Systematic Simulation-Based Predictive Synthesis of Integrated Optical Interconnect

Ian O'Connor, *Member, IEEE*, Faress Tissafi-Drissi, Frédéric Gaffiot, Joni Dambre, *Member, IEEE*,
Michiel De Wilde, Joris Van Campenhout, *Student Member, IEEE*, Dries Van Thourhout, *Member, IEEE*,
Jan Van Campenhout, *Member, IEEE*, and Dirk Stroobandt, *Member, IEEE*

Abstract—Integrated optical interconnect has been identified by the ITRS as a potential solution to overcome predicted interconnect limitations in future systems-on-chip. However, the multiphysics nature of the design problem and the lack of a mature integrated photonic technology have contributed to severe difficulties in assessing its suitability. This paper describes a systematic, fully automated synthesis method for integrated microsource-based optical interconnect capable of optimally sizing the interface circuits based on system specifications, CMOS technology data, and optical device characteristics. The simulation-based nature of the design method means that its results are relatively accurate, even though the generation of each data point requires only 5 min on a 1.3-GHz processor. This method has been used to extract typical performance metrics (delay, power, interconnect density) for optical interconnect of length 2.5–20 mm in three predictive technologies at 65-, 45-, and 32-nm gate length.

Index Terms—Integrated optical interconnect, multidomain design methods, synthesis.

I. INTRODUCTION

THE emergence of very high performance systems-on-chip (SoC) is necessary to achieve future required application performance in terms of resolution (audio, video, and computing) and CPU power/total MIPS (real-time encoding-decoding, data encryption-decryption). The shift to distributed multiprocessor architectures is the recognized route to such performance and, therefore, requires organized high-speed communication between processors. Metallic interconnect will be highly inefficient in this role due to unachievable trade-offs between design parameters (the main limitations due to metallic interconnects are inter-line crosstalk, latency, global throughput, connectivity, and power consumption).

It is believed that the concept of integrated optical interconnect is a potential technological solution to alleviate some of these issues involved in exchanging data between cores in SoC

architectures. However, the multiphysics nature of the design problem and the lack of a mature integrated photonic technology have contributed to severe difficulties in assessing the suitability of integrated optical interconnect for on-chip data transport. In turn, these difficulties have rendered unclear the research directions for the integrated photonic device community. The objective of optical interconnect assessment should thus be two-fold: first, to explore the performance capabilities of optical links with existing technological constraints; and second, to provide feedback to be taken into account in component and integration oriented research.

To enable the analysis previously described, it is necessary to define and quantify the specification set for the various subcomponents in an integrated optical link. This requires the development of models and tools to design and optimize optical interconnect in a range of technologies and under a range of specification conditions, in order to evaluate the properties of the technologies developed. Ultimately, the objective is to compare them to the performance of current and future electrical interconnect alternatives in terms of important performance metrics (mainly area, data rate, and power). A contiguous aim is to determine the specifications for which the proposed technologies would outperform traditional interconnect.

Optical-electrical interconnection comparison is not new: several authors have already published comparative studies between optical and electrical on-chip interconnect technologies [1]–[3]. Most of those consider expected technological evolutions to provide a roadmap of interconnect performance. This is done either by using analytical models, or by a simulation-based approach. Analytical models are usually based on ITRS[1] [4] or other expected future CMOS technology parameters and rough estimates of optical interconnect performance. Simulation models are either self-constructed, based on ITRS parameters, or using publicly available Predictive Technology Models (PTM[2] for gate lengths ranging down to 32 nm).

In this paper, as it was intended to be able to determine device-level specifications, it was necessary to generate accurate optical link performance estimates that can be traced back to individual component parameters. Part of the optical link consists of analog CMOS circuitry. The design thereof can be done based on analytical models, but it is very sensitive to MOSFET parasitics so it is crucial that these parameters are available. However, the data presented in the ITRS is intended for digital design applications. This implies that estimates on transistor parasitics

I. O'Connor and F. Gaffiot are with the Lyon Institute of Nanotechnology, Ecole Centrale de Lyon, F-69134 Ecully, France (e-mail: ian.oconnor@ec-lyon.fr).

F. Tissafi-Drissi was with the Lyon Institute of Technology, Ecole Centrale de Lyon, F-69134 Ecully, France. He is now with the Mixed-Signal Design Department, Dolphin Integration, F-38242 Meylan, France.

J. Dambre, M. De Wilde, J. Van Campenhout, and D. Stroobandt are with the Electronics and Information Systems Department, Ghent University, B-9000 Ghent, Belgium.

D. Van Thourhout is with the INTEC Department, Ghent University, B-9000 Ghent, Belgium.

[1]International Technology Roadmap for Semiconductors. [Online]. Available: http://public.itrs.net

[2][Online]. Available: http://www.eas.asu.edu/~ptm/

are very rough or in some cases not even available. Hence, for an accurate evaluation, it is best to use technologies for which simulation models are available.

For these reasons, we have developed a simulation-based approach for optical interconnect synthesis. By way of illustration, we apply our approach to predictive technology models with gate lengths of 65, 45, and 32 nm. In contrast to the model-based approaches, our choice implies that we presently cannot stretch our analysis beyond 32-nm gate lengths according to the availability of models. However, according to the ITRS, a 32-nm gate length technology (which is the smallest available transistor model on the PTM website) would be taken into production for high-end applications in 2008, and this is only the beginning of its lifetime. Therefore, we claim that our analysis does stretch sufficiently far into the future. Also, by extrapolating observed trends in performance metric evolution over technology nodes, some insight into the possible results for future technology nodes can be obtained.

The models and software developed were, therefore, made to be sufficiently generic in order to enable the exploration of: 1) performance capabilities of optical links with existing technological constraints; and 2) the impact of improvements in device performance on interface circuit sizing and overall link performance. This set of data allows tradeoffs to be understood, as well as their relevance in the context of overall link performance metrics.

This paper first describes some technological and device aspects in Section II. The generation of optical link parameters based on link specifications is based on analog synthesis techniques, and the individual interface circuit synthesis methods are described in Section III. To enable complete optical link simulation, a number of behavioral models were developed, and this work is outlined in Section IV, although details are outside the scope of this paper. Section V concentrates on describing the complete optical link synthesis method exploiting the simulation work and individual circuit synthesis methods, as well as the details of the investigation program. Extraction methods concerning interconnect density analyses are described in Section VI. To generate performance metric data for considered technologies and optical link lengths, the simulation-based quantitative assessment of optical interconnects at the physical link level is described in Section VII. Through identification of the impact of individual device parameter variations on particular performance metrics, an exploration of these variations is shown in Section VIII.

II. TECHNOLOGY AND DEVICE ASPECTS

Various technological solutions exist for using optics with a standard CMOS SoC [5]. Our approach is generic, in the sense that it can be applied to any type of integrated photonic technology, as long as a behavioral model is available for each component (see Section IV). In this paper, we have supposed a microsource-based link using heterogeneous integration [6] to achieve an above-IC optical transport layer (see Fig. 1). A CMOS driver circuit modulates the current flowing through the microsource laser, and thus the intensity of light emission.
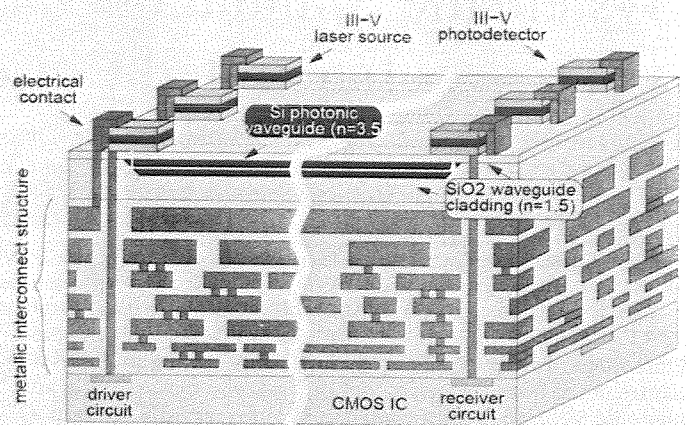

Fig. 1. Cross section of optical interconnect structure.

The microsource is coupled to a 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 back to an electrical digital (rail to rail) signal and subsequently distributed by a local electrical interconnect network, represented by an output capacitance (fanout) that the link can drive. Similarly, the overall input capacitance presented by the link to the digital signal generator is part of the specification set for the optical link.

