nature photonics

Article

https://doi.org/10.1038/s41566-024-01394-2

Inverse-designed low-index-contrast structures on a silicon photonics platform for vector–matrix multiplication

In the format provided by the authors and unedited

Contents

S1 Effective-Index Approximation Overview **2**

S2 Normalizing the scattering error **5**

S3 Device Fabrication **6**

S4 Measurement **7**

S5 Measurements Methodology **8**

S6 Analyzing the computational resources for 3D and p2DEIA structures **9**

S1 Effective-Index Approximation Overview

The ever-increasing interest in wave-based analog computing demands efficient numerical methods for the simulation and design of optically-large electromagnetic structures. One way to alleviate the computational effort is by approximating a structure with an effective model. An effective model is meant to replace the original problem with a simpler one that is crafted to capture the essential features of the electromagnetic features in question. The computational effort of the original problem might rapidly scale with the size of the structure, however, while the effective model of it might scale more favorably with regards to computational resources. One way of formulating an effective model is through dimensionality reduction. One example is the Eigenmode Expansion Method (EME) in which a waveguide of arbitrary cross-section (but uniform in z) is reduced to several guided modes, each propagating in 1D according to their index. This removes two dimensions of computational effort while leaving one to be handled via appropriate means such as FDTD. Furthermore, in the variational method for planar dielectric structures, one may replace the 3D structure with a two-dimensional effective index distribution for emulating the in-plane wave propagation in the actual 3D structure[1]. The effective index distribution is optimized by minimizing the functional form of Maxwell's equations for estimating the in-plane wave propagation.

In this work, the traditional effective model based on slab mode effective index (p2DEIA) is utilized for the inverse design of 3D planar structures based on silicon photonics. In a typical silicon planar structure as shown in Figure S1(a), the wave is confined within the silicon core in the form of a slab mode. One can change the propagation characteristics of a slab mode, by changing the thickness of the core via etching down into some desired areas. The resulting discontinuities between the etched and non-etched regions create interfaces through which the wave reflects and refracts and one can manipulate the light propagating from one side to the other by engineering the topology and shape of the interfaces. The computational cost of full-wave 3D simulations of Maxwell's equations might render a design process highly inefficient. Specifically, the computational cost grows in some methods such as the inverse design in which one needs to simulate the structure hundreds of times during the optimization process. On the other hand, in silicon planar structures, the wave mostly propagates in the form of the fundamental slab mode in different regions (see Figure S1(a)). Therefore, the wave function along the out-of-plane direction (y-direction in Figure S1(a)) is known provided the wave couples to a fundamental slab mode as it propagates through interfaces between etched and non-etched regions, that is, the coupling between the fundamental slab mode in each region and higher-order modes is negligible. Therefore, for simulating the wave propagation efficiently throughout the planar structure in Figure S1(a), one can simulate the *in-plane* slab mode propagation (the $xz-plane$ in Figure S1(a)) by an *effective* two-dimensional model. In the effective 2D model as illustrated in Figure S1(b), the regions and interfaces of the structure in Figure S1(a) are mapped onto a two-dimensional plane. Therefore, a 2D scalar wave satisfying Helmholtz's wave equation represents the in-plane propagation of the slab mode throughout the structure in Figure S1(a). The refractive index of the regions in the 2D model is assigned from the effective indices of the slab modes confined in the

