

Capturing the Effects of Spatial Process Variations in Silicon Photonic Circuits

Yufei Xing, Jiaying Dong, Umar Khan, and Wim Bogaerts*

Cite This: <https://doi.org/10.1021/acsp Photonics.2c01194>

Read Online

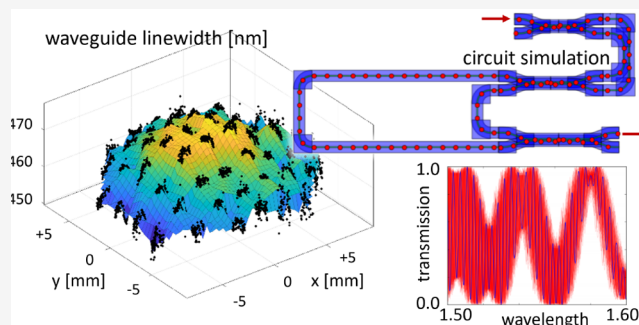
ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: Silicon photonic devices are very sensitive to process variation, and it is important for circuit designers that they can predict the effect of this variability during the design phase, and optimize their design for both performance and yield. This requires an accurate predictive model of the spatial variations induced by the fabrication process. We present in this paper a method to extract a granular map of the line width and thickness variation on a silicon photonics wafer. We propose a hierarchical model to separate the layout-dependent and location-dependent systematic process variation from the random process variation on different spatial levels. We identify the relative contributions to width and thickness variations and use this to construct a synthetic model for virtual wafers that can be used to analyze the effect on circuit behavior and eventually predict the yield of photonic circuits after fabrication. We observe that the main contribution to waveguide width and thickness variations are systematic, and that die-scale systematic line width variation is correlated with local pattern density.

KEYWORDS: silicon photonics, design optimization, variability, process variations, yield estimation



1. INTRODUCTION

Silicon photonics is rapidly scaling toward large-volume production of ever more complex photonic integrated circuits. This scaling is enabled by the use of CMOS manufacturing technologies and the high refractive index contrast between silicon and silicon oxide, which allows strong light confinement and a compact waveguide footprint in the silicon-on-insulator (SOI) layer.¹ However, this high material contrast also makes silicon photonic circuits sensitive to minute variations in the fabricated geometries (in particular the line width or thickness of submicron silicon waveguides). This is especially problematic as circuit complexity scales up, because variations propagate and accumulate in a circuit. As most photonic circuits are designed to process analog signals, this will lead to performance degradation and low-yield production.

There are multiple approaches to improve the yield in larger circuits. First of all, the fabrication processes themselves can be improved, by fine tuning the process steps or moving to more advanced (and more expensive) fabrication tools. Obviously, this is not always possible within the constraints of a specific foundry without significant investments. Second, the individual circuit building blocks can be optimized for performance, but also for increased robustness against fabrication variations. Such optimizations can be performed by experimental parameter sweeps, or by heavy numerical simulations. In the past few years, inverse design techniques have become more popular for this

purpose.^{2–4} Multiobjective optimizations can then be used to make these geometries more robust.⁵

A third method, which is complementary to the previous two, is to optimize the circuit itself for better performance in the presence of variability, without changing the fabrication process or the library of standard building blocks. As we know that variability cannot be fully eliminated, this technique is an essential part of the design toolbox to scale up to larger circuits. Of course, this requires that we can accurately predict the actual fabrication variations and their effect on the circuit performance during the design phase.⁶

Most effects of variability originate from the fabrication process, but the effects are measurable when characterizing the device, circuit, or system performance. This is illustrated in Figure 1.⁶ Process variations such as lithography exposure dose, resist age, etch plasma density, chemical mechanical polishing (CMP) slurry composition pressure distribution, and wear of the polishing pads lead to nonuniformity in device geometry such as width and thickness. The high index contrast of SOI

Special Issue: Optimized Photonics and Inverse Design

Received: August 4, 2022




Figure 1. Variability in photonic circuits is present at different levels. Variability information collected by the foundry (processes, geometries) must be translated to the functional level so designers can incorporate this information into their design and optimization flow. Reprinted (adapted) with permission from ref 6. Copyright (2018) John Wiley and Sons.

waveguides translates even small geometry variations into a non-negligible variation in the optical properties of the circuit elements, which then propagate into the behavior of the circuit and ultimately the performance of the system. Depending on the desired functionality, only a fraction of the fabricated chips will work within specifications. This fraction is usually called the *yield* of the fabrication process for this specific circuit design. As circuits increase in complexity, the effects of variability in the many building blocks accumulate, and the yield will go down. Passive wavelength filter circuits are exceptionally prone to the accumulation effects of fabrication variations, as they rely on precise control of the optical path length through a set of delay lines.⁷ A 1 nm width or thickness variation in these delay lines will lead to a shift in wavelength filter of 1–2 nm, which can already span multiple communication bands in a dense wavelength-division multiplexing (DWDM) link.

In this paper, we base ourselves on the fabrication process of IMEC's silicon photonics process.⁸ It is a so-called “thin SOI” process, with a waveguide silicon-on-insulator core layer of 220 nm thickness. A typical single-mode waveguide is designed as a

silicon strip of 450 nm width, which results in a waveguide of approximately 470 nm width after lithography and etching. The entire process flow has many additional steps, and also different waveguide geometries, but as these strip or wire waveguides are the most commonly used to define wavelength filtering circuits, this work only looks into the variability of this critical process step.

There are many other silicon photonics fabrication platforms, and most are based on a similar thin SOI stack with a waveguide layer of 200–300 nm thickness. Table 1 lists reported within-wafer fabrication variations for various silicon photonic platforms. The 6σ spread of both waveguide width w and thickness t variation within a wafer are usually in excess of 15 nm, an order of magnitude larger than the 1 nm variation described earlier. Ignoring the impact of process variation could lead to circuit failure even when active tuning capabilities are available in the platform. For instance, a tunable beam coupler implemented as a balanced Mach–Zehnder interferometer with phase tuners in the arms can only tune over the full range of 0–100% power coupling if the static beam splitters at the input

Table 1. Literature Results for Within-Wafer Fabrication Variations

process	within-wafer variations	
	$\sigma[\Delta_w]$ (nm)	$\sigma[\Delta_t]$ (nm)
IMEC 193 nm dry lithography ^{9,10}	2.59	2
IMEC wafer-scale corrective etching ¹¹		0.83
IMEC wafer-scale corrective etching ¹²		3.64
IMEC 192 nm immersion lithography ¹³	2.53	
IMEC 200 mm wafer, 193 nm dry lithography ¹⁴	0.78	
IMEC 300 mm wafer, 193 nm immersion lithography ¹⁴	2.65	
IMEC 193 nm lithography ¹⁵		2.4
AMF 248 nm lithography ¹⁶	2	
AMF 248 nm DUV lithography process ¹⁷	3.86	1.32
AMF 193 nm lithography ¹⁸	6.4	2.4
AIM 193 nm immersion lithography ¹⁹	~2.5	
GF 193 nm immersion lithography ²⁰	3.3	

and outputs have a perfect 50:50 splitting ratio. The design can be made more tolerant by adding extra stages to the MZI, which relaxes the requirement on the splitting ratio of the static beam splitters to 25:75 for a two-stage MZI²¹ and 15:85 for a three-

stage circuit.²² But if the beam splitters deviate more than a few percent from each other within the same tunable coupler circuit, it becomes very difficult to still cover the full 0–100% tuning range.²³

To assess the impact of variability on photonic circuits, the variability effects during processing and at the level of building block geometries need to be properly mapped onto the functional behavior of the circuits. Initial variability studies started from the assumption that the fabricated geometry parameters on a photonic chip such as line width and thickness^{24,25} or device behavior parameters such as coupling coefficients²⁶ can be modeled as totally random and independent of the absolute location on the wafer, or even positioning of devices relative to one another within the same circuit. However, it is well-known that this is an incorrect assumption: there is a spatial correlation between the properties of devices located closer together,^{27,28} and we have also shown that there is also a strong correlation between device properties and the density of the patterns in the immediate vicinity on the wafer.²⁹ Most fabrication processes try to minimize these so-called “loading” effects by adding tiling or dummy structures in between the circuits, but this is not always possible (e.g., in the slab area of a star coupler).

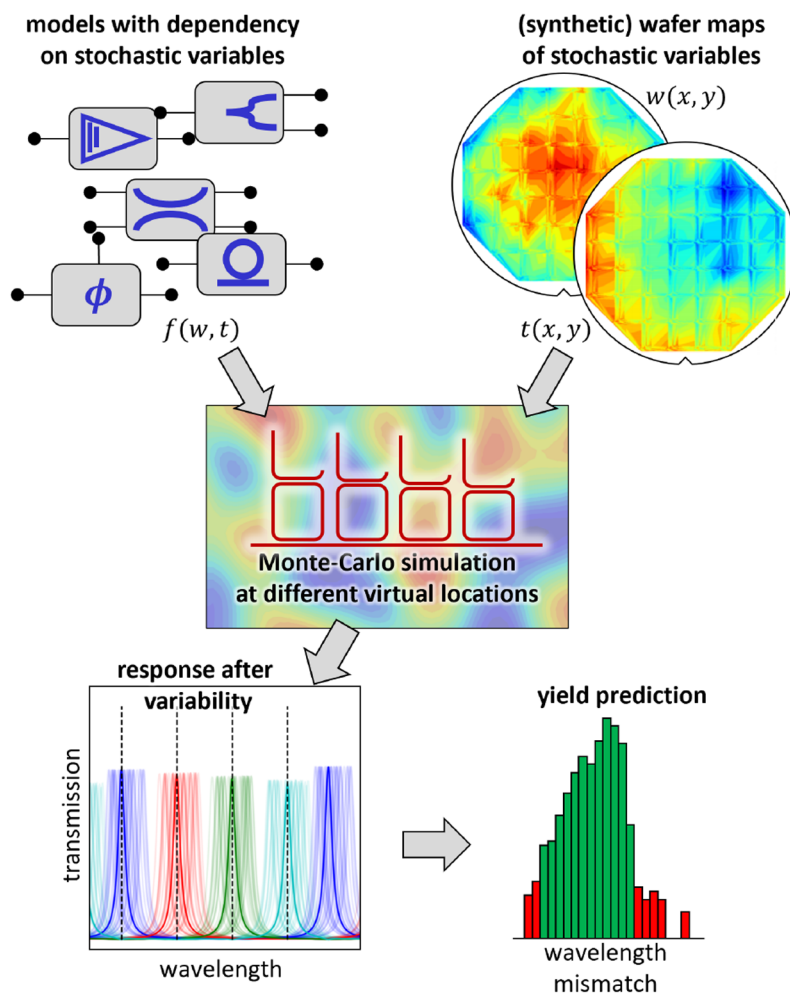


Figure 2. Simulation flow for layout-aware variability analysis.³⁰ The effect of variations in geometric parameters (w , t) is translated into variations of model behavior. Circuits are then projected onto wafer maps of w and t in different locations and simulated in Monte Carlo fashion, which results in a set of transmission curves that can be used for yield assessment.

As a result, not only does the actual circuit layout on the chip (i.e., the choice and positioning of the building blocks) influence the proximity of devices, but also the resulting lithography pattern will induce a systematic process variation that will be repeated on each die. Overall, we can identify effects of process variations that are location-dependent (where the circuit is placed on the wafer or within the die), layout-dependent (how the circuit schematics are translated into a mask layout), and purely random, i.e., independent of location. A realistic variability analysis and yield prediction should incorporate all these elements.

In ref 30, we earlier proposed a method to model these different effects taking into account the actual layout of a circuit, and a similar technique has been proposed in ref 17. Figure 2 illustrates how a circuit layout is projected onto wafer maps of geometric variations (e.g., width w or thickness t of a waveguide). Using behavioral models that contain a dependency on w and t , the transmission of the circuit at that wafer location can be calculated, taking into account the local geometries for each circuit building block. Using a Monte Carlo process using different wafer maps and wafer locations, the variability of the circuit response can be calculated, and from this, the yield can be estimated. As also shown in ref 30, this can then be used to rearrange the layout of the circuit such that the effect of both local and global variability is minimized.

