Highly Tailored Computational Electromagnetics Methods for Nanophotonic Design and Discovery

Dirichlet-to-Neumann maps in a domain decomposition scheme are discussed in this paper, as is the application of this method to rapid nanophotonic design optimization of a wavelength division multiplexer.

By Victor Liu, David A. B. Miller, *Fellow IEEE*, and Shanhui Fan, *Fellow IEEE*

ABSTRACT | The use of computational electromagnetics (CEM) techniques has greatly advanced nanophotonics. The applications of nanophotonics in turn motivates the development of efficient highly tailored algorithms for specific application domains. In this paper, we will discuss some specific considerations in seeking to advance CEM for nanophotonic design and discovery, with examples drawn from the design of aperiodic nanophotonic structures for on-chip information processing applications.

KEYWORDS | Aperiodic design; computational electromagnetics (CEM); nanophotonic design

I. INTRODUCTION

In nanophotonics, one seeks to discover new regimes of optical physics or to invent highly functional and compact optical devices by creating complex nanophotonic structures in which the dielectric properties of the structures vary strongly at wavelength or even deep subwavelength scales. Examples of nanophotonic structures include photonic crystals [1]–[4], plasmonic nanostructures [5]–[8], and metamaterials [9]–[11]. Moreover, whereas the majority of nanophotonic structures that have been explored are regular or periodic structures, there has been substantial recent interest in exploring aperiodic systems with even more structural complexity [12], [13].

The development of the field of nanophotonics has been greatly facilitated by the advances in computational electromagnetics (CEM), which is widely used for both device design and the understanding of the underlying physics. The CEM tools that have been used in nanophotonics include practically every major general-purpose CEM approach, such as finite-difference time-domain (FDTD) [14], finite-difference frequency-domain (FDFD) [15], finite-element method (FEM) [16], and rigorous coupled wave analysis (RCWA) [17] methods. However, the increasing complexity of structures is placing a greater demand on the speed and efficiency of computational approaches.

The need for efficiency is particularly critical both in exploring the physics of aperiodic complex systems and in systematic optimization of nanophotonic devices. From the engineering perspective, the fundamental problem is to design devices that satisfy a given specification, be it particular scattering profiles, transmission or reflection coefficients, or frequency dependencies. Here, device design is essentially a specification of the dielectric function over the entire domain of interest. Classically, the design of structures has been guided by physical intuition, usually combining basic elements such as waveguides and cavities in ways which are analyzable by simple mathematical models (e.g., coupled mode theory [18]). However, as both

the device structure and the performance criteria grow more complex, it becomes increasingly difficult to apply intuition-based design. Compared with the design of circuits, which from the electromagnetic perspective operates in the long wavelength regime, photonic design fundamentally differs due to the complicated and unavoidable coupling between design elements. Therefore, there is substantial value in the use of systematic optimization in an automated design process that does not require human intuition [19]–[24].

In this paper, we argue that a set of special-purpose CEM tools, which are specifically tailored to and optimized for a particular class of nanophotonic structures, may be of great value for nanophotonic design and discovery, provided that class is of large enough complexity to enable sufficient device performance. We illustrate the benefits of this approach by considering the design of an extremely compact photonic-crystal-based wavelength division multiplexer (WDM). We show that optimization of this system is greatly facilitated by the use of an ultraefficient algorithm that is specifically tailored for photonic crystal circuits. We also show that such efficient methods enable us to generate statistics regarding design parameter spaces, as illustrated by the example of a waveguide coupler. These statistical explorations have the potential to lead to the discovery of new physical laws governing large classes of systems.

We will present only 2-D simulation results for simplicity, but we note that these ideas generalize to 3-D. The study of 2-D structures is nonetheless relevant, as there have been significant experimental demonstrations of practically realizable 2-D systems [25], [26]. In addition, 2-D design optimization is often used to provide a reasonable starting point for 3-D design optimization. For example, the 2-D photonic crystal structures we present may be extended to 3-D using an effective index approximation, which is accurate to the order of a few percent for photonic crystal slabs [27].

II. WDM DESIGN AND CHOICE OF CEM METHOD

As an example to illustrate our numerical approach, we consider the design of an ultracompact photonic-crystal-based WDM. A wavelength division multiplexer, which separates signals from different wavelength or frequency channels, is of essential importance for optical communications. For on-chip applications, there is an additional need to design structures that occupy a small footprint.

Our WDM design is related to the theoretical framework outlined in [28]. Here, as shown in Fig. 1, an input waveguide is coupled to a large coupler region, which is subsequently coupled to three high-Q cavities that perform frequency selection, each of which is coupled to an output waveguide. An optimized structure is shown in Fig. 1(a),

with the coupler region highlighted in yellow, and each high-Q cavity formed by a single defect cell highlighted in red, green, and blue. The entire structure is embedded in a photonic crystal background composed of dielectric rods (black circles, radius $0.18a$ where a is the lattice constant, and $n = 3.4$). The transmission spectra for the optimized structure shown in Fig. 1(a) exhibit less than 5% reflection at each of the output frequencies.

This structure was designed starting from an initial structure shown in Fig. 1(b), which has transmission spectra that are far from ideal, after searching through approximately 73 000 structures. This design and its underlying physics were recently published in [29]. Here we focus on numerical and optimization aspects that enable such a design.

A. Conventional General Purpose Methods

The structure shown in Fig. 1 certainly can be simulated with general-purpose CEM tools. In fact, the most widely used numerical methods, such as FDTD and FEM, are capable of simulating very general classes of structures possessing arbitrary geometric complexity. This generality comes at a cost, however. In general-purpose CEM tools, the numerical representation of the field typically uses a grid that depends very little on the underlying physics of the structure. This translates into a relatively large number of unknowns, which typically results in the need to solve a large matrix system in frequency-domain simulations.

