Invited Talk
Optical Solutions for System-Level Interconnect

Ian O'Connor
Laboratory of Electronics, Optoelectronics and Microsystems
Ecole Centrale de Lyon
36 avenue Guy de Collongue
F-69134 Ecully, France
ian.oconnor@ec-lyon.fr

ABSTRACT

Throughput, power consumption, signal integrity, pin count and routing complexity are all increasingly important interconnect issues that the system designer must deal with. Recent advances in integrated optical devices may deliver alternative interconnect solutions enabling drastically enhanced performance. This paper begins by outlining some of the more pressing issues in interconnect design, and goes on to describe system-level optical interconnect for inter- and intra-chip applications. Inter-chip optical interconnect, now a relatively mature technology, can enable greater connectivity for parallel computing for example through the use of optical I/O pads and wavelength division multiplexing. Intra-chip optical interconnect, technologically challenging and requiring new design methods, is presented through a proposal for heterogeneous integration of a photonic "above-IC" layer followed by a design methodology for on-chip optical links. Design technology issues are highlighted and the paper concludes with examples of the use of optical links in clock distribution (with quantitative comparisons of dissipated power between electrical and optical clock distribution networks) and for novel network on chip architectures.

Categories and Subject Descriptors: B.4.3 **[Hardware]**: Input/ouput and data communications–*Interconnections (Subsystems)*

General Terms: Design

Keywords: Interconnect technology, optical interconnect, optical network on chip

1. INTRODUCTION

Due to continually shrinking feature sizes, higher clock frequencies, and the simultaneous growth in complexity, the role of interconnect as a dominant factor in determining circuit performance is growing in importance. Table 1, drawn from the 2001 ITRS [27] (International Technology Roadmap for Semiconductors), shows that by 2010, high performance integrated circuits will count up to two billion transistors per chip and work with clock frequencies of the order of 10GHz. Coping with electrical interconnects under these conditions will be a formidable task.

	2004	2007	2010	2016
Local wiring pitch (nm)	210	150	105	50
Chip size at production (mm^2)	310	310	310	310
Total interconnect length (m/cm^2)	6879	11169	16063	33508
On-chip local clock (GHz)	3.990	6.739	11.511	28.751
Number of metal levels	9	10	10	11
Maximum power (W)	160	190	218	288
Package pin-count (high-performance)	2263	3012	4009	7100

Table 1: Trends in some integrated circuit parameters

Timing is already no longer the sole concern with physical layout of metallic interconnect although it remains the predominant issue especially when considering progagation time constraints on wires where the length is of the order of the chip size. Several other non-ideal effects, opposed to digital design, are becoming increasingly difficult to ignore. 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. One way of reducing the influence of delay and noise constraints is to increase wire spacing or use wire shielding techniques, both of which cause interconnect resources to be used less efficiently and consequently result in routing congestion or even non-routability. Finally, 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.

A promising approach to the interconnect problem is the use of an optical interconnect layer. Such a layer could empower an enormous bandwidth increase, immunity to electromagnetic noise, a decrease in power consumption, synchronous operation within the circuit and with other circuits, and reduced immunity to temperature variations. Important constraints when developing the optical interconnect layer are the fact that all fabrication steps have to be compatible with future IC technology and that the additional cost incurred remains affordable. 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.

This paper is organised as follows. A classification of intercon-

SLIP'04 February 14–15, 2004, Paris, France

nect types is given in section 2. The optical technology, photonic devices and circuits comprising any optical interconnect structure are described in sections 3, 4 and 5 respectively. A generic design methodology for optical interconnect is proposed in section 6, leading to a discussion of EDA issues in section 7. Finally in section 8, some optical interconnect applications (clock distribution and network on chip) are analysed in detail.

2. INTERCONNECT CLASSIFICATION

For the purposes of this work, data links have been categorised into three broad domains, for which various analyses have been carried out concerning the suitability of optics to each domain: point-to-point (1-1 link); broadcast (1-n link); bus and switching network (n-n link).

2.1 Point-to-point (1-1) links

The basic idea behind using point-to-point optical links consists of replacing electrical global links with optical ones. Research has been carried out on analysing the benefits of introducing optical interconnect in critical data-intensive links, such as CPU-memory buses in processor architectures [6]. 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, even at the expense of greater silicon area and power. Hence, for optical interconnect to be accepted as a real alternative to metallic interconnect, performance gains of at least one order of magnitude must be demonstrated through circuit and device research advances, as well as through application targeting.

2.2 Broadcast (1-n) links

Another and potentially more profitable application of optical interconnect technology 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 [9] (CDN). In order to operate at high frequencies, CDNs require several hundreds of repeaters to drive the metallic tracks over the entire chip, resulting in using 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. 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.

2.3 Network links

Network, or n-n optical links, where wavelength routing may be used, would be targeted at (i) optical buses and possibly (ii) reconfigurable networks. System architectures are moving rapidly towards platform-based designs, whereby every functional block on the chip (DSP, analog/RF, video processors, memory etc.) interfaces to an interconnect network architecture for data communication. Global system on chip communication is around several tens of Gb/s.

A number of innovative interconnect architectures, often called networks on chip (NoC), have been recently proposed [3] [13] [7] to overcome the limitations of commercial bus-based platforms (e.g. latency due to line delay, non-scalability, time-sharing, non-reconfigurability, etc.). NoC architectures look much more like switching

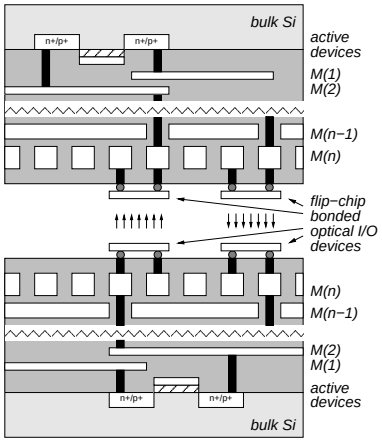

Figure 1: Inter-chip communication using flip-chip bonded optical I/O devices

telecommunication networks than conventional bus-based architectures. 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. Latency, connectivity, global throughput and reconfigurability constitute the main performance indicators of these networks.

Integrated optics may constitute an effective and attractive alternative for NoC. The superiority of optical interconnects in long distance links 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. But, above all, wavelength division multiplexing (WDM) may offer new and appealing solutions such as optical buses and reconfigurable networks (see section 8 for a discussion of optical NoCs).