One of the key challenges in this method is to generate a realistic wafer map of the geometry variations (we will use the width w and thickness t of a strip waveguide as typical parameters for the rest of the paper) that captures the statistics and the location-dependence of the different contributions to the variability. For this, we need to do two things: First, obtain an accurate and granular wafer map of the process variation. Second, build a realistic wafer model that decomposes variation into systematic and stochastic contributions on different length scales that allows us to predict typical variations in future fabrication runs.

In this paper, we discuss such an approach to collect and separate process variations on different spatial scales and apply them for variability and yield prediction for future designs. In section 2 we present a spatial classification scheme of process variations, and then section 2.7 takes us step by step through the process to separate the different contributions. Subsequently, we quantitatively analyze the process variation extracted from a 200 mm silicon photonics wafer in section 3 and apply the extracted model to synthesize virtual wafers which can be used for yield prediction.

2. METHOD: HIERARCHICAL SPATIAL VARIABILITY MODEL

2.1. Classification of Variations. For the purpose of our discussions in this paper, *variations* describe a deviation from the nominal properties or behavior of a photonic circuit that is not identical for every copy of the component or circuit, or at any given time. We can classify these variations in different ways.

2.1.1. Fabrication and Operational Variations. With *fabrication variations* we mean the differences between photonic circuits (and differences within one circuit) that are generated during the fabrication, and which are permanently incorporated in the chip. These usually take the form of variations in material distribution, such as the line width or thickness of a waveguide. *Operational variations*, on the other hand, are induced by external factors during the lifetime of the chip. Typical examples are temperature gradients, fluctuations in the supply voltage, or

aging effects. In this paper, we focus on fabrication variations, but some techniques can also be applied to operational variations. The two are not always fully decoupled: variations induced during fabrication can later give rise to operational variations.

2.1.2. Temporal and Spatial Variations. Variations can have spatial and temporal components. Spatial variability is the nonuniformity that depends on the device or circuit location on the wafer or the distance between two circuits, and can include for instance the local line width or thickness over a wafer. Temporal variations indicate changes that vary over time, and can range from near-instantaneous transient effects³¹ over microsecond-scale linear or nonlinear self-heating, to year-long effects such as dielectric material deterioration. There is also a temporal component in the fabrication process, which could relate to the season of fabrication, shelf life of chemicals, maintenance schedules of tools, or the wafers that were processed on the tool prior to the photonics batch. To complicate matters, during the fabrication process, temporal variations can give rise to spatial variations. An example is the aging of photoresist which can impact the lithography uniformity over the wafer or between wafers.

2.1.3. Systematic Variation and Random Variation. Process variations consist of systematic contributions and random contributions. Systematic or deterministic variations denote a repeatable deterministic pattern that can be deduced from a known variable, such as the circuit's location on the die or wafer. For example, for plasma deposition and etch processes, there can be a slow-varying variation in the plasma composition and density, which induces a deterministic variation in etch or deposition rate. Or it can be a repeated pattern on every die induced by pattern deviations on the photomask or optical proximity effects. The systematic variations can also be due to nonidealities in the lithographic system such as defocus and misalignment, errors in the photomask fabrication, or pattern-density-dependent effects in processes such as CMP or plasma etching.

Random or stochastic variations refer to nondeterministic effects where only the properties of stochastic distributions can be modeled. They are random in nature and cannot be correlated to time and location. Contributing effects are fluctuations in the lithographic exposure dose from die to die, microscale line width variations due to sidewall roughness,³² photoresist granularity,³³ lithographic shot noise, or atomic-scale oxide-thickness variation.³⁴

Systematic and random variations impact the device and circuit performance in different ways. Systematic variations can be predicted by the device or circuit locations on a wafer, and there can be a location-dependent correlation between devices and circuit (i.e., how the devices are placed inside the larger circuit). On the other hand, random variation is independent of circuit position and it is usually modeled as a set of independent fluctuations. In practice, the classification between systematic and random variation is not absolute. Often, the exact location of the design on the wafer or spatially related process features (such as local pattern density) are not available in full detail to the designers. So, it is impossible to predict the variation with certainty. In such cases, some components of the systematic variations can be treated as stochastic variables.

In addition, the model for the systematic variations might not describe the actual systematic variations with sufficient accuracy, which can result in an overestimation of the stochastic contribution. For example, we assume that the systematic

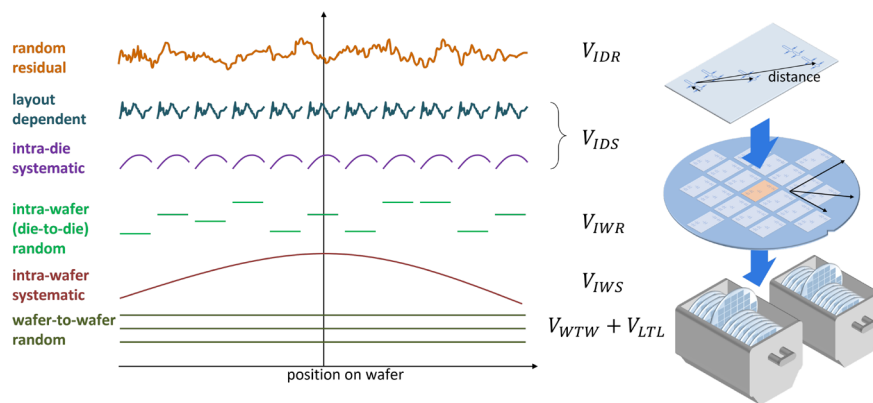


Figure 3. Decomposing spatial variability into die-level and wafer-level contributions. Within a lot, the wafer-to-wafer variations can be temporal and spatial. From lot to lot, the variations are mostly temporal.

width variation over a wafer is radially symmetrical and can be described by a bivariate polynomial. But this is an oversimplification, and as the real systematic variation cannot be fully described by this polynomial, the residuals will be misclassified as a contribution to the residual random variations.

2.2. Physical Origins of Spatial Variations. In this paper, we look in particular at spatial variations induced by the fabrication process. Silicon photonic chips are made using an elaborate sequence of deposition, lithography, etch, polishing, implantation, and annealing steps,⁸ and each of these contributes its own systematic and random variations to the fabricated chip. It is not straightforward to disentangle those contributions, so a model for variability needs to aggregate multiple of these effects.

In most semiconductor fabs, the chips are processed in batches (also called *lots*) containing up to 25 wafers. Some process steps are executed for all wafers in a lot simultaneously, while some steps are sequential. For example, a *chemical vapor deposition* (CVD) process heats multiple wafers together in a furnace with a reactive gas to deposit a thin layer on the surface of wafers.³⁵ But in processes such as lithographic patterning and *reactive ion etching* (RIE), each wafer is processed individually. For some steps, such as lithography, each die is individually illuminated in a stepper or scanner. These different processing steps induce a number of specific spatial patterns, at the level of the lot, the wafer, and the die. For this, we have revised the model for variability from ref 27 and ref 36, illustrated in Figure 3.

Lot-to-lot variations are mostly of a temporal nature, as they are processed sequentially. Tool drift, aging of chemicals (e.g., photoresist, developer), or a change of wafer supplier are typical contributions. Wafer-to-wafer variations can be attributed to both temporal effects and spatial effects, because some fabrication steps process wafers sequentially (lithography, etching) and some concurrently (annealing in an oven). These variations are induced for instance by tool priming, or the nonuniformity in a processing chamber. In theory, determining the systematic signature of lot-to-lot and wafer-to-wafer variations requires time-series models³⁵ and long-term monitoring of many wafers. For the discussion in this paper, we will not look deeper into the terms describing the *wafer-to-wafer* (V_{WTW}) and *lot-to-lot* (V_{LTL}) variations, and describe them as normally distributed stochastic variations.

Instead, we focus on the detailed contributions to variations within a wafer. Within a wafer, nonuniformity can be separated

into wafer-level effects and die-level effects. At the wafer level, we experience nonuniformity in layer thicknesses, photoresist spinning effects, and plasma distributions which in turn induce a change in etch rate or deposition rate. These effects are generally slowly varying over the wafer, and can be described by the *intra-wafer systematic* term V_{IWS} . On top of that, there is a periodic pattern of identical chips (dies) over the wafer. In large-scale production processes, these are defined using a stepper or scanner which projects the same photomask over the wafer multiple times. Fluctuations in exposure dose and imaging focus can induce a die-to-die *intra-wafer random* variation V_{IWR} .

Within each die, low-frequency variations in layer thickness, local pattern density, and errors in the photomask will lead to a systematic pattern V_{IDS} that repeats on each die. On top of that, intrinsic randomness (e.g., in layer thickness or waveguide sidewalls) will add a device-to-device *intradie random* variation V_{IDR} . We will now discuss each of these contributions in more detail.

2.3. Intra-wafer Systematic (IWS) Variations. Across a wafer, many steps contribute to the variation in the circuits. Interestingly, many of these nonuniformities follow a radial pattern by the nature of the process that imprinted them. Many processes have either a “center-fed” or “edge-fed” mechanism which imposes different boundary conditions near the edge. For example, the *post-exposure bake* (PEB) after lithography smoothens the standing wave-induced roughness on the sidewall. But the PEB temperature is often slightly higher in the middle of the wafer.³⁷ Similarly, chamber wall conditions affect the plasma during RIE etching, which causes etch rate nonuniformity. The radial plasma distribution is also affected by the nonuniform field profile of the magnetic RF fields that heat the plasma. During many processes, the wafers are rotated, which averages out azimuthal nonuniformity, but maintains the radial components.³⁵ In addition to radial process patterns, nonuniformities such as thermal gradients can result in linear or polynomial spatial variation.^{38,39} These variations across the wafer produce a slowly varying systematic intra-wafer variation V_{IWS} and combine both radial and polynomial contributions, that can be described as a bivariate polynomial of order n :

$$V_{IWS} = \sum_{i,j=0,1,\dots,n} p_{ij} \cdot x_w^i \cdot y_w^j \quad (1)$$

with x_w and y_w the location on the wafer, and p_{ij} the polynomial coefficients. For this study, we constrained ourselves to second-order polynomials ($n = 2$). Note that the observed radial

patterns are not unique to the processing of photonic circuits. As these effects affect the fabrication of electronic circuits just as well, each new generation of processing tools includes more refined controls to fine-tune the uniformity over the wafer. As many silicon photonic circuits are fabricated in slightly older fabs, often using 200 mm wafers, these effects can still be significant, even if they have been engineered away in the latest generation of tools.

2.4. Intra-wafer Random (IWR) Variation. The patterns on the wafers consist of multiple, nominally identical dies. These are images in a lithographic step-and-repeat process. This can induce die-to-die variations because of fluctuations in exposure dose, focus, and alignment between different lithographic layers. We define the *intra-wafer random* (IWR) variation as the random variation between dies. We model IWR as a normally distributed random variable:

$$V_{\text{IWR}} = \mathcal{N}(0, \sigma_{\text{IWR}}^2) \quad (2)$$

2.5. Intradie Systematic (IDS) Variation. Historically, variability from die-to-die and wafer-to-wafer were considered the dominant sources of variability for integrated circuits.⁴⁰ As an increasing level of integration increases the number of components, the die size has grown gradually, which reduces the controllability of the fabrication process at the die level and makes the intradie variability more significant.

A significant component of the intradie variation is systematic. Three main contributions can be considered. Nonuniformities in the optical lithography system can induce a slow variation over the die, which can also have a directional component when a scanner is used. The second source of systematic variations is the photomask. This is usually fabricated using an e-beam or laser writing technique, and this can induce variability as well. A third major origin of systematic intradie variations relates to the chip layout, and in particular the distribution of the pattern density. Plasma etch rates and polishing rate during CMP are both very dependent on the locally exposed area, and this then results in a local variation of line width or layer thickness.⁴¹ To alleviate this, fabs add dummy fillers in unpatterned spaces to reduce the nonuniformity in pattern density and control better the etch or polishing rate. We have discussed such correlation in more detail in ref 29. In theory, if we can trace back the IDS variation to all of its physical origins and study how they are correlated, we can derive the IDS as a function of its position $f_i(x_{\text{die}}, y_{\text{die}})$. In practice, due to the limited access and knowledge of the process, and the limited resolution with which we can observe variations, we can only identify partial correlations with causes such as the pattern density. We could treat the part that cannot be explained as a correlated random variation with a spatial correlation length. Then, we imitate the variation by a coherent noise map with a right correlation length L , such as Perlin or Simplex noise.⁴² The IDS variation can be expressed as

$$V_{\text{IDS}}(x_{\text{die}}, y_{\text{die}}) = \sum_i p_i f_i(x_{\text{die}}, y_{\text{die}}) + \text{Noise}(L) \quad (3)$$

where $(x_{\text{die}}, y_{\text{die}})$ are coordinates of the location on the die; p_i maps the impact of the parameter $f_i(x_{\text{die}}, y_{\text{die}})$, such as pattern density over the die, to the IDS variation $V_{\text{IDS}}(x_{\text{die}}, y_{\text{die}})$.

