of the latency while operating above the 10GHz mark.

In this context, integrated microphotonics is essential to the emergence of systems on chips including optical functions. The objectives of this domain include the creation or reproduction, at the submicron scale, of passive (guides, filters, switches) and active (sources, detectors, wavelength shifters) optical functions by integrating them directly on a chip. Improvements here benefit fully from advances in microelectronics technology, including scaling.

Another important constraint on this integration, mentioned previously, is the fact that all fabrication steps have to be compatible with future IC technology and also that the additional cost incurred remains affordable. Additionally, predictive design technology is required to quantify the performance gain of optical interconnect solutions, where information is scant and disparate concerning not only the optical technology, but also the CMOS technologies for which optics could be used (post-45nm node). In a similar spirit, design technology for optical interconnect has to be developed retaining as far as possible compatibility with state of the art EDA tools and methodologies, in order to bring optical interconnect to the IC design community.

4.4. Comparison

A rough comparison between the various alternative interconnect strategies is now presented. The following sections detail the advantages and disadvantages of the optical solution in particular (section 5), as well as potential applications (section 6).

	RF interconnect	3D interconnect	Optical interconnect
Bandwidth	High in channel Limited by transceivers	Better than planar interconnect	High in waveguide Limited by interface circuits
Latency	Could be limited by shared transmission medium	Good	Limited by interface circuits and optoelectronic devices
Timing behaviour	Requires clock data recovery circuit	Reduced skew	Reduced skew
Interconnect density	High (limited by transceivers)	High	Unknown (limited by interface circuits)
Power	High	Good	Good
Signal crosstalk and isolation	Poor	Poor	Good
Physical security	Poor	Poor	Good
Process modifications	No	Yes	Yes
Analogue signal transmission	Yes	Yes	Difficult
Design practices	Minor change	Minor change	Major change

Table 2 Comparison between interconnect alternatives

5. A priori advantages and disadvantages of optical interconnect

5.1. Advantages of optical interconnect

The expected advantages of the use of optical interconnect, whatever the application (data transport or clock distribution) are:

High speed

Passive optical waveguides and fibres enjoy a flat frequency response up to several THz[20]. Further, the frequency characteristics of waveguides and fibres are independent of their length. This makes it possible to imagine architectures with physically large synchronous zones and a high number of long, high-speed connections. However, speed is in practice limited by the emission and detection devices and associated circuitry. The laser source must firstly be biased above a threshold current in order to avoid turn-on/off

[20] Above this level, the dependency of group delay on frequency is no longer negligible and has a serious impact on the various frequency components of the signal. This causes unacceptable distortion levels. For a monomodal Si/SiO_2 waveguide for example, the relative effective index variation at 200THz is around 0.1%/THz, which induces a relative variation in group delay of around 0.15%/THz [LET2005].

delays[21]. In a traditional laser, this time quantity is approximated by

$$t_{turnon} = t_N \left(\frac{I_2 - I_1}{I_{2-I_{th}}} \right)$$

when switching from I_1 below the threshold current I_{th} to I_2 above I_{th} (τ_N is the carrier lifetime). Using typical values (τ_N=3ns, I_1=0mA, I_2=3I_{th}), t_{turnon}=1.2ns is clearly too high to be compatible with operation in the GHz range. The switching time is reduced when the laser is modulated between two current values above the threshold current. Hereby the laser obeys the small-signal transfer function

$$H(s) = \frac{1}{1 + \frac{sg}{w_R^2} + \frac{s^2}{w_R^2}} \frac{1}{1 + st}$$

whereby the first factor accounts for the intrinsic modulation response of the laser (with relaxation oscillation frequency ω_R and damping coefficient γ) and the second includes device parasitics (RC time constant τ). For typical values (ω_R=10GHz, γ=15GHz, τ=25ps), t_{switch}=0.3ns to reach 90% of the step to steady state.

The constraints on the operating frequency can thus be achieved, with regard to the source, with the provision of using a minimal level of current for its biasing. However, it is in fact the interface circuits that most severely limit the frequency response of the optical link as a whole. Indeed, as mentioned previously, the maximum operating frequency of these circuits does not exceed as a general rule f_{max}/10, where f_{max} is limited by the technology of the transistors to a value typically (for a CMOS process) a hundred times lower than the transmissible frequencies in the optical field. This leads to the remark that the use of III-V technologies, not only for sources and detectors, but also for interface circuits, would make it possible to reduce the difference between the operating frequencies of the circuits and of the optical devices.

Deterministic timing behaviour

A related advantage of optical interconnect relates to the temporal aspect of signal transmission. The time of flight of photons in an optical link is roughly the same as that of electrons in a purely metallic interconnection. The advantage of optical interconnect appears when one considers the influence of the repeaters necessary to long metallic interconnections. Repeater circuits introduce spatially varying behaviour into signal propagation because of the dispersion of technological and electrical parameters in the system. Since there is correlation between parameter variation and the layout co-ordinates of each inverter, and these co-ordinates follow a fairly regular pattern (in an H-tree for example), the characteristics of a complete branch (from generator to gate) constitute the accumulation of repeater performance variations from nominal behaviour. Since one of these performance characteristics is delay, this leads to different total propagation times between branches. This behaviour, inherently undeterministic before fabrication,

[21] This assumes the use of a source-detector type optical interconnect. However, modulator-based approaches may achieve higher performance in this respect.

constitutes one of the origins of "skew" between two signals being propagated in two lines of the same length and which must arrive at the same moment. In optical interconnect, the only elements having an influence on the propagation of the signal are the interface circuits. One can thus consider that the temporal behaviour of the signals in optical interconnect is very largely deterministic, which implies that the "skew" is significantly reduced and makes it possible to affirm that a synchronising clock signal transmitted optically has an extremely high temporal precision. However, this will be offset by the dependency of detector delay on the optical power, and by the dependency of the wavefront speed to the temperature-dependent waveguide optical index. To minimise the former, it will be necessary to tightly control variation in all sources of loss in the optical link (coupling, distance). For the latter, it will be necessary to evaluate thermal variations across the chip.

Increased interconnect density
Exploitation of WDM[22] techniques also makes it possible to increase the aggregate data rate of the optical link. Multiplexing several signals carried by different wavelengths is possible, and for n wavelengths one can consider that the waveguide represents an optical bus of n bits. If D represents the data rate of the link for a given signal, the aggregate data rate of the link will be equal to nD. Moreover, even if the geometry and the spacing of the waveguides are fixed by monomodal propagation conditions and the minimisation of optical coupling, and do not vary with the scaling factor (the *pitch*, p, between two waveguides is about 1µm), interconnect density can still be increased. Indeed, the equivalent interconnect pitch in the case of an optical bus is p/n. In the same way, routing complexity will be reduced, since an additional layer of interconnect adds degrees of freedom (this is of course true for any additional routing layer, whether the medium be optical waveguides, RF channels or metal wires).

Low power
Because of the elimination of the power-hungry repeaters present in metallic interconnections, optical interconnect can contribute to the reduction of power consumption, as well as that of the active area used for data transmission. However, even though part of the optical link is passive and should not (in principle) require power-costly signal regeneration, the complete optical link is actually active. The interface circuits are again of prime importance here. The power budget in an optical link is actually based on interface circuit power dissipation, which is determined

- at the receiver end by the receiver circuit[23]: this is a circuit design task aiming to minimise power dissipation and noise figure for a given frequency response;
- at the emitter end by the modulator circuit: this is determined firstly by the current required to bias the source above its threshold current; and secondly by the current modulation required to deliver sufficient light power to the receiver for a specified bit

[22] Wavelength Division Multiplexing
[23] Advanced research in D. Miller's group at Stanford University, USA, have demonstrated the concept of receiverless optical interconnect based on short communication pulses to pull-up or pull-down detector devices. At present this approach does not seem practical since it imposes stringent constraints on the matched high-speed performance of both pull-up and pull-down channels.

error rate. The latter value has to compensate for light power losses induced by coupling and waveguide attenuation.

This *a priori* advantage is thus not so clear and has to be evaluated. Conclusions from [CHO2004] show that *inter*-chip optical interconnect becomes more power favourable at higher bit rates, lower BER and longer transmission distance (43*cm* threshold in the case under consideration).

Transmitted signal isolation

By using optics to transport high-frequency digital signals, it is clear that crosstalk will be reduced. The contamination of "clean" signals transmitted over local metallic interconnect (either analogue or digital) will be reduced and probably eliminated concerning the digital signal in question since the only coupling point with local electrical transport is at emission or reception. Transport of such signals via optical interconnect will generally reduce the ambient noise in the substrate.

