Optimizing the topology of photonic crystals

E. E. Hart[a,*], A. Sóbester[a], K. Djidjeli[a], M. Molinari[a], K. S. Thomas[b], S. J. Cox[a]

[a] *CED Group, Faculty of Engineering and the Environment, University of Southampton, SO17 1BJ, U.K.*
[b] *DSSE Group, School of Electronics and Computer Science, University of Southampton, SO17 1BJ, U.K.*

Abstract

The performance of photonic crystal devices can depend strongly on their geometry. Alas, their fundamental physics offers relatively little by way of pointers in terms of optimum shapes, so numerical design search techniques must be used in an attempt to determine high performance layouts. We discuss strategies for solving this type of optimization problem, the main challenge of which is the conflict between the enormous size of the space of potentially useful designs and the relatively high computational cost of evaluating the performance of putative shapes. The optimization technique proposed here operates over increasing levels of fidelity, both in terms of the resolution of its non-parametric shape definition and in terms of the resolution of the numerical analysis of the performance of putative designs. This is a generic method, potentially applicable to any type of electromagnetic device shape design problem. We also consider a methodology for assessing the robustness of the optima generated through this process, investigating

*Corresponding author
Email address: eeh102@soton.ac.uk (E. E. Hart)

the impact of manufacturing errors on their performance. As an illustration, we apply this technology to the design of a 2D photonic crystal structure; the result features a large complete band gap structure and a topology that is different from previously published designs.
Keywords: topology optimization, evolutionary algorithms, photonic crystals, band gaps

---

1. The Need for a Generic Optimization Framework

Since Krauss et al. [1] experimentally showed the manufacturing feasibility of two-dimensional photonic lattice structures, there has been increasing interest in improvements in the design of 'standard' PhC geometries by changing their filling fraction and reducing their crystal symmetry [2], changing their geometry [3], as well as through research into irregular topologies and multi-axial PhC structures with up to 12-fold symmetric quasi-crystals [4]. Such design exploration and optimization is typically driven by an objective function derived from some desirable characteristic, such as the presence of both a transverse-electric (TE) and a transverse-magnetic (TM) cavity mode, a wide separation of high- and low-frequency bands, sharp density of state band edges, flat bands, extremely large (or narrow and clearly defined) gap-to-midgap frequency ratios, identical propagation constants for the two polarizations in waveguides, etc.

Regardless of the suitability metric or measure of performance selected to drive such search processes, a geometry modeling and optimization framework will be needed for the systematic generation and evaluation of putative design solutions. We report here on a study aimed at constructing and testing the prototype of such a framework. We show that an evolutionary heuristic based on a 'blank slate', global, non-parametric geometry description is capable of yielding designs of complexity comparable to those seen in nature [5, 6] (the evolutionary paradigm at the heart of the design process itself mirroring, to some extent, the evolutionary processes that yielded these natural PhCs).

As a simple illustration of the process we consider the problem of simultaneously maximizing the complete TM and TE gap-to-mid-gap ratio of a PhC structure. The selective pressure guiding the process in this example is the width of the complete band gap.

We next discuss the optimization framework employed to search through the space of possible device layouts, followed by our optimization results on an illustrative PhC device design problem in Section 3. In Section 4 we assess the sensitivity of the performance of the optimum design to variations in geometry. Conclusions and possible future directions are discussed in Section 5.

2. Shape- and Topology Optimization

2.1. Exploring and Exploiting the Design Space

The science of the geometrical design of continua has no universally agreed distinction between the concepts of 'shape' and 'topology', but for the purposes of optimization they are usually viewed as being related to a more localized *exploitation* of the search space (shape), as opposed to its global *exploration* (topology). Shape optimization usually deals with a range of geometries belonging to a parametric family – some objective (goal) function is optimized over the parameter space of the family. Topology optimization, sometimes also referred to as layout- or morphology optimization does away with the concept of a parametric family (indeed, it sometimes does away with parameters altogether) and considers a much broader range of possible solutions. Beyond such issues of semantics, it is important to note the relationship between search scope and geometry definition scope: while there are some grey areas, shape optimization generally targets local improvements around an initial, baseline design, while topology optimization is always global.