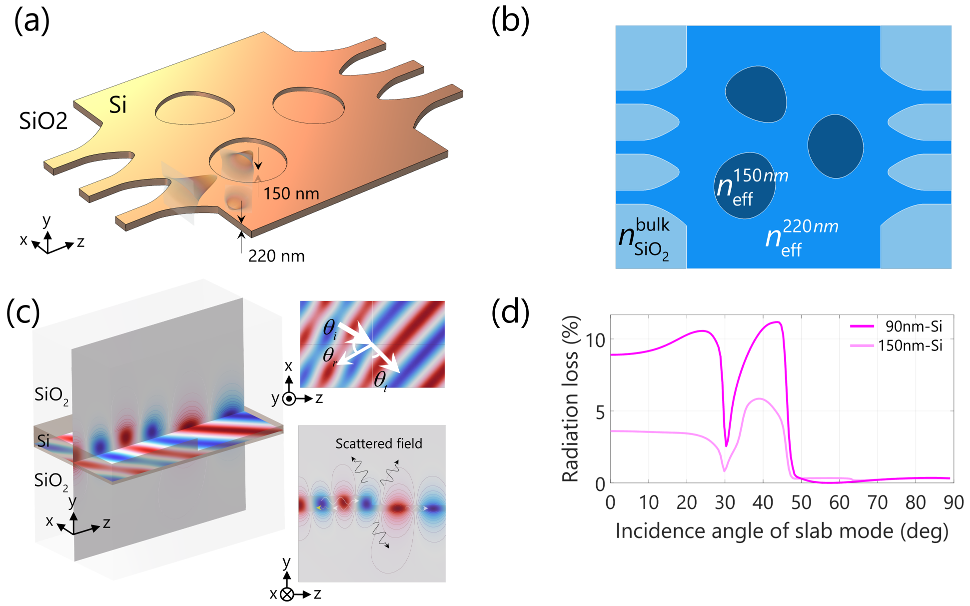

Fig. S1 Effective index approximation for planar structures. (a) The 3D model of a typical silicon planar structure with etched regions for manipulating the in-plane wave propagation. (b) Effective index approximation method for projecting the 3D model onto an effective 2D model (p2DEIA) for capturing the main features of the in-plane wave propagation. (c) The full-wave 3D simulation of a slab mode propagation impinging on an etched interface with incidence angle θ_i. On the xz plane, the primary in-plane wave propagation, which is captured by p2DEIA, is illustrated. On the yz plane the secondary out-of-plane scattering due to height discontinuity is shown that can not be captured by p2DEIA. (d) The energy of the out-of-plane scattering (radiation loss) vs. the angle of incidence, θ_i.

respective regions in the 3D structure. As a result of this modeling, the phase progression of the slab mode through the in-plane geometry is accurately modeled as long as the slab mode propagates without encountering an interface. In other words, the bulk phase progression can be accurately modeled using the p2DEIA.

Focusing on phase progression comes at a cost with regard to accuracy when an interface is encountered. The ground truth for interface behavior is the full 3D model as illustrated in Figure. S1(c). The wave propagation through the 3D interface exhibit a few effects not included in the 2D model such as the out-of-plane scattering as shown by the field distribution on the yz plane in Figure S1(c). In contrast, simulating plane wave illumination of a planar interface between two materials of differing effective indices using the p2DEIA model will result in the standard Fresnel reflection and transmission coefficients. As a meaningful way to understand the interface errors introduced by the p2DEIA model, we can compare the Fresnel reflection and transmission coefficients to the reflection and transmission of the fundamental slab modes in the full 3D model as a function of the angle of incidence.

The channeling to the secondary modes (out-of-plane scattering), which we call the radiation loss, causes the reflection and transmission of the fundamental slab modes to slightly deviate from the Fresnel reflection and transmission coefficients calculated using the 2D model. In Figure S1(d), we plot the fraction of the input power that channels to radiation loss i.e $L = 1 - |R|^2 - |T|^2$ vs. the angle of incidence for the 220 nm to 150 nm and the 220 nm to 90 nm interfaces.

As one can see, the amount of radiation loss scales directly with the height of the etch while following similar angular dispersion. Also, it peaks around the critical angle of incidence beyond which total internal reflection happens and the radiation loss becomes smaller. Therefore, the accuracy of the p2DEIA relates to the etch height, that is, the larger the etch height the more inaccurate the p2DEIA is. In our design, we try to control the error of the p2DEIA by using a relatively shallow etch of 70 nm (corresponding to a 150 nm-thick silicon slab). Note that while 150 nm presents a maximum radiation loss $\approx$ %6 just prior to the critical angle, the average radiation loss is lower when considering all possible angles.

The amplitude and phase of the transmission coefficients through interfaces with different etch depths are plotted in Figure S2. As one can see, beyond the critical

