

Computational Bounds For Photonic Design

Guillermo Angeris Jelena Vučković Stephen Boyd

December 2018

Abstract

Physical design problems, such as photonic inverse design, are typically solved using local optimization methods. These methods often produce what appear to be good or very good designs when compared to classical design methods, but it is not known how far from optimal such designs really are. We address this issue by developing methods for computing a bound on the true optimal value of a physical design problem; physical designs with objective smaller than our bound are impossible to achieve. Our bound is based on Lagrange duality and exploits the special mathematical structure of these physical design problems. For a multi-mode 2D Helmholtz resonator, numerical examples show that the bounds we compute are often close to the objective values obtained using local optimization methods, which reveals that the designs are not only good, but in fact nearly optimal. Our computational bounding method also produces, as a by-product, a reasonable starting point for local optimization methods.

1 Introduction

Computer-aided design of physical systems is growing rapidly in several fields, including photonics [MLP⁺18] (where it is known as inverse design), horn design [NUS⁺10], and mechanical design (aerospace, structures) [HG12]. These design methods formulate the physical design problem as a constrained nonconvex optimization problem, and then use local optimization to attempt to solve the problem. Commonly used methods include gradient descent, with adjoint-based evaluations of the gradient [LKBY13], methods that alternate optimizing over the structure and over the response [LV10], and the alternating directions method of multipliers (ADMM) [LV13], among others. These methods can be very effective, in the sense of producing what appear to be very good physical designs, for example when compared to classical design approaches.

Because they are local optimization methods, they do not guarantee that a globally optimal design is found, nor do we know how far from optimal the resulting design is. This paper addresses the question of how far a physical design is from globally optimal by computing a lower bound on the optimal objective value of the optimization problem. A

lower bound on the objective value can be interpreted as an impossibility result since it asserts that no physical design can have a lower objective than a number we compute.

Our bound is similar in spirit to analytical lower bounds, which give lower bounds as simple formulas in terms of gross quantities like temperature and wavelength, based on very simplified models and objectives, *e.g.*, the Reynolds number [Pur77], the Carnot efficiency limit [Fer36, §3.8], or the optical diffraction limit [BW13, §8.6]. There has been some additional work in bounding some other quantities and figures of merit for optical systems, including the local density of states [MPR⁺16, SFJM18] for different types of materials, via fundamental physical principles. In contrast, our method computes a (numerical) lower bound for the optimization objective for each design problem.

In this paper, we derive a parametrized family of lower bounds on the optimal objective for a class of physical design problems, using Lagrange duality. We can optimize over the parameter, to obtain the best (largest) lower bound, by solving the Lagrange dual problem—which is convex even though the original design problem is not. We illustrate our lower bound on a two-dimensional multi-mode resonator. Our lower bound is close to the objective obtained by a design using ADMM, which shows that the design, and indeed our lower bound, are both very close to the global optimum.

2 Physical design

2.1 Physical design problem

In physical design, we design a structure so that the field, under a given excitation, is close to some desired or target field. We parametrize the structure using a vector θ , and we denote the field by the vector z . In photonic design, for example, we choose the index of refraction at each rectangle on a grid, within limits, to achieve or get close to a desired electromagnetic field.

We can express this as the following optimization problem:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && 0 \leq \theta \leq \theta^{\max}, \end{aligned} \tag{1}$$

with variables $z \in \mathbf{R}^n$ (the field) and $\theta \in \mathbf{R}^n$, which describes the physical design. The data are the weight matrix $W \in \mathbf{R}^{n \times n}$, which is diagonal with positive diagonal entries, the desired or target field $\hat{z} \in \mathbf{R}^n$, the matrix $A \in \mathbf{R}^{n \times n}$, the excitation vector $b \in \mathbf{R}^n$, and the vector θ^{\max} of limits on the physical design parameter θ . The constraint equation $(A + \mathbf{diag}(\theta))z = b$ encodes the physics of the problem. We let p^* denote the optimal value of (1).

We can handle the case when the lower limit on the physical parameter is nonzero, for example, $\theta^{\min} \leq \theta \leq \theta^{\max}$. We do this by replacing the lower limit by 0, the upper limit by

$\theta^{\max} - \theta^{\min}$, and replacing A with $A + \mathbf{diag}(\theta^{\min})$. Additionally, the construction extends easily to the case where the field z , the matrix A , and the excitation b have complex entries.

When the coefficient matrix in the physics equation $(A + \mathbf{diag}(\theta))z = b$ is nonsingular, there is a unique field, $z = (A + \mathbf{diag}(\theta))^{-1}b$. In some applications, however, the coefficient matrix is singular, and there is either no field that satisfies the equations, or many. In the former case, we take the objective to be $+\infty$. In the latter case, the set of solutions is an affine set and simple least squares can be used to find the field that satisfies the physics equation and minimizes the objective.

An important special case occurs when we seek a mode (eigenvector) of a system that is close to \hat{z} . To do this we take $b = 0$ and subtract λI from the coefficient matrix, where λ is the required eigenvalue. We can handle the case of unspecified eigenvalues by a simple extension described later in problem (14), where λ also becomes a design variable, subject to a lower and upper bound.

In the problem (1), the physical design parameters enter in a very specific way: as the diagonal entries of the coefficient matrix of the physics equation. Many physics equations have this form for a suitable definition of the field z and parameter θ , including the time-independent Schrödinger equation, Helmholtz's equation, the heat equation, and Maxwell's equations in one dimension. (Maxwell's equations in two and three dimensions are included in this formalism via the simple extension given in problem (13).)

