Integration of Modeling Tools for Parallel Optical Interconnects in a Standard EDA Design Environment

Michiel De Wilde*[a], Olivier Rits[b], Wim Meeus[a], Hannes Lambrecht[b], Jan Van Campenhout[a]
[a]ELIS/[b]INTEC department, Ghent University, IMEC

Abstract

Optical interconnects on the PCB or rack level offer a way to overcome fundamental bandwidth limits associated with electrical interconnects. The optimal design of a system using guided-wave CMOS-to-CMOS parallel optical interconnects with flip-chipped optical devices requires adequate EDA support. In this paper, we present an integration and simulation approach for interconnects in system design tools. Important stochastic effects affecting the performance and reliability of the interconnect are furthermore treated. To conclude, we present our measurement setups for interconnect timing behaviour, error rate measurement and substrate noise sensitivity.

1 Introduction

The commercial deployment of guided-wave optical chip-to-chip interconnects where optical components are directly hybridized with CMOS is currently being investigated. The support from EDA design tools for the integration of such optical interconnects in a system design is still rather limited. In section 2, we discuss a clean way to integrate optical interconnect in the design flow imposed by EDA tools and the analog circuit-level simulation models that are required for this purpose.

Simulation of full optical links by combining models for all link parts can be useful for extracting and even optimizing important system-level link parameters, e.g., link power consumption or timing information. The latter information is necessary for digital simulation of the interconnect within a digital system where timing figures are annotated to the simulator.

For timing verification purposes, minimal and maximal timings for the interconnect are additionally required. Values for skew and jitter as well as estimated bit error rates cannot be extracted from a simulation of only the typical behaviour. To this end, we will need to incorporate unwanted stochastic effects in our modelling approach. In section 3, an overview is given of effects relevant for parallel optical interconnects.

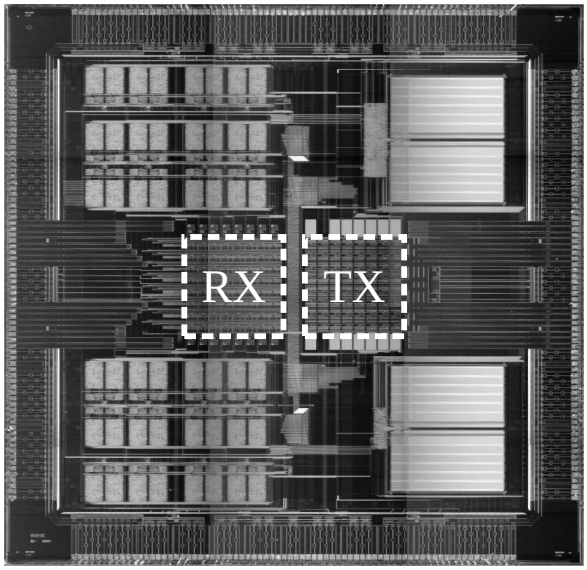

Figure 1: First test IC in 0.35μm CMOS. The photodiode and VCSEL arrays are to be flip-chip mounted onto the marked areas.

To verify the accuracy of our modelling, we will perform measurements on a demonstrator system for parallel optical interconnect. We discuss two setups that we have designed: one will allow the monitoring of individual link latency, skew, jitter and bit error rate, while another setup provides circuits for generating and measuring substrate noise and its impact on the bit error rate of the optical interconnect.

*Michiel De Wilde is a Research Assistant of the Fund for Scientific Research – Flanders (Belgium) (F.W.O.)

Research supported by the Interconnect By Optics project, IST-2000-28358 http://www.intec.rug.ac.be/io/

2 EDA Integration

To measure the quality of EDA integration of parallel optical interconnects, we can often compare the ease with which optical versus electrical chip-to-chip interconnections can be manipulated in these tools.

2.1 On-chip Interface

The first occurrence of physical inter-chip interconnect in the design flow is the instantiation of the external interface when designing an IC. For electrical interconnect this is the ring of bond pads, while for parallel optical interconnect this is the area where electro-optical conversion components are to be flip-chip mounted. These opto-mounting areas contain flip-chip mounting pads and driver and receiver circuits for arrays of electro-optical conversion components (figure 1).