2.6. Intradie Random (IDR) Variation. The *intradie random* (IDR) variation captures the residual device-to-device disparity or local mismatch. IDR variation includes intrinsic variability like atomic oxide thickness variation or *line-edge roughness* (LER) due to photoresist granularity. The IDR variation dominates at the submicron level, while the size of a

photonic device is usually beyond several tens of micrometers. The distance between devices is far beyond the correlation length of the discrete causes of the IDR variation. Therefore, we can model IDR variation as a normal distribution independent of the other variation components:

$$V_{\text{IDR}} = \mathcal{N}(0, \sigma_{\text{IDR}}^2) \quad (4)$$

2.7. Separating Spatial Variation Components. With the above hierarchical variability model, we need to separate the different contributions of variation from the measured or extracted data on a fabricated wafer. There are a few methods in the literature to separate similar variations for CMOS integrated circuits. Ref 39 uses filtering, spline fitting, and regression-based approaches to extract wafer-scale effect. Then, the die-scale effects are separated by a spatial Fourier transform method. This method requires a good choice of parameters in the separation procedure, which is neither easy nor intuitive and sometimes subjective. Also, the extracted systematic variation is an interpolated map which cannot be described by an analytic expression nor can it be related to process parameters immediately, making it difficult to apply the extracted model for yield prediction. Ref 43 found that, in CMOS fabrication, systematic variations on intra-wafer and intradie levels can both be approximately described by a paraboloid function, i.e., a second-order bivariate polynomial. From our observation of the measured photonics wafer, we find that the IWS variation indeed corresponds roughly to such a paraboloid profile, but that it might be more accurate to apply a higher-order radial polynomial model. However, for the IDS variation, we do not really see such a pattern in the variation: of course, it is still possible to fit a second-order polynomial to describe these variations, but it will result in large residuals which are not truly random. Instead, there is a strong contribution of variations that correlate strongly with the chip layout, or rather the local pattern density. Therefore, it makes sense to split off layout-dependent variations from the other IDS contributions.

The separation of contributions is depicted in Figure 4. Essentially, we first separate the systematic patterns (V_{IWS} and

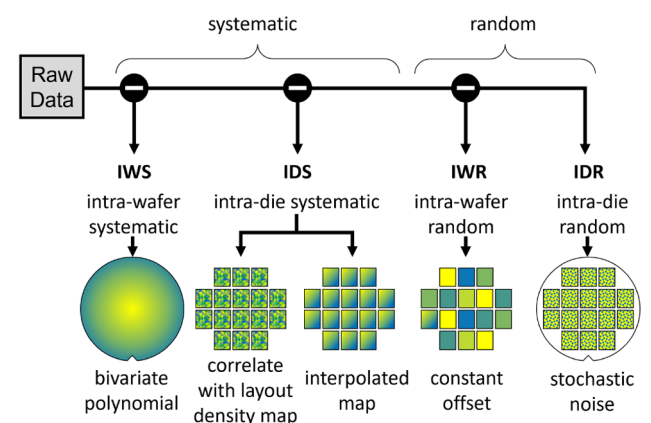


Figure 4. Workflow to separate variation on different spatial levels.

V_{IDS}) using a fitting process. Then we isolate the die-to-die random variations (V_{IWR}). Finally, the remaining term describes the residual random variations V_{IDR} . In the next section, we go step by step through this extraction based on a real set of data characterized on a wafer.

3. EXPERIMENTAL EXAMPLE ON ACTUAL WAFER DATA

A detailed model is only useful if it is populated with realistic parameters. For this, we performed an model extraction on a wafer fabricated in IMEC's 200 mm iSiPP50G technology platform,⁸ which uses 193 nm UV lithography. We focus primarily on the passive waveguide devices, and in particular on the variation in the width w and thickness t of fully etched silicon strip waveguides with design dimensions of $w = 450$ nm in an unprocessed SOI layer of $t = 220$ nm. After fabrication, the expected nominal dimensions are 470 nm \times 214 nm, taking into account lithographic biasing and material removal during processing. Note that the numbers used for this study relate to an older flavor of the iSiPP50G process and are no longer fully representative of the current state of the platform.

3.1. Measurements and Raw Data Extraction. As input data for our variability model, we need to characterize the waveguide width w and thickness t over the wafer, and perform this with a precision of 1 nm or better, given the sensitivity of the waveguides to small variations. Ideally, these values are extracted in situ, i.e., in the actual waveguides or patterns with similar geometric features. Direct measurements with scanning electron microscopy (SEM) do not have sufficient precision, while techniques such as scatterometry or ellipsometry require test sites with geometries that deviate in shape or pattern density from actual waveguides.

However, the sensitivity of the waveguides allows us to extract the variability through optical transmission measurement. In ref 44, and more recently in ref 45, we have developed an automated way for wafer-scale measurement of optical parameters such as effective index n_{eff} and group index n_g of waveguides. These optical parameters can then be mapped onto the variation of the waveguide line width w and thickness t .⁴⁴

The wafer maps we analyze in this paper are measured by a compact interferometric circuit described in ref 45, which is replicated many times over each die on the wafer. This allows us to extract a detailed map of subnanometer variations, averaged over the waveguide length within the circuit. The circuit and its distribution over the die and wafer are shown in Figure 5. It is a two-stage folder Mach–Zehnder interferometer with two inputs and two outputs connected to fiber grating couplers.

The 200 mm wafer consists of 52 complete dies of 21.84 mm \times 21.84 mm. We scattered 117 duplicates of this monitor circuit on a 5 mm \times 10 mm block within each die. As this fabrication is part of a multiproject wafer (MPW) run, we do not have access to the full die space and we can monitor about 10% of the total die space. Within our design space, we distribute some circuits in densely packed clusters (80 μm horizontally and 400 μm vertically), and then scatter the remaining circuits among the payload designs on the chip. Overall, this gives us more than 6000 instances of the circuit over the wafer.

We measure the wafer in a temperature-controlled cleanroom environment using an automated optical measurement setup. Light from a tunable laser is injected into the chip through the grating couplers and the output light is collected in a fiber and coupler to an optical power meter. Before and after the wafer measurement, we calibrate the tunable laser with a CO₂ and NH₃ gas cell. We also perform a stability test to rule out a drift in the laser.

The optical transmission of the circuits is measured from 1500 to 1600 nm, with a 20 pm resolution. As we measure the transmission between one input and two output ports, the entire

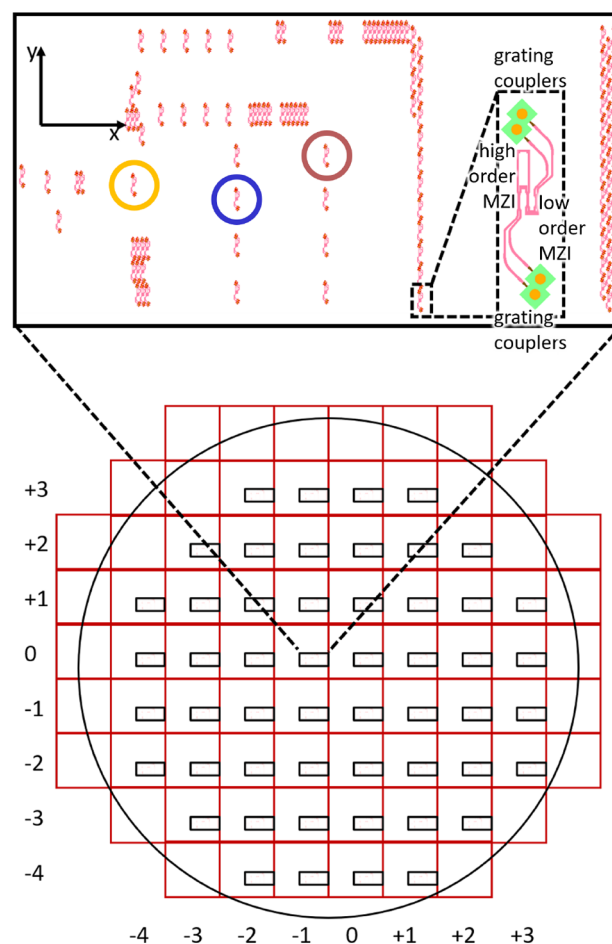


Figure 5. Location of the monitoring circuits on the die and the wafer. Inset is the layout of the cascaded MZI monitoring circuit.⁴⁵ Three locations are indicated, which are used in Figure 9 and Figure 10.

wafer measurement consists of 12168 measurements on the 6084 circuit samples (two measurements per test circuit, collecting the light from the two output ports). Of these, 5841 measurement pairs were valid (94%). Reasons for invalidation of measurements include the lack of one of the two transmission channels, saturation of the photodetector, or misalignment of the fiber.

The recorded spectra are then analyzed according to the procedure explained in ref 45 to extract n_{eff} and n_g , and then mapped onto the waveguide width w and thickness t according to the mapping strategy in ref 44. Among the valid samples, the average fitting error is 0.15 nm for width w and 0.08 nm for thickness t .

The derived wafer maps for w and t , with interpolations, are shown in Figure 6a and Figure 7a, respectively. We observe that the waveguide tends to be wider near the center of the wafer, narrowing down toward the perimeter, following a dome-like profile. The average width is 464.7 nm, close to the expected value of 470.0 nm, with a standard deviation of 4.6 nm. The maximum measured value on the wafer is 476.0 nm, while the minimum is 450.8 nm.

The thickness of the wafer varies like a slope from the southwest toward the northeast of the wafer. Near the edge, the thickness changes more abruptly, which could be attributed to the CMP process used for planarization. The average thickness is 210.3 nm where the target value is 214.0 nm, with a standard

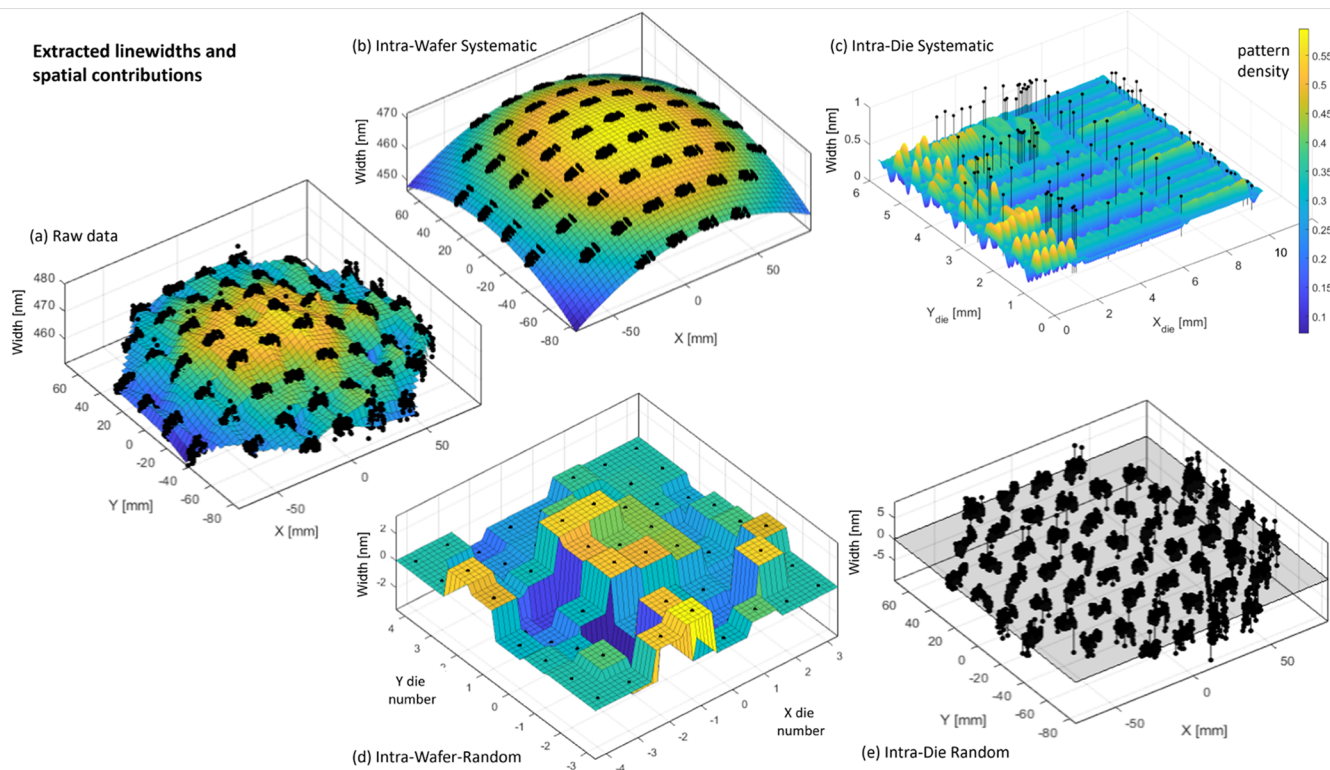


Figure 6. Maps of the extracted line width. (a) Raw data, (b) intrawafer systematic (IWS) variation, (c) intradie systematic (IDS) variation, (d) intrawafer random (IWR) variation, and (e) residual intradie random (IDR) variation.

