

Design Challenges in Silicon Photonics

Wim Bogaerts, *Member, IEEE*, Martin Fiers, *Member, IEEE*, and Pieter Dumon, *Member, IEEE*

Abstract—Silicon Photonics is rapidly gaining maturity in high-bandwidth optical communication, with applications in datacom, access networks and I/O for bandwidth-intensive electronics. Also, applications are emerging in spectroscopy and sensing. To get the best performance out of the photonics, co-integration with electronics is needed: side-by-side, stacked or on the same chip. However, the combination of photonics and electronics introduces a range of new problems on the design side: Codesign and cosimulation of complex photonic and electronic circuits, tolerance to variability, and verification algorithms that can handle photonic circuits. We will discuss these challenges and give an outlook on how tools need to evolve to address the needs of photonic-electronic IC designers.

Index Terms—silicon photonics, design

I. INTRODUCTION

Silicon photonics is rapidly gaining traction as a technology platform for complex photonic integrated circuits (PIC). The technology is enabled by the infrastructure developed for silicon electronics, making it possible to define optical waveguide circuits with unparalleled precision. Moreover, the material system with its high index contrast supports submicron waveguides with sharp bends, allowing for dense integration of optical building blocks on one chip. This makes silicon photonics the first platform that can accommodate thousands of photonic components on a single chip [1], [2].

The main industrial drive for silicon photonics is in optical interconnects[3], where optical links should replace electrical links on 10cm to 2km spans[4]. Here, the photonics serves the needs of the electronics, enabling ever larger data bandwidths. But the potential of silicon photonics, and photonic integrated circuits in general, extends to many other applications in spectroscopy, sensing [5], and medical diagnostics [6]. For some of these applications, integration with electronics is needed to provide control and read-out functionality. Many strategies exist to co-integrate the photonics and the electronics: on one chip, or by co-integrating chips. But functionally the photonics and the electronics behave as a single complex circuit.

However, there is a very wide gap between what the technology can deliver, and the functionality engineers can design and simulate in complex photonic-electronic circuits: Today's design methods and tools are not up to the task. This paper aims to provide some perspectives on the upcoming challenges in photonic design automation and discusses the different aspects of the design gap in silicon photonics.

We approach this from the point of view of *electronic design automation* (EDA) [7]. Today's electronics design methods allow for a functional approach based on proven component

libraries provided by *foundries* in *process design kits* (PDK). This foundry approach is already adopted by photonics at R&D level through several *multi-project-wafer* (MPW) services [8], [9], [10], [11], which start offering photonic PDKs that enable photonic designers to follow a similar design flow. Still, the functionality of these PDKs is limited, and real photonic-electronic codesign capability is still missing.

In this paper we will discuss our view on what constitutes the design gap of silicon photonics, and offer perspectives on the capabilities that are required to close this gap in a satisfactory way, such that it becomes possible to design silicon photonic+electronic circuits for different applications.

First, in section II we will give a short background on silicon photonics technologies, which we need to analyze in section III what is needed to enable a photonic+electronic codesign flow. We will highlight several key challenges: the richness of information in a photonic signal that propagates through a waveguide (section IV), the need for some multiphysics effects to be handled at the circuit level (section V), the extreme impact of both manufacturing and operation variability on the operation of some silicon photonic circuit elements (section VI), layout and routing of photonic circuits (section VII), handling of photonic parasitics (section VIII) and finally verification algorithms that can handle the peculiarities of optical waveguide circuits (section IX).

II. SILICON PHOTONICS

Silicon photonics is a logical evolution of the photonic IC technology that emerged in the 1980s. On a PIC, light is processed and routed on a chip by means of optical waveguides. There are several material systems for PICs: doped glass, III-V semiconductors, polymers, silicon, and others, each its own advantages, and all in active use today. Silicon's key differentiator is the high refractive index contrast in the *silicon-on-insulator* (SOI) platform. This enables tight confinement of light in submicron waveguide cores, and bends with a radius of only a few micrometer [12]. This allows integration of many more functional optical building blocks on a single chip.

Unlike in *complimentary metal-oxide-semiconductor* (CMOS) electronics, where all functions are executed with a combination of transistors, resistors, diodes and capacitors, the elementary building on a PIC are quite diverse and impose different requirements. First of all, light should be guided in low-loss *waveguides*. Distribution and routing also requires efficient splitters and waveguide crossings. Coupling light on and off chip to optical fibers is also a quite challenging task, especially with submicron silicon waveguides. This is made more difficult by the fact that optical fibers support two polarizations of light, but on-chip waveguides are typically polarization sensitive.

All authors are with the Photonics Research Group, Ghent University - IMEC, Department of Information Technology, Gent, Belgium. e-mail: (see <http://photonics.intec.ugent.be>).

Manuscript received October 1, 2013 ; revised December 2, 2013.

Light covers a wide spectrum. Most PIC technologies are designed to support the telecom wavelength bands in the near infrared around 1310nm and 1550nm. For applications in *wavelength-division multiplexed* (WDM) communication, sensing or spectroscopy, efficient wavelength filters are needed, consisting of optical delay lines or resonators [13].

Essential functions involve opto-electronic conversion. Converting optical signals into electrical ones can be done with a photodiode, which in silicon photonics is usually implemented in Germanium [14]. Electrical signals can be imprinted on an optical carrier using phase or amplitude modulation; this can be done by modifying electrical carrier densities in a silicon waveguide with an embedded diode or capacitor [15]. Lasers as a light source can be integrated monolithically in III-V technology platforms, but in silicon the light source has to be either bonded [16] or coupled off-chip.