To form a planar optical waveguide, the most common CMOS-compatible materials used are polymers Si_3N_4 and silicon. In a first approach, we assumed silicon as the core and SiO_2 as the cladding material. Si/SiO_2 structures are compatible with conventional silicon technology and are transparent for 1.3–1.55 μm wavelengths. Such waveguides also benefit from a high relative refractive index difference $\Delta \approx \left(n_1^2 - n_2^2\right)/2n_1^2$ between the core ($n_1 \approx 3.5$ for Si) and claddings ($n_2 \approx 1.5$ for SiO_2) and enable the realization of compact optical circuits, with bend radius of the order of a few micrometers [7] and low pitch to crosstalk ratio. To avoid modal dispersion, improve coupling efficiency, and reduce loss, single mode conditions are applied to the waveguide dimensions. For a wavelength of 1.55 μm in Si/SiO_2, the waveguide dimensions are around 0.5 μm × 0.22 μm (width × height).

Si_3N_4 waveguides are also contenders for optical interconnect. Since the index difference is lower, the propagation delay is also lower, but the penalty for this is higher crosstalk and bend loss, which implies a requirement for higher pitch and limited use of bends. The waveguide dimensions for these waveguides are around 0.4 μm × 0.8 μm (width × height). Data for both types of waveguide is given in the optical link simulation Section IV.

Integrated III–V sources can achieve >10-GHz 20% total efficiency (including coupling to waveguide) when biased above a 40 μA threshold current, and have a footprint of some 10 × 10 μm^2 [8]. Similarly, 50 μm^2 InP detectors exhibit 24 GHz 70% responsivity and low (< 10 fF) capacitance [9]. Heterogeneous integration/bonding of such devices on CMOS has also been proven [10].
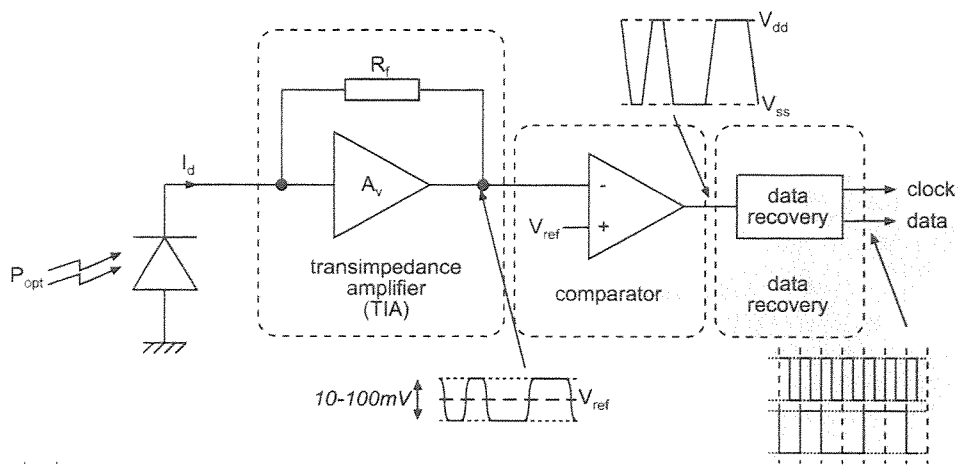
Fig. 2. Typical receiver circuit structure.

III. Synthesizable Model Library

Manual design of physically heterogeneous systems results in long design cycles and increasingly apparent bottlenecks in the overall design process [11]. This explains the growing awareness in industry that the advent of heterogeneous synthesis and optimization tools is a necessary step to increase design productivity by assisting or even automating the design process. The fundamental goal of heterogeneous synthesis is to quickly generate a first-time correct sized schematic from a set of specifications. This is critical since the heterogeneous design problem is typically under constrained with many degrees of freedom and with many interdependent (and often conflicting) performance requirements to be taken into account across physical domains. For this paper, this type of approach is essential since the objective is to explore the design of optical links subject to many sets of performance requirements and to many sets of technological and device-level constraints.

For this paper, we used our synthesis framework RuneII [12], which enables the specification and automated design of heterogeneous systems, using a hierarchical approach to partition the system at a behavioral description level, and domain-specific simulators to design individual blocks. It exploits the concept of synthesizable AMS IP [13], which extends the concept of digital and software IP to the analog and heterogeneous domains. In essence, all information necessary for the design of a given block is formalized in a technology independent way such that a generic framework (such as RuneII) can exploit the various facets of the AMS IP information to arrive at a sized system schematic satisfying all performance requirements in the target technology.

To carry out a systematic design analysis of optical interconnects, a set of design requirements and synthesizable models must be established. This section will detail the various synthesizable models developed (their final integration into a complete optical link synthesis method will be described in Section V).

A. Receiver Circuit Synthesizable Model

The classical structure for a receiver circuit is shown in Fig. 2: a transimpedance amplifier (TIA) converts the photocurrent of a few microamperes into a voltage of a few millivolts; a comparator generates a rail-to-rail signal; and a data recovery circuit eliminates jitter from the restored signal.
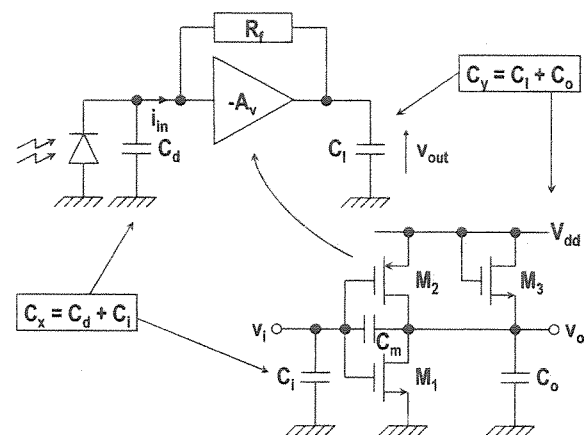

Fig. 3. Resistive feedback transimpedance amplifier and CMOS fast inverter structure used as internal amplifier in TIA.

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. For the comparator, scaling figures were used from a reference circuit manually designed at the 90-nm gate length node, and the input capacitance estimated for TIA synthesis.

It should be noted that photodetector capacitances in advanced devices can be around 10 fF or less—at this level, it is likely that the input capacitance of the TIA itself will play a large part in the determination of bandwidth. Bandwidth/power ratio maximization can be achieved in the following 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 [14] or capacitors as artificial loads or feedback [15];
- architectural exploration: use complex architectures such as bootstrap or common-gate input stages [16].

The basic transimpedance amplifier structure in a typical configuration is shown in Fig. 3, with the internal structure of the amplifier used in the resistive feedback configuration [17]. In
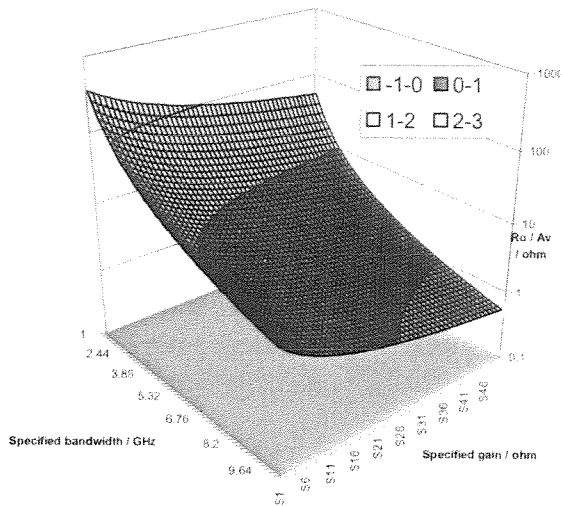
Fig. 4. R_o/A_v design space with varying bandwidth and transimpedance gain requirements.

Fig. 5. Evolution of TIA characteristics (power, area, noise) with technology node for Zg0. BW $= 1$ THzΩ, Cd $= 400$ fF, Cl $= 45$ fF.

the following analysis, C_x represents the aggregated capacitance at the TIA input node (including the photodetector capacitance C_d), C_y represents the aggregated capacitance at the TIA output node (including the load capacitance C_l), C_m represents the aggregated Miller capacitance between the TIA input and output nodes, R_f represents the feedback resistance value, R_o represents the fast inverter output resistance, and A_v represents the fast inverter voltage gain. 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

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

$$\omega_0 = \frac{1}{R_0 C_y} \sqrt{\frac{1 + A_v}{M_f (M_x + M_m + M_x M_m)}} \tag{2}$$

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

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

Taking into consideration their physical realization, amplifiers 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. 4 shows a plot of this quantity against the TIA specifications (bandwidth and transimpedance gain) for $C_x = C_d = 400$ fF and $C_y = C_l = 45$ fF.