However, optical coupling exists between waveguides. The optical power coupling coefficient, α, from a given waveguide (1) of cross-sectional area S_1, to an adjacent waveguide (2) of cross-sectional area S_2 is given by [HAU1984]:

$$a = \left|\frac{k_{21}}{b_0}\right|^2 \sin^2 b_0 l$$

where:

$$b_0 = \sqrt{\left(\frac{b_1 - b_2}{2}\right) + \left|k_{12}\right|^2}$$

$$k_{21} = -\frac{jw}{4}\int_{s_2} da(e_i - e)e_1 e_2^*$$

$$k_{12} = -\frac{jw}{4}\int_{s_1} da(e_i - e)e_2 e_1^*$$

α is thus dependent on the distance for which the waveguides are adjacent (l), the distance between the waveguides, the optical indices of the transmission medium ε_i and cladding ε, the wavelength of the signals. It is important to note that in theory conditions can be found under which 100% of the power is coupled into the adjacent waveguide.

Physically secure transmission

Secure data transmission is an increasingly critical issue: firstly for the transmission of personal data (banking, medical...) and professional data (sensitive information) between data-processing machines; and secondly for the transfer of protected algorithms from a memory towards a processor at the chip scale. In both cases, the metal pads of the chips, as well as the last level of metallic interconnections, are potentially easily accessible for setting up a tap into significant information. Optics is generally well adapted to secure transmission. It is physically more difficult (but not impossible) to tap a flow of photons, at a sufficiently weak level to avoid detection and at a sufficiently high level to be able to exploit the information. Thus to protect sensitive inter- and intra-chip transmissions it is possible to imagine optical input/output pads which avoid any electrical conversion, and

an upper, all-optical interconnect layer to improve the physical security of the transmission links. It is of course debatable if it really needs to be integrated but for some very high security links possibly even the last bonding links are an issue.

5.2. Limiting factors for optical interconnect

Process modifications (heterogeneous integration)
Materials have to be chosen with different constraints:

- efficient light detection: the quantum efficiency of the active devices is of fundamental importance to the global power budget of the link, one of the main comparison criteria between optical and metallic interconnect. Also, particular attention has to be paid to the receiver: the signal to noise ratio determines the minimal optical power at the detector
- efficient signal transport: attenuation and compactness are the main parameters for the choice of the passive waveguide.
- technological compatibility with standard CMOS processes: an industrial solution is conceivable only if the optical process is completely separated from the CMOS process (the development cost of a new CMOS technology is so high that it seems very difficult to propose a solution which would require a fundamental rethink of IC fabrication processes)

Different materials are available for the realisation of the optical passive guides. As an example, Si/SiO_2 structures are compatible with conventional silicon technology and silicon is an excellent material for transmitting light of wavelengths above 1.2µm (mono-mode waveguiding with attenuation as low as 0.8 dB/cm has been demonstrated [LEE2001][24]). The use of silicon waveguides makes it possible to imagine either monolithic (planar) or hybrid (3D) integration of the optical subsystem with CMOS systems. It is believed that the former solution is not realistic.

The integration of silicon waveguides at the front end of the CMOS process (i.e. before fabrication of the metallic interconnection layers) is certainly possible but other considerations have to be taken into account. At the transistor level, the routing of the waveguide is extremely difficult and requires routing space at the IC level. This approach will not therefore contribute to reducing routing congestion problems. Further, the problem of the active devices remains: silicon-based sources cannot yet (and for the foreseeable future) be considered to be mature, while the growth of III-V devices on silicon faces strong technological barriers[25]. The use of external sources and detectors bonded by flip-chip is unrealistic due to the high number of individual bonding

[24] However, the fabrication technology that was used in this communication is not compatible with high integration density. Oxide annealing over several microns was used to fabricate waveguides of submicron width. This process means that waveguide pitch is several microns, such that the overhead is too high to be considered for high density interconnect.

[25] The temperatures involved are generally too high (800°C). However, recent research [SEA2001] has successfully demonstrated hydrophilic bonding of an InP microdisk laser onto silicon at room temperature, through the use of two SiO_2 layers on both InP and silicon substrates. The highest temperature in this process was 200°C, compatible with CMOS IC fabrication steps.

operations required. Also, silicon-based devices can only work at low wavelengths (850nm), which translates to higher attenuation in the waveguides. This solution requires an extraordinary mutation in the CMOS process, and as such is highly unattractive from an economic point of view.

Hybrid integration of the optical layer on top of a complete CMOS IC, the approach adopted in the PICMOS project, is much more practical, and with more scope for evolving. The source and detector devices are not bound to be realised in the host material. For particular applications, the photonic source can be on- or off-chip: it seems likely that for some near-term applications, such as clock distribution, it is better to target off-chip signal generation for thermal reasons, even if it means higher assembly costs. It should be noted that this solution also applies to MCM[26] technology. The optical process is completely independent from the CMOS process, which is appealing from an industrial point of view. Disadvantages of this approach include the more complex electrical link between the CMOS subcircuits (source drivers and detector amplifiers) and inevitably more advanced technological solutions for bonding. The use of III-V technologies to realise the interface circuits, indicated as a potential alternative above, would further compound this problem.

Cannot transmit analogue signals directly
Direct analogue signal transmission, which is particularly sensitive to noise in the surrounding environment, would benefit from the signal isolation characteristics of optical interconnect. However, a fundamental problem with analogue signal transmission over optical links is that linear transmission, required for analogue signals, cannot be realised directly. This is due to the combination of logarithmic source emission characteristics and linear detection characteristics, and is generally solved in two ways: at the source, by pre-distortion techniques to compensate for the logarithmic characteristic; and at the receiver, by using logarithmic detectors. The former solution is difficult and costly, while the second is generally limited to low frequency applications. For integrated optical interconnect, where transmission distance and therefore attenuation are not likely to be tightly controlled, it would also be necessary to add a variable gain amplifier to the receiver chain in order to compensate for this distance-dependent loss in the optical link. This means in any case that the interconnect structure has to be adapted to the type of signal (analogue or digital) in a much more fundamental way than is the case with metallic interconnect.

Changes to design practices
The very use of optical interconnect, traditionally considered as a specialised technique, raises design questions that industrial IC designers are not equipped to answer. For example:

- for a given (long) transmission line and target performance (in terms of power, delay/latency, area etc.), is it advantageous to use optical interconnect instead of metallic interconnect or not? A "figure of merit" for both metallic and optical

[26] Multi-Chip Module

interconnect technology, such as that mentioned in section 3, would enable this in a first approach.

- even before having identified long transmission lines, should the interface circuits and clean power lines necessary to supply them be integrated into floorplanning tasks?
- should the interface circuits be standard cells or automatically optimised based on requirements?

Most of these questions can be resolved with the support of adequate CAD software, taking into account *a priori* interconnect evaluation, specification-driven interconnect design, waveguide routing and interface placement. This is a main target of WP1 in the PICMOS project.

6. Applications of optical interconnect

For the optical interconnect concept to be seriously considered as a viable alternative to metallic interconnect, performance gains of at least an order of magnitude must be shown by advances in research at circuit and device level, as well as by definition of potential applications.

For the purposes of this work, we have categorised these applications into three broad domains, for which various analyses can be carried out concerning the suitability of optics to each domain. These link types are: data transport (point-to-point and N-N – bus, networks); clock distribution (1-n link); and physical interfacing (input/output, control), where a particular property of optics, not accessible to electronic solutions, is exploited.

6.1. Data transport

The volume of data being transferred between on-chip functional blocks is rising fast due to (i) the need for higher resolution (sound, image (video) and computing) and (ii) the need for greater CPU power or total MIPS (real-time encoding/decoding, data encryption/decryption). This is engendering a move to distributed multi-processor architectures and therefore requires organised high-speed communication between processors. For all the reasons mentioned previously, metallic interconnect will be highly inefficient in this role. Optical interconnect as developed in the PICMOS project could provide a technological solution to this problem, but will have to compete against integrated radio frequency interconnect. Future systems requiring such high interconnect bandwidth will include high-performance human-machine interfaces as a driving force, such as gesture interpretation, speech recognition and synthesis for voice-controlled applications, facial/retinal/digital recognition for security applications.

In the following paragraphs, we describe the potential of optical interconnect in communication between IP[27] blocks in SoC (optical buses, NoC[28]); and in configurable hardware. Optical solutions can be considered in terms of the nature of transmission: point-to-point, bus, networks. The metrics relative to data transport include the temporal

[27] Intellectual property. This term is used in the semiconductor industry to generically describe a largely autonomous on-chip block composed of mainly digital circuitry running software content, communicating with other blocks via a standard interface.
[28] Network on Chip

characteristics (pure latency, skew, jitter, delay), power consumption, interconnect density.

The basic idea behind using point-to-point optical links consists of replacing electrical global links with optical ones[29]. Typically on today's more complex chips, hundreds or thousands of global links are necessary [DAV1998] (internal or local interconnect, i.e. connections within the IP block itself, are usually short but can carry signals of high frequency). Indeed, the ever-increasing number of transistors in a SoC has led to such a high level of complexity that IP re-use is mandatory. SoC architectures are thus evolving to structures integrating some hundreds of IP blocks, which communicate between each other through a common interconnect structure. The organisation of this data transfer is one of the main tasks in SoC design today.