Fabrication of silicon photonics can build on the infrastructure developed for CMOS electronics. This technology is now capable of defining deep submicron features with unprecedented accuracy in large volumes. the use of silicon as a base material makes reuse of these manufacturing tools possible, but photonics imposes different demands on the processes. Layer thicknesses can be different for waveguides than for transistors, and optical quality (e.g. unwanted absorption, scattering at roughness) is now an important additional requirement. Also, silicon photonic waveguides are extremely sensitive, which imposes nanometer-scale tolerances on the fabrication process. Because of process-induced variability, wavelength-selective filter functions are difficult to fabricate exactly on spec [17], and therefore will require some form of active compensation.

For these controlling and driving functions the photonics will require integration with electronics. This can be either monolithically, incorporating optical and electrical elements in the same wafer-scale process [18], or it can be through hybrid or 3-D integration. As photonic chips become more complex, the electronic functions will become more integrated, and this imposes requirements on the design side.

III. PHOTONIC VS. ELECTRONIC DESIGN

Today's photonic IC design is a largely fragmented landscape, largely separated from the world of EDA. With the increasing requirements for electronic control and driver circuitry, and the integration of optical links in high-speed interconnects, the design process between photonic and electronic circuits needs to be integrated. This will require an efficient codesign flow, including cosimulation between the photonic and electronic domain. Given today's ratio between the photonic IC and electronic IC industries, the most likely codesign scenario is an integration of photonic capabilities into one or more EDA environments.

Figure 1 sketches the stages in a photonic design flow. It starts at the same level as an electronic flow: specifying the system requirements. From that, a circuit is composed, and that needs to be simulated. On-chip photonic circuit simulation tools are only now coming of age [19], [20], [21], [22], [23], [24], and have not been stressed to handle potential complexity




Fig. 1. Today's photonic design flow: Ideally, one starts with system specification and designs on a progressively detailed level down to the physical layout. With verification and testing, one moves again to the behavioral level. However, it is often needed to switch from physical to behavioral and back in mid-flow, or even move back to an earlier stage in the flow to perform additional component optimization.

of the thousands of mixed active and passive photonic elements supported by silicon photonics technology. This does not even involve cosimulation with electronic circuits, as we discuss further in section 2.

There had never been a real need for large-scale photonic circuit simulators until the advent of silicon photonics. Also, in the first decade of silicon photonics, research was largely focused on physical simulations to construct the photonic building blocks needed for future circuits. This posed sufficient modeling and design challenges in its own right: the high refractive index contrast often requires brute-force full-wave electromagnetic simulations. Full-vectorial brute-force electromagnetic number crunching of physical geometries is still the most applied design method for individual building blocks [25], [26], [27], [28], [29], or alternatives based on mode expansion [30], [31] and propagation [32]. These computationally expensive methods replace the more approximate methods used for low-contrast material systems. And as already mentioned, the high contrast also introduces very low tolerance to any variation in geometry. This makes optimizing devices a computationally intensive task, and difficult to extract accurate behavioral models for use in circuit simulations. Therefore, photonic circuit design still involves a lot of physical design iterations of individual building blocks.

Electronic design automation, on the other hand, has evolved to facilitate design of complex circuitry consisting of a limited set of standard building blocks. To some extent, this has gone at the expense of all-out physical optimization: individual transistors are not being optimized in the circuit design stage. Rather, the transistors are being modeled using a behavioral model based on prior physical models, theory and parameters extracted from measurements. This is more efficient and has proven to be accurate, so why is this approach not applied in photonics?

One of the challenges lies in compiling accurate behavioral models for photonic building blocks. As already discussed in the previous section, photonic circuits consist of a variety of active and passive building blocks. In addition, most of

these blocks exhibit a behavior that is wavelength-dependent, and in many cases strongly dependent on temperature and other effects. These effects should be incorporated in the component model in order to reliably simulate the circuit. We'll discuss the wavelength-dependence in more detail in the next section, and the influence of temperature in section V. A problem with compact behavioral models is often their limited validity range: they only work for a given set of parameters. If the designer moves out of the valid parameter range (either intentionally, or because of variability effects), additional physical simulations might be required. Ideally, these should be triggered automatically, without requiring the user to set up a new simulation environment. Today, only a few tools enable this kind of hierarchical modeling, and only to a limited extent [33], [34], [35], [21].

The next step in the photonics design flow is generating the mask layouts for fabrication. This process is disjoint from the circuit design (if any circuit design is actually done outside of a conceptual schematic), and usually executed in a separate tool. Support for automatic place and routing of photonic blocks is minimal. The reason for this will be discussed in section VII. Quite often, the technology imposed restrictions and limitations on the component layout, and this can have an impact on the actual performance. Especially when designing custom layouts (which is still the most common way of photonic design), some physical simulation might also be needed at the layout level to evaluate the effects of the fabrication process, or to optimize functionality [36].

Verification of the mask layout is usually done with EDA verification tools, but this is not straightforward as photonic layout elements are very unlike typical manhattan-oriented electronic features, with many curvilinear features. Also, extracting photonic circuit elements from the layout is difficult: Layout-versus-schematic verification is a significant challenge. Verification challenges are discussed in section IX.

In the photonic design flow, it is often still necessary to come back to physical simulation in a later stage in the design flow. Once aspects like variability and parasitics are incorporated into the design flow, this will often require additional physical simulations. This back-and-forth switching between the different levels in the design flow is error prone, as it often requires a redefinition of the problem in different simulation tools. While this is not a fundamental problem of photonic design, good data exchange between tools will be essential for a successful design flow [33]. An integrated flow where transitions between design steps is largely automated is therefore desirable. This is an area where EDA has paved the way, and where some early photonics solutions are appearing [34], [37], [27].

