Hierarchical Design and Optimization of Silicon Photonics

Andrew Michaels, *Member, IEEE*, Ming C. Wu, *Fellow, IEEE*, and Eli Yablonovitch, *Fellow, IEEE*

***Abstract*—Silicon photonics is a rapidly maturing platform for optical communication and sensing. As systems leveraging silicon photonics have grown in size and complexity, so too has the demand for high performance silicon photonics components. In order to meet these demands, we propose a hierarchical approach to design and optimization of silicon photonics components. Our approach applies simple physical analysis to choose an effective starting geometry for a two-step gradient-based shape optimization. This optimization employs carefully chosen geometrical constraints in order to consistently produce robust, high performance devices which satisfy practical fabrication constraints of deep UV lithography. In order to demonstrate the versatility of method, we optimize a 3 dB coupler which achieves better than 0.04 dB excess loss over the O-band, a four port 3-dB coupler which achieves better than 0.41 dB excess loss and near 50:50 splitting over the O-band, and a fabrication-tolerant waveguide crossing which achieves better than 0.075 dB insertion loss over the O-band even when subject to $\pm 10\%$ silicon thickness variations. These results pave the way for high efficiency silicon photonic component libraries.**

***Index Terms*—Silicon photonics, nanophotonics, inverse design, shape optimization, 3 dB coupler, 2×2 splitter, waveguide crossing.**

I. Introduction

In RECENT years, there has been a trend towards designing large-scale silicon photonic systems. For example, integrated beam steering for LIDAR [1]–[3], optical switches [4]–[7], and optical computing platforms [8]–[10] require thousands of silicon photonic components. When cascading many components together as is common in these applications, the efficiency of each individual component becomes increasingly important. Typical silicon photonic components are often too lossy or, in the case of devices which rely on adiabatic transitions, too large for the given application. This situation is further exacerbated when the platform demands a larger operational bandwidth or a greater insensitivity to fabrication variations.

In response to this, shape and topology optimization methods have emerged as a promising solution [11]–[18]. Although demonstrated to be very powerful tools, these *inverse design* methods often suffer a glaring flaw: they are inherently *local* (rather than global) optimization techniques. Consequently, it is difficult to guarantee that a design "discovered" by the optimization method will end up meeting a desired specification. This becomes particularly evident when designing high performance devices with excess losses below 0.5 dB. In such cases, it becomes increasingly important that we choose a starting point for the optimization and geometric constraints which will lead to a final design with the desired high performance.

In spite of their importance, prior work on shape and topology optimization in nanophotonics has placed limited emphasis on how the starting condition and geometric constraints influence the final performance of an optimized device. In many works [11]–[17], starting structures with seemingly (or even intentionally) arbitrary sizes and shapes are chosen. These choices, while successful in seeding optimizations which yield drastic improvements in the figure of merit, do not always succeed in producing final devices which significantly outperform their hand-design counterparts. To overcome this, additional physical insight is essential.

In this manuscript, we will introduce a hierarchical approach to design and optimization of nanophotonic devices that minimizes guesswork and which consistently yields high performance designs. This methodology involves developing a starting geometry based on simple physical analysis and then applying a shape optimization process that strategically employs different geometric constraints in order to limit the likelihood of falling into undesirable local optima. Using this methodology, we are able to design common silicon photonic components with record high efficiencies and which are broadband and tolerant to fabrication variations.

II. Overview of Inverse Design Tools

A core component of our approach, which we describe in detail in the next section, is gradient-based shape or topology optimization. Specifically, in this work, we employ adjoint optimization with boundary smoothing based on the methods discussed in [19] extended to three dimensions. Because our goal is to design very high efficiency devices, it is important that our optimization leverage a simulation method that can handle sufficiently high resolutions in a reasonable amount of time. To this end, we use a custom implementation of the 3D Finite Difference Time Domain (FDTD) method which is fast, scales very well on parallel computing platforms, and makes efficient use of

Manuscript received April 24, 2019; revised July 18, 2019; accepted August 6, 2019. Date of publication August 19, 2019; date of current version September 12, 2019. This work was supported in part by the Advanced Research Projects Agency-Energy (ARPA-E) ENLITENED Program under Grant DE-AR0000849 and in part by the National Science Foundation (NSF) Partnership for Innovation–Technology Transition (PFI-TT) under Grant 1827633. *(Corresponding Author: Andrew Michaels.)*

The authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, CA 94720 USA (e-mail: amichaels@eecs.berkeley.edu; wu@eecs.berkeley.edu; eliy@eecs.berkeley.edu).

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

Digital Object Identifier 10.1109/JSTQE.2019.2935299

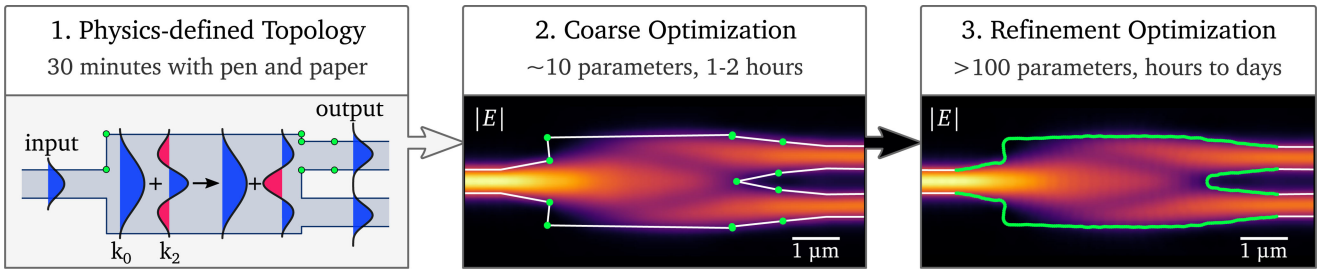

Fig. 1. Graphical overview of our hierarchical approach to design and optimization of electromagnetic structures. Our systematic approach involves developing a device topology based on simple physical analysis, coarsely optimizing this starting topology to improve its performance, and then refining the coarsely optimized structure with a second optimization that incorporates constraints. The end result is a device with high performance that can be fabricated.

memory. This enables us to optimize larger devices with higher resolutions. In our implementation, the simulation is driven by a ramped continuous wave source. This allows us to operate the FDTD algorithm as a frequency domain solver, simplifying the implementation of the adjoint method. Furthermore, using a custom FDTD implementation (as opposed to a commercial solver) gives us access to its internals, allowing us to ensure that forward and adjoint simulations are self-consistent, which is essential to calculating accurate device sensitivities.

In the spirit of enabling the photonics community to take advantage of the design methods we demonstrate in this manuscript, we open sourced our optimization tools, called *EMopt* [20]. EMopt includes the FDTD solver, mode solvers, and adjoint method implementation needed to optimize most silicon photonic components.

III. Hierarchical Design and Optimization

The way in which we apply inverse design tools has a tremendous impact on the quality of the end results. In this section, we introduce our *hierarchical approach to design and optimization* which systematizes gradient-based optimization of photonic devices. This approach consists of three main steps as depicted in Fig. 1 for the case of a 3 dB coupler: First we develop a starting structure based on physical intuition. Second, we run a coarse optimization of this starting structure which consists of a small number of degrees of freedom. Third, we run a refinement optimization with an increased number of degrees of freedom and constraints in order to produce an optimal structure that can be fabricated. We will discuss each step of this process in detail in the remainder of this section.

A. Physics-Defined Topology

The starting geometry we use to initialize an optimization is important as it can restrict the quality of both local and global optima. One reason for this is that we typically choose and fix the inputs and output of our device at the beginning of the design process. If we choose a starting structure with a certain size, then the final optimized structure will fit roughly within the same footprint. As a result, choosing a starting structure which is too small or too large to achieve the desired functionality will place limits on the maximum performance we can achieve.

In addition to influencing the quality of possible outcomes, the starting structure has a strong impact on which local optimum we fall into and whether that result is robust. An example of this is presented in the work by Su *et al.* [17]. Using randomized initial structures as a starting point, the authors optimize hundreds of grating couplers. Of the optimized results, only a comparatively small number of devices achieve a high coupling efficiency; the rest of the optimized structures exhibit lower performance, indicating that the corresponding optimizations fell into undesirable local optima. This observed sensitivity to initial conditions may give the impression that optimizing electromagnetic devices is imprecise and riddled with guesswork. Fortunately, we have recourse: strong physical intuition and understanding of Maxwell's equations. By leveraging our knowledge of wave mechanics, in many (if not all) cases, we can devise physically-motivated starting structures which lead to efficient optimized devices.

