5
ADVANCED RESEARCH IN ON-CHIP OPTICAL INTERCONNECTS

Ian O'Connor and Frédéric Gaffiot,
Ecole Centrale de Lyon - Laboratory of Electronics, Optoelectronics and Microsystems

5.1 The Interconnect Problem

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. The 2001 ITRS [1] (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. Timing is already no longer the sole concern with physical layout: power consumption, crosstalk and voltage drop drastically increase the complexity of the tradeoff problem. With decreasing device dimensions it is increasingly difficult to keep wire propagation delays acceptable. While dielectric constants below 2 (around 1.7-1.8) can be achieved through the use of nanoporous SiOC-like or organic (SilK type) materials with an "air gap" integration approach, integration complexity is higher and mechanical properties are weaker. In addition, the use of ultra low-k materials is physically limited by the fact that no material permittivity can be less than 1, that of air. Thus, even with the most optimistic estimates for RC time constants using low-resistance metals such as copper and low-k dielectrics, global interconnect performance required for future generations of ICs cannot be achieved with metal. Furthermore, since IC power dissipation is strongly linked to switching frequency, tomorrow's architectures will require power over the 100W mark to be able to operate in the 10GHz range and above. At this level, thermal problems will jeopardize system performance if not strictly controlled.

Analysis of electrical interconnect performance

The overall device scaling factor *s* makes it possible to determine the performance of interconnects. Each process shrink has a large impact on electrical parameters of metallic interconnections. Before deep submicron nodes (the threshold is widely accepted as being around the 0.35µm technology node), the gate delay was higher than the interconnect delay, such that each shrink led to an improvement of the maximum working speed of a system by a factor of *1/s*. In the present DSM era however, global interconnect delay has become larger than gate delay and consequently, interconnect has become the dominant factor determining speed.

Sakurai's equation [2] gives the propagation delay of a signal transmitted from an emitter gate to a receiver gate.

$$t_d = R_{out}\left(C_{out} + C_L\right) + R_{out} c l_W + 0{,}4\left[\left(crl_W^2\right)^{1,6} + \tau_{tof}^{1,6}\right]^{1/1,6} + 0{,}7 r l_w C_L \qquad (1)$$

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

In the case of local and intermediate interconnects, this equation reduces to:

$$t_d = R_{out}\left(C_{out} + C_L\right) + R_{out} c l_W \quad (2)$$

Equation (2) shows that the delay time is a combination of the gate output resistance, interconnect and load capacitances. Gate sizing makes it possible to reduce delay by increasing gate strength, at the cost of increased area and power consumption.

In the case of global links, Sakurai's formula shows that the delay time in the line becomes predominant. In order to limit the delay time in the metallic line, global links are routed on the upper metal layers where it is possible to increase the width and the thickness of the line and thus to reduce the lineic resistance. Reverse scaling (by reducing the thickness of the metal layer less than the scale factor) is commonplace, leading to high aspect ratios. Gate sizing makes it possible to minimize t_d, and it is possible to show that t_d varies with l_w^2. This increase of the delay time with the second power of the line length cannot be avoided. Repeater insertion makes it possible to make the delay vary with l_w, but this of course comes at the cost of a very large number of repeaters. In this scenario therefore, a relatively high percentage of silicon real estate and IC power consumption is devoted to interconnect rather than to data processing functions.

The problem facing us then is that evolutionary solutions will not be sufficient to meet the performance roadmap. To tackle the issues developed above, radically different interconnect approaches displaying a highly improved data-rate to power ratio must be developed. At present, the most prominent ideas are the use of integrated radio frequency or microwave interconnects [3], 3D (non-planar) integration [4] and optical interconnects [5]. This chapter focuses on the latter concept.

The optical alternative

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 the 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. Difficulties expected are obtaining a large enough optical-electrical conversion efficiency, reducing the optical transmission losses while allowing for a sufficient density of photonic waveguides on the circuit and reduction of the latency while operating above the 10GHz

mark. Sections 5.3, 5.4 and 5.5 describe, respectively, the issues involved in photonic waveguides, active devices and optoelectronic conversion circuits.

Identified applications

Optical links can be categorized into three broad domains, for which various analyses have been carried out and applications identified: single wavelength point-to-point (1-1 link); single and multiple wavelength broadcast (1-n link); multiple wavelength bus and switching (n-n link). The latter category is rather new and will be discussed in section 5.8.

Point-to-point (1-1) links

Typically on today's more complex chips, hundreds or thousands of global links are necessary [6]. 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 analyzing the benefits of introducing optical interconnect in critical data-intensive links, such as CPU-memory buses in processor architectures [7]. These analyses showed that point to point links do not present a sufficiently high performance gain to warrant their widespread use in future technologies. In essence, the bandwidth/power ratio for point-to-point optical links is higher than the electrical counterpart, but not high enough, when interface circuit power is taken into consideration. Instead, it is preferable to apply architectural modifications in order to enable bottlenecks to be overcome (in the given example application, the solution was to add more cache memory), even at the expense of greater silicon area and power. The benefits of optical interconnect in terms of physical cost in this situation are, in the long run, not viable for industrial manufacturers since the entire manufacturing process (from design to fabrication) would have to be changed: a very costly course of action. This proves that 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.

Broadcast (1-n) links

Another and potentially more profitable application of optical interconnect technology is in clock distribution networks [8] (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. An electrical alternative is global clock distribution at a relatively low frequency and local clock multiplication to generate the required clock speed. Disadvantages of this approach include inter-zone synchronization and clock multiplication lock time. 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. However, it would be illusory to believe that the optical clock signal could be routed down to the single-gate level: optoelectronic interface circuits are of course necessary and consume power. An example system realizing a clock distribution function, shown in Fig. 5.1, 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 is a particularly crucial parameter in the overall system.

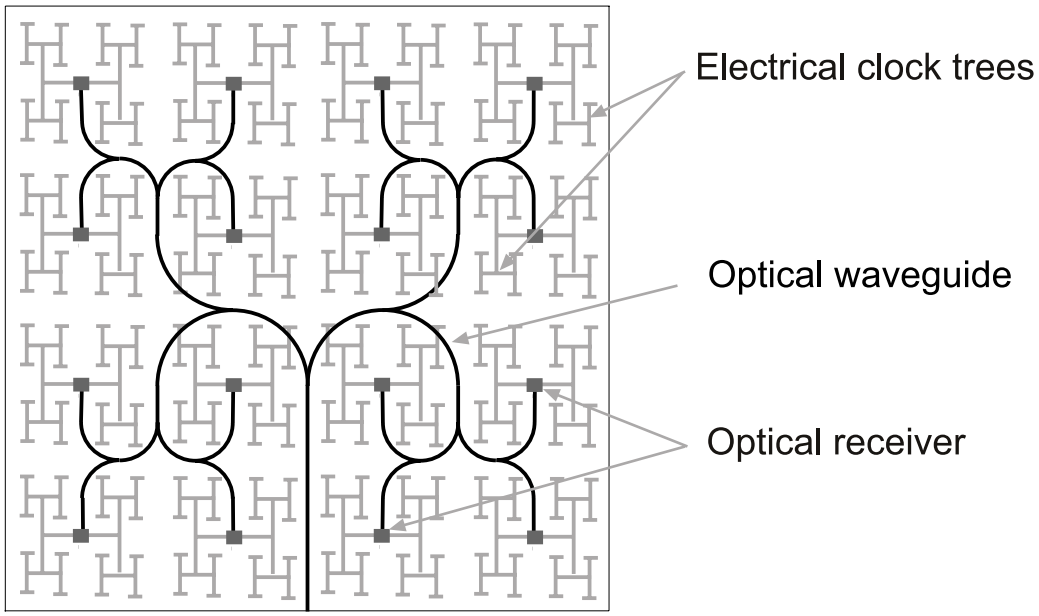

Figure 5.1. 1-16 point optical clock distribution tree

5.2 Top-Down Link Design

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.