Enabling a design flow for complex PICs will also need to integrate electronic design. This is a sine qua non if photonic+electronic codesign is to succeed. Today's EDA tools support much larger complexity, and integrate more tightly the transition from circuit design to layout. To bring photonic design into an electronic design flow, EDA tools need to accept the richness of photonics, and address some of the key challenges discussed in this paper. First steps in that direction have been made [38], but these do not necessarily scale to

all photonics applications. This tighter design integration can be accomplished by native support, or by interfacing with photonic design tools. Recent developments are progressing in that direction, with schematic editing, mask layout, verification and exporting to a photonic circuit simulator driven from a standard EDA environment [39].

In the following sections we describe a number of challenges in more detail, highlighting some fundamental differences between photonics and electronics that make codesign not possible today, unless with significant approximations. We do this from the perspective of EDA tools: what needs to be supported, on top of the existing electronic design capabilities, to design and simulate photonic integrated circuits.

IV. CHALLENGE: RICH SIGNALS

Typical simulations of electronic circuits are based on SPICE or similar engines, tracking voltages and currents through the circuit. For analog circuits, the time stepping, which can be adaptive, is typically one order of magnitude smaller than fastest signals that are expected in the circuit. Depending on the application, this is MHz to tens of GHz.

Photonic circuits handle light, which is basically a very high frequency electromagnetic wave, with an oscillation frequency around 200THz. Compared to electronic phenomena, as well as the data rates of high-speed communications (25-40Gbps) this is extremely fast. Simulating the propagation of an optical wave through a circuit simulation would require too much time and resources for practical use. Even when simulating an optical link at 40Gbps, it would require 10000 samples per bit to capture the waveform accurately.

Luckily, in most cases the exact shape of the optical waveform is not required, as the high-frequencies are generally used as carrier wavelengths that are modulated at "electronic" speeds. A photonic-electronic circuit simulation can therefore run with the same time stepping. For an optical point-to-point link (laser-modulator-waveguide-photodetector) it suffices to track the intensity envelope of the optical signal, which could be represented by a single number per time step in the circuit simulation. But amplitude is a relatively poor representation of light in an optical circuit, so more information (expressed in floating point numbers) needs to be exchanged to describe the optical signal. In practice, the required amount of data depends on which application and which phenomena one wishes to capture. This is illustrated in Figure 2 and in Table I.

As light is a wave, knowledge of the phase is essential to capture electromagnetic interference, needed for interferometers or optical resonators. Mechanisms such as phase modulation can also be modeled in the system. Now two numbers need to be communicated between components in the circuit simulation. When phase is used, it is often needed to know the wavelength as well, especially with the use of optical wavelength filters. In addition, as some waveguides can support multiple *Eigenmodes*, these add another two numbers that need to be handed over. Usually, a waveguide has at least 2 guided modes, corresponding to the TE and the TM polarization. For optical links with wavelength-division multiplexing, we therefore need these four numbers for each

TABLE I
AMOUNT OF INFORMATION (EXPRESSED IN FLOATING POINT NUMBERS) NEEDED PER TIME STEP WHEN PROPAGATING DIFFERENT TYPES OF SIGNALS.

Signal	Numbers / time step	Supports modeling of	Application domains
Optical Power P	1	Amplitude modulator Photodetector Power flow in waveguide	Simple point-to-point link
Power+phase $P+\phi$	2	Phase modulation Interferometric switches	Advanced modulation Coherent detection
Power+phase+Wavelength+mode $P+\phi+\lambda+M$	4	Resonant modulators Wavelength filters	links sensors with swept readout
Multi-wavelength $P+\phi+\lambda+M$	$4 / \lambda$	direct crosstalk	WDM links
Spectral band $P+\phi+M$	$100 \cdot 5'000$	indirect crosstalk Chirp, frequency shifters Slow nonlinearities	WDM links Spectrometers
Full wave	$10'000$ [@40Gbps]	Fast nonlinearities	all-optical signal processing




Fig. 2. Information in optical signals. On the horizontal axis, the amount of information per simulation time step is presented, while on the vertical axis the richness of phenomena is illustrated. The amount of information is expressed as 'numbers', i.e. how many floating-point numbers are needed to represent the instantaneous signal at any given time.

wavelength channel. The quantity of information that needs to be passed between components per time step in an optical circuit increases quite rapidly, but up till now the number is still manageable in electronic simulators and an optical waveguide connection between two components could be implemented as an array of signal lines. This approach can scale up to 8 or 16 WDM channels.

However, for accurate simulations a sufficiently fine sampling of the optical spectrum would be needed. this would allow to capture crosstalk in WDM links with good accuracy, but also model signal degeneration due to optical dispersion. Sampling the spectrum can rapidly multiply the amount of numbers needed with 10 or 100.

How to propagate these huge amount of numbers through a large photonic circuit? In electronic circuits, signal propagation can be accelerated by *modified nodal analysis* (MNA) [40]. Circuits can be reduced using Kirchhoff's conservation laws. Unfortunately, while these effort-flow techniques can be applied in other domains (thermal, fluidics, mechanics), photonic simulation does not follow the same formalisms. MNA requires signal quantities that meet power conservation

laws (e.g. Voltage $\times$ Current corresponds to Power). There is no equivalent set of quantities to describe optical signals, but there are techniques that allow conversion from scattering-wave formalisms to MNA [41]. One possible implementation is elaborated by Gunupudi et al. in [42].