In current SoC architectures, global data throughput between functional blocks can reach up to tens of gigabits per second, the load being shared by several communication buses. In the future the constraints acting on such data exchange networks will continue to increase: the number of IP blocks in an integrated system could be as high as several hundred and the global throughput could reach the Tbit/s scale. To provide this level of performance, the communication system itself is designed as an IP block into which the various functional units will be connected. This type of standardised hardware communication architecture is called a network on a chip (NoC). Every functional block on the chip (DSP, analogue/RF, video processors, memory etc.) can in this way interface to an interconnect network architecture for data communication.

The architecture of the interconnect network can go from a simple shared bus (easy to realise, sufficient for most applications, such as Amba) to a totally interconnected network. The latter can offer high performance, but is in practice limited to communication between a low number of IP blocks (typically 8) due to exponentially rising complexity and power requirements). Between these two extremes there are several "compromise/trade-off" network architectures such as crossbars, hierarchical or multi-level, ring networks. For high data-rate applications (processor-memory communication, multiprocessor and parallel computing), this compromise is relatively difficult to find [COL2003].

Research has been carried out on analysing the potential benefits of introducing optical interconnect in critical data-intensive links, such as CPU-memory buses in processor architectures [COL2000]. At the time, these analyses showed that point to point links do not present a sufficiently high performance gain to warrant their widespread use in future technologies. In essence, the bandwidth/power ratio for point-to-point optical links is higher than the electrical counterpart, but not high enough, when interface circuit power is taken into consideration. Instead, it is preferable to apply architectural modifications in order to enable bottlenecks to be overcome (in the given example application, the

[29] Inside an IP block, interconnects are relatively short and will not be concerned by the optical solution. The increase in frequency will lead on the other hand to a reduction of the physical size of the IP block, which should remain of constant complexity. If a global link is defined as an inter-IP block link, then this implies that the number of global links will increase appreciably with the technology node.

solution was to add more cache memory), even at the expense of greater silicon area and power. However, considering the technology nodes at which these analyses were carried out, one can note that the number of repeaters only accounted for approximately 5% of the total system area. At future technology nodes, the empirical Rent's Rule[30] indicates that this percentage increases considerably with technology node. Indeed, the most pessimistic scenario estimates that more than 25% of the active area of the system will soon be occupied by repeaters necessary only for data transmission (Figure 8 [BAN2001]). Thus the conclusion obtained a few years ago is perhaps not final.

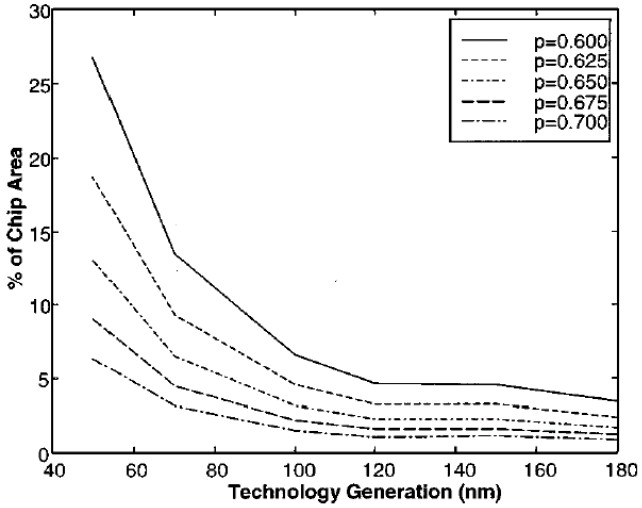

Figure 8 Silicon area percentage required by interconnect repeaters at different technology nodes

In any case, it seems more relevant today to consider an optical link in all its possible configurations and not simply for the point-to-point application. Thus it is necessary to consider the possible contribution of multiplexing techniques in time *and* in wavelength. Actually, wavelength could be exploited for routing functions. This approach enables us to consider all optical data transport links as N-N links of data rate D (which can be shared between the N sources), possibly with reconfigurable targets, and where the value of D will be determined by:

$$D = n_f n_l n_l D_e$$

where n_f represents the ratio between the optical transmission frequency and the data arrival rate, n_l the ratio between the number of physical optical and metallic lines, n_λ the number of wavelengths available in the system, and D_e the data arrival rate. The applications of these links would thus include point-to-point links (N=1), optical buses (N = 1) using TDM[31] ($n_f > 1$) and/or WDM ($n_l > 1$), and reconfigurable networks (N > 1, n_λ or $n_f > 1$). As in telecommunications, WDM provides a route to very high data rates, even if the individual devices cannot be modulated much faster than electrical bus data rates. A

[30] Rent's Rule expresses the number of inputs and outputs, T, of a digital block according to the number of gates in the block, N. It is written: $T=kN^p$, where the empirical constants k and p represent, respectively, the average fan-out of the gates and the level of complexity of the interconnections.
[31] Time Division Multiplexing

single waveguide could be used to replace a 64-bit bus, for example, where each individual signal makes use of a distinct wavelength. However, it should be noted that data resynchronisation would be necessary since the time of flight of signals would vary due to the use of different wavelength values.

6.1.1. Bus architectures

Commercial solutions for SoC development platforms are presently based on bus architectures [IBM1999] [VSI2000]. Architectures using metallic interconnect rely heavily on wide (64/128-bits) buses, as well as on the frequent use of switch boxes to dynamically define a communication route between two functional IP blocks. In bus-based architectures, a central bus arbiter ensures the definition of these routes. Once the arbiter has guaranteed the control of the bus to an IP block, this block has all the physical bus bandwidth and only propagation delays intervene in the data transmission. The bus bandwidth is thus shared in time by all IP blocks connected to the bus. Again, since the order of magnitude of the distance of communication is the chip die size, systematic use of repeaters (over the buses or within the switch boxes themselves) is necessary and increases power consumption.

For example, the AMBA bus is a *de facto* interconnection standard developed by ARM. It is based on the definition of a data transport bus of high (AHB) or average (ASB) performance for processor-type units, and of an input/output bus (APB) for peripheral-type units:

- ASB (Advanced System Bus): multi-master bus allowing pipeline operation with a bus D of 8/16/32 bits. The bus transfer cycle (read or write operations) lasts several clock cycles.
- AHB (Advanced High-Performance Bus): identical to the ASB with the possibility of operation in split mode and of authorising transfers in burst mode. Bus D can also be of up to 128-bit width.
- APB (Advanced Peripheral Bus): simple low power (and low data rate) interface bus, with memorisation of addresses and control. The bus transfer cycle always lasts two clock cycles.

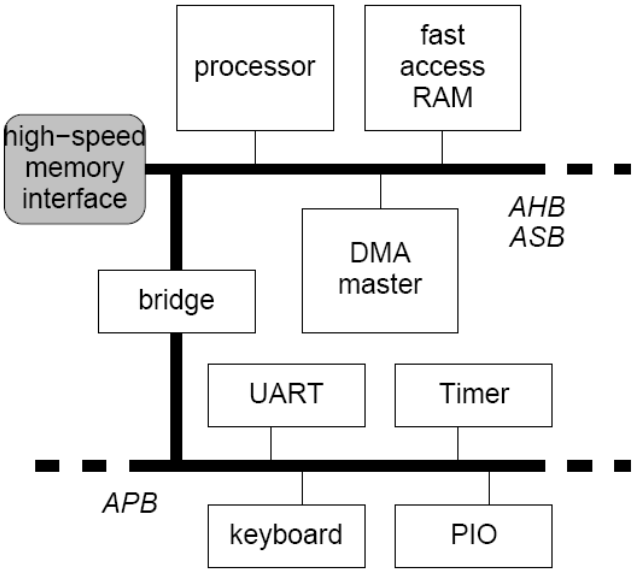

Figure 9 Example of interconnection by AMBA bus

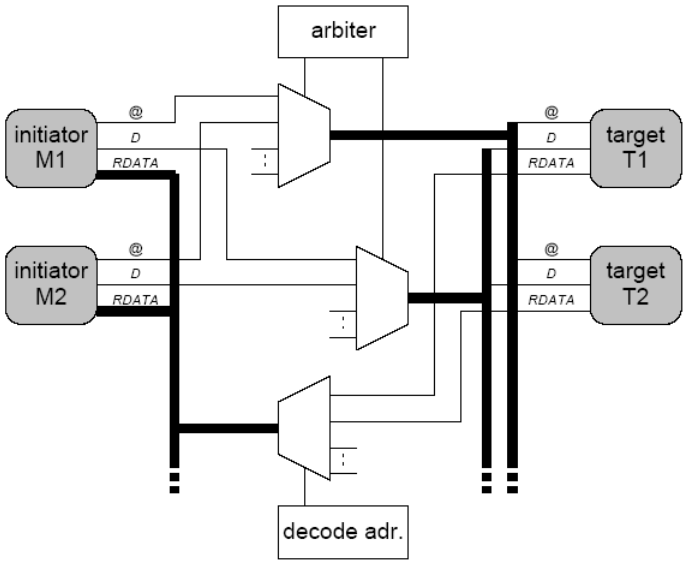

Figure 10 Arbiter and addressing with AMBA

The three buses are schematised in Figure 9, and are synchronous with the same clock. A master unit (typically a processor) can send a bus access request to the centralised bus arbiter, and thus request a slave unit (typically a memory block). A single and centralised address decoder carries out unit addressing (Figure 10). It is possible to find the optical equivalent of such a bus (Figure 11) by replacing the metal interconnections by high data-rate optical links (single or multi-wavelength). Thus the total data rate of the bus, shared by all the IP blocks over time, could be appreciably increased. This would make it possible to meet the needs for scaling by allowing the communication between a higher number of IP blocks over the same bus. However, the fundamental limit of this family of