Boolean physical design problem. A variation on the problem (1) replaces the physical parameter constraint $0 \leq \theta_j \leq \theta_j^{\max}$ with the constraint $\theta_j \in \{0, \theta_j^{\max}\}$, which limits each physical parameter value to only two possible values. (This occurs when we are choosing between two materials, such as silicon or air, in each of the patches in the structure we are designing.) We refer to this modified problem as the *Boolean physical design problem*, as opposed to the continuous physical design problem (1). It is clear that the optimal value of the Boolean physical design is no smaller than p^* , the optimal value of the continuous physical design problem.

2.2 Approximate solutions

The problem (1) is not convex and generally hard to solve exactly [BV04]. It is, however, bi-convex, since it is convex in z when θ is fixed, and convex in θ when z is fixed. Using variations on this observation, researchers have developed a number of methods for approximately solving (1) via heuristic means, such as alternating optimization over z and θ on the augmented Lagrangian of this problem [LV13]. Other heuristics can be used to find approximate solutions of the Boolean physical design problem. These methods produce what appear to be very good physical designs when compared to previous hand-crafted designs or classical designs.

2.3 Performance bounds

Since the approximate solution methods used are local and therefore heuristic, the question arises: how far are these approximate designs from an optimal design? In other words, how far is the objective found by these methods from p^* ? Suppose, for example, that a heuristic method finds a design with objective value 13.1. We do not know what the optimal objective p^* is, other than $p^* \leq 13.1$. Does there exist a design with objective value 10? Or 5? Or are these values of the objective impossible, *i.e.*, smaller than p^* ?

The method described in this paper aims to answer this question. Specifically, we will compute a provable lower bound L on the optimal objective value p^* of (1). In our example above, our method might compute the lower bound value $L = 12.5$. This means that no design can ever achieve an objective value smaller than 12.5. It also means that a design with an objective value of 13.1 is not too far from optimal, since we would know that $L = 12.5 \leq p^* \leq 13.1$.

A lower bound L on p^* can be interpreted as an *impossibility result*, since it tells us that it is impossible for a physical design to achieve an objective value less than L . We can also interpret L as a *performance bound*. The lower bound L does not tell us what p^* is; it just gives a lower limit on what it can be. (An upper limit U can be found by using any heuristic method, as the final objective value attained.)

We note that the lower bound L we find on p^* also serves as a lower bound on the optimal value of the Boolean physical design problem, since its optimal value is larger than or equal to p^* .

3 Performance bounds via Lagrange duality

In this section, we explain our lower bound method.

3.1 Lagrangian duality

We first rewrite (1) as

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 + I(\theta) \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b, \end{aligned} \tag{2}$$

where I is an indicator function, *i.e.*, $I(\theta) = 0$ when $0 \leq \theta \leq \theta^{\max}$ and $+\infty$ otherwise. The Lagrangian of this problem is

$$\mathcal{L}(z, \theta, \nu) = \frac{1}{2} \|W(z - \hat{z})\|_2^2 + I(\theta) + \nu^T((A + \mathbf{diag}(\theta))z - b), \tag{3}$$

where $\nu \in \mathbf{R}^n$ is a dual variable. The Lagrange dual function is

$$g(\nu) = \inf_{\theta, z} \mathcal{L}(z, \theta, \nu).$$

(See [BV04, Chapter 5].) It is a basic and easily proved fact that for any ν , we have $g(\nu) \leq p^*$ (see [BV04, §5.1.3]). In other words, $g(\nu)$ is a lower bound on p^* . While $g(\nu)$ always gives a lower bound on p^* , the challenge for nonconvex problems such as (1) is to *evaluate* $g(\nu)$. We will see now that this can be done for our problem (1).

3.2 Evaluating the dual function

To evaluate $g(\nu)$ we must minimize $\mathcal{L}(z, \theta, \nu)$ over z and θ . Since for each θ , $\mathcal{L}(z, \theta, \nu)$ is convex quadratic in z , we can analytically carry out the minimization over z . We have

$$\begin{aligned} g(\nu) &= \inf_{\theta} \inf_z \mathcal{L}(z, \theta, \nu) \\ &= \inf_{\theta} \left(-\frac{1}{2} \|W^{-1}((A + \mathbf{diag}(\theta))^T \nu - W^2 \hat{z})\|_2^2 - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2 + I(\theta) \right) \\ &= \inf_{0 \leq \theta \leq \theta^{\max}} -\frac{1}{2} \|W^{-1}((A + \mathbf{diag}(\theta))^T \nu - W^2 \hat{z})\|_2^2 - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2. \end{aligned} \quad (4)$$

We can see that this is true since the minimizer of the only terms depending on z ,

$$\operatorname{argmin}_z \left(\frac{1}{2} \|W(z - \hat{z})\|_2^2 + \nu^T (A + \mathbf{diag}(\theta))z \right),$$

can be found by taking the gradient and setting it to zero (which is necessary and sufficient by convexity and differentiability). This gives that the minimizing z is

$$z = \hat{z} - W^{-2}(A + \mathbf{diag}(\theta))\nu, \quad (5)$$

which yields (4) when plugged in.

The expression in (4) is separable over each θ_i ; it can be rewritten as

$$\begin{aligned} g(\nu) &= \inf_{0 \leq \theta \leq \theta^{\max}} -\frac{1}{2} \sum_{j=1}^n W_{jj}^{-2} ((A^T \nu)_j + \nu_j \theta_j - W_{jj}^2 \hat{z}_j)^2 - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2 \\ &= \sum_{j=1}^n \left(\inf_{0 \leq \theta_j \leq \theta_j^{\max}} -\frac{1}{2} W_{jj}^{-2} ((A^T \nu)_j + \nu_j \theta_j - W_{jj}^2 \hat{z}_j)^2 \right) - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2 \\ &= -\frac{1}{2} \sum_{j=1}^n W_{jj}^{-2} \max \left\{ (a_j^T \nu - W_{jj}^2 \hat{z}_j)^2, (a_j^T \nu + \nu_j \theta_j^{\max} - W_{jj}^2 \hat{z}_j)^2 \right\} - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2, \end{aligned} \quad (6)$$

where a_j is the j th column of A . In the last line, we use the basic fact that a scalar convex quadratic function achieves its maximum over an interval at the interval's boundary.