Materials

Materials have to be chosen with different constraints:

- *efficient light detection:* obviously, the active devices are of fundamental importance to the power budget of the link. Optics is suitable only if, for a given throughput, the global power consumption of the whole link is lower than the power consumption of classical metallic links. The quantum efficiency of the active devices is of prime importance in this context. 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 mature existing technologies (to ensure the required reproducibility, and homogeneity of the device parameters)
- *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 realization of the optical passive guides but we focus here on silicon/silica wave guides. Silicon is an excellent material for wavelengths above 1.2μm, and monomode waveguiding with attenuation as low as 0.8 dB/cm has been proven [9]. Moreover, the high refractive index difference between silicon and silica makes it possible to realize passive structures with dimensions compatible with DSM technologies (for example, it is possible to realize monomode waveguides less than 1μm wide). The realization of silicon/silica waveguides is (at least in principle) compatible with a standard CMOS process. The choice of silicon waveguides leads to the use of wavelengths greater than 1.2 μm. To capitalize on the maturity of devices and concepts developed by the telecommunications industry, the choice of the wavelength is in practice limited to 1.3-1.55μm windows.

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.

The integration of silicon waveguides at the front end of the CMOS process (i.e. before fabrication of the metallic interconnection layers) is certainly possible but other

considerations have to be taken into account. At the transistor level, the routing of the waveguide is extremely difficult and requires routing space at the IC level. Further, the problem of the active devices remains: silicon-based sources cannot yet (and for the foreseeable future) be considered to be mature, while the growth of III-V devices on silicon faces strong technological barriers. The use of external sources and detectors bonded by flip-chip is unrealistic due to the high number of individual bonding operations required. Also, silicon-based devices can only work at low wavelengths (850nm), which translates to higher attenuation in the waveguides. This solution requires an extraordinary mutation in the CMOS process, and as such is highly unattractive from an economic point of view.

Hybrid integration of the optical layer on top of a complete CMOS IC is much more practical, and with more scope for evolving. The source and detector devices are no longer bound to be realized in the host material. Fig. 5.2 shows a cross section of how a complete "above IC" photonic layer could be realized. 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.

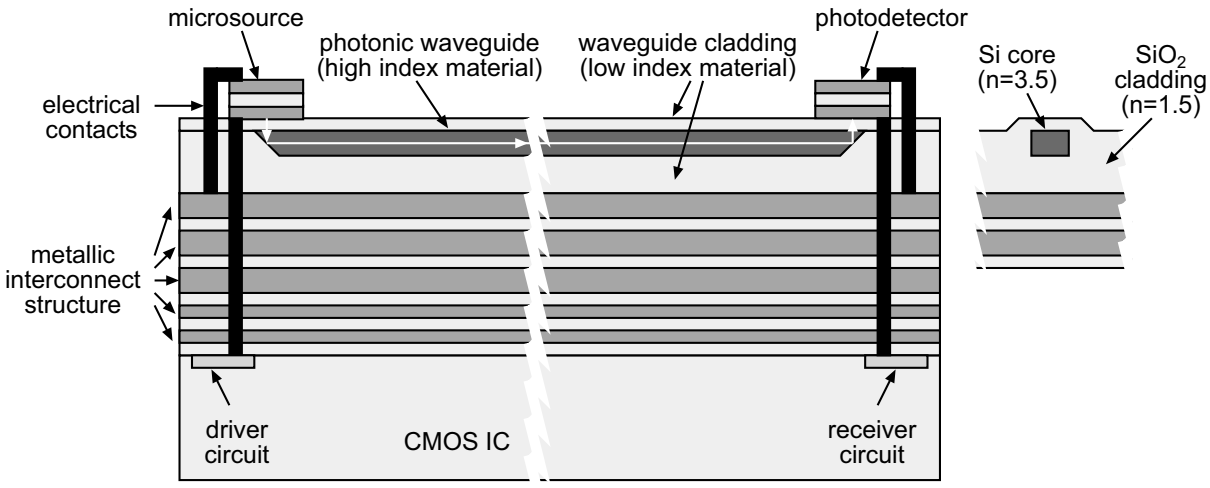

Figure 5.2. Cross section of hybridized interconnection structure

In the system shown in Fig. 5.2, 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.

To form a planar optical waveguide, silicon is used as the core and SiO_2 as the cladding material. Si/SiO_2 structures are compatible with conventional silicon technology and transparent for 1.3-1.55μm wavelengths. Such waveguides with high relative refractive index difference $\Delta \approx (n_1^2 - n_2^2)/2n_1^2$ between the core ($n_1 \approx 3.5$ for Si) and claddings ($n_2 \approx 1.5$ for SiO_2) allow the realization of a compact optical circuit, with bend radius of the order of a few μm [10]. To avoid modal dispersion, improve coupling efficiency and reduce loss, single mode conditions are applied to the waveguide dimensions. For a wavelength of 1.55μm, this means a waveguide width of 0.3μm.

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.

Design requirements

To make a reasoned comparison between electrical and optical interconnect, a set of design requirements must be established. We have already mentioned BER; we must add to this 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.

In order to be able to evaluate and optimize 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.

5.3 Passive Photonic Devices for Signal Routing

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. 5.3) is the sum of losses (in decibels) of all optical components:

$$L_{total} = L_{CV} + L_B + L_Y + L_{CR} \tag{3}$$

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. To provide an unambiguous comparison in terms of dissipated power between optical and electrical on-chip interconnect networks it is necessary to incorporate all of these quantities.

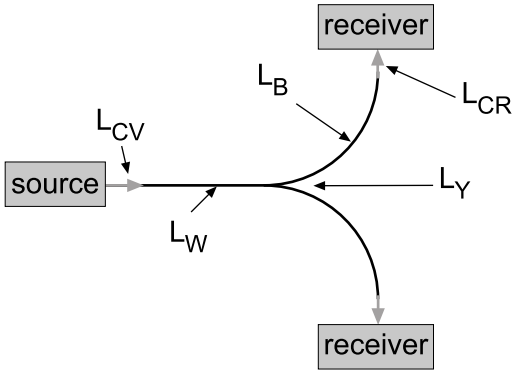

Figure 5.3. Losses in an optical link

In the present technology, there are several methods to couple the beam emitted from the laser into the optical waveguide. In the proposed system we assumed 50% coupling efficiency L_{CV} from the source to a single mode waveguide.

Transmission loss L_W describes the attenuation rate of the optical power, as light travels in the waveguide. 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 losses. To calculate the attenuation coefficient we used the Payne formula [11] associated with the Effective Index Method [12] [13]. For the waveguide fabricated by Lee [9] 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. For low Δ, the bending loss is very high, which prevents increasing the packing density. 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 realized. To assess the bending loss L_B we use the Marcuse method [14]. As can be seen from Fig. 5.4, the bending losses associated with a single mode strip waveguide are negligible if the radius of curvature is bigger then 2μm.