3. OPTICAL INTERCONNECT TECHNOLOGY

3.1 Inter-chip optical interconnect technology

Systems where integrated circuit components are physically separate can benefit from relatively mature optical technology. For example, parallel computing architectures require high-bandwidth data links between chips capable of transmitting, receiving and even switching data at aggregate data rates exceeding 1Tb/s. It has already been shown [12] that the figures given by the ITRS for electrical I/O are inadequate when calculating the number of I/O required while keeping the ratio of I/O to computational bandwidth (number of gates $\times$ clock frequency) constant. This is equally valid for board-to-board and chip-to-chip communication.

Using optical I/O alleviates pin density problems. The principal idea is that the CMOS IC is bonded to an optoelectronic chip using routine flip-chip bonding techniques (fig. 1). The optoelectronic chip thus bonded is actually an array of several thousands of modulators and diodes driven by information leaving and entering the chip vertically. Typical pitch between active array elements is 50μm, giving a possible 160000 I/O elements over a 2×2 cm^2 chip surface.

Future ideas are for WDM to be used between chips. Obviously this augments the channel capacity of each I/O link but active devices are more difficult to fabricate because they have to be able to change emission or detection wavelength in a fraction of a clock cycle without increasing surface area. The former constraint eliminates mechanical switching solutions, the latter eliminates multiple-circuit solutions.

Connection of the optical interconnect network and the electronic IC is a non-trivial aspect to the whole optical interconnect concept. Probably the most effective and proven technique is flip-chip bonding [15]. This involves the depositing of gold solder bumps on either the electronic or photonic IC, then alignment and finally bonding, usually using thermocompression. At the wafer-scale bonding level, advanced machines are capable of precision alignment down to the order of 1μm. In such cases, the solder bump can be made small (under 10μm diameter) so that the total capacitance of the link, including pads, is of the order of a few tens of fF, compatible with high-speed interface circuit operation.

More futuristic ideas concern auto-selective DNA bonding and epitaxial integration, where the III-V material is grown directly onto the silicon substrate. This is possible, but the temperatures involved are usually rather high (800C). However, recent research [26] has successfully demonstrated hydrophilic wafer bonding of an InP microdisk laser onto a silicon wafer at room temperature by means of SiO_2 layers on both InP and silicon substrates. The highest temperature in this process was 200C (for annealing, to increase bonding energy), which is compatible with CMOS IC fabrication steps.

3.2 Intra-chip optical interconnect technology

Various technological solutions may be proposed for integrating an optical transport layer in a standard CMOS system. The first choice is to specify where this optical layer has to be placed. Then one has to choose the different materials used for the active devices and the passive transport layer.

3.2.1 Hybrid or monolithic

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.

Hybrid integration of the optical layer on top of a complete CMOS IC 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. Fig. 2 shows a cross-section of how a complete "above IC" photonic layer could be realised. The photonic source shown 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 (multi-chip module) 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.

In the system shown in fig. 2, the microsource is coupled to the passive waveguide structure and provides a signal to an optical receiver (or possibly to several, as in the case of a broadcast function). At the receiver, the high-speed optical signal is converted to an electrical signal and subsequently distributed by a local electrical interconnect network.

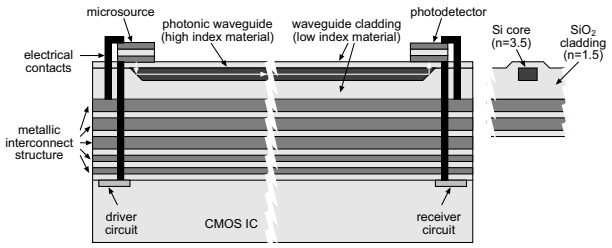

Figure 2: Cross-section of hybridised interconnection structure

3.2.2 Materials

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 but we focus here on Si/SiO_2 waveguides, where silicon is used as the core and SiO_2 as the cladding material. Si/SiO_2 structures are compatible with conventional silicon technology and silicon is an excellent material for wavelengths above 1.2μm. Monomode waveguiding with attenuation as low as 0.8 dB/cm has been proven [16]. The high relative refractive index difference $\Delta = (n_1^2 - n_2^2)/2n_1^2$ between the core ($n_1 \approx 3.5$ for Si) and cladding ($n_2 \approx 1.5$ for SiO_2) allow the realisation of a compact optical circuit with dimensions compatible with DSM technologies. For example, it is possible to realise monomode waveguides less than 1μm wide (waveguide width of 0.3μm for wavelengths of 1.55μm), with bend radii of the order of a few μm [24].

4. INTEGRATED OPTOELECTRONIC DEVICES

In this section a brief overview of integrated optoelectronic devices is given. Any optical link is composed of a interface (driver and receiver) circuits, active optoelectronic devices (light source and detector), and (in the case of guided optical links) passive photonic devices (waveguides and wavelength routing structures).

4.1 Passive photonic devices for signal routing

4.1.1 Waveguides

Optical system performance depends on the minimum optical power required by the receiver and on the efficiency of passive optical devices used in the system. The total loss in any optical link (represented in fig. 3) is the sum of losses (in decibels) of all optical components:

$$L_{total} = L_{CV} + L_B + L_Y + L_{CR} \quad (1)$$

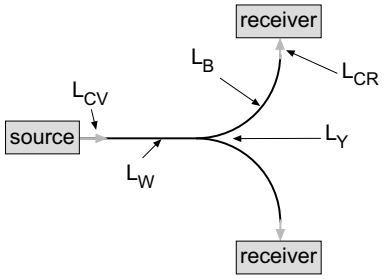

Figure 3: Losses in an optical link

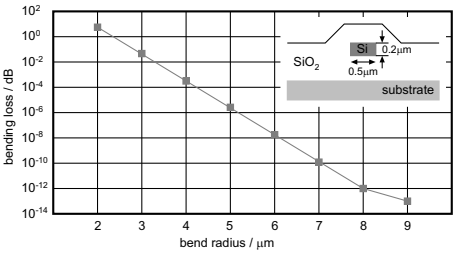

Figure 4: Simulated bending loss for Si/SiO_2 strip waveguide

where L_{CV} is the coupling coefficient between the photonic source and optical waveguide, L_W is the rectangular waveguide transmission loss, L_B is the bending loss, L_Y is the Y-coupler loss and L_{CR} is the coupling loss from the waveguide to the optical receiver.