IC cell library Just like electrical pads, the CMOS data of an opto-mounting area can be delivered to an IC designer as a design kit library with cells having a symbolic view, a layout view and a digital and analog simulation view. For IP protection reasons, the layout view can be reduced and the analog simulation model can be delivered as a precompiled netlist or a simplified behavioural model. In the latter case, driver and receiver circuits are broken down into parts, the behaviour of which can be approximated by a simple description in a behavioural language like Verilog-AMS or VHDL-AMS. Figure 2 shows the major parts of a receiver circuit.

Mixed-signal simulation EDA tools that can be configured to simulate a given design on a digital, analog or mixed-signal level choose either the analog or digital views of cell instances during hierarchical netlisting. This requires corresponding interface ports for analog and digital views of a cell—only global nets in the analog view need not to exist in the digital view. Fortunately, the analog model of an opto-mounting layout can be described by addressing only one electrical terminal per laser or photodiode device, as the other terminal is typically connected to a global net (ground or a power supply). This one-to-one relationship between input and output terminals makes the digital view easy: simple single signal propagation (with appropriate delay) suffices.

AC-JTAG IEEE 1149.1 boundary scan cells are typically used in pad cells for electrical interconnect to enable testing of chip logic and PCB interconnect. In the Interconnect by Optics project [1], our project partner Helix has also equipped driver and receiver circuits with boundary scan cells for configuration and testing purposes. Since the automatic gain control in the receiver circuits requires the optical signal to be DC-balanced, a test signal that alternates with the test clock is used during interconnection testing (hence the name *AC*-JTAG).

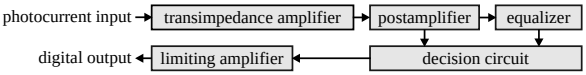

Figure 2: Building blocks of a photodiode receiver that can be easily modelled in analog HDL language.

2.2 Electro-optical conversion devices

Photonics in circuit simulators The analog HDL languages VHDL-AMS and Verilog-AMS explicitly support quantities other than voltages and currents. Light power and other analog natures can be represented in two ways: either as a value on an interface port (like a potential), or as a current between interface ports, a light source and a light sink. At first sight, the latter choice seems the most logical one. There is a problem though: the laws of Kirchhoff—enforced by circuit simulators on current-like branches—don't apply to optics. Light gets lost in media, and photons travelling in opposite senses through a waveguide do not compensate for each other like opposite currents in a wire. Therefore it is easier to represent light as a value on an interface port, and explicitly provide for light propagation in simulation model descriptions. Also note that bidirectionally used waveguides need separate interface ports for each direction.

VCSELs VCSELs *(Vertical Cavity Surface-Emitting Lasers)* are popular devices for parallel optical interconnect, as they can be directly produced in 2D arrays. Being semiconductor laser devices, they exhibit a very complex and dynamic behaviour in the interaction between carriers and photons. The rate equations of Moriki et. al. [2] (figure 3) describing these interactions require 3D spatial integration over the active region of the device, and are as such unsuitable for circuit-level simulators.

Starting from Moriki's description, Mena, Morikuni, et. al. have created a simpler model where some acceptable simplifications allow the separation of spatial and temporal integrations [3]. This model is further improved by Jungo, et. al. [4]. The spatial integration can then be done statically, yielding only temporal differential equations, which are suitable for circuit simulators. In this model, the optical output of a VCSEL is approximated as the excitation of light power in a few emis-

$$\frac{\partial N(\vec{r},t)}{\partial t} = \frac{\eta_i I(\vec{r},t)}{q} - \frac{N(\vec{r},t)}{\tau_n} - \sum_k G_k(\vec{r},t)\, S_k(t)\, \psi_k(\vec{r}) + \frac{L_{eff}^2}{\tau_n}\nabla^2 N(\vec{r},t) - \frac{I_l(N,t)}{q}$$

$$\frac{\partial S_k(t)}{\partial t} = -\frac{S_k(t)}{\tau_{pk}} + \frac{\beta_k}{\tau_n}\cdot\frac{1}{V}\int_{V_T} N(\vec{r},t)\,dv + \frac{1}{V}\int_{V_T} G_k(\vec{r},t)\, S_k(t)\, \psi_k(\vec{r})\,dv$$

Figure 3: VCSEL rate equations. N represents the carrier number in the active region and S_k and ψ_k are the photon number and transverse mode profile in the kth transverse mode, respectively. For a detailed discussion of all parameters, see [3].