We refer to this physically-motivated starting structure as the *physics-defined topology*. For many components in silicon photonics, a usable topology can be devised relatively easily by applying simple design techniques, such as modal analysis.

As an example, consider the canonical problem of designing a 3 dB coupler (for TE fields). The goal of a 3 dB coupler is to split power from an input waveguide equally between two output waveguides. One way of achieving this functionality is to take advantage of modal dispersion in a multimode waveguide [21]. The fundamental TE mode of a narrower waveguide connected to a wider multimode waveguide will excite the symmetric modes of that wider waveguide, as depicted on the left hand side of Fig. 1. If the multimode waveguide is not too wide, the majority of the power will reside in the first two even modes of the waveguide with effective indices n_0 and n_2, respectively. As these two modes co-propagate over a distance $\Delta z = \lambda/[2(n_0 - n_2)]$, the relative phase between the two modes becomes equal to π, leading to a splitting of the field intensity towards the outside of the multimode waveguide. With the field evenly split and concentrated near the outer edges of the waveguide, we can terminate the multimode waveguide and place two output waveguides to capture the divided optical power.

This physically intuitive approach translates into a relatively well performing initial structure, which is shown in Fig. 2. The structure is defined in 220 nm thick silicon using a 110 nm partial

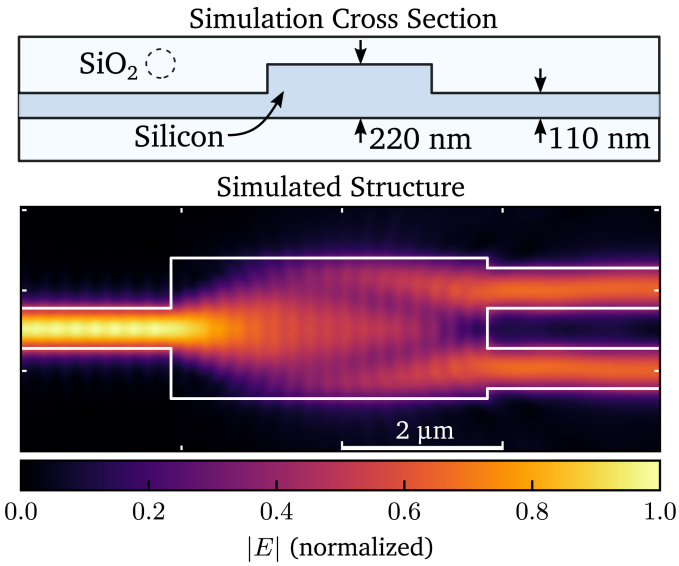

Fig. 2. (top) Diagram of the cross section of the silicon platform used for the 3 dB coupler. (bottom) Plot of a slice of the simulated electric field of a 3 dB coupler based on a multimode interference topology for a wavelength of 1310 nm. This serves as the starting point for gradient-based optimization.

etch. Based on a wavelength of 1310 nm and a multimode waveguide width of 1.75 μm, our previous analysis yields a device length of 3.94 μm. Despite the simplicity of this approach, the chosen topology yields a device with a reasonably high coupling efficiency of 93.3% (-0.301 dB). This gives us confidence that the overall size of the device and the chosen positions of the inputs and outputs are sufficient for obtaining a high efficiency device.

This example demonstrates the general thought process one can follow to develop a starting topology for many silicon photonic components. In fact, the same analysis can be applied to many other devices like waveguide crossings and 4 port 3 dB couplers, which we demonstrate later in this manuscript. Other devices, however, may require a different approach to discovering an effective topology; one such example, which we demonstrated previously in [22], is grating couplers. While developing the topology for an electromagnetic device may not always be as simple as this 3 dB coupler, we should still be able to apply our intuition of electromagnetics to come up with *something* which is superior to random guesses. To this end, the large amount of work on electromagnetic device design is tremendously useful.

B. Coarse Optimization

One of the primary goals of defining a physically-inspired device topology is to ensure that the inputs and outputs of the system are situated such that high efficiency designs exist within the design space. In many cases, the analysis used to develop this topology neglects some of the finer details of the device. For example, in our 3 dB coupler example, we chose the separation of the output waveguides to approximately match the intensity distribution formed by the beating modes within the larger multimode waveguide section of the device. In reality, these abrupt transitions between the multimode waveguide and the input and output waveguides are unlikely to be optimal. This presents us with an opportunity to apply shape optimization techniques to improve the coupler performance.

While we may be tempted to throw the full power of inverse design with many degrees of freedom at the problem, it is important that we proceed in a strategic manner. Typically when we run shape (or topology) optimizations with large numbers of degrees of freedom, we impose constraints to prevent the development of features that we are unable to fabricate. The combination of large numbers of degrees of freedom and constraints significantly increases the complexity of the design space. This can lead to a larger number of potentially undesirable local optima and a slower evolution towards the final optimized structure (an effect which we demonstrate in detail in Appendix A). This in turn makes it more difficult to make large modifications, like displacing a boundary of the structure over $\sim$ micron scale distances, to the structure.

To help mitigate these issues, we run an initial *coarse optimization* without constraints and only a few design variables ($\sim$10). By running an optimization on a coarsely parameterized structure, we are able to rapidly improve the general shape of the device. Convergence is typically significantly faster for simpler design spaces, and thus the coarse optimization allows us to both quickly verify the quality of our topology and evaluate the likelihood that the design process will lead to an efficient device. Furthermore, the result of the coarse optimization is an excellent starting point for a final optimization which includes more degrees of freedom and fabrication constraints.

In order to demonstrate the coarse optimization, we continue with our example of designing a 3 dB coupler. In this example, we choose the coarse parameterization to be the xy coordinates of the points which define the splitter, as depicted by the green dots in the left and middle parts of Fig. 1. In total there are seven points and hence fourteen design parameters. This choice allows the optimization to both manipulate the input and output waveguides and also to modify the size of the multimode waveguide section. Although in this manuscript we employ a polygon-based shape optimization method, any form of shape or topology optimization which affords control over the number of degrees of freedom that describe a device could be employed in our hierarchical approach.

The figure of merit used for this optimization is the coupling efficiency into the fundamental supermode of the two output waveguides which corresponds to equal splitting of optical power. All simulations are performed for a wavelength of 1310 nm with O-band operation in mind. The grid has a step size of 30 nm, providing a good trade-off between simulation speed and simulation accuracy. The optimization is terminated after 40 iterations at which point the figure of merit does not change appreciably.

In this optimization, we use the L-BFGS-B minimization algorithm which we found to yield rapid convergence. The progression of the figure of merit during the coarse optimization is shown in Fig. 3. Beginning with a modest value of 0.30 dB, the excess loss is quickly improved during optimization, reaching a low value of only 0.020 dB by the end of the optimization.

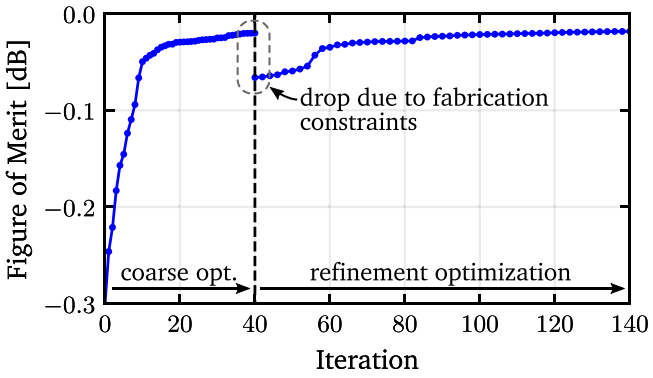

Fig. 3. To the left of the vertical dashed line, the figure of merit during the coarse optimization is plotted vs iteration. To the right of the vertical dashed line, the figure of merit during the refinement optimization is plotted. The visible decrease in performance at the beginning of the refinement optimization is result of imposing fabrication constraints.

