

Chapter 6

Inverting the Helmholtz Equation

6.1 Introduction

In this chapter, we formalize our idea for using an inverse problem based approach to photonic device design. We begin in section 6.2 by surveying the field of PBG design using inverse problem based approaches in the literature, comparing the various methods with the approach we have adopted. This will frame the results presented in these final chapters within the greater context of this emerging field of research. In section 6.3, we derive the inverse Helmholtz equation we need to solve in order to address the design problems motivated in section 5.3. We provide a simple proof-of-principle example in section 6.4 to show that the inverse equations are in fact correct, and also the importance of regularization for solving this problem. We conclude this chapter with a first look at what happens when we ask for a mode that is not supported by a physically realizable dielectric function. This problem will be explored more extensively in chapter 7.

6.2 Inverse Problem Based Design

As we reviewed in chapter 5, there are many applications envisioned for PBG devices. However, the traditional design paradigm is really based on *trial and error*, which is not as suitable for applications. In the existing paradigm, different geometries are proposed and studied (numerically or experimentally), and the effects catalogued. As the

field progresses, the understanding of the fundamental scientific principles increases, and we begin to collect together useful effects that can be adapted for useful applications. Devices are improved by repeatedly varying different parameters to existing designs. After much trial and error, more intuition regarding the design problem is developed, so that, with a lot of human ingenuity, one hopes that subsequent ‘trials’ will less frequently turn out to be ‘errors.’ The other problem is that despite the many excellent devices that have been developed using this approach, one can usually not be certain how much room for improvement exists for the device. In other words, we are uncertain how *optimally* it is designed given the current manufacturing/cost constraints. Even if we somehow know that we are not optimal, the traditional design approach also does not give us a systematic or algorithmic way that lead us to the better/best designs.

An inverse problem based method is the exact reverse of the traditional method. Rather than systematically varying the causes (dielectric function) to catalogue the resulting effects (optical properties), we purposefully choose the desired effect and look for the unknown cause. As explained in chapter 3, the starting point is the effect, rather than the cause. We will review some methods that claim to be inverse problem based, but are effectively automated trial and error methods.

In an applications driven design paradigm, we imagine designing a device that will optimally perform some desired function. The starting point of the design process is focused on the end application, rather than the best known solution to date. This paradigm shift thus lends itself readily to the inverse problem formulation. Given some desired performance criterion suitable for the application, one needs to first develop a performance metric $\mathcal{P}(H_i(r), \omega_i(r), \eta(r))$ that may be a function of multiple eigenmodes and eigenfrequencies, and even the structure itself. The goal of **optimal** PBG device design is to find a structure that maximizes this function. It should be clear from our use of the term optimal in chapter 4 that we mean the *global* rather than the local optimum in this context, though we will see that this is often not how the term is applied in the literature. We can generally assume that this \mathcal{P} metric is obtainable in that under any design approach, it is required for evaluating

the performance of various designs. A design method that produces these optimal devices is considered an optimal design method (and to be clear, in principle this does not necessarily need to be an inverse problem based approach). An optimal inverse problem based design method combines both of these features. The underlying assumption is that if we can somehow find the field that optimizes the objective function (which can be a difficult problem in itself), an inverse problem method will allow us to solve for the dielectric that produces the desired properties. This was the idea behind the work (in the 2D approximation) by Geremia et al. [67], and is in fact the predecessor to the work in this thesis. It turns out there were some fundamental errors with the way the inverse problem was formulated, and these are explored and discussed in appendix D.

A brief survey of the field

Despite increased interest and efforts at developing this inverse problem based design paradigm in recent years, there are still only a few papers that have been published. For a more comprehensive review of the inverse problem approach to PBG design, the reader is referred to an article by Burger et al. [68], but the emphasis there is on the topology optimization approach, particularly those using the level set method. The journal *Inverse Problems* has many electromagnetic type inverse problems, but few directly applicable to photonic crystals or PBG devices in particular. An early treatment by Popov et al. [69] in this journal of a 1D photonic crystal (i.e. alternating layers of dielectric) scattering problem shows that given the reflection coefficient, the index of refraction can be determined assuming ‘practically sensible conditions’, which is of course a regularization condition. The 2D scattering problem is treated by Ammari et al. [70], and both of these are quite heavy on the mathematics, and are more concerned with parameter estimation than design. The interest there is much more heavily biased towards the mathematics of the inverse problem rather than the physics, and this is typical of the papers one finds in that journal to date. The first direct application of inverse problem methods to the PBG community was by