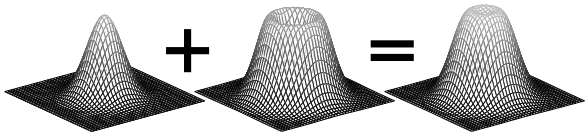

Figure 4: Far-field excitations of VCSEL mode profiles can be added together

sion profiles, corresponding to the different transverse modes of the VCSEL (figure 4). The model is rather complete, as all important effects like self-heating and spatial hole burning are included.

In our description in Verilog-AMS [5] this results in a device with an electrical cathode (ground) and anode (electrical input), and an array of optical output terminals to represent the light power emitted in each VCSEL mode. In an abstracted digital view, the digital signal enters the cathode and propagates to the first optical output, corresponding to the fundamental mode—the mode that emits the most light.

The electrical behaviour of a VCSEL is just a diode model with some associated parasitics and can hence be natively represented. Implementing the ordinary differential equations describing the optical behaviour should have been easy in an analog HDL. Nevertheless, we had to overcome a few problems:

- Directly solving the rate equations for steady state behaviour yields multiple solutions. This made the simulator fail at finding an initial operating point. The suggestion[3] to replace every positive-valued quantity X by a new quantity $Y = (X+C)^2$, solves this problem. Nevertheless, the Newton-Rhapson method that computes the initial operating point still failed for driving currents well above threshold. We alleviated this problem by providing a starting value for all carrier/photon-related quantities. Other approaches apply an approximated steady state solution with better convergence properties [6, 7].
- Some quantities that appear in the rate equations give rise to very large values, causing for instance Cadence Spectre (or NC-Sim) to panic during simulation, even if we configured it to accept high values for the guilty datatypes. This was resolved by rescaling all quantities to magnitudes that are common to voltages and currents.

Photodiodes A photodiode subjected to a reverse voltage bias converts incident light power to electrical current by an equation of the first degree. Only the dark current (Y axis intercept) and the responsitivity (slope) are relevant for the conversion of light. They are a function of the bias and the temperature. The electrical parasitics between the forthcoming conversion current and the electrical terminals are large enough to mask the effects caused by the exact impact location of photons on the photodiode.

The photodiode model therefore needs only one optical terminal. The interface and digital abstraction are otherwise similar to the VCSEL case, albeit with opposite signal propagation.

Parameter extraction The model parameters of a photodiode are generally extracted with few measurements or can be directly read from device datasheets. Obtaining VCSEL model parameters by measurements is a lot more complicated due to a strong nonlinearity, self-heating, multiple modes and the vast amount of model parameters (more than 60 for only two modes). VCSEL datasheets generally quantify most interesting device properties, but do not carry enough information to get any simulation model running.

We are currently working on a parameter extraction approach that does not require any knowledge of the VCSEL production process, but derives the necessary data from measurements. Measurements are hereby performed at different ambient temperatures and can include steady state light-current-voltage curves, scattering parameters, the optical spectrum, and the near/far field.

To achieve a better integration of VCSELs in EDA tools, a standardized VCSEL model for circuit simulation should be established in the future. Parameter extraction is a tiresome job and should ideally be performed only once, by the VCSEL supplier. This

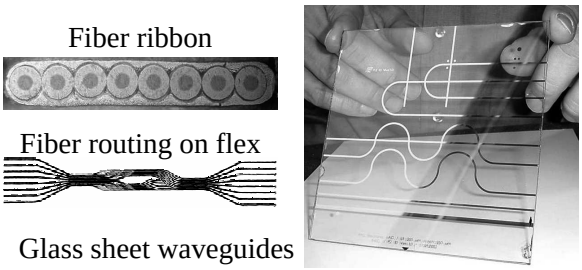

Figure 5: Different waveguide approaches

would furthermore enable an easy VCSEL comparison and the exchange of VCSELs in the simulation of a design. A good role model for this model adoption would be the BSIM series of field-effect transistor models [8], where a complex simulation model and parameter extraction procedure are clearly documented and well adopted.