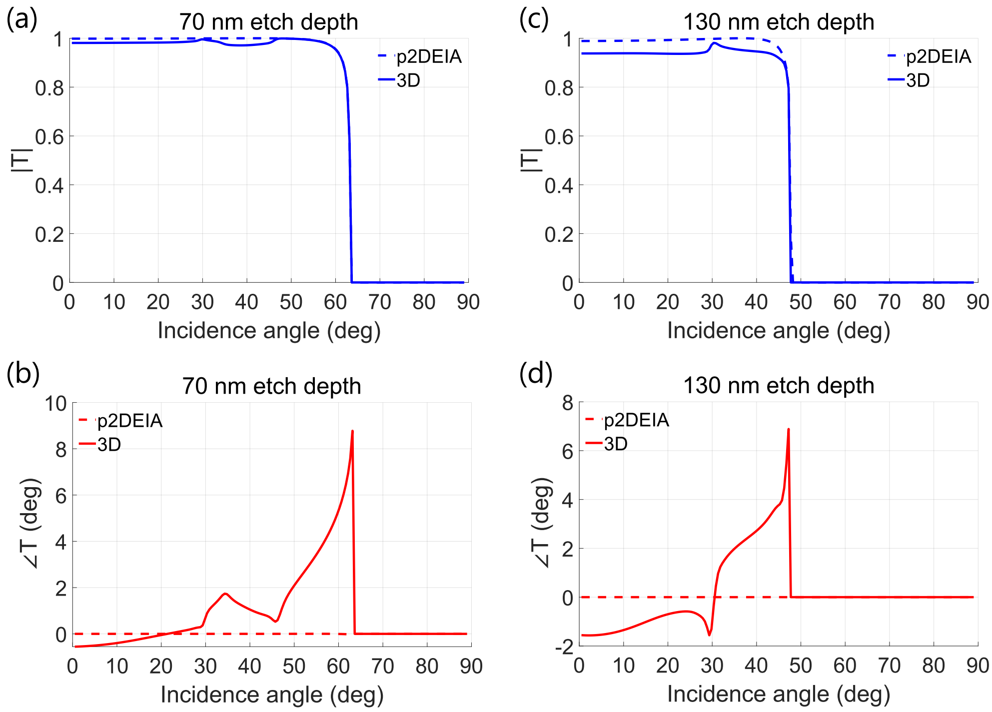

Fig. S2 Interface transmission coefficient (a,b) The magnitude and phase of the transmission coefficient for the 3D 220nm to 150nm-thick silicon interface (the solid curves) and for the corresponding projected interface in the p2DEIA model (dashed curves). (c,d) The corresponding plots for the 220nm to 90nm silicon interface.

angle the data for 3D and p2DEIA coincide since there is negligible radiation loss due to total internal reflection. In the p2DEIA model, the effective index contrast is greater for the 130nm etch than the 70nm etch and naturally, this will reduce the transmission magnitude. However, this difference is slight and transmission is near unity for both. Before the critical angle, one can observe a far greater discrepancy in the 130nm etch compared to the 70nm etch between the 3D model and p2DEIA due to the out-of-plane scattering modes. The reason for this transmission drop is poorer mode overlap (both shape and mode center) between the fundamental modes yielding increased reflection, 3D radiation loss, and scattering into higher-order guided modes (see Figure S3 for comparing the reflections for p2DEIA and 3D). Also, from the phase plots in Figure S2(b) and (d), one can observe that in the region near the critical angle the phase discrepancy between the 3D model and the p2DEIA spikes due to the heightened mismatch. However, before the critical angle, the phase error profile is smaller for 70 *nm*-etch height than the 130 *nm* case which further justifies the use of 70 *nm*-etch height for better accuracy.

For completeness, we also present analogous data for the reflection coefficients in Figure S3. As a generally feed forward system, this is of less importance, however it is clear that the reflection magnitude is less for the shallower etch depth.

In a practical design, the wave will interact with many interfaces between the input and output. Each interaction will introduce some power loss and a slight phase rotation the accumulation of which results in an overall loss factor and phase error.

S2 Normalizing the scattering error

We desire to consider a meaningful relative error for our devices. Consider an error defined on a realized experimental $N \times N$ matrix $\mathbf{T}_{\text{exp}}$ when compared to a similar target matrix $\mathbf{T}_{\text{targ}}$. A natural metric for the squared error is the square of the Frobenius norm, $\|\mathbf{T}_{\text{exp}} - \mathbf{T}_{\text{targ}}\|_F^2$. However, this metric makes it difficult to compare matrices of different sizes as a larger matrix will naturally have a larger norm than a smaller one of similar values. On the other hand, these matrices will likely *not* have similar values as the elements of a passive scattering matrix will naturally decrease with the size of the matrix. This potentially will make the error appear smaller for larger matrices. We seek an error metric that allows us to compare devices in spite of these complications.