Dobson and Cox, whose work focused on finding unit cell geometries that maximized the bandgap in the *TM* [71] and *TE* [72] polarizations. Their technique uses a gradient-based algorithm, and they found that their final ‘optimal’ designs were quite sensitive to the targeted mid-frequency of the bandgap (which of course indicates that they have only reached local optima). The method in the *TM* polarization required a structure with an existing bandgap, but the later work removed this restriction.

6.2.1 Genetic algorithms

Sanchez-Dehesa’s group has recently published some results based on an inverse problem method for designing high efficiency waveguide couplers [73] and waveguide demultiplexers [74]. The demultiplexer design for an incoming signal with wavelengths $1.5 \mu m$ and $1.55 \mu m$ gives $> 45dB$ crosstalk suppression and $> 75\%$ coupling efficiency. However, the method they use is a genetic algorithm, where the geometry is parameterized and the only variable is whether the cylindrical rods of fixed size and location remain or are removed. Of course, using a genetic algorithm (GA) does not solve an inverse problem as we have defined the term in this thesis. It is systematic trial and error, with the errors generally discarded. The other comment is that GA’s generally do give good results for globally optimizing arbitrary functions with many local maxima and minima, as long as the design space is kept small. Therefore, this method cannot guarantee the optimality of the design over all possible structures since the design space must be heavily parameterized.

The technique by Gheorma et al. expands on their approach by allowing ‘aperiodic’ structures, so the location of these cylindrical rods are not fixed [66]. They found that a straight-forward application of the GA did not converge well given the extra degrees of freedom, so they used an adaptive algorithm, but it is of course still not an inverse problem method as we define it. The key idea here that is consistent with our work is that they give up the notion of an overall lattice periodicity to achieve improved performance. This is a bit of a break from the traditional view of PBG devices, where the bandgap effect plays a vital role in the design. Our results on the

enlarged defect region cavity design have similar philosophical underpinnings. The other relevant result here is that they allude to the idea of not actually obtaining the mode they had designed for, but attribute this effect to numerical (i.e. series truncation) errors and other constraints rather than as a fundamental limitation. To our knowledge, this is the only other paper that makes this observation explicitly, although the importance of this property was not fully appreciated.

6.2.2 Topology optimization methods

Topology optimization methods have also been used for minimizing losses in a waveguide bend [75], T-junction [76], and to improve the directional transmission properties of a waveguide termination [77]. In contrast to the GA methods, the dielectric function is not parameterized, but discretized into finite elements, where each element can take on any value. Therefore any design within the discretization bandwidth is within the design space. This aspect is similar to the domain of our design space. The optimization routine uses the iterative Method of Moving Asymptotes (MMA) optimizer which is a local gradient based algorithm, and requires the determination of the sensitivities of the objective function to the design variables at each iteration. However, computing the various sensitivities directly is not feasible, so linear approximations to the model are used. Linearization is necessary even though the Helmholtz equation is linear in the dielectric function because the objective function is not linear in each of the discretized elements. Based on the perturbation theory of linear operators [78], we expect this function can be highly non-linear. The validity and consequences of the approximations made are not discussed, but given the accuracy issues with just solving the forward problem, more caution should be exercised here.

In addition, there are some known problems associated with the method. It has been observed that the design algorithm stops converging as the resolution of the grid is increased [79]. This effect is attributed to the numerical instability of calculating the gradient as the ill-conditioning increases. This is consistent with what we discussed in chapter 4 regarding gradient descent methods for ill-conditioned problems. Finally,

since this is a gradient based method, it is also limited to finding locally optimal designs.

6.2.3 Level set method

Closely related to topology optimization is the level set method [80]. The level set function is a way to define an interface between distinct media, so it is a more ideal way of describing the dielectric function. Briefly, the zero crossings of the level set function define the boundaries between ϵ_1 and ϵ_2 . The optimization is done by iteratively updating the level set function by solving a Hamilton-Jacobi equation where the velocity term is chosen to climb the gradient to the objective function. Similar approximation issues found with the topology optimization methods are encountered here. This technique was recently applied to maximizing the bandgap for a 2D square lattice with great success for the *TM* polarization, and to a lesser degree for the *TE* polarization [81].