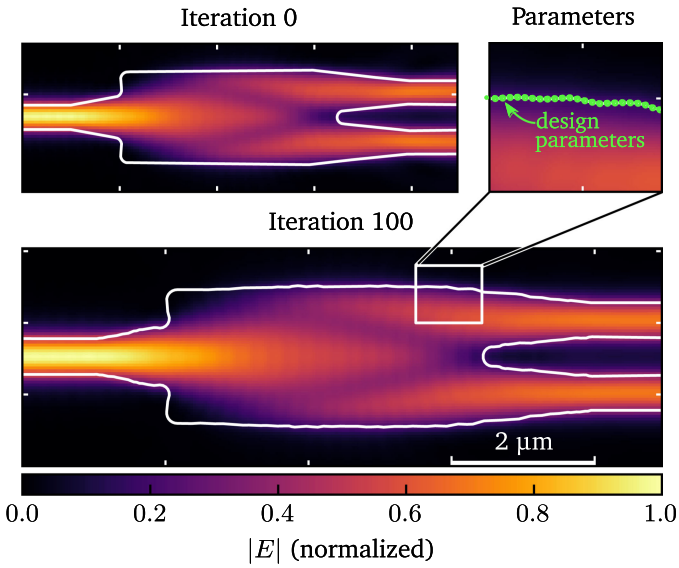

Fig. 4. Plots of the (top left) initial rounded structure and (bottom) final optimized structure for the refinement phase of the design process. The inset in the top right shows a section of the polygon which defines the boundary of the structure. In all three plots, the electric field amplitude is overlaid with an outline of the device boundaries.

This whole process took just over 20 minutes as detailed in Appendix B.

C. *Refinement Optimization*

Although the coarse optimization in our 3 dB coupler example successfully produces a structure with high coupling efficiency, the structure itself is not entirely practical as it contains sharp corners that are difficult to fabricate. Simply rounding off these sharp corners will inevitably result in a sub-optimal design. In order to rectify this issue, we can run a *refinement optimization* which includes more degrees of freedom and fabrication constraints. These additional degrees of freedom afford us more flexibility in the shape of the device, allowing us to further improve its performance while ensuring the device can be made.

When increasing the number of degrees of freedom which define a device boundary, we are prone to generating features which are too small to fabricate. In order to avoid this, it is common practice to incorporate constraints into the optimization [12], [16]. For waveguiding devices, we commonly constrain the radius of curvature which limits the formation of small features. Depending on the device, minimum gap size or line width constraints may also be imposed. Choosing such constraints is an extremely important part of the refinement optimization and can have a strong impact on the final performance of the optimized device, the manufacturability of the device, and the time it takes to run the optimization.

As a demonstration, we continue with our example 3 dB coupler. We choose the figure of merit for the refinement optimization to be

$$F(\mathbf{E}, \mathbf{H}, \vec{p}) = \eta(\mathbf{E}, \mathbf{H}) - \mathcal{P}_{\mathrm{roc}}(\vec{p}) - \mathcal{P}_{\mathrm{gap}}(\vec{p}) \qquad (1)$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic fields and $\vec{p}$ is the set of design parameters. The individual functions η, $\mathcal{P}_{\mathrm{roc}}$, and $\mathcal{P}_{\mathrm{gap}}$ are the mode match integral, radius of curvature constraint, and gap size constraint, respectively [19]. In this example, we choose a minimum radius of curvature of 120 nm and a minimum gap size of 200 nm to make the device compatible with readily available deep UV lithography. The penalty functions are weighted such that violation of the constraint at a single point reduces the figure of merit by up to 2%. By imposing the radius of curvature and gap constraints using a penalty function, the constraints become part of the optimization. This ensures that the structure we end up with is optimal given these constraints.

It is important to note that if we used the structure generated by the coarse optimization without modification, the sharp corners would significantly violate the radius of curvature constraint, causing corresponding penalty function to dominate the figure of merit. As a result, the refinement optimization would attempt to reduce the impact of the penalty function by rounding out the structure, potentially leading to very slow convergence and a susceptibility to falling into undesirable local optima. To avoid this, we can manually round off the corners of the coarse optimized structure (using a fillet operation) and then use this modified structure as the starting point for the refinement optimization. By doing so, we ensure that the radius of curvature (and gap) constraint is initially satisfied. In this case, the penalty function acts as a barrier preventing the optimization from entering regions of the design space which cannot be fabricated. Although this process may lead to an initial drop in device performance as shown in Fig. 3, the subsequent refinement optimization can make up for the lost performance.

The rounded 3 dB coupler used as the starting point for the refinement optimization is depicted in the top left of Fig. 4. In order to represent boundaries with rounded features, we increase the number of points in the polygon which defines the coupler structure as shown in the top right in Fig. 4. These additional points serve as the (significantly augmented) degrees of freedom for the refinement optimization and enable the coupler to take on a more sophisticated shape than was possible with the coarse optimization. The optimization is allowed to run for a total 100 iterations which yields a satisfactory improvement in

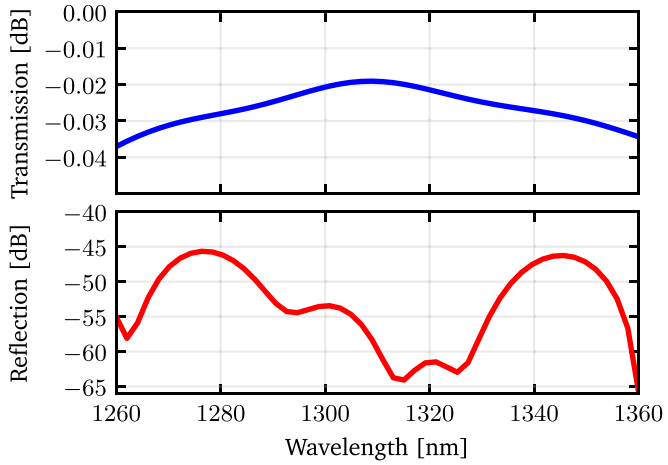

Fig. 5. Plots of the (top) O-band coupling efficiency and (bottom) reflection for the final optimized device. The optimized device achieves better than −0.04 dB transmission and −45 dB reflection over the full O-band.

performance in a modest amount of time (just over 9 hours). In this optimization, our simulations are run using a grid spacing of 20 nm which leads to improved accuracy and hence the increased runtime. Over the course of the refinement optimization, the figure of merit improves from −0.065 dB to under −0.02 dB as shown in Fig. 3. The structure that achieves this is depicted in the bottom of Fig. 4 and notably has a more fluid shape than the result of the coarse optimization.

Not only is the coupling efficiency at the design wavelength high, but it is also maintained over a broad wavelength range as shown in Fig. 5. The optimized coupler achieves an excess loss lower than 0.04 dB over the entire O-band (1260 nm to 1360 nm). Furthermore, the back reflection into the input waveguide is lower than −45 dB over the full wave length range as shown in Fig. 5. This performance exceeds other published results for both hand designed and numerically optimized 3 dB couplers [12], [23]–[28].

The final structure produced by the refinement optimization not only achieves high efficiency, but it also does so while satisfying the radius of curvature and gap size constraints that we imposed on the structure. As a result, we expect the simulated device to translate well to experiment with minimal additional loss. Notice that even if we modified our constraints (for example, if we improved our fabrication process and could hit smaller feature sizes), we only need to rerun our refinement optimization. To a certain degree, the device topology and coarse structure are independent of the particular constraints we impose. This presents an additional advantage of our hierarchical approach to nanophotonic design: it allows us to more easily design component libraries which are compatible with a wider range of lithographic requirements.

IV. Application to Silicon Photonics

The hierarchical approach to photonic design that we introduced in the previous section presents a promising strategy for designing silicon photonic components which demand a high level of performance. In order to further demonstrate this method, we have designed broadband four port 3 dB couplers and efficient fabrication-tolerant waveguide crossings. For each example, we develop a physically-inspired topology for the device, parameterize it, and run coarse and refinement optimizations. The end results of these optimizations are high performance components that can be fabricated with deep UV lithography.

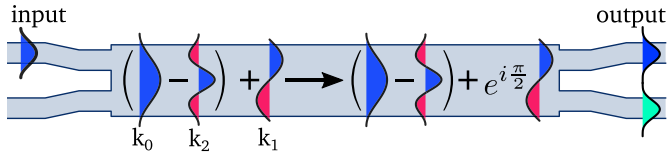

Fig. 6. Graphical explanation of the basic operating principle of the four port 3 dB coupler. Light from either of the two input waveguides primarily excites the first three modes of a multimode waveguide. These modes copropagate and interfere leading to a field pattern which splits evenly at the output.

In these examples, we employ the silicon photonic platform depicted in Fig. 2 which consists of a 220 nm top silicon layer, that is patterned with a 110 nm shallow etch. The devices are clad top and bottom with silicon dioxide. This partially-etched platform is chosen as it is compatible with lower-loss ridge waveguides, and the thicknesses are commonly employed in the nanophotonics community.

A. Four Port 3-dB Coupler

A key component in many integrated photonic systems is the four port 3 dB coupler (or 2×2 splitter). This four port splitter takes light from either of its two input waveguides and splits the optical power evenly between its two ouput waveguides. Because the device has an inherently asymmetric operation, it is more difficult to design compared to the three port optical splitter that we used as our previous guiding example.