structures lies in the arbitration, which cannot be done in the optical domain. Actually, this function introduces latency at each communication request, and it would increase with the number of IP blocks sharing the bus. Thus an optical bus with centralised arbitration will not solve all the problems associated with the wire data buses.

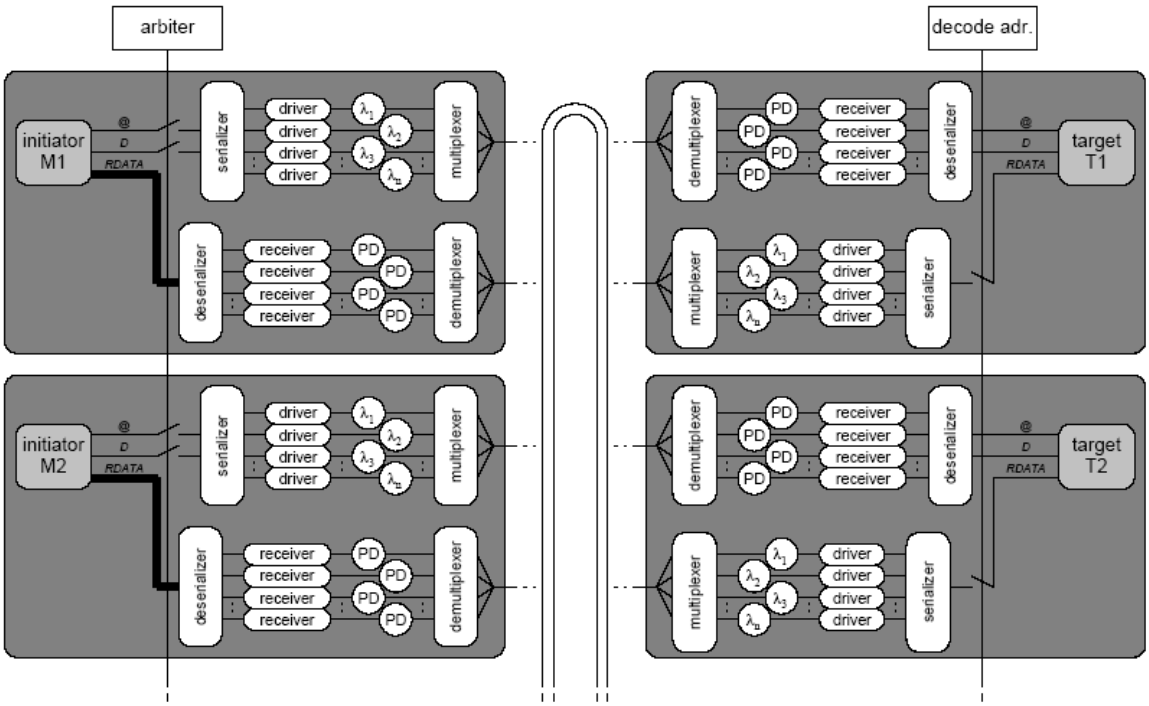

Figure 11 Possible optical bus structure using TDM and WDM

6.1.2. Networks on chip

In the future, the limitations of bus-based architectures (latency due to arbiting and line delay, non-scalability, time-sharing, non-reconfigurability, etc.) will render them obsolete. Future architectures of integrated systems will require new concepts for on-chip data exchange. The ever increasing number of transistors in a chip will lead to such complexity that IP reuse will be mandatory: a system is designed by integrating some hundreds of pre-designed complex functional blocks, the designer concentrating mainly on the organisation of data transfer between these blocks.

A number of innovative interconnect architectures (networks on chip - NoC), have been recently proposed [BEN2002] [GUE2000] [DAL2001] to overcome the limitations of commercial bus-based platforms by providing each IP block, interfaced towards the network, with one or more reconfigurable channels of high-speed communication, represented by a switch box (Figure 12). NoC architectures look much more like switching telecommunication networks than conventional bus-based architectures. The principal difference between networks on chip and buses lies in the presence of routing nodes in the networks, operating in an asynchronous way, able to buffer the data and to decide the best routing itinerary to use for the data transfer. Thus, the inherent

"intelligence" of the elements within the network makes it possible to eliminate the bus arbiter and therefore the associated latency. In counterpart, the latency of the network itself is higher in networks on chip than in buses.

Depending on the target application (multiprocessor SoCs or systems in which different functional blocks process heterogeneous signals), NoCs may have different structures, such as rings, meshes, hypercubes or random networks. Indeed, MPSoC applications (where the layout is regular and the data-rates are high and fairly constant) can benefit from a fixed topology whereas SoCs (where the layout is very irregular and the data rates very variable), will require a topology "à la carte". Latency, connectivity, global throughput and reconfigurability constitute the main performance indicators of these networks.

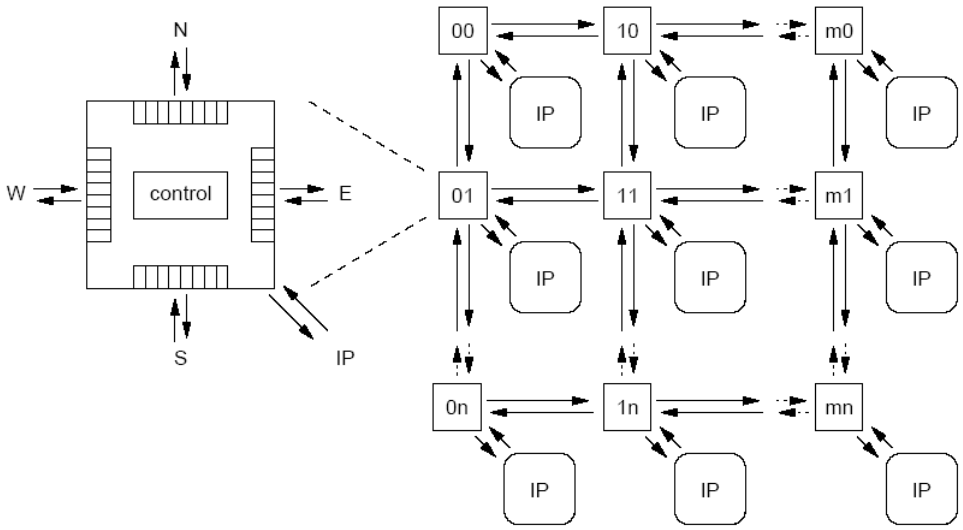

Figure 12 General NoC architecture showing IP blocks and interconnected switchboxes

With respect to the synchronisation of operation of the processor cells, at the system level the natural evolution for processor architectures is that each cell, integrating network interface logic to communicate with other cells on- and off-chip, operates on a locally synchronous basis, but the whole is globally asynchronous. The structure then operates as an integrated local network and communication between cells by data packets becomes advantageous with respect to dedicated global wiring in terms of structure, performance and modularity.

Integrated optics may constitute an effective and attractive alternative for NoC. The superiority of optical interconnects in networks is established, and possibly, some advantages of optical propagation may overcome the limitations of classical technologies for data exchanges at the integrated system scale. Some of the physical advantages of optical interconnects may be of prime interest for NoC: flat frequency response (i.e. signal attenuation does not depend on frequency), limitation of crosstalk, no repeaters and power consumption. By using an optical switching approach (by modifying the effective

material index by some electrical phenomenon, by modifying the coupling between guides by optical resonance, or by modifying the coupling angle in a mechanical way[32]), reconfigurable networks could be realised in the optical domain (for FPGA applications for example). The main advantages of such networks would be power reduction and higher integration density. In particular, switch boxes, a key element for reconfigurable networks, could also be realised by using compact micro-resonators (about 10 x 10 μm^2); capable of selecting and redirecting a signal based on its wavelength. Such networks would be entirely passive; i.e. no power would be required to transport the data, whatever the communication route necessary (but of course power would be required to convert the data to and from the optical domain). However, such a scheme would imply a shift in the routing paradigm from a centralised arbiter acting on the switch boxes, to one acting on the block interfaces to select the wavelength(s) to be used. Also, tuneable and thermally stable microlasers would be required.