6.2.4 Analytical inversion of waveguide modes

Of the various papers in the literature, the work by Englund et al. [64] is probably the most similar to ours in spirit, and received positive review from the community [82]. They designed high- Q small-mode volume cavities by analytical solution of the inverse problem. By restricting the set of target modes to an expansion of the waveguide modes with a slowly varying envelope, they reduce the inverse problem to a 1D problem and analytically solve for the dielectric function along the waveguiding axis. The off-axis structure along the line defect region is reconstructed from the solution of the dielectric function along the line by assuming cylindrical defects. They solved for both a Gaussian and a sinc function envelope modulated defect mode with Q 's of 1.6×10^6 and 4.3×10^6 , and mode volume (in $(\frac{\lambda}{n})^3$) of 0.85 and 1.43 respectively. The results while promising have some limitations that were not addressed. Particularly for the sinc function design, the output mode actually did not resemble the target mode. In fact, comparing their figures 6(c) and 6(f), the sinc function mode looks

more like the Gaussian mode. This is actually consistent with our finding, as we believe this general inability to reach an arbitrary target is a ubiquitous problem.

6.2.5 Our approach

The first distinctive feature is that we derive the inverse problem from first principles, i.e. *ab initio*. We make no additional approximations or linearizations beyond a reduction of the problem from 3D to 2D. In contrast to the topology optimization methods, we solve the full inverse problem exactly, rather than restricting ourselves to local improvements to existing designs using approximate methods. The philosophy behind the approach is that if you can specify what you want, our method will tell you how to get it. Approximate methods limit the set of functionalities you can specify, because they must be of the correct form.

The other distinctive feature of our approach is that we do not parameterize the dielectric function at all in our design space. Any geometry commensurate with our chosen bandwidth is within the design space. Using the convex optimization regularization tool, we remove non-physically realizable values of the dielectric only. We can include additional fabrication constraints without increasing the computational domain if we so desire. Thus, we are not limited as in the GA approaches or the waveguide expansion approach. On this point, it is more akin to the topology optimization method. A criticism of our approach might be that our designs are not binary valued (as with level set methods), so they are not compatible with current fabrication technology. As discussed in appendix C, any discretized dielectric function should really be interpreted via its underlying continuous function. In many cases, we would argue that these should not actually be considered binary valued anyway. Secondly, even without this limitation, this is somewhat intentional. Consider a device requiring continuous valued dielectric functions that can improve performance by several orders of magnitude. Existence of such a design would still be important information to have. This may provide the right incentive to push fabrication technology improvements in that direction. Of course graded index fibers are a sim-

ple example of such a dielectric function, so they are not fundamentally impossible designs. Conversely, if existing designs already yield close to optimal performance, then it is not worth spending more effort into looking for improvements. Either way, there is value to our formulation. In that sense, it was our goal to construct a design method that simply gives you the best device possible without particular regard for current fabrication limitations, because those can improve with time and ingenuity. It is possible to derive new insight into the problem, which is another contribution of this method to the field. Furthermore, some forms of fabrication constraints can still be modeled into the design problem as well, and we demonstrate that with our enlarged defect cavity problem. We now derive the inverse equations for the design problem.

6.3 Inverse Helmholtz Equation

Using the plane wave basis, the derivation of the inverse Helmholtz equation uses the same mathematical principles that we used in obtaining the matrix form of the forward problem (explicitly derived in section 2.2 and applied to defects using the supercell method in section 2.3). Our restriction to TE polarized modes within a 2D analysis described in chapter 2 is still in effect here. We make use of the completeness and orthogonality of the plane waves here as well, but rather than solving for the $h_{\mathbf{k}}$ coefficients given an $\eta(r)$, we solve for the $\eta_{\mathbf{k}}$ coefficients given some desired or target field distribution $\mathbf{H}_m(\mathbf{r})$ instead.