As a specific example, consider the structure shown in Fig. 1. This structure can be simulated with methods such as FDFD [15], where the electromagnetic fields are discretized on a Yee grid to set up a system matrix A. The field distribution (x) for a particular source configuration (b) can then be obtained by solving the linear system

$$Ax = b. \tag{1}$$

In the FDFD simulation of the structure shown in Fig. 1, one typically uses 20 grid points per lattice constant, or 400 grid points per unit cell of the photonic crystal, to obtain sufficient accuracy for the simulation. The structure in Fig. 1 then has 134 400 unknowns, and the system matrix A has a size of 134 400^2. Since A is sparse, solving the linear system in (1) can be performed straightforwardly on a standard workstation in a matter of minutes. Thus, to simulate a few structures, the standard general-purpose CEM tools are quite adequate.

Problems arise, however, when there is a need to simulate large numbers of structures. In our examples here, at the speed of 1 min per structure, simulating over 73 000 structures would take several months to complete. The optimization quickly becomes prohibitively expensive. In such situations, standard CEM tools are in fact not
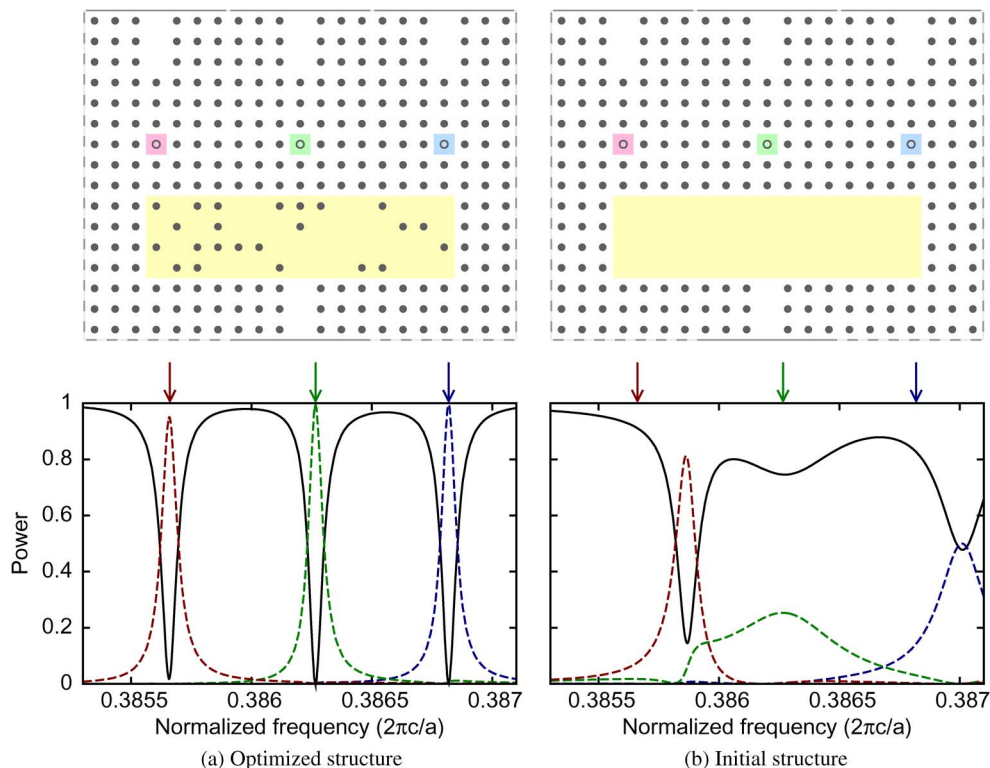
Fig. 1. *(a) Our optimized WDM structure. Filled circles indicate dielectric rods. The coupler region is highlighted in yellow and the cavities are highlighted in red, green, and blue with hollow circles. (b) The initial unoptimized structure. The response spectra for each structure are shown below the respective structure. The solid black lines are the reflection spectra in the input waveguide. The colored dashed lines are the transmission spectra for each output waveguide, with the colors matching the color of the cavity. The target frequencies are indicated by arrows.*

adequate at all, even if, at first glance, the numerical cost for simulating a single structure appears relatively modest.

B. Special-Purpose Methods

Based on the discussions outlined above, we see that there is a critical need to develop numerical algorithms that are suitable for efficiently simulating large numbers of structures. Our strategy is to develop a more efficient representation of the electromagnetic fields for the specific class of structures of interest. With a more efficient representation, the number of unknowns in the linear system can be significantly reduced, which certainly speeds up the solution of a single structure. Moreover, as we will see in Section III, reducing the number of unknowns enables the use of matrix algorithm techniques. As a result, once a structure is simulated, rigorous solution of subsequent structures can be done at a fraction of the computational cost compared with that of the initial structure.

For the photonic crystal structure shown in Fig. 1, we have chosen to use the method of Dirichlet-to-Neumann (DtN) maps. The method is also called the finite-element tearing and interconnecting (FETI) method [30], and more generally falls into the class of nonoverlapping domain decomposition methods [31]. We use a particular method

tailored for photonic crystals, the details of which can be found in [32]. Here we provide only a brief overview. In the DtN method, one only stores the fields at the edge of every unit cell of the photonic crystal, as illustrated in Fig. 2(a), as opposed to everywhere in the interior as in FDFD [Fig. 2(b)]. The system matrix, which essentially describes the coupling between these fields, is then set up in two steps. In the first step, one solves for the eigenmodes in an individual unit cell to obtain its DtN map. For the circular rod or hole features in our example, the DtN map can be computed using a cylindrical wave expansion. This step is a precomputing step; the DtN map is computed once and stored for each unit cell type.