With this simple expression for the dual function, we can now generate lower bounds on p^* , by simply evaluating it for any ν . We note that g is also the dual function of the Boolean physical design problem.

3.3 Dual optimization problem

It is natural to seek the best or largest lower bound on p^* , by choosing ν that maximizes our lower bound. This leads to the dual problem (see [BV04, §5.2]),

$$\text{maximize } g(\nu),$$

with variable ν . We denote the optimal value as d^* , which is the best lower bound on p^* that can be found from the Lagrange dual function. The dual problem is always a convex optimization problem (see [BV04, §5.1.2]); to effectively use it, we need a way to tractably maximize g , which we have in our case, since the dual problem can be expressed as the convex quadratically-constrained quadratic program (QCQP)

$$\begin{aligned} & \text{maximize} && -(1/2)\mathbf{1}^T t - \nu^T b + (1/2) \|W\hat{z}\|_2^2 \\ & \text{subject to} && t_j \geq W_{jj}^{-2} (a_j^T \nu - W_{jj}^2 \hat{z}_j)^2, \quad j = 1, \dots, n \\ & && t_j \geq W_{jj}^{-2} (a_j^T \nu + \nu_j \theta_j^{\max} - W_{jj}^2 \hat{z}_j)^2, \quad j = 1, \dots, n, \end{aligned} \tag{7}$$

with variables t and ν . This problem is easily solved and its optimal value, d^* , is a lower bound on p^* .

The dual optimization problem (7) can be solved several ways, including via ADMM (which can exploit the fact that all subproblems are quadratic; see [BPC⁺11a]), interior point methods (see [BV04, §11.1]), or by rewriting it as a second-order cone program (SOCP) (see [LVBL98]; this can also be done automatically by modeling languages such as CVXPY [AVDB18]) and then using one of the many available SOCP solvers, such as SCS [OCPB16a, OCPB16b], ECOS [DCB13], or Gurobi [GO18]. We also note that the dual problem does not have to be perfectly solved; we get a lower bound for *any* value of the dual variable ν .

In this paper, we used the Gurobi solver to solve a (sparse) program with $n = 63001$, which took approximately 8 minutes to solve on a two-core Intel Core i5 machine with 8GB of RAM. By further exploiting the structure of the problem, giving good initializations, or by using less accurate methods when small tolerances are not required, it is likely that these problems could be solved even more quickly, for larger systems.

3.4 Initializations via Lagrange dual

The solution of the Lagrange dual problem can be used to suggest starting points in a heuristic or local method for approximately solving (1).

Initial structure. Let ν^* be a solution of the dual problem (7). We can take as initial structure θ^0 which minimizes (4), *i.e.*,

$$\theta_j^0 \in \operatorname{argmax}_{\theta_j \in \{0, \theta_j^{\max}\}} (a_j^T \nu^* + \nu_j^* \theta_j - W_{jj}^2 \hat{z}_j)^2.$$

This choice of initial structure is feasible for (1) and, in fact, is feasible for the Boolean physical design problem as well.

Initial field. One way to obtain an initial field is to simply solve the physics equation for θ^0 , when the physics coefficient matrix is nonsingular. When it is singular, but the physics equation is solvable, we compute z as the field that minimizes the objective, subject to the physics equation. This gives a feasible field, but in some cases the resulting point is not very useful. For example when $b = 0$, and the coefficient matrix is nonsingular, we obtain $z^0 = 0$.

Another possibility is to find the minimizer of the Lagrangian with the given structure and an optimal dual variable value, *i.e.*,

$$z^0 = \underset{z}{\operatorname{argmin}} \mathcal{L}(\theta^0, z, \nu^*).$$

The value is already given in (5):

$$z^0 = \hat{z} - W^{-2} (A + \mathbf{diag}(\theta^0))^T \nu^*.$$

This initial field is not feasible, *i.e.*, it does not satisfy the physics equation, but it seems to be a very good initial choice for heuristic algorithms.

4 Multi-scenario design

In this section we mention an extension of our basic problem (1), in which we wish to design one physical structure that gives reasonable performance in N different scenarios. The scenarios can represent different operating temperatures, different frequencies, or different modes of excitation.

We will index the scenarios by the superscript i , with $i = 1, \dots, N$. Each scenario can have a different weight matrix W^i , a different target field \hat{z}^i , a different physics matrix A^i , and a different excitation b^i . We have only one physical design variable θ , and N different field responses, z^i , $i = 1, \dots, N$. We take as our overall objective the sum (or average) of the objectives under the scenarios. This leads to the problem

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \sum_{i=1}^N \|W^i(z^i - \hat{z}^i)\|_2^2 \\ & \text{subject to} && (A^i + \mathbf{diag}(\theta))z^i = b^i, \quad i = 1, \dots, N \\ & && 0 \leq \theta \leq \theta^{\max}, \end{aligned} \tag{8}$$

with variables θ (the structure) and z^i (the fields under the N different scenarios).

Our bounding method easily generalizes to this multi-scenario physical design problem.

