

Layout-Aware Variability Analysis, Yield Prediction, and Optimization in Photonic Integrated Circuits

Wim Bogaerts , Senior Member, IEEE, Yufei Xing, Student Member, IEEE, and Umar Khan , Member, IEEE

Abstract—We present a simulation framework for evaluating the effect of location-dependent variability in photonic integrated circuits. The framework combines a fast circuit simulator with circuit layout information and wafer maps of waveguide width and layer thickness variations to estimate the statistics of the circuit performance through Monte Carlo simulations. We illustrate this with ring resonator filters, a design sweep of Mach–Zehnder lattice filters, and the tolerance optimization of a Mach–Zehnder interferometer, and show how variability aware design can be essential for future photonic circuit design, especially in a fabless ecosystem where details of the foundry processes are not available to the designers.

Index Terms—Integrated optics, design automation, circuit simulation, yield estimation.

I. INTRODUCTION

SILICON Photonics has rapidly gained adoption as one of the most promising technologies for photonic integrated circuits (PIC) [1]. It is especially attractive because it combines the use of industrial CMOS manufacturing technology with the potential of large-scale integration [2]. Like with CMOS electronics, the high cost of the fabrication facilities has spurred an ecosystem of foundries and fabless users, where circuit designers become disconnected from the fabrication process [3], [4].

At the same time, circuit complexity is rapidly growing. Especially in silicon photonics, this is made possible by the extremely high refractive index contrast, allowing submicrometer-scale waveguides with tight bends [5]. However, this high index contrast introduces its own set of challenges. In particular, it makes the submicron waveguide components extremely sensitive to very small variations in the geometry. Even nanometer-scale deviations in waveguide width or layer thickness can induce non-negligible changes in the optical behavior [6]. This sensitivity introduces significant challenges for the design of photonic circuits [4]: as circuits increase in size, the effects of variability

accumulate, and the overall circuit yield will rapidly decrease as individual devices do not meet specifications, or the properties of devices within a circuit are not properly matched. This is especially true in interferometric wavelength filters based on delay lines, such as Mach-Zehnder interferometers (MZI), ring resonators and arrayed waveguide gratings (AWG) [7]. In wavelength filters with submicrometer silicon waveguides, a width variation of 1 nm can give rise to a 1 nm wavelength shift of the filter response. For thickness variations, this effect is even twice as strong, close to 2 nm wavelength shift per 1 nm thickness change. Taking into account that even in a good silicon photonics process the waveguide width and thickness can vary with several nanometers over a wafer, fabrication variability can induce non-negligible performance variations. For instance, in dense wavelength division multiplexing (WDM) telecommunication systems wavelength channels can be spaced as closely as 0.2 nm (or 25 GHz in the C-band); wavelength shifts of a few nanometer are really unacceptable. By integrating active tuning elements, such as heaters, some of the imperfections introduced by variability can be corrected. But this active compensation introduces more complexity, and requires additional footprint and power consumption. To compensate for a 1 nm width change, a temperature increase of approximately 10 °C is required.

Therefore, it is always preferable to minimize the effects of variability at the design stage. This is especially needed in an ecosystem based on foundries, where the fabless users typically have only a limited insight into the process specifics of the fabrication, but still want to design circuits that meet specifications. As a result, in the past few years variability modelling of photonic circuits has been recognized as a significant need. It is important that this happens at the circuit level, using efficient behavioral models, rather than full-vectorial electromagnetic simulations which are too computationally intensive. With such efficient models, the most straightforward method is to apply random variations on the model parameters and perform Monte-Carlo simulations to assess the impact of the variations [8]. When comparing these simulation results with the original performance specifications for the circuits, the overall yield can be predicted. More advanced stochastic methods, such as polynomial chaos expansion [9]–[11] and stochastic collocation [12] can replace the brute-force Monte-Carlo simulations with more efficient simulations that incorporate the stochastic properties of the circuit elements.

However, purely random circuit variations only give a qualitative estimate of the yield: by just looking at the circuit parameters the impact of the actual layout is ignored. It can be

Manuscript received January 9, 2019; revised February 26, 2019; accepted March 15, 2019. Date of publication March 27, 2019; date of current version May 1, 2019. This work was supported in part by the Flemish Research Council (FWO-Vlaanderen) under Grant G013815N, and in part by the Flemish Fund for Innovation and Development (VLAIO) and Luceda Photonics through the MEPIC project. (Corresponding author: Wim Bogaerts.)

The authors are with the Photonics Research Group, Department of Information Technology, Ghent University—IMEC, Ghent 9052, Belgium, and also with the Center for Nano and Biophotonics (NB-Photonics), Ghent University, Ghent 9000, Belgium (e-mail: wim.bogaerts@ugent.be; yufei.xing@ugent.be; umar.khan@ugent.be).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSTQE.2019.2906271




Fig. 1. Variability in the fabrication process can propagate all the way to the performance at system level. During the processing of silicon waveguides typical parameters that are monitored are the width and layer thickness. This has a direct effect on the propagation constant of the waveguides, or the coupling coefficients of directional couplers. In turn these will influence the circuit behavior and ultimately the yield of the overall circuit.

intuitively understood that circuit elements close together are more likely to have similar properties than circuit elements far apart on the chip. Therefore, layout-aware circuit modeling and variability analysis can give a more accurate estimate of the circuit performance [13]. To realize this, three essential elements are required:

- 1) Compact models of all building blocks in the circuits, and the sensitivity of the model parameters to the global variables, such as waveguide width and thickness changes.
- 2) A representative model of how these global variables vary over a wafer, and from wafer to wafer. Variations of width and thickness usually have different contributions, and such a wafer map therefore has multiple components with different length scales.
- 3) The layout of the circuit, with the positions (and orientation) of the individual building blocks.

We have implemented a simulation engine that combines these three elements for layout-aware variability analysis, by adding *variability extensions* to the commercial circuit simulator *Caphe* by *Luceda Photonics* [14]. In this paper, we will discuss the basic operation of the *Caphe* simulator with variability extensions, and we will illustrate it through three examples of wavelength filters, showing how variability modelling can guide the design process and make circuits more robust against process variations. In Section II we discuss the origin of variability in photonic circuits. Then, the operation of our simulation engine is discussed in more detail in Section III. We illustrate the functionality with three examples in Section IV: We design a Mach-Zehnder lattice filter, analyze its performance under conditions of waveguide width and thickness variability, and then optimize the number of stages to obtain the highest yield. A second example illustrates the problem of device matching, by modelling a 4-channel wavelength demultiplexer based on ring resonators. In a third example, we apply our technique to a design of an MZI which is already tolerant to long-range width variations, and we illustrate how our tool can differentiate between long-range and short-range variations. Finally, in Section V, we discuss the strengths and weaknesses of the technique in more detail and offer some perspectives for future developments.

II. MODELLING OF VARIABILITY IN PHOTONIC CIRCUITS

Accurately capturing the effects of variability in photonic circuits has been studied for well over a decade. the problem is not

unique to photonics; in electronics, the effects of device variability on the performance of circuits has been known for quite some time, and support for variability modelling has been built into electronic design automation (EDA) tools. Still the most commonly used technique is *corner analysis*, where circuits are tested with models for *worst-case* and *best-case* scenarios (corners). Usually, this translates into *slow* and *fast* transistors. While this has proven its use in electronics for many decades, it is gradually becoming inadequate, as for more advanced electronic nodes better understanding of the statistics is needed, beyond *best-case* and *worst-case*. For photonics, corner analysis is also not very applicable. While it might be possible to define best-case and worst-case behaviour for many photonic devices (e.g. responsivity of a photodetector, or propagation losses in waveguides), many properties are not inherently good or bad, such as the effective index of a waveguide. Instead, it is important to know the distribution of the properties, and how well some of these properties vary from device to device.

A. Variability Contributions

To effectively model variability in photonic circuits, one has to go back to the origins of the variability. This is illustrated in Fig. 1. Most variations in a photonic circuit originate in the fabrication process. Patterns on a photomask (which itself also carries some variability) are illuminated using photolithography and transferred into the substrate material through plasma etching. Layers are deposited using a multitude of techniques, and planarized using chemical/mechanical polishing (CMP). All these steps are not 100% reproducible from wafer to wafer and lot to lot, and not always uniform over an entire wafer.

These process variations result into small changes in the geometries and material distributions in the photonic circuit, such as layer thicknesses, doping profiles, sidewall slope, etc. In silicon waveguides, the most influential geometric properties are the width and the thickness of the guided silicon layer. The first is largely determined by the lithography and etching, while the second originates often from the silicon-on-insulator (SOI) wafer bonding. Therefore, the distribution over the wafer of these two properties can be very different. Changes in geometry are usually monitored during fabrication using process control monitor (PCM) structures.

The geometry variations translate directly into a variation in optical properties. For instance, waveguide width and thickness

variations of a silicon waveguide can be mapped directly onto the effective index n_{eff} and group index n_g of the guided mode [15]. For other devices, such as directional couplers, a similar mapping can be done to the coupling coefficients.