Let us consider the *average relative squared error* of the matrix terms $T_{i,j}$. The average relative squared error is $\|\mathbf{T}_{\text{exp}} - \mathbf{T}_{\text{targ}}\|_F^2/(N^2\Delta_{\text{typ}}^2)$ where Δ_{typ} is some typical complex distance value and N^2 is the number of terms in matrix. In many instances in physics and engineering, the typical distance value is taken to be the *target* value itself. However, such a scheme over-emphasizes the importance of small target values. Indeed, if a single target term would be zero, any experimental realization would have infinite relative error even if practically it performed quite well. Rather, we note that we are interested in target matrices that are unitary (or close to it) and will have elements with magnitudes on the order of $|\mathbf{T}_{i,j}| \approx 1/\sqrt{N}$. Since we are interested in the *complex* error, a realized experimental value that is 180 degrees out of phase would possess a typical absolute error distance, Δ_{typ}, of $2/\sqrt{N}$. Therefore, the average relative squared error becomes $\text{error}^2 = \|\mathbf{T}_{\text{exp}} - \mathbf{T}_{\text{targ}}\|_F^2/(4N)$.

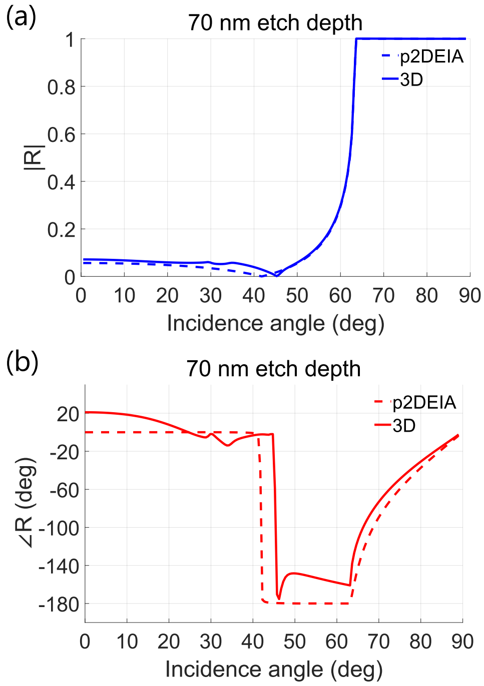

Fig. S3 Interface reflection coefficient (a,b) The magnitude and phase of the reflection coefficient for the 3D 220nm to 150nm-thick silicon interface (the solid curves) and for the corresponding projected interface in the p2DEIA model (dashed curves). (c,d) The corresponding plots for the 220nm to 90nm silicon interface.

This metric has the following useful properties. Suppose that the target matrix, $\mathbf{T}_{\text{targ}}$, is unitary. If $\mathbf{T}_{\text{exp}} = \text{-}\mathbf{T}_{\text{targ}}$, then error^2 = 1.00. If $\mathbf{T}_{\text{exp}} = \mathbf{0}$, then error^2 = 0.25. If $\mathbf{T}_{\text{exp}} = \mathbf{T}_{\text{targ}}$, then error^2 = 0.00 regardless of any zero terms.

S3 Device Fabrication

The designed metastructures are fabricated using Advanced Micro Foundry (AMF) 180 nm SOI process with a 2 µm thick buried oxide using their 180nm process. The photonic signal is routed on chip using single mode waveguides with 500 nm width and 220 nm height. This process provides two levels of etch (70 *nm* and 130 *nm*), of which we used the most shallow. Typically, this level of etch is used to define devices such as grating couplers and ridge waveguides, and we are clearly using it for a non-standard purpose.

Our inverse design process needed to have appropriate inline filters to prevent Design Rule Constraint (DRC) violations. Despite this, the geometry resulting from the inverse design process had a handful of minor DRC violations which were manually corrected with no effect on the simulated device performance. It is worth noting that the fabricated design required no DRC waivers and can be considered "standard."

This success is largely due to the relatively large, rounded structures required by this low-contrast method.

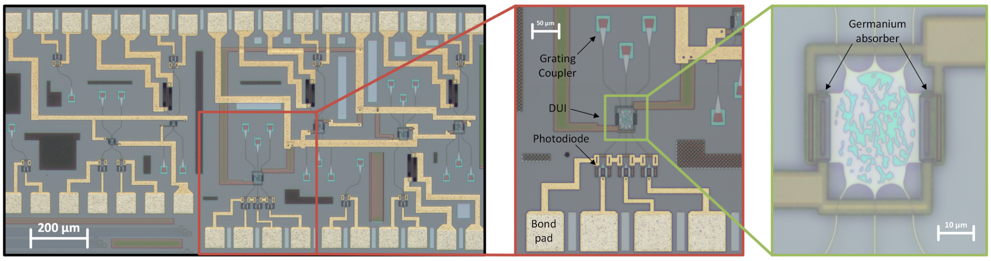

Fig. S4 Micrograph of the physical experiment showing several 2×2 and 3×3 kernels and various calibration structures. The 3×3 kernel is highlighted.