Just as with its three port counterpart, multimode interference effects lead to a good topology for the four port 3 dB coupler. Specifically, our device topology consists of a larger multimode waveguide with two symmetrically situated waveguides connected both at its input and its output. Light incident on the multimode waveguide from one of the two input waveguides will excite higher order modes which copropagate; based on the relative phases of these higher order modes, the field can be effectively split between two output waveguides [21].

For the purpose of choosing the approximate dimensions of our initial design, it is convenient to consider the first three TE modes (with wavenumbers k_0, k_1, and k_2) of the multimode waveguide section of the device as depicted in Fig. 6. It is important to note that the fundamental and second order modes have even symmetry while the first order mode has odd symmetry. If we add the fundamental mode and second order mode with the correct phase, we obtain a symmetric field with two positive lobes. Adding this resulting field to the first order (odd) mode of the multimode waveguide, we can concentrate the optical power in the top half of the multimode waveguide. This field combination roughly corresponds to the input of the multimode waveguide which is excited by one of the input waveguides as depicted in Fig. 6.

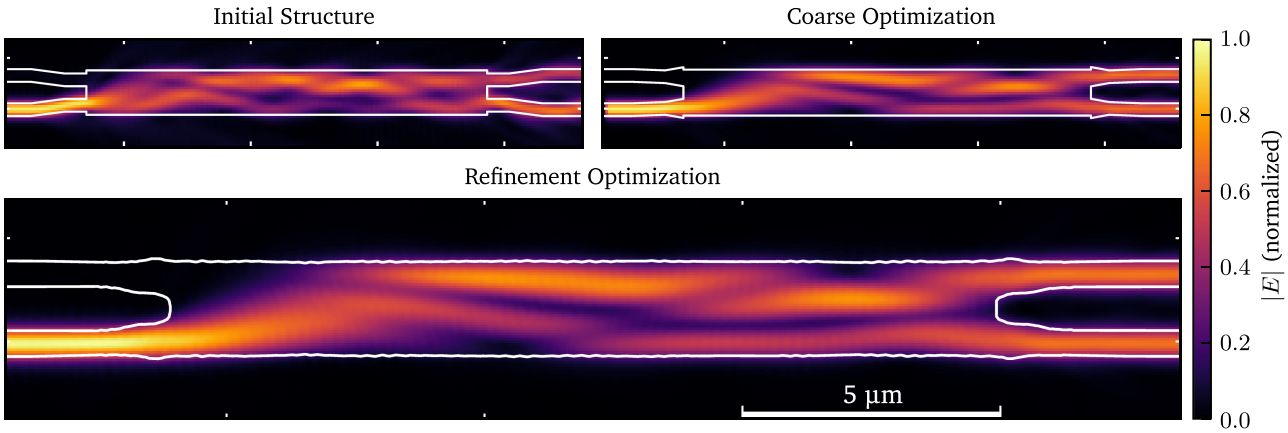

Fig. 7. Plot of the electric fields of the initial, coarse optimized, and refinement optimized four port 3 dB coupler. The outline of the optimized structure is overlaid on top of the electric field plots, showing the evolution of the device as a result of the optimization process.

As the three modes copropagate, the relative phases between the modes will change. The relative phase between the fundamental and second order mode is $\varphi_{20} = (k_0 - k_2)\, z$ and the phase between the fundamental and first order mode is $\varphi_{10} = (k_0 - k_1)\, z$. Notice that if we were to add the even symmetry field profile generated by adding the fundamental and second order mode to the odd-symmetry first order mode with a $\pi/2$ phase difference, the resulting distribution of power inside the multimode waveguide will split evenly away from its center. This condition occurs when the waveguide has a length such that $\varphi_{20} = 2\pi p$ and $\varphi_{10} = (2q+1)\pi/2$ where p and q are integers. These relative phase shifts correspond to waveguide lengths of $L_{2\pi} = p\lambda/(n_0 - n_2)$ and $L_{\pi/2} = (2q+1)\lambda/4(n_0 - n_1)$ where n_0, n_1, and n_2 are the effective indices of the three modes and λ is the free-space wavelength. While we ideally would like to find values for p and q such that $L_{2\pi} = L_{\pi/2}$, this condition is not strictly possible to satisfy because the effective indices are typically irrational numbers. Fortunately, we can choose p and q such that $L_{2\pi} \approx L_{\pi/2}$. The resulting approximate length is the length of the multimode waveguide which makes up our four port splitter. At the output of the multimode waveguide, we place two symmetrically situated waveguides. The evenly split fields within the multimode waveguide couple into these two output waveguides.

Our ultimate goal is to design a four port splitter which maintains high coupling efficiency and even splitting over the whole O-band. As such, we choose initial dimensions which are suited to operate at 1310 nm. The input and output waveguides are chosen to be 500 nm wide and are separated by 500 nm. The multimode waveguide section of the device is chosen to be 1.75 μm which yields a device with a modest footprint that can be accurately simulated in a reasonable amount of time. Based on these parameters, we find that $L_{2\pi} \approx L_{\pi/2} \approx 15.8\ \mu$m.

The starting structure produced by this simple analysis and its simulated electric field is plotted in the top left of Fig. 7. Although the fields are effectively coupled into the output waveguides, there is visible imbalance in the fraction of power exiting in the top and bottom waveguides, which is reflected by a coupling efficiency of -1.74 dB. This is likely a result of a few different factors. First, our analysis considered only the first three modes of the multimode waveguide section. In reality, higher order modes exist and will contribute to the fields in the device in a way unaccounted for. Next, our analysis largely neglected the exact size and position of the inputs and outputs which both have a large impact on the overall performance on the device. Despite these non-idealities which hinder the performance of the starting structure, we can still expect the device to be an appropriate size and shape that will enable efficient 3 dB splitting.

In order to improve this initial design, adherent to our hierarchical approach, we run a coarse optimization of the structure. In this coarse optimization, we define the structure using the six degrees of freedom: the input width, input separation, output width, output separation, multimode waveguide width, and multimode waveguide length. In order to ensure that the device works for both input ports identically, we force it to be symmetric about the xz plane. The figure of merit for this optimization is the coupling efficiency (approximated by the mode match integral) into the desired output field at 1310 nm which consists of the sum of the first two super modes of the two-output-waveguide system with a relative $\pi/2$ phase shift applied between the modes. Finally, all simulations during the coarse optimization use a 40 nm grid spacing which speeds up the process significantly at the expense of a minor reduction in simulation accuracy. As with our three port splitter, the structure is optimized using the L-BFGS-B minimization algorithm.

The coarse optimization is allowed to run until the figure of merit has converged which takes 35 iterations. The result of this is shown in the top right of Fig. 7. Due to the relatively small number of parameters, the coupling efficiency of the structure increases rapidly from it's initial value of -1.75 dB to nearly -0.25 dB in only two iterations. During the remainder of the optimization, the device is fine tuned until it achieves final coupling efficiency of only -0.176 dB.

Next, we move on to the refinement optimization, which has two purposes. First, it should produce an optimal structure that does not contain any sharp corners or small features which cannot be resolved using deep UV lithography. Second, it should ensure the power leaving the two output ports is as balanced (close to 50/50) as possible over the full O-band.

The latter goal, in particular, requires some additional attention. Naively, we might expect that simply maximizing the coupling efficiency at a few wavelengths spread over the O-band would be enough to ensure even splitting. Unfortunately, this turns out to not be the case, because we will inherently be operating near a maximum in the mode match equation which occurs when the power is exactly equally split between the two output waveguides. When we are near a maximum, modest changes to the relative power in the two waveguides lead to only small changes in the mode match integral. In practice, the relative power in the two outputs needs to change by more than a percent or two before the mode match figure of merit starts to decrease significantly.

One way to circumvent this issue is to apply a penalty to the figure of merit which increases when the outputs become imbalanced. A tidy form of this penalty function is the mode match with respect to the set of fields that are orthogonal to the desired set of output fields. This function has a minimum at the point in the design space where the overlap with the desired fields is maximized and the outputs are balanced. Deviating from this point causes the function to increase and thus more heavily penalize the figure of merit. With this in mind, our complete figure of merit for the refinement optimization is

$$F(\mathbf{E},\mathbf{H},\vec{p}) = \frac{1}{N_\lambda}\sum_{\lambda}\big[\eta(\mathbf{E}_\lambda,\mathbf{H}_\lambda) - \alpha\eta_{\mathrm{ortho}}(\mathbf{E}_\lambda,\mathbf{H}_\lambda) - \mathcal{P}_{\mathrm{roc}}(\vec{p})\big] \quad (2)$$

where η is the mode match integral calculated with respect to the desired and orthogonal fields, $\mathcal{P}_{\mathrm{roc}}(\vec{p})$ is the radius of curvature penalty function, λ denotes the simulation wavelength, N_λ is the total number of wavelengths, and α sets the weight of the orthogonal mode match penalty function. In this case, we use $\alpha = 5\times10^3$ which initially penalizes the figure of merit by about 15%. This figure of merit is just the average of multiple figures of merit computed at different wavelengths which behaves similarly to a minimax figure of merit when working with high efficiency devices since the maximum coupling efficiency possible is equal to one.