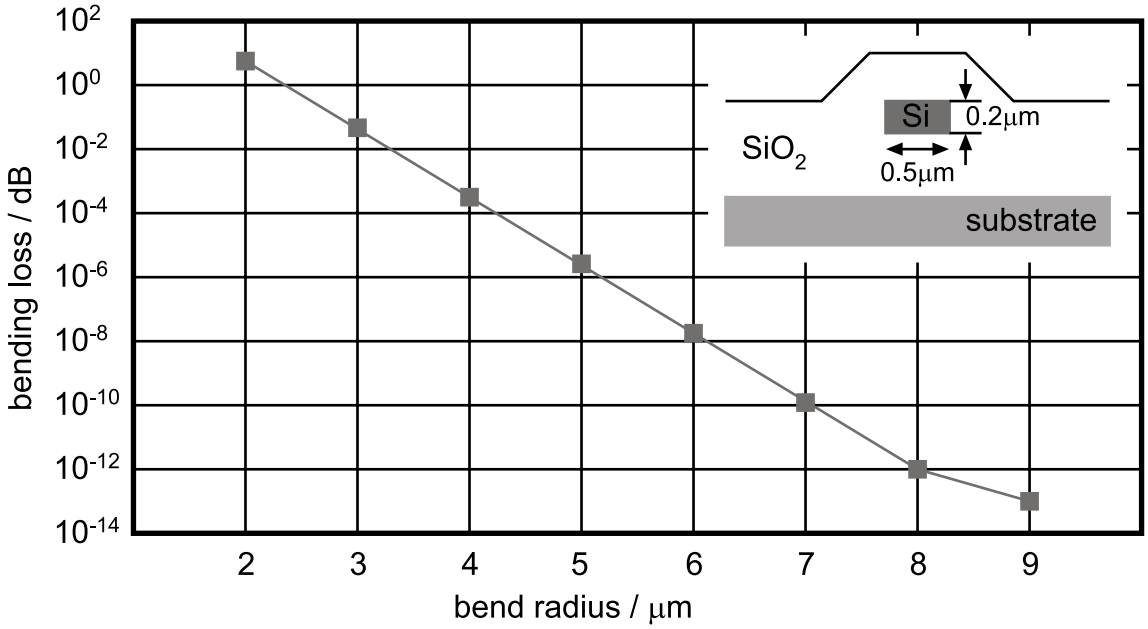

Figure 5.4. Simulated bending loss for Si/SiO_2 strip waveguide

The Y-junction loss L_Y depends on the reflection and scattering attenuation into the propagation path and surrounding medium. Different Y-branch structures have been analyzed by several methods [15] [16]. 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 [17].

Using currently available materials and methods it is possible to achieve an almost 100% coupling efficiency from waveguide to optical receiver. In the proposed system the coupling efficiency L_{CR} from the waveguide to the optical receiver is assumed to be 87% [18].

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. 5.5) 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).

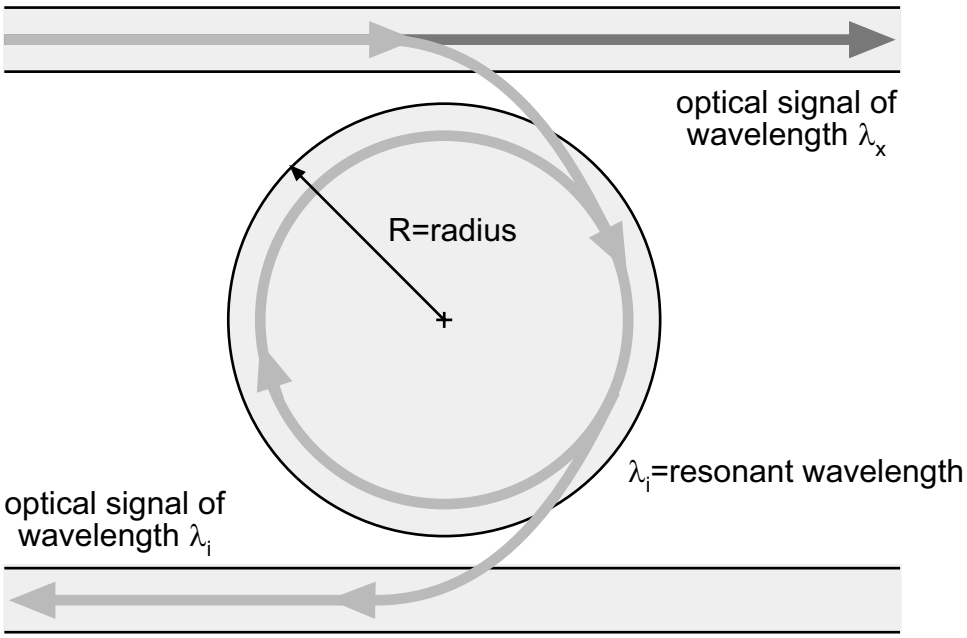

Figure 5.5. Micro-disk realization of an add-drop filter

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 NxN switching networks have been devised [19], 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.

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 realized 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. Other functions can be created using photonic crystals, such as couplers [20], 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.

5.4 Active Devices for Signal Conversion

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 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 [22].

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, by stimulated emission via a current above a few microamperes. The active layer is formed by multiple quantum wells surrounded by III-V compound materials, and the whole forms the optical cavity of the desired wavelength. Above and below are Bragg reflectors, with deep proton implant to confine the current injected via the anode. 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 [23], and an 850nm, 70μA threshold current, 2.6μm diameter CMOS compatible VCSEL [24] have been reported).

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. 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 will be covered in section 5.6.
The structure of a microdisk laser is shown in Fig. 5.6 [25]. The active region of the disk is supported by upper and lower posts. Small cavity volume and strong optical confinement through semiconductor/air boundaries leads to low threshold currents. Current injection via the top contact causes carriers to diffuse to the disk edge and consequently produce optical gain. Lasing oscillation is generated by "whispering gallery" modes (so-called because of how the energy is distributed) rounding inside the disk edge. These modes, defined by the disk radius and representing the emission wavelengths, can be calculated using FDTD (Finite Difference Time Domain) simulations.

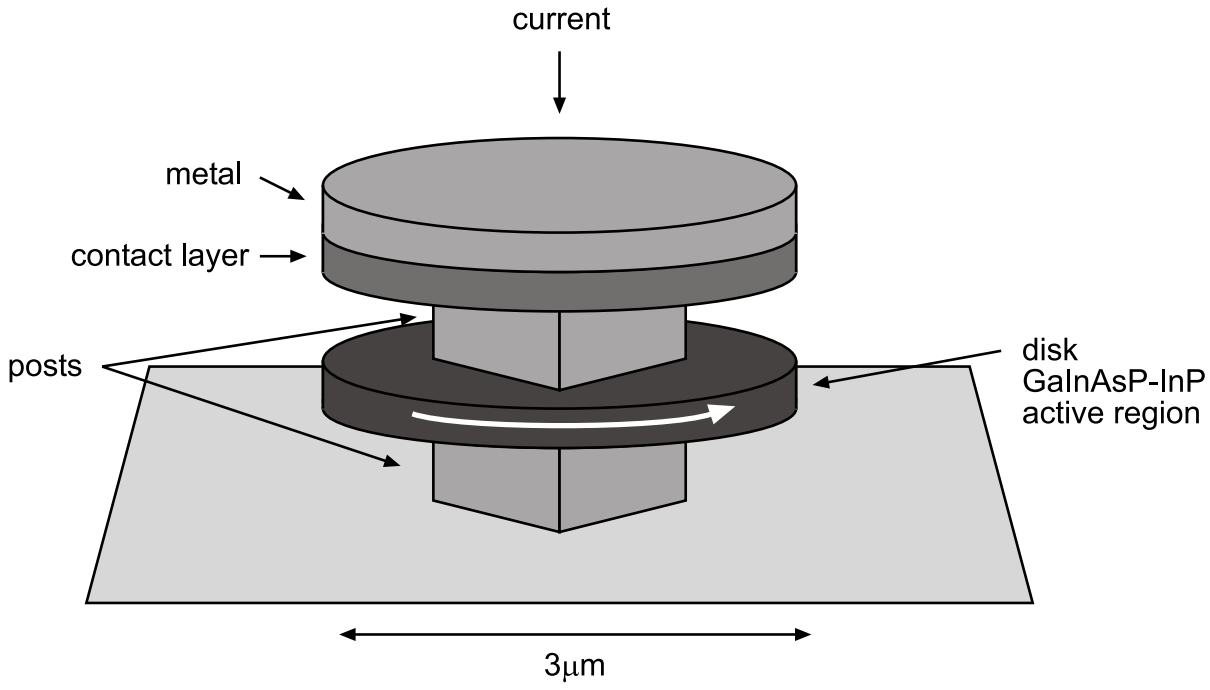

Figure 5.6. Structure of a microdisk laser

Photonic crystals are also capable of being used as microsources. While they are potentially smaller than microdisks and with better control of emission directivity and coupling, their mode behavior is complex and difficult to evaluate, and structures designed for lightwave frequencies are difficult to fabricate and to characterize. Research is ongoing in this area.