On the other hand, it is possible to reduce large photonic circuits by aggregating passive linear circuit elements. These can be described as a node with a scattering matrix (S-matrix) between the input and output ports [23], [24]. Typically, photonic circuits consist of larger aggregates of passive circuit elements (waveguides, splitters, wavelength filters, ...) that have a linear, but wavelength dependent, behavior. As long as no active or nonlinear element is present, an aggregate of passive elements can be represented by a single S-matrix, which only needs to be calculated once [43], [23].

The signals need to be passed on between components, which requires connections through optical pins or ports. Again, these concepts are different than in electronics, but more similar to planar microwave/RF design: optical interfaces of components typically consist of an open-ended waveguide to which another waveguide needs to be attached. For this attachment to work, the other waveguide should match the component's optical port in waveguide cross section and orientation, to ensure a seamless interface that does not excite reflections or unwanted scattering. Also, as already mentioned above, optical waveguides can carry multiple modes, so this modal information should also be available in the ports: the scattering matrix should be a generalized scattering matrix that covers every combination of port and mode. These concepts are used for both RF and optical waveguide circuit design, but of course there are the vast differences in frequency/wavelength, which resulted in tools that are optimized for either one or the other.

There is no straightforward solution to incorporate the richness of photonic signals into an electronic circuit simulator. Simulating electrical circuits in a photonic simulator (e.g. based on scattering matrices) is possible, but these will never meet the performance of today's SPICE simulators, and have no track record to displace the proven tools. The solution is interfacing simulators for the electronic and photonic domain.

In a first approximation this can be done using waveform exchange. A complete mixed-signal solution where separate photonic and electronic simulations are run in lockstep would be a good solution, but will be a challenge to implement.

V. CHALLENGE: MULTIPHYSICS

Photonics design is essentially multiphysics design. Photonic circuit elements interact in various ways with the light but are also affected by other quantities. Apart from the optical domain, two domains that should be handled as well at the circuit level are temperature and electronic carrier densities.

Silicon photonic waveguide elements are significantly affected by temperature. This means that outside temperature variations should be compensated by thermal tuners (i.e. heaters). This can require a local temperature control with an accuracy better than 0.1K. When applied to a wavelength-selective device this corresponds to a wavelength shift in the response of about 10pm, which can be enough to detune the device [44], [45]. In addition to outside temperature fluctuations, there can also be self-heating: a modulator based on a diode or capacitor will generate heat during operation. Lasers are especially power hungry, but also driver electronics can create hot spots that affect the photonics. Handling the thermal problem at the circuit level is not trivial, as it requires advance knowledge of the physical layout of all components, and solving the thermal distribution across the chip in sync with the circuit simulation.

Electrical carriers also affect the photonic circuits. In silicon the change of free carrier density will result in a change of both the refractive index of the silicon as well as its optical absorption [46]. This can be used intentionally (e.g. in electro-optic modulators) to modulate signals, but it can also be unintentional. Accidental doping of an optical waveguide, or incorrect annealing steps might change the local carrier density. The upside of carrier-induced effects is that they are typically locally confined, and that they can be directly controlled by the electrical circuits connected to the optical circuit. However, carriers can also be generated by the light itself. Even though silicon is quite transparent for optical wavelengths beyond $1.2\mu\text{m}$ it can still absorb two photons at the same time (two-photon absorption or TPA). This effect is weak, but scales with the intensity of the light. For high optical power, TPA will generate carriers that will change the optical properties of the circuit. In addition, the thermalization and recombination of the carriers will heat up the waveguide from the inside out, and will also induce a change in refractive index. On top of TPA, most silicon waveguides also suffer from some linear losses by defects in the lattice structure and at the etched sidewalls. Similar to TPA, this absorption induces changes in refractive index. All these secondary nonlinear effects are very difficult to include in a circuit model, especially because they have relaxation times on very different length scales: Free carriers thermalize and recombine on ns-scale, while thermal effects dissipate on a μs scale.

On top of carrier-induced and thermal nonlinearities, silicon also has an intrinsic nonlinear response to the optical field. The Kerr effect also induces a change in refractive index, and like




Fig. 3. Variability in photonic design. Both fabrication related as well as operation sources of variability affect the physical structure. This should be propagated to the circuit design level to allow an accurate prediction of the actual yield of the entire device.

TPA this scales with the optical intensity in the waveguide [47]. Unlike carrier-based or thermal nonlinearities, the Kerr effect is virtually instantaneous. While this makes it somewhat easier to handle in a circuit simulation (little or no memory effects) it can result in very different responses to optical pulses: short pulses with a high peak powers will provoke a different response than longer pulses with a lower peak power, even with the same pulse energy.

Nonlinear optical phenomena significantly complicate circuit modeling: because of the nonlinearity, new frequencies will be generated in the optical spectrum. In a WDM system with a set of fixed wavelengths, nonlinearities might cause new wavelengths to appear, which should somehow be captured by the circuit model, even if only to ascertain that they will not impact the performance of the circuit.

VI. CHALLENGE: VARIABILITY

The high refractive index contrast of silicon photonics makes it possible to confine light in a submicron waveguide core. However, the *effective* refractive index of the waveguide core is strongly governed by the actual geometry of the cross section. For a rectangular cross section, this is the width and height, but also the slope of the sidewalls.

To illustrate this more clearly, let us look at a ring resonator. This is an optical resonator where the resonance condition depends on the optical roundtrip length of the ring waveguide, and this in turn depends on the actual geometry of the waveguide. Small width or height variations will induce a significant shift in the resonance wavelength: for a width change of 1nm, the wavelength will shift approximately 1nm, which can span more than one channel in a WDM system.