In this optimization, we co-optimize our splitter at three different wavelengths spanning a 50 nm range: 1285 nm, 1310 nm, and 1335 nm. For all three wavelengths, the structure is identical and a maximum 120 nm radius of curvature is enforced. The forward and adjoint simulations of the structure are carried out with a uniform grid spacing of 30 nm. Given the size of the simulation domain, this gives a good trade off between speed and accuracy. The final results presented later in this section are produced using higher resolution simulations which have a 20 nm grid spacing to ensure higher accuracy.

The optimization is run for a total of 100 iterations (which took a little over 35 hours on 256 cores of our Intel Xeon based cluster) at which point the figure of merit stops increasing appreciably. Initially, the figure of merit begins at −0.87 dB, significantly lower than the final coarse optimization result. This is primarily a consequence of the orthogonal field mode match constraint which is not initially well satisfied. The refinement optimization is able to rapidly improve the figure of merit to below one tenth of a decibel in 20 iterations, eventually settling on a final value of −0.061 dB.

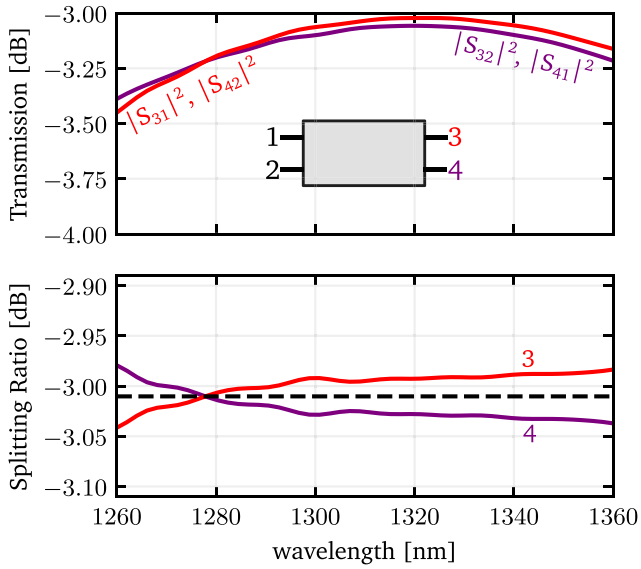

Fig. 8. Plot of the optimized four port 3 dB coupler performance. On top, the transmission into the two outputs is plotted versus wavelength. On the bottom, the splitting ratio (defined as the fraction of output power in each output waveguide) is plotted versus wavelength. The black dashed line corresponds to 50%, which is the desired splitting ratio.

The performance of the final optimized structure is shown in the top half of Fig. 8. As desired, the total excess loss remains low over the full O-band, reaching a maximum value of 0.41 dB at 1260 nm. This is further reflected in a 0.1 dB bandwidth of 61 nm. In addition to low excess loss, a splitting ratio very close to 50% is achieved over a broad range of wavelengths as shown in the bottom half of Fig. 8. Specifically, the relative fraction of output power in either of the two waveguides deviates from 50% by no more than a few hundredths of a decibel over the full O-band.

These results are superior to previously reported values for multimode interference and directional coupler designs [16], [29]–[32] and are consistent with adiabatic-transition-based designs [33]–[37] which are typically at least an order of magnitude larger than our design. Such high performance makes these optimized 2×2 splitters ideal for large scale integrated photonics applications that demand low loss splitters without sacrificing large amounts of die area.

B. Fabrication-Tolerant Waveguide Crossing

One of the great advantages of optical waveguides compared to conventional electrical interconnects is that they can intersect each other with minimal crosstalk. This functionality turns out to be essential to many silicon photonic integrated circuits which are only able to leverage a single layer of silicon. In large systems in particular, waveguide crossings must be very low loss over the relevant bandwidth, even in the presence of any fabrication errors.

As a final demonstration of our hierarchical approach, we have designed a waveguide crossing tolerant to fabrication variations.

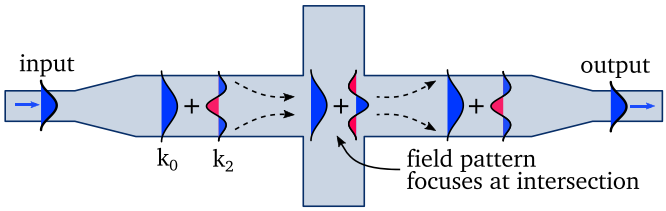

Fig. 9. Graphical explanation of the basic operating principle of a waveguide crossing. Light from the fundamental mode of an input waveguide is coupled primarily into the first two even modes of a wider multimode waveguide. These two modes co-propagate and interfere to form a focusing intensity pattern near the center of the crossing waveguide. This focusing effect allows us to place an intersecting perpendicular waveguide without significantly altering the propagation of light in the horizontal multimode waveguide.

In particular, we have optimized a crossing which maintains exceptionally low insertion loss over the O-band even when subject to $\pm 10\%$ thickness variations in the top silicon layer. Our starting point for this optimization is a standard waveguide crossing design based on multimode interference effects. The topology we choose, depicted in Fig. 9, operates in a manner identical to the three port 3 dB coupler but with the exception that the desired crossing length is four times the splitter length. This produces a focusing field pattern within the multimode waveguide which limits interaction with an intersecting waveguide.

For our starting structure, we use a 1.7 μm wide multimode waveguide which, based on the two-mode analysis, corresponds to a crossing length of 15 μm. In order to ensure that the focusing distance within the crossing waveguide is sufficiently long (which is mediated by the taper length), we make the input and output taper lengths equal to the width of the multimode waveguide (1.7 μm). These tapers connect the multimode crossing waveguide to the 500 nm input and output waveguides. The simulated fields of this starting structure are shown in the top left of Fig. 10. As desired, the power is effectively coupled from the input to output waveguide with little visible light coupling into the perpendicular waveguide.

Unfortunately, this standard waveguide crossing design is neither very broadband nor insensitive to fabrication variations, and in particular extreme silicon thickness variations. Therefore, our goal in running the coarse and refinement optimizations is to improve its performance when subject to such variations. In order to do so, our figure of merit for the coarse optimization will be the average (mode matched) transmission into the output waveguide at 1310 nm for three different top silicon thicknesses: the desired 220 nm thickness, a 10% increase in thickness, and a 10% decrease in thickness. For each thickness, we define the crossing structure using the same 110 nm partial etch depth.

For the purpose of the coarse optimization, we parameterize the structure using only four degrees of freedom: the length of the input and output tapers, the length of the crossing waveguide, and width of the crossing waveguide at the end of the taper, and the width of the crossing waveguide at the waveguide intersection.

The structure itself is simulated on a uniform 40 nm grid for the sake of speed. The optimization is performed using the L-BFGS-B minimization algorithm and is allowed to iterate until the figure of merit changes by less than 10^{-6}, which occurs after 24 iterations. Despite having only four degrees of freedom to manipulate, the optimization is able to rapidly improve the average transmission to -0.0269 dB and the total variation is drastically reduced to only 0.0102 dB at 1310 nm.

The structure produced by this coarse optimization and its simulated electric field is shown in the top right of Fig. 10. Although this coarsely optimized structure already achieves very good performance, there may be more room for improvement. Furthermore, the coarse result contains sharp corners that need to be eliminated. As with our previous examples, this is easily accomplished using a refinement optimization. To initialize this refinement optimization, we use the coarse optimization result with rounded corners and an increased number of vertices (one point every 40 nm). The positions of these vertices are the design parameters of the optimization. The figure of merit of the refinement optimization is the same as the coarse optimization with the addition of a radius of curvature penalty term. Finally, all simulations in this final optimization use a 22 nm grid spacing to ensure accurate results.

