Chapter 8

Results

8.1 Waveguide Dispersion

We begin with the W1 waveguide, using a $9a \times 1a$ supercell geometry with a bulk hole radius of $0.3a$, and select 11 points along the $\Gamma - M$ propagation direction (labeled $\mathbf{q}$) and truncate after 19 Brillouin zones. This gives us 171 plane waves per $\mathbf{q}$-point for a total of 1881 plane waves. We formulate the Helmholtz operator as in eqn. (2.30) and solve to obtain the dispersion curve in the $\Gamma - M$ direction. We choose the curve with the anomalous dispersion characteristic, and as the primary objective[1] we seek to flatten the entire curve by 50% (see figure 8.1). As an additional objective in anticipation of mode matching with a high-Q small-mode volume cavity, we also require that the waveguide mode be more localized. The inverse problem is exactly as set up in chapter 3, and shown in eqn. (6.8). We once again use the defect-free lattice as $\eta_0(r)$, and since we do not have any additional constraints, we can use the simpler Tikhonov scheme. The required dispersion relation $\omega_i(q_i)$ is embedded into the set of β_i's. The target eigenmodes were constructed out of the initial eigenmodes by compressing $H(r)$ in the y-direction by 50% and padding the edges with zeros. This mode was then smoothed out by using the Fourier coefficients (numerically integrated) up to the truncation bandwidth and then renormalized. The resultant

[1]There is still no preliminary data on the specific form of the envelope function of a photon emission from an atom strongly coupled to a PBG-cavity, so at this point we do not know the precise form of the dispersion curve we will need. As our demonstration of arbitrary design, we chose to flatten the anomalous curve because both anomalous dispersion and the flattening of dispersion curves have greater relevance in the PCW community.

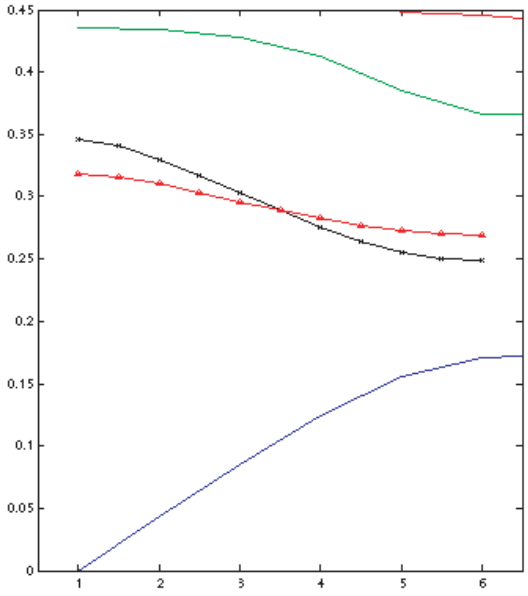

Figure 8.1: Dispersion relation for the W1 waveguide. The black x is the dispersion curve of the nominal W1 waveguide, while the red triangles show the desired curve.

target field is now simply the truncated fourier series using those coefficients.

8.2 Enlarged Defect Cavity

The cavity problem has the added complication of the fabrication constraint (eqn. (5.1)). This design problem cannot be solved using the Tikhonov regularization scheme, since a small norm does not exclude the undesired designs. Our investment in a more sophisticated regularization scheme allows us to incorporate these additional design constraints as part of the regularization procedure (provided that they are in the required convex form). We select a $5a \times 5a$ supercell geometry with a resolution of 12 points per lattice constant for a total of 3721 plane waves used. A smaller supercell was chosen because we did not want the algorithm to simply decrease the holes at the outer layers. We anticipate the algorithm may do that because the e-m field is negligible there (hence making the supercell approximation valid). The smaller geometry also allows us to perform calculations at a higher resolution so we can resolve