Temperature can be used to tune wavelength-selective devices, so in principle, heating a ring resonator can compensate for some fabrication imperfections. However, it takes a 12K temperature increase to compensate for 1nm fabrication variation. Therefore, fabrication precision should be very accurate, but also, at the component and circuit design stage tolerances should be taken into account [44], [48].

Therefore, the effects of variability should be considered in the design stage. It is possible to optimize components

at the physical level for tolerance against geometric variations, especially if the distributions and the correlations of these variations are known through careful monitoring of the fabrication process [17]. This typically requires intensive 3-D electromagnetic simulations, and a careful analysis to translate geometric variability into functional variability. Also, at the circuit level variability can be addressed in the design phase. For instance, wavelength filters can be made athermal by adding compensation circuits [49], [50], or multiple copies of a device can be added to the circuit as a form of redundancy.

A key challenge is the efficient simulation of circuits with variability: taking into account statistics and correlation between effects, a given distribution of variable parameters must be propagated through the circuit. This could be done through Monte-Carlo simulations, but careful choice of the parameter space is needed to avoid an unmanageable number of simulations. Getting closure for the variability in photonic circuits covers much more effects than the exploration of slow and fast corners used in electronic design.

Ultimately, the outcome of a variability analysis is a prediction of the circuit yield. The propagation of variability from nominal devices to circuit yield is illustrated in Fig. 3.

VII. CHALLENGE: MASK LAYOUT

Photonic layouts are quite different from electronic layouts. As the propagation of electromagnetic waves is entirely governed by the geometry, this needs to be accurately controlled. While electronic building blocks usually consist of rectangles (without considering assist features or optical proximity corrections), photonic device layouts consist of curvilinear shapes to define smooth waveguide bends.

Silicon photonic waveguides are typically fabricated in a single device layer of crystalline or amorphous silicon embedded in oxide. However, different patterning layers can be defined with different etch depths [13], and sometimes additional material is added to the device layer [51]. Given the high sensitivity to geometric variations, very accurate overlay between these different patterns is required. Also, parasitic effects of patterning must be corrected with suitable proximity corrections [52]. Optimizing the lithography conditions for a specific type of pattern is not really possible. Unlike critical layers in electronic designs, photonic waveguide layers often contain a mix of isolated and dense features of different dimensions and shapes. Also, local and global density plays an important role in pattern fidelity [53]. Today's silicon photonic circuits all rely on tiling to guarantee a uniform pattern density, but some photonic components rely on larger unetched areas (e.g. arrayed waveguide gratings and echelle gratings [13]).

Another difficulty is placement and routing (P&R). Photonic waveguides need to observe a minimum bend radius to avoid high propagation losses. Also, sufficient spacing is needed between routing waveguides to avoid evanescent coupling. At the interfaces with the component ports, the routing waveguides also need to match the port's waveguide geometry, as well as angle, to avoid backreflections and scattering. Given that optical waves have a phase, in some cases it is also needed that the length of the interconnections can be set exactly, to control




Fig. 4. Different parasitic light paths can severely affect the performance of a photonic integrated circuit.

interferences. These are not straightforward, but probably be enabled by existing custom electronic P&R tools.

However, probably the largest challenge is that the photonic circuit is typically only defined in a single layer. But routing a complex circuit on a single layer is not straightforward, and often impossible without crossing interconnections. Fortunately, photonic waveguides can be engineered to cross without excessive loss and crosstalk [54], but routing algorithms need to take this into account and minimize the number of crossings. While it is possible to make multi-layer circuits, efficient optical vias are difficult to implement in a high-contrast material system, this comes with serious performance and integration penalties.

VIII. CHALLENGE: PARASITICS

Today's waveguide circuits are usually designed by composing them of ideal building blocks. However, just as in electronics, there are parasitic mechanisms that can disturb the operation of the circuit. Some of these parasitics could already be simulated at the design stage. Most of the unwanted effects include light that does not end up in the place where it is supposed to be. This parasitic light can interfere constructively or destructively, complicating the calculation of the actual net result of the parasitics.

Some examples of photonic parasitic effects are shown in Fig. 4. They include (back)reflections at discontinuities [55], [56], distributed scattering at roughness, unintended coupling to adjacent components and excessive absorption and stray light propagating in cladding modes or the substrate. In addition, optical nonlinearities can be considered as parasitics, as devices models might not take into account local high intensities within a building block.

IX. CHALLENGE: VERIFICATION

At the end of the design flow the layout should be verified before it is sent for manufacturing. Like in electronics, the first line of verifications is a design rule check of the mask layers, to verify minimum critical dimensions, overlays and density checks. Such rules are already in use in today's photonic design flows, making use of the tools available for electronic DRC [57], although some photonic design tools support DRC specifically for photonics [58]. However, photonics might have need for additional DRC rules that are not as easily




Fig. 5. A number of illustrations on how waveguides can be connected or not. (a) Two properly connected waveguides; (b) physical contact but no angular matching will cause loss and reflections; (c) two adjacent waveguides might be evanescently coupled (intentionally or unintentionally) without physical contact and the coupling can depend on the wavelength; (d) Two waveguides can cross, but have little or no coupling.

implemented. As photonic layouts don't follow manhattan orientations, guaranteeing linewidth uniformity along curvilinear paths needs to be enabled. As discussed in section VI, linewidths need to be controlled down to nanometer level.