Approximate equations for the small signal characteristics and bias conditions of the circuit allow a first cut sizing of the amplifier. The solution is then fine tuned by a local numerical optimization method driven by a sum of weighted cost functions
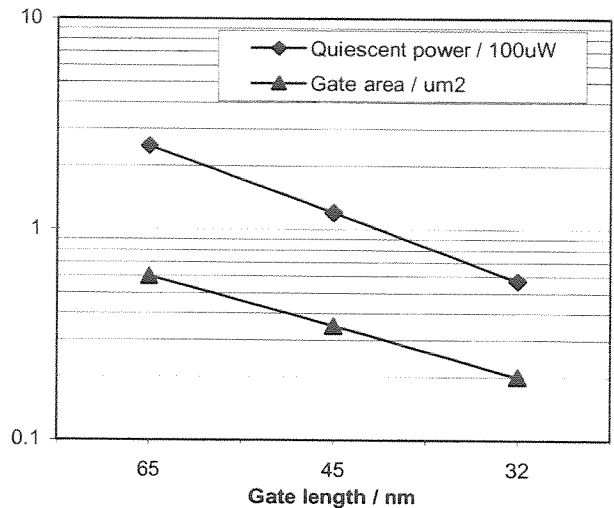
to achieve the desired values of A_v and R_o, using simulation for exact results [18]. Parasitic capacitances are extracted from the resulting transistor sizes and are used to update the values of C_x and C_y. The transimpedance amplifier sizing and amplifier optimization steps are subsequently run again with continually updating values of A_v, R_o, C_x, and C_y until the process converges.

Using this methodology and predictive BSIM3v3 and BSIM4 models for technology nodes from 65 nm down to 32 nm, we generated design parameters for 1-THzΩ transimpedance amplifiers to evaluate the evolution in critical characteristics with technology node. Fig. 5 shows the results of transistor level simulation of fully generated photo-receiver 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 considered technology nodes 65–32 nm, $s^2 = 4$, which is approximately verified through the sizing procedure (quiescent power scales by 1/4.4, gate area scales by 1/3). This methodology also allows us to find a particular specification to a given tolerance, as shown in Fig. 6. This shows the active area and power of the generated TIA for bandwidths of 1–5 GHz (with $Z_{g0} = 1$ kΩ and $Q = 1/\sqrt{2}$ in the 45-nm technology).

B. Driver Circuit Synthesizable Model

The basic current modulation configuration of the source driver circuit is shown in Fig. 7. While more advanced driver schemes exist (such as voltage-pulsed drivers [19]), the current modulation circuit is sufficient to evaluate the characteristics of the overall link in a systematic way. 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 50 μA appear achievable [8]). Device M_1 modulates the current flowing through the source, and consequently the output optical power injected into the waveguide. The reduced dimensions of the
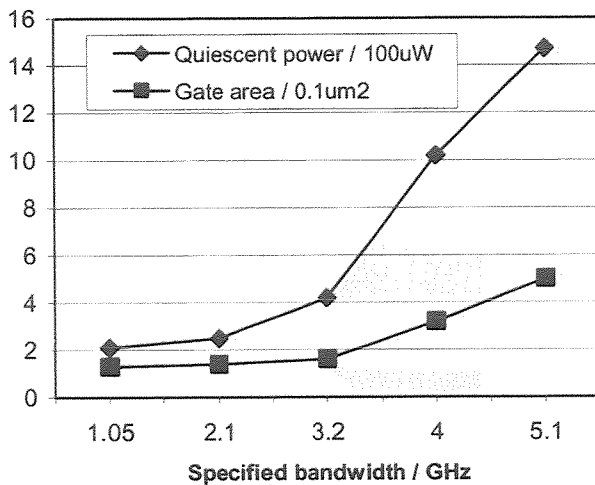
Fig. 6. Power and area against bandwidth requirement for TIAs at the 65-nm technology node for $L_g = 65$ nm, $C_d = 400$ fF, $C_1 = 45$ fF.
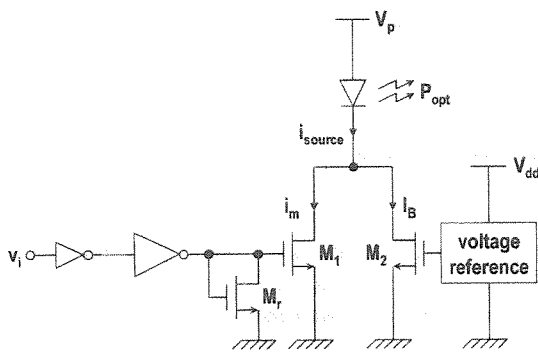

Fig. 7. Basic current modulation source driver circuit.

TABLE I
OPTICAL DEVICE PARAMETERS

	Parameter	Min. value	Max. value	Units
Source	Total efficiency	*0.1*	0.3	mW/mA
	Area	10x10	*100x100*	μm²
	Threshold current	1.5e-5	*1.5e-3*	A
	Bandwidth	10		GHz
	Parameter	**Min. value**	**Max. value**	**Units**
Detector	Responsivity	*0.5*	0.7	mA/mW
	Area	10x10	*100x100*	μm²
	Capacitance	1e-15	*1e-13*	F
	Dark current	1e-18		A
	Noise current	1e-15		A
	Bandwidth	20		GHz
	Parameter	**SOI**	**Si₃N₄**	**Units**
Waveguide	Guide index	3.45	2.0	
	Cladding index	1.46	1.46	
	Height	0.22	0.8	μm
	Width	0.5	0.4	μm
	Pitch	1.1e-6	4e-6	m
	Loss	2.7	1.5	dB/cm
	Delay	13.3	6.7	ps/mm
	Excess bend loss	0.027	1.52	dB/90°
	Bend radius	2	10	μm

Fig. 8. Optical link schematic.

microsource (with respect to commercial VCSELs) translate to lower parasitic capacitance and good bandwidth performance.

This circuit is sized simply by numerical bisection using the required values for the modulation current i_m and the bias current I_B (a unique solution exists).

IV. SIMULATION

To enable complete link simulation in an EDA framework, it is necessary to develop behavioral models for the optoelectronics devices and passive waveguides. For all behavioral models, the choice of an appropriate level of description is prerequisite to developing and using the models in the required context. Essentially, the description level falls into one of two categories: functional modeling or structural modeling. A functional model will describe the behavior of a device according to its specifications and behavioral equations, without defining the structure of the device. A structural model will describe the behavior of a device according to its internal structure and physical parameters, without necessarily satisfying the specification criteria (which do not have to be formalized in this approach).

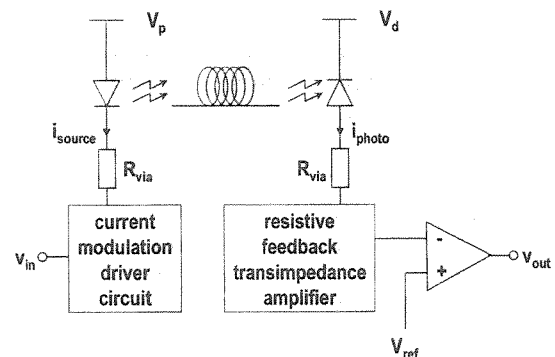Ideally, both functional and structural models should exist for the devices considered. However, this is not absolutely necessary and careful consideration was given to choosing the appropriate description level for each device. Since the source behavior is arguably the most complex and likely to exhibit nonlinear behavior (thermal roll-off, temperature changes) important to complete link simulation, it was decided to model this element at a structural level. The waveguide and detector were modeled at a functional level.

The models were all implemented in the OVI-96 Verilog-A subset of Verilog-AMS, an extension of the IEEE 1364–1995 Verilog hardware description language. This extension is an industry standard for analog simulation model description and can be simulated with a number of general-purpose circuit simulators (we use Spectre). This way the optical and photonic devices can be simulated together with the interface circuitry and with the rest of the optical link given adequate simulation models. This enables interesting optimization strategies (e.g., joint power optimization) and the analysis of link performance sensitivity to various parameter variations as well as temperature changes.

Detailed description of these models [20] is outside the scope of this paper, but the device parameters for optical interconnect varied in this analysis are shown in Table I, with minimum
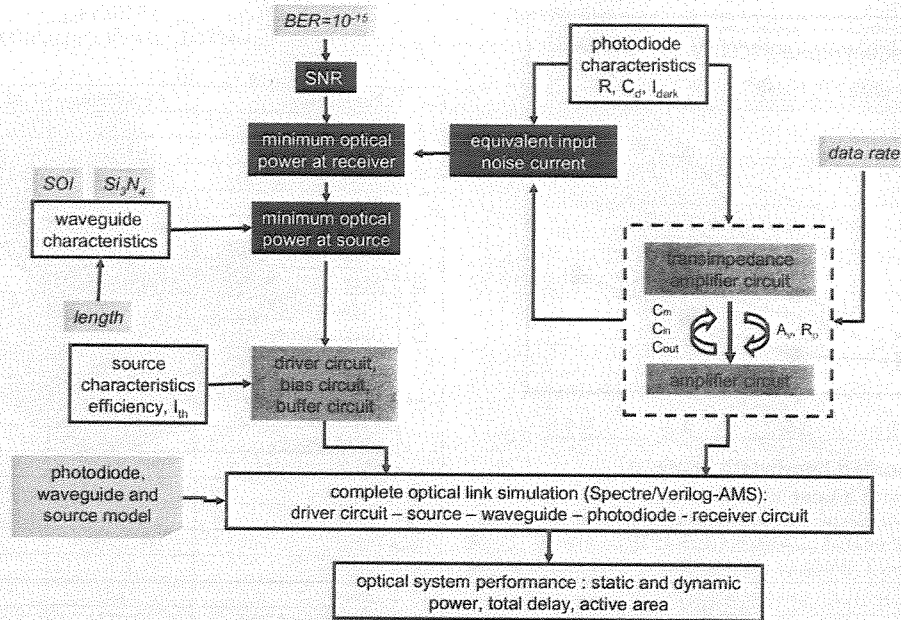
Fig. 9. Optical link sizing method.

and maximum values defining the limits of the parameter variation. These limits are based on discussions with experts in the field [21], and on data available in the literature on sources [8], waveguides [7], and detectors [9]. The values in bold italics represent the (pessimistic) nominal values.