2.3 Optical path

In guided-wave parallel optical interconnect, the optical path consists of multimode waveguides interconnecting VCSELs and photodiodes. We can discriminate two waveguide approaches (figure 5): routed fiber, where plastic optical fibers (POFs) are bundled by ribbonisation or flex routing, and an integrated approach, where waveguides are produced in a transparent printed circuit board (PCB) layer.

The physical design of the waveguide routing requires a substantial extension of the functionality of traditional PCB design tools. Different experimental technologies require different specialized software support, and a detailed discussion would be beyond the scope of this text. Instead we focus here on a simulation model for waveguides.

Simulation The circuit-level model for the optical path has to connect the outgoing mode excitations of the modelled VCSELs to the optical input of the modelled photodiodes. The propagation of each VCSEL mode profile through the entire optical path can be calculated independently of all other VCSEL modes. At the input port of each modelled photodiode, the contributions of each propagated VCSEL mode can therefore be simply added.
For the optical path model, we can distinguish between two categories of light propagation alongside the path:

- Interface jumps: the input or output coupling of a waveguide and connector interfaces
- Dispersion and losses in waveguide caused by

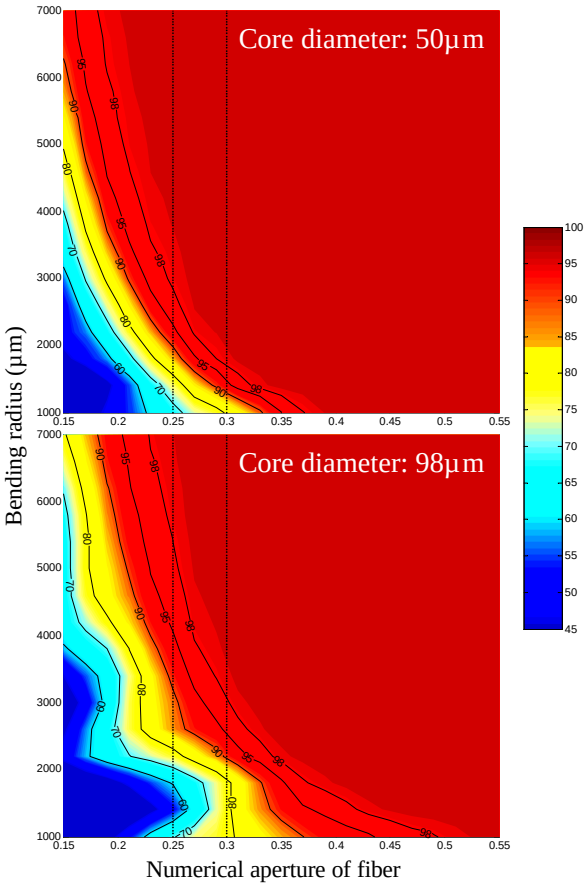

Figure 6: Transmission loss of a step index fiber after a 90 degree bend in function of bending radius and numerical aperture. The input ray pattern is the far field characteristic of a typical Avalon Photonics VCSEL.

curves, rough sidewalls and light absorption

For large interface jumps, like the coupling of a VCSEL mode profile in a waveguide, we calculate the diffraction pattern of the emission on the waveguide interface. The Fresnel diffraction formula is applicable as we are generally still in the near field of the VCSEL. For a fiber-photodiode coupling we can skip the calculations if all light reaches the photodiode as the light profile is not needed at that point for further calculations.

The dimensions of multimode waveguides are large enough to apply a ray tracing model for light propagation. Figure 6 illustrates the losses of a step index fiber in function of its numerical aperture and bending radius for a 90 degree bend directly after a VCSEL-fiber interface.

The waveguide modelling that we now use is ade-

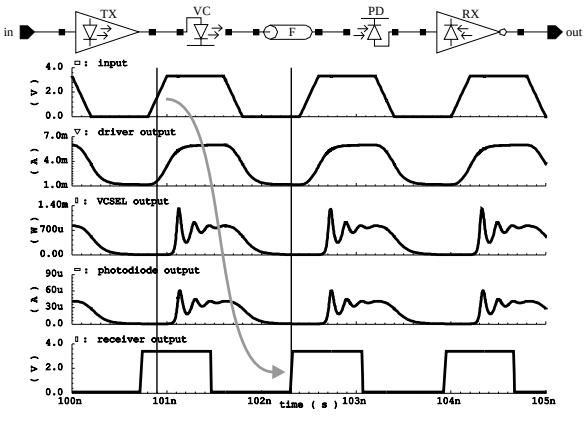

Figure 7: Illustration of an analog simulation of an optical link (using fictional model parameters)