There are specialized methods that have less well defined place along this continuum of scopes between local and global geometry design (such as level set methods – see [7] for a recent PhC design example – and inverse design techniques – see, e.g., [8], or the two-dimensional PhC cavity loss minimization work of Englund et al.[9]), but most widely used local and global optimization methods sit clearly near the extremes of the scale.

There is a relatively long history of local geometry improvements in the design of electromagnetic media and, most germanely to the present study, in the design of PhC devices. One of the earliest efforts was that of Cox et al. [10], which set the tone for much subsequent geometry design work. They performed local optimization aimed at maximizing band gaps through generalized gradient ascent. Later developments include increasing the efficiency of local searches through the use of adjoint formulations of the governing equations – see Kiziltas et al. [11] (who implemented the method in a sequential linear programming framework) and Jensen and Sigmund [12] for examples of this approach. A more recent adjoint sensitivity-driven optimization example is the work of Jensen [13], who designed non-linear optical devices using updates based on the method of moving asymptotes.

Most of these local searches share a common feature: an initial guess is required (see Jiao et al.[14] for a further example, involving the design of a mode separator), which will commit the search to a specific basin of attraction on the objective landscape. Such guesses may be available in certain applications, but others, such as the design of PhC lattice structures for specific band structures or maximum complete band gaps, may demand a more global approach, due to the lack of obvious starting points.

As noted by Joannopoulos et al. [15], a lattice of isolated high-ϵ regions favours band gaps in the TM polarization, while TE band gaps are favoured in a connected lattice. This observation provides a useful handhold for any local optimization process seeking to maximize TE *or* TM band gaps, as the initial, sensible baseline designs required by such methods can be readily constructed following this philosophy – the work of Men et al. [16] provides a recent example of this school of thought (see also the global optimality conjectures of Sigmund and Hougaard [17]). However, such guidelines are yet to be formulated for complete gaps (beyond the simple combination of these two types of geometries [18]), which leaves the designer equipped with a local search technique faced with the following dilemma. Should the search for the widest total band gap be launched from either TE- or TM-favouring lattices or should a plethora of random starting points serve as baseline designs for a multi-start local search? Neither strategy is particularly appealing, especially considering the sheer size of the shape and topology design space. The former approach carries the risk of exploiting the local basins of attraction that are fairly likely to surround designs with broad TE or TM gaps, but never finding the truly interesting global optima of the total band gap width landscape. The random restart technique, on the other hand, while having a broader scope, may waste much time on trying to improve on hopeless starting points and might thus end up being ultimately quite inefficient.

Here we propose a multi-stage, hybrid heuristic, which, to a large extent, circumvents this unhappy dichotomy. It combines a population-based, fully global, evolutionary search (more on the philosophy behind this in the next section), with a greedy exploitation of promising basins of attraction as the second stage of the algorithm. We shall demonstrate via a simple case study (the maximization of the total band gap of a periodic structure) that this hybrid technique combines the exploratory ability of an evolutionary search with the fast basin optimization of a greedy local search. For an effective implementation of this philosophy the description of the geometry must be coupled tightly with these two optimization algorithms – we discuss this next.

2.2. A Biological Paradigm

Evolutionary algorithms are simplified models of the natural process of evolution by natural selection – in fact, they first emerged as instruments of such biological simulations in the late 60s and early 70s [19]. Virtual individuals, each representing a candidate solution to the design problem, evolve through multiple simulated generations under the selective pressure of a fitness function derived from the objective function of the optimization problem.

'Real' evolution produces increasingly well adapted life forms through the non-random survival of randomly varying replicators. This artificial evolution subjects the individuals (designs, whose variables are encoded in artificial 'chromosomes') to random variations through simulated crossover, mutation and other similar operators. The objective function (or, technically, the fitness function derived from it) governs the non-random survival part: a carefully controlled selection step weeds out the weaker (worse objective function value) individuals, guiding the population towards the better parts of the objective landscape. 'Carefully' here refers to the importance of getting the selective pressure right: brutally removing all weaker individuals from the gene pool runs the risk of premature convergence of the population towards a local optimum, while too gentle a selection step at the end of each generation may slow down the process and cost much computing time when the evaluation of the fitness function involves expensive analysis.