There are currently several methods to couple the beam emitted from the laser into the optical waveguide. In this analysis we assumed 50% coupling efficiency L_{CV} from the source to a single mode waveguide.

Transmission loss L_W describes the attenuation per unit distance of the optical power. Due to small waveguide dimensions and large index change at the core/cladding interface in the Si/SiO_2 waveguide the side-wall scattering is the dominant source of loss. For the waveguide fabricated by Lee [16] with roughness of 2nm the calculated transmission loss is 1.3dB/cm.

The bending loss L_B is highly dependent on the refractive index difference Δ between the core and cladding medium. In Si/SiO_2 waveguides, Δ is relatively high and so due to this strong optical confinement, bend radii as small as a few μm may be realised. As can be seen from fig. 4, the bending losses associated with a single mode strip waveguide are negligible if the radius of curvature is bigger then 2μm.

The Y-junction loss L_Y depends on the reflection and scattering attenuation into the propagation path and surrounding medium. For high index difference waveguides the losses for the Y-branch are significantly smaller then for the low Δ structures and the simulated losses are less then 0.2dB per split [23].

Using currently available materials and methods it is possible to achieve an almost 100% coupling efficiency from waveguide to optical receiver. In this analysis the coupling efficiency L_{CR} from the waveguide to the optical receiver is assumed to be 87% [25].

4.1.2 Resonators

Microdisks are resonating structures, and are most commonly used in "add-drop" filters (so-called because of their capacity to add or subtract a signal from a waveguide based on its wavelength). The filter itself (fig. 5) is composed of one or more identical disks evanescently side-coupled to signal waveguides. The electromagnetic field is propagated within the structure only for modes corresponding to specific wavelengths, where these resonant wavelength values are determined by geometric and structural parameters (substrate and microdisk material index, thickness and radius of microdisk).

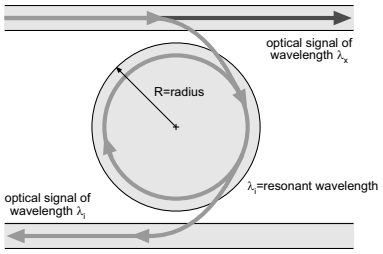

Figure 5: Micro-disk realisation of an add-drop filter

The basic function of a microresonator can be thought of as a wavelength-controlled switching function. If the wavelength of an optical signal passing through a waveguide in proximity to the resonator does not correspond to the resonant wavelength, then the electromagnetic field continues to propagate along the waveguide and not through the structure. If, however, the signal wavelength is close enough to the resonant wavelength (tolerance is of the order of a few nm, depending on the coupling strength between the disk and the waveguide), then the electromagnetic field propagates around the structure and then out along the second waveguide. Switching has occurred, based on the physical properties of the signal.

A fairly obvious application of this device is in optical crossbar networks. More elaborate N×N switching networks have been devised [17], but experimental operation has yet to be proven. The advantages of such structures lie in the possibility of building highly complex, dense and passive on-chip switching networks.

4.1.3 Photonic crystals

Photonic crystals are nanostructures composed of, in the 2D case, ultra-small cylinders periodically arranged in a background medium. 3D photonic crystals also exist but are much more difficult to fabricate from a technology point of view. Typically for 2D photonic crystals, the cylinders are realised in a low-index material (such as SiO_2, or air), the background being a high-index material (such as Si). For light of certain wavelengths, such structures have a photonic band gap, leading to optical confinement. By introducing line defects (i.e. by removing one or more rows of cylinders), single-mode waveguides can be created. The extraordinary versatility of photonic crystals is illustrated by the possiblility of other functions such as couplers [19], multiplexers, demultiplexers, microresonators (using point defects instead of line defects) and even lasers. Photonic crystals are certainly good candidates for microscale optical integrated circuits due to their small size (a typical value for waveguide pitch is 0.5μm) and massive fabrication potential. However, attenuation is an order of magnitude higher than that in planar waveguides (6dB/mm [21]), although good progress has recently been made in this area.

4.2 Active Devices for Signal Conversion

4.2.1 III-V sources

Fundamental requirements for integrated semiconductor lasers are small size, low threshold lasing operation, single-mode operation (i.e. only one mode is allowed in the gain spectrum). From the viewpoint of mode field confinement and mirror reflection, two

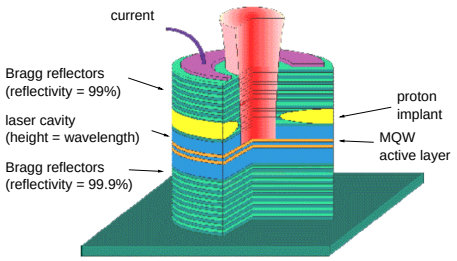

Figure 6: Vertical cavity surface emitting laser (VCSEL) structure

types of microcavity structures exist: multiple reflection (VCSELs and photonic crystals), and total internal reflection (microdisks). An overview of microcavity semiconductor lasers can be found in [2].

4.2.2 VCSELs

VCSELs (Vertical Cavity Surface Emitting Lasers) are without doubt the most mature emitters for on-chip or chip-to-chip interconnections. As their name indicates, light is emitted vertically at the surface as shown in fig. 6, by stimulated emission via a current above a few microamperes.

VCSELs are intrinsically single-mode due to their small cavity dimensions. They also have a very low threshold current, low divergence and arrays of VCSELs are easy to fabricate. However, the internal cavity temperature can become quite high, and this is important because both wavelength and optical gain are dependent on the temperature.

Commercial VCSELs, when forward biased at a voltage well above 1.5V, can emit optical power of the order of a few mW around 850nm, with an efficiency of some 40%. Threshold currents are typically in the mA range. It is clear that effort is required from the research community if VCSELs are to compete in the on-chip optical interconnect arena, to increase wavelength, efficiency and threshold current. Long wavelength, and low-threshold VCSELs are only just beginning to emerge (for example, a 1.5μm, 2.5Gb/s tuneable VCSEL [8], and an 850nm, 70μA threshold current, 2.6μm diameter CMOS compatible VCSEL [18] have been reported).

4.2.3 Microsources

Integrated microlasers differ from VCSELs in that light emission is in-plane to be able to inject light directly into a waveguide with minimum loss. The structure of a microdisk laser is shown in fig. 7 [10]. Such devices, to be compatible with dense photonic integration, must satisfy the requirements of small volume and high optical confinement, with low threshold current and emitting in the 1.3-1.6μm range. This wavelength implies the necessary use of InP and related materials, which leads to heterogeneous integration: bonding issues arise which were covered in section 3.1.

