Chapter 1

ON-CHIP OPTICAL INTERCONNECT FOR LOW-POWER

Ian O'Connor and Frédéric Gaffiot

Laboratory of Electronics, Optoelectronics and Microsystems
Ecole Centrale de Lyon, France

ian.oconnor@ec-lyon.fr, frederic.gaffiot@ec-lyon.fr

Abstract   It is an accepted fact that process scaling and operating frequency both contribute to increasing integrated circuit power dissipation due to interconnect. Extrapolating this trend leads to a red brick wall which only radically different interconnect architectures and/or technologies will be able to overcome. The aim of this chapter is to explain how, by exploiting recent advances in integrated optical devices, optical interconnect within systems on chip can be realised. We describe our vision for heterogeneous integration of a photonic "above-IC" communication layer. Two applications are detailed: clock distribution and data communication using wavelength division multiplexing. For the first application, a design method will be described, enabling quantitative comparisons with electrical clock trees. For the second, more long-term, application, our views will be given on the use of various photonic devices to realize a network on chip that is reconfigurable in terms of the wavelength used.

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

Introduction

In the 2003 edition of the ITRS roadmap [17], the interconnect problem was 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". Continually shrinking feature sizes, higher clock frequencies, and growth in complexity are all negative factors as far as switching charges on metallic interconnect is concerned. Even with low resistance metals such as copper and low dielectric constant materials,

bandwidths for long interconnect will be insufficient for future operating frequencies. Already the use of metal tracks to transport a signal over a chip has a high cost in terms of power: clock distribution for instance requires a significant part (30-50%) of total chip power in high-performance microprocessors.

A promising approach to the interconnect problem is the use of an optical interconnect layer, which could 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. Another important constraint 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 section 1, we will describe the "above-IC" optical technology. Sections 2 and 3 describe an optical clock distribution network and a quantitative electrical-optical power comparison respectively. A proposal for a novel optical network on chip in discussed in section 4.

1.   Optical interconnect technology

Various technological solutions may be proposed for integrating an optical transport layer in a standard CMOS system. In our opinion, the most promising approach makes use of hybrid (3D) integration of the optical layer above a complete CMOS IC, as shown in fig. 1.1. The basic CMOS process remains the same, since the optical layer can be fabricated independently. The weakness of this approach is in the complex electrical link between the CMOS interface circuits and the optical sources (via stack and advanced bonding).

In the system shown in fig. 1.1, a CMOS source driver circuit modulates the current flowing through a biased III-V microsource through a via stack making the electrical connection between the CMOS devices and the optical layer. III-V active devices are chosen in preference to Si-based optical devices for high-speed and high-wavelength operation. The microsource is coupled to the passive waveguide structure, where
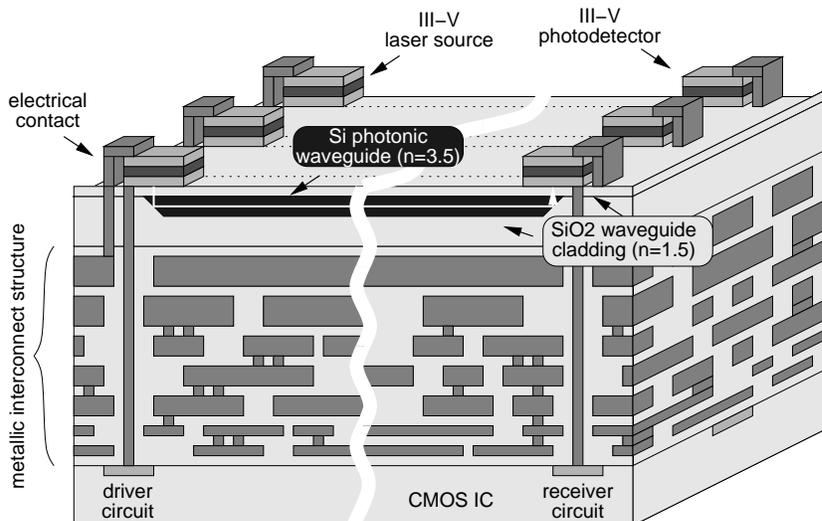
Figure 1.1. Cross-section of hybridised interconnection structure

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 transmitting wavelengths above $1.2\mu m$ (mono-mode waveguiding with attenuation as low as 0.8 dB/cm has been demonstrated [10]). The waveguide structure transports the optical signal to a III-V photodetector (or possibly to several, as in the case of a broadcast function) where it is converted to an electrical photocurrent, which flows through another via stack to a CMOS receiver circuit which regenerates the digital output signal. This signal can then if necessary be distributed over a small zone by a local electrical interconnect network.