However, the larger verification challenge is *layout-versus-schematic*. Extracting the functional behavior from a photonic layout is difficult without resorting to full electromagnetic simulations. And even simply checking connectivity might be hard, as illustrated in Fig. 5. Waveguides are only properly connected when there is physical contact, a matching of the cross section and the angle. When there is no proper angular alignment, reflections and scattering will occur. But on the other hand, two waveguides do not need to be connected to have evanescent coupling. Oppositely, crossing waveguides might be engineered to have no coupling.

This analysis is made even more complicated when taking into account wavelength. Wavelength selective components might couple one wavelength efficiently while blocking another, but how to extract this functionality from the layout? When wavelength channels in a WDM circuit are separate signal lines, it is difficult to discriminate those in the layout.

X. CONCLUSIONS

In this paper, we picture an overview of the design challenges faced by silicon photonics, from a viewpoint of enabling photonic-electronic codesign. This technology finally enables large-scale integration of photonic components on a chip, but this implies silicon photonics designers need to address unprecedented circuit complexity, while at the same time cope with extremely small tolerances and high sensitivity. EDA tools, with their installed base and proven workflow for electronic design, provide an ideal platform to accommodate the designer. But the differences between photonics and electronics will require that customized solutions for photonics are developed and integrated into existing workflows. Mixed-signal simulation of photonic circuits and electronic circuits (first analog, but later even digital) will require the integration of photonic-capable circuit simulators with existing electrical simulators. Similarly, interfaces to physical electromagnetic solvers will be needed, as photonic design cannot always be captured in an abstract model. EDA tools should also facilitate

curvilinear design and verification tools need to be extended to support photonic concepts for functional verification.

REFERENCES

- [1] R. Soref, "The past, present, and future of silicon photonics," *J. Sel. Top. Quantum Electron.*, vol. 12, no. 6, pp. 1678–1687, 2006.
- [2] T. Baehr-Jones, T. Pinguet, P. Lo Guo-Qiang, S. Danziger, D. Prather, and M. Hochberg, "Myths and rumours of silicon photonics," *Nature Photonics*, vol. 6, no. 4, pp. 206–208, Apr. 2012.
- [3] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, 2009.
- [4] J. Doyle and A. Knights, "The evolution of silicon photonics as an enabling technology for optical interconnection," *Laser Photon. Rev.*, vol. 6, no. 4, pp. 504–525, Jul. 2012.
- [5] X. Fan, I. White, S. Shopova, H. Zhu, J. Suter, and Y. Sun, "Sensitive optical biosensors for unlabeled targets: A review," *Analytica Chimica Acta*, vol. 620, no. 1–2, pp. 8–26, Jul. 14 2008.
- [6] Y. Li, S. Verstuyft, G. Yurtsever, S. Keyvaninia, G. Roelkens, D. Van Thourhout, and R. Baets, "Heterodyne laser doppler vibrometers integrated on silicon-on-insulator based on serrodyne thermo-optic frequency shifters," *Appl. Opt.*, vol. 52, no. 10, pp. 2145–2152, 2013.
- [7] L.-T. Wang, Y.-W. Chang, and K.-T. T. Cheng, *Electronic design automation: synthesis, verification, and test*. Morgan Kaufmann, 2009.
- [8] P. Munoz, J. D. Domenech, I. Artundo, J. H. den Bested, and J. Capmany, "Evolution of fabless generic photonic integration," in *Transparent Optical Networks (ICTON), 2013 15th International Conference on*. IEEE, 2013, pp. 1–3.
- [9] ePIXfab, "The silicon photonics platform," <http://www.epixfab.eu/>.
- [10] OpSIS, "Optoelectronic systems integration in silicon," <http://opsisfoundry.org/>.
- [11] JEPPIX, "Joint european platform for inp-based photonic integrated components and circuits," <http://www.jepix.eu/>.
- [12] W. Bogaerts and S. K. Selvaraja, "Compact Single-Mode Silicon Hybrid Rib/Strip Waveguide With Adiabatic Bends," *IEEE Photon. J.*, vol. 3, no. 3, pp. 422–432, Jun. 2011.
- [13] W. Bogaerts, S. Selvaraja, P. Dumon, J. Brouckaert, K. De Vos, D. Van Thourhout, and R. Baets, "Silicon-on-insulator spectral filters fabricated with cmos technology," *J. Sel. Top. Quantum Electron.*, vol. 16, no. 1, pp. 33–44, 2010.
- [14] L. Vivien, J. Osmond, J.-M. Fédéli, D. Marris-Morini, P. Crozat, J.-F. Damlencourt, E. Cassan, Y. Lecunff, and S. Laval, "42 ghz p.i.n germanium photodetector integrated in a silicon-on-insulator waveguide," *Opt. Express*, vol. 17, no. 8, pp. 6252–6257, Apr. 2009.
- [15] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon optical modulators," *Nature Photonics*, vol. 4, no. 8, pp. 518–526, Aug. 2010.
- [16] G. Roelkens, L. Liu, D. Liang, R. Jones, A. Fang, B. Koch, and J. Bowers, "III-V/silicon photonics for on-chip and inter-chip optical interconnects," *Laser Photon. Rev.*, vol. 4, no. 6, pp. 751–779, Nov. 2010.
- [17] S. K. Selvaraja, W. Bogaerts, P. Dumon, D. Van Thourhout, and R. Baets, "Subnanometer Linewidth Uniformity in Silicon Nanophotonic Waveguide Devices Using CMOS Fabrication Technology," *J. Sel. Top. Quantum Electron.*, vol. 16, no. 1, pp. 316–324, Jan.-Feb. 2010.
- [18] A. Mekis, S. Gloeckner, G. Masini, A. Narasimha, T. Pinguet, S. Sahni, and P. De Dobbelaere, "A Grating-Coupler-Enabled CMOS Photonics Platform," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 597–608, May-Jun. 2011.
- [19] Ghent University - imec, "Caphe circuit simulator," <http://www.caphesim.eu>.
- [20] Photon Design, "PICWave, a photonic integrated circuit (pic) design and simulation tool," <http://www.photond.com/products/picwave.htm>.
- [21] Lumerical Solutions, Inc., "Interconnect," <http://www.lumerical.com/tcad-products/interconnect/>.
- [22] Filarete, "ASPIC photonic circuit simulator," <http://www.aspicdesign.com/>.
- [23] M. Fiers, T. Van Vaerenbergh, K. Caluwaerts, D. Vande Ginste, B. Schrauwen, J. Dambre, and P. Bienstman, "Time-domain and frequency-domain modeling of nonlinear optical components at the circuit-level using a node-based approach," *JOSA B*, vol. 29, no. 5, pp. 896–900, 2012.
- [24] D. Melati, F. Morichetti, A. Canciamilla, D. Roncelli, F. Soares, A. Bakker, and A. Melloni, "Validation of the building-block-based approach for the design of photonic integrated circuits," *J. Lightwave Technol.*, vol. 30, no. 23, pp. 3610–3616, 2012.