Variation in the device parameters propagates into the circuits where the devices are being used. Here, the variations accumulate, and because circuits can connect devices that are spaced quite apart on the chip, it is not trivial to assess the impact, and how the system performance will be affected. For instance, variations in the effective index of waveguides will result in a shift of the response of wavelength filters, but when multiple delay lines in the filter experience different variations, the mismatched phases can cause ripple on the transmission spectrum, and increase the crosstalk. Similarly, filters based on directional couplers will have increased crosstalk and insertion loss if the couplers deviate from their designed specifications.

B. Modelling Variability

The classical approach to model variability is the Monte-Carlo process, where a large number of simulations are run with stochastic variations in the parameters. The advantage of Monte-Carlo simulations is the simplicity of implementation and the support for large numbers of independent variables. However, the Monte-Carlo method is computationally expensive, just because of the large number of simulations.

Therefore, over the past decade, novel techniques have emerged. For instance, in stochastic collocation, the simulation models are replaced with a cheaper surrogate model that captures the stochastic properties of the original model in a minimum of expensive simulations, after which a Monte-Carlo simulation can be run on the surrogate model [12]. This can be applied to component simulation, but also to circuits.

In polynomial chaos expansion (PCE), the component or circuit model is extended to incorporate not just the original model parameters, but also the moments of their distribution [16], [17]. A new, larger model is created which can solve the circuit in a single simulation, including the stochastic moments of the performance metrics. While extremely powerful, this technique requires good a-priori knowledge of the probability density functions (PDF) of the model parameters. Also, it assumes pure stochastic variables, which makes it harder to incorporate knowledge of the circuit layout into the technique.

C. Location-Aware Variability Modelling

Because the fundamental sources of variability in photonic circuits are present on different length scales, it is important to incorporate the actual layout information into the variability analysis. An effective technique to photonic layouts with variability modelling was first introduced in [13]. It starts from a wafer map of geometry variations (width and thickness) and the layout of the photonic circuit. The photonic circuit is projected onto the wafer map, and the local width and thickness are extracted for each circuit component. The circuit models have waveguide width and thickness as parameters, and can be evaluated based on these local values. This is then repeated for different positions on the wafer, and for different wafer maps, to obtain a valid Monte-Carlo simulation.

This approach requires that the circuit models support these global geometry variables, or at least can evaluate their actual parameters (e.g. effective index, loss) from these global geometry parameters. In its simplest form, the perturbation of a circuit model parameter C by a global parameter X at position p could be implemented as a linear perturbation:

$$C(p) = C_0 + \frac{\partial C}{\partial X} \Delta X(p) \quad (1)$$

For this, the sensitivity of C to X should be known. This can be characterized through measurements or simulations.

Such layout-aware variability modelling requires that the layout information is coupled to the corresponding circuit models. This can be done either by annotating the layout files (GDSII) with the model and its parameters [13], or by using a cell-driven design tool where the layout and models are embedded (parametric) cells that are generated by code or stored in a database [19], [20].

III. THE CAPHE VARIABILITY EXTENSIONS (CAPHEVE)

We have implemented such a simulation scheme on top of the IPKISS design framework by Luceda Photonics [20]. This design tool combines layout, connectivity and circuit model into parametric cells. It also has a built-in photonic circuit simulator, *Caphe*, that supports both frequency domain and time-domain simulations, with efficient circuit models that can be custom-written in Python [14]. The circuit design flow for a photonic circuit is depicted in the top part of Fig. 2. Starting from a component library in a process design kit (PDK), a circuit is composed of parametric building blocks, and a mask layout is generated [4], [20]. The resulting circuit is then simulated and optimized until it meets the specifications.

To extend this design flow with variability analysis and yield prediction, we did not want to impose any restrictions on how the fab has constructed its circuit models, and how designers generate their layout. The IPKISS framework, which is written in Python, can be easily extended with additional functionality [20], [21], so we created the necessary data structures and processes without perturbing the original framework and without requiring the fab to change their circuit models.

To describe the sensitivity, we annotated the existing circuit models with a variability matrix, describing how every circuit model parameter varies with changes in local width and thickness. These annotations can be of the linear form in Eq. (1), higher-order polynomial expressions, or a custom Python function. The sensitivity data structure is added to the existing models. The actual sensitivity data is not generated automatically; if this is not supplied by the fab, it is up to the designer to evaluate this by running simulations or experimentally characterize fabricated devices. By default, the sensitivity of component parameters is set to zero.

A second data structure which is injected is a set of sampling points for each component. This determines where in the layout the global variables such as width and thickness are evaluated. For most (small) building blocks, a single sampling point is sufficient, but longer waveguides are automatically sampled with a regular spacing (which is parametric).




Fig. 2. Adding layout-aware variability modelling to the photonic circuit design flow. Above the dotted line, the classical photonic circuit design flow is depicted [4], starting from PDK blocks and composing circuits first as a schematic and then as a layout. Only the nominal circuit response is simulated. We extend the PDK models with sensitivity data (either from measurement or simulation) and generate wafer maps of global variables (waveguide width, thickness) [13], [18]. We can then perform Monte-Carlo simulations by placing the circuit on different wafer positions (and different virtual wafers). From the many circuit responses we can then extract the yield, i.e. the fraction of circuits that meet the specifications set out in the system requirements.

Finally, representative wafer maps of the global variables are needed. Because waveguide width and thickness variations can have multiple causes, wafer maps can have variations over different length scales. A hierarchical model captures wafer-to-wafer, die-to-die and intra-die variability of both systematic and stochastic nature [18], [22].

The CapheVE framework combines these three elements with the existing circuit models and layout. First it positions the circuit on the wafer, then evaluates the local width and thickness for each sample point, and for building blocks with multiple sample points these values are aggregated. Using the sensitivity matrix, the updated model parameters for each instance are calculated and a circuit simulation is launched. This is repeated for multiple positions on the wafer. In this process, the original circuit design is not altered.

Based on the results, plotting and data analysis routines from scientific Python libraries can be used to evaluate the impact of the variability or predict the yield of the circuit after fabrication. Because the whole process is scriptable from Python, this simulation routine can easily be embedded in an optimization loop to optimize a circuit for yield, rather than for ultimate performance.

IV. EXAMPLES

To illustrate the use of this tool, we performed variability analysis on a number of interferometric wavelength filter circuits, because such circuits are extremely sensitive to effects that perturb the relative phase between different light paths.

We designed these filters using the design kit of the iSiPP50G silicon photonics technology platform of IMEC. This platform features a 220 nm thick silicon waveguide layer on a 2 μm buried oxide layer. The 200 mm wafers are sourced from a commercial supplier and can have a variation of a few nanometer on the

silicon waveguide layer. Due to the wafer fabrication processes, this variation changes slowly over the wafer, often with a radial pattern. The waveguides are defined with a 193 nm projection lithography system and transferred into the silicon layer with a plasma etching process. Unless otherwise specified, we use a nominal waveguide width of 450 nm. Of course, the lithography step can induce small variations in the width, and the plasma etching can also affect the width due to variations in plasma composition and density. These are either induced by gradients in the etch chamber (usually with a radial pattern) or by pattern density variations on the wafer, which can change the local concentration of etch residues in the plasma.

When we model the wafer maps for waveguide width and thickness variations in the following examples, we use different techniques, as shown in Fig. 3. For the thickness variations we use an interpolated map measured on an actual wafer. For the waveguide width variations we use an OpenSimplex noise generator, which gives a uniform, isotropic random variation with a given amplitude and correlation length, similar to Perlin noise [23].

The three example circuits consist of a combination of waveguides and directional couplers. The waveguide model has a first-order dispersion model described by an effective index $n_{eff0} = n_{eff}(\lambda_0)$ at a center wavelength λ_0 and a group index $n_{g0} = n_g(\lambda_0)$. The wavelength dependent effective index is calculated as:

$$n_{eff}(\lambda) = n_{eff0} - \Delta\lambda \frac{n_{g0} - n_{eff}(\lambda_0)}{\lambda_0} \quad (2)$$

with $\Delta\lambda = (\lambda - \lambda_0)$. The model parameters $n_{eff}(\lambda_0)$ and $n_g(\lambda_0)$ have been precalculated for different design widths of the waveguide. And for each design width, the sensitivity of the

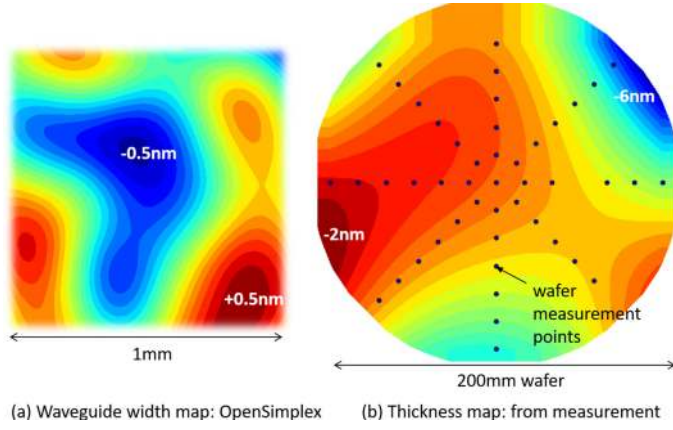


Fig. 3. Wafer maps of width and thickness variations used for the examples. (a) The width map uses an OpenSimplex noise generator, which gives uniform, isotropic fluctuations with a 1 nm peak-to-peak variation. (b) The thickness map was extracted from a wafer through ellipsometry and interpolated. It shows a deviation of $[-6 -2]$ nm from the nominal thickness of 220 nm.

model parameters n_{eff0} and n_{g0} are calculated as a second-order Taylor-expansion in width and thickness:

$$C(p) = C_0 + \frac{\partial C}{\partial X} \Delta X(p) + \frac{1}{2} \frac{\partial^2 C}{\partial X^2} \Delta X(p)^2 \quad (3)$$

For width and thickness sampling, sample points every $10 \mu\text{m}$ along the waveguide center line are added. This sampling distance is parametric, so it can be increased and decreased depending on the need, which depends on the typical correlation lengths of the variability maps we use.