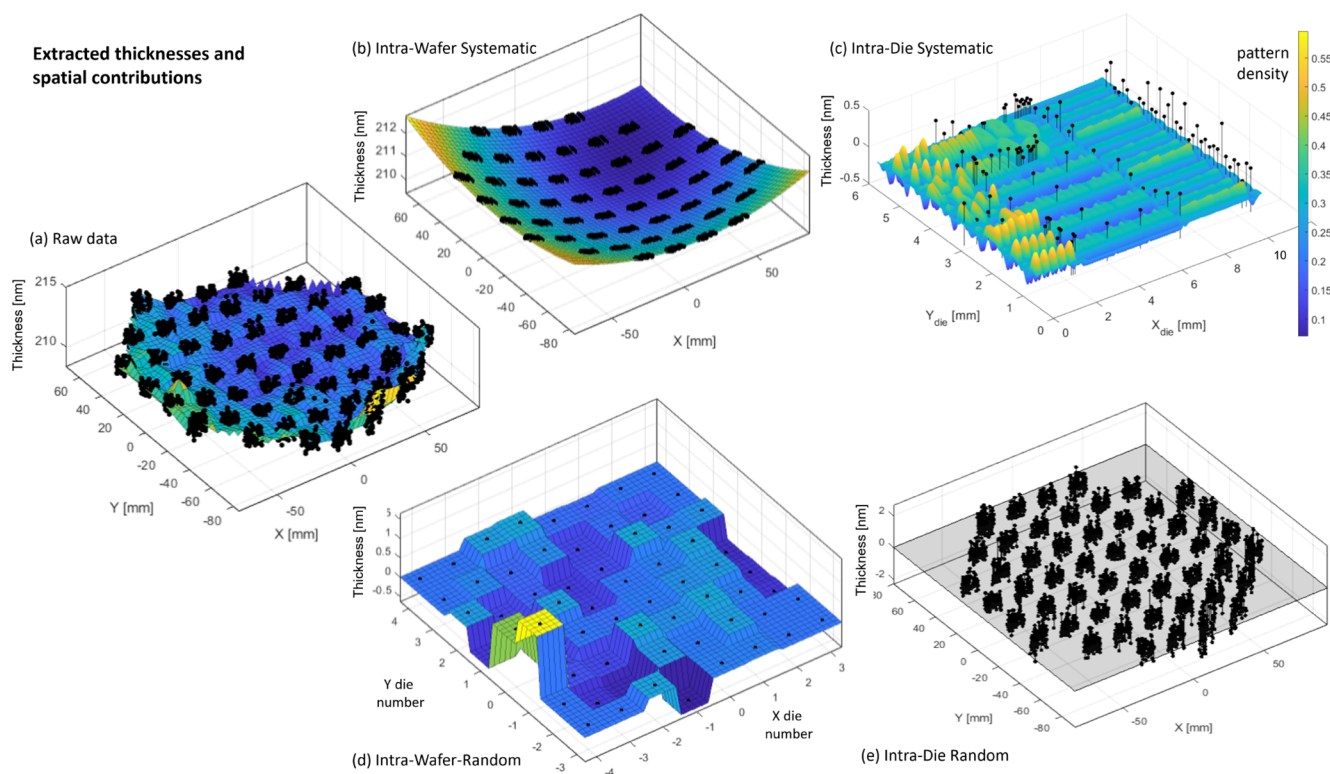


Figure 7. Maps of the extracted thickness. (a) Raw data, (b) intrawafer systematic (IWS) variation, (c) intradie systematic (IDS) variation, (d) intrawafer random (IWR) variation, and (e) residual intradie random (IDR) variation.

deviation of 0.8 nm. The maximum value on the wafer is 214.3 nm, while the minimum is 208.4 nm.

Figure 8 displays the histogram of both w and t , and Table 2 lists the key statistical properties. Obviously, both distributions are not simply normal distributions. This can already be

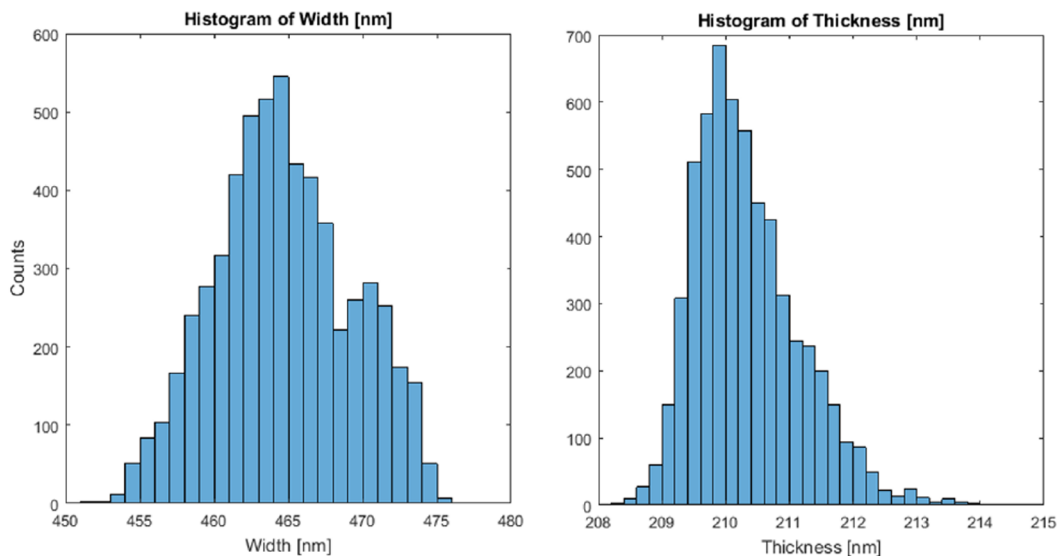


Figure 8. Histogram of measured width and thickness over the wafer.

Table 2. Statistics of Measured Width and Thickness

	width	thickness
mean [nm]	464.7	210.3
standard deviation [nm]	4.6	0.8
max [nm]	476.0	214.3
min [nm]	450.8	208.4
max–min [nm]	25.2	5.9

explained by the fact that there are many more dies on the perimeter than in the center of the wafer.

We now process these width and thickness maps to separate the various systematic and random contributions.

3.2. Intra-wafer Systematic Variation. As in the workflow of Figure 4, we start with identifying the intra-wafer systematic variation. There are two ways to derive this IWS variation: either we fit all the data on the wafer with the IWS model, or we fit a simple wafer map for each unique sample on a die, and then get the total IWS wafer map by averaging these individual wafer maps. The first approach should work even in the presence of short-range variations in the raw data. A model with low spatial frequencies (like the polynomial model proposed in ref 1) should filter out high-frequency variations and will not induce a significant estimation error. However, we prefer the second method to collect the IWS map by averaging simple wafer maps: In the case where our sampling points are not uniformly distributed over the wafer, this method will not overemphasize these locations to obtain the IWS map.

Figure 9 illustrates how, for each of the 117 copies of the test circuit in the design, we fit a polynomial function to the extracted width/thickness values in each die on the wafer. We then average out these 117 fitted polynomial curves to obtain the V_{IWS} wafer map. This averaging is also quite simple to implement: in the case of bivariate polynomial maps, it amounts to averaging the corresponding coefficients of the individual polynomials.

The resulting IWS variability maps for waveguide width and thickness are shown in Figure 6b and Figure 7b, respectively. We see that both these effects already encompass the main contributions to the width and thickness variation. The range of IWS variation among the measured samples is 16.46 nm, while for the thickness, the IWS variation spans 2.54 nm among

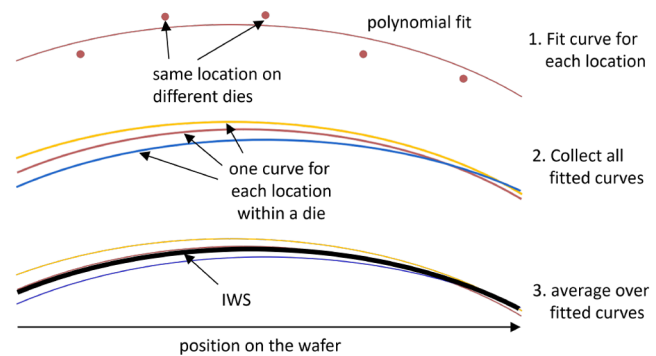


Figure 9. Fitting IWS variations. The extracted line width or thickness value is measured over the wafer on the same location in each die (e.g., the red, yellow and blue points indicated in Figure 5). To each of these 117 sets, a bivariate (x, y) polynomial is fitted. Then, the IWS map is obtained by averaging over these polynomials.

the measured samples. In terms of relative contribution to the total variation (which is measured by the variance), we find that IWS contributed for 76.1% and 64.4% for w and t respectively.

The width variation (Figure 6b) shows a very distinct radial pattern with its peak near the center of the wafer. Eq 1 can describe such a dome-like shape. As already mentioned, the radial symmetry is most likely attributed to the plasma profile (and therefore etch chemistry) over the wafer during the RIE process.

As we could already deduce from the raw data, the IWS variation of the thickness (Figure 7b) behaves more like a slanted plane leaning from the southwest toward the northeast of the wafer. The physical origin of this sloped profile across the wafer is not entirely clear and needs further study.

While this analysis for the IWS variations is based on a single wafer, an inspection of the data collected for ref 44 and ref 46 show very similar patterns.

One open question here is whether our model using a bivariate second-order polynomial is sufficient to describe the IWS variations. Higher-order polynomials might provide a better fit, especially toward the edges of the wafer.

3.3. Intradie Systematic Variation. The second contribution we separate out is the intradie systematic (IDS) variation, which also encapsulates effects induced by the circuit layout on the photomask. For that, we subtract the fitted IWS from the raw data. The intradie systematic (IDS) variation describes the variations that are repeated in each die. Figure 10 illustrates how

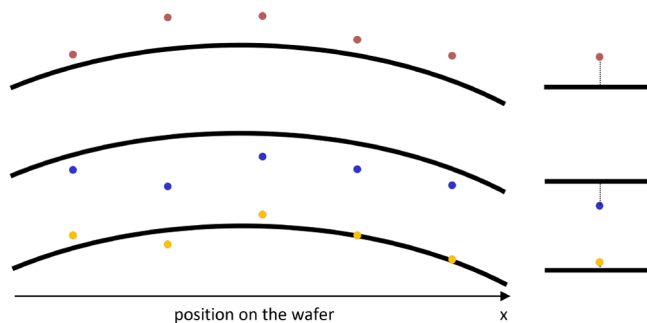


Figure 10. Left: Black solid curves indicate IWS variation. Points with the same color are parameters measured on identical locations on different dies. Right: The average offset between the measured parameter and IWS variation is the IDS variation.

this is extracted: For each of the 117 device locations, we calculate, over the 52 dies on the wafer, the average deviation from the IWS map that we fitted earlier.