In the second step, one then sets up the system matrix and solves the resulting linear system. For every structure that needs to be simulated, the matrix is set up by recalling the DtN maps precomputed for each individual unit cell of the structure, and by enforcing the field continuity conditions along each edge between adjacent unit cells.

The DtN method is specifically tailored for simulating photonic crystal devices with many repeated unit cells. In our optimization of the structure in Fig. 1, the coupler region consists of two types of unit cells: cells with or without a rod. As a result, to describe the coupler region,

Fig. 2. *(a) Layout of discretization points used in our DtN implementation. Only 5–7 field values are needed per edge, leading to about 10–14 unknowns per unit cell. Contrast with a typical FDFD discretization (b) using 20 points per direction of each unit cell leading to 400 unknowns per unit cell.*

only two different kinds of DtN maps need to be pre-computed and stored. For the structure in Fig. 1, we found that approximately ten unknowns per unit cell adequately represented the system, which is supported by the error analysis for a similarly sized transmission problem in [33] showing less than 10^{-4} relative error. We also show convergence of the resonance frequency of a defect cavity in Fig. 3 for an increasing number of field values per edge. For five unknowns per edge, the error in the resonance frequency is approximately 2×10^{-4}. The dimensions of the resulting system matrix are 3315^2 and are small enough such that A^{-1} can be stored in local and easily accessible memory, which is a capability of crucial importance for optimization purposes. In contrast, we note that for the same physical structure, the FDFD method results in an A^{-1} that is a dense matrix of a size of $134\,400^2$, which certainly cannot be stored in the same way.

Although our discussion centers around using DtN maps for numerical simulations, it is not the only fast method for photonic crystal circuit simulations. Other methods such as Wannier basis functions [34], [35] or multiple multipole methods [36], [37] have their own unique advantages, but the basic principles outlined here are equally applicable. We observe that all these special-purpose methods require significantly more precomputation time than the general purpose methods such as the FDFD. However, the speed advantage when iterating through many structures more than offsets this initial cost.

III. OPTIMIZATION ENABLED BY SPECIAL-PURPOSE METHODS

Optimization of nanophotonic structures is complex due to the large number of degrees of freedom usually present in a dielectric function specification. These optimization problems typically have nonsmooth objective function landscapes with many local minima. Any successful optimization strategy must be able to effectively handle such a difficult landscape. The combination of global and local optimization techniques is therefore essential. Here, in Sections III-A and B, we provide a general discussion of optimization techniques, and emphasize those aspects that are enabled by the use of special-purpose CEM tools as outlined in Section II. The optimization techniques that are specific to WDM filter design are then discussed in Section III-C.

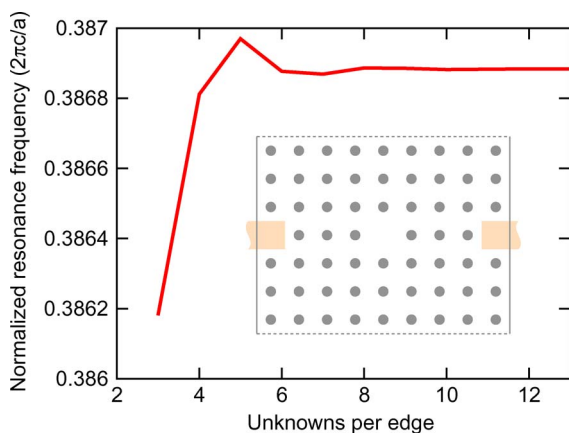
Fig. 3. *Convergence of the resonance frequency of a single rod defect cavity (shown in inset) with respect to the number of unknowns per unit cell edge. The structure is composed of dielectric rods in air with the same parameters as those in Fig. 1.*

A. Global Optimization

In a global optimization method, in each optimization iteration one makes a "drastic" structural change, in the

sense that the properties of the structure vary significantly after the change. The intent of a global optimization method is to explore as large a parameter space as possible, so that the design process does not get stuck in a local optimum that does not have sufficiently high performance.

In our WDM design problem, a change where a single rod was either added or removed from a lattice site (a "rod flip") in the coupler region can be considered a "drastic" structural change, since such a rod flip typically creates a significant change in the system response. Denoting the system matrices before and after the rod flip as A and $\tilde{A}$, respectively, since only a single unit cell is modified, $\tilde{A}$ is a low-rank modification of A. Because A^{-1} is available in its entirety, the matrix inversion lemma (2) can be used to efficiently compute low-rank updates for the updated structure. If we write the modification as $\tilde{A} = A + U\Delta V^T$, where Δ is a small square matrix (rank 20 for our structures), then

$$\tilde{A}^{-1} = A^{-1} - A^{-1}U(\Delta^{-1} + V^T A^{-1} U)^{-1} V^T A^{-1}. \quad (2)$$

It should be emphasized that treating A^{-1} as a black box solver without storing it in its entirety is impractical after several modifications to the system since the number of solves required to recursively apply (2) grows exponentially with respect to the number of modifications. Thus, the ability to store A^{-1} is crucial for enabling efficient solutions over large parts of design spaces.

With the ability to efficiently iterate through structures, many combinatorial optimization algorithms can be applied, the simplest being Monte Carlo sampling. In the context of our optimization problem, this would involve testing a sequence of structures where subsequent structures differ by a single rod flip chosen at random. More sophisticated methods such as simulated annealing and genetic algorithms can also benefit from the low-rank update technique discussed here.

B. Local Optimization

In contrast to the global optimization methods, in a local optimization method one modifies a current design by fine tuning it. The intent here is to locate the local optimum that is in the vicinity of the current design. Typically, a local optimization method involves calculation of the device sensitivity with respect to design parameters, and application of a gradient-based method that updates the device parameters based on the sensitivities.