6.3.1 Point defects

Working out the derivation for the point defect design, we expand the field in the Fourier basis:

$$\mathbf{H}_m(\mathbf{r}) = \sum_{\mathbf{\kappa}} a_{\mathbf{\kappa}}^{(m)} e^{i\mathbf{\kappa} \cdot \mathbf{r}} \hat{\mathbf{z}} \quad (6.1)$$

Starting from the Helmholtz equation again, we left multiply with a plane wave and integrate.

$$\int e^{-i\boldsymbol{\gamma} \cdot \mathbf{r}} \left\{ \frac{\omega_m^2}{c^2} \mathbf{H}_m(\mathbf{r}) = \nabla \times (\eta(\mathbf{r}) \nabla \times \mathbf{H}_m(\mathbf{r})) \right\} d^2\mathbf{r}$$

$$\frac{\omega_m^2}{c^2} a_\gamma^{(m)} = \int e^{-i\boldsymbol{\gamma} \cdot \mathbf{r}} \left[\nabla \times \sum_{\mathbf{k}} \eta_{\mathbf{k}} e^{i\mathbf{k} \cdot \mathbf{r}} \nabla \times \sum_{\boldsymbol{\kappa}} a_{\boldsymbol{\kappa}}^{(m)} e^{i\boldsymbol{\kappa} \cdot \mathbf{r}} \right] d^2\mathbf{r}$$

$$\frac{\omega_m^2}{c^2} a_\gamma^{(m)} = \sum_{\mathbf{k}, \boldsymbol{\kappa}} a_{\boldsymbol{\kappa}}^{(m)} \eta_{\mathbf{k}} \boldsymbol{\kappa} \cdot (\mathbf{k} + \boldsymbol{\kappa}) \int e^{i(\mathbf{k} + \boldsymbol{\kappa} - \boldsymbol{\gamma}) \cdot \mathbf{r}} d^2\mathbf{r}$$

$$\equiv b_\gamma^{(m)}$$

Choosing to collapse the delta function with $\boldsymbol{\kappa} = \boldsymbol{\gamma} - \mathbf{k}$

$$b_\gamma^{(m)} = \sum_{\mathbf{k}} \left[a_\gamma^{(m)} - \mathbf{k} \boldsymbol{\gamma} \cdot (\boldsymbol{\gamma} - \mathbf{k}) \right] \eta_{\mathbf{k}} \quad (6.2)$$

$$\equiv \sum_{\mathbf{k}} \mathbf{A}_{\boldsymbol{\gamma}\mathbf{k}}^{(m)} \eta_{\mathbf{k}} \quad (6.3)$$

In anticipation of using Tikhonov regularization for solving the inverse problem, we split up the dielectric function into two parts: an ‘initial’ geometry and a small corrective part:

$$\sum_{\mathbf{k}} \mathbf{A}_{\boldsymbol{\gamma}\mathbf{k}}^{(m)} \delta\eta_{\mathbf{k}} = \frac{\omega_m^2}{c^2} a_\gamma^{(m)} - \sum_{\mathbf{k}'} \mathbf{A}_{\boldsymbol{\gamma}\mathbf{k}'}^{(m)} \eta_{\mathbf{k}'}^{(0)} \quad (6.4)$$

$$\sum_{\mathbf{k}} \mathbf{A}_{\boldsymbol{\gamma}\mathbf{k}}^{(m)} \delta\eta_{\mathbf{k}} \equiv \beta_\gamma^{(m)} \quad (6.5)$$

We use β and b to distinguish between the two cases. Just as we did with the Helmholtz operator $\Theta^{(n)}$ in equation (2.24), we have explicitly included a superscript m on \mathbf{A} , a , b , and β to emphasize their dependence on the desired mode $\mathbf{H}_m(\mathbf{r})$. The importance of \mathbf{A} ’s dependence on the desired mode will be explored in section 7.2. The procedure for solving the inverse problem given some desired mode is to first

express it in the Fourier basis, and then forming the \mathbf{A} matrix using eqn. (6.2) and (6.3), and finally solve either eqn. (6.3) or (6.5) depending on which regularization scheme we choose to utilize.