The directional couplers are modelled as a small circuit consisting of four waveguides and a zero-length coupler, as illustrated Fig. 4. The waveguide segments model the phase propagation in the coupler, while the zero-length coupler provides a parametric coupling between the waveguides, with a $\pi/2$ phase delay for the cross coupling. The power coupling K in the zero-length coupler is modelled with a sinusoidal behavior as function of the total coupler length L [24]:

$$K(\lambda) = \sin^2(\kappa_0(\lambda) + \kappa'(\lambda) \cdot L) \quad (4)$$

The wavelength dependent behavior of the specific coupling $\kappa'(\lambda)$ is modelled as a second-order Taylor expansion around a center wavelength λ_0 :

$$\kappa'(\lambda) = \kappa'(\lambda_0) + \Delta\lambda \cdot \frac{\partial \kappa'}{\partial \lambda}(\lambda_0) + \frac{1}{2} \Delta\lambda^2 \cdot \frac{\partial^2 \kappa'}{\partial \lambda^2}(\lambda_0) \quad (5)$$

Similarly, the contribution of the bends is modelled as the lumped coupling κ_0 , which is also expanded as a Taylor series:

$$\kappa_0(\lambda) = \kappa_0(\lambda_0) + \Delta\lambda \cdot \frac{\partial \kappa_0}{\partial \lambda}(\lambda_0) + \frac{1}{2} \Delta\lambda^2 \cdot \frac{\partial^2 \kappa_0}{\partial \lambda^2}(\lambda_0) \quad (6)$$

This results in six model parameters for the directional coupler, which are extracted from measurement of a design sweep of directional couplers with different lengths.

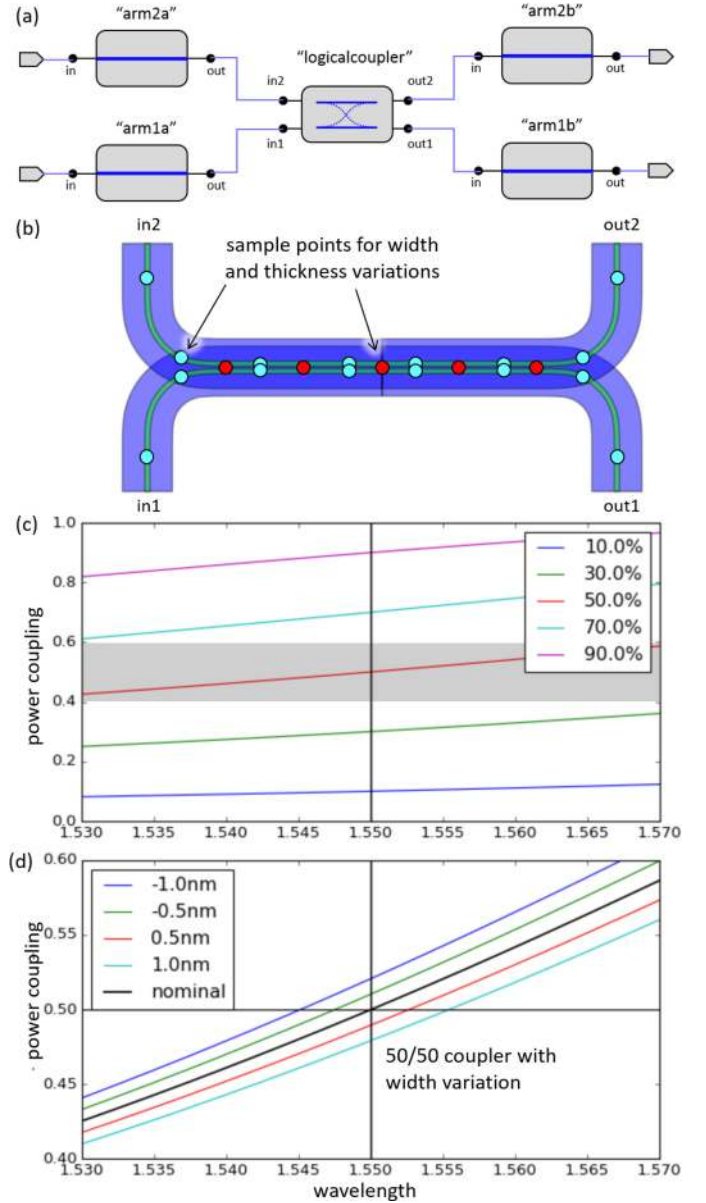


Fig. 4. Directional coupler circuit block. (a) Internally, the model of the directional coupler consists of 4 waveguide segments which capture the propagation, and a zero-length 'logical' coupler which models the actual coupling. The sensitivity of the coupler parameters and the waveguide parameters is modelled using a combination of eigenmode expansion and finite-difference time-domain simulations (FDTD) (b) To model the effect of variability, the width and thickness variations are sampled on multiple locations in the layout, both for the waveguides (cyan dots) and for the logical coupler (red dots). (c) Nominal coupler response for a coupler designed at different coupling ratios at a wavelength of 1550nm. (d) Response of the directional coupler with a nominal 50/50 splitting ratio but with different width variations, up to quite extreme cases of ± 10 nm.

A. Four-Channel Ring Resonator Demultiplexer

As a first example, we simulate the performance of a four-channel wavelength demultiplexer that consists of four ring resonators on a common bus waveguide [13], [25]. As illustrated in Fig. 5, the ring resonators are constructed from two directional couplers and two short waveguide segments. The directional coupler are designed with a nominal power coupling of 10% (at

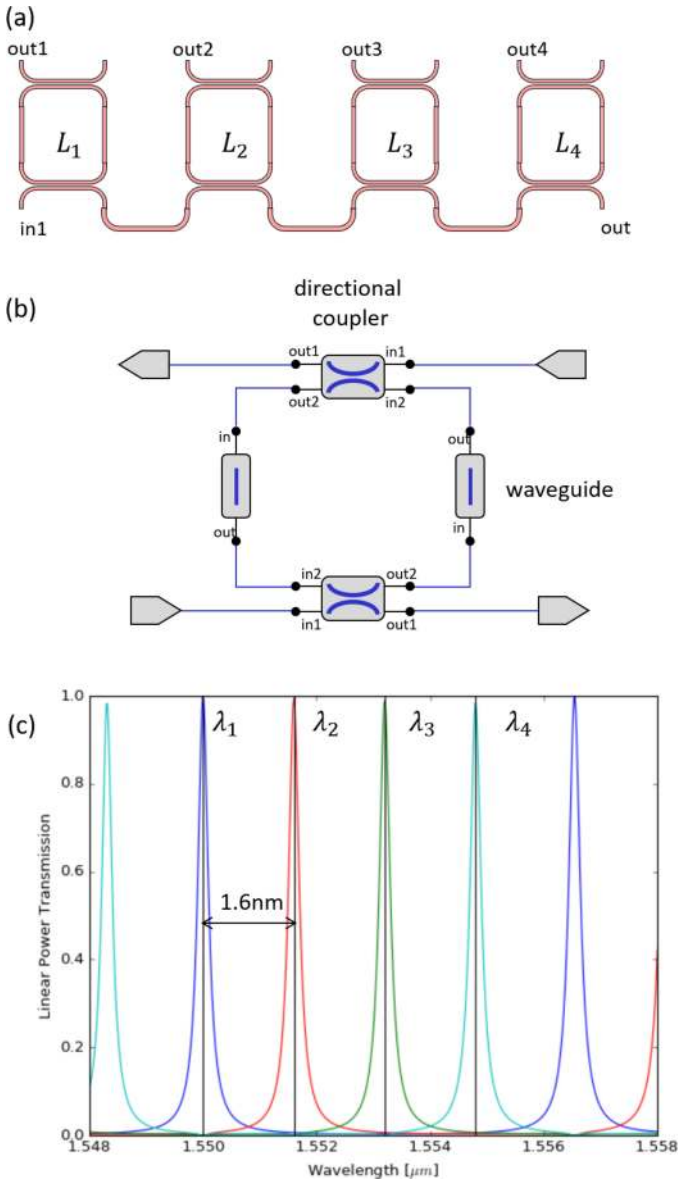


Fig. 5. (a) 4-channel wavelength demultiplexer based on ring resonators. (b) Each ring resonator consists of two directional couplers connected with straight waveguides of different length, to obtain resonances spaced 1.6 nm in wavelength (200 GHz in frequency). (c) Simulated transmission of the nominal devices at the four drop channel outputs.

the center wavelength of 1550 nm). The resonance wavelength of the rings is controlled by adjusting the lengths of the short waveguide segments in the ring, and the four rings are designed to have a wavelength spacing of 1.6 nm, which corresponds to 200 GHz channel spacing near 1550 nm. The nominal transmission of the four rings is also shown in Fig. 5c. For the operation of the wavelength demultiplexer, the channel spacing is important: when the channel spacing is sufficiently accurate, we can compensate wavelength shifts by collectively tuning all rings (e.g. using a heater/cooler in the package of the chip).