S4 Measurement

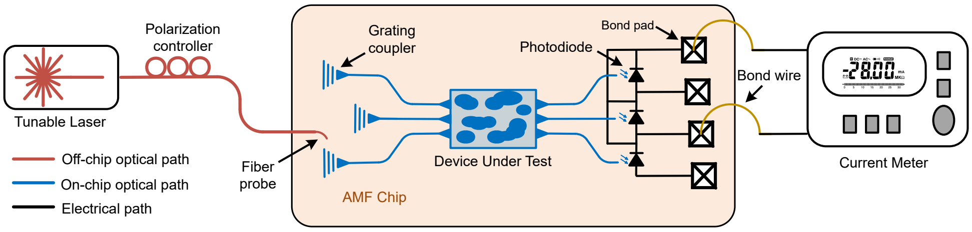

Fig. S5 Measurement setup. A tunable laser is used to illuminate one of several grating couplers. The light moves through the inverse-designed structure. Finally, the power is measured at each of the three photodiodes with a current meter.

The fabricated $N \times N$ inverse-designed kernel is characterized using PDK's I/O devices. For the inputs, a grating coupler per waveguide, each with an efficiency of about 40 % is positioned to facilitate light coupling to the chip. At the output, photodiodes with a responsivity of about 0.8 A/W are utilized to convert optical energy into electrical current, which is then measured using a current meter. A tunable laser is used to test the performance of the fabricated inverse-design devices. As shown in Figure S5, the infrared light is launched to the on-chip grating coupler using a fiber probe. The light polarization is adjusted using a polarization controller to match the grating coupler's polarization response and maximize the coupled optical power to the on-chip waveguide.

S5 Measurements Methodology

By assuming that the inverse-designed structure and all I/O devices function within the linear region, a single individual input is excited one at a time using a tunable laser operating around 1550 nm, and the corresponding response is recorded. Utilizing the superposition property of the system, the responses to N inputs were aggregated, resulting in the $N \times N$ transfer function. The complex vector is encoded in optical intensity and phase. To evaluate the intensity response of the kernel, the input light is coupled to one input, and the optical power of N outputs is captured using individual photodiodes. As for the phase measurement, direct phase measurements in silicon photonics are exceedingly difficult to perform. The primary challenge is that small variations between single-mode waveguides caused by slight over/under etching of the waveguide sidewalls will result in a perturbation to the waveguide index. Over any significant length, this can easily result in phase shifts between two waveguides. These etching artifacts are dependent on the etch loading and position on the wafer. This makes an absolute phase reference very difficult to define and fabricate. In fact, this is the very challenge that our device wishes to solve. By allowing the light to flow through a series of refractive interfaces, defined by shallow etches, we have significantly reduced the dependence on the absolute side wall position. Since our device operates through a relatively few number of refractions at the etch interfaces, it is primarily dependent on the angles and shapes of these interfaces, and less on their absolute position. It is worth noting that our structure works based on constructive and destructive interactions between the waves from various input ports and from the waves refracted through various internal structures. In other words, in contrast to an MZI mesh, these various channels share the geometry in which the wave propagates. As is well known, the gold standard for measuring material thicknesses and indices in deposited materials is Variable Angle Spectral Ellipsometry (VASE). The details of VASE are not relevant here, but the general principle is. By using a broadband source and a spectrometer, one can collect thousands of measurements from a device. The spectral scattering parameters are spectrally constrained by geometry and material in conjunction with the Kramers-Kronig relations. Once the basic model has been defined, there are only a handful of parameters to adjust including material thicknesses and material resonance parameters. If you can match complex spectroscopic data with a small number of parameters, then you can be assured that your model is correct, and the device is functioning as intended. With confidence in the model, one can then calculate the phase of light at any point, even deep within the material stack, which would be experimentally inaccessible. In short, in many cases, the best way to obtain phase is through a spectroscopic model. A similar reasoning can be considered here. Obviously, we can directly evaluate the magnitude of each scattering parameter using standard power measurements as was shown in the main text to a good agreement. However, as a means of estimating phase, we evaluate our device spectroscopically as a means of verifying our model. In our case, our model is defined by a geometry, but once that geometry has been confirmed, then phase can be inferred from the model, rather than direct measurement. It is worth noting that the nature of inverse design is that the fields share space within the device such that even a power measurement represents interference between various paths through the structure and these measurements reinforce each other. Given the linear nature of electromagnetic waves, a linear combination of complex inputs will produce the expected vector of complex outputs. As a final note, it is worth noting that while it is true that direct phase measurement in silicon photonics platforms is hard to conduct due to the reasons mentioned earlier, controlling and adjusting relative phases between the output ports is achievable and commonly performed using y-junctions between pairs of output waveguides in order to utilize interference to convert the relative phase information into intensity output of the y-junctions which is then measured by a photodiode located at the end of the junction. Incorporating phase shifters in the arms of the y-junction provides the ability to control, adjust and compensate for the phase differences between the pairs of outputs, thus mitigating the effect of unavoidable fabrication imperfections mentioned above. Studying this topic can be the subject of the next phase of our investigation.