The extracted IDS variation for the waveguide width is shown in Figure 6c. We observe that the IDS width variations are very locally correlated. The maximum IDS width variation is 1.52 nm, and the minimum is -2.52 nm, giving a range of 4.04 nm difference and $(0.88 \text{ nm})^2$ variance, corresponding to 3.6% of total width variance. The effect of layout and local pattern density plays an important role here: We observe that the waveguide width is systematically larger on the West side of the die, which contains designs with arrayed waveguide gratings (AWG) that contain large areas that are almost completely etched (the waveguide arrays) and areas that are not etched at all (the star couplers).⁷ In ref 29 we used the same wafer to perform a detailed analysis of the pattern density effects on the waveguide line width. We have observed a moderately strong correlation between the IDS line width variation and the pattern density over a radius of 50–300 μm , with a peak around 65 μm . Figure 11 shows the die layout (before and after tiling) and the extracted IDS contributions at the individual extraction points on the die. While the tiling maintains a uniform density over most of the chip, the main points of deviation can be observed near the right, where the density of test circuits is very high, and near the left, where we placed several arrayed waveguide gratings (AWG) with unpatterned star couplers that significantly disrupt the average pattern density.

The IDS thickness variation, plotted in Figure 7c, has a maximum of 0.40 nm and a minimum of -0.51 nm, resulting in a range of 0.91 nm. The corresponding variance of $(0.14 \text{ nm})^2$ contributes 2.5% to the total thickness variance. As can be seen more closely in Figure 11, there is no significant correlation between the pattern density and the extracted IDS device thickness. We did not observe an association between pattern density and IDS thickness variation. Also, we did not observe a correlation between the IDS width and IDS thickness variation.

3.4. Intra-wafer Random Variation. After we remove both systematic components from the data, the residual contains the stochastic contributions at both wafer (IWR) and die (IDR)

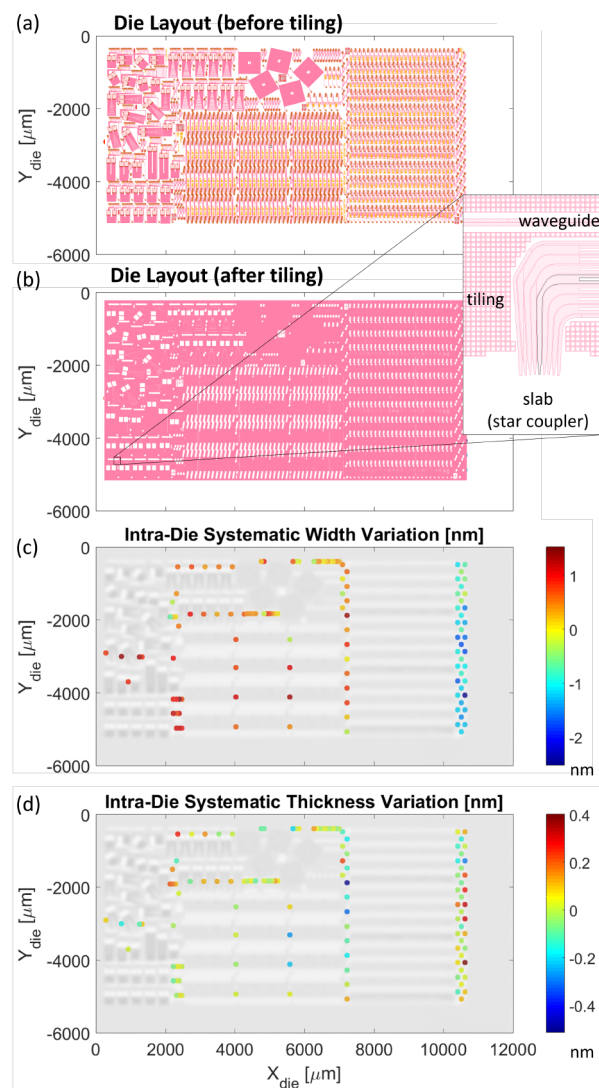


Figure 11. Extracted intradie systematic (IDS) width and thickness variations correlated to the device layout on the die. (a) The original circuit layout, (b) the circuit layout with dummy tiling to homogenize pattern density, (c) intradie systematic width variation, and (d) intradie systematic thickness variation.

levels. We model both models with a normal distribution over the wafer and die, respectively. By definition, both contributions average out as zero over all the dies on the wafer.

The intra-wafer random variation can also be called the die-to-die random variation. The IWR width variation is dominated by effects that differ from die to die, which can be caused by lithographic focus or exposure fluctuations when using a stepper or scanner. As can be seen from the plot in Figure 6d, the IWR width variation shows no spatial correlation. It ranges between -3.40 nm and $+3.83$ nm, with a variance of $(1.54 \text{ nm})^2$ contributing 11.3% to the total width variation.

The IWR thickness variation (Figure 7d) shows an interesting deviation from the systematic flat plane on the wafer level. We observe that prominent variations are present near the edge of the wafer. This offset might be caused by the faster polishing rate near the periphery of the wafer. However, the strong deviation near the edge is an indication that our IWS model might not be sufficiently detailed. Except for the dies near the edge, other dies




Figure 12. Histogram of the residual IDR width and thickness variation.

on the wafer have minimal IWR variation with similar values compared with the fitting error (0.08 nm).

3.5. Intradie Random Variation. The remaining residual is the local random variation between the individual circuits or devices. Indeed, a correlation tests on the residuals (plotted in Figure 6e and Figure 7e) for the width and thickness variation shows no spatial correlation. The histograms plotted in Figure 12 show that these residuals can be modeled well with a normal distribution. The IDR width variation has a standard deviation of 1.38 nm, while the IDR thickness variation has a standard deviation of 0.37 nm. The corresponding variances contribute 8.9% and 20.8% to the total variance of w and t , respectively.

3.6. Discussion. The separation of wafer-scale variations into spatial contributions helps us to understand the location dependency of the process variation. Clearly, the effects of process variations are not simply random and cannot be simply described as a normal distribution, even with a certain local correlation length. The location dependence clearly has different length scales, both for systematic and stochastic contributions.

The absolute and relative magnitudes of the different contributions are listed in Table 3. As we separate the different

Table 3. Comparison of Different Spatial Levels of Variation

	range [nm]		st. dev. σ [nm]		percentage of total variance σ^2 [%]			
	w	t	w	t	w		t	
IWS	16.46	2.54	4.02	0.65	76.1	79.8	64.4	66.8
IDS	4.04	0.91	0.88	0.14	3.6		2.5	
IWR	6.85	2.07	1.54	0.28	11.3	20.2	12.4	33.2
IDR	18.46	4.50	1.37	0.37	8.9		20.8	
total	25.2	5.9	4.59	0.82				

types of variations as independent components of the total variation, we can (to a good approximation) determine their relative importance by the second-order moment of their distribution, which we approximate by the variance over the population of sample points on the wafer. For the width variations, the systematic effects (IWS and IDS) have a total variance of $(4.02 \text{ nm})^2 + (0.88 \text{ nm})^2 = (4.12 \text{ nm})^2$ and contribute to $\approx 80\%$ of the total variance over the wafer. For the

thickness, the systematic variance accounts for $(0.65 \text{ nm})^2 + (0.14 \text{ nm})^2 = (0.67 \text{ nm})^2$, which is equivalent to around two-thirds of the total measured variance. The remaining one-third of random variations in the thickness is quite large, but as already discussed, a better model of the IWS thickness variations might be able to capture the strong deviations at the perimeter of the wafer.

The large systematic contributions in width and thickness variations emphasize the importance of location-dependent variability analysis. The analysis also provides insight into how we might improve or compensate for these effects in future designs (for layout-dependent variations) and process developments.

Spatially the largest contributions to the range of variation in both width and thickness can be found on the wafer level. To a small circuit on a die, these wafer-level variations could be modeled like a common-mode variation that impacts all elements of the entire circuit in the same way, with maybe an additional linear gradient within the circuit footprint. Shorter-range die-scale variations impact the circuit components differently. For example, different arms in an interferometer could feel a different local variation of line width and thickness.

As mentioned earlier, a 1 nm change in line width w can lead to approximately 1 nm spectrum shift, and this sensitivity is twice as strong for the thickness t . The total width and thickness variation over this wafer can therefore induce more than 30 nm of variation in wavelength response over the wafer.

4. PREDICTION WITH A SPATIAL VARIABILITY MODEL

Based on the extracted data, we can now construct a spatial model to be used in the layout-aware variability simulation workflow shown in Figure 2. We use the Caphe circuit simulator by Luceda Photonics, enhanced with *variability extensions* that project wafer-level variability maps of line width and thickness on the parameters of the layout-aware circuit models.³⁰

As an example, and as a test of the validity of our model, we simulate the same two-stage Mach–Zehnder interferometer as used for the extraction of n_{eff} and n_g . The circuit, shown in Figure 13, consists of waveguides and directional couplers. Both are modeled with a wavelength-dependent model. The waveguide is described by a linear dispersion of the propagation constant

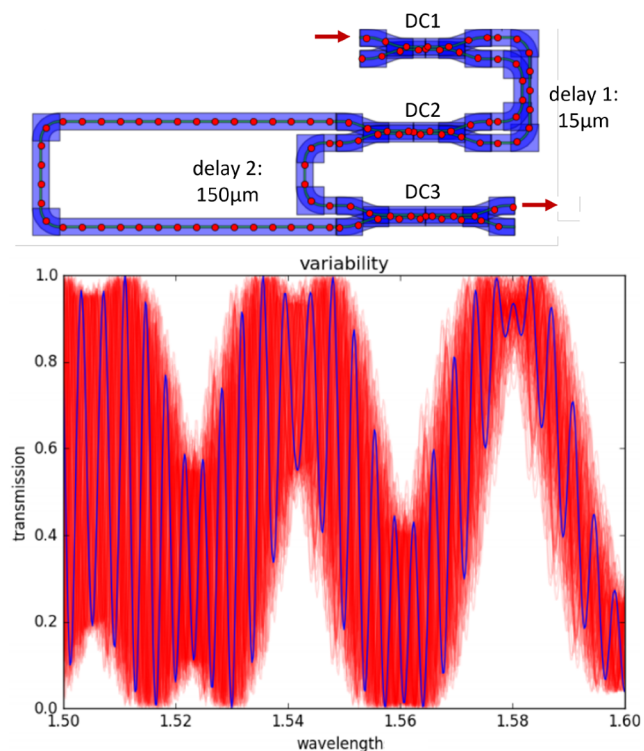


Figure 13. Two-stage MZI circuit from Figure 5 used for variability simulation. The circuit model parameters of the waveguides and directional couplers are locally adjusted based on the line width and thickness at the sampling points (red dots). Bottom: Transmission simulation of 117 copies within a single die.

($n_{\text{eff}}(\lambda_0)$ and n_g), while the directional coupler is described by a combination of a constant coupling $\kappa_0(\lambda)$ and a length-proportional coupling $\kappa'(\lambda) \cdot L$. The wavelength dependence of κ_0 and κ' are described by a second-order Taylor expansion, resulting in six model parameters.

The mapping of the waveguide width w and thickness t to n_{eff} and n_g of the waveguides has been described in ref 44. The sensitivity of the directional coupler model parameters to the line width and thickness variations are calculated using a combination of eigenmode expansion and FDTD simulations.

4.1. Comparison of Wafer Model with Measurements.

To assess the validity of our workflow, we simulate this MZI circuit in the same locations on the virtual wafer, using the spatial model that we just extracted. We then compare the results of the model with the extracted values from the physical wafer. Figure 13 shows a Monte Carlo transmission simulation of the circuit on 117 locations with a single die of the wafer, using the width and thickness maps in Figure 15a,b.

From those simulated transmission curves we now extract first the n_{eff} and n_g of the waveguides by fitting the transmission calculated with circuit model to the measured transmission using a global multivariate fitting procedure as described in ref 45. These extracted values can then be mapped onto the line width w and thickness t of the waveguide using a homomorphic transformation that maps (n_{eff}, n_g) onto (w, t) . This mapping is precalculated using a finite-element electromagnetic mode solver. This procedure is described in ref 44.