This enabling design technology was essential to carry out simulation in an EDA environment of the complete optical link, associating these behavioral models with transistor-level circuit schematics. The nonlinear behavior of the microsource laser was modeled (enabling visualization of physical limits) and converges systematically in Spectre. Achieving complete link simulation was a necessary step to enable subsequent simulation-based link synthesis (using interface circuit design variables) over a range of target technologies and specification sets to extract link performance data. The iterative optimization step is facilitated by the low simulation time required for the complete link (a few seconds for ten data bits on a 1.3-GHz processor).

V. POINT-TO-POINT LINK SYNTHESIS

The objective of our paper was to carry out transistor-level sizing of the receiver and of the driver circuits according to complete link specifications. The optical link under consideration is represented by Fig. 8 (the CMOS structures used at transistor level are fixed in this synthesis procedure).

A. Synthesis Procedure

The synthesis approach implemented consists of creating scenarios allowing the specification of each model, evaluation, and design methods, and of communication between the different blocks using synthesizable AMS IP. A device library containing the synthesizable models of each device in the optical link based on the UML language was developed to allow the modeling of this hierarchical synthesis problem. The procedure used to au-

tomatically synthesize an optical point-to-point link is shown in Fig. 9.

The process starts by defining the photodetector characteristics and the required data rate. Using the method described in Section III-A, the transistor-level schematic for the transimpedance amplifier is automatically generated and linked to a manually scaled comparator circuit.

The value of the root mean square (rms) noise power i_n is extracted from simulation of the schematic and updated for each synthesis loop using the Morikuni formula [19] in the transimpedance amplifier noise calculations

$$i_N^2 = \left(2q\left(I_{\text{gate}} + I_{\text{dark}}\right) + \frac{4kT}{R_f}\right)\frac{C}{4D}$$
$$+ 4kTT\frac{C^2}{16\pi^2 DE}\frac{\left(2\pi C_T\right)^2}{g_m} \quad (4)$$

where

$$C = 1 + g_m R_f$$
$$D = R_o\left(C_x + C_y\right) + R_f\left(C_x + C_m\right) + g_m R_f R_o C_m$$
$$E = R_f R_o\left[\left(C_x + C_y\right)C_m + C_x C_y\right]. \quad (5)$$

For a given bit error rate (BER) specification and noise signal associated with the photodiode and transimpedance circuit, we can then calculate the minimum optical signal power ΔP_s required by the receiver to operate at the given error probability

$$\Delta P_s = i_n \times \text{SNR} \quad (6)$$

where

$$\text{SNR} = \left(\text{erfc}^{-1}\left(2\text{BER}\right)\right)^2. \quad (7)$$

Here SNR represents the linear signal-to-noise ratio (absolute value, not in decibels). BER, defined as the rate of error occurrences, is one of the main criteria in evaluating the performance

of digital transmission systems. In our analyses, we fixed BER at 10^{-18} bits^{-1} (this corresponds to 1 error/3.17 years for a single link at 10-Gb/s communication, or 1 error/18 days for a 64-bit data bus at the same data rate.)

The value of the power which needs to be emitted by the laser source is evaluated from the calculated value of the minimum optical power at the receiver, and from the power losses induced by the waveguide structure (length and intrinsic loss, number of bends and loss/90° bend) and coupling. These figures depend to a large extent on the materials used.

The final sizing step is to calculate the driver and associated bias and buffer circuits using the emitted power value and the source characteristics in conjunction with the method described in Section III-B. This then enables the simulation of the complete optical link, using transistor-level schematics for the interface circuits and behavioral models for the photonic devices. From the simulation results, the performance criteria can be extracted.

Using this approach, the synthesis problem is considered to be complete, such that no constraint partitioning is required. In fact, the constraints are derived directly from system specifications, and thus constraint exploration is achieved directly by the user, as shown in the investigation strategy detailed in Section VII.

B. Metrics

In order to be able to evaluate and optimize link performance criteria correctly, a clear definition of the performance metrics is required.

First, 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 as explained before, 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 photo-receiver front-end design methodologies.

For interconnect 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 photo-receiver: 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.

The limitations of this analysis are the following.
- PTM models do not take noise into account particularly well, which means that no real noise analysis can be carried out. However, in the sizing process, this problem was circumvented by using the Morikuni formula to estimate noise at block level.
- No automatic layout generation tools were used as it is not in general possible to achieve optimal layout for high-speed analog circuits. Parasitic capacitances were, therefore, extracted from layout estimations rather than from real layouts.

C. Specifications

Table II shows the sets of specifications used for analysis and interface circuit sizing. Three predictive technologies were considered for this analysis: PTM models [23] for 65-, 45-, and 32-nm gate length technology nodes.

VI. INTERCONNECT DENSITY ANALYSES

In this paper, we supposed that only one photonic routing layer would be available above IC. The impact of this hypothesis is that interconnect links can only be defined in a single routing direction (this does not exclude the occasional use of bends to avoid routing obstacles) due to unacceptable crosstalk (around 10% power loss) for straight waveguide crossings, and unacceptable area penalties (around $10 \times 10 \ \mu m^2$) for optimized crossings (1% power loss at 1.55-μm wavelength) [24].

Direct optical link bundles (see Fig. 10) were considered to be used for all inter-processor communication, except between physically adjacent processors (electrical buses used in this case). Normally of course, a design engineer would not create a fully connected architecture, but this is still a useful example to explore and quantify achievable interconnect density.

A number of parameters were defined in order to carry out interconnect density analyses, and are summarized in Table III.

If we assume that the overall chip is square, then the total chip size $= c_x^2$. If each IP block is also square then the IP block size $= c_x^2/p_x^2$. Since the coupling from sources to waveguides and from waveguides to detectors is electromagnetic, we used the

TABLE II
LINK SPECIFICATION SET

Parameter	Scenario	PTM65	PTM45	PTM32
BER	bits^{-1}	10^{-18}	10^{-18}	10^{-18}
ITRS max. frequency	bit/s	2.98×10^9	5.20×10^9	1.10×10^{10}
Link length	mm	{2.5,20}	{2.5,20}	{2.5,20}
Activity rate		1	1	1
Ambient temperature	°C	70	70	70
Vdd (CMOS)	V	1.2	1.1	1.0

TABLE III
OPTICAL INTERCONNECT ROUTING DENSITY PARAMETERS

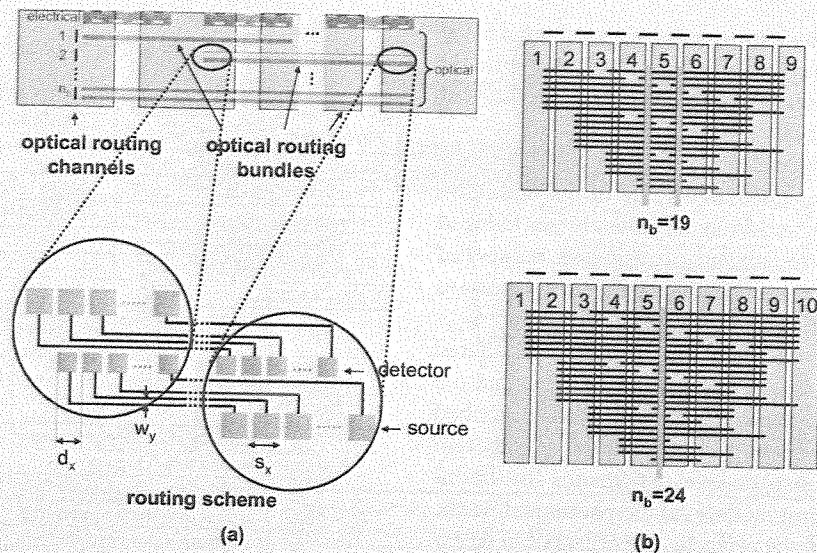|		
p_x	number of IP blocks in row	{2; 10}
c_x	side of total chip (m)	0.02
s_x	x-direction source pitch (m)	{10e-6; 100e-6}
s_y	y-direction source pitch (m)	{10e-6; 100e-6}
d_x	x-direction detector pitch (m)	{10e-6; 100e-6}
d_y	y-direction detector pitch (m)	{10e-6; 100e-6}
w_y	waveguide pitch (m)	{1.1e-6; 4e-6}
w_b	waveguide bend radius (m)	{2e-6; 10e-6}

Fig. 10. Routing scheme for processor matrix showing: (a) organization of routing bundles and channels and detail of source/detector matrices at each end of a bidirectional link and (b) example routing schemes and optical routing channel count for odd (9) and even (10) IP block number cases.

routing scheme shown in Fig. 10(a) to prevent any waveguide routing below an optoelectronic device and avoid in this way any potential crosstalk issues. If $s_x \neq d_x$, then using minimum device spacing will result in data skew and should be evaluated or avoided by nonoptimal spacing of the smaller device fixing $s'_x = d'_x = \max(s_x, d_x)$. This approach will also result in a delay penalty due to longer waveguide lengths.