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 [26].

5.5 Conversion 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. The objective therefore is to attain the maximum speed/power ratio using dedicated circuit design methodologies.

Driver circuits

The basic current modulation configuration of the source driver circuit is shown in Fig. 5.7. 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 [27] 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.

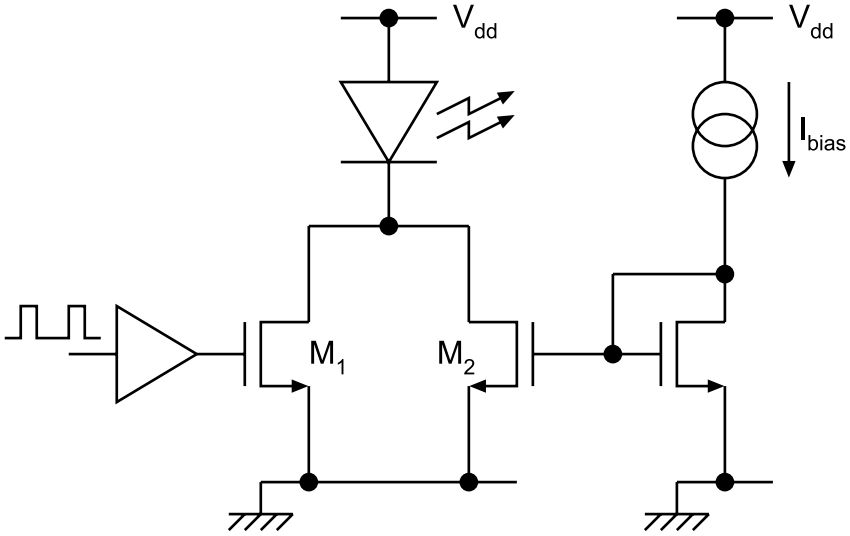

Figure 5.7. Basic current modulation source driver circuit

Receiver circuits

The classical structure for a receiver circuit is shown in Fig. 5.8: 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.

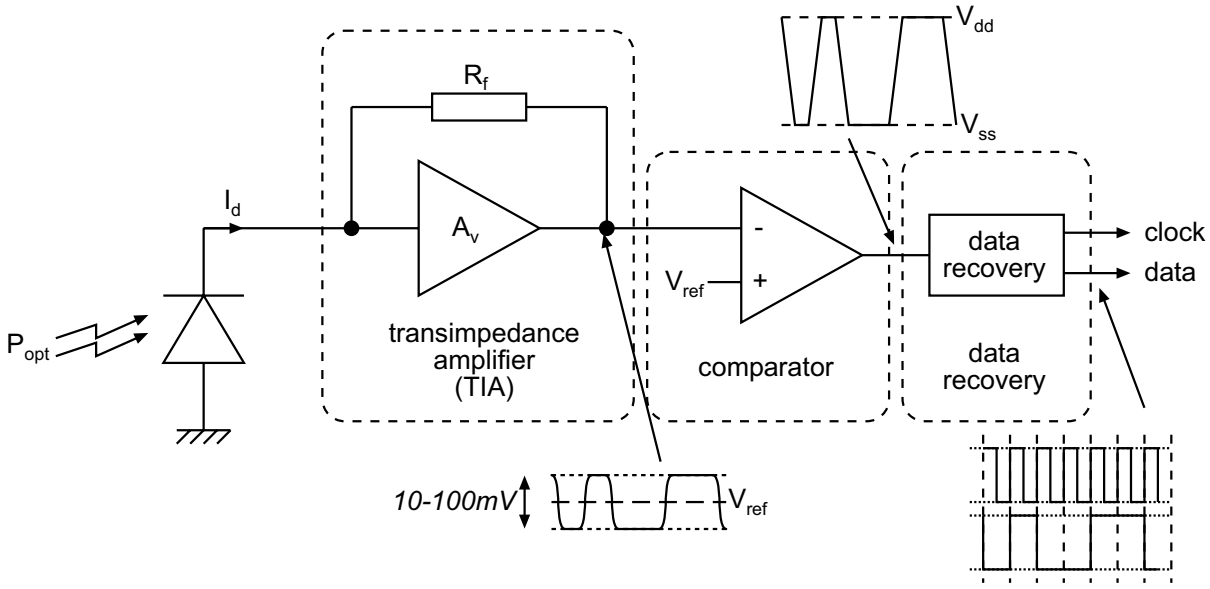

Figure 5.8. Typical photoreceiver circuit

Of these, the TIA is arguably the most critical component, since it has to cope with a generally large photodiode capacitance situated at its input. Bandwidth/power ratio maximization can be achieved in several ways:

- parametric optimization: for a given transimpedance structure, find the combination of component parameters necessary for maximum bandwidth
- structural modification: for a given preamplifier architecture, make structural modifications, usually by adding elements such as inductors for shunt peaking [28] or capacitors as artificial loads or feedback [29]
- architectural exploration: use complex architectures such as bootstrap or common-gate input stages [30]