2. An optical clock distribution network

In this section we present the structure of the optical clock distribution network, and detail the characteristics of each component part in the system: active optoelectronic devices (external VCSEL source and PIN detector), passive waveguides, interface (driver and receiver) circuits. The latter represent extremely critical parts to the operation of the overall link and require particularly careful design.

An optical clock distribution network, shown in fig. 1.2, 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. Hence the primary tree is optical, while the secondary tree is electrical. It is not feasible to route the optical signal all the way down to the individual gate level since each drop point requires a receiver circuit which consumes area and power. The clock signal is thus routed optically to a number of drop points which will cover a zone over which the last part of the clock distribution will be carried out by the electrical secondary clock tree. The size of the zones is determined by calculating the power required to continue in the optical domain and comparing it to the power required to distribute over the zone in the electrical domain. The number of clock distribution points (64 in the figure) is a particularly crucial parameter in the overall system.

The global optical H-tree was optimised to achieve minimal optical losses by designing the bend radii 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_3 in fig. 1.2) is 625μm, which leads to negligible pure bending loss.
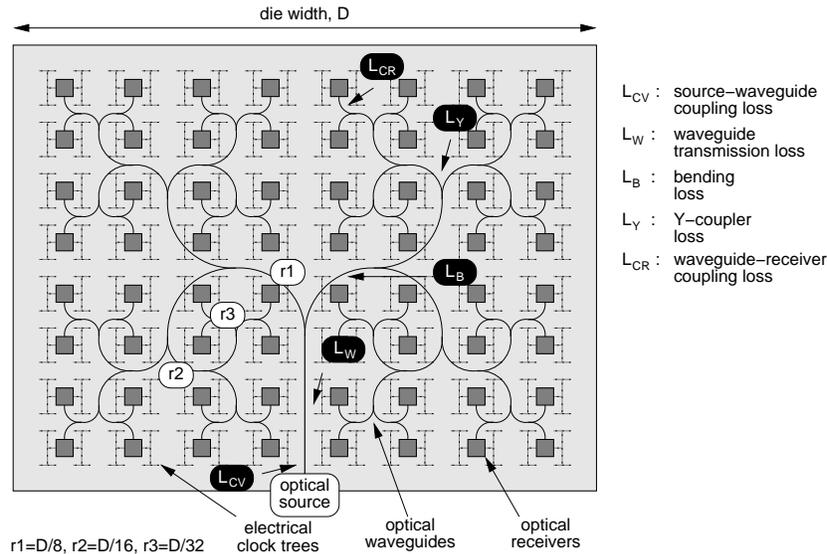

Figure 1.2. Optical H-tree clock distribution network (OCDN) with 64 output nodes. r_{1-3} are the bend radii linked to the chip width D

VCSEL sources

VCSELs (Vertical Cavity Surface Emitting Lasers) are certainly the most mature emitters for on-chip or chip-to-chip interconnections. 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. However, fundamental requirements for integrated semiconductor lasers in optical interconnect applications are small size, low threshold lasing operation and single-mode operation (i.e. only one mode is allowed in the gain spectrum). Additionally, the fact that VCSELs emit light vertically makes coupling less easy. It is clear that significant effort is required from the research community if VCSELs are to compete seriously in the on-chip optical interconnect arena, to increase wavelength, efficiency and threshold current in the same device. Long wavelength, and low-threshold VCSELs are only just beginning to emerge (for example, a 1.5μm, 2.5Gb/s tuneable VCSEL [5], and an 850nm, 70μA threshold current, 2.6μm diameter CMOS compatible VCSEL [11] have been reported). Ultimately however, optical interconnect is more likely to make use of integrated microsources as described in section 4, as these devices are intrinsically better suited to this type of application.

PIN photodetectors

In order to optimise the frequency and power dissipation performance of the overall link, photodetectors must exhibit high quantum efficiency, large intrinsic bandwidth and small parasitic capacitance. The photodetector performance is measured by the bandwidth efficiency product.

Conventional III-V PIN devices suffer from two main limitations. On one hand, their relatively high capacitance per unit area leads to limitations in the design of the transconductance amplifier interface circuit. On the other hand, due to its vertical structure, there is a tradeoff between its frequency performance and its efficiency (the quantum efficiency increases and the bandwidth decreases with the absorption intrinsic layer thickness) [9].