NoC architectures are currently based on 2D mesh topologies and use in general packet switching with intelligent routing techniques such as wormhole routing. It is not clear that a direct replacement of electrical links between switchboxes by optical interconnect will achieve a significant performance gain. Such an approach would require conversion between optical and electrical domains ***at each switchbox***. This would lead to an explosion in the number of active devices and interface circuits required and thus be rather inefficient in terms of power and delay. Also, it would not deal satisfactorily with the inherent latency issue in NoCs, introduced by the concept of distributed routing nodes and packet switching. However, the main advantage of this approach would be to retain compatibility with distributed NoC architectures and routing paradigms. To mitigate the high penalties associated with optoelectronic conversion, one could envisage conversion of destination addresses at each switchbox, or transmission of the destination address over a physically separate (electrical or optical) medium (Figure 13), rather than conversion of address ***and*** data information at each switchbox. The fundamental difference between this approach and that of existing approaches is that buffering (temporary storage of information at a switchbox node) is not possible in the optical domain. With a wormhole routing strategy for example, a phase to set up the routing itinerary, where data transmission is halted until the itinerary is established, would be required.

[32] Lucent (http:///www.lucent.com) recently unveiled an 8x8 optical switchbox (WaveStar™ LambdaRouter) to switch lightwaves among fibers in fiber-optic networks based on a mechanical switching technique. It is composed of 8 rows of 8 micro-mirrors tilted by electrical charges.

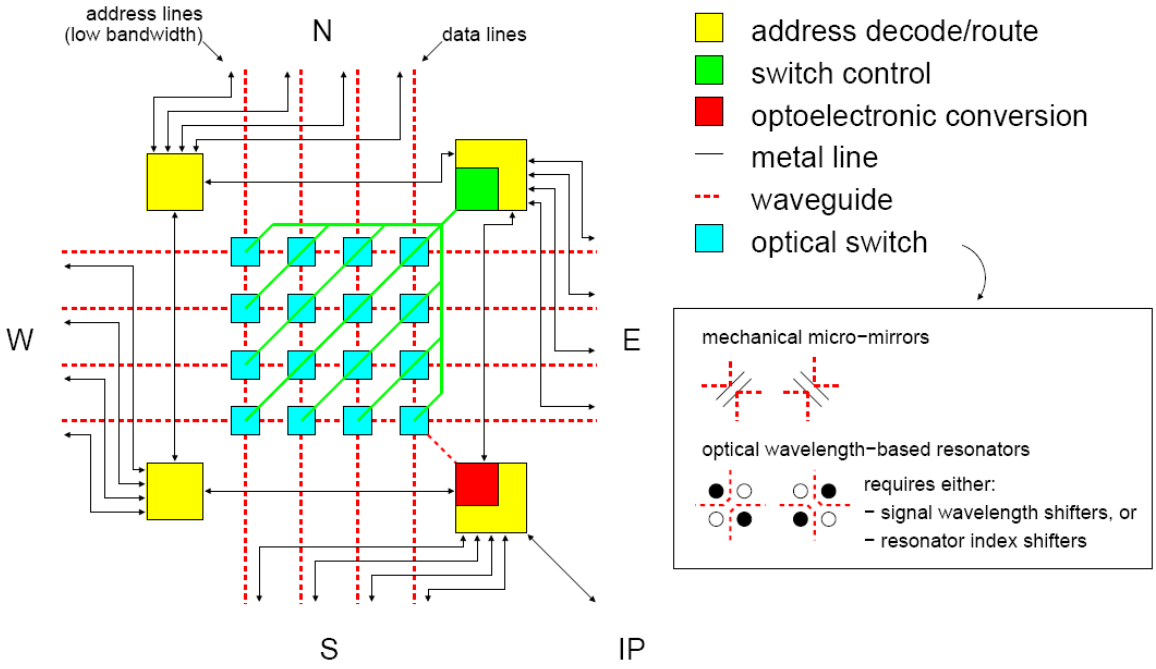

Figure 13 Switchbox architecture exploiting partial optoelectronic conversion

Such an approach would not exploit to the full all the advantages of routing in the optical domain. In particular, wavelength division multiplexing (WDM) may offer new and appealing solutions such as optical buses and reconfigurable networks. The concept of a scalable, completely passive network based on microresonators was demonstrated in [BRI2004], where the address is coded in the transmission wavelength. This must make it possible to increase the rate of connectivity while reducing the active device count, and consequently the power consumption and area used. However, this would imply that the switching technique would no longer be packet-based, but circuit-based. Inevitably this would cause more fundamental changes in design techniques than the previous approach.

6.1.3. Reconfigurable hardware

The level of granularity involved in the considered applications is an important parameter to take into account. Switchboxes in NoCs route between blocks of very coarse granularity (IP blocks of several hundred Kgates), while at the other end of the spectrum, switchboxes in FPGAs route between blocks of very fine granularity (configurable logic blocks of complexity ranging from a few tens of gates to Kgates). For FPGAs, interconnect represents a huge proportion of the total resources available. As an example, for an FPGA tile consisting of active logic (LUT, multiplexers and flip-flops), programming memory and interconnect, the active logic represents at best 20Kgates, the memory 80Kgates and the interconnect 700Kgates. It is clear that such a high proportion causes typical interconnect-related problems (power, delay, and area).

The same type of structure as that mentioned previously (switchbox exploiting partial optoelectronic conversion) could be used in FPGAs. Possibly this would be a better application since reconfiguration of the network does not have such stringent time

constraints as in NoCs. The latency involved with mechanical switching for instance would be less of a problem in FPGAs than in NoCs since configuration of the FPGA is carried out before putting the programmed function into operation and there are no particularly tight demands on configuration time. This is not so true for dynamically reconfigurable FPGAs, but still acceptable switching times can be considered to be far higher than in NoCs.

It should be stressed however that two fundamental difficulties remain with the use of optics in FPGAs. Firstly, it would certainly not be efficient to replace an electrical switchbox by an optical switchbox to realise a configurable link between blocks that are spatially close. Performance gain is likely to be achieved only be for "long" hops between configurable blocks. Thus, as in the previous applications, the choice between which communication domain to use to connect structures would be based on the distance between them.

The second difficulty is that not all electrical switchbox configurations could be easily realised in a single device in the optical domain. Routing could be realised by microresonators or micromirrors, but other configurations require splitting. Total connectivity for example (an incoming signal at one face is broadcast to the three other faces) would require a 1-3 splitter, thus increasing the complexity of the optical switchbox.

As no optical switching functions are being developed in the PICMOS project, it is impossible to consider any of the applications requiring switchboxes (NoC and FPGA) as a suitable test vehicle for the project. It is more realistic to consider long, fixed-function data links as a first application. If several are considered in parallel, this would be close to achieving an optical bus (without WDM) and will enable the evaluation of any interaction between the individual interconnect lines. This is important information for several applications and such a test vehicle would be a solid building block contributing towards the more advanced data transport applications requiring additional technology developments.

For these reasons it is recommended that a parallel long-haul optical data link be considered as a primary candidate for the application test vehicle.

6.2. Clock distribution

Another group of potential applications, where optical interconnect technology could profitably replace that of metallic interconnect, is where switching activity is high and where the number of receivers (and therefore repeaters) is also high. Both characteristics are present in clock distribution networks (CDN) [FRI2001].

In order to operate at high frequencies, CDNs require several thousands of repeaters to switch charges on the metallic tracks over the entire chip. This results in the use of a high portion of overall IC power (up to 40-50%). This mode of operation also leads to stringent constraints on the design of the clock tree, since an unbalanced tree will result in serious clock skew and consequently system failure. An electrical alternative is global

clock distribution at a relatively low frequency and local clock multiplication to generate the required clock speed. Disadvantages of this approach include inter-zone synchronisation and clock multiplication lock time.