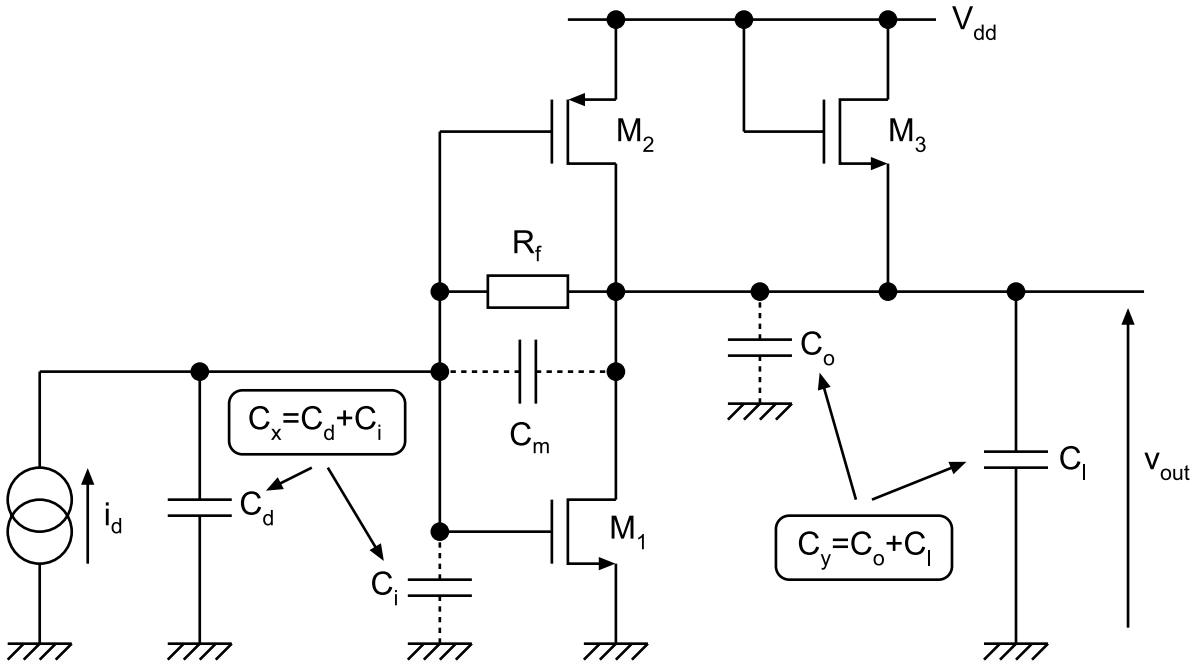

Figure 5.9. CMOS transimpedance amplifier structure

The basic transimpedance amplifier structure in a typical configuration is shown in Fig. 5.9 [31]. The bandwidth/power ratio of this structure can be maximized by using small-signal analysis and mapping of the individual component values to a filter approximation of Butterworth type, which gives:

$$Z_{g0} = \frac{R_0 - R_f A_v}{1 + A_v} \tag{4}$$

$$\omega_0 = \frac{1}{R_0 C_y} \sqrt{\frac{1 + A_v}{M_f \left(M_x + M_m + M_x M_m \right)}} \tag{5}$$

$$Q = \frac{\sqrt{M_f \left(M_x + M_m \left(1 + M_x \right) \right) \left(1 + A_v \right)}}{1 + M_x \left(1 + M_f \right) + M_m M_f \left(1 + A_v \right)} \tag{6}$$

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, normalizing all expressions to the time constant $\tau=R_oC_y$. 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 realization 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. 5.10(a) shows a plot of this quantity against the TIA specifications (bandwidth and transimpedance gain) for $C_x = C_d = 500fF$ and $C_y = C_l = 100fF$. 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 optimization, using simulation for exact results [32].

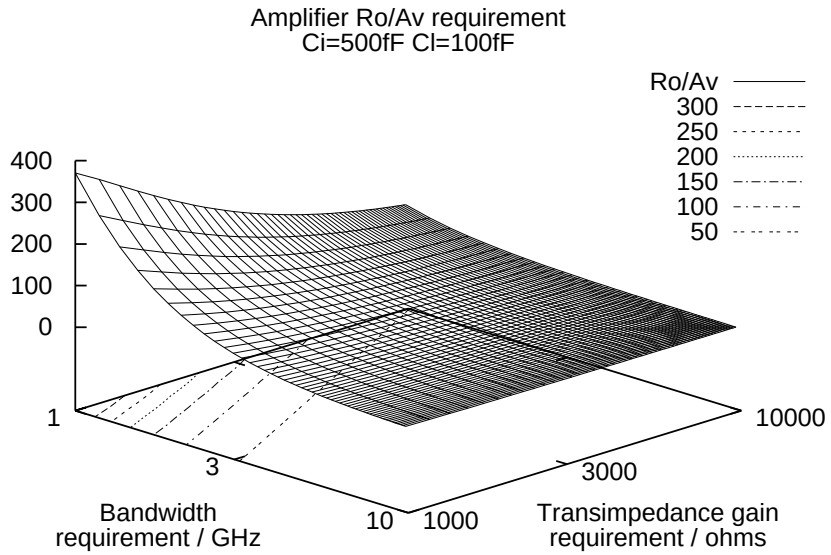

(a) R_o/A_v design space with varying bandwidth and transimpedance gain requirements

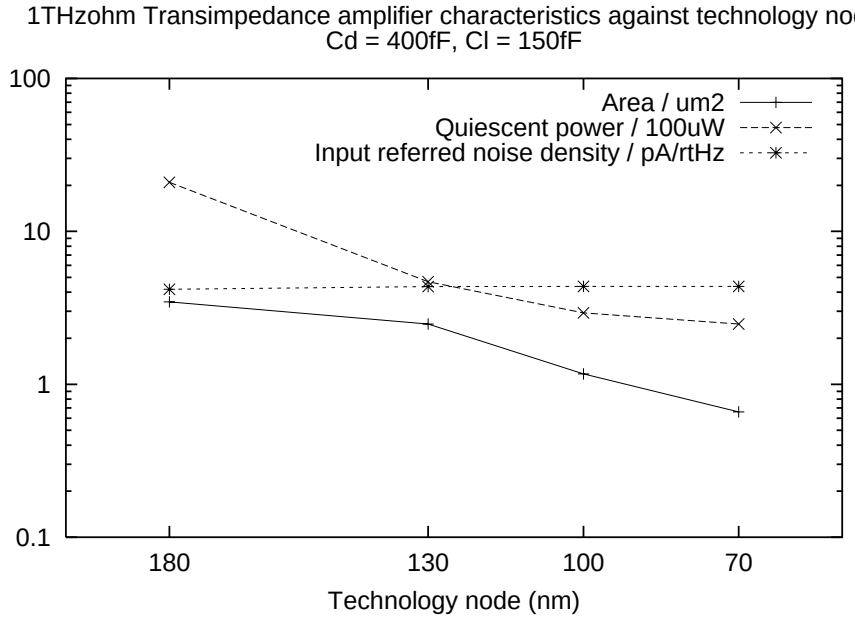

(b) Evolution of TIA characteristics (power, area, noise) with technology node

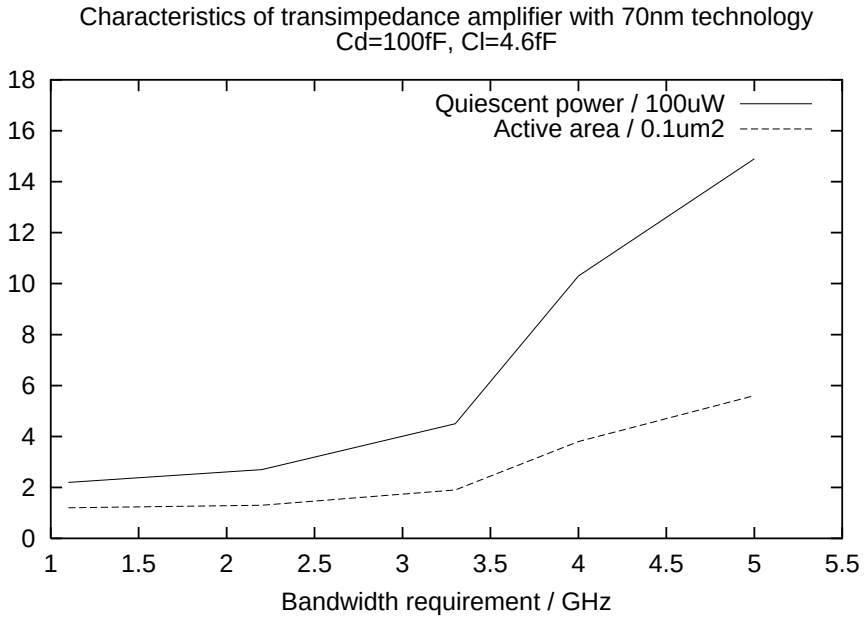

(c) Power and area against bandwidth requirement for TIAs at the 70nm technology node

Figure 5.10. Transimpedance amplifier characteristics

Using this methodology and predictive BSIM3v3 models for technology nodes from 180nm down to 70nm [33], we generated design parameters for 1THzΩ transimpedance amplifiers to evaluate the evolution in critical characteristics with technology node. Fig. 5.10(b) shows the results of transistor level simulation of fully generated photoreceiver circuits at each technology node. According to traditional "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. This methodology also allows us to find a particular specification to a given tolerance, as shown in Fig. 5.10(c). This shows the active area and of the generated TIA for bandwidths of 1GHz - 5GHz (with $Z_{g0}=1k\Omega$ and $Q=1/\sqrt{2}$).