Figure 6 shows the results of a variability analysis. The top result (Fig. 6a) shows the performance of the ring demultiplexer over a wafer when we only consider waveguide width

variations, with an amplitude of 1 nm. In this design, the rings are widely spaced, with a center-to-center distance of 200 μm . We see that the transmission spectra are spread over approximately half the channel spacing, which is consistent with the aforementioned width sensitivity of 0.9 nm/nm. However, when we look at the histogram of channel spacings between every two adjacent channels in the same circuit, we see that there is a huge spread. Because the rings are spaced far apart, the width deviations within a single circuit are effectively decorrelated. As a result, the channel wavelengths within the same circuit are not evenly spaced, which ruins the functionality as a demultiplexer.

In Fig. 6b we have brought the four rings much closer together. Now the waveguide widths of the rings are correlated. While we see a similar spread in the transmission spectra over the wafer, the histogram of the channel spacings shows a much improved performance, with most channels spaced within 0.2 nm of the desired 1.6 nm channel spacing. This confirms quantitatively that the common-sense practice of grouping devices together for better matching is a valid strategy.

Finally, we also assessed the performance of the closely packed ring circuit when we add thickness variations. Figure 6c shows that the transmission spectra are now spread out over multiple channel spacings, which is logical given the large spread in thickness over the wafer. However, because thickness variations are mostly a long-range effect, this has very little impact on the relative channel spacing. Because the spread in channel spacing remains small, the demultiplexers remain useful if we assume a collective thermal tuning approach. The free spectral range (FSR) of the rings is not much affected by the thickness change, as this depends on the group index n_g of the waveguide, and this quantity is less dependent on the waveguide thickness variation than the effective index. The change in channel spacing induced by a collective thickness variation of the four rings is of the order of 5 pm.

B. Mach-Zehnder Lattice Filter

In a second example we construct a Mach-Zehnder lattice filter [7], [26] consisting of a cascade of directional couplers with different coupling coefficients and delay lines with the same length, as shown in Fig. 7. The lengths of the couplers are calculated from the filter coefficients of a Kaiser window, which are then mapped onto the coupling coefficients for the directional couplers [26]. The filter performance is determined by the exact coupling ratios, which distribute the light through the different paths in the filter. Also, all stages should introduce the same phase and time delay to guarantee the constructive/destructive interference. Errors in the coupling and phase delays will cause higher insertion loss and crosstalk levels, and a shift of the filter passband [27].

We designed this filter to have a pass band at 1550 nm wavelength, with a free spectral range of 800 GHz (6.4 nm) and a passband width of 80 GHz (0.64 nm). We allow for a wavelength shift of the passband of 40 GHz (0.32 nm), an insertion loss of -1 dB, a crosstalk level of -15 dB, and a nonuniformity within the passband of 1 dB.

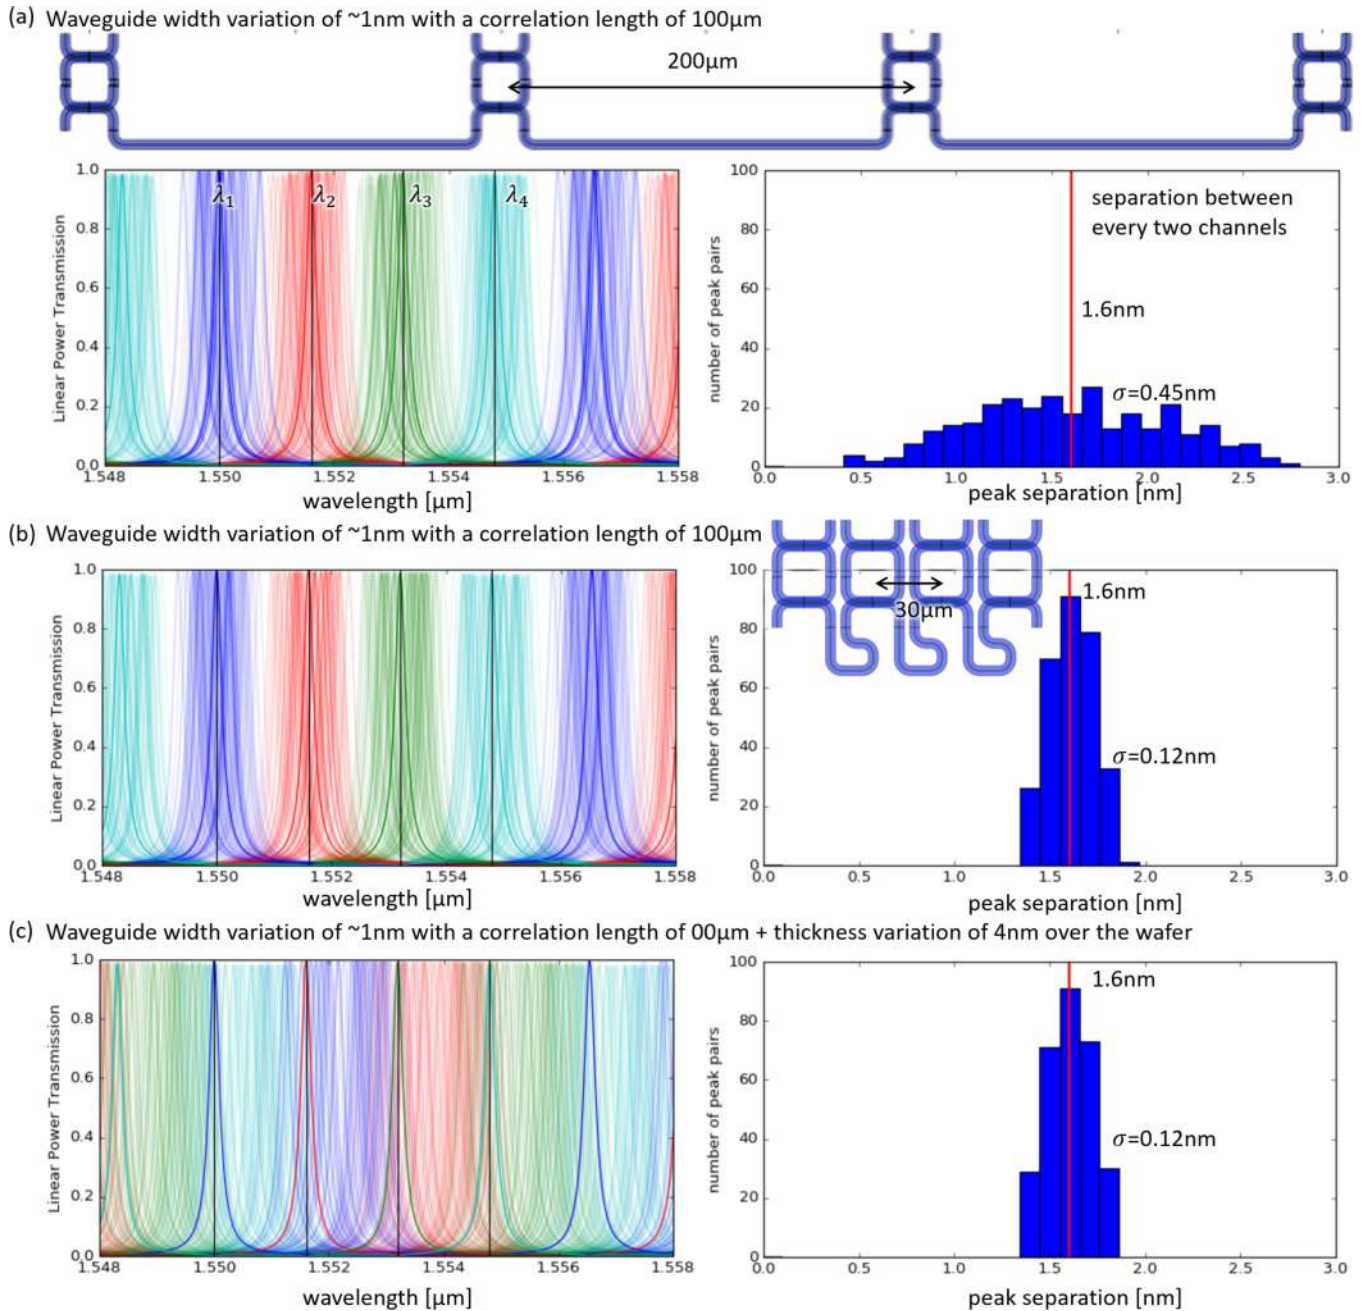


Fig. 6. Monte-Carlo simulations of the 4-channel ring demultiplexer in Fig. 5, using the wafer maps for width and thickness in Fig. 8. (a) Transmission of the four outputs with only the effect of the waveguide width variation, and the separation between every two adjacent channels in the same circuit. While the circuit is designed to have a uniform 1.6 nm channel spacing, the separation in the actual circuits is much larger. (b) Same circuit, but with the rings spaced close together in the layout. As a result, the width variations between adjacent rings is much smaller, and the channel separation is much better controlled, even if the absolute positioning of the spectra is similar as in the wider spaced circuit. (c) The same circuit as (b) but now simulated with both width and thickness variations. The large thickness variation causes extreme shifts of the absolute wavelength positions, but because thickness variation are a long-range effect, this has little impact on the relative channel spacing.