We suppose that the system matrix A and the solution of the linear system x depend upon a particular structural parameter λ. (As an example, in our WDM structure, λ can be the dielectric constant of a single rod in the coupler region.) For a device performance metric J (e.g., the transmission coefficient through a single port), we are tasked with calculating its sensitivity $J' = dJ/d\lambda$ with respect to λ. J' can be straightforwardly calculated from $x' \equiv \partial x/\partial \lambda$ using the chain rule

$$J' = [\nabla_x J]x'. \quad (3)$$

On the other hand, typically $A' \equiv \partial A/\partial \lambda$ is known. Therefore, taking the derivative with respect to λ on either side of (1), and assuming for simplicity that $b' = 0$, we have

$$x' = -A^{-1}A'x. \quad (4)$$

Combining (3) and (4), one can then calculate the sensitivity of a device performance metric with respect to any structural parameter change. This method is commonly known as the adjoint variable method (AVM) [38].

We note that using the AVM requires only two applications of A^{-1} per design parameter: one to obtain x, and another to obtain x'. Thus, for a small number of design parameters, storing A^{-1} is not necessary. Many implementations of the AVM, for instance [15], simply solve the linear system of (1) several times. On the other hand, the ability to store A^{-1} means that sensitivities with respect to many design parameters can be calculated efficiently by only matrix–vector multiplication. To illustrate this capability, the sensitivity of rods within the coupler region of the final WDM design with respect to radius was computed, and is shown in Fig. 4. The intensity of the color of each rod is proportional to the sensitivity, with respect to the rod's radius, of the output power in the left-most waveguide of Fig. 1 at its resonance frequency. Red indicates a positive derivative (tending to increase the output power with increasing radius) and blue indicates a negative derivative. The entire computation took less than 2 seconds using a precomputed A^{-1}, whereas a computation from scratch took several minutes. Here again, we see the advantage of our special-purpose CEM approach which provides access to A^{-1}.

For a locally optimal design, the sensitivity of the device performance to design parameters is identically zero, and it might appear that there is no reason to perform sensitivity analysis. However, in the discretely constrained design space of the present example, it is almost guaranteed that any structure is not locally optimal. This discreteness is not inherently artificial, since any physically realizable structure is discretely constrained by the sizes of its constituent molecules. For nanophotonic structures, this atomic length scale is not negligible compared to the dimensions of structural features, and so there is still great value in computing sensitivities.

Fig. 4. *Sensitivity of transmission with respect to the radius of the dielectric rods within the yellow coupler region of Fig. 1(a). Greater red intensity in a rod indicates a greater tendency to increase the power through the left-most output waveguide in Fig. 1, and blue indicates a decrease, when the radius of the rod increases.*

C. Issues Specific to WDM Design

Sections III-A and B provided some of the general considerations of optimization in the context of nanophotonic design. In this section, we discuss some of the problems that are specific to the design of WDM filters.

For our choice of performance metric, mentioned in Section III-B, we define an error metric that describes the deviation in the performance of a candidate structure from the desired performance specifications. The error metric for each frequency was chosen to be the ℓ_2-norm of the difference between specified and simulated outgoing power in each waveguide

$$J_\omega(x) = \sum_{\substack{i-\text{th} \\ \text{waveguide}}} \frac{1}{2} \left| p_{\omega,i}^{\text{sim}}(x) - p_{\omega,i}^{\text{spec}}(x) \right|^2. \quad (5)$$

The total error metric is the sum of each single-frequency error metric

$$J(x) = \sum_\omega J_\omega(x). \quad (6)$$

This choice of error metric is attractive for its simplicity and differentiability, though other choices could have been used as well. In general, different types of design problems may require specific formulations of the error metric for the best results.

Motivated by its success in similar design problems [35], simulated annealing was used to optimize the WDM design by minimizing the error metric as rods were flipped. Simulated annealing is similar to a simple random search over the parameter space by flipping rods one at a time, except that each rod flip is kept with a certain probability instead of kept unconditionally (rod flips that are not kept are undone). In our implementation, we accept a rod flip unconditionally if the error metric of the new structure is strictly less than that of the previous structure considered. For flips which increase the error metric relative to the previous structure, we accept the rod

flip with a probability P exponentially distributed with respect to the error metric increase

$$P = \begin{cases} 1, & J_{\text{new}} < J_{\text{prev}} \\ \exp\left(-\dfrac{J_{\text{new}} - J_{\text{prev}}}{T}\right), & J_{\text{new}} \geq J_{\text{prev}} \end{cases} \quad (7)$$

where T is an effective "temperature" which is lowered as the optimization process progresses.

The design of frequency-selective devices presents unique subtleties for systematic optimization. For our device that contains high-Q cavities, global optimization iterations that flip rods in the coupler region tend to change the resonant frequency of these high-Q cavities by an amount that is comparable to the cavity linewidth. On the other hand, in actual applications, the filter needs to provide high transmission at frequencies that are dictated by a system specification. To illustrate this effect, Fig. 5
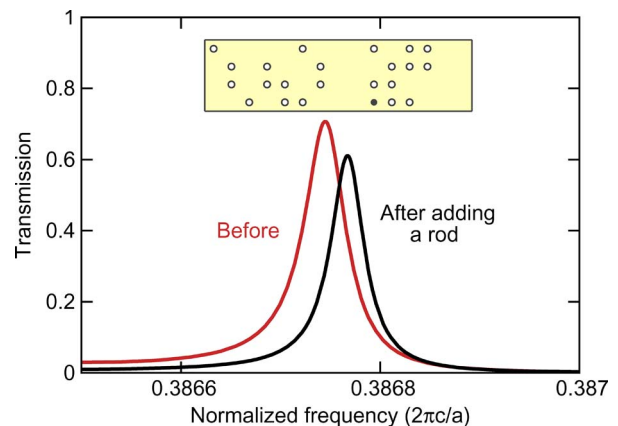

Fig. 5. *Illustration of the shift of high-Q cavity transmission spectra through the left-most waveguide of Fig. 1, after adding a rod to the coupler region of a structure encountered during the simulated annealing process. The structure is shown in the inset and the added rod is indicated by a solid circle. Notice that the resonance frequency shift is comparable to the resonance linewidth.*

shows the transmission spectra of two structures encountered in the optimization process which differ by a single rod flip. The evaluation of device performance at the specified output frequencies of the WDM is no longer accurate due to this shift. This need to lock operating frequencies to specified design frequencies appears in many design problems and merits special discussion.