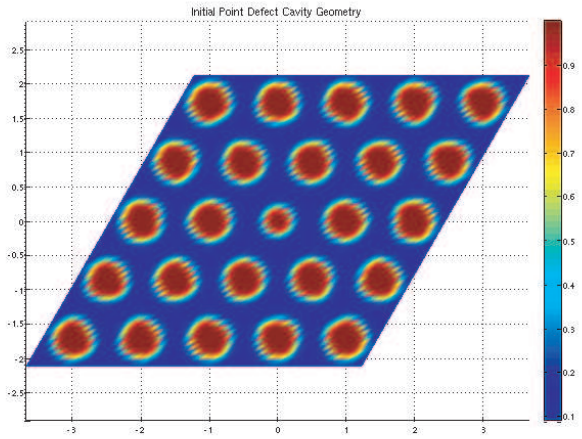

Figure 8.2: 5×5 h_1 point defect cavity starting geometry.

finer details in the inverted geometry. In particular, rather than strictly using the analytical expression for the Fourier coefficients, we use the FFT coefficients obtained for the discontinuous dielectric function and broaden it by convolution with a sharp Gaussian. Figure 8.3 shows the cross-section of the initial underlying dielectric function used. By defining the dielectric this way, we avoid Gibb's phenomenon with the Fourier series truncation, which is particularly important because the regularization scheme we have chosen bounds the maximum and minimum values of the dielectric to physically realizable values. Since we are testing a violation of the donor/acceptor mode rule, we also want to get accurate values of $\overline{\eta}_{\mathcal{D}}$ and $\overline{\eta}_{\mathcal{B}}$. We solve for the mode of a structure with nominal bulk hole size of $r_{bulk} = 0.4a$ and $r_{defect} = 0.3a$ (see figure 8.2)[2], and use that localized mode as the target mode.

In contrast to the waveguide dispersion problem, here we only care about the properties of the field, and not so much about the frequency. We can simply scale the 'lattice constant' of the structure to accommodate a real-world frequency that we would wish to use in a lab. We can therefore allow the frequency of the mode

[2]The actual 'hole radii' of the continuous function becomes smaller than these nominal values because of the convolution with the Gaussian. At this resolution, this was the largest nominal radius configuration we could accommodate without significant overshoot of the underlying dielectric function.

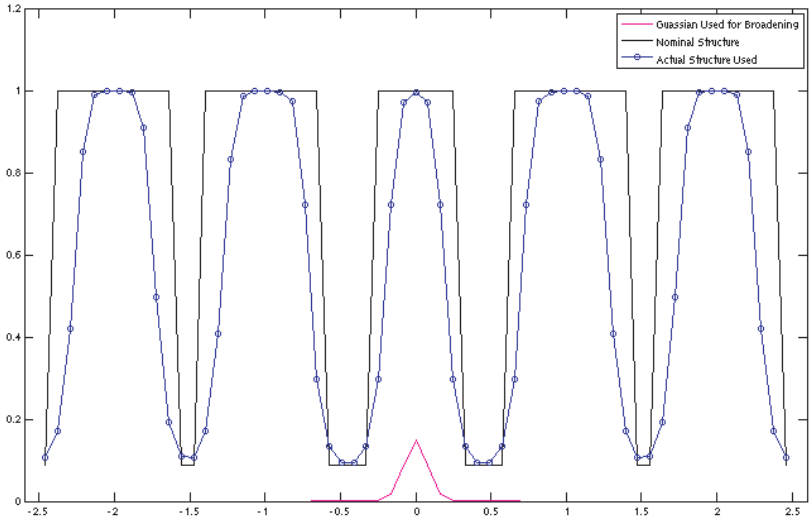

Figure 8.3: Cross-section view of dielectric function. The nominal structure and the actual structure is compared. Convolution with the Gaussian (pink curve) removes the overshoot due to truncation, but bandwidth limitations prevent an analysis of the nominal structure. Otherwise, the design goal and constraints are analogous.

to be another optimization variable, with the constraint that the frequency remain within the bandgap. This is another advantage to this scheme, as we no longer need to have a correct target frequency (as in [67]) in order to obtain the correct solution. To accommodate for the fabrication constraints and the variable frequency,

the regularized problem of eqn. (6.9) has to be modified slightly:

$$
\begin{aligned}
&\min_{\widetilde{\eta}} \left| \widetilde{A}\widetilde{\eta} \right|^2 \\
&\text{subject to } \eta_{min} \preceq \mathcal{F}^{-1}\eta \preceq \eta_{max} \\
&\eta^r_{\mathcal{D}} \geq \overline{\eta}_{\mathcal{D}} \\
&\eta^r_{\mathcal{B}} < \overline{\eta}_{\mathcal{D}} \\
&\frac{\omega^2_{min}}{c^2} \leq \frac{\omega^2}{c^2} \leq \frac{\omega^2_{max}}{c^2} \\
&\text{where } \widetilde{A} = [\ A \mid -b\ ] \quad , \\
&\widetilde{\eta} = \left[\ \eta^T \mid \frac{\omega^2}{c^2}\ \right]^T \quad , \\
&\eta^r_{\mathcal{D}} \equiv \frac{1}{\mathcal{A}_{\mathcal{D}}} \sum_{\mathcal{D}} \mathcal{F}^{-1}\eta \quad \text{, and} \\
&\eta^r_{\mathcal{B}} \equiv \frac{1}{\mathcal{A}_{\mathcal{B}}} \sum_{\mathcal{B}} \mathcal{F}^{-1}\eta
\end{aligned}
\tag{8.1}
$$

where the superscript T denotes the transpose. We choose $\eta_{max} = 1$ and $\eta_{min} = 0.0796$, and $\frac{\omega^2_{min}}{c^2} = 1$ and $\frac{\omega^2_{max}}{c^2} = 2.45$.

8.2.1 The direct solution

When we attempt to solve the two problems, we find as expected large residual norms, meaning that our target modes (and dispersion curve) are not supported. We can solve the forward problem using the solution to the inverse problem as our design. This gives a dispersion relation shown in figure 8.4, and a cavity mode that has only 72.9% overlap with the original target mode.

On the one hand, the results are quite encouraging, and in particular, the cavity mode is surprisingly good given our discussion about the acceptor/donor mode regime. On the other hand, the performance is not 'optimal' in the sense that we have not reached our target perfectly in either case. The question will always remain whether we have done the best we can do with a particular design, and that, as discussed in chapter 7, is difficult to overcome. However, we now show an approach that can be

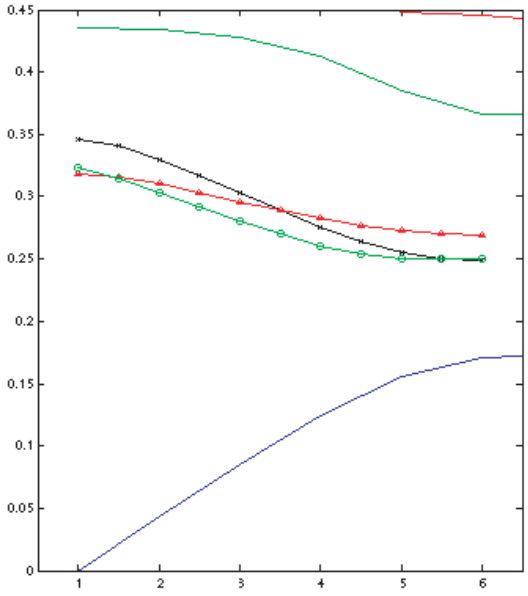

Figure 8.4: Initial result of the waveguide design. The green circles show the obtained dispersion curve, while the red triangles show the target curve. The black x show the original dispersion curve.

applied to these two problems that does yield even better results.

8.3 Iterative Approach

The method we will use to overcome these limitations involves an iterative approach. While we still cannot make any optimality claims, it does achieve the specified objective to an astonishingly high degree. We know that the root problem is the sparsity of the valid eigenmode landscape as discussed in section 7.3.1. The idea with the iterative approach is to use known valid eigenmodes as starting points of the inverse problem. We elaborate on the details as applied to each problem below.

8.3.1 Dispersion design

Given that our desired specifications are not physically realizable, it is time to make some design compromises. For the waveguide dispersion design problem, we relax the requirement that the eigenmodes be exactly as specified, and focus on the dispersion relation requirement (i.e. the eigenvalues). After the initial step, we solve the forward problem for the actual supported waveguide mode, and use that as the new target mode of our next inverse problem iterate. The target eigenvalues, however, remain as originally specified at each iteration. The iterative procedure proceeds as follows:

1. Use target modes and frequencies to formulate $\widetilde{A}$ and $\widetilde{b}$.
2. Solve the inverse problem. Pick a regularization parameter that gives a reasonable looking geometry.
3. Solve the forward problem using the interpreted result (as in section 6.5) from the solution of the inverse problem.
4. Look for the eigenmodes of interest. Check the frequencies of the waveguide modes. Continue if not satisfactory.
5. Use the new eigenmodes as the new target modes, but keep the original target frequencies.
6. Repeat until successive designs no longer improve.