5.6 Bonding Issues

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 [34]. 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 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 (800°C). However, recent research [35] 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 200°C (for annealing, to increase bonding energy), which is compatible with CMOS IC fabrication steps.

5.7 Link Performance (Comparison of Optical and Electrical Systems)

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 analyzed in both systems at the 70nm technology node. Power dissipation figures for electrical and optical CDNs were calculated based on the system performance summarized in Tables 5.1 and 5.2. 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.

Technology [μm]	0.07
V_{dd} [V]	0.9
T_{ox} [nm]	1.6
Chip size [mm^2]	400
Global wire width [μm]	1
Metal resistivity [μΩ-cm]	2.2
Dielectric constant	3
Optimal segment length [mm]	1.7
Optimal buffer size [μm]	90

Table 5.1. Electrical system performance

Wavelength λ [μm]	1.55
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

Table 5.2. Optical system performance

The first input to this program is the set of technology parameters for the process of interest, in particular the feature size, dielectric constant and metal resistivity according to the ITRS roadmap. In the next step, the resistance, capacitance and inductance values for a given metal layer are calculated, as well as the electrical parameters of minimum size inverters. Based on this information, it is then possible to determine the optimal number and size of buffers needed to drive the clock network. For such system the program creates the SPICE netlist where the interconnect is replaced by RC or RLC distributed lines coupled by buffers designed as CMOS inverters. Berkeley BSIM3v3 [33] parameters were used to model the transistors used in the inverters. The power dissipated in the system is extracted from transistor-level simulations.

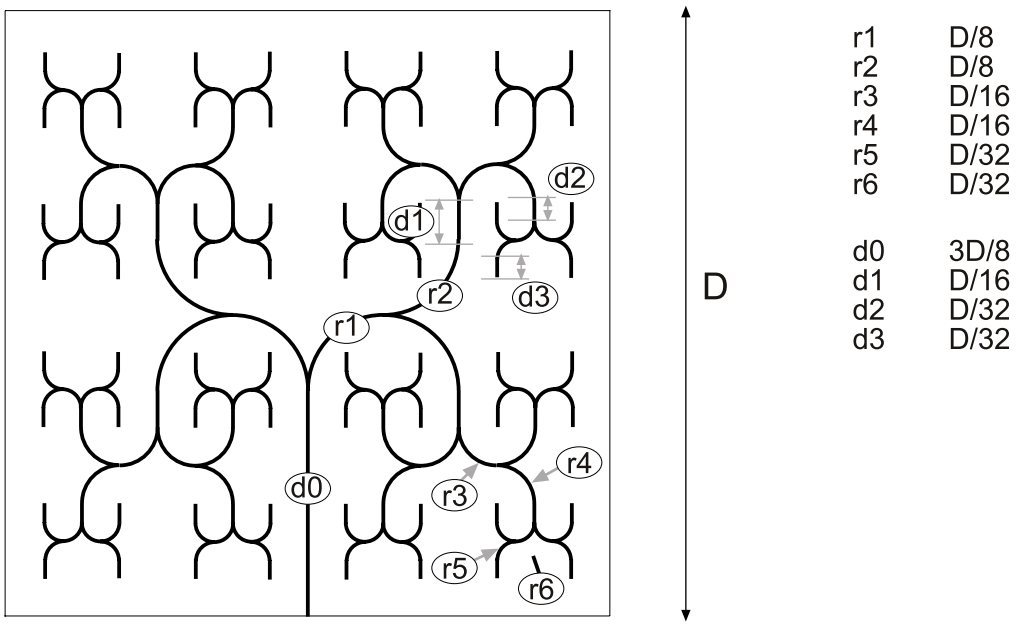

Figure 5.11. Optical H-tree network with 64 output nodes. $r_1..r_6$ are the radius of curvatures and $d_0..d_3$ are the lengths of straight lines linked to the chip width D.

In the optical clock distribution network (Fig. 5.11) there are two main sources of electrical power dissipation: (i) power dissipated by the optical receivers and (ii) energy needed by the optical source to provide the required optical output power. To estimate the electrical power dissipated in the system we used the methodology shown in Fig. 5.12. We assumed the use of an external VCSEL source, rather than an integrated microsource. The global optical H-tree was optimized 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. 5.11) is 625μm, which leads to negligible pure bending loss.

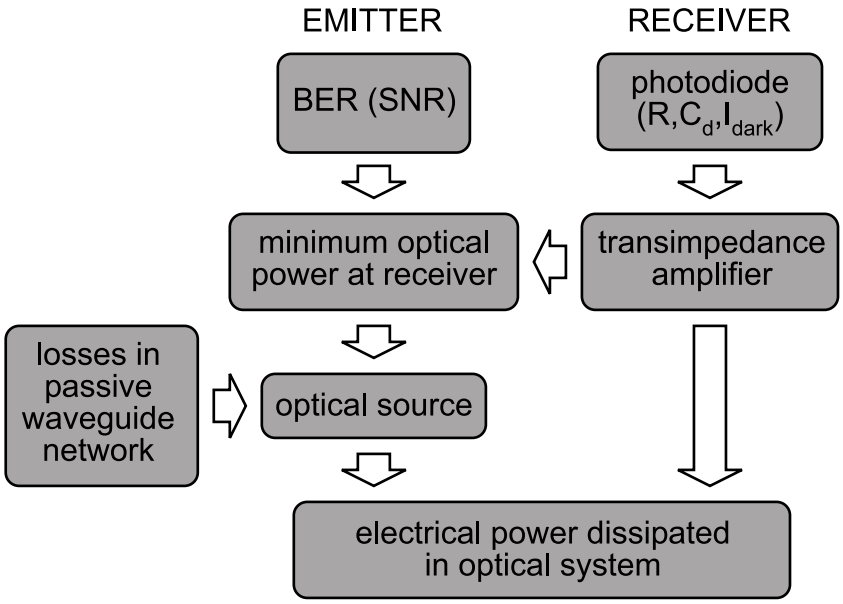

Figure 5.12. Methodology used to estimate the electrical power dissipation in an optical clock distribution network

First, based on the given photodiode parameters (C_d, R, I_{dark}), the method described in section 5.5 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 [36] in the preamplifier noise calculations.

To estimate the optical power emitted by the VCSEL we took into account the previously calculated minimal required signal by the receiver and the losses incurred throughout passive optical waveguides. The electrical power dissipated in optical clock networks is the sum of the power dissipated by the number of optical receivers and the energy needed by the VCSEL to provide the required optical power. The electrical power dissipated by the receivers has been extracted from transistor-level simulations. To estimate the energy needed by the optical source, we use the laser light-current characteristics given by Amann [37].

Number of nodes in H-tree	4	8	16	32	64	128
Loss in straight lines [dB]	1.3	1.3	1.3	1.3	1.3	1.3
Loss in curved lines [dB]	1	1.31	1.53	1.66	1.78	1.85
Y-dividers [dB]	6	9	12	15	18	21
Loss in Y-couplers [dB]	0.4	0.6	0.8	1	1.2	1.4
Output coupling loss [dB]	0.6	0.6	0.6	0.6	0.6	0.6
Input coupling loss [dB]	3	3	3	3	3	3
Total optical loss [dB]	12.3	15.8	19.2	22.5	25.8	29.1
Min. receiver power [dBm]	-22.3	-22.3	-22.3	-22.3	-22.3	-22.3
Laser optical power [mW]	0.1	0.25	0.5	1.1	2.30	4.85

Table 5.3. Optical power budget for 20mm die width at 3GHz

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 summarized in Table 5.3. Fig. 5.13 shows the optical power emitted by the VCSEL necessary to provide a given BER for various waveguide transmission losses. The comparison in terms of dissipated power between the optical and electrical global clock distribution networks is shown in Fig. 5.14. 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.