We can now compare the values we calculate using our wafer model and circuit simulation from the values we originally extracted directly from measurements. We expect that they will be similar but not identical: In our measurements, we cannot

sample inside a single circuit, and the extracted values of w and t include the effects of local variability within the circuit. The wafer model we have constructed will also project interpolated local variability on the circuit simulation, but this will be different from the original wafer.

Figure 14 now shows the histograms of the difference between the originally extracted line width and thickness, and the values




Figure 14. Histograms of the error between the line width w and thickness t extracted from the simulations on the virtual wafer, and the values extracted from the actual measurements.

for line width and thickness we obtained from our circuit simulation with our wafer model. We see that these differences are subnanometer in magnitude, indicating that not only are the short-range variations within the circuit on the fabricated wafer tiny, but also they follow the interpolated model, and therefore variations on a short length scale have only a minor contribution.

4.2. Synthesized Wafer Models. We can also use the extracted spatial variation parameter to generate new wafer maps with similar statistics but different parameters. Since the spatial variability is additive, we can construct the different contributions independently, and then add up the variations in the end. All contributions can be pregenerated, except for the pattern-density contribution to the *intradie systematic* (IDS) variations.

The generated virtual wafer does not have to be the same as the fabricated wafer since every fabricated wafer is different. Nonetheless, the virtual wafer should mimic the real wafer in at least two ways:

- It should have similar statistics of global variables over the wafer as the real wafer.
- It should exhibit a similar spatial correlation of global variables as the real wafer.

The second criterion is important, as it ensures that yield estimation can capture the correlation between neighboring locations on a chip. A realistic virtual map of a wafer should be useful to help the designer generate process-tolerant circuits.

The *intra-wafer systematic* (IWS) contribution is again generated as a bivariate second-order polynomial, of which the coefficients are similar but not identical with the ones extracted from our original wafer. Collecting data from additional wafers will make it possible to get a more precise distribution of these polynomial coefficients.

The *intradie systematic* (IDS) describes short-range correlation and variation within a die, repeated for all dies over the wafer. The IDS model consists of two parts: a part correlated to the pattern density of the layout, and a systematic part that includes global trends within a die. This part is modeled as a polynomial fit.

The layout-dependent variation is based on the local pattern density. For this, the layout over the wafer should be known. In

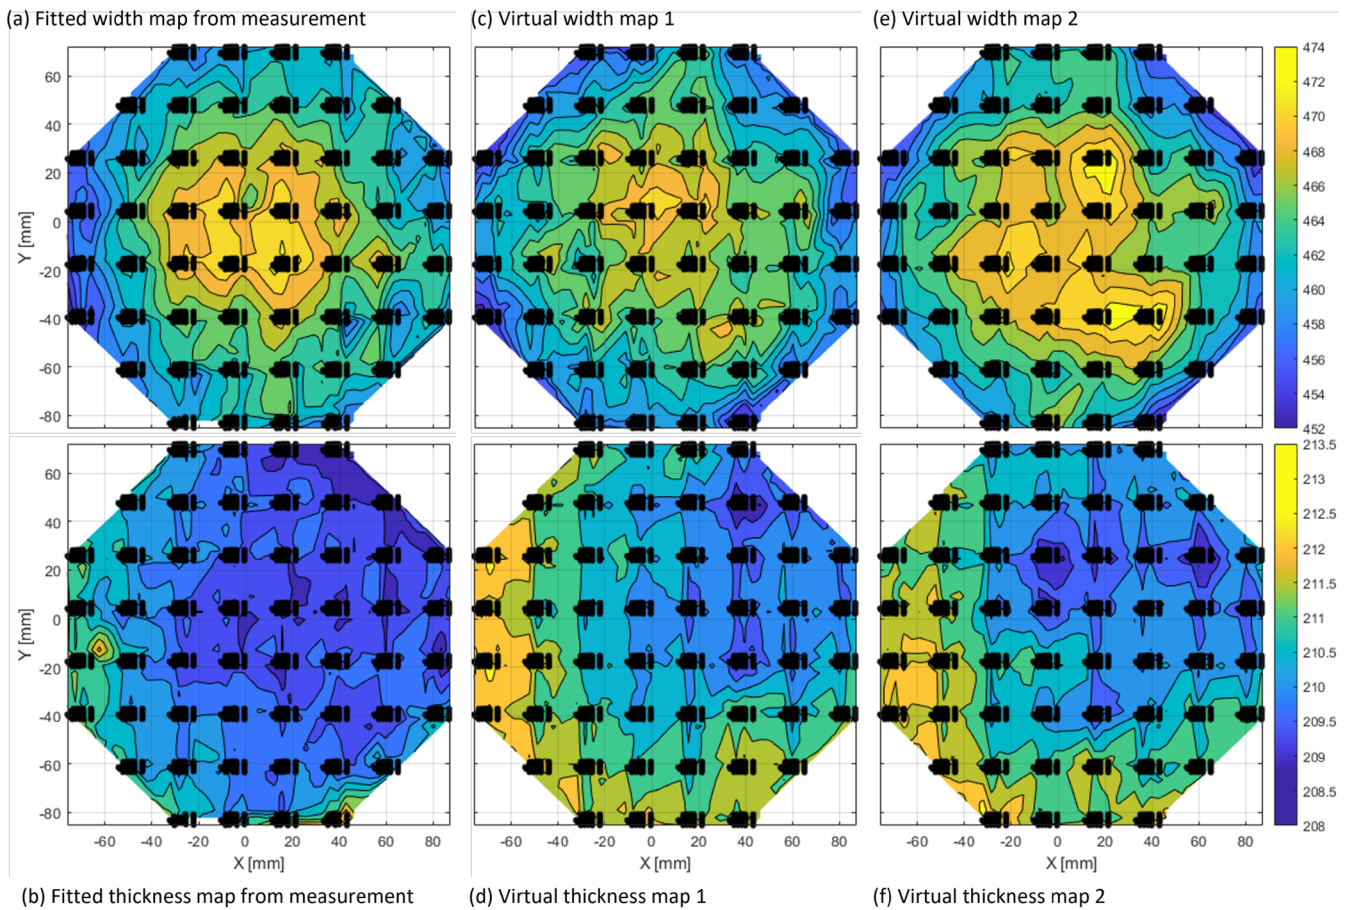


Figure 15. Virtual wafer maps models for width w and thickness t for use in layout-aware variability simulations. (a,b) Maps extracted from the measurements. (c–f) Maps synthesized using similar statistics as the extracted maps, combining the contributions on the different spatial levels.

our case, we already know the entire layout, as we base it on the previously fabricated circuits. To calculate the effect of local pattern density, we take the layout patterns used to define the waveguides, add the tiling patterns in the empty areas, and then convert it to a grayscale bitmap with a pixel size corresponding to 500 nm. On this bitmap, we perform a Gaussian blur with a radius of 65 μm , which corresponds to the peak correlation length of layout-dependent variations.²⁹ The Gaussian blur corresponds to a 2D convolution with a Gaussian kernel function. For the line width, the correlation with the pattern density is 62%, which means that the additional 38% contribution needs to be modeled by a random contribution, which is modeled using a coherent noise map (Simplex or Perlin noise⁴²) with a radius of 500 μm . For the thickness variations we observe no correlation (<25%) to the pattern density, and therefore the entire contribution is modeled by a similar random noise map. These random contributions will require some future refinement to determine the statistics, the length scale and the amplitude. It is also possible that there are additional layout-dependent effects originating from interactions between different process layers (full and partial etch, depositions, doping, ...).

Using these blurred bitmap images, we interpolate the locally averaged density for all the sample points in the circuit (red dots in Figure 2), and then calculate the deviation from the nominal value in each point by subtracting it from the reference density. The systematic part of the layout-dependent line width variation is proportional to this local density deviation: 3% in pattern density will induce ≈ 1 nm change in line width.

The *intra-wafer random* (IWR) variation explains the difference in die-to-die averages. The generation of intra-wafer random variation maps is quite straightforward: since no spatial correlation is observed in IWR, we assign each die a random value that follows a zero-mean Gaussian distribution. We set the standard deviation the same as what we extracted from the measured wafer. The dies are not correlated.

In contrast, the *intra-die random* (IDR) variation denotes device-to-device variations that are not location-dependent, but that are still spatially correlated, at least over a short-range: line width and thickness will vary continuously along a waveguide; there should be no discontinuities. As we observed that the device-to-device random variation seemed to be spatially uncorrelated, we assume the correlation length of such randomness is smaller than the device-to-device distance. Therefore, we use a spatial coherent noise function with a correlation length of 100 μm , which is smaller than the circuit footprint of our Mach–Zehnder test circuit.

The resulting maps for width and thickness variation are shown in Figure 15. Compared to the map extracted from the originally fabricated wafer, we see that there are significant differences, but there are also strong similarities: There is a clear radial pattern, and there is a similar die-to-die variation.

It is useful to look at one of the generated *intra-die systematic* (IDS) maps, shown in Figure 16. For the line width variation map, we see a clear layout-dependent effect, which is logical as we incorporated the actual layout into the generation of the map.

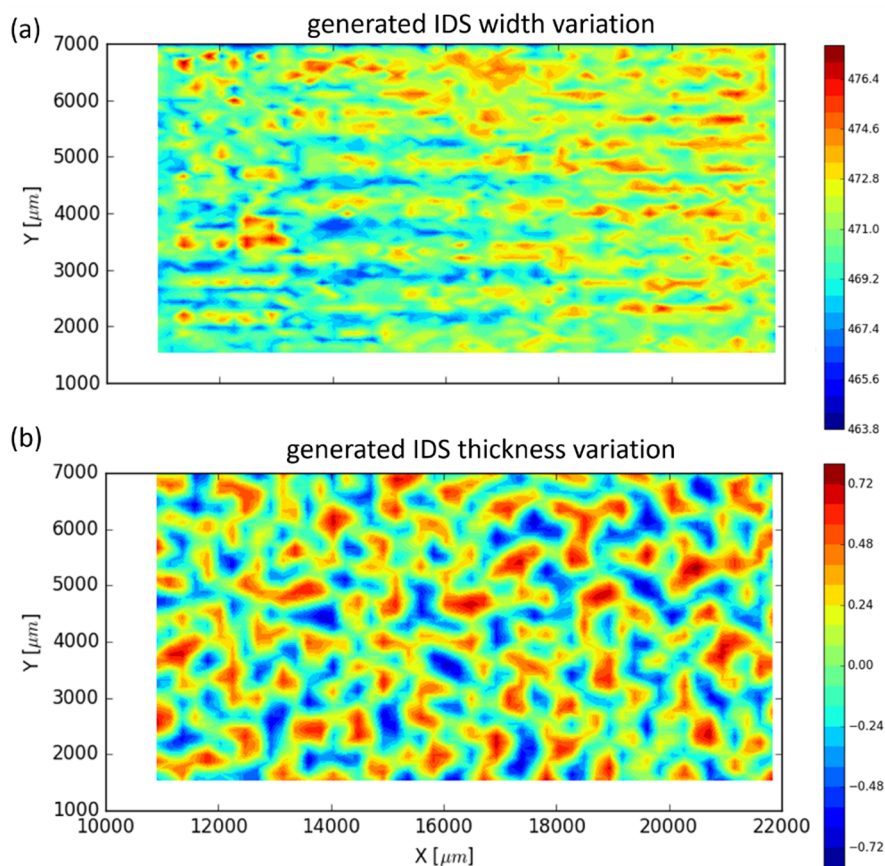


Figure 16. Generated virtual die maps for the IDS contribution on the width and thickness variation. These maps are generated for the same layout as used in the test chips (see Figure 11b) and contain both a layout-dependent part and a layout-independent systematic part. (a) Width variation for a 470 nm waveguide, and (b) thickness variation. Note that the thickness variation has little or no dependence on the layout.

For the thickness, the variation is purely random, but with a clear spatial correlation length.

To test the usefulness of these new synthetic wafer maps for the characterization of photonic circuits, we use the same technique as for the synthesized wafer map based on the original data. We sample the virtual wafers at the same locations where the test circuits are located on the fabricated wafer (Figure 5). Then we simulate the transmission of each circuit on the virtual wafers with different maps, and from the transmission curve, we extract first n_{eff} and n_g , and then the local w and t for each instance of the circuit.

The result is shown in Figure 17. The distributions of w and t on the virtual wafers are not identical, but similar to the ones from the fabricated wafer. The mean and standard deviation of these distributions is similar for both line width and thickness, as can be seen in Table 4.