Frequency locking requires the combination of global and local optimization strategies. In our example, we follow each rod flip by a subsequent gradient-descent iteration which adjusts the refractive index of the rods in the high-Q cavities in order to lock the transmission peak to the specified frequency locations. As mentioned in Section III-B, the gradients are readily available through the AVM, and the computation time is dominated by the repeated use of the low-rank updates as described in Section III-A. In our optimization, for each rod flip, the local gradient optimization requires us to further evaluate up to 16 additional structures. Although combining global and local optimization requires an order of magnitude more computation time for the same number of rod configurations in our example, the simulated annealing process for WDM design did not produce any operating device if the cavities were not tuned after each rod flip operation.

In usual applications of simulated annealing, the optimization is run until convergence is achieved, wherein the error metric cannot be reduced further and the structure no longer changes. In our example, on the other hand, full convergence was not strictly necessary to discover an acceptable design. The device shown in Fig. 1, among others, was encountered in the simulated annealing iteration well before it converged.

IV. STATISTICAL EXPLORATIONS

The ability to efficiently evaluate large numbers of structures also enables the collection of statistics over the parameter space. Such a capability is important from a fundamental perspective. There has been a substantial amount of work aiming to understand the statistics of transmission or reflection coefficients in an ensemble of random structures. On the other hand, in many cases, it has been quite difficult to obtain reliable numerical statistical data over a large enough ensemble. Knowing the statistics is also important from a practical point of view. We certainly would like to know how difficult it is to accomplish a certain engineering task given a particular set of structures. And, one way to measure the difficulty is to calculate the percentage of the structures in the set that actually function as desired.

We illustrate this capability of generating statistics by exploring a coupler device between photonic crystal waveguides of different widths. In practice, one is concerned with coupling of a single-mode photonic crystal waveguide with a width of half a wavelength, to a fiber mode that is many times larger. Moreover, one would like to perform such a mode coupling using a device that is as compact as possible and for which the coupling needs to occur over a sufficient bandwidth.

In our design problem, we consider a photonic crystal consisting of a square lattice of air holes. Both of the waveguides consist of rows of missing holes. We place an aperiodic coupler region in between, and we seek to design an efficient coupler by exploring structures where a hole is either present or missing at each lattice site in the coupler region. The goal is to efficiently couple the fundamental


(a) General structure   (b) Optimal structure

Fig. 6. *(a) Schematic of waveguide coupler devices embedded in a photonic crystal background of air holes ($r = 0.4a$) in dielectric ($n = 3.481$). The class of structures contains 7×5 lattice sites in the optimization region (yellow) and the waveguides are indicated by the curves on the left and right sides. (b) An optimized broadband waveguide coupler device design.*

Fig. 7. *Transmission spectrum for the structure shown in Fig. 6. The black curve is the transmitted power and the dashed curves correspond to reflected powers into the fundamental and next higher order mode of the wide input waveguide. The field pattern for $\omega = 0.2675 \times 2\pi c/a$ is shown in the inset.*

modes of both waveguides, therefore we consider structures that are symmetric in the coupler region.

For the class of coupler structures shown in Fig. 6(a) that consists of 7×5 lattice sites (yellow region), taking into account the symmetry, there are a total of 2^{21} such structures. The optimal structure in this design space is shown in Fig. 6(b). From the spectra for this structure shown in Fig. 7, we see that such a structure exhibits high transmission over a sizable bandwidth. The field distribution at a frequency at the center of the high transmission band is shown in the inset.

Using DtN maps with low-rank updates, the simulation time for a single coupler structure is well below a second, enabling us to generate statistics of the device performance for the entire design parameter space. Fig. 8(a) shows the distribution over all designs for a coupler region of 7×5, of the transmission coefficient between the two waveguides, over all designs for a frequency near the middle of the passband. It is interesting to note that in the entire design space, devices with the desired near unity transmission coefficient are exceedingly rare, and that the precise frequency selected for this histogram did not affect the shape of the distribution. Knowledge of these statistics is of use in engineering since parameters such as the size of a required optimization region can be estimated from these statistical trends.

Since a set of couplers with a smaller coupler region is a subset of the 7×5 coupler as considered above, we can plot the statistics for couplers that are smaller as well. The statistics for structures with 6×5 and 5×5 coupler regions are shown in Fig. 8(b) and (c), respectively. We may also consider the transmission coefficient from the next higher order symmetric mode of the wide waveguide to the single-mode waveguide. These statistics for the same

structures as in Fig. 6(a) are shown in Fig. 8(d). The most striking feature of all these data is the similarity in the shapes of the distributions. All of them feature a sharp peak at $T = 0$. Away from the peak, over a large range of transmission coefficients, the histograms appear linear on a log scale, with a sharp roll-off near $T = 1$. This suggests there may be some universal features in this space of structures.