Here, the motivation for the iterative approach is in obtaining a new starting field to define the updated inverse problem. This approach is inspired by the fact that we then start with an eigenmode that is known to be valid, even though we intend to perturb the eigenvalues. We know that the eigenmodes will be altered (if the inversion is successful), so that the target eigenmode-eigenvalue pair will likely be invalid still. However, without better insight, the prior valid eigenmode seems as good as any other. Using our strategy, at each iteration, the inverse problem tries to balance the

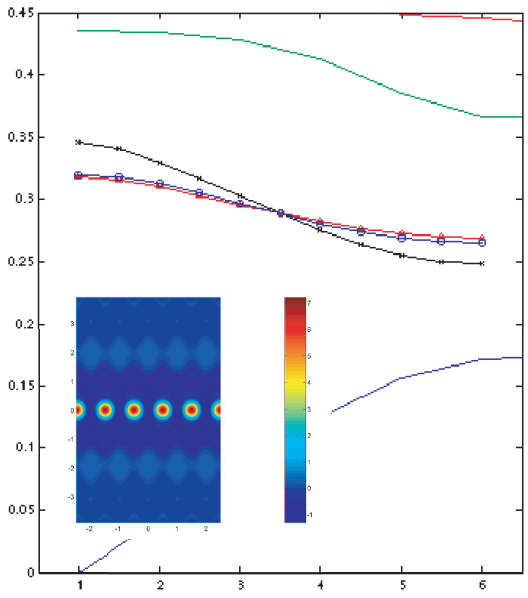

Figure 8.5: Final result of the waveguide design. The blue circles show the final dispersion curve, while the red triangles show the target curve. The black x show the original dispersion curve. The inset shows the waveguide mode still well localized.

demand to change frequencies with the need to accommodate the current eigenmodes. We let the result of the inverse problem tell us what the fields need to look like for the next iteration, since we have no way of knowing what they should look like anyway. The inverse problem 'machinery' itself will dictate how the dielectric needs to alter in order to shift the eigenvalues in the desired manner, and we will simply track the evolution of the eigenmode as it evolves. Each iteration brings us closer to the desired frequencies, and we accommodate the evolution of the fields by updating the target mode at each iteration.

The final waveguide design achieves a dispersion curve that compressed the original curve by 57% instead of the target of 50% (see figure 8.5). The field mode remains well localized, although we had not put any explicit restrictions on its form after the first iterate. This is not entirely surprising, since we know in general that small shifts in frequencies (eigenvalues) can be generated through small changes in the dielectric (as demanded by the Tikhonov regularization), which then bound the changes in the

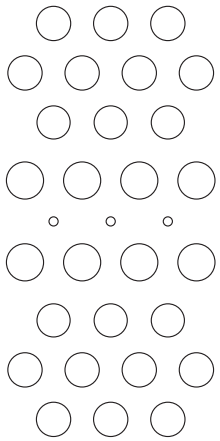

Figure 8.6: Final waveguide design for desired dispersion curve. See text for a description of the dimensions.

corresponding eigenmodes [78]. Figure 8.6 shows the final design. The missing central hole now reappears with a radius of $0.08a$, and the first neighboring row is displaced in the transverse direction towards the center by $0.15a$, with an increased hole radius to $0.325a$ from $0.3a$.