For a given number p_x of IP blocks, we can then calculate the following:

- the number of optical routing channels necessary for total connectivity (excluding adjacent IP block communication which will be carried out by electrical interconnect);
- the maximum number of links in each bundle, based on waveguide pitch and optoelectronic device dimensions.

A. Number of Optical Bundles

The number of necessary optical routing channels can be calculated by examining the connection scheme between IP blocks and determining the number of channels at the plane where the maximum connection density occurs. A number of cases can be considered in order to derive a general model for the number of necessary channels, as shown in Fig. 10(b). In fact, two basic cases can be extracted: for an even number of IP blocks, a single dense plane exists, whereas for an odd number of IP blocks, two dense planes exist.

In general[3]

$$n_b = p_x(p_x - 2) - \sum_{i=2}^{p_x - 1} i - \sum_{i=2}^{\text{floor}(p_x - 1/2)} p_x - 2i. \quad (8)$$

This is shown in graphical form in Fig. 11.

B. Number of Optical Links Per Bundle

The total number of optical links in a row depends on the side of the total chip c_x, the number of rows (equal to p_x if the chip is square and the IP blocks are also square), the pitch between

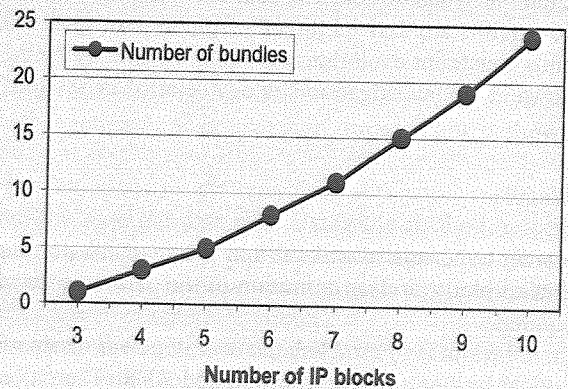[3]*Floor* is the function that rounds a real number to the integer closer to zero.

Fig. 11. Number of channels for varying number of IP blocks.

waveguides w_y, the source and detector pitches (s_x, s_y, d_x, d_y), and waveguide bend radius.

The maximum average number of aligned sources or detectors over the width of an IP block, n_{sd} is calculated as

$$n_{sd} = \text{floor}\left(\frac{\frac{c_x}{p_x}}{\frac{(s_x + d_x)}{2}}\right). \quad (9)$$

The maximum number of links in a row, n_l is calculated as

$$n_l = \text{floor}\left(\frac{\frac{c_x}{p_x}}{\max(s_y, d_y) + 2w_b + n_{sd}w_y}\right)^* n_{sd}. \quad (10)$$

The maximum number of bidirectional links in a bundle n_{lc} is calculated as

$$n_{lb} = \text{floor}\left[\frac{\text{floor}\left(\frac{n_l}{n_b}\right)}{2}\right]. \quad (11)$$

Fig. 12 shows a plot of the maximum achievable number of bi-directional links in a bundle for varying number of IP blocks in a row and varying source/detector sizes using this routing scheme. It can be seen for example that 64-bit, "crosstalk-free" optical communication buses between 8 IP blocks in a single
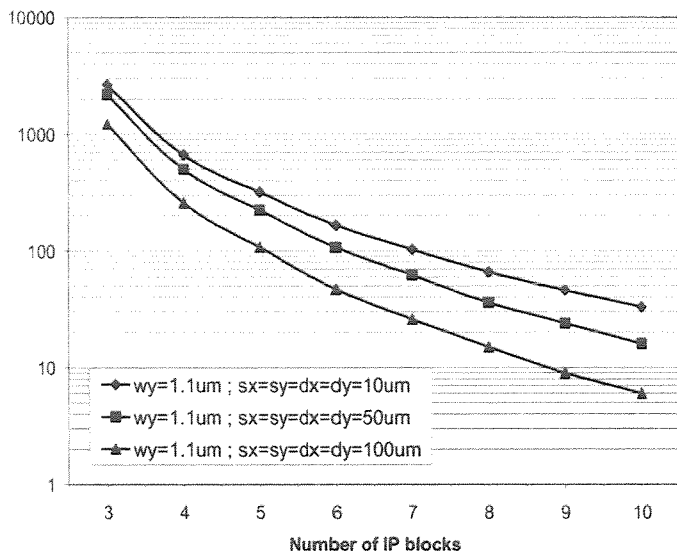
Fig. 12. Number of optical links per bundle for varying number of IP blocks (waveguide pitch $= 1.1 \ \mu$m).

row can only be achieved for source/detector sizes of around $10 \times 10 \ \mu$m^2 and waveguide pitch of $1.1 \ \mu$m.

C. Via Stack Calculations

In order to calculate the area penalty (at the CMOS level) of the via stack for passing a relatively high level of current (mainly driver modulation current to the source), we calculate the via matrix area A, defined by host technology characteristics and the maximum current passing through the stack. The area required by the via matrix is $A = N(a+b)^2$, where A represents area $(\mu$m$^2)$, N represents the number of contacts in the matrix, a and b represent minimum contact width $(\mu$m$)$ and spacing $(\mu$m$)$, respectively.

The overall surface area of a given metal, Na^2, is proportional to the maximum current $I_{\max}$, following $N = k_m I_{\max}/a^2$, where k_m is a constant associated with the via metal. Rewriting the equation for A

$$A = k_m I_{\max} \left(1 + \frac{b(2a + b)}{a^2} \right). \qquad (12)$$

Based on a set of industrial 130-nm gate length design rules for which damascene copper is used for all metal layers, we can empirically determine that for the upper-level vias (worst case) $k_m = 0.466 \ \mu$m^2/mA.

Even if a and b scale with shrinking process rules, it is trivial to show that a *does not* scale with the shrinking factor, as long as a and b retain the same ratio to each other. In the aforementioned 130-nm gate length process, $a = 0.36 \ \mu$m and $b = 0.54 \ \mu$m for via. Hence

$$A \left(um^2 \right) = 2.913^* I_{\max} (mA). \qquad (13)$$

For example, if $I_{\max}$ is situated at around 5 mA, the total via stack area as seen at the circuit level is $14.56 \ \mu$m^2.
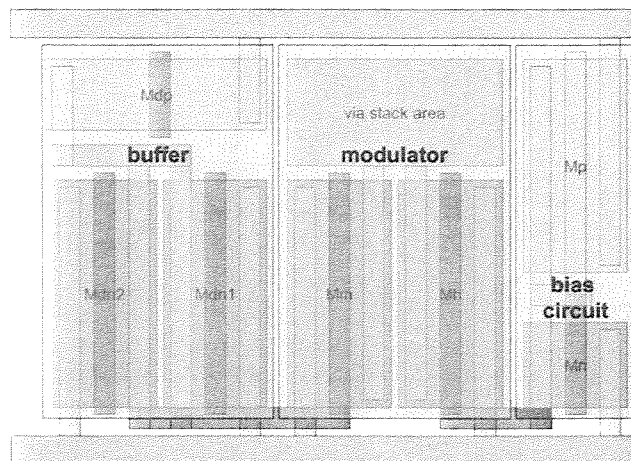

Fig. 13. Skeleton layout for current modulation driver circuit.

D. CMOS Interface Circuit Area Calculations

In order to validate the area work, the final step is to estimate the real circuit area necessary from gate area sizes. In the absence of satisfactory high-speed analog layout synthesis tools, our approach was to define a skeleton layout[4] based on worst case transistor sizes and derive technology independent layout rules from this. An important consideration in this work is the actual size of the via stack and its position in each interface circuit layout. The resulting estimated layout was used both to calculate circuit area and also to estimate parasitic capacitances on all circuit nodes.