The simulation of an 8-stage filter is shown in Fig. 8. We also show the Monte-Carlo results of 277 samples on the wafer maps in Fig. 8, plotting both the pass band (blue) and the rejection band (red). We see that the nominal curve meets the specifications well, but the result of the variability analysis over the wafer paints a more dire picture.

Figure 9 shows the passband transmission curves, but now classified in three categories:

- Good filter devices that meet all the specifications (blue).
- Filter devices that meet the specifications, except for the absolute wavelength positioning of the passband.
- Filter devices that fail one or more of the criteria on insertion loss, ripple or crosstalk.

We see that the absolute wavelength positioning is a difficult criterion to meet. This can be intuitively understood from the sensitivity to thickness: the wafer map in Fig. 3 shows that the

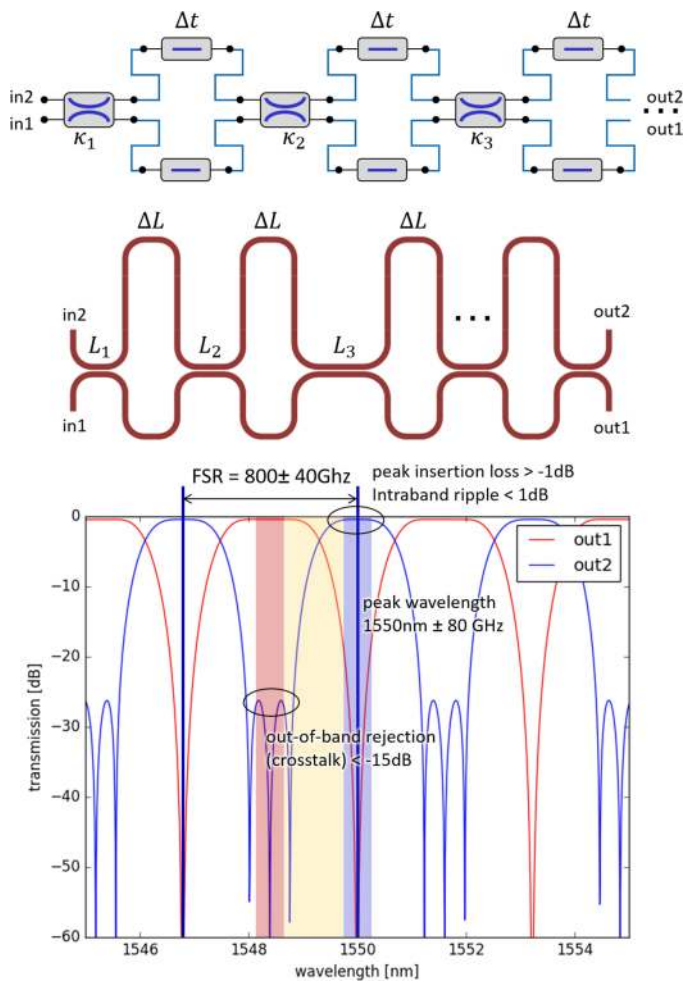


Fig. 7. Construction of a Mach-Zehnder lattice filter [26]. Based on filter coefficients, the coupling coefficients for the directional couplers are calculated, and then translated into the coupling lengths. In this example, we set specifications on the insertion loss, intra-band ripple, crosstalk, free spectral range (FSR) and the absolute peak wavelengths.

thickness over the entire wafer is at least 2 nm off the ideal value, which corresponds to a wavelength shift of 3-4 nm, which is half the designed FSR. With the absolute wavelength criterion, our yield is below 10%, and this is due to devices that have shifted a full FSR. When we ignore this criterion, we find that we get an overall yield of 77%.

With MZI lattice filters, adding additional stages should make the passband window more box-like, with a steeper roll-off and lower crosstalk. This is shown in Fig. 10a, which shows the design of the normalized passband when adding additional stages to the filters, increasing from 2 to 14. When translated into an actual circuit layout, the dispersion of the waveguides and the couplers has to be taken into account, and this will change the passband response. Figure 10 shows the passband response of these different filters, and we see that the response approximates the theoretical response close to the center wavelength of 1550 nm, but we see a sharp increase in crosstalk levels for shorter wavelengths. The effect of dispersion is actually most detrimental for filters with more stages: small variations in

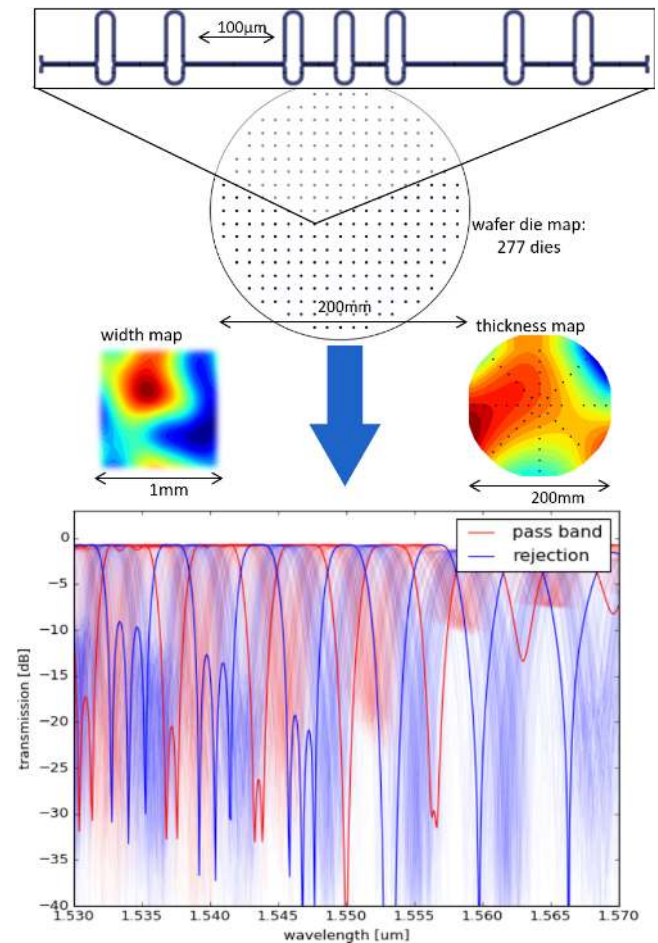


Fig. 8. Simulating the effect of waveguide width and thickness variability on an 8-stage MZI lattice filter. We applied a width variation map from Fig. 3, and simulated 277 Monte-Carlo samples. The plot shows the transmission spectra of the filter passband (blue) and rejection band (red).

coupling strengths will have a more profound impact when there are more couplers.

On top of wavelength dependence, the filter response will degrade as a result of width and thickness variations, as these also induce changes in the couplers. So while we can expect an improvement of the filter response when we increase the number of stages, we also expect a deterioration for more stages when the effects of variability start to kick in. This is clear from Fig. 10c, where we plot the yield predictions after a wafer-level simulation for each of these different lattice filters. We plot both the yield with absolute wavelength positioning, and the yield without absolute wavelength requirements. For only 2 stages, we basically have a simple Mach-Zehnder interferometer, for which it is hard to meet the design specifications. The yield goes up for 4 and 6 stages, but then goes down again for larger circuits, and for 10 stages there are no more working devices. Fig. 10d shows the transmission curves of the 6-stage device simulated over the wafer sampling points.

This shows that variability analysis should be an essential step in optimizing filter circuits. While experienced designers are able to identify empirically the optimum number of filter

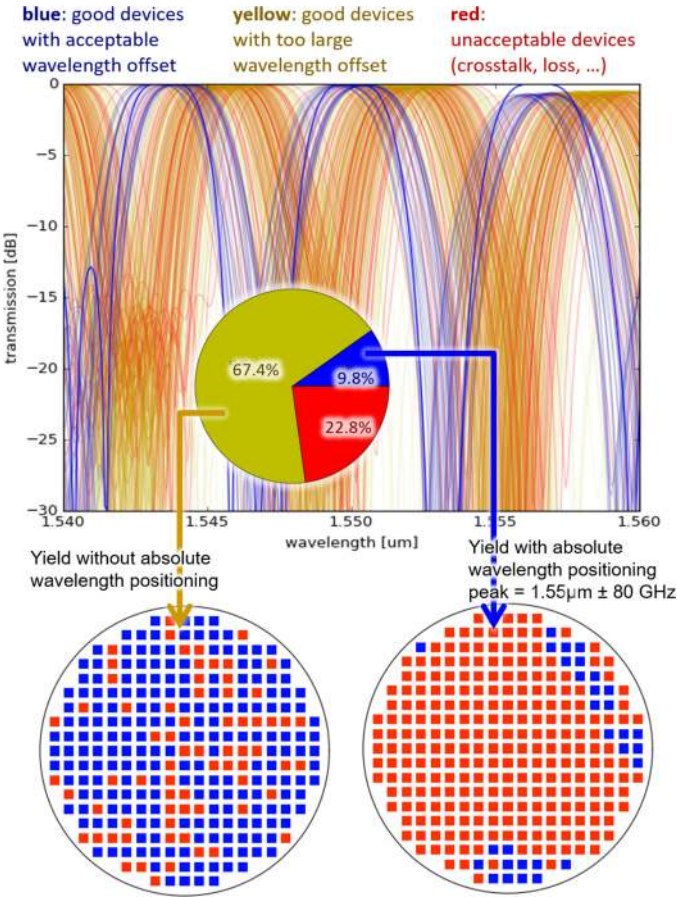


Fig. 9. Yield simulation of the MZI lattice filter with 8 stages designed to the specifications in Fig. 7 and simulated with the Monte-Carlo parameters described in Fig. 8. We differentiate between devices that pass all criteria, devices which are OK except for the peak wavelength position, and devices which fail to pass criteria such as crosstalk and insertion loss. The peak wavelength is a difficult criterion in the presence of width or thickness variations, as the effective index of silicon waveguides is particularly sensitive to geometry variations.

stages for certain specifications (or through a parameter sweep in fabrication), with location-aware variability analysis we can identify the sweet spot during the design stage.