Dual optimization problem. As before, define a_j^i to be the j th column of A^i and allow ν^i to be the Lagrange multiplier for the i th constraint, then the new dual problem is,

$$\begin{aligned} \text{maximize} \quad & -(1/2)\mathbf{1}^T t - \sum_{i=1}^N (\nu^i)^T (b^i) + (1/2) \sum_{i=1}^N \|W^i \hat{z}^i\|_2^2 \\ \text{subject to} \quad & t_j \geq \sum_{i=1}^N (W_{jj}^i)^{-2} \left((a_j^i)^T \nu^i - (W_{jj}^i)^2 \hat{z}_j^i \right)^2, \quad j = 1, \dots, n \\ & t_j \geq \sum_{i=1}^N (W_{jj}^i)^{-2} \left((a_j^i)^T \nu^i + \nu_j^i \theta_j^{\max} - (W_{jj}^i)^2 \hat{z}_j^i \right)^2, \quad j = 1, \dots, n, \end{aligned} \quad (9)$$

which is also a convex QCQP. This new dual optimization problem can be derived in a similar way to the construction of §3.

Initial structure and fields. Similar initializations hold for (8) as do for (1). We can find an initial θ^0 given by

$$\theta_j^0 \in \operatorname{argmax}_{\theta_j \in \{0, \theta_j^{\max}\}} \left(\sum_i (W_{jj}^i)^{-2} \left((a_j^i)^T (\nu^i)^* + (\nu^i)_j^* \theta_j - (W_{jj}^i)^2 \hat{z}_j^i \right)^2 \right), \quad (10)$$

while we can find feasible initial fields by solving the physics equations for each scenario, or as the minimizer of the Lagrangian,

$$(z^i)^0 = \hat{z}^i - (W^i)^{-2} (A^i + \mathbf{diag}(\theta^0))^T (\nu^i)^*, \quad (11)$$

for $i = 1, \dots, N$, which gives infeasible fields (often, however, these fields are good initializations).

5 Numerical example

5.1 Physics and discretization

We begin with Helmholtz's equation in two dimensions,

$$\nabla^2 f(x, y) + \left(\frac{\omega}{c(x, y)} \right)^2 f(x, y) = 0, \quad (12)$$

where $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ is a function representing the wave's amplitude, $\nabla^2 = \partial_x^2 + \partial_y^2$ is the Laplacian in two dimensions, $\omega \in \mathbf{R}_+$ is the angular frequency of the wave, and $c : \mathbf{R}^2 \rightarrow \mathbf{R}_+$ is the speed of the wave in the material at position (x, y) , which we can change by an appropriate choice of material. For this problem, we will allow the choice of any material that has a propagation speed between $0 < c^{\min}(x, y) \leq c(x, y) \leq c^{\max}(x, y)$, such that f is close to \hat{f} , some desired field.

Throughout, we will also assume Dirichlet boundary conditions for convenience (that is, $f(x, y) = 0$, whenever (x, y) is on the boundary of the domain), though any other boundary conditions could be similarly used with this method.

We discretize each of c , f , and ∇^2 in equation (12) using a simple finite-difference approximation over an equally-spaced rectilinear grid. (More sophisticated discretization methods would also work with our method.) Specifically, let (x_i, y_i) for $i = 1, \dots, n$ be the discretized points of the grid, with separation distance h (e.g., $y_{i+1} - y_i = x_{i+1} - x_i = h$). We then let z and \hat{z} , both in \mathbf{R}^n , be the discretization of f and \hat{f} , respectively, over the grid,

$$z_i = f(x_i, y_i), \quad \hat{z}_i = \hat{f}(x_i, y_i).$$

Using this discretization, we can approximate the second derivative of f at the grid points as,

$$\partial_x^2 f(x_i, y_i) \approx \frac{f(x_i + h, y_i) - 2f(x_i, y_i) + f(x_i - h, y_i)}{h^2} = \Delta_x z,$$

for some matrix Δ_x , and similarly for ∂_y^2 , whose finite approximation we will call Δ_y . We can then define a complete approximate Laplacian as the sum of the two matrices,

$$\Delta = \Delta_x + \Delta_y.$$

We also similarly discretize $c(x, y)$ as

$$\theta_i = \frac{1}{c(x_i, y_i)^2},$$

where $\theta \in \mathbf{R}^n$. The constraints on $c(x, y)$ become

$$\theta^{\min} = \frac{1}{c^{\max}(x_i, y_i)^2} \leq \theta_i \leq \frac{1}{c^{\min}(x_i, y_i)^2} = \theta^{\max}.$$

We can now write the fully-discretized form of Helmholtz's equation as

$$(\Delta + \omega^2 \mathbf{diag}(\theta))z = 0$$

or, equivalently,

$$\left(\frac{1}{\omega^2} \Delta + \mathbf{diag}(\theta) \right) z = 0.$$

So the final problem is, after replacing θ with $\theta - \theta^{\min}$,

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && ((1/\omega^2)\Delta + \mathbf{diag}(\theta^{\min}) + \mathbf{diag}(\theta)) z = 0 \\ & && 0 \leq \theta \leq \theta^{\max} - \theta^{\min}. \end{aligned}$$

This has the form of problem (1), with

$$A = \frac{1}{\omega^2} \Delta + \mathbf{diag}(\theta^{\min}), \quad b = 0.$$

Note that the design we are looking for—one that supports non-vanishing modes at each frequency—will, in general, have a singular (or indeterminate) physics equation. More specifically, the final design's physics equations will each have a linear set of solutions, from which we pick the one that minimizes the least squares residual in the objective.




Figure 1: The three target resonator regions.