- [25] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Johnson, "MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method," *Comp. Phys. Comm.*, vol. 181, pp. 687–702, January 2010.
- [26] Lumerical Solutions, Inc., "FDTD solutions," <http://www.lumerical.com/tcad-products/fdtd/>.
- [27] Phoenix BV, "Photonic design automation," <http://www.phoenixbv.com/>.
- [28] RSoft, "Photonics and network design software," <http://www.rsoftdesign.com/>.
- [29] Optiwave, "Design software for photonics," <http://www.optiwave.com>.
- [30] Photon Design, "FIMMProp-3D, a bidirectional optical propagation tool," <http://www.photond.com/products/fimmprop.htm>.
- [31] P. Bienstman and R. Baets, "Optical modelling of photonic crystals and vcsels using eigenmode expansion and perfectly matched layers," *Opt. Quantum Electron.*, vol. 33, no. 4/5, p. 327, 2001.
- [32] R. Scarmozzino, A. Gopinath, R. Pregla, and S. Helfert, "Numerical techniques for modeling guided-wave photonic devices," *J. Sel. Top. Quantum Electron.*, vol. 6, no. 1, pp. 150–162, 2000.
- [33] W. Bogaerts, Y. Li, S. Pathak, A. Ruocco, M. Fiers, A. Ribeiro, E. Lambert, and P. Dumon, "Integrated design for integrated photonics: from the physical to the circuit level and back," *Proc. SPIE*, vol. 8781, pp. 878 102–878 102–11, 2013.
- [34] M. Fiers, E. Lambert, S. Pathak, P. Dumon, B. Maes, P. Bienstman, and W. Bogaerts, "Improving the design cycle for nanophotonic components," *J. Comp. Sci.*, vol. 4, no. 5, pp. 313 – 324, 2013.
- [35] C. Arellano, S. Mingaleev, I. Koltchanov, A. Richter, J. Pomplun, S. Burger, and F. Schmidt, "Efficient design of photonic integrated circuits (pics) by combining device-and circuit-level simulation tools," in *SPIE OPTO*. International Society for Optics and Photonics, 2013, pp. 862 711–862 711.
- [36] W. Bogaerts, P. Bradt, L. Vanholme, P. Bienstman, and R. Baets, "Closed-loop modeling of silicon nanophotonics from design to fabrication and back again," *Opt. Quantum Electron.*, vol. 40, no. 11-12, pp. 801–811, Sep. 2008.
- [37] Ghent University - imec, "IPKISS parametric design framework," <http://www.ipkiss.org>.
- [38] A. Mekis, S. Abdalla, P. De Dobbelaere, D. Foltz, S. Gloeckner, S. Hovey, S. Jackson, Y. Liang, M. Mack, G. Masini, R. Novais, M. Peterson, T. Pinguet, S. Sahni, J. Schramm, M. Sharp, D. Song, B. Welch, K. Yokoyama, and S. Yu, "Scaling CMOS photonics transceivers beyond 100 Gb/s," *Proc. SPIE*, vol. 8265, 2012, Conference on Optoelectronic Integrated Circuits XIV, San Francisco, CA, JAN 25-26, 2012.
- [39] L. Chrostowski and M. Hochberg, "Silicon photonics design," *Self-published via Lulu.com*, 2013.
- [40] C.-W. Ho, A. Ruehli, and P. Brennan, "The modified nodal approach to network analysis," *IEEE Trans. Circ. Sys.*, vol. 22, no. 6, pp. 504–509, 1975.
- [41] R. Longoria, "Wave-scattering formalisms for multiport energetic systems," *J. Franklin Inst.*, vol. 333, no. 4, pp. 539–564, 1996.
- [42] P. Gunupudi, T. Smy, J. Klein, and Z. Jakubczyk, "Self-consistent simulation of opto-electronic circuits using a modified nodal analysis formulation," *IEEE Trans. Adv. Pack.*, vol. 33, no. 4, pp. 979–993, 2010.
- [43] C. Arellano, S. Mingaleev, E. Sokolov, I. Koltchanov, and A. Richter, "Time-and-frequency-domain modeling (tfdm) of hybrid photonic integrated circuits," in *SPIE OPTO*. International Society for Optics and Photonics, 2012, pp. 82 650K–82 650K.
- [44] P. De Heyn, J. De Coster, P. Verheyen, G. Lepage, M. Pantouvaki, P. Absil, W. Bogaerts, J. Van Campenhout, and D. Van Thourhout, "Fabrication-tolerant four-channel wavelength-division-multiplexing filter based on collectively tuned si microrings," *J. Lightwave Technol.*, vol. 31, no. 16, pp. 3085–3092, Aug 2013.
- [45] D. Dai, L. Yang, and S. He, "Ultrasmall thermally tunable microring resonator with a submicrometer heater on Si nanowires," *J. Lightwave Technol.*, vol. 26, no. 5-8, pp. 704–709, March 2008.
- [46] R. Soref and B. Bennett, "Electrooptical effects in silicon," *J. Quantum Electron.*, vol. 23, no. 1, pp. 123 – 129, 1987.
- [47] M. Dinu, F. Quochi, and H. Garcia, "Third-order nonlinearities in silicon at telecom wavelengths," *Appl. Phys. Lett.*, vol. 82, no. 18, pp. 2954–2956, May 5 2003.
- [48] A. Krishnamoorthy, X. Zheng, G. Li, J. Yao, T. Pinguet, A. Mekis, H. Thacker, I. Shubin, Y. Luo, K. Raj, and J. Cunningham, "Exploiting CMOS Manufacturing to Reduce Tuning Requirements for Resonant Optical Devices," *IEEE Photon. J.*, vol. 3, no. 3, pp. 567–579, Jun. 2011.
- [49] B. Guha, B. B. C. Kyotoku, and M. Lipson, "CMOS-compatible athermal silicon microring resonators," *Opt. Express*, vol. 18, no. 4, pp. 3487–3493, Feb. 15 2010.
- [50] S. Dwivedi, H. D'heer, and W. Bogaerts, "A Compact All-Silicon Temperature Insensitive Filter for WDM and Bio-Sensing Applications," *Photon. Technol. Lett.*, vol. 25, no. 22, pp. 2167–2170, 2013.
- [51] D. Vermeulen, S. Selvaraja, P. Verheyen, G. Lepage, W. Bogaerts, P. Absil, D. Van Thourhout, and G. Roelkens, "High-efficiency fiber-to-chip grating couplers realized using an advanced cmos-compatible silicon-on-insulator platform," *Opt. Express*, vol. 18, no. 17, pp. 18 278–18 283, 2010.
- [52] S. Selvaraja, P. Jaenen, W. Bogaerts, D. Van Thourhout, P. Dumon, and R. Baets, "Fabrication of photonic wire and crystal circuits in silicon-on-insulator using 193nm optical lithography," *J. Lightwave Technol.*, vol. 27, no. 18, pp. 4076–4083, 2009.
- [53] S. K. Selvaraja, K. De Vos, W. Bogaerts, P. Bienstman, D. Van Thourhout, and R. Baets, "Effect of device density on the uniformity of silicon nano-photonic waveguide devices," in *2009 IEEE Leos Annual Meeting Conference Proceedings, Vols 1 and 2*, ser. IEEE Lasers and Electro-Optics Society (LEOS) Annual Meeting. New York: Ieee, 2009, pp. 311–312.
- [54] W. Bogaerts, P. Dumon, D. Van Thourhout, and R. Baets, "Low-loss, low-crosstalk crossings for SOI nanophotonic waveguides," *Opt. Lett.*, vol. 32, no. 19, pp. 2801–2803, October 2007.
- [55] A. Canciamilla, M. Torreggiani, C. Ferrari, F. Morichetti, R. Costa, and A. Melloni, "Backscatter in integrated optical waveguides and circuits," *Proc. SPIE*, vol. 7218, p. 72180N, 2009.
- [56] E. Kleijn, M. Smit, and X. Leijtens, "Analysis of parasitic effects in PICs using circuit simulation," *Proc. SPIE*, vol. 8781, pp. 878 104–878 104–9, 2013.
- [57] Mentor Graphics, "Calibre - IC verification," <http://www.mentor.com>.
- [58] Design Workshop, "Dw2000 - photonic physical layout and verification software," <http://www.designw.com/PDA-Overview.php>.