Metal-semiconductor-metal (MSM) photodetectors offer an alternative over conventional PIN photodetectors. An MSM photodetector consists of interdigitated metal contacts on top of an absorption layer. Because of their lateral structure, MSM photodetectors have very high bandwidths due to their low capacitance and the possibility to reduce the carrier transit time. However, the responsivity is usually low compared to PIN photodetectors [4]. MSM photodiodes with bandwidth greater than 100GHz have been reported.

Waveguides

Optical waveguides are at the heart of the optical interconnect concept. In the Si/SiO$_2$ approach, the high relative refractive index differ-

ence $\Delta = (n_1^2 - n_2^2)/2n_1^2$ between the core ($n_1 \approx 3.5$ for Si) and cladding ($n_2 \approx 1.5$ for SiO$_2$) allows 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 [15].

However, the performance of the complete optical system 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 is the sum of losses (in decibels) of all optical components:

$$L_{total} = L_{CV} + L_W + L_B + L_Y + L_{CR} \qquad (1.1)$$

where

- L_{CV} is the coupling coefficient between the photonic source and optical waveguide. 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.
- L_W is the rectangular waveguide transmission loss 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 (fig. 1.3a). For the waveguide fabricated by Lee [10] with roughness of 2nm the calculated transmission loss is 1.3dB/cm.
- L_B is the bending loss, 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. 1.3b, the bending losses associated with a single mode strip waveguide are negligible if the radius of curvature is larger then 3μm.
- L_Y is the Y-coupler loss, and 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 than for low Δ structures and the simulated losses are less then 0.2dB per split [14].
- L_{CR} is the coupling loss from the waveguide to the optical receiver. 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 is assumed to be 87% ($L_{CR} = 0.6dB$) [16].

Figure 1.3a. Simulated transmission loss for varying sidewall roughness in a $0.5\mu m \times 0.2\mu m$ Si/SiO$_2$ strip waveguide

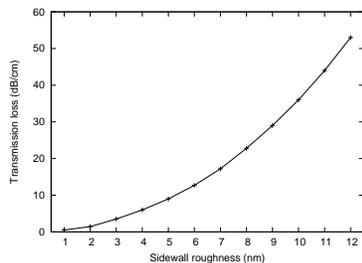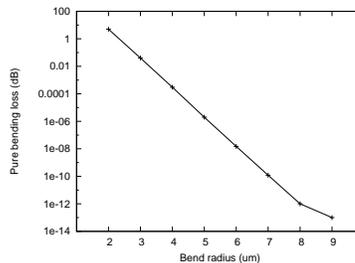*Figure 1.3b.* Simulated pure bending loss for various bend radii in a $0.5\mu m \times 0.2\mu m$ Si/SiO$_2$ strip waveguide

Interface circuits

High-speed CMOS optoelectronic interface circuits are crucial building blocks to the optical interconnect approach. The electrical power dissipation of the link is defined by these circuits, but it is the receiver circuit that poses the most serious design challenges. The power dissipated by the source driver is mainly determined by the source bias current and is therefore device-dependent. On the receiver side however, most of the receiver power is due to the circuit, while only a small fraction is required for the photodetector device.

Driver circuits. Source driver circuits generally use a current modulation scheme for high-speed operation. The source always has to be biased above its threshold current by a MOS current sink to eliminate turn-on delays, which is why low-threshold sources are so important (figures of the order of $40\mu A$ [7] have been reported). A switched current sink modulates 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.

Receiver circuits. A typical structure for a high-speed photoreceiver circuit consists of: a transimpedance amplifier (TIA) to convert the photocurrent of a few μA into a voltage of a few mV; a comparator to generate a rail-to-rail signal; and a data recovery circuit to eliminate jitter from the restored signal. Of these, the TIA is arguably the most critical component for high-speed performance, since it has to cope with a generally large photodiode capacitance situated at its input.

8

The basic transimpedance amplifier structure in a typical configuration is shown in fig. 1.4 [8]. 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 approximation of Butterworth type.

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

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

$$Z_{g0} = -\left(\frac{R_f - R_o/A_v}{1 + 1/A_v}\right) \quad \begin{array}{l} M_f = R_f\,/\,R_o \\ M_i = C_x\,/\,C_y \\ M_m = C_m\,/\,C_y \end{array}$$

$$C_x = C_d + C_i \qquad C_y = C_l + C_o$$

Figure 1.4. CMOS transimpedance amplifier structure

It is then possible to develop a synthesis procedure which, from desired transimpedance performance criteria (gain Z_{g0}, bandwidth and pole quality factor Q) and operating conditions (photodiode and load capacitances, C_d and C_l respectively) generates component values for the feedback resistance R_f and the voltage amplifer (voltage gain A_v and output resistance R_o). Circuits with high R_o/A_v ratio ($\approx 1/\sum g_m$) require the least quiescent current and area and this quantity constitutes therefore an important figure of merit in design space exploration (fig. 1.5a). To reach a sized transistor-level circuit, approximate equations for the small-signal characteristics and bias conditions of the circuit are sufficient to allow a first-cut sizing of the amplifier, which can then be fine-tuned by numerical or manual optimisation, using simulation for exact results. The complete process is described in [13].