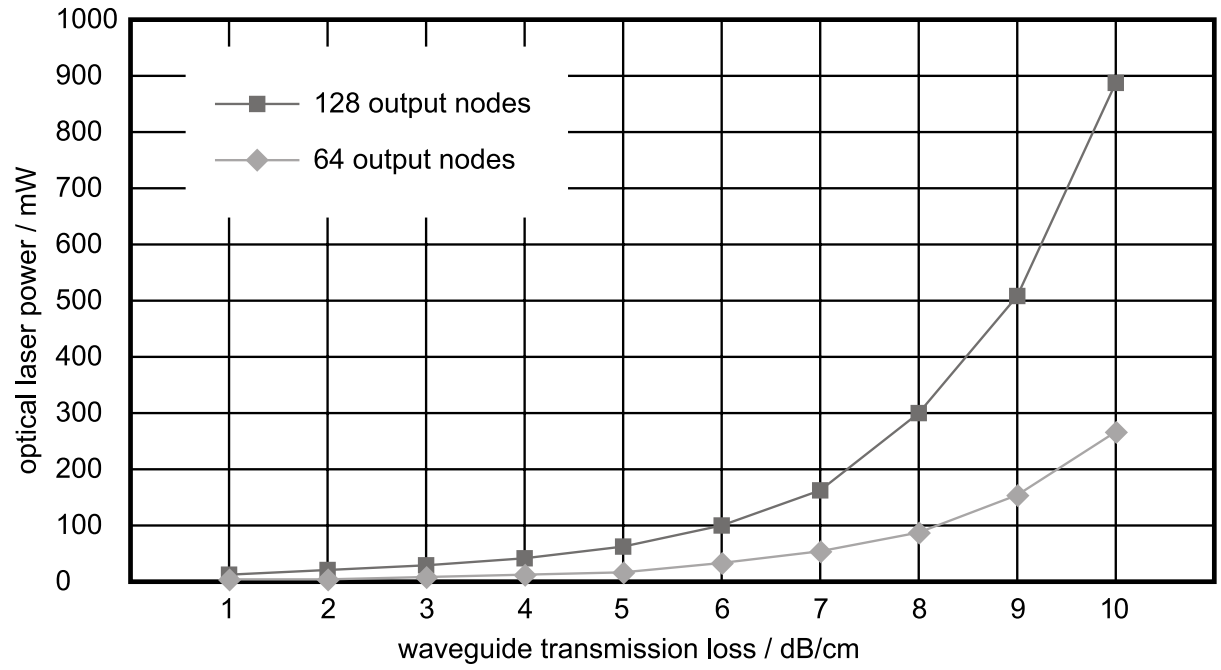

Figure 5.13. VCSEL optical output power required by the H-tree to provide a given BER for varying waveguide transmission loss

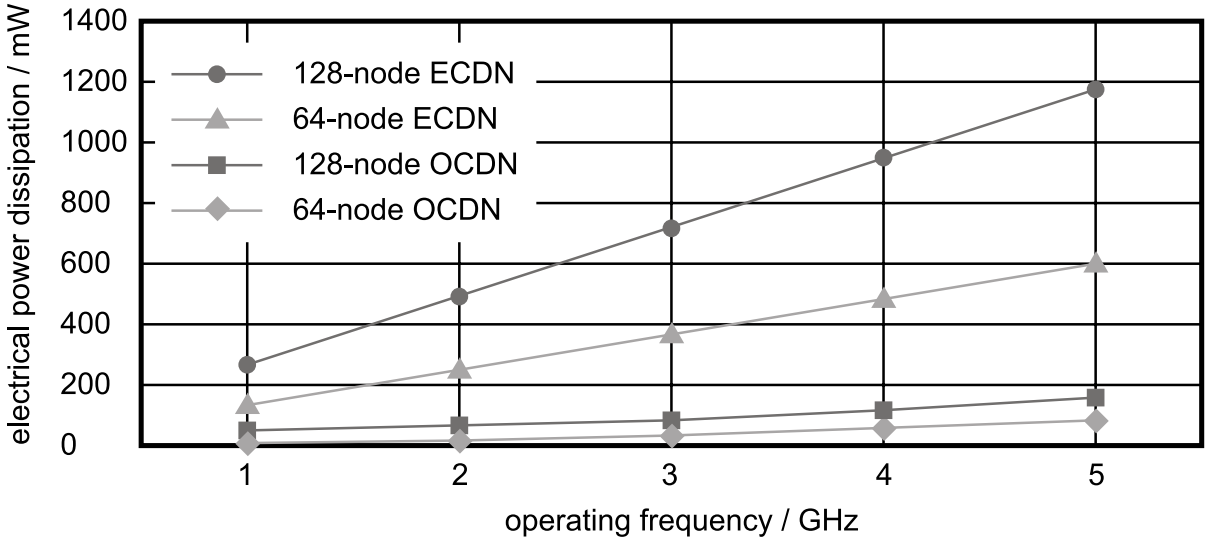

Figure 5.14. Electrical power dissipated by electrical (ECDN) and optical (OCDN) clock distribution networks for varying operating frequency

5.8 Research Directions

Integrated optical interconnect is one potential technological solution to reduce the power required to move volumes of data between circuit blocks on integrated circuits. But it only makes sense to use this technology for *global high-speed* data links. And if the use of this

technology implies a hard breach in terms of process technology and design methodologies, then architectural solutions may be an easier way to reduce power, by optimizing layout for the application such that the need for global high-speed data links is alleviated.
It is difficult to predict between these scenarios. Probably future ICs will make use of advances in both areas. One parameter that is likely to make the difference in favor of optical interconnect is new research into the possibility of on-chip wavelength division multiplexing (WDM).

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 ofGb/s. Commercial solutions for SoC development platforms are now based on bus architectures [38] [39]: metallic interconnect architectures rely heavily on wide (64/128-bits) buses, as well as frequent use of switch boxes to dynamically define a communication route between two functional IP blocks. Again, since the order of distance of communication is the chip die size, systematic use of repeaters (over the buses or within the switch boxes themselves) is necessary and increases power consumption.
In the future, limitations (latency due to line delay, non scalability, time sharing, non reconfigurability, etc.) of bus-based architectures will appear: future architectures of integrated systems will require new concepts for on-chip data exchange. The ever increasing number of transistors in a chip will lead to such complexity that IP reuse will be mandatory: a system is designed by integrating some hundreds of pre-designed complex functional blocks, the designer concentrating mainly on the organization of data transfer between these blocks. A number of innovative interconnect architectures, often called networks on chip (NoC), have been recently proposed to overcome the limitations of bus-based platforms [40] [41] [42]. NoC architectures look much more like switching 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 a 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.

Optical buses

As in telecommunications, WDM provides a route to very high data rates, even if the individual devices cannot be modulated much faster than electrical bus data rates. A single waveguide could be used to replace a 64-bit bus, for example, where each individual signal makes use of a distinct wavelength.