quate for graded-index fiber. Due to its parabolic index profile, dispersion in the light propagation is limited to some fraction of a picosecond for a short optical path (less than a meter). The dispersive effect can be safely neglected in this case. However, for a step index fiber of 1m length, a numerical aperture of 0.25 and a core index of 1.5 the difference in propagation between the shortest and longest meridional ray amounts to a significant 164ps. For this kind of waveguide, or almost all integrated approaches, dispersion modelling will be required. Recent work by Gerling et. al. [9] allows to model this dispersion as a filter in a circuit simulator; the step response of the waveguide is hereby approximated as a sum of a few exponentials, resulting in a compact representation in analog HDL languages.

In the digital abstraction of the optical path, always the first VCSEL output terminal–carrying the digital signal—is propagated to photodiode input ports.

2.4 Simulation of full interconnects

The combination of circuit-level models for all parts of the parallel optical interconnect allows a circuit simulator to calculate the timing behaviour of quantities anywhere in the optical interconnect. Figure 7 illustrates an analog simulation with fictional model parameters.

Link optimization The interconnect simulation can be used to extract important system-level characteristics, e.g., power consumption or the latency of an electrical-optical-electrical path as a function of drive currents or the path length. This additionally results in a methodology for the optimisation of the parallel optical interconnect. The sensitivity of important system-level characteristics to incremental model parameter changes is hereby approximated using simulations. Furthermore, using parameter extraction, the sensitivity of the model parameters in their own right is calculated as to changes in the interconnect constitution, e.g., fiber numerical aperture or VCSEL-fiber distance. If both sensitivity analyses are combined to directly determine the dependency of system level characteristics on changes in the interconnect constitution, sets of changes with a global positive effect on the interconnect may be found.

Digital simulation with timing annotation For the different components of the optical interconnect, we have always provided both an accurate analog and abstract digital view. This allows the designer, as in design kits for electrical interconnect, to choose between either analog or digital simulation without extra work.

The digital simulation in its own is quite trivial: only a direct propagation of digital signals is done. We can however attribute minimal, typical and maximal delays to the interconnect through backannotation in Standard Delay Format (SDF) files. Tools for static timing verification can use these data to verify the correct operation of the digital design. The typical delay is easily extracted from a digital simulation. However, to be able to extract minimal and maximal timings we will need to do some additional modelling work: deviations from typical behaviour need to be taken into account. In the following section, an overview of the most important random effects is given.

3 Random Effects in Parallel Optical Interconnect

Most random effects related to optical links result in a change of the bit error rate (BER), latency, skew and jitter. From the point of view of a systems designer, it is important, besides an acceptable BER, to obtain low skew and jitter: the relation of their size to the signalling rate determines the complexity of power and area hungry circuits for timing correction.

3.1 Overview of Random Effects

Effect classification Random effects on guided-wave parallel optical interconnect can be divided in the following categories:

- **Static deviatons** are time-independent differences between instances of the same design, or between component instances in one system. They include process fluctuations in the processing of CMOS, VCSELs, photodiodes and waveguides and mechanical tolerances on connectors and mechanical interfaces. Resulting effects are changes in latency and signal skew. For skew reduction,

it is important that these random deviations occur as uniformly as possible over the parallel interconnect.

- **Random noise processes** are disturbances of a fundamental physical nature. They cause an increase of BER and jitter figures. The VCSEL exhibits relative intensity noise (RIN) and phase noise, both resulting in intensity noise after a dispersive waveguide. The input of the receiver circuit is also subject to thermal noise.

- **Coupled noise** is the collection of disturbances caused by the inadvertent coupling of signals. We distinguish between optical crosstalk, ground and power supply noise, and substrate noise.

Optical Path Misalignment Misalignment is the error on the position of a waveguide at its interface with another waveguide or an electro-optical conversion device. A slight misalignment readily causes a bad waveguide coupling. A larger error can even result in optical crosstalk, when inadvertent coupling through the wrong waveguide occurs, yielding an increase of BER and jitter.