4.2.4 Detectors

Conventional PIN devices have relatively small area per unit capacitance values, meaning that the optical responsivity bandwidth product is low. This is a problem for high-speed operation in optical interconnect since transimpedance amplifier interface circuits cannot support high photodetector capacitance values. Current research is focusing on thin-film metal-semiconductor-metal (MSM) photodetectors, due to their improved area per unit capacitance [5].

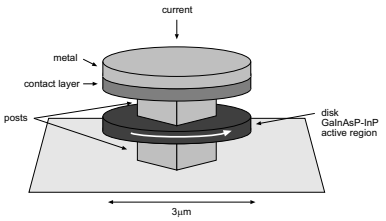

Figure 7: Structure of a microdisk laser

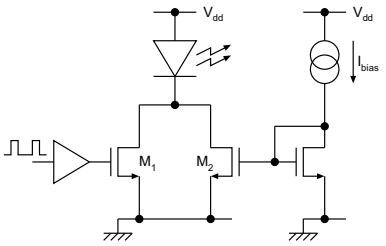

Figure 8: Basic current modulation source driver circuit

5. INTERFACE CIRCUITS

Between electronic data processing and photonic data transport lie crucial building blocks to the optical interconnect solution: high-speed optoelectronic interface circuits. On the emitter side, the power dissipated by the source driver is largely governed by the bias conditions required for the source itself. Advances in this area thus follow to a large extent improvements resulting from device research. On the receiver side however, things are rather different: most of the receiver power is due to the circuit. Only a small fraction is required for the photodetector device.

5.1 Driver circuits

The basic current modulation configuration of the source driver circuit is shown in fig. 8. The source is biased above its threshold current by M_2 to eliminate turn-on delays, and as the bias current value is the main contributing factor to emitter power, reducing the source threshold current is a primary device research objective. Figures of the order of 40μA [11] have been reported. Device M_1 serves to modulate the current flowing through the source, and consequently the output optical power injected into the waveguide. As with most current-mode circuits, high bandwidth can be achieved since the voltage over the source is held relatively constant and parasitic capacitances at this node have reduced influence on the speed.

5.2 Receiver circuits

The classical structure for a receiver circuit is shown in fig. 9: a transimpedance amplifier (TIA) converts the photocurrent of a few μA into a voltage of a few mV; a comparator generates a rail-to-rail signal; and a data recovery circuit eliminates jitter from the restored signal. Of these, the TIA is arguably the most critical component fpr high-speed performance, since it has to cope with a generally large photodiode capacitance situated at its input.

The basic transimpedance amplifier structure in a typical configuration is shown in fig. 10 [14]. The bandwidth/power ratio of this structure can be maximised by using small-signal analysis and mapping of the individual component values to a filter approxima-

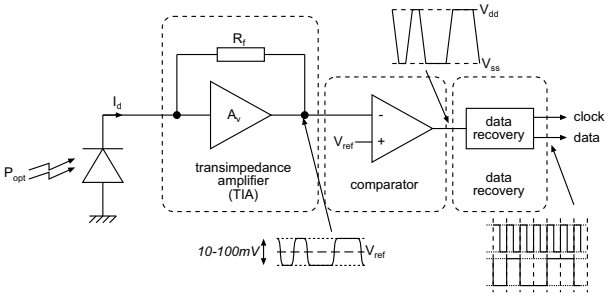

Figure 9: Typical photoreceiver circuit

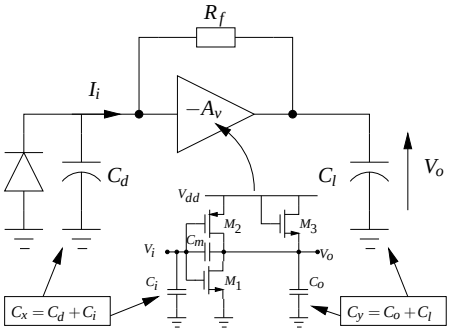

Figure 10: CMOS transimpedance amplifier structure

tion of Butterworth type, which gives:

$$Z_{g0} = \frac{R_o - R_f A_v}{1 + A_v} \tag{2}$$

$$\omega_0 = \frac{1}{R_o C_y} \sqrt{\frac{1 + A_v}{M_f (M_x + M_m + M_x M_m)}} \tag{3}$$

$$Q = \frac{\sqrt{M_f (M_x + M_m (1 + M_x))(1 + A_v)}}{1 + M_x (1 + M_f) + M_m M_f (1 + A_v)} \tag{4}$$

where the multiplying factors $M_f = R_f / R_o$, $M_i = C_x / C_y$ and $M_m = C_m / C_y$ are introduced, normalising all expressions to the time constant $\tau = R_o C_y$. By rearranging these equations, it is then possible to develop a synthesis procedure which, from desired transimpedance performance criteria (Z_{g0}, bandwidth and Q) and operating conditions (C_d, C_l) generates component values for the feedback resistance R_f and the voltage amplifer (A_v and R_o).

Taking into consideration the physical realisation of the amplifier, those with requirements for low gain and high output resistance (high R_o / A_v ratio) are the easiest to build, and also require the least quiescent current and area. Fig. 11 shows a plot of this quantity against the TIA specifications (bandwidth and transimpedance gain) for $C_x = C_d = 500\,fF$ and $C_y = Cl = 100\,fF$.

Approximate equations for the small-signal characteristics and bias conditions of the circuit allow a first-cut sizing of the amplifier. The solution can then be fine-tuned by numerical or manual optimisation, using simulation for exact results [22].