The presence of universal features in the ensemble of transmission spectra through an ensemble of random media has been a subject of great interest in mesoscopic physics [39]–[41]. Here, we note that the universal features

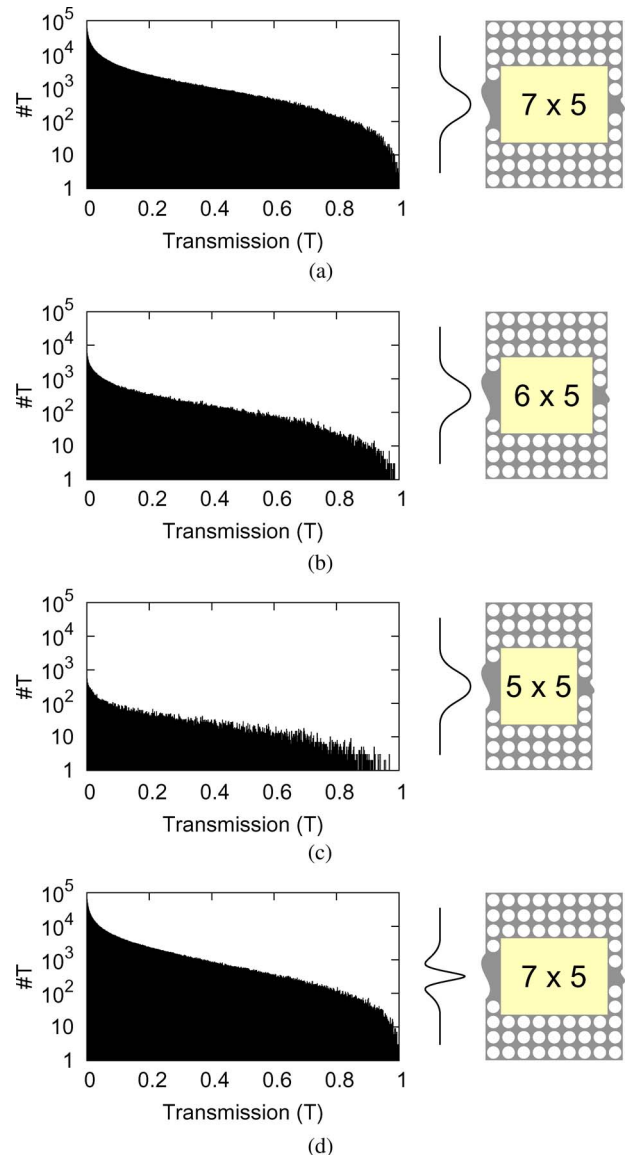**Fig. 8.** *(a) Transmission coefficient histogram for the fundamental mode of both waveguides for all structures in the design space of Fig. 6 at frequencies near the middle of the passband. (b) and (c) The same as (a) for a coupler region of size 6×5 and 5×5, respectively. (d) The same as (a) except for the transmission coefficient of the next higher order even mode of the wide waveguide.*

seen here in fact cannot be fit with the standard random matrix theory. The systematic exploration of large parameter spaces, as enabled by efficient numerical algorithms, may lead to a deeper understanding of mesoscopic physics.

V. FINAL REMARKS AND CONCLUSION

We have presented several examples illustrating the advantages of using a highly tailored numerical method for computational design and discovery problems in nanophotonics. Using a compact problem formulation that allows the inverse of the system matrix to be stored enables many global and local optimization methods as well as full design space explorations. These new capabilities greatly facilitate the design of novel classes of structures such as aperiodic systems, and are of potential interest to fundamental physics by enabling the analysis of ensembles of aperiodic systems. ∎

REFERENCES

[1] E. Yablonovitch. (1987, May). Inhibited spontaneous emission in solid-state physics and electronics. *Phys. Rev. Lett.* [Online]. *58(20)*, pp. 2059–2062. Available: http://link.aps.org/doi/10.1103/PhysRevLett.58.2059

[2] S. John. (1987, Jun.). Strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett.* *58(23)*, pp. 2486–2489. Available: http://link.aps.org/doi/10.1103/PhysRevLett.58.2486

[3] J. D. Joannopoulos, P. R. Villeneuve, and S. Fan. (1997). Photonic crystals: Putting a new twist on light. *Nature.* [Online]. *386(6621)*, pp. 143–149. Available: http://dx.doi.org/10.1038/386143a0

[4] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light*. Princeton, NJ: Princeton Univ. Press, 2008.

[5] S. A. Maier and H. A. Atwater. (2005). Plasmonics: Localization and guiding of electromagnetic energy in metal/dielectric structures. *J. Appl. Phys.* [Online]. *98(1)*. Available: http://link.aip.org/link/?JAP/98/011101/1

[6] E. Ozbay. (2006). Plasmonics: Merging photonics and electronics at nanoscale dimensions. *Science* [Online]. *311(5758)*, pp. 189–193. Available: http://www.sciencemag.org/content/311/5758/189.abstract

[7] H. Wang, D. W. Brandl, P. Nordlander, and N. J. Halas. (2007). Plasmonic nanostructures: Artificial molecules. *Accounts Chem. Res.* [Online]. *40(1)*, pp. 53–62. Available: http://pubs.acs.org/doi/abs/10.1021/ar0401045pMID: 17226945

[8] J. A. Schuller, E. S. Barnard, W. Cai, Y. C. Jun, J. S. White, and M. L. Brongersma. (2010). Plasmonics for extreme light concentration and manipulation. *Nature Mater.* [Online]. *9(3)*, pp. 193–204. Available: http://dx.doi.org/10.1038/nmat2630

[9] D. R. Smith, J. B. Pendry, and M. C. K. Wiltshire. (2004). Metamaterials and negative refractive index. *Science.* [Online]. *305(5685)*, pp. 788–792. Available: http://www.sciencemag.org/content/305/5685/788.abstract

[10] N. Engheta and R. Ziolkowski, *Metamaterials: Physics and Engineering Explorations*. Piscataway, NJ: Wiley/IEEE Press, 2006.