The optical loss caused by a bad waveguide coupling increases the skew, as at the receiving end the latency of the signal reconstruction circuitry can shift over some hunderds of picoseconds depending on the incident optical power. Maximal allowable losses in the optical path are of the order of 10dB. A XY-misalignment of only 10μm of a fiber-VCSEL coupling at a distance of 120μm results in a loss of 0.5dB (figure 8).

Both with POF and integrated waveguides, a global error on the positioning of the fiber-device interface impacts all links simultaneously. The generic misalignment with six degrees of freedom creates a gradient of losses over the parallel links, resulting in skew. In the POF approach, the termination of individual fibers in connectors is also subject to a heap of mechanical tolerances: the orientation of the POF axis, the perpendicularity of the fiber butt, the protrusion out of the connector and the eccentricity of the fiber core centre.

Temperature Sensitivity of VCSELs VCSELs are rather temperature dependent: the change in optical output power is of the order of a significant 0.1dB per degree centigrade. Self-heating of VCSELs is already provided for in the simulation model, and the ambient temperature range can easily be swept by a simulator. An additional effect that can be significant for parallel interconnects is an unbalanced heating of the VCSEL

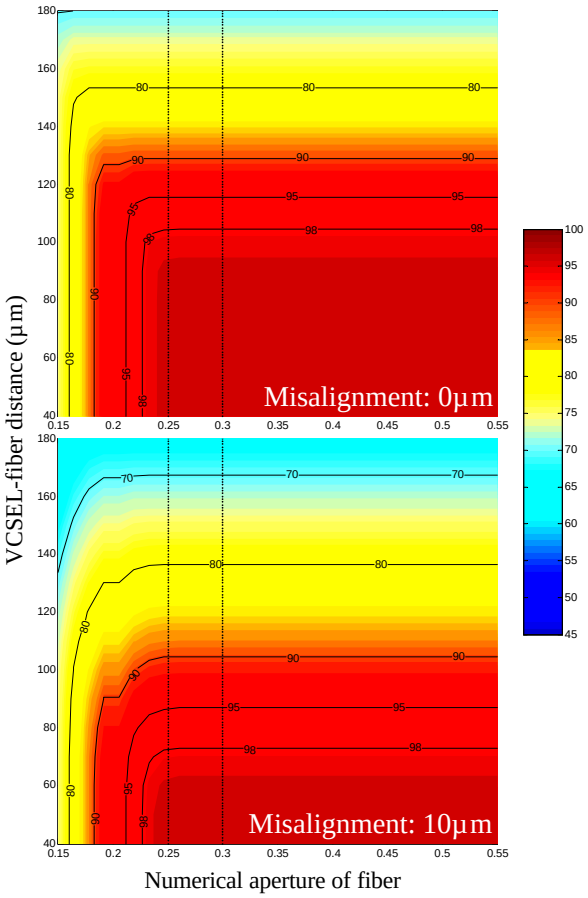

Figure 8: Coupling efficiency between a VCSEL with a diameter of 8μm and a fiber with a core size of 50μm in function of fiber numerical aperture and VCSEL-fiber distance.

array, where the central devices become hotter than the outer ones, yielding more signal skew.

Electrically Coupled Noise The opto-mounting layouts, containing analog driver and receiver circuits for the electro-optical conversion devices, are typically implemented in the middle of an otherwise digital IC. This is a very hostile operating environment: the simultaneous switching noise (SSN) of the digital gates can cause disturbances in the analog circuits. Especially the transimpedance amplifier of the receiver circuit has to deal with very small photocurrents (around 50μA).

Digital noise can reach the analog circuits through power lines or the substrate. Supply noise contains SSN voltage drops, caused by inductance and resistance in the power supply connections, and voltage ringing between this inductance and on-chip capacitance. Propagation of this noise to the power supply

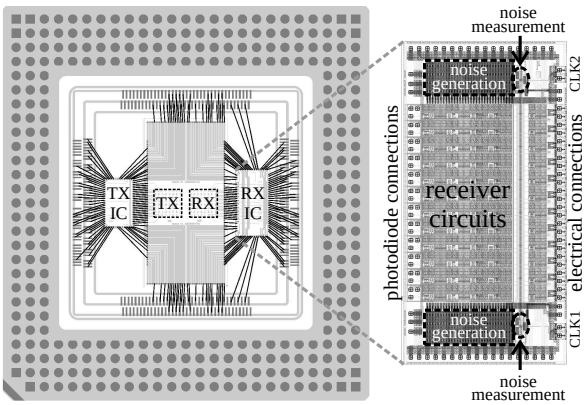

Figure 9: Test assembly with driver and receiver ICs in 0.18μm CMOS. The receiver IC is augmented with substrate noise generation and measurement circuits.