Using this methodology with industrial transistor models for technology nodes from 350nm to 180nm and predictive BSIM3v3/BSIM4 models for technology nodes from 130nm down to 45nm [3], we generated design parameters for $1THz\Omega$ transimpedance amplifiers to evaluate the evolution in critical characteristics with technology node. Fig. 1.5b shows the results of transistor level simulation of fully generated photoreceiver circuits at each technology node.

Figure 1.5a. TIA R_o/A_v design space with varying bandwidth and transimpedance gain requirements
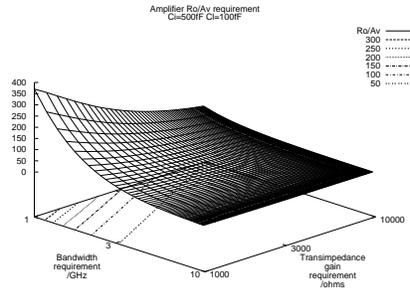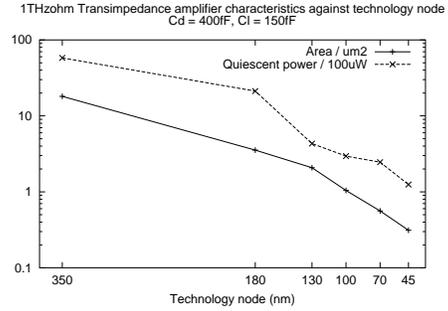
Figure 1.5b. Evolution of TIA characteristics (power, area, noise) with technology node

3. Quantitative power comparison between electrical and optical clock distribution networks

Design methodology

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. 1.6.
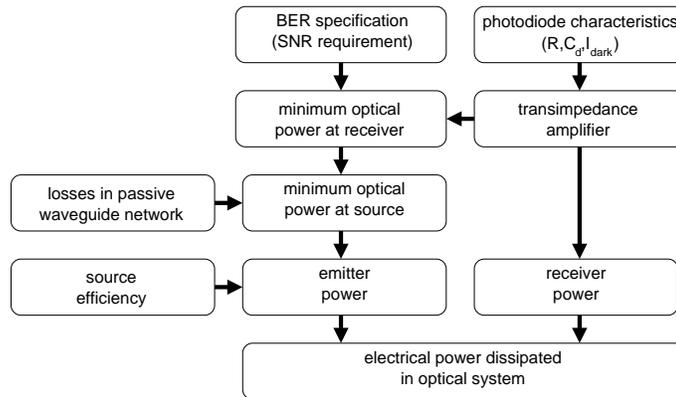
Figure 1.6. Methodology used to estimate the electrical power dissipation in an optical clock distribution network

The first criterion for defining the performance of the optoelectronic link is the required signal transmission quality, represented by the bit error rate (BER) and directly linked to the photoreceiver signal to noise ratio. For an on-chip interconnect network, a BER of 10^{-15} is acceptable. To calculate the required signal power at the receiver, the characteristics of the receiver circuit have to be extracted from the transistor-level schematic, which is generated from the photodetector characteristics (responsivity R, C_d, dark current I_{dark}) and from the required operating frequency using the method described in section 2. For the given BER and for the noise signal associated with the photodiode and transimpedance circuit the minimum optical power required by the receiver to operate at the given error probability can be calculated using the Morikuni formula [12].

With this figure, and knowing the layout and therefore the optical losses that will be incurred in the waveguides, the minimum required optical power at the source can be estimated. The total 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, laser light-current characteristics given by Amann [1] were used.

Design performance

Our aim in this work was to quantitatively compare the power dissipation in electrical and optical clock distribution networks for a number of cases, including technology node prediction. For both electrical and optical cases we used technology parameters from the ITRS roadmap (wire geometry, material parameters). For transistor models we used predictive model parameters from Berkeley (BSIM3V3 down to 70nm and BSIM4 down to 45nm). 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 [18]. For optical performance predictions we used existing technology characteristics while for the optoelectronic devices we took datasheets from two real devices and used these figures.