6.3.2 Line defects

For the waveguide dispersion, recall that there is a separate eigenvalue problem for each wavevector \mathbf{q}_i along the propagation direction of interest. For each \mathbf{q}_i , we must perform the same steps as above on the waveguide form of the Helmholtz equation (eqn. (2.29)) to obtain the following inverse problem:

$$b_{\boldsymbol{\gamma}}^{(\mathbf{q}_i, m)} = \sum_{\mathbf{k}} \left[a_{\boldsymbol{\gamma} - \mathbf{k}}^{(\mathbf{q}_i, m)} \boldsymbol{\gamma} + \mathbf{q}_i \cdot (\boldsymbol{\gamma} - \mathbf{k} + \mathbf{q}_i) \right] \eta_{\mathbf{k}} \quad (6.6)$$

$$\mathbf{A}_{\mathbf{k}, \mathbf{k}'}^{(\mathbf{q}_i, m)} \equiv a_{\mathbf{k} - \mathbf{k}'}^{(\mathbf{q}_i, m)} (\mathbf{k} - \mathbf{k}' + \mathbf{q}_i) \cdot (\mathbf{k} + \mathbf{q}_i) \quad (6.7)$$

The solution $\eta_{\mathbf{k}}$ will have to simultaneously satisfy all $N_{\mathbf{q}}$ of these inverse equations, where $N_{\mathbf{q}}$ is the number of wavevectors we will include in the dispersion curve. This can be formally expressed by a vertical concatenation of the $\mathbf{A}^{(\mathbf{q})}$ matrices and $\mathbf{b}^{(\mathbf{q})}$ vectors.

$$\tilde{\mathbf{A}} \equiv \begin{bmatrix} \mathbf{A}^{(\mathbf{q}_1)} \\ \mathbf{A}^{(\mathbf{q}_2)} \\ \vdots \\ \mathbf{A}^{(\mathbf{q}_n)} \end{bmatrix}, \quad \tilde{\mathbf{b}} \equiv \begin{bmatrix} b^{(\mathbf{q}_1)} \\ b^{(\mathbf{q}_2)} \\ \vdots \\ b^{(\mathbf{q}_n)} \end{bmatrix}, \quad \text{and} \quad \tilde{\boldsymbol{\beta}} \equiv \begin{bmatrix} \beta^{(\mathbf{q}_1)} \\ \beta^{(\mathbf{q}_2)} \\ \vdots \\ \beta^{(\mathbf{q}_n)} \end{bmatrix} \quad (6.8)$$

For the remainder of this chapter, unless required for clarity, we will omit many of the cumbersome subscripts and superscripts on \mathbf{A} , b , and β with the understanding that, with a truncated basis, they can be treated like matrices and vectors.

6.4 Proof of Principle

In this section, we will work through a contrived problem as a proof of principle demonstration, but also demonstrate the steps one would take in performing such a design procedure. We first define our computational domain. We will choose a hexagonal lattice of cylindrical air holes¹ of radius $r = 0.3a$ in dielectric using a $7a \times 7a$ supercell, and include 19 reciprocal lattice vectors before truncating the Fourier series. The total number of plane waves is 931 in this example, and a is the lattice constant. As our illustration, we will ‘design’ a point defect cavity geometry where the central air hole is refilled with a dielectric material. This is referred to as the h_1 defect. Using the set of 931 plane waves, we would normally construct an $\mathbf{H}_m(\mathbf{r})$ with some suitably desired properties. In this case, we obtain the ‘desired mode’ $\mathbf{H}_m(\mathbf{r})$ by explicitly solving the forward problem.

For this simple geometry, we can use the analytical expression for $\eta_{\mathbf{k}}$. We follow the method found in [3] to evaluate the transform. The reference gives the defect-free coefficients:

$$\eta_{\mathbf{G}} = \eta_d \delta(|\mathbf{G}|) + (\eta_a - \eta_d) \frac{2\pi r^2}{\sqrt{3}} \left(\frac{2J_1(|\mathbf{G}|r)}{|\mathbf{G}|r} \right)$$

where $\eta_d = \frac{1}{11.56}$ is the reciprocal dielectric constant of the dielectric (value is typical of a semiconductor like AlGaAs), $\eta_a = 1$ represents air, and J_1 is the Bessel function. Evaluating only when $\mathbf{k} = \mathbf{G}$ gives us the bulk symmetry. Adding the defect means we need the expansion for $\delta\eta(r)$ corresponding to filling in the air hole. A straightforward modification of the above gives the required coefficients.