The optimization is run until the value of this figure of merit decreases by less than 10^{-6}, which occurs after 26 iterations. The final result is shown at the bottom of Fig. 10. Despite the fact that we modified the coarse optimized structure by rounding off its corners, the initial drop in performance at the start of the refinement optimization is very small. This is explained by the fact that our coarse optimized design is designed to be very insensitive to thickness variations, and we expect this insensitivity to apply to other small non-thickness-related changes to the structure as well. During the course of the refinement optimization, the minimum coupling efficiency improves from -0.0410 dB to -0.0292 dB which is approximately a 30% improvement. This means that 30% more crossings can be used along a path for the same amount of loss.

Compared to the initial maximum loss of more than 0.5 dB, the final optimized result is both highly efficient and exceptionally insensitive to thickness variations. Fig. 11 highlights that this improved tolerance to variations is a direct consequence of the improved bandwidth of the optimized structure. In the top half of Fig. 11, the transmission of a typical hand-designed waveguide crossing is plotted versus wavelength. While a crossing with the desired silicon thickness achieves a high peak coupling efficiency of >-0.01 dB, this performance quickly drops to nearly -0.2 dB at the edges of the O-band. When the thickness varies by $\pm 10\%$, the device's usefulness is largely eliminated because the minimum coupling efficiency at 1310 nm drops to -0.4 dB and further to -0.8 dB at the edges of the O-band.

Our optimized design has significantly improved performance across the O-band by comparison, both with and without thickness variations as shown in the bottom half of Fig. 11. Note that both larger plots in Fig. 11 share the same y axis scale. It is immediately apparent that the sensitivity to wavelength and silicon thickness variations is reduced by roughly an order of magnitude. The inset in the bottom of Fig. 11 shows a zoomed in version of the transmission of the optimized structure. When subject to $\pm 10\%$ variations in the silicon thickness, our

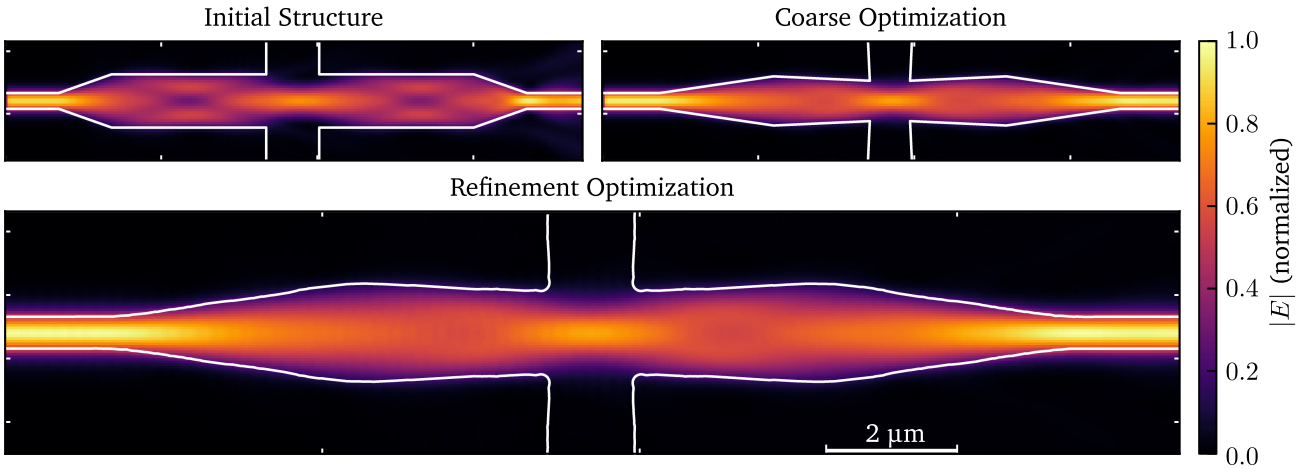

Fig. 10. Plots of the (top left) starting waveguide crossing structure, (top right) result of the coarse optimization, and (bottom) result of the refinement optimization. In all three cases, the fields are simulated with a wavelength of 1310 nm. In all three plots, the magnitude of the electric field taken from a slice running through the center of the device is overlayed with an outline of the structure.

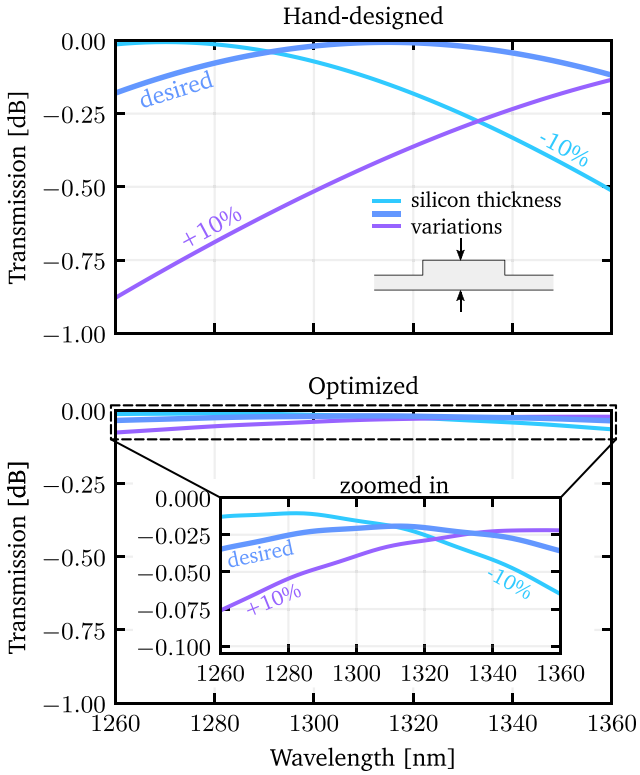

Fig. 11. Plots of the performance of the optimized waveguide crossing (bottom) compared to a typical hand-designed waveguide crossing (top). When subject to ±10% SOI thickness variations, the optimized design maintains a maximum loss across the O-band that is nearly 10 times smaller than a typical waveguide crossing.

optimized design achieves better than −0.075 dB transmission, better than −40 dB back reflection, and better than −38 dB cross-talk over the entire O-band.

This broadband behavior and insensitivity to fabrication variations comes only at the expense of a small reduction in peak coupling efficiency (−0.0065 dB to −0.019 dB). Compared to previously published results [38]–[44], our optimized crossing has comparable peak coupling efficiency and significantly higher bandwidth. To our knowledge, our optimized crossing achieves better O-band performance when subject to variations than any previously reported design. Even though our choice of thickness variations may be extreme, we expect the improved bandwidth of the optimized crossing to also reduce its sensitivity to other fabrication variations.

V. Conclusion

In order to design high efficiency silicon photonic components, we have proposed a hierarchical approach to design and optimization which leverages both simple physical analysis and numerical optimization techniques in a systematic way. Our approach consists of three steps. First, we define a *physics-defined topology* which roughly accomplishes our design goal and serves as the starting point for optimization. Next, a starting geometry based on this topology is rapidly improved using a *coarse optimization* with a limited number of degrees of freedom. Finally, the result of the coarse optimization is optimized further using a *refinement optimization* which incorporates a large number of degrees of freedom and design constraints. The result of this final optimization is a high efficiency device which can be fabricated. Using this approach, we have demonstrated how we can successfully design three port 3 dB couplers with better than 0.04 dB excess loss over the entire O-band, four port 3 dB couplers with better than 0.41 dB excess loss and near 50:50 splitting over the whole O-band, and finally waveguide crossings with better than 0.075 dB insertion loss over the O-band even when subject to ±10% silicon thickness variations.

These three examples demonstrate the effectiveness with which our approach can be applied to essential silicon photonic components. The application of our strategy does not stop here: many silicon photonics components remain which could be further improved by applying our hierarchical design methodology. We believe this method will pave the way for high performance and fabrication-tolerant silicon photonic component libraries.