[11] N. Liu, H. Guo, L. Fu, S. Kaiser, H. Schweizer, and H. Giessen. (2008). Three-dimensional photonic metamaterials at optical frequencies. *Nature Mater.* [Online]. *7(1)*, pp. 31–37. Available: http://dx.doi.org/10.1038/nmat2072

[12] I. L. Gheorma, S. Haas, and A. F. J. Levi. (2004). Aperiodic nanophotonic design. *J. Appl. Phys.* [Online]. *95(3)*, pp. 1420–1426. Available: http://link.aip.org/link/?JAP/95/1420/1

[13] A. Gopinath, S. V. Boriskina, B. M. Reinhard, and L. D. Negro. (2009, Mar.). Deterministic aperiodic arrays of metal nanoparticles for surface-enhanced raman scattering (SERS). *Opt. Exp.* [Online]. *17(5)*, pp. 3741–3753. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-17-5-3741

[14] A. Taflove and S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method,* 3rd. Norwood, MA: Artech House, 2005.

[15] G. Veronis, R. W. Dutton, and S. Fan. (2004, Oct.). Method for sensitivity analysis of photonic crystal devices. *Opt. Lett.* [Online]. *29(19)*, pp. 2288–2290. Available: http://ol.osa.org/abstract.cfm?URI=ol-29-19-2288

[16] J.-M. Jin, *The Finite Element Method in Electromagnetics*. Chichester, U.K.: Wiley, 2002.

[17] M. G. Moharam, D. A. Pommet, E. B. Grann, and T. K. Gaylord. (1995, May). Stable implementation of the rigorous coupled-wave analysis for surface-relief gratings: Enhanced transmittance matrix approach. *J. Opt. Soc. Amer. A* [Online]. *12(5)*, pp. 1077–1086. Available: http://josaa.osa.org/abstract.cfm?URI=josaa-12-5-1077

[18] H. Haus and W. Huang. (1991, Oct.). Coupled-mode theory. *Proc. IEEE.* [Online]. *79(10)*, pp. 1505–1518. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=\&arnumber=104225\&isnumber=3217

[19] C. Y. Kao, S. Osher, and E. Yablonovitch. (2005). Maximizing band gaps in two-dimensional photonic crystals by using level set methods. *Appl. Phys. B.* [Online]. *81*, pp. 235–244, DOI: 10.1007/s00340-005-1877-3. Available: http://dx.doi.org/10.1007/s00340-005-1877-3

[20] J. S. Jensen and O. Sigmund. (2005, Jun.). Topology optimization of photonic crystal structures: A high-bandwidth low-loss t-junction waveguide. *J. Opt. Soc. Amer. B* [Online]. *22(6)*, pp. 1191–1198. Available: http://josab.osa.org/abstract.cfm?URI=josab-22-6-1191

[21] A. Gondarenko, S. Preble, J. Robinson, L. Chen, H. Lipson, and M. Lipson. (2006, Apr.). Spontaneous emergence of periodic patterns in a biologically inspired simulation of photonic structures. *Phys. Rev. Lett.* [Online]. *96(14)*. Available: http://link.aps.org/doi/10.1103/PhysRevLett.96.143904

[22] Y. Jiao, S. Fan, and D. Miller. (2006, Mar.). Systematic photonic crystal device design: Global and local optimization and sensitivity analysis. *IEEE J. Quantum Electron.* [Online]. *42(3)*, pp. 266–279. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=\&arnumber=1593859\&isnumber=33492

[23] A. Mutapcic, S. Boyd, A. Farjadpour, S. Johnson, and Y. Avniel. (2009, Apr.). Robust design of slow-light tapers in periodic waveguides. *Eng. Optim.* [Online]. *41(20)*, pp. 365–384. Available: http://www.ingentaconnect.com/content/tandf/geno/2009/00000041/00000004/art00005

[24] J. Lu, S. Boyd, and J. Vučković. (2011, May). Inverse design of a three-dimensional nanophotonic resonator. *Opt. Exp.* [Online]. *19(11)*, pp. 10 563–10 570. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-19-11-10563

[25] S.-Y. Lin, E. Chow, V. Hietala, P. R. Villeneuve, and J. D. Joannopoulos. (1998, Oct.). Experimental demonstration of guiding and bending of electromagnetic waves in a photonic crystal. *Science.* [Online]. *282(5387)*, pp. 274–276. Available: http://www.sciencemag.org/content/282/5387/274.abstract

[26] L. H. Gabrielli, J. Cardenas, C. B. Poitras, and M. Lipson. (2009, Jul.). Silicon nanostructure cloak operating at optical frequencies. *Nature Photon.* [Online]. *3(8)*, pp. 461–463. Available: http://dx.doi.org/10.1038/nphoton.2009.117

[27] M. Qiu. (2002). Effective index method for heterostructure-slab-waveguide-based two-dimensional photonic crystals. *Appl. Phys. Lett.* [Online]. *81(7)*, pp. 1163–1165. Available: http://link.aip.org/link/?APL/81/1163/1

[28] C. Jin, S. Fan, S. Han, and D. Zhang. (2003, Jan.). Reflectionless multichannel wavelength demultiplexer in a transmission resonator configuration. *IEEE J. Quantum Electron.* [Online]. *39(1)*, pp. 160–165. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=\&arnumber=1158819\&isnumber=25962

[29] V. Liu, Y. Jiao, D. A. B. Miller, and S. Fan. (2011, Feb.). Design methodology for compact photonic-crystal-based wavelength division multiplexers. *Opt. Lett.* [Online]. *36(4)*, pp. 591–593. Available: http://ol.osa.org/abstract.cfm?URI=ol-36-4-591

[30] M. Vouvakis, Z. Cendes, and J.-F. Lee, "A FEM domain decomposition method for photonic and electromagnetic band gap structures," *IEEE Trans. Antennas Propag.*, vol. 54, no. 2, pp. 721–733, Feb. 2006.