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 tables 1.1a and 1.1b.

Table 1.1a. Electrical CDN characteristics

Electrical system parameter	
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

Table 1.1b. Optical CDN characteristics

Optical system parameter	
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 coefficient (%)	50
Photodiode capacitance (fF)	100
Photodiode responsivity (A/W)	0.95

What follows is the results of comparisons of the power dissipation in electrical and optical clock distribution networks. This was quantitatively carried out for varying chip size, operating frequency, number of clock distribution points, technology node, and finally sidewall roughness. This latter performance characteristic is the only non system-driven characteristic, but it gives some important design information to technology groups working on optical interconnect.

Fig. 1.7a shows a power comparison where we vary square die size from 10 to 37 mm width. This analysis was carried out for the 70nm node at a distribution frequency of 5.6GHz (which is the clock frequency associated with this node) and 256 drop points. Electrical CDN power rises almost linearly with die size, which is understandable since the line lengths increase and therefore require more buffers to drive them. Optical CDN power rises much more slowly since all that is really changing is transmission loss and this has a quite minor effect on the overall power dissipation.

When we vary clock frequency for constant chip width, fig. 1.7b we observe a similar effect for the electrical CDN. Again, the number of buffers has to increase since the segment lengths have to be reduced in order to attain the lower RC time constants. For the optical CDN, what is changing is the receiver power dissipation. The transimpedance amplifier requires a lower output resistance in order to operate at higher frequencies and this translates to a higher bias current.

In fig. 1.7c, we vary the number of drop points and see that both electrical and optical CDN power dissipation rises, but optical rises much

faster than electrical. There are two reasons for this: firstly, every time the number of drop points is doubled, so is the number of receivers and this accounts for a large part of the power dissipation; secondly, the number of splitters is doubled, which in turn means that the power at emission also has to be doubled. These two factors cause the optical power to catch up with the electrical power at around 4000 drop points.

Fig. 1.7e shows a comparison for varying technology node. Not only the technology is changing here, we are also changing the clock frequency associated with the node. We can see that at the 70nm node there is a five-fold difference between electrical and optical clock distribution. As the technology node advances, this difference becomes even more marked.

A final analysis, fig. 1.7f, shows how technological advances are required to improve system performance, concerning in this case waveguide sidewall roughness. 5nm roughness translates to a transmission loss of around 8dB/cm, which in turn corresponds to a power dissipation figure of around 500mW for the 70nm node at 5.6GHz and 20mm chip width. Looking at the 2nm roughness point, achieved at MIT [10] and corresponding to a transmission loss of 1.3dB/cm, we obtain a power dissipation figure of about 10mW, a fifty-fold decrease in the overall power dissipation by going from 5nm roughness to 2nm roughness. This demonstrates the importance of optimising the passive waveguide technology for the whole system.

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 1.2.

Table 1.2. Optical power budget for 20mm die width at 3GHz

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

We can conclude from this analysis that power dissipation in optical clock distribution networks is lower than that of electrical clock distribution networks, by a factor of five for example at the 70nm technology

Figure 1.7a. Comparison of power dissipation in electrical and optical clock distribution networks for varying chip size (70nm technology, 5.6GHz, 256 drop points)
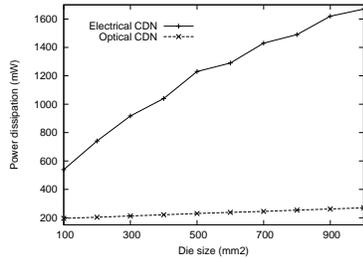

Figure 1.7b. Comparison of power dissipation in electrical and optical clock distribution networks for varying clock frequency (70nm technology, 400mm^2, 256 drop points)
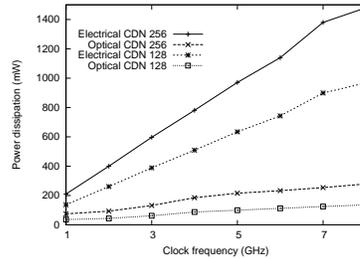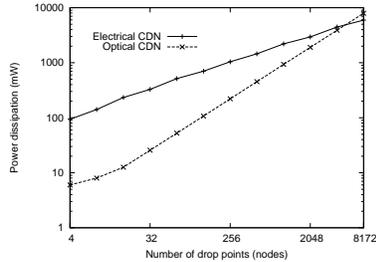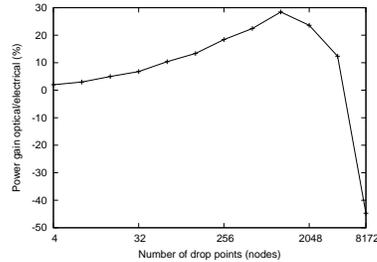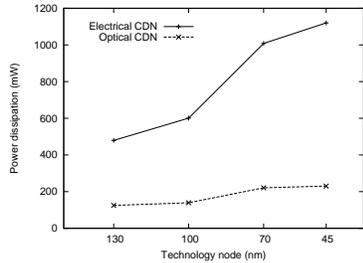