$$\eta_{\mathbf{k}} = (\eta_d - \eta_a) \frac{2\pi r^2}{N_1 \times N_2 \sqrt{3}} \left(\frac{2J_1(|\mathbf{k}|r)}{|\mathbf{k}|r} \right)$$

This is now evaluated for all \mathbf{k} ’s, and the factor $N_1 \times N_2$ in the denominator is due to integrating over the size of the entire supercell. We show in figure 6.1 the underlying dielectric function. Using these coefficients, we construct the Helmholtz

¹Recall the discussion in section 2.4 about the convergence issues of the plane wave method. Thus we are actually considering the underlying continuous function of the nominal geometry. Refer to appendix C for additional details.




Figure 6.1: Underlying continuous dielectric function of the nominal h_1 defect geometry.

operator using eqn. (2.28) and solve the eigenvalue problem. The localized mode is located at the 50th band. Figure 6.2 shows the magnetic field of the localized mode. We now take this $\mathbf{H}(\mathbf{r})$ and using only information about this field, attempt to get back the original dielectric function. We form the inversion matrix A following eqn. (6.2). Of course, after the discussion in chapter 3, it should not be surprising that this problem is ill-posed. Nevertheless, just to illustrate, we can try to perform a QR factorization to invert the A matrix and solve for $\eta_{\mathbf{k}}$. The result is shown in figure 6.3, and as we expect, it looks like noise that has been amplified, bearing no resemblance to the actual dielectric function that produced this mode.

6.4.1 Tikhonov solution

We now use Tikhonov regularization to solve the inverse problem. As a demonstration, we imagine that we are only given this localized mode, and assume we have no knowledge of what photonic crystals or bandgaps are at all. Solving the regularized problem gives a solution shown in figure 6.4. The structure is a marked improvement over the QR solution, and definitely suggests creating a periodic lattice of air holes




Figure 6.2: Real part of the magnetic field intensity for localized defect mode for the h_1 defect geometry.

with a central defect, although there are still some areas near the edges that do not quite resemble the exact solution. This can be partly attributed to the fact that the solution norm is sizeable $|\eta_{\mathbf{k}}| = 0.37$. To improve on the result, we now invoke the perturbative form of the inverse equation (eqn. (6.5)), and use the defect-free bulk lattice as an initial geometry. This finally gives us the same geometry that we had started out with, demonstrating that we have indeed solved the inverse problem. This solution is shown in figure 6.5.

6.5 Simulating Design Errors

In the previous section, the entire scenario is, of course, rather artificial, since we knew (by explicit construction) that some geometry must exist that will produce the target mode. In an actual design problem, we would not be certain of that *a priori*. We model this uncertainty by adding a small noise term to the target (h_1) field. From the discussion in chapter 3, we know that we are not guaranteed the existence of a solution in an inverse problem, which means that some desired modes just simply




Figure 6.3: Solution to h_1 inverse problem using a QR factorization to invert A . Notice the values of the dielectric function far exceed the original function.

cannot be supported. In this example, even if the perturbed field is not supported, we would still like to recover the h_1 geometry, because we know that the h_1 geometry reproduces the target field minus the small noise term. In the following example, the noise corresponds to a 1% perturbation.

Using the perturbed field as the target field, we proceed to form the inverse problem as before with a bulk lattice starting point and solve using Tikhonov regularization again. Using the same regularization parameter of $\lambda = 3 \times 10^{-3}$ as before yielded a noisy solution similar to the QR factorized solution. Clearly, much more regularization is required in this case. In figure 6.6, we show the solution $\delta\eta(r)$ using $\lambda = 3.73$. The residual norm in this case was 0.3143. We show both the real and imaginary parts of $\delta\eta(r)$, since there are some fairly significant contributions to the imaginary part of the dielectric. Looking at the real part of $\delta\eta(r)$, we find the dominant feature is as we expected, which is to make the central air hole more dielectric-like. Notice also in the dielectric function some fluctuations in the area surrounding the defect, which we see is an attempt to accommodate the added noise term in the target field. Of course, our model assumes that the dielectric is real-valued, so the imaginary parts




Figure 6.4: Solution to h_1 inverse problem using Tikhonov regularization with no assumptions about the dielectric geometry. Solution suggests creating a periodic lattice of air holes with a central defect.

are particularly problematic. Again, because we knew *a priori* the ‘correct’ solution, we can safely disregard the ripples in the real part and focus only on refilling the central hole.