S6 Analyzing the computational resources for 3D and p2DEIA structures

A natural question is what is quantitatively gained by performing the inverse-design method using the p2DEIA vs the full-wave 3D model. Since the inverse design may not be computationally feasible for large structures in the model, we will employ a heuristic analysis.

It is true that due to the planar nature of the metastructure, the height of the device is relatively constant and is equal to the small thickness of the silicon slab. However, in comparing the full-wave 3D model with the effective 2D model in terms of computational complexity there are two main differences (1) Despite being constant the thickness of the silicon must be discretized with sufficient cells, particularly near the etched interfaces for accurate estimation of high-k fields. Furthermore, the height of the cladding must be large enough such that it allows for capturing the out-of-plane scattering modes without getting the Perfectly-Matched Layers (PMLs) to interact with the high-k near fields that are around the etched interfaces. Therefore, for the 3D model, one would need half a wavelength for discretizing the silicon thickness and allocate at least one wavelength worth of thickness on each side of the slab for the height of the cladding. Also, each of the top and bottom PML layers would need at least half a wavelength worth of thickness. Therefore, optically speaking for discretizing the height of the simulation domain one needs three and a half wavelengths worth of thickness (two for cladding plus one half for silicon slab plus one for two PML layers on top and bottom). Since each wavelength needs to be discretized with at least 10 mesh cells (assuming linear shape functions), there is an additional constant factor of 35 per polarization for DOF for field calculations in the 3D model compared to the p2DEIA. (2) In p2DEIA, we only consider a scalar field for the out-of-plane magnetic field, whereas in the 3D model one would need to solve for all the polarization components that are $\{E_x, E_y, E_z\}$. This will introduce an additional factor of 3 for DOF in the 3D model. Note: In the p2DEIA model, we solved for $\{E_x, E_y\}$ rather than H_z due to the FEM formulation in COMSOL Multyphysics®. However, we believe the fundamental analysis remains true if one just solves for H_z. Therefore, assuming the in-plane DOF is the same for both p2DEIA and 3D, the 3D model requires 105 times more DOF for EM field calculations to account for the height of the device and two more polarization components. Finally, we conclude that despite both 3D and p2DEIA scale quadratically with the in-plane size of the device in terms of DOF i.e., both exhibit DOF$\sim \mathcal{O}(d^2)$ where d represent the in-plane length scale, the 2D model needs 105 times less DOF than the 3D for field calculations.

The inverse design utilizes the Adjoint method in which the gradient is calculated from a series of simulations in which both the "input" and "output" ports are illuminated. Given a chosen scattering parameter objective, whether the p2DEIA or 3D model is being optimized, the same number of simulations per optimization iteration is required. Therefore, to first order, in order to compare the computational effort required to perform an optimization, we need only compare the costs of solving for the optical fields in both cases. However, this is only approximate, as are many details which can affect the convergence in both cases.