5.2 Problem data

In this example, we will design a 2D resonator with modes that are localized in the boxes found in figure 1, at each of three specified frequencies. More specifically, let S^i be the indices at frequency i corresponding to the boxes shown in figure 1. We define the target field for frequency i as

$$\hat{z}_j^i = \begin{cases} 1, & j \in S^i \\ 0, & j \notin S^i. \end{cases}$$

We set the weights within the box containing the mode to be one and set those outside the box to be larger:

$$W_{jj}^i = \begin{cases} 1, & j \in S^i \\ 5, & j \notin S^i. \end{cases}$$

We specify three frequencies (*i.e.*, $N = 3$),

$$\omega = (30\pi, 40\pi, 50\pi),$$

at which to generate the specified modes by picking the propagation speed of the wave at each discretization point of the domain. We constrain the allowed propagation speed by picking

$$\theta_j^{\min} = 1, \quad \theta_j^{\max} = 2, \quad j = 1, \dots, n.$$

Our discretization uses a 251×251 grid, so $n = 251^2 = 63001$, with $h = 1/n$.

5.3 Physical design

We use ADMM to approximately solve the physical design problem, as in [LV13], using penalty parameter $\rho = 100$. We initialized the method using the feasible structure and fields

from §5.4, though similar designs are achieved with simple initializations like $\theta = \theta^{\min}$ and $z^i = 0$, for $i = 1, 2, 3$. We stop the algorithm when the physics constraint residual norm drops below a fixed tolerance of 10^{-2} . The resulting locally optimized design is shown in figure 2 and the associated fields are shown in figure 3. In particular, after local optimization, we receive some θ and z with

$$\theta^{\min} \leq \theta \leq \theta^{\max}, \quad \|(A + \mathbf{diag}(\theta))z - b\|_2 \leq 10^{-2},$$

and then evaluate

$$p = \frac{1}{2} \sum_{i=1}^3 \|W^i(z^i - \hat{z}^i)\|_2^2,$$

which gives $p = 5145$.

Our non-optimized implementation required around 1.5 seconds per iteration and took 332 iterations to converge to the specified tolerance, so the total physical design time is a bit under 9 minutes on a 2015 2.9GHz dual core MacBook Pro. Our implementation used a sparse-direct solver; an iterative CG solver with warm-start would have been much faster.

5.4 Dual problem

We solved problem (9) using the Gurobi [GO18] SOCP solver and the JuMP [DHL17] mathematical modeling language for Julia [BEKS17]. Gurobi required under ten minutes to solve the dual problem, about the same time required by the physical design. This time, too, could be very much shortened; for example, we do not need to solve the dual problem to the high accuracy that Gurobi delivers.

The optimal dual value found is $d^* = 4733$, with the initial design and fields suggested by the optimal dual solution shown in figure 2 and figure 3, respectively. This tells us that

$$4733 = d^* \leq p^* \leq p = 5145,$$

which implies that our physical design objective value is no more than $(5145 - 4733)/4733 \approx 8.7\%$ suboptimal. (We strongly suspect that p^* is closer to our design's value, 5145, than the lower bound, 4733.)

6 Further extensions

There are several straightforward extensions of the above problem, which may yield useful results in specific circumstances. All of these problems have analytic forms for their Lagrange dual functions, and all forms generalize easily to their multi-frequency counterparts. Additionally, we explicitly derive the dual functions for some extensions which require a little more care.




Figure 2: *Left.* Initial design suggested by the dual solution. *Right.* Optimized physical design.

Equality-constrained parameters. Sometimes, it might be the case that a single design parameter might control several points in the domain of z —for example, in the case of Maxwell’s equations in two and three dimensions (see the appendix for more details), or when the domain’s grid size is much smaller than the smallest features that can be constructed.

Let $S_k \subseteq \{1, \dots, n\}$ for $k = 1, \dots, m$ be a partition of indices, $\{1, \dots, n\}$. In other words, we want S_k for $k = 1, \dots, m$ to satisfy,

$$\bigcup_{k=1}^m S_k = \{1, \dots, n\}$$

and $S_k \cap S_l = \emptyset$ whenever $k \neq l$. These sets S_k will indicate the sets of indices which are constrained to be equal—conversely, indices that are not constrained to be equal to any other indices are represented by singleton sets.

We can then write the new optimization problem as

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && \theta_i = \theta_j, \text{ for all } i, j \in S_k, \quad k = 1, \dots, m \\ & && 0 \leq \theta \leq \theta^{\max}. \end{aligned} \tag{13}$$

To compute the Lagrange dual, let I' be an indicator function with $I'(\theta) = 0$ whenever $0 \leq \theta \leq \theta^{\max}$ and $\theta_i = \theta_j$ for all $i, j \in S_k$ for $k = 1, \dots, m$. Otherwise, $I'(\theta) = +\infty$. We can write the new problem as problem (2) with the same Lagrangian as the one given in (3), replacing I with I' in both expressions.