C. Optimizing a Tolerant Mach-Zehnder Interferometer

As a final example, we show how layout-aware variability analysis can evaluate the effect on local variations on circuits that have been designed to be tolerant to those variations. For instance, [28] describes the design process for a Mach-Zehnder interferometer that can be made tolerant to waveguide width variations, thickness variations or temperature variations. The concept can be extended to multi-stage lattice filters like the one described in the previous section.

If we take a simple Mach-Zehnder interferometer with arm lengths L_1 and L_2 , it will have a sinusoidal response in wavelength, with constructive interference at

$$n_{eff}(L_1 - L_2) = m \cdot \lambda_0 \quad (7)$$

This we see in Fig. 11a. Because the waveguides are sensitive width variations, this peak wavelength will shift when the

width w changes from its nominal value. The sensitivity to width variations is

$$\frac{d\lambda_0}{dX} = \frac{\lambda_0}{n_g} \frac{dn_{eff}}{dX} \quad (8)$$

The principle behind the tolerant MZI is to use waveguides with a different width in both arms to decrease this wavelength shift. The different design widths mean that both arms now have a different effective index and group index. So the condition Eq. (7) for constructive interference at λ_0 becomes:

$$n_{eff1}L'_1 - n_{eff2}L'_2 = m \cdot \lambda_0 \quad (9)$$

But the two waveguides not only have a different index, the sensitivity of the effective index and group index to width variations is different. Broader waveguides will usually have less sensitivity to waveguide width variations. This means that the shift of peak wavelength λ_0 can now be written as

$$\frac{d\lambda_0}{dw} = \frac{\lambda_0}{n_{g1}L'_1 - n_{g2}L'_2} \left(L'_1 \frac{dn_{eff1}}{dw} - L'_2 \frac{dn_{eff2}}{dw} \right) \quad (10)$$

If we make the assumption that both arms will experience the same waveguide width change, we can make the MZI response tolerant to global waveguide width changes by forcing Eq. (10) to zero. We get the following condition:

$$\frac{L'_1}{L'_2} = \frac{dn_{eff2}/dw}{dn_{eff1}/dw} \quad (11)$$

With Eq. (9) and (11) we have two equations to determine the two lengths L'_1 and L'_2 . The remaining free parameter is the interference order m , which can be chosen to approximate a desired free spectral range.

In this example, we choose the waveguide widths to be $w_1 = 400$ nm and $w_2 = 500$ nm. We aimed for a free spectral range as close as possible to 10 nm, which results in an interference order $m = 105$. Because we use the original waveguide width of 450 nm for the directional coupler, a taper is needed to interface the directional coupler with the different arm waveguide. To make sure both arms have the same set of taper, we add a section of narrow waveguide to the wide arm, and a section of wide waveguide to the narrow arm. this way, each arm contains the same three tapers, and the only difference is in the different waveguides.

To assess the impact of both long-range (global) waveguide width variations and short-range (local) width variations, we provide a wafer map for waveguide width that consists of two components

- an OpenSimplex map with a correlation length of 1 cm and an amplitude of 10 nm to model long-range variations, much larger than the circuit,
- an OpenSimplex map with a correlation length of 100 μ m and an amplitude of 2 nm to mimic more random short-range fluctuations in waveguide width.

The results of the Monte-Carlo simulations over a virtual wafer are shown in Fig. 11. We see that the standard MZI with a single waveguide width (Fig. 11a) shifts over a wide range. The tolerant filter in Fig. 11b has a much lower variation, but it is still significant, and much larger than the sensitivity suggested in

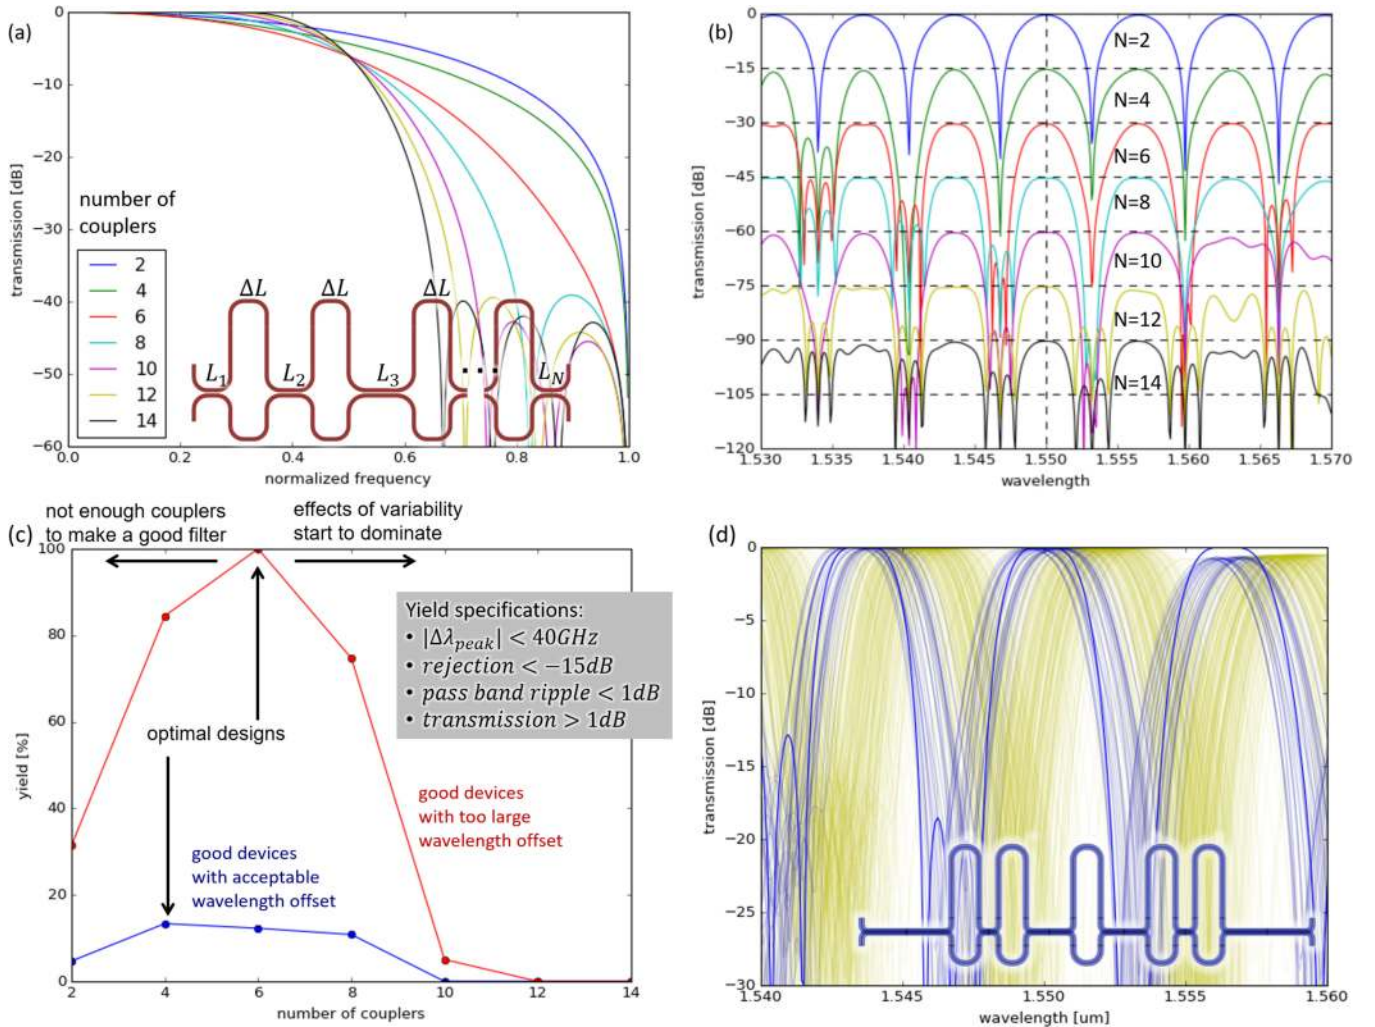


Fig. 10. Optimization of the lattice filter specified in Fig. 7 to maximize the yield under variability of width and thickness. (a) Normalized baseband response of lattice filters with increasing number of couplers (taps), resulting in an increasingly box-like profile. (b) Simulated nominal filter response of the implemented filter circuit for these different filters. For clarity, the responses are separated along the y-axis by -15 dB, with each new 0 dB reference indicated with dashed lines. These also serve as a visual reference for the -15 dB specification on crosstalk. (c) Fraction of filters that pass the design criteria in a virtual wafer Monte-Carlo simulation as outlined in Fig. 8. (d) Transmission spectra of the 6-stage lattice filter, which is the best filter when ignoring the requirement of absolute peak wavelength position.