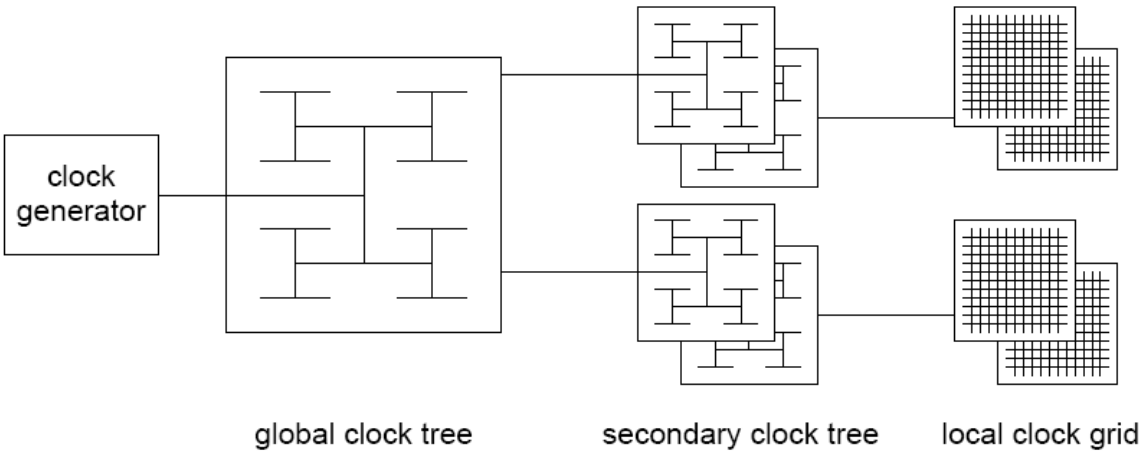

Figure 14 Clock tree hierarchy

By replacing the electrical clock distribution tree by an optical one, the need for repeaters or clock multiplier circuits would be eliminated, thus reducing power consumption and clock skew. Since controlling the clock signal is critical to the operation of high performance systems, a great deal of engineering goes into its proper generation and distribution. Techniques for synchronous VLSI often utilise a buffered tree and grid-like structure with several levels of hierarchy as shown in Figure 14. It represents the most efficient solution in terms of skew covering the symmetric topology on the upper metal levels, and is required to cope with the inherently large fanout associated with the clock signal. The clock distribution topology is generally partitioned into global, secondary and local clock networks. The global clock network, which distributes the clock signal across the chip, is usually placed on the upper metal layer and utilises a tree-like structure. Such structures are fully symmetrical (H-tree or X-tree), which leads theoretically to elimination of clock skew due to unequal clock path lengths. However, strict adherence to the H-tree structure requires a highly structured design and hence is directly applicable only to highly structured repetitive circuits. The optional secondary clock trees connect global and local clocks and are placed on intermediate metal layers. The local clocks provide the clock signal to the clocking registers. They are usually placed on the first metal layers and utilise grid-like structures. A grid provides a regular structure so the clock signal may be easily distributed near to every location where it may be needed. The granularity of the grid is directly connected to the distribution of the clock loads. Additionally, the clock grid ensures a universal availability of the clock signal, good process-variation tolerance (thus reducing risk of local skew) and an easy design process. Disadvantages with this approach are that the grid wires increase the total capacitive load of the clock network, which in turn increases the power consumption. Additionally the clock grid is an area-inefficient solution.

As an example, the Itanium IA-64 (Intel 64-bit microprocessor) core contains 25.4 million transistors placed on a 464mm^2 chip in a 0.18µm, six-metal layer CMOS process, and operates at 800MHz. The architecture of the IA-64 clock distribution network follows the three distribution (global, regional and local) structure previously described. An on-chip phase-locked loop is used to generate the clock signal. The H-tree global distribution network, which occupies the two highest-level metal layers, distributes the clock signal to eight deskew buffers (DSK). In order to minimise the effects of any capacitive and inductive coupling between the clock lines and adjacent signal lines, each wire of the global clock is fully shielded laterally by power and ground lines. The regional clock distribution encompasses the DSK, the regional clock driver (RCD) and the regional clock grid. The regional clock distributes the clock signal to 30 separate clock regions (regional clock grids). The regional clock grid is implemented using 4.1% of metal 4 (x-direction) and 3.5% of metal 5 (y-direction). Local clock distribution constitutes the final segment of the clock distribution. It consists of local clock buffers (LCBs) taking the input directly from the regional clock grid and the opportunity-time-borrowing (OTB) delay clock generation. The overall power dissipation for the chip is 130W and the measured global skew is around 28ps.

According to the ITRS, the first CMOS integrated circuits operating at a clock rate above 10GHz should appear by 2008. At this frequency, metallic interconnects will require a large proportion of circuit power and area, leading in turn to thermal problems. Again, the technology developed by the PICMOS project could be used in this context, but architectural (globally asynchronous locally synchronous) and 3D integration solutions are the competitors here.

It is however to be noted that the development of globally asynchronous architectures will be limited by the capability of manufacturers to migrate existing applications to this new technology. The majority of existing systems are synchronous and architectural modifications of this nature during technology migration is a high-risk process and not necessarily very attractive from the point of view of the manufacturers. Moreover, the existence of a global synchronising signal makes it possible to avoid clock data recovery circuits. These complex and costly circuits also introduce latency into the transmission of the data.

For clock distribution in general, the important parameters are:

- temporal precision ("skew", i.e. relative precision between several signals; and "jitter", i.e. instantaneous absolute precision)
- level of coupling of the clock signal to other signals (clock feedthrough)
- level of power consumption and its effects on the thermal characteristics of the system

This last parameter probably constitutes the strongest argument in favour of the abandonment of the metal-based clock distribution. By replacing the electrical clock distribution tree by an optical tree, one would eliminate the need for repeaters and clock multiplier circuits, thus reducing power consumption. By the same reasoning, the random

temporal behaviour ("skew" and "jitter") will be reduced if not eliminated, and the use of a non-electric link would lead to the reduction of clock signal coupling.

However, it would be illusory to believe that the optical clock signal could be routed down to the single-gate level: optoelectronic interface circuits are of course necessary and consume far more area and power than single logic gates. The number of digital gates that can be controlled by a branch of the optical clock distribution is a critical parameter that must be determined. In a previous work [OCO2004] it was found that for the 70nm technology node, power dissipation in an optical clock tree can be estimated under certain conditions to be equivalent to that of an electrical clock tree at around 4000 distribution points. For an 878Mtransistor chip of 310mm^2, this roughly translates to a single drop point delivering the clock signal to a 220Ktransistor zone of 77500µm^2, which is only just feasible. This analysis will be updated for more advanced technology nodes and more precise optical device performance data in the PICMOS project.

Generally there are three main approaches of optical clock distribution: unfocused free-space, focused free-space and guided wave.

The main idea of unfocused free-space clock system is that an off-chip source (attractive from a thermal point of view) broadcasts to the entire chip, with detectors placed at the desired points. The free-space optical system provides flexible distribution with relatively few physical contact points on the board region. However, the unfocused system is very inefficient since only a small fraction of the optical energy is absorbed on the photosensitive areas of the detectors and the rest is wasted. This inefficient use of optical energy results in stringent requirements for amplification of the detected signals on the chip.

In the focused free-space optical system, the optical source sends the signal to the desired detector locations through the use of a focusing element, thus reducing the need for global masking of the unnecessary regions. The disadvantage of the focused interconnect technique is the very high degree of alignment precision that has to be achieved and maintained to ensure that the focused spots are the appropriate places on the chip. Substrate-mode guided-wave distribution, which represents a compact approach to focused free-space distribution, confines the optical signal within the substrate in the transverse direction. Redirection of the optical signal can occur through multiple gratings, mirrors and microlenses located along individual optical paths.

In a guided wave system, the optical source is coupled to a symmetrical passive waveguide network structure (for clock distribution this can be an H-tree structure). In the waveguide, light travels in a high refractive index medium surrounded by lower index medium. In particular, two types of optical waveguides can be used: multi-mode and mono-mode waveguides. Optical waveguides would provide the clock signal to n-number of optical receivers, where the high speed optical signal is converted to a logic level electrical signal (Figure 15).