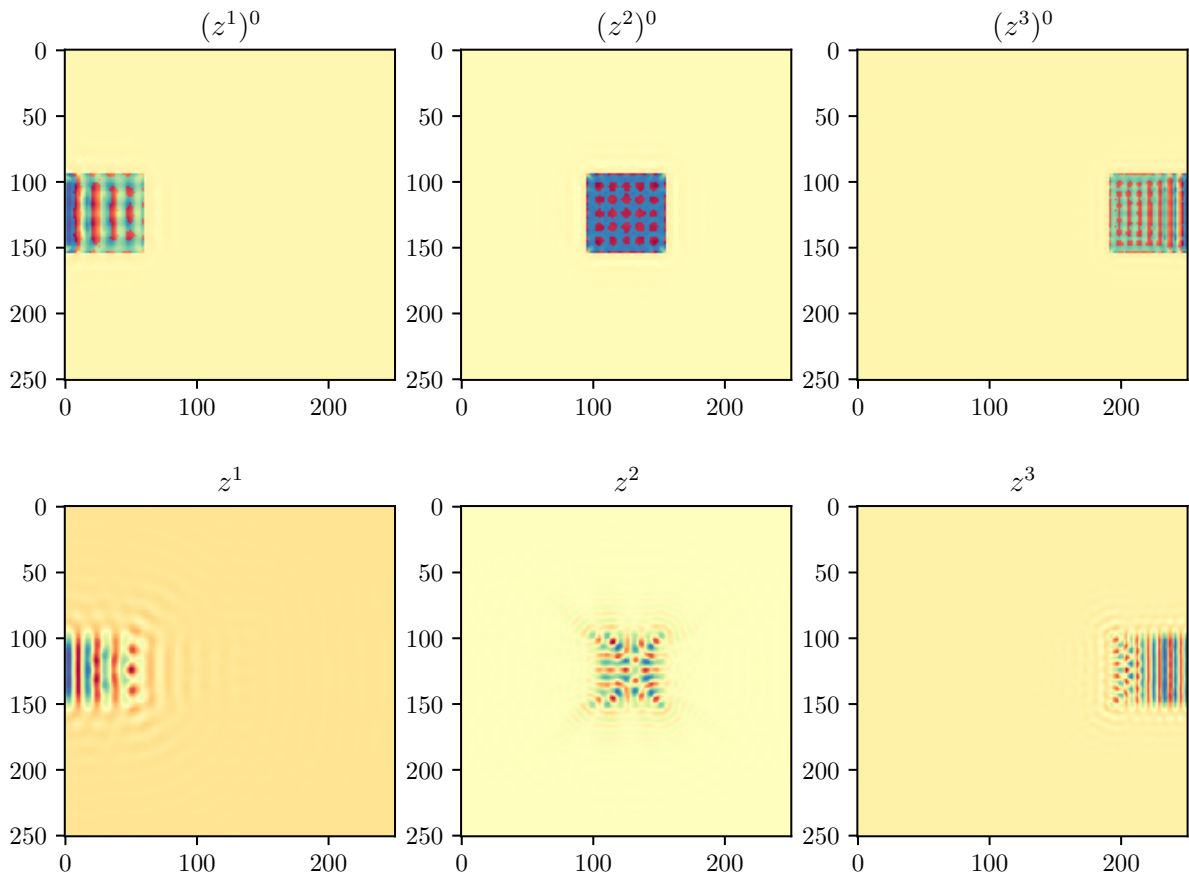


Figure 3: *Top row.* Fields suggested by solution to the dual problem. *Bottom row.* Fields in ADMM physical design. Columns show the three frequencies.

Minimization over z is identical to (4) and minimization over θ is similar minus the fact that for each k , the indices found in S_k are all constrained to be equal. Since the sum of convex quadratics is still a convex quadratic and, as before, since convex quadratics achieve minima at the boundary of an interval, we have

$$g(\nu) = -\frac{1}{2} \sum_{k=1}^m \max \left\{ \sum_{j \in S_k} W_{jj}^{-2} (a_j^T \nu - W_{jj}^2 \hat{z}_j)^2, \sum_{j \in S_k} W_{jj}^{-2} (a_j^T \nu + \nu_j \theta_j^{\max} - W_{jj}^2 \hat{z}_j)^2 \right\} - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2,$$

as the final Lagrange dual function. The corresponding dual problem can be written as a convex QCQP.

Field constraints. In the case where (1) has field constraints, *i.e.*,

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && (z_j - h_j)^2 \leq (z_j^{\max})^2, \quad j = 1, \dots, n \\ & && 0 \leq \theta \leq \theta^{\max}, \end{aligned}$$

for some $h \in \mathbf{R}^n$, the construction also parallels the one given in §3. The resulting dual optimization problem, in comparison to problem (7), cannot be written as a QCQP—it is, instead, a more general SOCP.

Regularizers. It is also possible to add a separable regularization term for θ , the parametrization of the device; for example, in the case where we would want to bias specific θ_j towards either 0 or θ_j^{\max} .

If we have a family of concave functions, $r_j : \mathbf{R} \rightarrow \mathbf{R}$ such that our regularizer can be written as a function of the form

$$\theta \mapsto \sum_{j=1}^n r_j(\theta_j),$$

(one such example is a linear function of θ), then the problem becomes

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 + \sum_{j=1}^n r_j(\theta_j) \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && 0 \leq \theta \leq \theta^{\max}. \end{aligned}$$

By using the fact that r_j is concave and therefore achieves a minimum over an interval at the boundary of the interval, it is possible to derive a bound that parallels (6).

Parameter perturbations. In some cases (*e.g.*, when considering temperature perturbations), it might be very natural to have a physical constraint of the form

$$(A + D \mathbf{diag}(\theta))z = b,$$

where $D \in \mathbf{R}^{n \times n}$ is a diagonal matrix that is not necessarily invertible. In other words, our new problem is

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + D \mathbf{diag}(\theta))z = b \\ & && 0 \leq \theta \leq \theta^{\max}. \end{aligned}$$

Directly applying the method from §3 yields a similar explicit form for g as given in (6).