[28]. The origin of these fluctuations are in the local variations. With a size of $100 \mu\text{m} \times 130 \mu\text{m}$. The tolerant device size is larger than the correlation length of the width variations. On top of that, both arms are on opposite sides of the circuit. So, while the global variations cancel out, the local width variations still contribute to the wavelength shift.

Using layout-aware variability analysis, we can now look into possible improvements. Reorganizing the circuit in such a way that the waveguides are closer together should reduce the effect of local variability. Such a layout is shown in Fig. 11c: nesting the waveguides not only reduces the device footprint, but also runs the waveguides of both arms largely parallel. In terms of waveguide lengths and building blocks, this circuit is equivalent to the circuit with opposing arms. However, the Monte-Carlo simulations demonstrate that this nested layout performs much better in the face of waveguide width variability.

Again, this is an effect that an experienced designer might anticipate, so he might choose a nested layout accordingly. However, layout-aware variability analysis makes it possible to quantitatively assess the impact on the circuit yield.

V. DISCUSSION

The three examples show that layout-aware variability analysis can expose potential problems in circuit due to variability that cannot be quantitatively predicted with other techniques. It also makes it possible to optimize layouts for reduced sensitivity to fabrication variables such as width and thickness.

The applicability of this technique could have some limitations. A reliable analysis is only possible if the necessary data is available. First of all, the sensitivity of the device model parameters is needed. Ideally, this is provided by the device designer or the fab. In case this data is not provided, the circuit designer

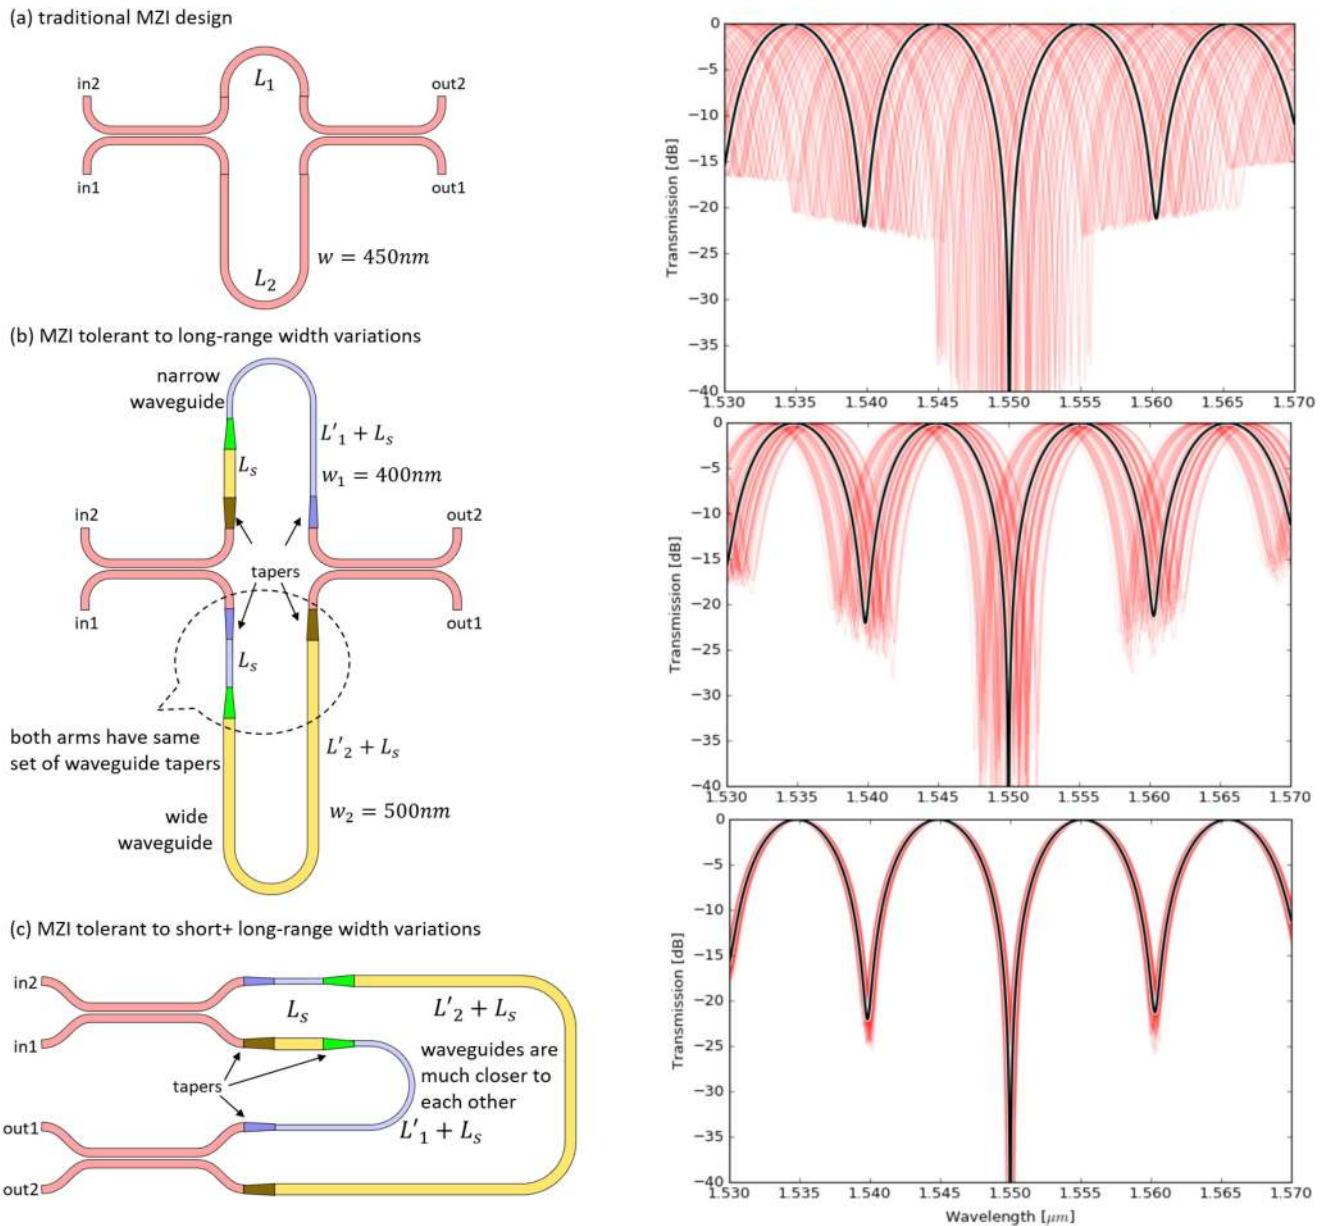


Fig. 11. Testing the design of a Mach-Zehnder interferometer that is tolerant to waveguide width variations [28]. The devices were simulated over a virtual wafer with both long-range (5 cm) width variations with an amplitude of 10 nm, and short-range ($100 \mu\text{m}$) variations with an amplitude of 2 nm. (a) A regular MZI, without any design for tolerance to waveguide width variations. The transmission peaks shift with 0.9 nm/nm width change. (b) A tolerant MZI, designed according to [28], using waveguides of different widths, including short compensation sections to balance the use of tapers in both arms. This device is tolerant to long-range width variations where both arms feel the same effect, but not to short-range variations where the arms feel a different effect. (c) A folded version of the tolerant device, where the arms are brought closer together, shows a good tolerance to both long-range and short-range width variations.

could invest the effort to simulate or measure these effects, but this is not always possible. In some cases, fabs do not provide the designer with the actual layout/geometry of a building block. Such 'black-box' building blocks are substituted by the actual geometry only at the fabrication stage, in order to protect the intellectual property of the fab. In this situation, the designer cannot simulate the building block with modified width /thickness variations, or other variables such as doping implant dose and energy.

Similarly, the Monte-Carlo analysis is only meaningful if the circuit designer has a representative model of the

wafer-scale variability of the relevant external variables, such as waveguide width and thickness. Because these numbers could give unnecessarily detailed insight into a fab's fabrication process quality, fabs are generally reluctant to release such data, and variations are often specified as guaranteed bounds that are not really representative of the actual spatial distribution.

A possible solution would be to map the physical variables such as waveguide width and thickness onto a new set of variables that without a direct physical meaning. Also, the sensitivities of the building blocks could be remapped to these abstract

variables. This way, the fab can provide, as part of the PDK, a set of building blocks with models, including sensitivity, and some generator functions to generate the maps of these new abstract variables. The drawback of this technique is that it is not possible for a device designer to test new geometries against the variability of the fab, as the transformation of the abstract variables into the physical dimensions is not disclosed.