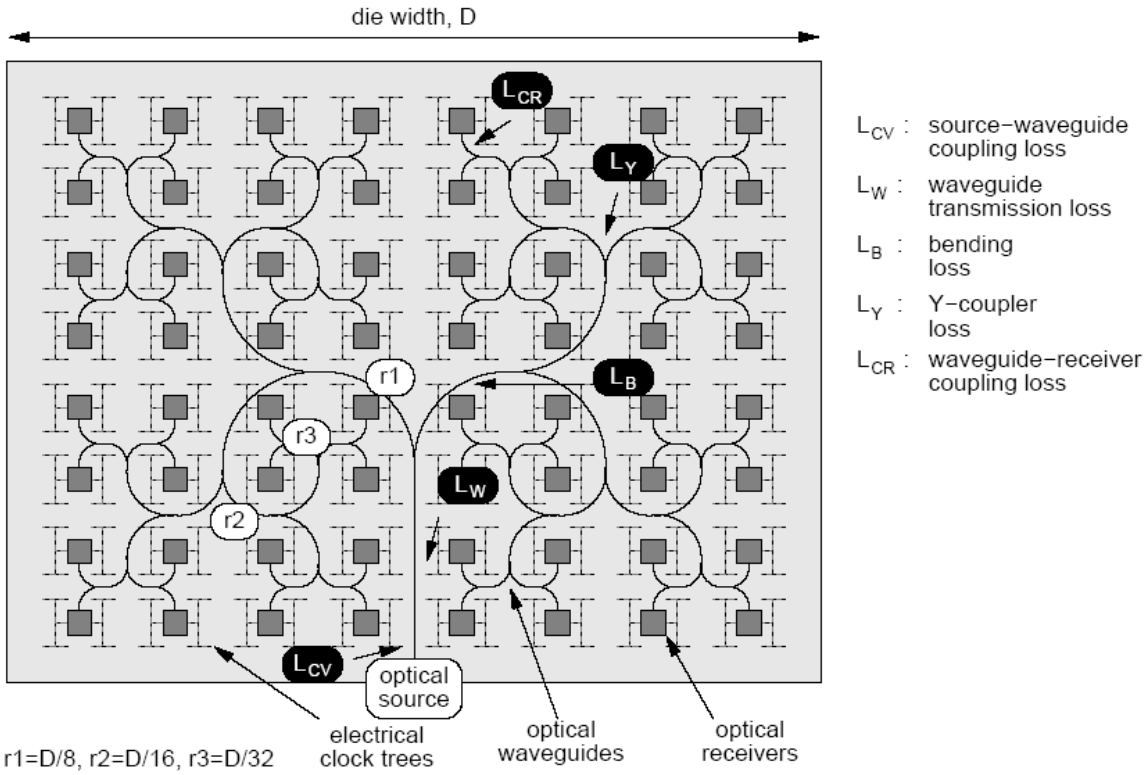

Figure 15 Possible structure for optical clock distribution network

Clock distribution has long been considered to be an ideal application for optical interconnect technology. It is also completely feasible within the scope of the PICMOS project, since all devices necessary to optical clock distribution are being developed within the project. However, recent work in both academia [TOS2004] and in industry [KOB2004] has shown that a) the advantages of optical clock distribution are not so great as could be expected, and b) architectural solutions such as GALS may render clock distribution obsolete before optical technology can even prove the feasibility of clock distribution in the optical domain. However, one fundamental optical function is present in clock distribution: splitting. While this function is absent from most data transport applications, it should still be considered to be a potentially useful function. It could be used in broadcast operations for example, where a processor broadcasts a request to all on-chip memories to verify the coherence of certain data blocks.

For these reasons it is recommended that clock distribution not be chosen as the primary application vehicle, but that a splitting function be included in order to retain sufficient flexibility with respect to niche data transport applications.

6.3. Physical interfacing

Finally, physical interfacing is another potential application for the PICMOS technology: either in off-chip communication to alleviate pin density problems, or in galvanic isolation for power applications, enabling electrical separation of power and control circuits.

6.3.1. I/O and packaging

Pin density is a growing problem [KRI1998]. As shown previously in Table 1, several thousands of pins will soon be necessary for high-performance microprocessors. It is likely that the footprint of an optical pin will be of the same order of magnitude as that of an electrical pin, since both are limited by the same alignment and manipulation issues inherent to bonding technology. However, matrices of detectors similar to CMOS vision systems could be used, one advantage being that pins could be placed anywhere on the chip rather than be placed around the edge as is the case when limited by wire bonding [LIU2002]. Also TDM and WDM techniques could be employed to increase equivalent pin density. For TDM techniques, this would require non-CMOS multiplexers and SERDES[33], while WDM would require optical multiplexers and multiple wavelength emitters. Applications could also be found in massively parallel computing architectures and in telecommunications or FTTH[34].

6.3.2. Galvanic isolation for power applications

Single chip or MCM power applications are faced with a major difficulty in that the electrical connection between power circuits and control circuits is liable to saturate and destroy the control circuits. This is due to the circuit power supply, which due to inductive groundbounce can be subject to strong fluctuations. If these fluctuations are allowed to return to the control circuits, they can be sufficiently large to break the control circuit transistors.

Galvanic isolation avoids this by using an optical coupling scheme which electrically separates power and control circuits. The control signal drives the optical emitter, while the receiver converts the optical signal back to an electrical signal to drive the power circuitry. The supply lines to emitter and receiver circuits are of course separate and are in fact the supply lines to the control and power circuits respectively. Optical interconnect could be used in integrated power applications, and more widely for I/O or packaging roles.

Concerning I/O, no off-chip communication technology is being specifically developed within PICMOS. This application has also already been extensively investigated in the IO project and it is therefore recommended that, in order to break new ground, the PICMOS project application test vehicle should focus on purely intra-chip interconnect. It is also recommended that joint conclusions between IO and PICMOS be drawn at the end of PICMOS. The power application can be considered to be an interesting variant to a data link application recommended previously. The main differences consist of less stringent constraints on bandwidth (mainly one-way transport of control signals), and of a higher importance attributed to temperature sensitivity and to supply voltage variations. For these reasons it is recommended that the application test vehicle also specifically take into account temperature and supply voltage dependencies.

[33] SERializer-DESerializer
[34] Fiber To The Home

7. Impact on design technology

The introduction of optical interconnect technology as an above-IC technique for standard CMOS has huge technological challenges, but the importance of EDA tools and models cannot be overestimated since these are equally necessary to help introduce the technology into existing design flows and methodologies. The main difficulty with design technology for integrated optical interconnects (and for that matter any other non-electronic technology) is that specific non-electronic tools are needed but must be compatible with existing EDA flows.

Possible modifications of an interconnect-centric design flow developed in [JCO2001] to incorporate the use of optical interconnect are shown in Figure 16. Essentially the choice, for a given net, of whether to use optical interconnect or not is based on several criteria (such as the physical characteristics of the net and of the transmitted signal: length, data rate, signal activity, noise sensitivity; and state of congestion in the physical routing process). The physical characteristics of the net are not known before final placement, but early constraints before system refinement are useful to avoid repetitive design loops. These constraints can be generated early in the design cycle through the definition of fixed-complexity (constant number of gates) zones, where ideally intra-zone links are metallic and inter-zone links are optical. This implies that only distance information is taken into account, which is why ***a posteriori*** partitioning should also be carried out after final placement to correct erroneous net allocation to the electrical or optical interconnect structures.

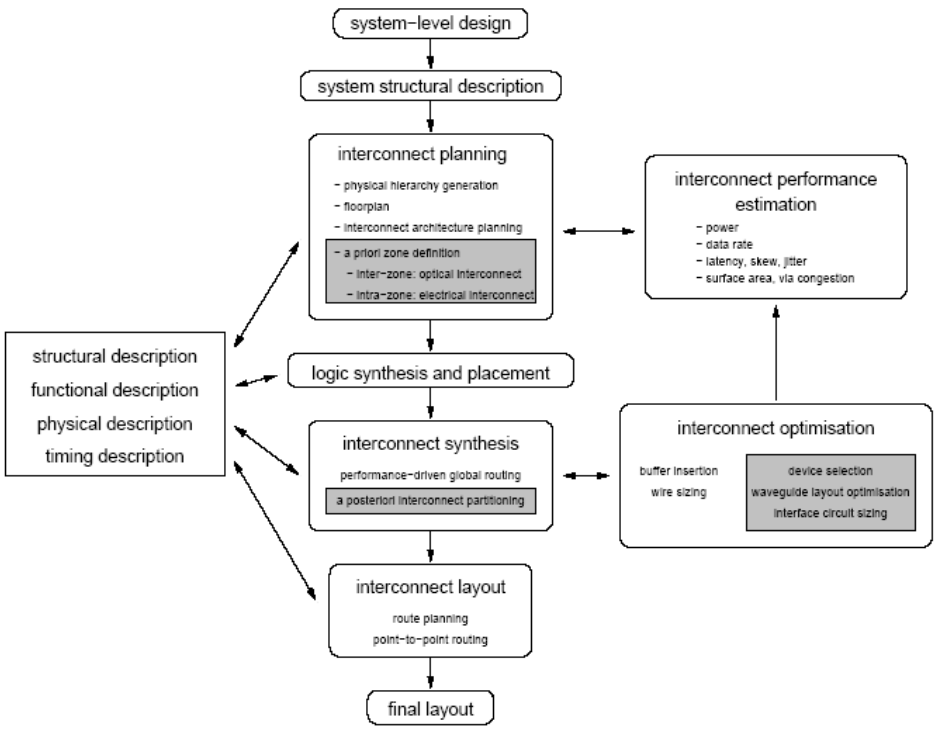

Figure 16 Interconnect-centric design flow with electrical/optical interconnect partitioning

In order for such a flow to work, particular tools, models and IP are required and WP1 in the PICMOS project will go some way to delivering these. In a first instance, accurate behavioural models are required to evaluate the performance of integrated optical interconnect, in short simulation times in order to be compatible with high-level system simulation. They should enable the evaluation of all metrics to objectively compare optical with electrical interconnect (power, bandwidth, latency, surface area, noise…). To enable the design of optical interconnect, multi-domain synthesis tools are also required. This is especially critical for the layout and for the interface circuit blocks, which cannot be designed manually for each optical link. An alternative would be to develop a library of standard interface circuits, offering a sufficiently wide range of performance characteristics to allow a near-optimal choice for all uses of optical interconnect. Finally, early prediction techniques are necessary to make the choice between optical or electrical interconnect. One approach is to develop a solid figure of merit to evaluate tradeoffs between the more fundamental performance characteristics.