Indeterminate eigenvalue. In the case where we want z to be a mode of the device with some unspecified eigenvalue λ with upper and lower limits $\lambda^{\min} \leq \lambda \leq \lambda^{\max}$, we can write the problem as

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \lambda I + \mathbf{diag}(\theta))z = b \\ & && \lambda^{\min} \leq \lambda \leq \lambda^{\max} \\ & && 0 \leq \theta \leq \theta^{\max}. \end{aligned} \tag{14}$$

To construct the dual, note that the Lagrangian of this problem is similar to the Lagrangian of problem (1),

$$\mathcal{L}(z, \theta, \lambda, \nu) = \frac{1}{2} \|W(z - \hat{z})\|_2^2 + I(\theta) + \nu^T ((A + \lambda I + \mathbf{diag}(\theta))z - b).$$

We will define the partial Lagrangian, \mathcal{L}^p to be the infimum of \mathcal{L} with respect to z and θ , leaving λ and ν as free variables. The solution to the partial minimization of \mathcal{L} is given in (6),

$$\begin{aligned} \mathcal{L}^p(\lambda, \nu) &= \inf_{z, \theta} \mathcal{L}(z, \theta, \lambda, \nu) \\ &= -\frac{1}{2} \sum_{j=1}^n W_{jj}^{-2} \max_{\theta_j \in \{0, \theta_j^{\max}\}} (a_j^T \nu + (\lambda + \theta_j) \nu_j - W_{jj}^2 \hat{z}_j)^2 - \nu^T b + \frac{1}{2} \|W \hat{z}\|_2^2. \end{aligned}$$

As $\mathcal{L}^p(\lambda, \nu)$ is concave in λ , it achieves its minimum at the boundaries of the domain of λ . So, since

$$g(\nu) = \inf_{\lambda^{\min} \leq \lambda \leq \lambda^{\max}} \mathcal{L}^p(\lambda, \nu),$$

we can write,

$$g(\nu) = \min_{\lambda \in \{\lambda^{\min}, \lambda^{\max}\}} \mathcal{L}^p(\lambda, \nu)$$

which is the minimum over a (finite) number of concave functions. The corresponding dual problem can then be expressed as a convex QCQP.

7 Conclusion

This paper has derived a set of lower bounds for a general class of physical design problems, making it possible to give (a) an easily-computable certificate that certain objectives cannot be physically achieved and (b) a bound on how suboptimal (relative to the global optimum) a given design could be. Additionally, as a side-effect of computing this lower bound, we also receive an initialization for any heuristic approach we might take for approximately solving (1) or its multi-frequency version (8).

Additionally, it seems feasible to obtain asymptotic bounds with respect to physical parameters (*e.g.*, with respect to the size of the device) via this approach, since the optimization problem in (7) can easily be written in an unconstrained form. In other words, picking any $\nu \in \mathbf{R}^n$ will yield *some* lower bound, and an appropriate choice might yield scaling laws that could be useful as general rules-of-thumb in inverse design.

Acknowledgements

We thank the Gordon and Betty Moore Foundation and Google for financial support. The authors would also like to thank Rahul Trivedi and Logan Su for useful discussions and help with debugging both code and derivations.

References

8 Appendix

8.1 Optimization using ADMM

We can approximately minimize (1) via the alternating direction method of multipliers, as in [LV13]. The method proceeds by forming the augmented Lagrangian of (1) and minimizing over each available variable, before updating a dual variable after each iteration.

ADMM iteration. We form the augmented lagrangian of problem (1) as in [BPC⁺11b].

$$\mathcal{L}^{\text{aug}}(z, \theta; \nu) = \frac{1}{2} \|W(z - \hat{z})\|_2^2 + \frac{\rho}{2} \|(A + \mathbf{diag}(\theta))z - b + \nu\|_2^2,$$

where $\rho > 0$ is a penalty parameter we set. Minimizing over each of z and θ (with the constraint $0 \leq \theta \leq \theta^{\text{max}}$) yields the following update rules

$$\begin{aligned} z^{(k+1)} &= (W^2 + \rho A(\theta^{(k)})^T A(\theta^{(k)}))^{-1} (W^2 \hat{z} + \rho A(\theta^{(k)})^T (b - \nu^{(k)})) \\ \theta_i^{(k+1)} &= S((b_i - a_i^T z^{(k+1)} - \nu_i^{(k)}) / z_i^{(k+1)}, \theta_i^{\text{max}}), \quad i = 1, \dots, n \\ \nu^{(k+1)} &= \nu^{(k)} + A(\theta^{(k+1)}) z^{(k+1)} - b, \end{aligned}$$

where we have defined $A(\theta) = A + \mathbf{diag}(\theta)$ and $S(x, u) = \min\{\max\{x, 0\}, u\}$ is the clamp function with upper limit, u , and we arbitrarily define $S(x/0, u) = 0$ for any x , though any value in $[0, u]$ would similarly suffice.

It can be shown that, if a feasible field exists for some θ , then, as $k \rightarrow \infty$, the iterates converge to a locally-optimal design θ and feasible field z , for an appropriately large choice of penalty parameter ρ [GGC18]. In practice, we find that ADMM is fairly robust and converges for a large range of values of ρ , though some choices appear to increase convergence speed.

8.2 Formulations of physical problems

Here, we describe ways of mapping the photonic inverse design problem into extensions of problem (1).

Maxwell's equations in three dimensions. Ampere's law and Faraday's law in Maxwell's equations, for a specific frequency ω , can be written as

$$\nabla \times H = -i\omega \varepsilon E + J \tag{15}$$

$$\nabla \times E = i\omega \mu H, \tag{16}$$

over some compact region of space $\Omega \subset \mathbf{R}^3$, with appropriate boundary conditions for H and E . Here, $E, H, J : \Omega \rightarrow \mathbf{C}^3$ are the electric field, magnetic field, and the current density, respectively, $\varepsilon, \mu : \Omega \rightarrow \mathbf{R}_+$ are the permittivity and permeability of the space (which we can often control by an appropriate choice of material), respectively. The bold \mathbf{i} —to avoid confusion with the index i —is the imaginary unit with $\mathbf{i}^2 = -1$. We will also assume that we can choose any permittivity and permeability that satisfy $\varepsilon^{\text{min}}(x) \leq \varepsilon(x) \leq \varepsilon^{\text{max}}(x)$ and $\mu^{\text{min}}(x) \leq \mu(x) \leq \mu^{\text{max}}(x)$ at each point of the region $x \in \Omega$.