The tools developed here for layout-aware variability analysis are not limited to evaluating fabrication-related variability. It could also be used to model the effects of operational variability. For instance, the effect of temperature gradients on a component could be evaluated, or the impact of the placement of a hot component (e.g. a laser diode) on the chip.

The main weakness of our variability analysis tool is that it still relies on Monte-Carlo simulations. While these simulations can be parallelized, they can still become computationally expensive, especially with larger circuits. The small filter circuits used in this example consisted of 20–50 components (including waveguides, tapers, and couplers), and were simulated for 2000 wavelength points. A Monte-Carlo simulation run for one of these circuits required between 5 and 10 minutes, depending on the complexity of the circuit, running on a single processor core of a powerful laptop with a Core-i7 processor. So while computationally expensive, parallelization on moderate cluster infrastructure of a few tens or hundreds of cores brings the simulation time down to less than 1 minute, which is acceptable for interactive design.

VI. CONCLUSION

Variability analysis and yield prediction at circuit level will become more important as circuit complexity increases and designers become disconnected from the foundry fabrication processes. We developed a set of variability extensions for the commercial circuit simulator Caphe, to evaluate the effect of process variability on photonic integrated circuits. This is especially critical for wavelength filter design in high-contrast waveguide systems such as silicon photonics. The CapheVE engine combines knowledge of the circuit layout with virtual wafer maps of fabrication variables such as waveguide width and thickness and circuit models with sensitivity data to these variables. Using Monte-Carlo simulations over these virtual wafers, the circuit performance can be evaluated, or the circuit layout can be optimized to maximize fabrication yield.

REFERENCES

- [1] X. Chen *et al.*, “The emergence of silicon photonics as a flexible technology platform,” *Proc. IEEE*, vol. 106, no. 12, pp. 2101–2116, Dec. 2018.
- [2] P. Dong, Y.-K. Chen, G.-H. Duan, and D. T. Neilson, “Silicon photonic devices and integrated circuits,” *Nanophotonics*, vol. 3, no. 4/5, pp. 215–228, 2014.
- [3] A. Rahim, T. Spuesens, R. Baets, and W. Bogaerts, “Open-access silicon photonics: Current status and emerging initiatives,” *Proc. IEEE*, vol. 106, no. 12, pp. 2313–2330, Dec. 2018.
- [4] W. Bogaerts and L. Chrostowski, “Silicon photonics circuit design: methods, tools and challenges,” *Laser Photon. Rev.*, vol. 1700237, pp. 1–29, 2018.
- [5] D. X. Xu *et al.*, “Silicon photonic integration platform—Have we found the sweet spot?” *IEEE J. Sel. Topics Quantum Electron.*, vol. 20, no. 4, Jul./Aug. 2014, Art. no. 8100217.
- [6] 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,” *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 316–324, Jan./Feb. 2010.
- [7] W. Bogaerts *et al.*, “Silicon-on-insulator spectral filters fabricated with CMOS technology,” *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 33–44, Jan./Feb. 2010.
- [8] G. Fishman, *Monte Carlo: Concepts, Algorithms, and Applications*. Berlin, Germany: Springer, 2013.
- [9] D. Cassano, F. Morichetti, and A. Melloni, “Statistical analysis of photonic integrated circuits via polynomial-chaos expansion,” *Adv. Photon.*, no. 2, 2013, Paper JT3A.8.
- [10] T.-W. Weng *et al.*, “Uncertainty quantification of silicon photonic devices with correlated and non-Gaussian random parameters,” *Opt. Express*, vol. 23, no. 4, pp. 4242–4254, 2015.
- [11] A. Waqas, S. Member, D. Melati, and A. Melloni, “Sensitivity analysis and uncertainty mitigation of photonic integrated circuits,” *J. Lightw. Technol.*, vol. 35, no. 17, pp. 3720–3728, 2017.
- [12] Y. Xing, D. Spina, A. Li, T. Dhaene, and W. Bogaerts, “Stochastic collocation for device-level variability analysis in integrated photonics,” *Photon. Res.*, vol. 4, no. 2, pp. 93–100, 2016.
- [13] Z. Lu *et al.*, “Performance prediction for silicon photonics integrated circuits with layout-dependent correlated manufacturing variability,” *Opt. Express*, vol. 25, no. 9, pp. 9712–9733, 2017.
- [14] M. Fiers *et al.*, “Time-domain and frequency-domain modeling of nonlinear optical components at the circuit-level using a node-based approach,” *J. Opt. Soc. Amer. B*, vol. 29, no. 5, pp. 896–900, 2012.
- [15] Y. Xing, J. Dong, S. Dwivedi, U. Khan, and W. Bogaerts, “Accurate extraction of fabricated geometry using optical measurement,” *Photon. Res.*, vol. 6, no. 11, pp. 1008–1020, 2018.
- [16] M. S. Eldred, “Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design,” in *Proc. 11th AIAA Non-Deterministic Approaches Conf.*, May 2009, pp. AIAA 2009–2274.
- [17] D. Spina *et al.*, “Efficient variability analysis of electromagnetic systems via polynomial chaos and model order reduction,” *IEEE Trans. Compon., Packag. Manuf. Technol.*, vol. 4, no. 6, pp. 1038–1051, Jun. 2014.
- [18] Y. Xing, J. Dong, U. Khan, and W. Bogaerts, “Hierarchical model for spatial variations of integrated photonics,” in *Proc. IEEE Int. Conf. Group IV Photon.*, 2018, pp. 91–92.
- [19] M. Guiney and E. Leavitt, “An introduction to open access an open source data model and API for IC design,” in *Proc. Des. Autom., Asia South Pac. Conf.*, 2006, pp. 434–436.
- [20] W. Bogaerts, M. Fiers, M. Sivilotti, and P. Dumon, “The IPKISS photonic design framework,” in *Proc. Opt. Fiber Commun. Conf.*, 2016, pp. 2–4.
- [21] W. Bogaerts and Y. Li, “Integrated design for integrated photonics: From the physical to the circuit level and back,” *Proc. SPIE*, vol. 8781, 2013, Art. no. 878102.
- [22] K. Qian, B. Nikolić, and C. J. Spanos, “Hierarchical modeling of spatial variability with a 45 nm example,” *Proc. SPIE*, vol. 7275, 2009, Art. no. 727505.
- [23] K. Perlin, “An image synthesizer,” *SIGGRAPH Comput. Graph.*, vol. 19, no. 3, pp. 287–296, 1985.
- [24] Y. Xing, U. Khan, A. Riberio, and W. Bogaerts, “Behavior model for directional coupler,” in *Proc. Symp. IEEE Photon. Soc. Benelux*, Delft, The Netherlands, 2017, pp. 128–131.
- [25] Y. Xing *et al.*, “From parameter extraction, variability models to yield prediction,” *Latin Amer. Opt. Photon.*, Lima, Peru, 2018, Paper W3E.1.
- [26] K. Jinguji and M. Oguma, “Optical half-band filters,” *J. Lightw. Technol.*, vol. 18, no. 2, pp. 252–259, Feb. 2000.
- [27] W. Bogaerts, U. Khan, and Y. Xing, “Layout-aware yield prediction of photonic circuits,” in *Proc. IEEE Int. Conf. Group IV Photon.*, 2018, pp. 93–94.
- [28] S. Dwivedi, H. D’Heer, and W. Bogaerts, “Maximizing fabrication and thermal tolerances of all-silicon fir wavelength filters,” *IEEE Photon. Technol. Lett.*, vol. 27, no. 8, pp. 871–874, Apr. 2015.



Wim Bogaerts received the Ph.D. degree in the modelling, design and fabrication of silicon nanophotonic components from Ghent University, Ghent, Belgium, in 2004.

He is currently a Professor with the Photonics Research Group, Ghent University—IMEC, Ghent, Belgium. During this work, he started the first silicon photonics process on IMEC's 200 mm pilot line, which formed the basis of the multiproject-wafer service ePIXfab. His current research focuses on the challenges for large-scale silicon photonics: Design methodologies and controllability of complex photonic circuits and main research interests include telecommunications, information technology and applied sciences. In 2014, he co-founded Luceda Photonics to develop software solutions for silicon photonics design, in particular the IPKISS design framework. Since 2016, he is again a Full-Time Professor with the Ghent University, looking into novel topologies for large-scale programmable photonic circuits, supported by a consolidator grant of the European Research Council (ERC).

Dr. Bogaerts is a member of Optical Society of America (OSA) and SPIE.



Umar Khan received the Ph.D. degree from Tyndall National Institute, University College Cork, Cork, Ireland, for his work on "Design and implementation of micro-structures with refractive index contrast for optical interconnects and sensing applications." After completing his Ph.D., he joined Photonics Research Group, Ghent University—IMEC, Belgium, as a Postdoctoral Researcher. His current research interests include compact models, variability aware design, parameters extraction, and programmable circuits.



Yufei Xing received the bachelor's degree in optical engineering from Zhejiang University, Hangzhou, China, in 2011 and the master's degree from European Master of Science in Photonics, Ghent University, Ghent, Belgium, in 2013. He is currently working toward the Ph.D. degree with the Photonics Research Group, Ghent University—IMEC, Ghent, Belgium. His current research interests include variability aware design, compact circuit models, and parameter extraction.