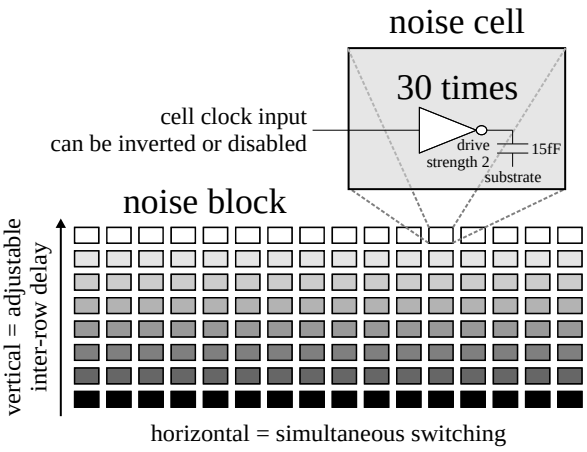

Figure 10: Noise block

of the analog part can be eliminated if separate supply lines are provided for analog and digital circuits.

The substrate, however, is coupled to the digital ground net through substrate contacts in every digital gate. All current hereby injected into the substrate causes fluctuations of the substrate potential. As a result, the threshold voltage of the transistors of the analog circuits changes as well (this is the *body effect*). This noise effect increases jitter and BER. For the most sensitive amplifiers, a differential design can suppress the common-mode substrate noise, alleviating the problem.

3.2 Measurement of Stochastic Effects

Timing parameters Our first test IC (figure 1) implements 8x8 optical chip-to-chip interconnection in both directions. The IC is developed in 0.35μm CMOS technology and operates at a link bitrate of 1.25Gbps. 24 of the 64 links are used for various testing purposes, and the other 40 links are individually resychronized to the local clock at the receiver side.

The setup allows us to measure the link latency and monitor the low-frequency jitter behaviour over time. While this measurement is relatively coarse—100ps resolution—all links of the optical interconnect array are simultaneously evaluated, thereby revealing effects that can have a gradient over the array (like misalignment or centralized heating). A circuit for measuring the BER of any channel is also present on-chip.

Substrate noise Our second test setup (figure 9) consists of an arrangement where a transmitter and a receiver chip are wirebonded to a substrate carrying 8x8 VCSEL and photodiode arrays. Both ICs have been processed in 0.18μm CMOS technology and contain 12 channels operating at 2.5Gbps.

To quantify the effect of substrate noise on the BER, the receiver IC is equipped with substrate noise generation and measurement circuits. At both short sides of the receiver circuit array, a programmable noise block is provided, accompanied with a noise measurement circuit.

Figure 10 shows a noise block in detail. It is inspired from the noise generation circuit of Nagata et. al. [10]. Each noise block contains 8 rows of 16 noise cells. The input clock of the noise block propagates over the rows with an adjustable delay and drives all cells within a row simultaneously. A noise cell consists of 30 digital inverters, each driving a capacitance to the substrate of 15fF. Every cell can be individually programmed to switch its inverters with the clock, the inverted clock, or not at all. This approach yields a highly customizable noise source.

We measure the substrate noise by wavetracing the difference between the voltage of the local metal ground and the substrate voltage. To this end, an adjustable voltage and a fixed bias voltage are coupled using NMOS capacitors to respectively the metal ground and the substrate. These voltages are the inputs of a sense amplifier circuit (figure 11), which compares both after a trigger signal. The wavetracing process requires repetitive measurements: by adjusting both the triggering moment relative to the noise generation clock and the adjustable bias voltage, the noise waveform can be reconstructed.