Following transistor-level synthesis of the interface circuits under worst case conditions (65-nm technology for 20-mm link length, 1.5-mA source threshold current and 10% efficiency, 2.7-dB/cm waveguide loss, 100-fF photodiode capacitance and 50% responsivity), the total gate area of the interface circuits was calculated. The total circuit area was estimated from this using the skeleton layout template shown in Fig. 13 (including via stack area) and the lambda rules described in the Appendix.

For the receiver circuit, a similar analysis to that of the driver circuit was carried out using the skeleton layout shown in Fig. 14 to enable prediction of actual circuit size from transistor gate dimensions.

The total area was estimated using this method as $55.5 \ \mu$m^2. This worst case figure for overall CMOS area is considerably less than the area required for the smallest active device. As such, the active devices remain the limiting factor for area concerns rather than interface circuits.

VII. INVESTIGATION PROGRAM RESULTS

A. Gate Area Analysis for Varying Technologies

The link sizing method described in Section V was applied according to the specifications for the PTM 65-, 45-, and 32-nm technologies. Fig. 15 shows the results in terms of gate area (i.e., transistor channel dimensions only), extracted as the sum of all transistor gate channel areas $W \cdot L$. These results show that the gate area metric approximately verifies the scaling law

[4]A skeleton layout defines relative positions of transistors, without associating any information concerning actual transistor sizes.
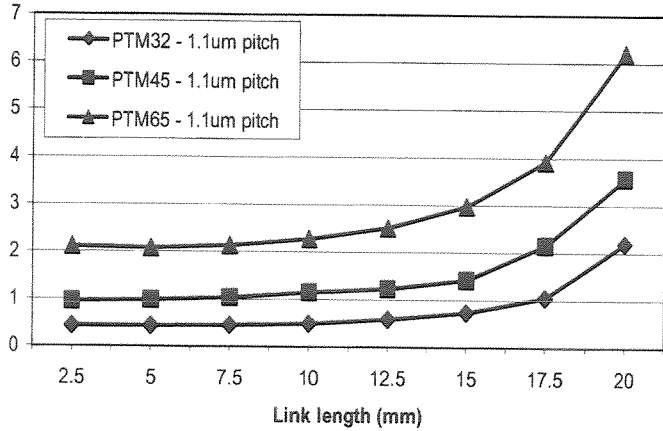
Fig. 14. Skeleton layout for receiver circuit.

Fig. 16. Total link delay (ps) for varying optical interconnect length and technologies.

Fig. 15. Total CMOS gate area (μm^2) for varying optical interconnect length and technologies.

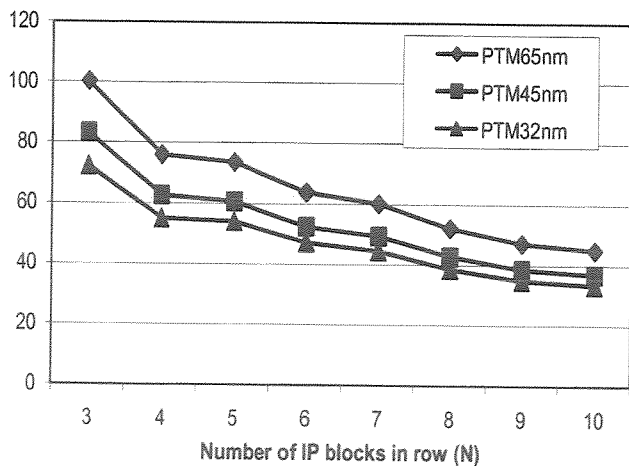Fig. 17. Average static power (in milliwatts) versus interconnect length for PTM 65-, 45-, and 32-nm technologies.

$(A_{32 \text{ nm}} \approx A_{45 \text{ nm}} \cdot s^2 \approx A_{65 \text{ nm}} \cdot s^2 \cdot s^2)$, where s is equal to 0.7 (scaling factor between technology generations).

B. Delay Analysis

The link sizing method described in Section V was applied according to the specifications for the PTM 65-, 45-, and 32-nm technologies. The 50% propagation delay was extracted from simulation, as well as a point of reference in the form of the intrinsic waveguide delay, using $t_{TOF} = 13.3$ ps/mm for the Si/SiO$_2$ waveguides[5] Fig. 16 shows the delay results for varying link lengths.

It can be seen that: 1) the *circuit* delay (i.e., the difference between the total delay and the intrinsic waveguide delay) decreases with smaller gate lengths and 2) the same quantity also decreases with longer interconnect. This latter effect is due to higher modulation current I_m, required to compensate higher overall waveguide loss, being able to drive the source capacitance faster.

[5]TOF: Time of Flight.

C. Power Analysis

The link sizing method described in Section V was applied according to the specifications for the PTM 65-, 45-, and 32-nm technologies. The average static power was extracted from transient simulations using

$$\bar{P} = \frac{I_{\text{source}_0} + I_{\text{source}_1}}{2} \cdot V_p + \frac{I_{\text{det}_0} + I_{\text{det}_1}}{2} \cdot V_d$$
$$+ \frac{I_{cct_0} + I_{cct_1}}{2} \cdot V_{dd} \quad (14)$$

where I_{source}, I_{det} and I_{cct} represent the currents flowing through the source, detector, and circuit voltage supplies of V_p, V_d, and V_{dd}, respectively. Fig. 17 shows the average static power results for varying link lengths.

The average switching energy is calculated from rising and falling edge transitions (extracted from simulations as the integral of supply currents in edge transitions). Fig. 18 shows the average switching energy results for varying link lengths.

Using static power and switching energy information, the total power can be calculated (see Fig. 19).

Fig. 18. Average switching energy (in picojoules) versus interconnect length for PTM 65-, 45-, and 32-nm technologies.

Fig. 19. Total power (in milliwatts) versus interconnect length for PTM 65-, 45-, and 32-nm technologies.

Fig. 20. Total power dissipation (W) for varying number of interconnected IP blocks.

The figures for aggregated maximum number of links per optical bundle and associated bundle lengths (including additional length due to device staggering) were then used in conjunction with the total power dissipation figures to calculate the total power in one row of interconnected IP blocks, for varying number of IP blocks (from 2 to 10) (see Fig. 20). It can be
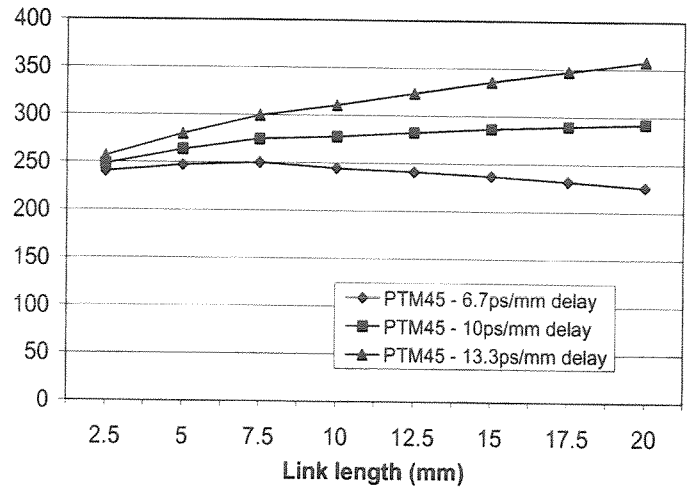
Fig. 21. Total link delay (in picoseconds) for varying interconnect length and waveguide intrinsic delay: PTM45 nm.

seen that the total power decreases with a rising number of IP blocks. This is due to the fact that, while there are more links, each link is shorter. Power increases exponentially with length and is dominated by the source power. For a 10×10 matrix of IP blocks, and for the most long-term technology node (32-nm printed gate length), this analysis gives a total power of 33.3 W (source threshold currents of 1.5 mA). Compared to the figure given by the ITRS for maximum power at this technology node (167 W), this works out as equal to 20% of total power.

These results show the capacity of the method to synthesize optical links subject to technological specifications, both CMOS and optical. The generation of each data point requires approximately 5 min on a 1.3-GHz processor.

VIII. DESIGN SPACE EXPLORATION

By carrying out variations on the photonic device characteristic set (within the bounds established in Table I) and rerunning the investigation program, it is possible to establish the impact of individual device characteristics on link performance. This section presents some conclusions from this analysis, as an illustration of the type of feedback our approach can give to photonic device engineers.