In this section, we compare the computational resources needed to solve for the optical fields in the full-wave 3D and p2DEIA models. As for the case in the study, we consider the silicon photonics metastructure designed to perform 10×10 vector-matrix product (see the main text). We examine the structure by considering one random input vector of modes over the input ports. We employ the finite element method (FEM) to simulate time-harmonic optical fields at the design wavelength $\lambda_0 = 1.525\ \mu m$ using an Intel(R) Xenon(R) Gold 6130 CPU @ 2.10GHz with 16 cores and 225GB of RAM. For solving the linear systems, we employ the Generalized Minimum Residual Method (GMRES) iterative solver[2] and parallel direct solver PARDISO[3] using COMSOL 5.6. Table 1 includes information about the domain

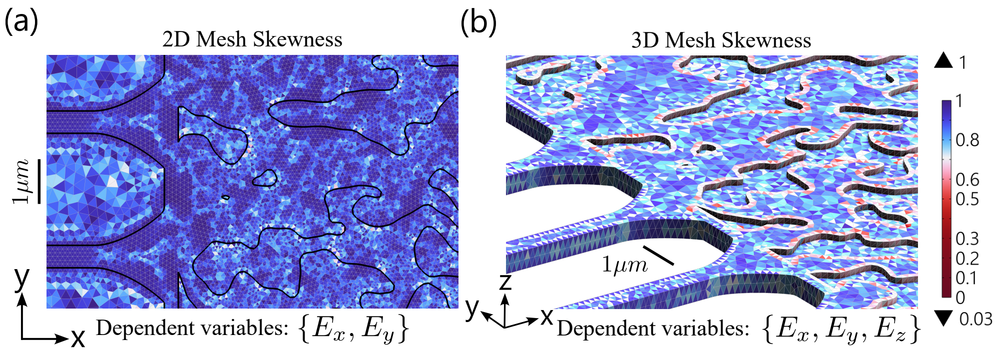

Fig. S6 Mesh skewness (a) Mesh plot for the p2DEIA model of the 10×10 silicon photonics metastructure. (b) The corresponding plot for the 3D model.

discretization, resources, and time needed to solve the linear systems as detailed below:

- **Mesh volume(area):** The volume(area) of the simulation domain of the 3D(p2DEIA) structure.
- **Mesh cells (order)**: The total number of cells discretizing the simulation domain. For the 3D simulation, the majority of the cells are tetrahedral. For the p2DEIA

simulation, the cells are triangular. Here *order* refers to the order of the shape function associated with individual cells.

- **Average quality**: We considered the mesh skewness as the quality factor as shown in Fig.S6. It quantifies how deformed a cell is compared with a reference cell that is optimal for discretization. The skewness merit varies between 0 (bad) and 1 (good). The average of the skewness merit over the simulation domain is reported.
- **Degrees of freedom (DOF)**: The number of unknowns that characterize the optical fields. It depends on the number of electric field components considered as the dependent variables, the shape function, and the number of cells. According to the case in the study, for the p2DEIA simulation, we consider $\{E_x, E_y\}$ and for the 3D model we consider all three components, $\{E_x, E_y, E_z\}$.
- **Physical memory (RAM)**: The size of the physical memory that is occupied by the linear systems.
- **Time**: The amount of time that is needed to solve the linear systems.

EM Model (Solver)	Volume or Area	Cells (Order)	Average Quality	DOF	RAM	Time
p2DEIA (PARDISO)	788 μm^2	181K (Cubic)	0.85	1.92M	8.01GB	4.9sec
p2DEIA (GMRES)	788μm^2	181K (Cubic)	0.85	1.92M	3.4GB	16min
3D (GMRES)	3058μm^3	926K (Cubic)	0.63	18.42M	186GB	5h 45min
3D (PARDISO)	3058μm^3	926K (Cubic)	0.63	18.42M	400GB	N.A.

Table 1 Comparison of computational resources required for performing one FEM numerical simulation with a random excitation between p2DEIA and full-wave 3D models using direct (PARDISO) and iterative (GMRES) solvers.

The data shows the p2DEIA model can significantly speed up the inverse design process by fast evaluation of the optical field in each iteration while using fewer computational resources.

References

[1] Hammer, M. & Ivanova, O. V. Effective index approximations of photonic crystal slabs: a 2-to-1-d assessment. *Optical and Quantum Electronics* **41**, 267–283 (2009).

[2] Saad, Y. & Schultz, M. H. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing* **7**, 856–869 (1986).

[3] Schenk, O. & Gärtner, K. Two-level dynamic scheduling in pardiso: Improved scalability on shared memory multiprocessing systems. *Parallel Computing* **28**, 187–197 (2002).