8.3.2 Cavity design

For the cavity design problem, we found that implementing the additional constraints (Eqn. (5.1)) in one step also takes us out of the valid eigenfunction space. However, unlike the waveguide design, there are no further constraints to relax. We have already allowed the frequency to take on any value within the bandgap. Once again, the problem is that the target mode (given the additional constraints) is not valid. It is clear that as we add more stringent constraints, the hypervolume of valid eigenmode space decreases. What was a valid mode becomes invalid as we lose access to the particular dielectric function. As we discussed in section 7.3.1, what we would like to do when our desired mode is invalid is to find a mode 'nearby' that is a valid one. We mentioned that the residual norm makes this a non-optimal process, but if we are 'close enough', then we increase our probability of finding the correct one. This is accomplished by iteratively ramping up to the final design constraint, rather

than requiring in one step $\overline{\eta}_{\mathcal{D}} > \overline{\eta}_{\mathcal{B}}$. At each iteration, we decrease the maximum allowed value of $\overline{\eta}_{\mathcal{B}}$. We use two indicators to guide how large a stepsize we take. First is the magnitude of the residual norm, since we know when it is large that the inverse problem has deviated significantly from the eigenvalue problem. The second is by comparing the predicted eigenfrequency (as given by the solution to the inverse problem) to the actual eigenfrequency (as calculated by the solution to the subsequent forward problem). We reduce the step size accordingly. There is no proof of convergence to this adaptive scheme, but we have not found a finer mesh to produce better results with the extra iterations. The algorithm of our approach is outlined as follows:

1. Use target mode to formulate A and b.
2. Update constraint $\overline{\eta_{\mathcal{B}}} < \xi_i, \xi_i \equiv \xi_{i-1} - \delta$.
3. Solve the inverse problem. If the residual norm is greater than some threshold, decrease δ and backtrack.
4. Solve the forward problem using the result from the inverse problem.
5. Look for the eigenmode of interest. The frequency of the mode should be close to the frequency obtained by the solver. If the frequency error is greater than some threshold, decrease δ.
6. Use the new eigenmode as the new target mode. Record the overlap of the new eigenmode with the original eigenmode.
7. Repeat until design constraint satisfied.

Using the iterative approach, the 'next nearest neighbor' remains available, whereas in a single step approach, many neighbors simultaneously become invalid. The residual norm metric performs worse the farther away we need to go. Using this approach, we find a structure that supports a mode with a 93.6% overlap with the original mode while satisfying our donor mode regime constraint (see figure 8.7).

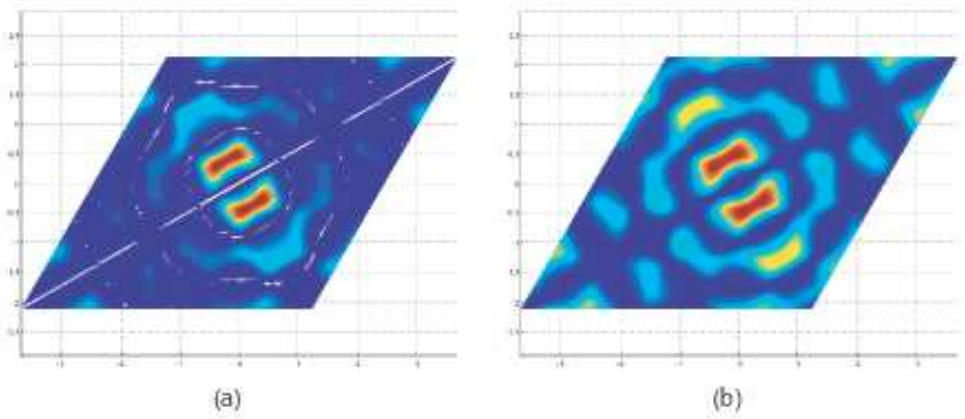

Figure 8.7: Magnetic field intensity of (a) original eigenmode and (b) eigenmode supported by final design.

Figure 8.8 shows the final design, and the first most noticeable feature is that the bulk holes now have a reduced value of η closer to 0.7, and in some regions it is down to 0.5. Additionally, the size of these features are also noticeably smaller. Six of the eight nearest holes from the center now take on a 'crescent moon' shape so it looks as though there is an effective lateral displacement for those holes, while the other two remain the same. For this structure, the final average bulk dielectric value $\overline{\eta_{\mathcal{B}}}$ is 0.1819, significantly reduced from the initial $\overline{\eta_{\mathcal{B}}}$ of 0.3962. The unaffected holes correspond to regions where the field intensity remains high, which is consistent with our intuition. The obtained eigenmode shows some leakage in the bulk area, but otherwise clearly resembles our target mode.