Reconfigurable networks

By using a WDM approach, reconfigurable networks could be realized in the optical domain, leading to power reduction and higher integration density. Switch boxes, a key element for reconfigurable networks, could also be realized by using compact micro-resonators (about $10x10\mu m^2$), capable of selecting and redirecting a signal based on its wavelength. Such networks would be entirely passive; i.e. no power would be required to transport the data, whatever the communication route necessary. However, such a scheme would imply a shift in the routing paradigm from a centralized arbiter acting on the switch boxes, to one acting on the block interfaces to select the wavelength(s) to be used. Also, tunable and thermally stable microlasers would be required.

Acknowledgements

The authors would like to thank Xavier Letartre for valuable discussions on active integrated devices for optical interconnect.

References

[1] Semiconductor Industry Association, International Technology Roadmap for Semiconductors 2001 Edition, http://public.itrs.net, 2001

[2] Sakurai, T. and Tamaru, K., Simple Formulas for Two- and Three–Dimensional Capacitances, *IEEE Tran. Electron Devices*, 30, 183, 1983

[3] Chang, M.F. et al., RF/wireless interconnect for inter- and intra-chip communications, *Proc. IEEE*, 89, 456, 2001

[4] Banerjee, K. et al., 3-D ICs: A novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration, *Proc. IEEE*, 89, 602, 2001

[5] Miller, D.A.B., Rationale and challenges for optical interconnects to electronic chips, *Proc. IEEE*, 88, 728, 2000

[6] Davis J.A., De V.K., Meindl J.D., A stochastic wire-length distribution for gigascale integration (GSI) - Part I: Derivation and Validation, *IEEE Trans. Electron Dev.*, 45, 580, 1998

[7] Collet, J.H. et al., Architectural approach to the role of optics in mono- and multi-processor machines, *Applied Optics*, 39, 671, 2000

[8] Friedman E.G., Clock distribution networks in synchronous digital integrated circuits, *Proc. IEEE*, 89, 665, 2001

[9] Lee, K.K. et al., Fabrication of ultralow-loss Si/SiO_2 waveguides by roughness reduction, *Optics Letters*, 26, 1888, 2001

[10] Sakai, A., Hara, G. and Baba, T., Propagation Characteristics of Ultrahigh-Δ Optical Waveguide on Silicon-on-Insulator Substrate, *Jpn. J. Appl. Phys. - Part 2*, 40, 383, 2001

[11] Payne, F.P. and Lacey, J.P.R., A theoretical analysis of scattering loss from planar optical waveguides, *Optical and Quantum Electronics*, 26, 977, 1994

[12] Nishihara, H., Haruna, M. and Suhara, T., Optical Integrated Circuits, McGraw-Hill, 1988

[13] Kim, C.M., Jung, B.G. and Lee, C.W., Analysis of dielectric rectangular waveguide by modified effective-index method, *IEE Electron. Letters*, 22, 296, 1986

[14] Marcuse, D., Light Transmission Optics, Van Nostrand Reinhold, New York, 1972

[15] Chu, F.S. and Liu, P.L., Low-loss coherent-coupling Y branches, *Optics Letters*, 16, 309, 1991

[16] Rangaraj, M., Minakata, M. and Kawakami, S., Low loss integrated optical Y-branch, *IEEE J. Lightwave Technology*, 7, 753, 1989

[17] Sakai, A., Fukazawa, T. and Baba, T., Low Loss Ultra-Small Branches in a Silicon Photonic Wire Waveguide, *IEICE Trans. Electron.*, E85-C, 1033, 2002
[18] Schultz, S.M., Glytsis, E.N. and Gaylord, T.K., Design, Fabrication, and Performance of Preferential-Order Volume Grating Waveguide Couplers, *Applied Optics-IP*, 39, 1223, 2000
[19] Little B.E, Chu S.T., Pan W. and Kokubun Y., Microring resonator arrays for VLSI photonics, *IEEE Phot. Tech. Lett.*, 12, 323, 2000
[20] Martinez, A., Cuesta, F. and Marti, J., Ultrashort 2-D photonic crystal directional couplers, *IEEE Phot. Tech. Lett.*, 15, 694, 2003
[21] Notomi, M. et al., Structural tuning of guiding modes of line-defect waveguides of silicon-on-insulator photonic crystal slabs, *IEEE J. Quantum Electronics*, 38, 736, 2002
[22] Baba, T., Photonic Crystals and Microdisk Cavities Based on GaInAsP-InP System, *IEEE J. Selected Topics in Quantum Electronics*, 3, 808, 1997
[23] Filios, A. et al., Transmission performance of a 1.5-μm 2.5-Gb/s directly modulated tunable VCSEL, *IEEE Phot. Tech. Lett.*, 15, 599, 2003
[24] Liu, J.J. et al., Ultralow-threshold sapphire substrate-bonded top-emitting 850-nm VCSEL array, *IEEE Phot. Lett.*, 14, 1234, 2002
[25] Fujita, M., Sakai, A. and Baba, T., 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, 673, 1999
[26] Cho, S.Y. et al., Integrated detectors for embedded optical interconnections on electrical boards, modules and integrated circuits, IEEE J. Sel. Topics in Quantum Electronics, 8, 1427, 2002
[27] Fujita, M., Ushigome, R. and Baba, T., Continuous wave lasing in GaInAsP microdisk injection laser with threshold current of 40μA, *IEE Electron. Lett.*, 36, 790, 2000
[28] Mohan, S.S. et al., Bandwidth extension in CMOS with optimized chip inductors, *IEEE J. Solid-State Circuits*, 35, no. 3, March 2000
[29] Kuo, C.W. et al., 2 Gbit/s transimpedance amplifier fabricated by 0.35μm CMOS technologies, *IEE Electron. Letters*, 37, 1158, 2001
[30] Graeme, J., *Photodiode Amplifiers*, McGraw-Hill, Boston, 1996, chap. 4
[31] Ingels, M. and Steyaert, M. S. J., A 1-Gb/s, 0.7μm CMOS optical receiver with full rail-to-rail output swing, *IEEE J. Solid-State Circuits*, 34, 971, 1999

[32] O'Connor, I. et al., Exploration paramétrique d'amplificateurs de transimpédance CMOS à bande passante maximisée, in *Proc. Traitement Analogique de l'Information, du Signal et ses Applications*, 73, Paris, 2002
[33] Cao, Y. et al., New paradigm of predictive MOSFET and interconnect modeling for early circuit design, in *Proc. Custom Integrated Circuit Conference*, 2000
[34] Krishnamoorthy, A.V. and Goossen, K.W., Optoelectronic-VLSI: Photonics integrated with VLSI circuits, *IEEE J. Sel. Topics in Quantum Electronics*, 4, 899, 1998
[35] Seassal, C. et al., InP microdisk lasers on silicon wafer: CW room temperature operation at 1.6μm, *IEE Electron. Lett.*, 37, 222, 2001
[36] Morikuni, J.J. et al., Improvements to the standard theory for photoreceiver noise, *IEEE J. Lightwave Technology*, 12, 1174, 1994
[37] Amann, M.C., Ortsiefer, M. and Shau, R., Surface-emitting Laser Diodes for Telecommunications, in *Proc. Symp. Opto- and Microelectronic Devices and Circuits*, Stuttgart, 2002
[38] IBM, The core connect bus architecture, http://www.chips.ibm.com/products/coreconnect, 1999
[39] VSI Alliance, http://www.vsi.org
[40] Benini, L. and De Micheli, G., Networks on Chip: A New SoC Paradigm, *IEEE Computer*, 35, 70, 2002
[41] Guerrier, P. and Greiner, A., A generic architecture for on-chip packet-switched interconnections, in *Proc. Design, Automation and Test in Europe 2000*, 250, 2000.
[42] Dally, W.J. and Towles, B., Route packets, not wires: On-chip interconnection networks, in *Proc. 38th Design Automation Conference*, 2001