Figure 1.7c. Comparison of power dissipation in electrical and optical clock distribution networks for varying number of drop points (70nm technology, 5.6GHz, 400mm^2)

Figure 1.7d. Comparison of power dissipation in electrical and optical clock distribution networks for varying number of drop points (70nm technology, 5.6GHz, 400mm^2)

Figure 1.7e. Comparison of power dissipation in electrical and optical clock distribution networks for varying technology nodes
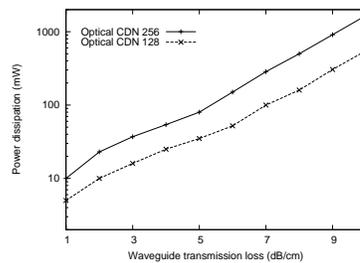

Figure 1.7f. Evaluation of power dissipation in optical clock distribution networks for varying waveguide sidewall roughness (70nm technology, 5.6GHz, 400mm^2)

node. This factor will in the future become larger due to two reasons: firstly due to improvements in optical fabrication technology; and secondly with the rise in operating frequencies. However, this figure is prob-

ably not sufficient to convince semiconductor manufacturers to introduce such large technological and methodological changes for this application. To improve the figure, weak points can be identified for each main part of an integrated optical link. For the source, the efficiency between electrical and optical power conversion is relatively low. This needs to be improved and one area is possibly in integrated microsources. For the waveguide structures, most of the losses need to be improved, especially transmission loss and coupling loss. Sidewall roughness especially has a direct and considerable impact on the power dissipation in the global system. Finally at the receiver end, the transimpedance amplifier power dissipation is too high. Better circuit structures must be devised, or the photodetector parasitic capacitance needs to be reduced.

4.   Optical network on chip

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 Tb/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).

Using wavelength division multiplexing (WDM) techniques, photonics and optoelectronics may offer new solutions to realise reconfigurable optical networks on chip (ONoC). An ONoC, as an electronic router with routing based on wavelength λ, is actually a circuit-switching based topology and can thus ensure data exchanges between IP blocks with very low contention. The advantages of using an optical network are many: independence of interconnect performance from distance and data rate, crosstalk reduction, connectivity increase, interconnect power dissipation reduction, increase in the size of isochronous tiles, use of communication protocols. Figure 1.8 shows a 4×4 ONoC with all electronic interfaces: photodetector and laser in III-V technology and optical network in SOI technology, using similar heterogeneous integration techniques as described in section 1. Intellectual property (IP) blocks shown can be processor cores, memory blocks, functional units etc. with standard interfaces to the communication network. This is a multi-domain device with high speed optoelectronic circuits (modulation of the laser current and photodetectors) and passive optics (waveguides and pas-

sive filters). In the figure, 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 figure 1.3. 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 network is in principle scalable to an infinite number of connections. In practice, this number is severely limited by lithography and etching precision. For a 5nm tolerance on the size of the microdisk, corresponding to state of the art CMOS process technology, the maximum size of the network is 8×8.
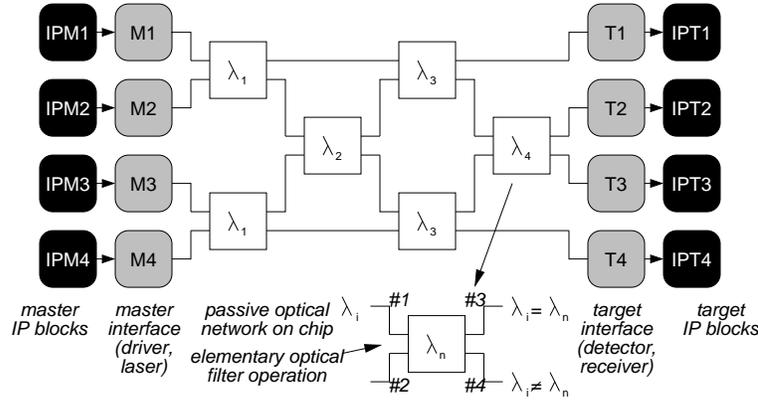

Figure 1.8. Architecture of 4x4 optical network on chip

Table 1.3. Truth table for 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

The basic element of the network is an optical filter, described in the next section. The ports $\sharp 1 - \sharp 4$ correspond to inputs/outputs of the optical filter. Its operation is the same as an electronic cross-bar: the cross function (output in $\sharp 4$) is activated when the injected wavelength in $\sharp 1$ does not correspond to a resonant ring wavelength and the bar

function is activated (output in ♯3) when the injected wavelength in ♯1 corresponds to a resonant ring wavelength. Operation is symmetrical: the same phenomena happens if the wavelength injection is placed in the port ♯4.