Using this methodology and predictive BSIM3v3 models for technology nodes from 180nm down to 70nm [4], we generated design parameters for $1THz\Omega$ transimpedance amplifiers to evaluate the evolution in critical characteristics with technology node. Fig. 12 shows the results of transistor level simulation of fully generated photoreceiver circuits at each technology node. According to tra-

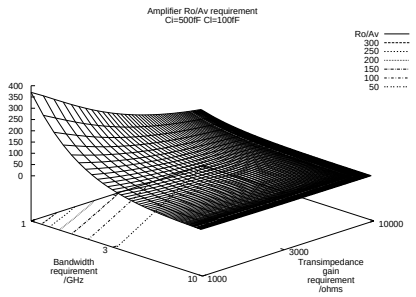

Figure 11: TIA R_o / A_v design space with varying bandwidth and transimpedance gain requirements

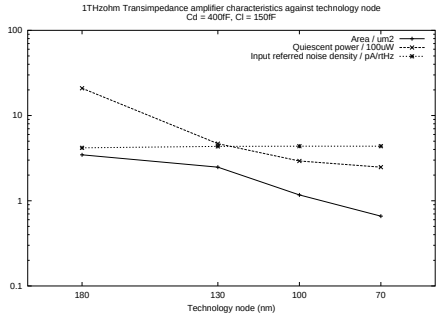

Figure 12: Evolution of TIA characteristics (power, area, noise) with technology node

ditional "shrink" predictions, which consider the effect of applying a unitless scale factor of $1/s$ to the geometry of MOS transistors, the quiescent power and device area should decrease by a factor of $1/s^2$. Between the 180nm and 70nm technology nodes, $s^2 \approx 6.61$, which is verified through the sizing procedure.

6. LINK DESIGN METHODOLOGY

6.1 Design requirements

To make a reasoned comparison between electrical and optical interconnect, a set of design requirements must be established. The main criterion in evaluating the performance of digital transmission systems is the resulting bit error rate (BER), which may be defined as the rate of error occurrences. Typical BER figures required by Gigabit Ethernet and by Fiber Channel is 10^{-12} or better. For an on-chip interconnect network, a BER of 10^{-15} is acceptable. It should be noted here that BER is not commonly considered in IC design circles, and for good reason: metallic interconnects typically achieve BER figures better than 10^{-45}! However, future operating frequencies are likely to change this, since the combination of necessarily faster rise and fall times, lower supply voltages and higher crosstalk increases the probability of wrongly interpreting the signal that was sent.

To the list of design requirements we must also include the ubiquitous power/speed/area trinity found in any digital system. Power and speed can be compared directly, while area (in the 3D scenario) is more difficult to evaluate since we are essentially aiming at adding a photonic layer of the same size as the chip itself. What is important therefore is the average achievable area/bit ratio.

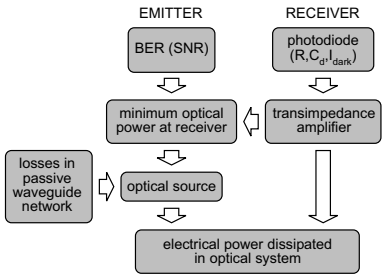

Figure 13: Methodology used to estimate the electrical power dissipation in an optical clock distribution network

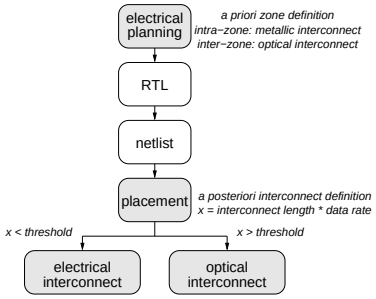

Figure 14: Design flow modifications for integrated optical interconnect

6.2 Design methodology

In order to be able to evaluate and optimise link performance criteria correctly, predictive models and design methodologies are required. Concerning the power aspects, the aim is to establish the overall power dissipation for an optical link at a given data rate and BER. The calculation is essentially conditioned by the receiver, since the BER defines the lower limit for the received optical power. This lower limit can then be used to calculate the required power coupled into waveguides by optical sources, the required detector efficiency (including optical coupling) and acceptable transmission losses. Power can then be estimated from source bias current and photoreceiver front-end design methodologies.

For integration density aspects, source and detector sizes must be taken into account, while the width, pitch and required bend radius of waveguides is fundamental to estimating the size of the photonic layer. On the circuit layer, the additional surface due to optical interconnect is in the driver and receiver circuits, as well as the depassivated link to the photonic layer. The circuit layout problem is compounded by the necessity of using clean supply lines (i.e. separate from digital supplies) to reduce noise (for BER).

The data rate is essentially governed by the bandwidth of the photoreceiver: high modulation speed at the source is generally more easily attainable than similar detection speed at the receiver. This is essentially due to the photodiode parasitic capacitance at the input of the transimpedance amplifier.

Apart from these concerns, functional aspects also have to be considered. For example, using the same signal to drive two nodes is not trivial (as is the case in electrical interconnect) since the layout of a 1-2 splitter is crucial to the equal distribution of power to each node. More fundamentally, dividing the power has a direct influence on the power required at the source in order to achieve the lower power limit at the receiving nodes.

In an optical link there are two main sources of electrical power dissipation: (i) power dissipated by the optical receiver(s) and (ii) energy needed by the optical source(s) to provide the required optical output power. To estimate the electrical power dissipated in the system we developed the methodology shown in fig. 13.

First, based on the given photodiode parameters (C_d, R, I_{dark}), the method described in section 5.2 is used for the design of the transimpedance amplifier. Next, for a given system performance (BER) and for the noise signal associated with the photodiode and transimpedance circuit we calculate the minimum optical power required by the receiver to operate at the given error probability, using the Morikuni formula [20] in the preamplifier noise calculations.

To estimate the optical power emitted by the source we take into account the previously calculated minimal required signal by the receiver and the losses incurred throughout the passive optical waveguide structure. The electrical power dissipated in the optical link is the sum of the power dissipated by the number of optical receivers and the energy needed by the source to provide the required optical power. The electrical power dissipated by the receivers can be extracted from transistor-level simulations. To estimate the energy needed by the optical source, we use the laser light-current characteristics given by Amann [1].

7. DESIGN TECHNOLOGY ISSUES

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. In a nutshell, 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.

Figure 14 shows a proposal for the modification of a typical design flow to incorporate the use of optical interconnect. Essentially the choice of whether to use optical interconnect or not is based on the distance between two points to be connected and the data rate of the signals between the points. Of course, this information is 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 correction* should be carried out after final placement.

This section describes the methods that are typically used in the optical domain and then details how such tools can be integrated into a standard EDA environment. We are aiming at remaining compatible with top-down and bottom-up methodologies characteristic of IC design, enabling IP reuse. This essentially requires a definition of hierarchical description levels for all components, with short simulation time at system level and high simulation accuracy at device level.