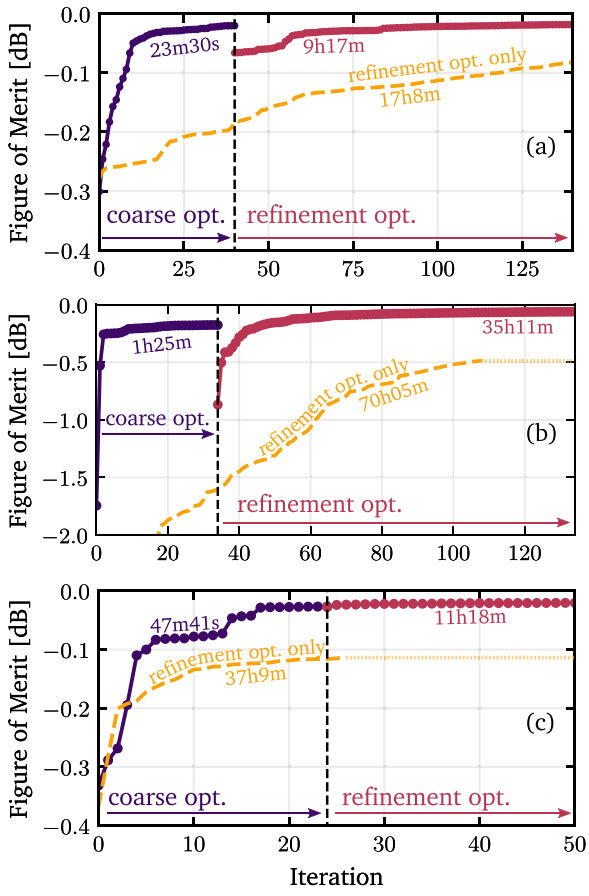

Fig. 12. Plots of the figure of merit vs iteration number for the (a) 3 dB coupler, (b) 2×2 splitter, and (c) waveguide crossing optimizations. In each case, the coarse optimization is plotted in purple to the left of the vertical dashed line and the refinement optimization is plotted in red to the right of the vertical dashed line. Imposed on this plot, the dashed yellow trace shows the figure of merit for a refinement optimization which is run without first running a coarse optimization. In all three cases, skipping the coarse optimization results in a much longer runtime and a significantly lower final performance.

Appendix A
Significance of the Coarse Optimization

A key step in our hierarchical design approach is the coarse optimization which we have found to significantly improve both the speed of the optimization and the quality of the final design. In order to see why the coarse optimization is beneficial, it is instructive to re-apply our methodology to our three examples, but this time omitting the coarse optimization (i.e., directly optimized the physically-motivated starting structure with a refinement optimization).

The resulting optimizations are depicted in Fig. 12. In each of the three plots, the figure of merit versus iteration number is plotted for the original coarse optimization, original refinement optimization, and the refinement-only optimization. Furthermore, the optimization run time is labeled in each case. In the case of the 3 dB coupler, which is the simplest of the optimizations we considered, omitting the coarse optimization results in very slow convergence. In double the time it takes for the original coarse and refinement optimization to finish, the refinement-only optimization only manages to reach just above -0.1 dB.

TABLE I
Computational Cost of Optimizations

	Yee Cells	Iterations	Simulations	Run Time
3 dB Coupler Coarse	1.84×10^6	40	102	23m30s
3 dB Coupler Refinement	7.76×10^6	100	414	9h17m
2×2 Splitter Coarse	3.90×10^6	34	92	1h25m
2×2 Splitter Refinement	9.25×10^6	100	306 ($\times3$)	35h11m
Crossing Coarse	1.69×10^6	22	70 ($\times3$)	47m41s
Crossing Refinement	10.17×10^6	18	114 ($\times3$)	11h18m

The situation is worse in the case of the 2×2 splitter and waveguide crossing. In the case of the 2×2 splitter, omitting the coarse optimization causes the optimization to converge on a local optimum that is nearly 0.5 dB lower than the original hierarchical approach. Furthermore, convergence to this final local optimum takes twice as long. Similarly, in the case of the waveguide crossing, omitting the coarse optimization causes the refinement optimization to converge on a local optimum that is around 0.1 dB lower than the optimum found with the hierarchical approach and takes three times longer to do so.

These three examples highlight the importance of the coarse optimization, and the hierarchical approach as a whole. Without the coarse optimization, we readily observe slower convergence to a worse solution.

It is worth noting that the coarse optimization comes with one potential disadvantage: it requires shape or topology optimization methods which allow us to directly choose the number and type of degrees of freedom that describe the device. This poses a potential challenge to density (or gray scale) methods and level set methods. The refinement optimization, on the other hand, is not constrained in this way and could benefit from the growing variety of shape and topology methods in the literature. We hope to explore this further in future work.

Appendix B
Computational Time

A key detail that one needs to consider when running an optimization is its computational cost. The size of the device one is trying to design, the desired accuracy of the underlying simulations, and the complexity of the devices operation (and hence the figure of merit) all have an important impact on the time and resources that are required to run an optimization.

The three examples we have presented in this work can serve as a rough benchmark for how long an optimization will take to run. Table I shows the size, number of iterations, number of simulations, and run time for each of our examples. All three examples were run on a 10 Gbps infiniband connected cluster consisting of 32 Intel Xeon E5-2670 CPUs with 16 GB of DDR3 memory per CPU. Considering the age of these machines, we expect improved runtime on more modern hardware.

In general, a minimum of two simulations are required per iteration in order to calculate the figure of merit and gradient, however on average more than one figure of merit and gradient

calculation is required as a result of the L-BFGS-B minimization algorithm that we employ. Based on the time it takes to run a single simulation we can get a rough sense of how long an optimization will take.

REFERENCES

[1] J. K. Doylend *et al.*, "Two-dimensional free-space beam steering with an optical phased array on silicon-on-insulator," *Opt. Express*, vol. 19, pp. 21595–21604, Oct. 2011.

[2] J. Sun, E. Timurdogan, A. Yaacobi, E. S. Hosseini, and M. R. Watts, "Large-scale nanophotonic phased array," *Nature*, vol. 493, pp. 195–199, Jan. 2013.

[3] C. V. Poulton *et al.*, "High-performance integrated optical phased arrays for chip-scale beam steering and LiDAR," presented at the Conf. Lasers Electro-Opt., May 2018, Paper ATu3R.2.

[4] S. Han, T. J. Seok, N. Quack, B.-W. Yoo, and M. C. Wu, "Large-scale silicon photonic switches with movable directional couplers," *Optica*, vol. 2, pp. 370–375, Apr. 2015.

[5] D. Nikolova *et al.*, "Scaling silicon photonic switch fabrics for data center interconnection networks," *Opt. Express*, vol. 23, pp. 1159–1175, Jan. 2015.

[6] T. J. Seok, N. Quack, S. Han, R. S. Muller, and M. C. Wu, "Large-scale broadband digital silicon photonic switches with vertical adiabatic couplers," *Optica*, vol. 3, pp. 64–70, Jan. 2016.

[7] E. Bernier *et al.*, "Large-scale silicon photonic switch," presented at the Opt. Fiber Commun. Conf. Expo., Mar. 2018, Paper Th1J.1.

[8] J. L. O'Brien, "Optical quantum computing," *Science*, vol. 318, pp. 1567–1570, Dec. 2007.

[9] X. Qiang *et al.*, "Large-scale silicon quantum photonics implementing arbitrary two-qubit processing," *Nature Photon.*, vol. 12, pp. 534–539, Sep. 2018.

[10] Y. Shen *et al.*, "Deep learning with coherent nanophotonic circuits," *Nature Photon.*, vol. 11, pp. 441–446, Jul. 2017.

[11] J. Lu and J. Vučković, "Objective-first design of high-efficiency, small-footprint couplers between arbitrary nanophotonic waveguide modes," *Opt. Express*, vol. 20, pp. 7221–7236, Mar. 2012.

[12] C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, and E. Yablonovitch, "Adjoint shape optimization applied to electromagnetic design," *Opt. Express*, vol. 21, pp. 21693–21701, Sep. 2013.

[13] Y. Elesin, B. S. Lazarov, J. S. Jensen, and O. Sigmund, "Time domain topology optimization of 3d nanophotonic devices," *Photon. Nanostructures Fundam. Appl.*, vol. 12, pp. 23–33, Feb. 2014.

[14] B. Shen, P. Wang, R. Polson, and R. Menon, "An integrated-nanophotonics polarization beamsplitter with $2.4 \times 2.4\ \mu m^2$ footprint," *Nature Photon.*, vol. 9, pp. 378–382, May 2015.

[15] L. F. Frellsen, Y. Ding, O. Sigmund, and L. H. Frandsen, "Topology optimized mode multiplexing in silicon-on-insulator photonic wire waveguides," *Opt. Express*, vol. 24, pp. 16866–16873, Jul. 2016.

[16] A. Y. Piggott, J. Petykiewicz, L. Su, and J. Vučković, "Fabrication-constrained nanophotonic inverse design," *Sci. Rep.*, vol. 7, May 2017, Art. no. 1786.

[17] L. Su *et al.*, "Fully-automated optimization of grating couplers," *Opt. Express*, vol. 26, pp. 4023–4034, Feb. 2018.

[18] S. Molesky *et al.*, "Inverse design in nanophotonics," *Nature Photon.*, vol. 12, pp. 659–670, Nov. 2018.

[19] A. Michaels and E. Yablonovitch, "Leveraging continuous material averaging for inverse electromagnetic design," *Opt. Express*, vol. 26, pp. 31717–31737, Nov. 2018.