Two high-speed clock inputs can drive either noise block or measurement circuit, and the noise blocks can be individually powered down. The resulting system

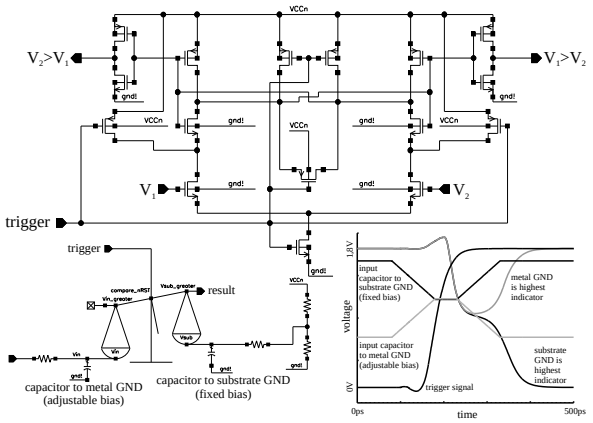

Figure 11: Substrate noise measurement circuits. On a positive trigger edge, the immediate voltage on a capacitor to metal ground (with an adjustable mean voltage) is compared to the immediate voltage on a capacitor to substrate ground (with a fixed mean voltage). Even in worst case, a voltage difference as small as 10mV during 50ps can be detected.

enables us to generate substrate noise, to measure its impact on the BER of the receiver circuits and to make a wave trace of the substrate voltage both nearby and far from the noise source.

4 Conclusion

In this paper, we have discussed a method to integrate support into EDA tools for CMOS-to-CMOS parallel optical interconnects with directly on-chip flip-chipped optics. Section 2 described the software support necessary for easy integration and multi-level simulation of this kind of interconnect. In section 3, random effects that have an impact on the performance and reliability of the parallel links have been treated. Finally, two test setups have been presented. One enables the verification of simulation models and the extraction of BER, latency, skew and low-frequency jitter figures. The other setup provides means for the generation and measurement of substrate noise and its impact on the BER of the optical interconnect.

References

[1] "Information society technologies programme contract IST-2000-28358: Interconnect by optics," http://www.intec.rug.ac.be/io/overview.html.

[2] K. Moriki, H. Nakahara, T. Hattori, and K. Iga, "Single transverse mode condition of surface-emitting injection lasers," *Electronics and Communications in Japan, Part 2*, vol. 71, no. 1, pp. 81–90, 1988.

[3] P. Mena, J. Morikuni, S. Kang, A. Harton, and K. Wyatt, "A comprehensive circuit-level model of vertical-cavity surface-emitting lasers," *IEEE Journal of Lightwave Technology*, vol. 17, no. 2, pp. 2612–2632, Dec. 1999.

[4] M. Jungo, *Series in Quantum Electronics*, 1st ed. Konstanz: Hartung-Gorre Verlag, 2003, vol. 30, ch. Spatiotemporal VCSEL Model for Advanced Simulations of Optical Links, ISBN 3-89649-831-2.

[5] M. De Wilde, O. Rits, R. Bockstaele, J. Van Campenhout, and R. Baets, "A circuit-level simulation approach to analyse system-level behaviour of VCSEL-based optical interconnects," in *Photonics Fabrication Europe*, vol. 4942. Brugge, Belgium: SPIE, Oct. 2002.

[6] M. Jungo, D. Erni, and W. Baechtold, "Quasi-analytic steady-state solution of VCSEL rate equations including spatial hole burning and carrier diffusion losses," *International Journal of Numerical Modelling*, vol. 16, no. 2, pp. 143–159, Mar. 2003.

[7] F. Mieyeville, P. Bontoux, I. O'Connor, and F. Gaffiot, "Design and optimization of optical links based on VHDL-AMS modeling," in *International Workshop on Behavioral Modeling and Simulation*, Orlando, Florida, Oct. 2000, pp. 62–72.

[8] "BSIM homepage," http://www-device.eecs.berkeley.edu/∼bsim3/.

[9] J. Gerling, O. Strübbe, and J. Schrage, "Improved time domain simulation of optical multimode intrasystem interconnects," in *Design, Automation and Test in Europe Conference*, Munich, Germany, Mar. 2003, pp. 11 110–11 111.

[10] M. Nagata, T. Ohmoto, J. Nagai, T. Morie, and A. Iwata, "Test circuits for substrate noise evaluation in CMOS digital ICs," in *Proceedings of the 2001 conference on Asia South Pacific design automation*, Yokohama, Japan, 2001, pp. 13–14.