A. Improvement of Waveguide Materials

Overall link delay (see Fig. 21) can be reduced by about a factor of 2 for 20-mm interconnect length by shifting to lower index waveguide materials. Interestingly, it can be seen from the figure that the waveguide delay increase for longer lengths is compensated for by faster modulation with higher currents, leading in some cases to a *decrease* in interconnect delay with length. It should be possible, in this configuration, to optimize the system such that overall delay is constant for all interconnect length. This would enable synchronous data communication and reduce dependence on data recovery circuits.

B. Reduction of Source Threshold Current

Source threshold current has a significant impact on static power (see Fig. 22). Approximately 5 mW can be saved by
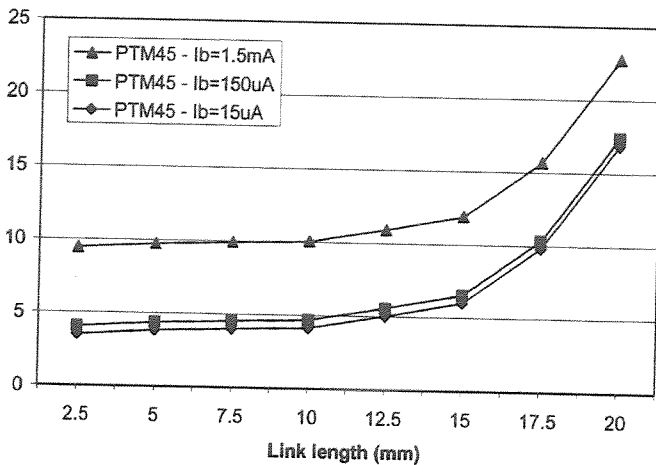
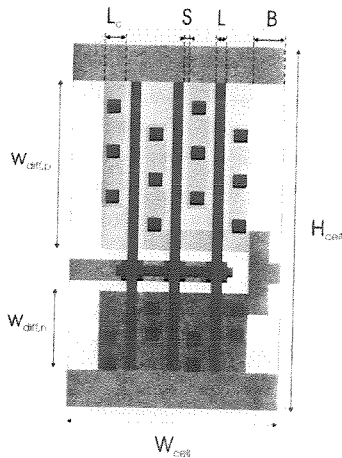Fig. 22. Average static power (mW) for varying interconnect length and source threshold current: PTM45 nm.

Fig. 23. Schematic λ-rule layout used for inverter active area and standard cell area approximations.

TABLE IV
SUMMARY OF PSEUDO-LAYOUT PARAMETERS AND EQUATIONS TO DERIVE CMOS CELL AREA ESTIMATES

Symbol	Interpretation	Value or equation
H_{cell} (λ)	Standard cell height	69
$W_{a,max}$ (λ)	Maximal total active width, i.e. maximal available cell height for MOSFET diffusion regions	47
W_n (λ)	NMOS gate width, expressed as a multiple of λ	Set by designer
W_p (λ)	PMOS gate width, expressed as a multiple of λ	Set by designer
F	Number of fingers in NMOS and PMOS transistors (equal)	ceil($(W_n+W_p) / W_{a,max}$)
$W_{diff,n}$ (λ)	Width of individual NMOS fingers, set equal to width of NMOS diffusion region.	W_n / F
$W_{diff,p}$ (λ)	Width of individual PMOS fingers, set equal to width of PMOS diffusion region.	W_p / F
L (λ)	MOSFET channel length	2
S (λ)	Spacing between gate poly and metal	1
L_c (λ)	Width of metal above diffusion region contact	4
B (λ)	Border at right and left edges of cell	6
$L_{diff,ext}$ (λ)	Length of MOSFET diffusion region at the edges	$L_c + S = 5$
$L_{diff,int}$ (λ)	Per MOSFET length of diffusion region if between fingers	$0.5 (L_c + 2S) = 3$
$L_{diff,tot}$ (λ)	Total length of diffusion region (for active area calculation)	$2L_{diff,ext}+2 (F-1) L_{diff,int}$ = 6F - 1
L_{act} (λ)	Total active area length	$L_{diff,tot} + F\ L = 8F - 1$
W_{cell} (λ)	Standard cell width	$L_{act} + 2B = 8F + 11$
A_{gate} (λ^2)	Total gate area	F L $(W_{diff,n}+W_{diff,p})$ = 2F $(W_{diff,n}+W_{diff,p})$
A_{active} (λ^2)	Total active area	$L_{act} (W_{diff,n}+W_{diff,p})$ =(8F-1)$(W_{diff,n}+W_{diff,p})$
A_{cell} (λ^2)	Total cell area	$W_{cell} H_{cell} = 69\ W_{cell}$

bringing the threshold current down to 150 μA from 1.5 mA. While the additional reduction down to 15 μA appears negligible in the context of overall link static power, it should be noted that, when multiplied by the total number of links on the chip, even a small individual power saving can represent a significant reduction at the chip scale

IX. CONCLUSION

This paper has described a methodology and toolset for the systematic design space exploration of integrated optical interconnect. In particular, it allows the automated design of optical links, resulting in a significant reduction in design time and, more importantly in this context, the capacity to carry out a detailed analysis of the available design space. This technology is particularly useful for repetitive design of fixed optical link structures subject to varying design constraints, technology parameters, and performance requirements. Further, as designer variability is removed from the overall design process, the resulting set of designs is likely to be relatively coherent and this facilitates interpretation of data and extrapolation to future design scenarios.

We have illustrated the direct application of our approach for link synthesis and technology performance characterization by analyzing optical link performance for a single set of photonic components and three CMOS technology generations. We have also hinted at the ways our toolset can be applied to generate useful feedback from system designers to component designers.

APPENDIX
λ-RULE SET FOR LAYOUT CALCULATIONS

In a λ-rule layout, all distances and design rules are expressed in so-called λ-units. This means that the layout grid unit is called λ, and it is set to half the minimal gate length ($\lambda = L/2$). All distances are expressed as a number of grid units. If the λ-based design rules are not made too tight, a λ-rule layout can be ported across several technology generations. In this paper, the specific design rules used were taken from an existing standard cell library and reexpressed as a function of λ.

For determining the finger widths, we have simplified our calculations by making all fingers equally wide. Hence, we are disregarding finger width restrictions from the layout grid, since we expect this will have very little impact on the simulation results.

As an example of how a scalable λ-rule layout template can be used to estimate layout area, we give a complete illustration for the case of a sample inverter (see Fig. 23). Table IV summarizes the design rules and equations that can be used to derive CMOS area estimates.

REFERENCES

[1] J. H. C. Collet et al., "Architectural approach to the role of optics in mono- and multi-processor machines," Appl. Opt., vol. 39, pp. 671–671, 2000.

[2] E. G. Friedman, "Clock distribution networks in synchronous digital integrated circuits," Proc. IEEE, vol. 89, no. 5, pp. 665–665, May 2001.

[3] M. Kobrinsky et al., "On-Chip optical interconnects," Intel Technol. J., vol. 8, no. 2, pp. 129–141, 2004.

[4] SEMATECH, Austin, TX, "International technology roadmap for semiconductors," 2005.

[5] J. H. Collet, F. Caignet, F. Sellaye, and D. Litaize, "Performance constraints for on chip optical interconnects," IEEE J. Sel. Topics Quantum Electron., vol. 9, no. 2, pp. 425–432, Mar./Apr. 2003.

[6] P. Rojo-Romeo, J. V. Campenhout, D. V. Thourhout, C. Seassal, P. Regreny, A. Kazmizerczak, C. Seassal, X. Letartre, G. Hollinger, D. V. Thourhout, R. Baets, L. D. Cioccio, and J. M. Fedeli, "'Heterogeneous integration of electrically driven microdisk based laser sources for optical interconnects and photonic ICs," Opt. Express, vol. 14, no. 9, pp. 3864–3864, 2006.

[7] A. Sakai, G. Hara, and T. Baba, "Propagation characteristics of ultra-high-Δ optical waveguide on silicon-on-Insulator substrate," Japan J. Appl. Phys.—Pt. 2, vol. 40, no. 383, pp. L383–L385, 2001.

[8] R. Fujita, R. Ushigome, and T. Baba, "Continuous wave lasing in GaInAsP microdisk injection laser with threshold current of 40 uA," Electron. Lett., vol. 36, no. 9, pp. 790–790, 2000.

[9] P. R. A. Binetti, X. J. M. Leijtens, M. Nikoufard, R. Orobtchouk, T. Benyattou, T. d. Vries, Y. S. Oei, and M. K. Smit, "A compact detector for use in photonic interconnections on CMOS ICs," in Proc. Symp. IEEE/LEOS Benelux Chapter, 2005, pp. 233–236.

[10] J.-M. Fedeli, R. Orobtchouk, C. Seassal, and L. Vivien, "Integration issues of a photonic layer on top of a CMOS circuit," SPIE Photon. Eur., vol. 6125, pp. 97–111, 2006.