5. DISCUSSION

Virtual wafer maps of line width and thickness variations make it possible to predict the performance of circuits without the need for resource-intensive electromagnetic device simulations. The only requirement is circuit models that also include the sensitivity of the model parameters to the variables that are measured over the wafer, such as line width and thickness, and a corresponding virtual wafer map of these variables.

However, as we made clear in this paper, the virtual wafer maps are only useful if they represent the fabrication process correctly. This means that they should be continuously

monitored and updated. It is therefore important to include the necessary characterization structures of all fabricated wafers that allow for a direct wafer-to-wafer and lot-to-lot comparison. As the collection of statistics improves, the wafer models can be refined.

In this work, we only looked at two key variables: line width and thickness of fully etched patterns in an SOI layer, and their variation over a wafer. These two variables are probably the most critical for most silicon photonics circuits, as they have a strong impact on the propagation constant of the waveguides, and therefore on the wavelength response of interferometers and resonators. But this does not mean that they are the only variables that matter: most silicon photonics platforms combine multiple etch depths in different process steps, and each of these will require a virtual wafer map. Also, processes for active devices such as modulators (e.g., dopant implantations) can be modeled with the same technique. In a first approximation, these different process modules can be treated independently, and the variability maps can just be added up. However, different process steps can likely interact, which means that the variability maps for line width, thickness, etch depth, dopant concentration, or other physical properties cannot be treated fully independently, but maps of later process steps could be correlated to the maps of steps earlier in the flow, which should be taken into account during the map generation process. Identification and characterization of these effects, so they can be incorporated in more advanced variability models, will require additional test structures, either based on in-line

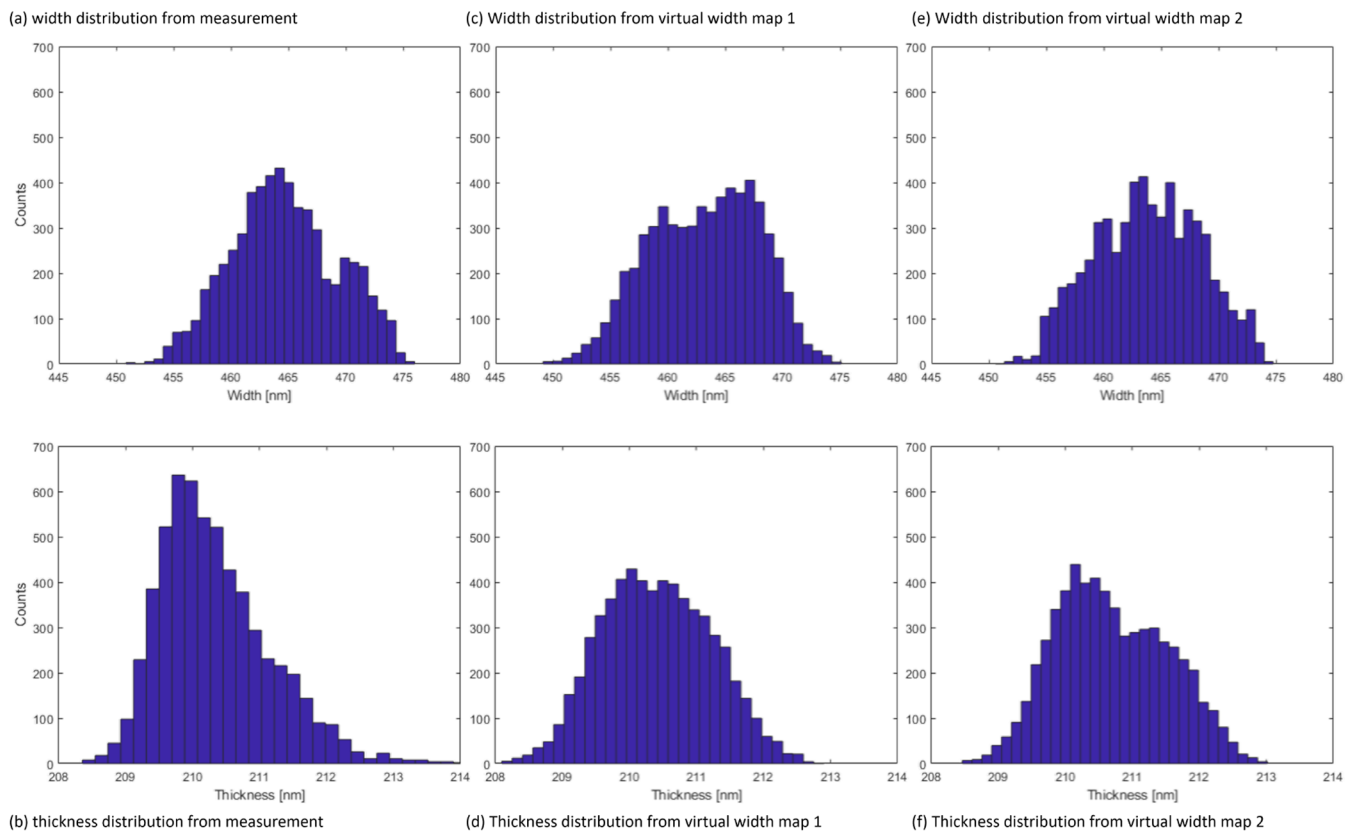


Figure 17. Histograms of the extracted line widths and thicknesses using layout-aware circuit simulation using the originally extracted wafer map, and the two synthesized wafer maps from Figure 15.

Table 4. Comparison of the Extracted Width and Thickness Deviations on the Fabricated and Generated (Virtual) Wafer Maps from Figure 15

	fabricated wafer	virtual wafer 1	virtual wafer 2
Width [nm]			
mean	464.67	465.55	463.18
standard dev.	4.59	4.44	4.71
Thickness [nm]			
mean	210.34	210.41	210.35
standard dev.	0.82	0.81	0.83

measurement techniques (e.g., SEM inspection, scatterometry, ellipsometry, ...) or dedicated electrical or optical test circuits.

As mentioned earlier and in ref 30, this technique requires two sets of information on top of the standard building blocks and their models provided in the process design kit (PDK). The first set comprises the sensitivities of the device model parameters to changes in the different variables (width, thickness, ...). The second set consists of representative wafer maps for these variables. The sensitivity of a geometry to a change in geometry can be determined by deliberately changing fabrication conditions or performing a sweep of the design parameters in the layout. However, in most cases, the sensitivity can be quite accurately estimated from simple electromagnetic simulations, as we mostly care about relative changes. Obtaining good wafer maps is somewhat more difficult. This can be done by characterizing test structures, like we have done here, and breaking down the spatial contributions to build a generative model for virtual wafer maps. Ideally, such a generative model for the key variables in the fabrication process should be

supplied by the foundry. However, this kind of information is often considered sensitive, as it could provide competitors with in-depth insight into the process. An intermediate solution would be to “anonymize” the variables, for instance through principal component analysis, which maps the variables with physical meaning (width and thickness) onto a new set of abstract variables. This would reduce the insight into the actual sources of variability, but still allow designers to optimize circuits for fabrication yield. The drawback of such remapping is that the sensitivity to the new variables can no longer be calculated by the designer, but should be provided by the foundry. In practice, this would limit design optimizations to circuits of standard PDK building blocks.

Note that our analysis assumes the use of optical deep-UV lithography techniques (usually at 248 or 193 nm wavelength), as these are the commonly used methods for volume fabrication. For research and prototyping purposes, electron-beam (e-beam) lithography is also commonly used. As this is a serial writing process, the statistics will be very different, and line width variations can be very dependent on the writing strategies, the order in which waveguide structures are written, and the proximity correction software used to prepare the patterns. Still, some statistics could be extracted for such processes, but some die-to-die correlations that we expect with optical lithography will no longer be present. Wafer-level variations related to layer thickness will be mostly unaffected by the lithography process. So even though the technique becomes less relevant for yield prediction when using e-beam lithography, it could still be used to assess some impact of variability in this process.

To actually calculate the yield, we use a simple, Monte Carlo method and run brute-force circuit simulations in which we apply the effect of these variability. This method is simple, but it does require a large number of simulations to get an accurate picture of the trends that affect success or failure. For pure yield estimation, which is a simple pass/fail criterion, a lower number of tests can be used: for an expected yield of 50%, only ≈ 34 samples are needed, and to verify an expected yield of 99%, approximately 1000 samples are needed for a 95% confidence level. The Monte Carlo method is very well suited to this problem: because the wafer maps combine many contributions with different statistical properties, it is not always compatible with faster stochastic methods that assume certain properties from the variable distributions. Also, the circuit simulations are fast compared to full-scale electromagnetic simulations, and can easily be parallelized. For example, assume a circuit simulation of 5 min. This corresponds to a frequency simulation on 200 wavelength points of a circuit with ≈ 1000 building blocks, or a 100 000-step time-domain simulation of the same circuit, on a single core of a typical 2021 business laptop. If we would launch the 1000 simulations to test the 99% yield hypothesis on a 64-thread workstation or cluster, it would take $a \approx 80$ min to complete. For a yield of 85%, it would require only ≈ 120 samples, which can be completed in 10 min. When there is no up-front yield estimate, the termination criterion can be made adaptive based on the observed yield during the Monte Carlo simulation (e.g., requiring a minimum number of failed circuits as a stop criterion.)

6. CONCLUSION

In this paper, we discussed a technique to describe wafer-scale variability in photonic devices and then use this information to calculate its effect on circuits. The calculation can then be used to make their design more robust. The challenge is to predict the circuit behavior in the presence of variability during the design phase, and accurately capture the different contributions of variability. We proposed a hierarchical spatial variability model and applied this to analyze measured data on a fabricated 200 mm Silicon photonic wafer.

We focused on the most critical parameters for passive waveguide devices, namely the width and thickness of submicrometer silicon wire waveguides. After extracting these values using a compact interferometric circuit, we separated the variability on various spatial levels into systematic and random contributions, for which we constructed a model that could then be used to construct new “virtual” wafers.

We applied the model and the workflow to process the measurements on a 200 mm wafer fabricated by IMEC’s standard passive silicon photonics platform using 193 nm lithography. The result shows that the intrawafer systematic variation is the major source of variation for both line width and thickness. We observed that the width variation has a systematic dome-like profile across the wafer. Thickness nonuniformity across the wafer looks like a slanted plane with a few mismatches around the wafer edge. On the die level, we found that repeated systematic width patterns are closely related to the local pattern density. Our analysis showed that the intradie systematic width variations are affected by the pattern density within a radius of $\sim 200 \mu\text{m}$. Our findings help to identify the process variation and create new design rules to alleviate the impact of the nonuniformity.

AUTHOR INFORMATION

Corresponding Author

Wim Bogaerts – Photonics Research Group, Ghent University–IMEC, 9052 Ghent, Belgium; Center of Nano and Biophotonics, 9052 Ghent, Belgium; orcid.org/0000-0003-1112-8950; Phone: +32 (0)9 264 3324; Email: wim.bogaerts@ugent.be

Authors

Yufei Xing – Photonics Research Group, Ghent University–IMEC, 9052 Ghent, Belgium; Center of Nano and Biophotonics, 9052 Ghent, Belgium; Present Address: Shanghai Industrial μ Technology Research Institute, 235 Chengbei Road, 201800 Shanghai, China; orcid.org/0000-0001-8772-2038

Jiaxing Dong – Photonics Research Group, Ghent University–IMEC, 9052 Ghent, Belgium; Center of Nano and Biophotonics, 9052 Ghent, Belgium; Present Address: VPIphotonics, Carnotstrasse 6, 10587 Berlin, Germany.; Present Address:atlanTTic Research Center, University of Vigo, Campus de Vigo As Lagoas, Escola de Enxeñaría de Telecomunicación, Rúa Maxwell, s/n, 36310 Marcosende, Pontevedra, Spain.; orcid.org/0000-0001-8957-2467

Umar Khan – Photonics Research Group, Ghent University–IMEC, 9052 Ghent, Belgium; Center of Nano and Biophotonics, 9052 Ghent, Belgium; orcid.org/0000-0001-5760-7485

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acsp Photonics.2c01194>

Funding

This work was supported by the Flemish Research Foundation (FWO-Vlaanderen) through Grant No. G013815N, and by the Flemish Agency for Innovation (VLAIO) and Luceda Photonics through the MEPIC project.