[20] A. Michaels, "Emopt: Electromagnetic optimization toolbox." 2018. [Online]. Available: https://github.com/anstmichaels/emopt

[21] R. M. Jenkins, R. W. J. Devereux, and J. M. Heaton, "Waveguide beam splitters and recombiners based on multimode propagation phenomena," *Opt. Lett.*, vol. 17, pp. 991–993, Jul. 1992.

[22] A. Michaels and E. Yablonovitch, "Inverse design of near unity efficiency perfectly vertical grating couplers," *Opt. Express*, vol. 26, pp. 4766–4779, Feb. 2018.

[23] Y. Zhang *et al.*, "A compact and low loss Y-junction for submicron silicon waveguide," *Opt. Express*, vol. 21, pp. 1310–1316, Jan. 2013.

[24] S. H. Tao *et al.*, "Cascade wide-angle Y-junction 1×16 optical power splitter based on silicon wire waveguides on silicon-on-insulator," *Opt. Express*, vol. 16, pp. 21456–21461, Dec. 2008.

[25] M. H. Tahersima *et al.*, "Deep neural network inverse design of integrated nanophotonic devices," *Sci. Rep.*, vol. 9, 2019, Art. no. 1368.

[26] Q. Wang, J. Lu, and S. He, "Optimal design method of a low-loss broadband Y branch with a multimode waveguide section," *Appl. Opt.*, vol. 41, pp. 7644–7649, Dec. 2002.

[27] L. H. Frandsen *et al.*, "Ultralow-loss 3-dB photonic crystal waveguide splitter," *Opt. Lett.*, vol. 29, pp. 1623–1625, Jul. 2004.

[28] Y. Wang, S. Gao, K. Wang, and E. Skafidas, "Ultra-broadband and low-loss 3 dB optical power splitter based on adiabatic tapered silicon waveguides," *Opt. Lett.*, vol. 41, pp. 2053–2056, May 2016.

[29] H. Wei, J. Yu, X. Zhang, and Z. Liu, "Compact 3-dB tapered multimode interference coupler in silicon-on-insulator," *Opt. Lett.*, vol. 26, pp. 878–880, Jun. 2001.

[30] K. Solehmainen, M. Kapulainen, M. Harjanne, and T. Aalto, "Adiabatic and multimode interference couplers on silicon-on-insulator," *IEEE Photon. Technol. Lett.*, vol. 18, no. 21, pp. 2287–2289, Nov. 2006.

[31] F. Wang, Y. Jianyi, C. Limei, J. Xiaoqing, and W. Minghua, "Optical switch based on multimode interference coupler," *IEEE Photon. Technol. Lett.*, vol. 18, no. 2, pp. 421–423, Jan. 2006.

[32] M. Yong, P. Seoijin, W. Liwei, and H. S. Tiong, "Ultracompact multimode interference 3-dB coupler with strong lateral confinement by deep dry etching," *IEEE Photon. Technol. Lett.*, vol. 12, no. 5, pp. 492–494, May 2000.

[33] H. Yun, W. Shi, Y. Wang, L. Chrostowski, and N. A. F. Jaeger, "2×2 adiabatic 3-dB coupler on silicon-on-insulator rib waveguides," *Proc. SPIE*, vol. 8915, 2013, Art. no. 89150V.

[34] J. Xing, Z. Li, Y. Yu, and J. Yu, "Design of polarization-independent adiabatic splitters fabricated on silicon-on-insulator substrates," *Opt. Express*, vol. 21, pp. 26729–26734, Nov. 2013.

[35] J. Xing *et al.*, "Silicon-on-insulator-based adiabatic splitter with simultaneous tapering of velocity and coupling," *Opt. Lett.*, vol. 38, pp. 2221–2223, Jul. 2013.

[36] K. Hassan, C. Durantin, V. Hugues, B. Szelag, and A. Glière, "Robust silicon-on-insulator adiabatic splitter optimized by metamodeling," *Appl. Opt.*, vol. 56, pp. 2047–2052, Mar. 2017.

[37] H.-C. Chung and S.-Y. Tseng, "Robust silicon 3-dB coupler using inverse engineering based optimization," *Jpn. J. Appl. Phys.*, vol. 57, Jun. 2018, Art. no. 08PC01.

[38] W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *Opt. Lett.*, vol. 32, pp. 2801–2803, Oct. 2007.

[39] P. J. Bock *et al.*, "Subwavelength grating crossings for silicon wire waveguides," *Opt. Express*, vol. 18, pp. 16146–16155, Jul. 2010.

[40] Y. Zhang, A. Hosseini, X. Xu, D. Kwong, and R. T. Chen, "Ultralow-loss silicon waveguide crossing using Bloch modes in index-engineered cascaded multimode-interference couplers," *Opt. Lett.*, vol. 38, pp. 3608–3611, Sep. 2013.

[41] Y. Ma *et al.*, "Ultralow loss single layer submicron silicon waveguide crossing for SOI optical interconnect," *Opt. Express*, vol. 21, pp. 29374–29382, Dec. 2013.

[42] Y. Zhang *et al.*, "A compact and low-loss silicon waveguide crossing for O-band optical interconnect," *Proc. SPIE*, vol. 8990, Mar. 2014, Art. no. 899002.

[43] Y. Liu, J. M. Shainline, X. Zeng, and M. A. Popović, "Ultra-low-loss CMOS-compatible waveguide crossing arrays based on multimode Bloch waves and imaginary coupling," *Opt. Lett.*, vol. 39, pp. 335–338, Jan. 2014.

[44] P. Dumais, D. J. Goodwill, D. Celo, J. Jiang, and E. Bernier, "Three-mode synthesis of slab Gaussian beam in ultra-low-loss in-plane nanophotonic silicon waveguide crossing," in *Proc. IEEE 14th Int. Conf. Group IV Photon.*, Aug. 2017, pp. 97–98.

Andrew Michaels received the B.S. degree in engineering physics from Cornell University, Ithaca, NY, USA, in 2014 and the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley in 2019.

Ming C. Wu (S'82–M'83–SM'00–F'02) received the B.S. degree in electrical engineering from National Taiwan University, Taipei City, Taiwan, in 1983 and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, Berkeley, CA, USA, in 1986 and 1988, respectively. From 1988 to 1992, he was a Member of Technical Staff with AT&T Bell Laboratories, Murray Hill. From 1993 to 2004, he was a Professor with the University of California, Los Angeles. Since 2005, he has been a Professor with the University of California, Berkeley, where he is currently a Nortel Distinguished Professor. He is also co-director of the Berkeley Sensor and Actuator Center. He is the author of more than 500 papers in journals and conferences and 29 U.S. patents. His research interests include optical micro-electro-mechanical systems (MEMS), MOEMS, optofluidics, silicon photonics, nanophotonics, and optoelectronics. He was a Packard Foundation Fellow (1992–1997). He was a member of the IEEE Photonics Society Board of Governors from 2013 to 2016. He was a recipient of the IEEE Photonics Socieity William Streifer Scientific Achievement Award in 2016, the Paul F. Forman Engineering Excellence Award in 2007, and the C.E.K. Mees Medal in 2017, both from the Optical Society.

Eli Yablonovitch received the Ph.D. degree in applied physics from Harvard University, Cambridge, MA, USA, in 1972. He is the Director of the NSF Center for Energy Efficient Electronics Science (E3S), a multi-University Center headquartered at Berkeley.

He introduced the idea that strained semiconductor lasers could have superior performance due to reduced valence band (hole) effective mass. With almost every human interaction with the Internet, optical telecommunication occurs by strained semiconductor lasers. In his photovoltaic research, he introduced the $4n^2$ ("Yablonovitch Limit") light-trapping factor, in worldwide use for almost all commercial solar panels. Based on his mantra that "a great solar cell also needs to be a great LED," his startup company Alta Devices, Inc., holds the world record for solar cell efficiency, now 29.1% at 1 sun. He is regarded as a Father of the Photonic BandGap concept, and he coined the term "Photonic Crystal." The geometrical structure of the first experimentally realized photonic bandgap, is sometimes called "Yablonovite." His startup company Ethertronics, Inc., shipped over two billion cellphone antennas.

He has been elected to the NAE, NAS, NAI, AmAcArS, and as Foreign Member, UK Royal Society. Among his honors are the OSA Ives/Quinn Medal, the Benjamin Franklin Medal, the IEEE Edison Medal, the Buckley Prize of the American Physical Society, and the Isaac Newton Medal of the U.K. Institute of Physics.