[11] G. Gielen et al., "Analog and digital circuit design in 65 nm CMOS: End of the road," Proc. Design Autom. Test Eur. (DATE), pp. 36–36, 2005.

[12] I. O'Connor, F. Tissafi-Drissi, G. Revy, and F. Gaffiot, "UML/XML-based approach to hierarchical AMS synthesis," in Proc. Forum Design Lang, 2005, pp. 98–100.

[13] M. Hamour, R. Saleh, S. Mirabbasi, and A. Ivanov, "Analog IP design flow for SoC applications," in Proc. Int. Symp. Circuits Syst. (ISCAS), 2003, pp. IV–676.

[14] S. S. Mohan et al., "Bandwidth extension in CMOS with optimized chip inductors," IEEE J. Solid-State Circuits, vol. 35, no. 3, pp. 346–355, Mar. 2000.

[15] C. W. Kuo et al., "2 Gbit/s transimpedance amplifier fabricated by 0.35 μm CMOS technologies," IEE Electron. Lett., vol. 37, no. 19, pp. 1158–1158, Sep. 2001.

[16] J. Graeme, Photodiode Amplifiers. Boston, MA: McGraw-Hill, 1996.

[17] M. Ingels and M. S. J. Steyaert, "A 1-Gb/s, 0.7 μm CMOS optical receiver with full rail-to-rail output swing," IEEE J. Solid-State Circuits, vol. 34, no. 7, pp. 971–971, Jul. 1999.

[18] I. O'Connor, F. Mieyeville, F. Tissafi-Drissi, G. Tosik, and F. Gaffiot, "Predictive design space exploration of maximum bandwidth CMOS photoreceiver preamplifiers," in Proc. IEEE Int. Conf. Electron., Circuits Syst. (ICECS), 2003, pp. 483–486.

[19] E. Kyriakis-Bitzaros et al., "Particular aspects of drivers for VCSELs operating at multi-Gb/s," J. Semicond. Technol. Sci., vol. 2, no. 1, pp. 82–86, Mar. 2002.

[20] I. O'Connor et al., "D1.2. definition and quantification of specifications for different subcomponents," Deliverable D1.2, 2006 [Online]. Available: http://picmos.intec.ugent.be

[21] C. S. Seassal et al., "Comparison of proposed architectures for μ-source and choice of design to be used for optical link," Deliverable D3.2, 2005 [Online]. Available: http://picmos.intec.ugent.be

[22] J. J. Morikuni, A. Dharchoudhury, Y. Leblebici, and S. M. Kang, "Improvements to the standard theory of photoreceiver noise," J. Lightw. Technol., vol. 12, no. 7, pp. 1174–1184, Jul. 1994.

[23] Y. Cao, T. Sato, D. Sylvester, M. Orshansky, and C. Hu, "New paradigm of predictive MOSFET and interconnect modeling for early circuit design," in Proc. Custom Integr. Circuits Conf., 2000, pp. 201–204.

[24] T. Fukazawa et al., "Low loss intersection of Si photonic wire waveguides," Japan J. Appl. Phys., vol. 43, no. 3, pp. 646–647, 2004.

Ian O'Connor (S'95–M'98) was born in Cambridge, U.K., in 1969. He received the European M.Sc. degree in electronics engineering from the University of Essex, Essex, U.K., the Ph.D. degree in electronics from the University of Lille, Lille, France, in 1997, and the professoral dissertation (Habilitation à Diriger des Recherches) from Ecole Centrale de Lyon, Ecully, France, in 2005.

He joined Ecole Centrale de Lyon in 1998 as an Associate Professor, and is currently head of the Design Group at the Lyon Institute of Nanotechnology, of which he is also one of the vice-directors. From 1993 to 1997, he was a Research and Teaching Assistant with the Institut Supérieur d'Electronique du Nord, Nord, France, and from 1997 to 1998, he was a Senior CAD Engineer with Philips Semiconductors, Southampton, U.K. His research interests include design methods and tools for physically heterogeneous systems on chip, and their application to novel architectures based on nonconventional devices. He has authored or coauthored well over 50 book chapters, journal publications, and conference papers and has presented invited papers at several major conferences.

Faress Tissafi-Drissi was born in Fes, Morocco, in 1977. He received the Ph.D. degree in microelectronics from Ecole Centrale of Lyon, Ecully, France, in 2004.

He is currently with Dolphin Integration, Meylan, France, as a Mixed-Signal Circuit Designer for systems on chip. He works particularly on the development of embedded memories for portable applications and scientific research in nanophotonics–optoelectronics areas.

Frédéric Gaffiot was born in 1959. He received the M.S. degree in applied physics from Ecole Normale Supérieure de Cachan, Cachan, France, and the Ph.D. degree in microelectronics from Ecole Centrale de Lyon, Ecully, France, in 1987.

He is currently with Ecole Centrale de Lyon. He was primarily an Associate Professor (1985–2000) and has been a full-time Professor since 2001. His research interests involve analog and mixed-signal circuit design, behavioral modeling, multiphysics design, optical interconnects, system on chip design, and wireless sensor networks.
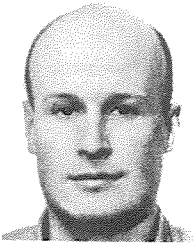
Joni Dambre (S'99–M'04) was born in Gent, Belgium, in 1973. She received the M.S. degree in electrotechnical engineering and the Ph.D. degree from Ghent University, Gent, Belgium, in 1996 and 2003, respectively.

She is currently a Postdoctoral Researcher with the Department of Electronics and Information Systems (ELIS), Ghent University. Her research interests include platform and technology evaluation, system level design and prediction techniques, and novel interconnection technologies.

Dr. Dambre is a member of ACM.

Michiel De Wilde was born in Leuven, Belgium in 1977. He received the M.S. degree in computer engineering and the Ph.D. degree in electronics and information systems from Ghent University, Ghent, Belgium, in 2000 and 2007, respectively.

His research interests include design, modelling and system-level evaluation of short-range parallel optical interconnects, quantification of noise phenomena in integrated circuits, evaluation of FPGA architectures, and design of spatial indexing methods.
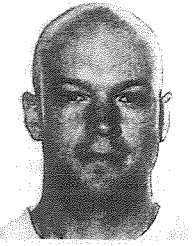
Joris Van Campenhout (S'01) received the M.S. degree in physics engineering from Ghent University, Gent, Belgium, in 2002, where he is currently pursuing the Ph.D. degree in electrical engineering.

His research interests include the design, fabrication, and characterization of silicon-based photonic integrated circuits. Over the last years, he has been working on the heterogeneous integration of InP-based microcavity lasers—based on photonic-crystal and microdisk structures—on a silicon platform.

Dries Van Thourhout (S'99–M'00) received the M.S. degree in physical engineering and the Ph.D. degree from Ghent University, Gent, Belgium, in 1995 and 2000, respectively.

From October 2000 to September 2002 he was with Lucent Technologies, Bell Laboratories, Crawford Hill, NJ, working on the design, processing, and characterization of InP/InGaAsP monolithically integrated devices. In October 2002, he joined the Department of Information Technology (INTEC), Ghent University, where he is currently a member of the permanent staff of the Photonics Group. His research interests focus on the design, fabrication, and characterization of integrated photonic devices. Main topics involve silicon nanophotonic devices, heterogeneous integration of InP-on-silicon, and integrated InP-based optical isolators. He is also working on the development of new fabrication processes for photonic devices, e.g., based on focused ion beam etching and die-to-wafer bonding. He holds three patents, has authored and coauthored over 30 journal papers, and has presented invited papers at several major conferences.

Jan Van Campenhout (M'95) was born in Vilvoorde, Belgium, in 1949. He received the M.S. degree in electromechanical engineering from Ghent University, Gent, Belgium, in 1972, and the M.S.E.E. and Ph.D. degrees from Stanford University, Stanford, CA, in 1975 and 1978, respectively.

He is currently a Professor teaching courses in computer architecture, electronics, and digital design with the Faculty of Applied Sciences, Ghent University. His current research interests include the study and implementation of various forms of parallelism in computer systems, and their application in programming language support, computer graphics, and robotics. He is currently the head of the ELIS Department in the Faculty of Engineering.

Dr. Van Campenhout is a member of Sigma Xi, KVIV, and ACM.

Dirk Stroobandt (S'91–M'98) received the M.S. and Ph.D. degrees, both in electrotechnical engineering, from Ghent University, Gent, Belgium, in 1994 and 1998, respectively.

Since October 2002, he has been a Professor with the Electronics and Information Systems Department, Ghent University. His research interests include *a priori* estimations of interconnection lengths in electronic systems and its applications to CAD, computer architecture evaluation, and design optimization.

SPECIAL SECTION ON SYSTEM-LEVEL INTERCONNECT PREDICTION

◆IEEE