[31] A. Toselli and O. Widlund, *Domain Decomposition Methods-Algorithms and Theory*, New York: Springer-Verlag, 2005.

[32] Z. Hu and Y. Y. Lu. (2008). Efficient analysis of photonic crystal devices by Dirichlet-to-Neumann maps. *Opt. Exp.* [Online]. *16(22)*, pp. 17383–17399. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-16-22-17383

[33] Y. Huang, Y. Y. Lu, and S. Li. (2007). Analyzing photonic crystal waveguides by Dirichlet-to-Neumann maps. *J. Opt. Soc. Amer. B.* [Online]. *24(11)*, pp. 2860–2867. Available: http://www.opticsinfobase.org/josab/abstract.cfm?URI=josab-24-11-2860

[34] K. Busch, S. F. Mingaleev, A. Garcia-Martin, M. Schillinger, and D. Hermann. (2003). The

wannier function approach to photonic crystal circuits. *J. Phys., Condens. Matter* [Online]. *15(30)*, pp. R1233–R1256. Available: http://stacks.iop.org/0953-8984/15/i=30/a=201

[35] Y. Jiao, S. Fan, and D. A. B. Miller. (2005). Demonstration of systematic photonic crystal device design and optimization by low-rank adjustments: An extremely compact mode separator. *Opt. Lett.* [Online]. *30(2)*, pp. 141–143. Available: http://ol.osa.org/abstract.cfm?URI=ol-30-2-141

[36] E. Moreno, D. Erni, and C. Hafner. (2002, Sep.). Modeling of discontinuities in photonic crystal waveguides with the multiple multipole method. *Phys. Rev. E.* [Online].

66(3), p. 036618. Available: http://link.aps.org/doi/10.1103/PhysRevE.66.036618

[37] J. Smajic, C. Hafner, and D. Erni. (2003, Mar.). On the design of photonic crystal multiplexers. *Opt. Exp.* [Online]. *11(6)*, pp. 566–571. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-11-6-566

[38] N. Georgieva, S. Glavic, M. Bakr, and J. Bandler. (2002, Dec.). Feasible adjoint sensitivity technique for em design optimization. *IEEE Trans. Microw. Theory Tech.* [Online]. *50(12)*, pp. 2751–2758. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=\&arnumber=1097993\&isnumber=24077

[39] P. A. Lee and A. D. Stone. (1985, Oct.). Universal conductance fluctuations in metals. *Phys. Rev. Lett.* [Online]. *55(15)*, pp. 1622–1625. Available: http://link.aps.org/doi/10.1103/PhysRevLett.55.1622

[40] B. Kramer and A. MacKinnon. (1993). Localization: Theory and experiment. *Rep. Progr. Phys.* [Online]. *56(12)*, pp. 1469–1564. Available: http://stacks.iop.org/0034-4885/56/i=12/a=001

[41] C. W. J. Beenakker. (1997, Jul.). Random-matrix theory of quantum transport. *Rev. Mod. Phys.* [Online]. *69(3)*, pp. 731–808. Available: http://link.aps.org/doi/10.1103/RevModPhys.69.731

ABOUT THE AUTHORS

Victor Liu received the B.S. degree in electrical engineering from the California Institute of Technology (Caltech), Pasadena, in 2007 and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 2009, where he is currently working toward the Ph.D. degree in electrical engineering.

His current research interests include nanophotonics and computational electromagnetics.

David A. B. Miller (Fellow, IEEE) received the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 1979.

He was with Bell Laboratories from 1981 to 1996, as a Department Head from 1987. He is currently the W. M. Keck Professor of Electrical Engineering, and a Co-Director of the Stanford Photonics Research Center at Stanford University, Stanford, CA. His research interests include physics and devices in nanophotonics, nanometalics, and quantum-well optoelectronics, and fundamentals and applications of optics in information sensing, switching, and processing. He has published more than 230 scientific papers and the text *Quantum Mechanics for Scientists and Engineers* (Cambridge, U.K.: Cambridge Univ. Press, 2008), and holds 69 patents.

Dr. Miller has received numerous awards. He is a Fellow of the Optical Society of America (OSA), the American Physical Society (APS), and the Royal Societies of Edinburgh and London, and a Member of the National Academy of Sciences and the National Academy of Engineering . He holds two honorary degrees.

Shanhui Fan (Fellow, IEEE) received the Ph.D. degree in theoretical condensed matter physics from the Massachusetts Institute of Technology (MIT), Cambridge, in 1997.

He was a Research Scientist at the Research Laboratory of Electronics at MIT prior to his appointment at Stanford University, Stanford, CA, where he is currently an Associate Professor of Electrical Engineering. He has authored or coauthored more than 220 refereed journal articles, has given over 170 invited talks, and holds 39 U.S. patents. His current research interests include computational and theoretical studies of solid-state and photonic structures and devices, especially photonic crystals, microcavities, and nanophotonic circuits and elements.

Dr. Fan is a Fellow of the Optical Society of America (OSA), the American Physical Society (APS), and The International Society for Optical Engineers (SPIE). He received a National Science Foundation Career Award in 2002, a David and Lucile Packard Fellowship in Science and Engineering in 2003, the National Academy of Sciences Award for Initiative in Research in 2007, and the Adolph Lomb medal from OSA in 2007.