8.2.1 Constant permeability

In many physical design problems, μ is also a constant that is independent of our material choice (*e.g.*, in the case where we are choosing between silicon or air, under small magnetic field) and constant through space (*i.e.*, $\mu(x) = \mu_c$ for $x \in \Omega$). Assuming this is true, we can write

$$\nabla \times \nabla \times E = \mathbf{i}\omega\mu_c\nabla \times H = \omega^2\mu_c\varepsilon E + \mathbf{i}\omega\mu_c J,$$

by taking the curl of (16) and plugging in (15). Rearranging gives,

$$-\nabla \times \nabla \times E + \omega^2\mu_c\varepsilon E = -\mathbf{i}\omega\mu_c J. \quad (17)$$

All we require is a discretization of E , ε , J , and the linear operator $-(\nabla \times \nabla \times \cdot)$. There are several standard ways of doing this (*e.g.*, the Yee lattice, see [Che95, §4.6.4]), though any method which discretizes the linear operator in the space will suffice. Let $z \in \mathbf{C}^{3n}$ be the optimization variable corresponding to the discretized field with $z^i \in \mathbf{C}^n$ being the field along each of the three axes, $i = 1, 2, 3$. Then, we can rewrite and discretize (17) as

$$\underbrace{(-\nabla \times \nabla \times + \omega^2\mu_c\varepsilon)}_A \underbrace{E}_z = \underbrace{-\mathbf{i}\omega\mu_c J}_b.$$

Here, each of $A \in \mathbf{C}^{3n \times 3n}$ and $b \in \mathbf{C}^{3n}$ are the corresponding discretizations of the variables they are below.

The next question is: how can we deal with the scalar permittivity term? One simple way is to allow $\theta \in \mathbf{C}^{3n}$ —which roughly corresponds to the discretized version of $\omega^2\mu_c\varepsilon$ —to have a component along each axis, which we will call θ^i for $i = 1, 2, 3$, and to then constrain all axes to be equal—*i.e.*, $\theta^1 = \theta^2 = \theta^3$. Using this idea, we can then write $\mathbf{diag}(\theta)z$, as a discretization of $\omega^2\mu_c\varepsilon E$. Note that, without the equality constraint, θ would be allowed to vary arbitrarily along each axis.

Finally, we set θ^{\max} to be the largest possible value of $\omega^2\mu_c\varepsilon$ at each point in the discretization (with a similar case for θ^{\min}), which lets us write the final program as a special case of (13),

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && \theta^1 = \theta^2 = \theta^3, \\ & && \theta^{\min} \leq \theta^i \leq \theta^{\max}, \quad i = 1, 2, 3. \end{aligned}$$

8.2.2 Arbitrary permeability

In the case where we are also allowed to vary the permeability throughout the space, we can discretize the equations in a similar way. The resulting system will have roughly double the size, but is still—usually, depending on the choice of discretization—relatively sparse.

First, we can write equations (15) and (16) in the suggestive form

$$\begin{bmatrix} \nabla \times & 0 \\ 0 & \nabla \times \end{bmatrix} \begin{bmatrix} E \\ H \end{bmatrix} + \mathbf{i}\omega \begin{bmatrix} \varepsilon I & 0 \\ 0 & -\mu I \end{bmatrix} \begin{bmatrix} E \\ H \end{bmatrix} = \begin{bmatrix} J \\ 0 \end{bmatrix}, \quad (18)$$

where $I \in \mathbf{R}^{3 \times 3}$ is the identity matrix. From here, we can perform a similar trick as in §8.2.1, by rewriting and discretizing (18) in the following way:

$$\left(-\mathbf{i}\omega^{-1} \underbrace{\begin{bmatrix} \nabla \times & 0 \\ 0 & \nabla \times \end{bmatrix}}_A + \begin{bmatrix} \varepsilon I & 0 \\ 0 & -\mu I \end{bmatrix} \right) \underbrace{\begin{bmatrix} E \\ H \end{bmatrix}}_z = \underbrace{-\mathbf{i}\omega^{-1} \begin{bmatrix} J \\ 0 \end{bmatrix}}_b,$$

where $A \in \mathbf{C}^{6n \times 6n}$, $z \in \mathbf{C}^{6n}$, and $b \in \mathbf{C}^{6n}$ are the discretized versions of the expressions above each. We will write z_E^i for $i = 1, 2, 3$ to be the i th component of the discretization of the E -field, with a similar definition for z_H^i .

As before, let $\theta_E^i \in \mathbf{R}^n$ and $\theta_H^i \in \mathbf{R}^n$ be the discretization of the permittivity and permeability, respectively, along each axis $i = 1, 2, 3$, with θ being the concatenation of each component and field over all points in the discretization. To ensure that the permittivity and permeability all remain scalar quantities, we simply constrain each entry of θ to be equal along all axes at each discretization point, which yields a problem which is a special case of (13):

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|W(z - \hat{z})\|_2^2 \\ & \text{subject to} && (A + \mathbf{diag}(\theta))z = b \\ & && \theta_E^1 = \theta_E^2 = \theta_E^3, \\ & && \theta_H^1 = \theta_H^2 = \theta_H^3, \\ & && \theta_E^{\min} \leq \theta_E^i \leq \theta_E^{\max}, \quad i = 1, 2, 3, \\ & && \theta_H^{\min} \leq \theta_H^i \leq \theta_H^{\max}, \quad i = 1, 2, 3, \end{aligned}$$

where θ_E^{\min} is defined to be the minimum value of ε at each discretization point with a similar definition for θ_E^{\max} , θ_H^{\max} , θ_H^{\min} .