The most surprising feature of this iterative method is that the performance (as defined by the overlap integral of the obtained mode and the original target mode) was not a monotonically decreasing function (see figure 8.9). If we examine the algorithm closer, we see that after the first iteration, the original target mode no longer enters the picture, at least not directly. Whatever the inverse problem solution gives us, that becomes our next target mode. Granted, we use the mode that most resembles our target, but otherwise it does not play a role. Considering that we would have missed out on obtaining this solution if we had stopped a few iterations earlier, this

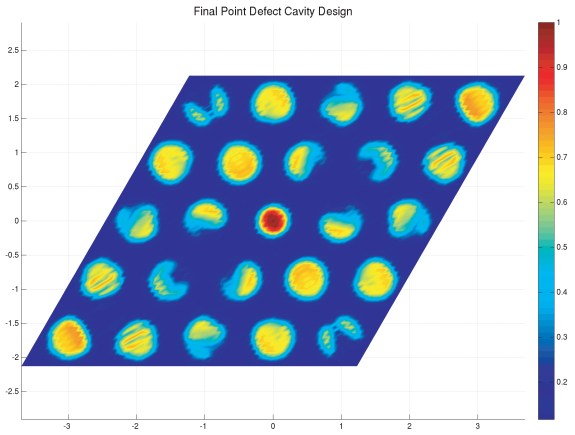

Figure 8.8: Final nanocavity design for acceptor eigenmode in donor configuration.

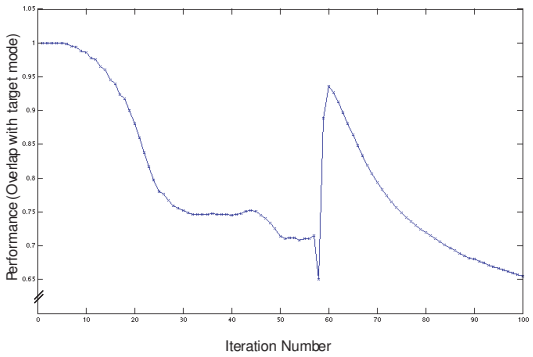

Figure 8.9: Plot of the performance at each intermediate iterate of the design process. Notice that this function is not monotonic, which also provides a glimpse into the topology of an objective function in design space.

result is somewhat troubling. However, it does highlight the difference between our methodology and the other inverse problem approaches that use gradient type schemes [76, 77], where there would have been no way of arriving at this design. It also reveals that in general, the performance function is not a simple functional of the dielectric, demonstrating real limitations of gradient based algorithms given the topology we found. This result further highlights the need to develop a better alternative to the residual norm as a metric for solving the inverse problem.

8.4 Concluding Remarks

We have set out to find a general turn-key methodology that allows arbitrary (and by extension optimal) design of desired PBG structures. Despite the limitations due to the convergence issues of the model, we retained the *ab initio* approach to elucidate fundamental and ubiquitous barriers to optimal PBG device design due to the topology of the space of valid eigenmodes. Having rigorously shown the non-existence of a solution, we attempted to look for designs that best approximates the desired functionality. We showed that the residual norm metric prevents a claim of finding the optimal 'next-best' design. While other inverse problem methods, particularly GA's and level set methods, can more readily yield designs that can be fabricated, they are confined to local improvements to existing design. Our approach, though currently non-optimal, can be used to provide these local methods with new starting points that lie 'closer' to the optimal design.

Even if the residual norm metric issue is worked out, our method still does not replace these other approaches because of the convergence issues, and the problem scales poorly with increased dimension. We were unable to take advantage of the ideas for fast convergence [4] because we do not know *a priori* where the boundary between high and low index materials will be in order to assign a tensorial value in that discretized element. However, we can make use of gradient based methods to go from an optimal continuous function design to the requisite binary valued design while minimizing sacrifice in performance.

Despite these general difficulties, we were able to extend our method to obtain remarkably good results that would be difficult to obtain with any other scheme. As we have shown, the topology of the problem is such that a gradient based algorithm would have missed our final design. The work we have developed here thus represents an important addition to the inverse problem based toolbox of design methods for the photonic problem.