7.1 Hierarchical design

Figure 15 shows the hierarchical levels involved in optical link design. The definition of hierarchical levels of components in an optical link implies the necessary use of multi-domain behavioural modelling, compatible with electrical simulation methods. VHDL-AMS[1] [30](IEEE standard 1076.1) is a language that is adapted to

[1]VHSIC (Very High Speed Integrated Circuit) Hardware Description Language Analog and Mixed-Signal

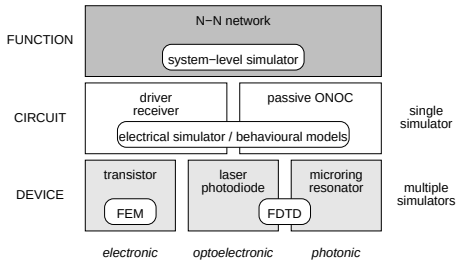

Figure 15: Hierarchical levels involved in optical network on chip and associated simulators

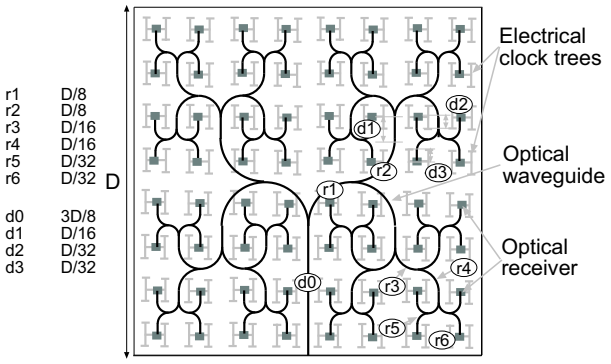

Figure 16: Optical H-tree clock distribution network (OCDN) with 64 output nodes. r_{1-6} are the bend radii and d_{0-3} are the lengths of straight lines linked to the chip width D.

these needs since it allows simulation in the digital, analogue and mixed-signal domains incorporating a transposition of non-electric (e.g. optical, thermal, mechanical) physical phenomena to analogies of electrical behaviour.

A library of models has been created with several "black boxes" which describe the behaviour of a photonic component with its spectral response R. A photonic device can be considered as a filter with several inputs and outputs. All the inputs have the same format type with three values: E_{opt}, optical field transmitted by the laser, ϕ, the central wavelength value, and $\Delta\lambda$, the width of the impulsion of the light (the input is supposed to be of gaussian form). All outputs have the same format type with one essential value : $E(t,\omega(i))$ which represents the value of the amplitude. This library enables the creation of photonic circuits by virtual component assembly.

7.2 Physical optical design

To optimise the photonic devices it is necessary to go down to a low level of simulation. The main family of methods used to analyse the behaviour of optical devices is the Finite Element Method type (FEM). We have developed an in-house program implementing the Finite-Difference Time-Domain algorithm [28]. This algorithm calculates the electric and magnetic field (i.e. the "value" of the light wave) in the structure. These electromagnetic fields give information about the propagation of the wave in a structure.

The FDTD algorithm is analogous to that used for modelling scalar wave propagation. The physical structure is meshed with several cells. Maxwell's equations are discretised by using a finite difference scheme. An example of this discretisation is shown with equation set 5 (for the 2D case), where (i,j) represents the position of the cell on the mesh where the electromagnetic field is to be calculated (E the electric field and H the magnetic field). Since FDTD is a temporal method, a new field value is calculated at each iteration (variable n in equation 5) in all the cells of the mesh. With this method, an optimisation of the device is possible (dimension, injection value, ...).

$$\begin{cases} H^{n+1/2}_{x_{i,j}} = H^{n-1/2}_{x_{i,j}} + C(E^{n}_{z_{i,j-1/2}} - E^{n}_{z_{i,j+1/2}}) \\ H^{n+1/2}_{y_{i,j}} = H^{n-1/2}_{y_{i,j}} + C(E^{n}_{z_{i+1/2,j}} - E^{n}_{z_{i-1/2,j}}) \\ E^{n+1}_{z_{i,j}} = E^{n}_{z_{i,j}} + CN_{i,j}(H^{n+1/2}_{y_{i+1/2,j}} - H^{n+1/2}_{y_{i-1/2,j}} \\ \qquad\qquad + H^{n+1/2}_{x_{i,j-1/2}} - H^{n+1/2}_{x_{i,j+1/2}}) \end{cases} \quad (5)$$

8. APPLICATIONS

8.1 Clock distribution

An example system realising a clock distribution function, shown in fig. 16, requires a single photonic source coupled to a symmetrical waveguide structure routing to a number of optical receivers. At the receivers the high-speed optical signal is converted to an electrical one and provided to local electrical networks. The number of clock distribution points (64 in this example) is a particularly crucial parameter in the overall system.

We assumed the use of an external VCSEL source, rather than an integrated microsource. The global optical H-tree was optimised to achieve minimal optical losses. The radius of curvatures are designed to be as large as possible. For 20mm die width and 64 output nodes in the H-tree at the 70nm technology mode, the smallest radius of curvature (r_5, r_6 in fig. 16) is 625μm, which leads to negligible pure bending loss.

To provide a clear comparison in terms of dissipated power between the optical and electrical interconnect networks it is necessary to estimate the electrical power dissipated in both systems. As an example, the power dissipated in clock distribution networks was analysed in both systems at the 70nm technology node. Power dissipation figures for electrical and optical CDNs were calculated based on the system performance summarised in table 2. The power dissipated in the electrical system can be attributed to the charging and discharging of the wiring and load capacitance and to the static power dissipated by the buffers. In order to calculate the power we used an internally developed simulator, which allows us to model and calculate the electrical parameters of clock networks for future technology nodes [29].

Fig. 17 shows the optical power emitted by the VCSEL necessary to provide a given BER for various waveguide transmission losses. For a BER of 10^{-15} the minimal power required by the receiver is -22.3dBm (at 3GHz). Losses incurred by passive components for various nodes in the H-tree are summarised in table 3. The comparison in terms of dissipated power between the optical and electrical global clock distribution networks is shown in fig. 18. It can be seen that the power dissipated by the electrical system is highly dependent on the operating frequency, while in the optical system it remains almost the same. The difference between the power dissipated in both systems is clearly higher if we increase the frequency and number of nodes in H-trees. For a classical 64-node H-tree at 5GHz frequency the power consumption in the optical CDN should be 5 times lower than in an electrical network.