A general limitation of the Tikhonov scheme is that the solution will need to be ‘interpreted’ to look for the most reasonable or feasible solution. In this case, if we did not know what the correct geometry should have been, we would first have dropped the imaginary parts, since we cannot do anything about those anyway, and then started the next iteration by filling in the central hole (since it is the most prominent feature). We would then redo the forward problem with the central hole refilled, and compare with our target mode. If we were still not satisfied, we could solve for the inverse problem again, but using the solution from the last iterate as the ‘initial geometry’. This is the strategy we will use for the PCW dispersion design problem. We would still not be guaranteed the solution, and we may well find the desired field is not supported. To rigorously confirm this, we make use of the convex optimization scheme developed in chapter 4.




Figure 6.5: Solution to h_1 inverse problem using Tikhonov regularization with an initial defect-free lattice geometry. (a) Actual solution $\delta\eta(r)$ to the inverse equation. (b) Full reconstructed solution $\eta(r) = \eta_0(r) + \delta\eta(r)$, which look identical to figure 6.1. The regularization parameter used was $\lambda = 3 \times 10^{-3}$, and the residual norm was below 10^{-9} .

6.5.1 Convex optimization regularization (COR)

The problem we need to solve, first shown without explanation in eqn. (4.1) becomes the following:

$$\begin{aligned} \min_{\eta} |A\eta - b|^2 \\ \text{subject to } \eta_{min} \preceq \mathcal{F}^{-1}\eta \preceq \eta_{max} \end{aligned} \quad (6.9)$$

where \mathcal{F}^{-1} is the inverse fourier transform operator. The symbol \preceq means component-wise less than or equal to, so each discretized value of the real-spaced dielectric function² must lie between η_{min} and η_{max} . We explicitly set the imaginary part of the dielectric function to zero by enforcing $\eta_{\mathbf{k}}^* = \eta_{-\mathbf{k}}$ and using the transformation in section 4.2.1. Notice that we are using b instead of β in the objective function, which means we assume no knowledge of the defect-free lattice. We use the convex optimization algorithm as described in chapter 4 to solve the constrained minimization

²The dielectric function used to generate this forward problem does not use the analytical expression for $\eta_{\mathbf{k}}$ because truncation leads to overshoot. As shown in figure 6.1, $\eta_{max} > 1$, which is strictly unfeasible. Therefore, the actual values of η can exceed the bounds here. Refer to section 8.2 for how this is handled.




Figure 6.6: Solution to the noisy h_1 inverse problem using Tikhonov regularization. (a) Real part of $\delta\eta(r)$. (b) Imaginary part of $\delta\eta(r)$. The regularization parameter used was $\lambda = 3.73$, and the residual norm was below 0.314.

problem. The solution is shown in figure 6.7 with a residual norm of 0.2489. There is a slight discrepancy between our solution and the original geometry within the region where the target mode has very little intensity. Intuitively, this is sensible if we think of the inverse problem in terms of its signal to noise ratio. Where the mode has little to no intensity, there is insufficient signal to overcome the added noise to reconstruct the desired dielectric completely. Without the added noise to the target field, the COR reconstructs the dielectric function perfectly, as with the Tikhonov regularization scheme with a residual norm at the machine precision level. The difference is that we did not require prior knowledge of the defect-free lattice using COR. We defer a more thorough comparison of the two schemes until section 7.4.1. With the added noise, both schemes gave reasonably close approximations to the original geometry. The interpretation we ought to make here is that the noisy mode is not supported by any physically realizable geometry. We can claim this rigorously because the globally minimized residual norm to the solution of eqn. (6.9) is non-zero. Therefore, no other geometry exists that can reduce the norm further (or exactly solves the inverse problem). This means that we cannot track the noise that has been introduced, and we will explore the role of the residual norm more closely in the next chapter.




Figure 6.7: Solution to the noisy h_1 inverse problem using CO regularization.