Microresonators

Microring resonators are ideal device candidates for integrated photonic circuits. Because they render possible the addition or extraction of signals from a waveguide based on wavelength in a WDM flow, they can be considered as basic building blocks to build complex communication networks. The use of standard SOI technology leads to high compactness (structures with radii as small as $4\mu m$ have been reported) and the possibility of low-cost photonic integration. Figure 1.9 shows the structure of an elementary add-drop filter based on microring resonators. The size of the structure is typically a few hundred μm^2. It consists of two identical disks evanescently side-coupled to two signal waveguides which are crossed at near right angles to facilitate signal directivity. The microdisks make up a selective structure: the electromagnetic field propagates in the rings for discrete propagation modes corresponding to specific wavelengths. The resonant wavelengths depend on geometric and structural parameters (indices of the substrate and of the microrings, thickness and diameter of the disks).
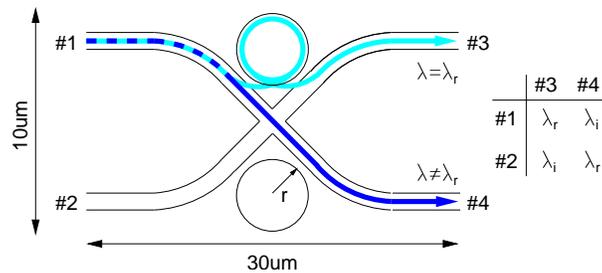

Figure 1.9.   Micro-disk realisation of an add-drop filter

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 (for example injected at port ♯1) is close enough to a resonant wavelength λ_1 (tolerance is of the order of a few nm, depending on the coupling strength between the disk and the waveguide), then the electromagnetic field is coupled into the microrings and then out along the second waveguide

(in the example, the optical signal is transmitted to the output port ♯3, as shown in fig. 1.10a). If the wavelength of the optical signal does not correspond to the resonant wavelength, then the electromagnetic field continues to propagate along the waveguide and not through the structure (in the example, the optical signal would then be transmitted to the output port ♯4, as shown in fig. 1.10b). This device thus operates as an elementary router, the behaviour of which is summarised in the table in fig. 1.9.

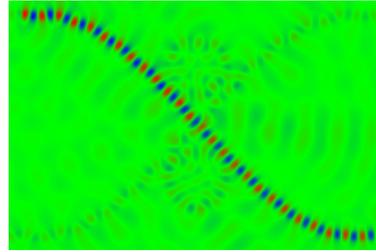*Figure 1.10a.* FDTD simulation of add-drop filter in on-state
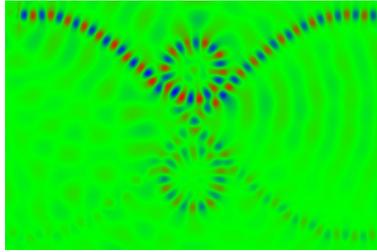
Figure 1.10b. FDTD simulation of add-drop filter in off-state

First structures have been realised and preliminary results are promising. Fig. 1.11a shows an IR photograph of the structure in the cross state (top) and in the bar state (bottom), while fig. 1.11b represents the transmission coefficient on the cross output: the transmitted power on the cross output reaches 100% for wavelengths corresponding to the resonant frequencies of the microdisk.

Figure 1.11a. Infra-red photograph of structure in both cross (top) and bar (bottom) states
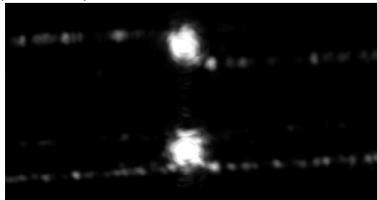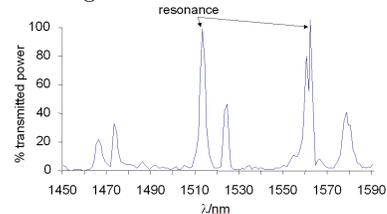

Figure 1.11b. Transmission coefficient on cross output for varying wavelength

Microsource lasers

From the viewpoint of mode field confinement and mirror reflection, microdisk lasers operate on the principle of total internal reflection, as opposed to multiple reflection, as is the case in VCSELs for example.