8.2 Optical network on chip

An optical network on chip (ONoC) can realise the same functions as an electronic router with routing based on wavelength λ. Moreover, there are some advantages to using optical waveguides: i) design simplification (absence of crosstalk and impedance matching, waveguide length independence of performance of optical in-

Electrical system	
Technology (nm)	70
V_{dd} (V)	0.9
T_{ox} (nm)	1.6
Chip size (mm^2)	400
Global wire width (μm)	1
Metal resistivity (Ω-cm)	2.2
Dielectric constant	3
Optimal segment length (mm)	1.7
Optimal buffer size (μm)	90
Optical system	
Wavelength λ (nm)	1550
Waveguide core index (Si)	3.47
Waveguide cladding index (SiO_2)	1.44
Waveguide thickness (μm)	0.2
Waveguide width (μm)	0.5
Transmission loss (dB/cm)	1.3
Loss per Y-junction (dB)	0.2
Input coupling coeffi cient (%)	50
Photodiode capacitance (fF)	100
Photodiode responsivity (A/W)	0.95

Table 2: Electrical and optical system performance

Number of nodes in H-tree	16	32	64	128
Loss in straight lines (dB)	1.3	1.3	1.3	1.3
Loss in curved lines (dB)	1.53	1.66	1.78	1.85
Loss in Y-dividers (dB)	12	15	18	21
Loss in Y-couplers (dB)	0.8	1	1.2	1.4
Output coupling loss (dB)	0.6	0.6	0.6	0.6
Input coupling loss (dB)	3	3	3	3
Total optical loss (dB)	19.2	22.5	25.8	29.1
Min. receiver power (dBm)	-22.3	-22.3	-22.3	-22.3
Laser optical power (mW)	0.5	1.1	2.30	4.85

Table 3: Optical power budget for 20mm die width at 3GHz

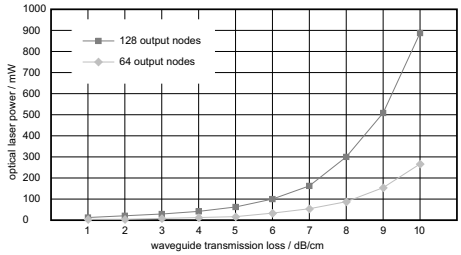

Figure 17: VCSEL optical output power required by the H-tree to provide a given BER for varying waveguide transmission loss

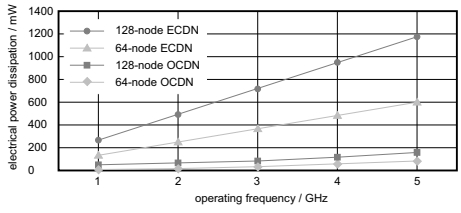

Figure 18: Electrical power dissipated by electrical (ECDN) and optical (OCDN) clock distribution networks for varying operating frequency

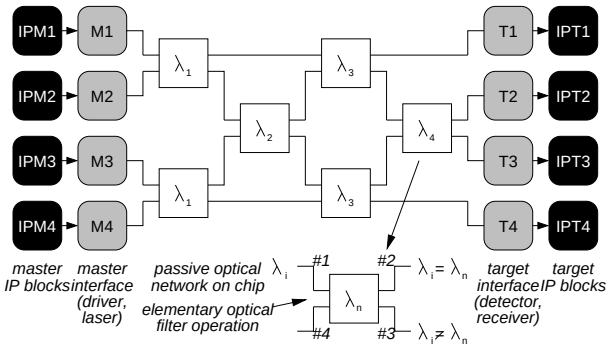

Figure 19: Architecture of 4x4 optical network on chip

	T_1	T_2	T_3	T_4
M_1	λ_2	λ_3	λ_1	λ_4
M_2	λ_3	λ_4	λ_2	λ_1
M_3	λ_1	λ_2	λ_4	λ_3
M_4	λ_4	λ_1	λ_3	λ_2

Table 4: Truth table for optical network on chip

terconnects, frequency independence of optical interconnects), ii) architectural advantages (large synchronous zones, architectures with large numbers of long high-speed connections), iii) timing (removal of timing skew in signal, predictability of the timing of signals, precision of the timing of the clock signal), iv) other physical benefi ts (reduction of power dissipation in interconnects, "high" interconnect density).

Figure 19 shows a 4×4 ONoC with all electronic interfaces: photodetector and laser in III-V technology and optical network in SOI technology. Intellectual property (IP) blocks shown can be processors, memories, etc. This is a multi-domain device with high speed optoelectronic circuits (modulation of the laser current and photodetectors) and passive optics (waveguides and passive fi lters). In the fi gure, M are masters (processor, IP, ...) which can communicate with targets T (memory, ...). The network is comprised of 4 stages, each associated with a single resonant wavelength. The operation of the 4×4 network is summarised in the table of fi gure 4. This system is a fully passive circuit-switching network based on wavelength routing and is a non-blocking network. From M_i to T_j, there exists only one physical path associated with one wavelength. At any one time, single-wavelength emitters can make 4 connections and multi-wavelength emitters can make 12 connections.

The basic element is an optical fi lter, described in section 4.1.2. The ports $\sharp1 - \sharp4$ correspond to inputs/outputs of the optical fi lter. Its operation is the same as an electronic cross-bar: the cross function (output in $\sharp c$) is activated when the injected wavelength in $\sharp a$ does not correspond to a resonant ring wavelength and the bar function is activated (output in $\sharp b$) when the injected wavelength in $\sharp a$ corresponds to a resonant ring wavelength. Operation is symmetrical : the same phenomena happens if the wavelength injection is placed in the port $\sharp d$.

Behavioural models enable us to verify the operation of the 4×4 ONoC at high level simulation. An injection of 4 wavelengths is realised (λ_1, λ_2, λ_3, and λ_4) at the port $\sharp1$ at the same moment (shown in fi gure 20). The input signal format is a matrix. Figure 20 is a 3-dimensional representation with wavelength on the X-axis (representing the 4 channels), time on the Y-axis and power (normalised) on the vertical axis. Each injected wavelength has two pulses (Gaussian) in time. The behavioural simulation analyses the

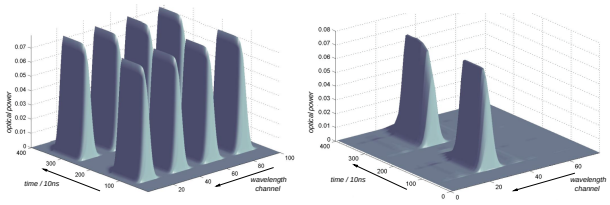

Figure 20: Simulation of 4x4 optical network on chip

4 outputs T_1, T_2, T_3 and T_4 (T_2 shown in fig. 20). As predicted in table 4, only λ_3 is detected at the output T_2.