When implementing this type of algorithm to a geometry optimization problem (such as our search for the PhC featuring the widest band gap), the key question is: what is the geometry description formulation that both works well when driven by an evolutionary algorithm and, at the same time, allows a thorough exploration of the search space? A possible answer comes from the structural design community.

A class of methods with a long pedigree in the design of linearly elastic solids subjected to structural loads involves a finite element discretization of the design domain for the purposes of determining a stress field. This stress field is then used by an iterative heuristic to reshape the domain, often by removing under-utilized elements or by applying some other geometry variation step based on the discretization. Bendsøe and Rozvány pioneered this methodology (see [20] and [21] respectively to gain an overview of their contributions). Sandgren et al. introduced a genetic algorithm-driven variation on this, where the discretised design domain is encoded in a binary string (chromosome) and the simulated evolution process is put in charge of generating new designs. The appeal of this technique lies in the simplicity of the parameterisation: each bit represents a cell, a value of '1' meaning the cell contains structural material, '0' referring to a cell that has been removed [22].

More recently the PhCs community has also embraced this binary shape encoding methodology, with Shen et al.[23], Drupp et al.[24] and Preble et al.[25] reporting band gap width maximization results on square lattices. Kiziltas et al. [11] used a variation on this encoding methodology recently in a local optimization context – instead of a 'black and white' encoding, their design emerges from a greyscale image, where each discrete cell can take its dielectric from a range of values – the local optimization is followed here by post-processing aimed at reducing the number of different dielectric values to two.

In the light of the earlier discussion on starting points, it can be seen that one attraction of this geometry description method is that if no obvious starting topology is available, the evolutionary algorithm can simply be started from a random population. Equally, if we are in the possession of a broadly defined topology that is likely to mark the neighborhood of the global optimum, this topology (and, perhaps, its variants) can be used to *seed* the initial population, incorporating our knowledge of the problem into the gene pool of the evolutionary search.

2.3. Controlling the Size of the Design Space

Successful geometry optimization can only be achieved through a careful control of the following trade-off. A more flexible geometry description is usually only achievable at the expense of increasing the size of the design space. This, in turn, leads to an exponential increase in the cost of exploring this design space (the so-called 'curse of dimensionality'). On the other hand, restricting the search to a narrower range of topologies and/or shapes runs the risk of excluding 'unexpected' good designs, which would, otherwise, turn out to be the global optimum of the objective function landscape.

A way of reducing the size of the search space to a manageable level is to only consider the space of symmetrical topologies. This is the approach we shall take here. In our illustrative example we use the bit-string definition method to describe one-eighth of an eight-fold symmetrical unit cell.

A multi-stage search strategy is another way of achieving a workable compromise between global exploration and local exploitation of the design space. The evolutionary search begins by exploring the search space of a low resolution geometry description – that is, the design space is discretized at a coarse level. This is intended to establish the rough topology of the design, which will then form the starting point of the next stage: a *greedy* local search. The fundamental principle of a greedy search is simple: the effect of adding a *tile* (a square element of high dielectric material, whose size is determined by the current resolution of the search) or removing one is considered across the current (baseline) shape and, at the end of each iteration, only one such change is preserved – greedily, we implement the change that yields the greatest immediate improvement in terms of the width of the band gap.

Partway along the spectrum between the global 'blank slate' evolutionary search and the local greedy search we define a third type: a partial evolutionary optimization step. This works on the principle of identifying key areas of the current best geometry and evolving those areas alone, under the selective pressure of the objective. The selection of these areas of interest amounts, to some extent, to an arbitrary trimming of the design space, though the progress of the overall search strategy may provide useful clues as to which regions can be 'frozen' relatively safely. For instance, if the shape of a particular feature of the geometry is in a continuous state of flux over several stages of the optimization process, whereas other features remain largely unchanged, the search process can be focused on the areas that are suffering comparatively large variations.