This fact gives this type of source two distinct advantages over VCSELs for on-chip optical interconnect. Firstly, light emission is in-plane (as opposed to vertical), meaning that emitted light can be injected directly into a waveguide with minimum loss [6]. Secondly, for communication schemes requiring multiple wavelengths, it is easier from a technological point of view to control the radius of such a device than it is to control the thickness of an air gap in a VCSEL. In any case 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. Although these devices are not as mature as VCSELs, they seem extremely promising for optical interconnect applications. An overview of microcavity semiconductor lasers can be found in [2].

Demonstration of principle

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 ♯1 at the same moment (shown in figure 1.12). The input signal format is a matrix. Figure 1.12 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 4 outputs T_1, T_2, T_3 and T_4 (T_2 shown in fig. 1.12). As predicted in table 1.3, only λ_3 is detected at the output T_2.
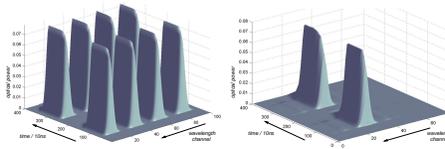

Figure 1.12. Simulation of 4x4 optical network on chip

5. 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 chapter, 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. The feasibility of on-chip optical interconnect is no longer really in doubt. We have given

some partial results to quantitavely demonstrate the advantages of optical clock distribution. Although lower power can be achieved (of the order of a five-fold decrease), more work is required to explore new solutions that benefit from advances both at the architectural and at the technological level. Also the existing basic building blocks need to be integrated together to physically demonstrate on-chip optical links. Research is well under way in several research groups around the world to do this. Looking further ahead, the use of multiple wavelengths in on-chip communication networks and in reconfigurable computing is an extremely promising and exciting field of research.

References

[1] M. Amann, M. Ortsiefer, and R. Shau: 2002, 'Surface-emitting Laser Diodes for Telecommunications'. In: *Proc. Symp. Opto- and Microelectronic Devices and Circuits.*

[2] T. Baba: 1997, 'Photonic Crystals and Microdisk Cavities Based on GaInAsP-InP System'. *IEEE J. Selected Topics in Quantum Electronics* **3**.

[3] Y. Cao, T. Sato, D. Sylvester, M. Orchansky, and C. Hu: 2000, 'New Paradigm of Predictive MOSFET and Interconnect Modeling for Early Circuit Design'. In: *Proc. Custom Integrated Circuit Conference.*

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

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

[6] M. Fujita, A. Sakai, and T. Baba: 1999, '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**.

[7] M. Fujita, R. Ushigome, and T. Baba: 2000, 'Continuous wave lasing in GaInAsP microdisk injection laser with threshold current of $40\mu A$'. *IEE Electron. Lett.* **36**.

[8] M. Ingels and M. S. J. Steyaert: 1999, 'A 1-Gb/s, 0.7μm CMOS Optical Receiver with Full Rail-to-Rail Output Swing'. *IEEE J. Solid-State Circuits* **34**(7).

[9] I. Kimukin et al.: 2002, 'InGaAs-Based High-Performance p-i-n Photodiodes'. *IEEE Phot. Tech. Lett.* **26**(3).

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

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

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

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

[14] A. Sakai, T. Fukazawa, and T. Baba: 2002, 'Low Loss Ultra-Small Branches in a Silicon Photonic Wire Waveguide'. *IEICE Tran. Electron.* **E85-C**.

[15] A. Sakai, G. Hara, and T. Baba: 2001, 'Propagation Characteristics of Ultrahigh-Δ Optical Waveguide on Silicon-on-Insulator Substrate'. *Jpn. J. Appl. Phys. – Part 2* **40**.

[16] S. Schultz, E. Glytsis, and T. Gaylord: 2000, 'Design, Fabrication, and Performance of Preferential-Order Volume Grating Waveguide Couplers'. *Applied Optics-IP* **39**.

[17] Semiconductor Industry Association: 2003, 'International Technology Roadmap for Semiconductors'.

[18] G. Tosik, F. Gaffiot, Z. Lisik, I. O'Connor, and F. Tissafi-Drissi: 2004, 'Power dissipation in optical and metallic clock distribution networks in new VLSI technologies'. *IEE Elec. Lett.* **4**(3).