Wim Bogaerts completed his studies in engineering (applied physics) at Ghent University, Belgium, in 1998 and joined the Department of Information Technology (INTEC) at both Ghent University and the Interuniversity Microelectronics Center (IMEC) where he received the Ph.D. degree in 2004. In the Photonics Research Group he specialized in the modeling, design and fabrication of nanophotonic components. He is currently a Professor in the field of silicon photonics at Ghent University, and coordinates the activities between imec and Ghent University in silicon photonics design. Wim Bogaerts is a member of the IEEE Photonics Society, the Optical Society of America (OSA) and SPIE.



Martin Fiers completed his studies in electrical engineering at Ghent University, Belgium in 2008 and joined the department of information technology (INTEC) at the same university for a Ph.D. in Photonic Reservoir Computing. Martin received his Ph.D. degree in 2013. He is now developing software for designing and modeling nanophotonic circuits. His main interests are the modeling of nanophotonic components and machine learning.



Pieter Dumon received the degree in electrical engineering from Ghent University, Gent, Belgium, in 2002, where he received a Ph.D. degree in electrical engineering at Ghent University in 2007, for his work in wavelength filters in silicon photonic wires. He currently coordinates ePIXfab, a initiative for multi-project wafer fabrication in photonics. He is a member of IEEE.