9. CONCLUSION

Integrated optical interconnect is one potential technological solution to alleviate some of the more pressing issues involved in moving volumes of data between circuit blocks on integrated circuits. In this paper, we have shown how novel integrated photonic devices can be fabricated above standard CMOS ICs, designed concurrently with EDA tools and used in clock distribution and NoC applications. More work is required for system-level optical interconnect prediction to explore new solutions that benefit from advances both at the architectural and at the technological level.

10. REFERENCES

[1] M. Amann, M. Ortsiefer, and R. Shau. Surface-emitting laser diodes for telecommunications. In *Proc. Symp. Opto- and Microelectronic Devices and Circuits*, 2002.

[2] T. Baba. Photonic crystals and microdisk cavities based on GaInAsP-InP system. *IEEE J. Selected Topics in Quantum Electronics*, 3, 1997.

[3] L. Benini and G. De Micheli. Networks on chip: A new SoC paradigm. *IEEE Computer*, 35, 2002.

[4] Y. Cao, T. Sato, D. Sylvester, M. Orchansky, and C. Hu. New paradigm of predictive MOSFET and interconnect modeling for early circuit design. In *Proc. Custom Integrated Circuit Conference*, June 2000.

[5] S. Cho et al. Integrated detectors for embedded optical interconnections on electrical boards, modules and integrated circuits. *IEEE J. Sel. Topics in Quantum Electronics*, 8, 2002.

[6] J. Collet et al. Architectural approach to the role of optics in mono- and multi-processor machines. *Applied Optics*, 39, 2000.

[7] W. Dally and B. Towles. Route packets, not wires: On-chip interconnection networks. In *Proc. 38th Design Automation Conference*, 2001.

[8] A. Filios et al. Transmission performance of a 1.5-μm 2.5-Gb/s directly modulated tunable VCSEL. *IEEE Phot. Tech. Lett.*, 15, 2003.

[9] E. Friedman. Clock distribution networks in synchronous digital integrated circuits. *Proc. IEEE*, 89, 2001.

[10] M. Fujita, A. Sakai, and T. Baba. Ultrasmall and ultralow threshold GaInAsP-InP microdisk injection lasers: Design, fabrication, lasing characteristics and spontaneous emission factor. *IEEE J. Sel. Topics in Quantum Electronics*, 5, 1999.

[11] M. Fujita, R. Ushigome, and T. Baba. Continuous wave lasing in GaInAsP microdisk injection laser with threshold current of 40μA. *IEE Electron. Lett.*, 36, 2000.

[12] K. W. Goossen. Optoelectronic/VLSI. *IEEE Trans. on Advanced Packaging*, 22(4), Nov. 1999.

[13] P. Guerrier and A. Greiner. A generic architecture for on-chip packet-switched interconnections. In *Proc. Design, Automation and Test in Europe*, 2000.

[14] M. Ingels and M. S. J. Steyaert. A 1-Gb/s, 0.7μm CMOS optical receiver with full rail-to-rail output swing. *IEEE J. Solid-State Circuits*, 34(7), July 1999.

[15] A. V. Krishnamoorthy and K. W. Goossen. Optoelectronic-VLSI: Photonics integrated with VLSI circuits. *IEEE J. Sel. Topics in Quantum Electronics*, 4(6), November/December 1998.

[16] K. Lee et al. Fabrication of ultralow-loss Si/SiO_2 waveguides by roughness reduction. *Optics Letters*, 26, 2001.

[17] B. E. Little, S. T. Chu, W. Pan, and Y. Kokubun. Microring resonator arrays for VLSI photonics. *IEEE Phot. Tech. Lett.*, 12(3), Mar. 2000.

[18] J. Liu et al. Ultralow-threshold sapphire substrate-bonded top-emitting 850-nm VCSEL array. *IEEE Phot. Lett.*, 14, 2002.

[19] A. Martinez, F. Cuesta, and J. Marti. Ultrashort 2-D photonic crystal directional couplers. *IEEE Phot. Tech. Lett.*, 15, 2003.

[20] J. Morikuni et al. Improvements to the standard theory for photoreceiver noise. *IEEE J. Lightwave Technology*, 12, 1994.

[21] M. Notomi et al. Structural tuning of guiding modes of line-defect waveguides of silicon-on-insulator photonic crystal slabs. *IEEE J. Quantum Electronics*, 38, 2002.

[22] I. O'Connor, F. Mieyeville, F. Tissafi-Drissi, G. Tosik, and F. Gaffiot. Predictive design space exploration of maximum bandwidth CMOS photoreceiver preamplifiers. In *Proc. IEEE International Conference on Electronics, Circuits and Systems*, 2003.

[23] A. Sakai, T. Fukazawa, and T. Baba. Low loss ultra-small branches in a silicon photonic wire waveguide. *IEICE Tran. Electron.*, E85-C, 2002.

[24] A. Sakai, G. Hara, and T. Baba. Propagation characteristics of ultrahigh-Δ optical waveguide on silicon-on-insulator substrate. *Jpn. J. Appl. Phys. – Part 2*, 40, 2001.

[25] S. Schultz, E. Glytsis, and T. Gaylord. Design, fabrication, and performance of preferential-order volume grating waveguide couplers. *Applied Optics-IP*, 39, 2000.

[26] C. Seassal et al. InP microdisk lasers on silicon wafer: CW room temperature operation at 1.6μm. *IEE Electron. Lett.*, 37, 2001.

[27] Semiconductor Industry Association. International Technology Roadmap for Semiconductors, 2001.

[28] A. Taflove. *Computational electrodynamics - The Finite-Difference Time-Domain Method*. Artech House, 1995.

[29] G. Tosik, F. Gaffiot, Z. Lisik, I. O'Connor, and F. Tissafi-Drissi. Optical versus metallic interconnections for clock distribution networks in new VLSI technologies. In *Proc. IEEE Power and Timing Modeling, Optimisation and Simulation Conference*, 2003.

[30] A. Vachoux, J. M. Bergé, O. Levia, and J. Raillard. *Analog and Mixed Signal Hardware Description Languages*. Kluwer Academic Publishers, 1997.