Notes

The authors declare no competing financial interest.

REFERENCES

- (1) Chen, X.; Milosevic, M. M.; Stankovic, S.; Reynolds, S.; Bucio, T. D.; Li, K.; Thomson, D. J.; Gardes, F.; Reed, G. T. The emergence of silicon photonics as a flexible technology platform. *Proceedings of the IEEE* **2018**, *106*, 2101–2116.
- (2) Borel, P. I.; Harpoth, A.; Frandsen, L. H.; Kristensen, M.; Shi, P.; Jensen, J. S.; Sigmund, O. Topology optimization and fabrication of photonic crystal structures. *Opt. Express* **2004**, *12*, 1996–2001.
- (3) Molesky, S.; Lin, Z.; Piggott, A. Y.; Jin, W.; Vucković, J.; Rodriguez, A. W. Inverse design in nanophotonics. *Nat. Photonics* **2018**, *12*, 659–670.
- (4) Hughes, T. W.; Minkov, M.; Williamson, I. A. D.; Fan, S. Adjoint method and inverse design for nonlinear nanophotonic devices. *ACS Photonics* **2018**, *5*, 4781–4787.
- (5) Rehman, S. U.; Langelaar, M. System robust optimization of ring resonator-based optical filters. *Journal of Lightwave Technology* **2016**, *34*, 3653–3660.
- (6) Bogaerts, W.; Chrostowski, L. Silicon photonics circuit design: methods, tools and challenges. *Laser & Photonics Reviews* **2018**, *12*, 1700237.
- (7) Bogaerts, W.; Selvaraja, S. K.; Dumon, P.; Brouckaert, J.; De Vos, K.; Van Thourhout, D.; Baets, R. Silicon-on-insulator spectral filters fabricated with CMOS technology. *IEEE Journal on Selected Topics in Quantum Electronics* **2010**, *16*, 33–44.
- (8) Pantouvaki, M.; Srinivasan, S. A.; Ban, Y.; De Heyn, P.; Verheyen, P.; Lepage, G.; Chen, H.; De Coster, J.; Golshani, N.; Balakrishnan, S.;

- Absil, P.; Van Campenhout, J. Active components for 50 Gb/s NRZ-OOK optical interconnects in a silicon photonics platform. *Journal of Lightwave Technology* **2017**, *35*, 631–638.
- (9) Selvaraja, S. K. S.; Bogaerts, W.; Dumon, P.; Van Thourhout, D.; Baets, R. Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology. *IEEE J. Sel. Top. Quantum Electron.* **2010**, *16*, 316–324.
- (10) Selvaraja, S. K. *Wafer-Scale Fabrication Technology for Silicon Photonic Integrated Circuits*. Ph.D. thesis, University of Gent, 2011.
- (11) Kumar Selvaraja, S.; Rosseel, E.; Fernandez, L.; Tabat, M.; Bogaerts, W.; Hautala, J.; Absil, P. SOI thickness uniformity improvement using wafer-scale corrective etching for silicon nano-photonics device. In *The 2011 Annual Symposium of the IEEE Photonics Benelux Chapter*; IEEE, 2011; pp 289–292.
- (12) Kumar Selvaraja, S. SOI thickness uniformity improvement using corrective etching for silicon nano-photonics device. In *IEEE International Conference on Group IV Photonics GFP*; IEEE, 2011; pp 71–73.
- (13) Dan-Xia Xu; Schmid, J. H.; Reed, G. T.; Mashanovich, G. Z.; Thomson, D. J.; Nedeljkovic, M.; Xia Chen; Van Thourhout, D.; Keyvaninia, S.; Selvaraja, S. K. Silicon photonic integration platform—Have we found the sweet spot? *IEEE J. Sel. Top. Quantum Electron.* **2014**, *20*, 189–205.
- (14) Selvaraja, S. K.; Winroth, G.; Locorotondo, S.; Murdoch, G.; Milenin, A.; Delvaux, C.; Ong, P.; Pathak, S.; Xie, W.; Sterckx, G.; Lepage, G.; Van Thourhout, D.; Bogaerts, W.; Van Campenhout, J.; Absil, P. 193 nm immersion lithography for high-performance silicon photonic circuits. In *27th Optical Microlithography Conference as part of the SPIE Advanced Lithography Symposium*; SPIE, 2014; p 90520F.
- (15) Wang, X.; Shi, W.; Yun, H.; Grist, S.; Jaeger, N. A. F.; Chrostowski, L. Narrow-band waveguide Bragg gratings on SOI wafers with CMOS-compatible fabrication process. *Opt. Express* **2012**, *20*, 15547.
- (16) Beausoleil, R. G.; Faraon, A.; Fattal, D.; Fiorentino, M.; Peng, Z.; Santori, C. Devices and architectures for large-scale integrated silicon photonics circuits. Optoelectronic Integrated Circuits XIII. *Proc. SPIE* **2011**, 794204.
- (17) Lu, Z.; Jhoja, J.; Klein, J.; Wang, X.; Liu, A.; Flueckiger, J.; Pond, J.; Chrostowski, L. Performance prediction for silicon photonics integrated circuits with layout-dependent correlated manufacturing variability. *Opt. Express* **2017**, *25*, 9712.
- (18) Siew, S. Y.; Li, B.; Gao, F.; Zheng, H. Y.; Zhang, W.; Guo, P.; Xie, S. W.; Song, A.; Dong, B.; Luo, L. W.; Li, C.; Luo, X.; Lo, P. G.-Q. Review of silicon photonics technology and platform development. *Journal of Lightwave Technology* **2021**, *39*, 4374–4389.
- (19) Fahrenkopf, N. M.; McDonough, C.; Leake, G. L.; Su, Z.; Timurdogan, E.; Coolbaugh, D. D. The AIM Photonics MPW: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits. *IEEE J. Sel. Top. Quantum Electron.* **2019**, *25*, 1–6.
- (20) Giewont, K.; et al. 300-mm monolithic silicon photonics foundry technology. *IEEE J. Sel. Top. Quantum Electron.* **2019**, *25*, 1–11.
- (21) Wang, M.; Ribero, A.; Xing, Y.; Bogaerts, W. Tolerant, broadband tunable 2×2 coupler circuit. *Opt. Express* **2020**, *28*, 5555–5566.
- (22) Miller, D. A. B. Perfect Optics with Imperfect Components. *Optica* **2015**, *2*, 747–750.
- (23) Bandyopadhyay, S.; Hamerly, R.; Englund, D. Hardware error correction for programmable photonics. *Optica* **2021**, *8*, 1247.
- (24) Xing, Y.; Spina, D.; Li, A.; Dhaene, T.; Bogaerts, W. Stochastic collocation for device-level variability analysis in integrated photonics. *Photonics Research* **2016**, *4*, 93.
- (25) Weng, T.-W.; Zhang, Z.; Su, Z.; Marzouk, Y.; Melloni, A.; Daniel, L. Uncertainty quantification of silicon photonic devices with correlated and non-Gaussian random parameters. *Opt. Express* **2015**, *23*, 4242.
- (26) Weng, T. W.; Melati, D.; Melloni, A. I.; Daniel, L.; et al. Stochastic simulation and robust design optimization of integrated photonic filters. *Nanophotonics* **2017**, *6*, 299–308.
- (27) Qian, K.; Nikolić, B.; Spanos, C. J. Hierarchical modeling of spatial variability with a 45nm example. In *Design for Manufacturability through Design-Process Integration III*; SPIE, 2009; p 727505.
- (28) Chrostowski, L.; Wang, X. X.; Flueckiger, J.; Wu, Y.; Wang, Y.; Fard, S. T. Impact of fabrication non-uniformity on chip-scale silicon photonic integrated circuits. In *Conference on Optical Fiber Communication*; Technical Digest Series: Washington, D.C., 2014; p Th2A-37.
- (29) Xing, Y.; Dong, J.; Khan, U.; Bogaerts, W. Correlation between pattern density and linewidth variation in silicon photonics waveguides. *Opt. Express* **2020**, *28*, 7961–7968.
- (30) Bogaerts, W.; Xing, Y.; Khan, U. Layout-aware variability analysis, yield prediction, and optimization in photonic integrated circuits. *IEEE J. Sel. Top. Quantum Electron.* **2019**, *25*, 1–13.
- (31) Tzintzarov, G. N.; Ildefonso, A.; Teng, J. W.; Frounchi, M.; Djikeng, A.; Iyengar, P.; Goley, P. S.; Khachatryan, A.; Hales, J.; Bahr, R.; Buchner, S. P.; Mcmorrow, D.; Cressler, J. D. Optical single-event transients induced in integrated silicon-photonics waveguides by two-photon absorption. *IEEE Trans. Nucl. Sci.* **2021**, *68*, 785–792.
- (32) Bogaerts, W.; Baets, R.; Dumon, P.; Wiaux, V.; Beckx, S.; Taillaert, D.; Luyssaert, B.; Campenhout, J. V. Nanophotonic Waveguides in Silicon-on-Insulator Fabricated With CMOS Technology. *J. Lightwave Technol.* **2005**, *23*, 401–412.
- (33) Croon, J.; Storms, G.; Winkelmeier, S.; Pollentier, I.; Ercken, M.; Decoutere, S.; Sansen, W.; Maes, H. Line edge roughness: characterization, modeling and impact on device behavior. In *Technical Digest: International Electron Devices Meeting*; IEEE, 2002; pp 307–310.
- (34) Cheng, B.; Roy, S.; Roy, G.; Adamu-Lema, F.; Asenov, A. Impact of intrinsic parameter fluctuations in decanano MOSFETs on yield and functionality of SRAM cells. *Solid-State Electron.* **2005**, *49*, 740–746.
- (35) May, G. S.; Spanos, C. J. *Fundamentals of Semiconductor Manufacturing and Process Control*; Wiley-Interscience, 2006.
- (36) Stine, B.; Boning, D.; Chung, J. Analysis and decomposition of spatial variation in integrated circuit processes and devices. *IEEE Transactions on Semiconductor Manufacturing* **1997**, *10*, 24–41.
- (37) Steele, D. A.; Coniglio, A.; Tang, C.; Singh, B.; Nip, S.; Spanos, C. J. Characterizing post-exposure bake processing for transient- and steady-state conditions in the context of critical dimension control. *Proceedings of SPIE*; SPIE, 2002; Metrology, Inspection, and Process Control for Microlithography XVI; Vol. 4689, p 517.
- (38) Huang, K.; Kupp, N.; Carulli, J. M.; Makris, Y. Process monitoring through wafer-level spatial variation decomposition. *Proceedings—International Test Conference*; IEEE, 2013.
- (39) Boning, D. S.; Chung, J. E. Statistical metrology: understanding spatial variation in semiconductor manufacturing. In *Microelectronic Manufacturing Yield, Reliability, and Failure Analysis II*; SPIE, 1996; Vol. 2874, pp 16–26.
- (40) Chiang, C.; Kawa, J. *Design for Manufacturability and Yield for Nano-scale CMOS*; Springer, 2007.
- (41) Ouma, D. O. *Modeling of Chemical Mechanical Polishing for Dielectric Planarization*. Ph.D. thesis, Massachusetts Institute of Technology, 1998.
- (42) Perlin, K. An image synthesizer. *ACM Siggraph Computer Graphics* **1985**, *19*, 287–296.
- (43) Qian, K. *Variability Modeling and Statistical Parameter Extraction for CMOS Devices*. Ph.D. thesis, University of California, Berkeley, 2015.
- (44) Xing, Y.; Dong, J.; Dwivedi, S.; Khan, U.; Bogaerts, W. Accurate extraction of fabricated geometry using optical measurement. *Photonics Research* **2018**, *6*, 1008.
- (45) Xing, Y.; Wang, M.; Ruocco, A.; Geessels, J.; Khan, U.; Bogaerts, W. A Compact silicon photonics circuit to extract multiple parameters for process control monitoring. *OSA Continuum* **2020**, *3*, 379–390.
- (46) Xing, Y.; Dong, J.; Khan, U.; Bogaerts, W. Hierarchical model for spatial variations of integrated photonics. In *IEEE 15th International Conference on Group IV Photonics (GFP) 2018*, IEEE: Cancun, Mexico, 2018; pp 1–2.