8. Conclusion

This document has covered the description and proposal of a number of potential applications for the optical interconnect technology under development in PICMOS. Individual recommendations pertaining to each application's suitability as test vehicle for the PICMOS technology have been given at the end of each description. To summarise, these applications can be compared with the following criteria:

- interest of application (are there clear difficulties already identified with this application in the electronic domain, and is it a necessary application or are there already other workarounds)
- is the application feasible in the PICMOS project
- in the case of negative points, what are the particularly interesting characteristics of the application that might be addressed in the chosen application test vehicle

application	difficulties	necessary	feasible in PICMOS	particular characteristics
data transport				
- data links/bus	++	+	++	
- NoC	+	++	--	switchbox
- FPGA	++	++	--	switchbox
clock distribution	++	-	+	splitter
physical interfacing				
- I/O	++	++	+	link with IO project
- galvanic isolation	+	-	+	temperature

As such it is recommended that the chosen application test vehicle concentrates on parallel long-haul optical data links. It would be useful to integrate splitters into some data links to prove the possibility of broadcast functions and open the route to the realisation of optical multi-way switchboxes. Capabilities of measuring the influence of strong temperature variations (Δ=100°C) on link performance (breakdown, delay,

attenuation and power, but also wavelength shift) should also be possible to evaluate the potential of the PICMOS technology for power applications.

In any case, quantitative answers to the question of the suitability of optical interconnect for data transport problems are required for the semiconductor industry. Integrated optical interconnect is one potential technological solution to alleviate problems associated with moving volumes of data between IP blocks on integrated circuits. But it only makes sense to use this technology for relatively long and high-speed data links, for which other technologies are also in competition (RF and 3D). Also, if the use of the optical interconnect technology implies a hard breach in terms of process technology and design methodologies, then architectural solutions may be an easier way to achieve improvement. The objective would be to optimise the layout for the application such that the need for global high-speed data links is alleviated or even eliminated.

Probably future ICs will make use of advances in both areas. One parameter that is likely to make the difference in favour of optical interconnect is new research into the possibility of on-chip wavelength division multiplexing (WDM). As the aim of the PICMOS project is also to follow the optical technology roadmap such that WDM techniques become possible in the future, both fabrication and design technology should

There is therefore a need to associate imagination at an architectural (interconnect architectures) and structural level (physical devices) with multi-domain and multi-abstraction level modelling (SystemC-Matlab-VHDLAMS-FDTD) and synthesis techniques to get these answers. These themes are part of the PICMOS project and are being addressed in WP1.

Acknowledgements

The authors wish to acknowledge the participation of R. Wilson (STMicroelectronics) and D. Van Thourhout (IMEC) for valuable discussions and comments.

References

BAN2001 K. Banerjee, S.J. Souri, P. Kapur, and K.C. Saraswat, "3-D ICs: A novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration," Proceedings of the IEEE, **89**(5), (2001)

BEN2002 L. Benini and G. De Micheli, "Networks on chip: A new SoC paradigm," IEEE Computer, **35**, (2002)

BER2003 C. Berger, M. Kossel, C. Menolfi, T. Morf, T. Toifl, and M. Schmatz, "High-density optical interconnects within large-scale systems," Proc. SPIE **4942**, 222 (2003)

BRI2004 M. Brière, E. Drouard, F. Tissafi-Drissi, F. Mieyeville, I. O'Connor and F.Gaffiot, "SystemC modeling of an Optical Network-on-Chip using VCI protocol," Proc. GSPx, Santa Clara, CA, USA, September 2004

CAI2001 F. Caignet, S. Delmas-Bendhia, E. Sicard, "The challenge of signal integrity in deep-submicrometer CMOS technology," Proc. IEEE, **89**(4), 556 (2001)

CHA2001 M. Chang et al., "RF/wireless interconnect for inter- and intra-chip communications," Proc. IEEE, **89**(4), (2001)

CHO2004 H. Cho, P. Kapur, K.C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," IEEE J. Lightwave Techn., **22**(9), 2021 (2004)

COL2000 J.H. Collet et al., "Architectural approach to the role of optics in mono- and multi-processor machines," Applied Optics, **39**, (2000)

COL2003 J.H. Collet, F. Caignet, F. Sellaye, D. Litaize, "Performance constraints for onchip optical interconnects," IEEE J. Sel. Top. Quantum Electron., **9**(2), 425 (2003)

DAL2001 W. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in Proc. 38th Design Automation Conference, (2001)

DAM2004 J. Dambre, D. Stroobandt, J. Van Campenhout, "Toward the accurate prediction of placement wire length distributions in VLSI circuits," IEEE Trans. VLSI Systems, **12**(4), (2004)

DAV1998 J. Davis, V. De, and J. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI) - Part I: Derivation and validation," IEEE Trans. Electron Dev., **45**, (1998)

FRI2001 E. Friedman, "Clock distribution networks in synchronous digital integrated circuits," Proc. IEEE, **89**(5), (2001)

GUE2000 P. Guerrier and A. Greiner, "A generic architecture for on-chip packet-switched interconnections," in Proc. Design, Automation and Test in Europe, (2000)

HAU1984 H.A. Haus, "Waves and fields in optoelectronics," Prentice-Hall, Englewood Cliffs, NJ, USA (1984)

HAV2001 R.H. Havemann and J.A. Hutchby, "High-performance interconnects: An integration overview," Proc. IEEE, **89**(5), (2001)

HO2001 R. Ho, K.W. Mai, and M.A. Horowitz, "The future of wires," Proc. IEEE, **89**(4), (2001)

IBM1999 IBM, "The core connect bus architecture," (1999)

ITR2003 Semiconductor Industry Association, "International Technology Roadmap for Semiconductors," (2003)

JCO2001 J. Cong, "An interconnect-centric design flow for nanometer technologies," Proc. IEEE, **89**(4), 505 (2001)

KAP2002 P. Kapur, J. P. McVitie, K. C. Saraswat, "Technology and reliability of constrained future copper interconnects - Part I: Resistance modeling," IEEE Trans. on Electron. Devices, **49**(4), 590 (2002)

KEE2003 G. A. Keeler *et al.*, "The Benefits of Ultrashort Optical Pulses in Optically-Interconnected Systems," IEEE J. Sel. Top. Quantum Electron., **9**, 477 (2003)

KOB2004 M. Kobrinsky *et al.*, "On-Chip Optical Interconnects," Intel Technology Journal, http://developer.intel.com/technology/itj/2004/volume08issue02/, (2004)

KRI1998 A.V. Krishnamoorthy and K.W. Goossen, "Optoelectronic-{VLSI}: Photonics integrated with VLSI circuits," *IEEE J. Sel. Topics in Quantum Electronics*, **4**(6) (1998)

LEE2001 K. Lee *et al.*, "Fabrication of ultralow-loss Si/SiO_2 waveguides by roughness reduction," Optics Letters, **26** (2001)

LEM2004 B.E. Lemoff *et al.*, "MAUI: enabling fiber-to-the-Processor with parallel multiwavelength optical interconnects," IEEE J. Lightwave Techn., **22**(9), 2043 (2004)

LET2005 X. Letartre, LEOM - Ecole Centrale de Lyon, Ecully (France), Personal communication (2005)

LIU2002 Y. Liu, "Heterogeneous integration of OE arrays with Si electronics and microoptics," IEEE Trans. Advanced Packaging, **25**(1), 43 (2002)

MIL2000 D.A.B. Miller, "Rationale and challenges for optical interconnects to electronic chips," Proc. IEEE, **88** (2000)

OCO2004 I. O'Connor and F. Gaffiot, "On-chip optical interconnect for low-power," *Ultra-low power electronics and design*, E. Macii, Ed., Kluwer Academic Publishers (2004)

TOS2004 G. Tosik, F. Gaffiot, Z. Lisik, I. O'Connor and F. Tissafi-Drissi, "Power dissipation in optical and metallic clock distribution networks in new VLSI technologies," *IEE El. Lett.*, **4**(3), 198 (2004)

SAK1983 T. Sakurai and K. Tamuru, "Simple formulas for two- and three-dimensional capacitances," IEEE Trans. on Electron. Device, **30** (1983)

SEA2001 C. Seassal *et al.*, "InP microdisk lasers on silicon wafer: CW room temperature operation at 1.6µm," IEE Electron. Lett., **37** (2001)

VSI2000 VSI Alliance, "http://www.vsi.org"