The three types of design exploration and exploitation steps outlined above can be repeated iteratively over a series of grids of increasing resolution. This gradual ramping up of the geometrical flexibility is another means of controlling the size of the design space, a strategy, whose feasibility is illustrated in the next section by means of its application to a specific problem: the maximum band gap design of a square lattice GaAs ($\epsilon = 11.4$) structure in air ($\epsilon = 1$) with a lattice constant of $a = 1$.

3. Iterating Towards an Optimum

3.1. Analysis

The example we have chosen to illustrate our evolutionary search technique with requires the solution of Maxwell's equations. The traditional method used for this is the plane wave expansion method (PWEM) [26] [27] [28] as it is simple to implement and provides good results. However the PWEM creates large dense matrices, is slow to converge, does not scale well to complex systems and is thus computationally expensive. In this work we used a finite difference method (FDM), as it creates sparse matrices that will require less computation and memory (which makes optimization feasible).

The governing equations for the FDM are obtained by rearranging the Maxwell equations [15]. The TM and TE modes are, respectively:

$$\begin{aligned} \nabla \cdot \left(\frac{1}{\mu} \nabla E \right) + \lambda \epsilon E &= 0, \\ \nabla \cdot \left(\frac{1}{\epsilon} \nabla H \right) + \lambda \mu H &= 0, \end{aligned} \tag{1}$$

where μ is the magnetic permeability, which for the derivation is deliberately not set to 1. When the unit cell is a square, the domain is:

$$\Omega = \left\{ (x, y) \cdot 0 \leq x, y, \leq 1 \right\}. \tag{2}$$

Using the method of Yang [29], which is based on the earlier work done by Bierwirth *et al.* [30], the finite differences for the problem on a square can be derived. The work of Varga [31] has also strongly influenced the derivation. Bierwirth *et al.*, considers a graded finite difference mesh with spacings north (n), south (s), east (e) and west (w), and four material regions characterised by μ_k, ϵ_k, $k = 1, ..., 4$. An inner rectangle $ABCD$ is considered that lies midway in the mesh. Figure 1 shows part of a finite difference grid labelled with the notation of Bierwirth *et al.*

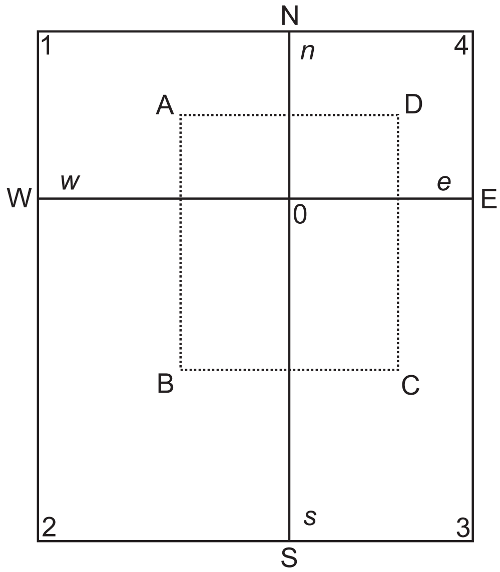

Figure 1: Part of a finite difference grid.

Consider the TE equation (1) and approximate the integral:

$$\int_{ABCD} \left[\nabla \cdot \left(\frac{1}{\epsilon} \nabla H \right) + \lambda \mu H \right] d\Omega. \tag{3}$$

Using Green's theorem to convert the first term into a surface integral gives:

$$\oint_{ABCD} \left[\frac{1}{\epsilon} \nabla H \right] \cdot \mathbf{n} ds + \lambda \oint \oint_{ABCD} \mu H d\Omega = 0. \tag{4}$$

Approximating the gradient on the line AB gives:

$$\nabla H \cdot \mathbf{n} = \frac{\partial H}{\partial x} \approx \frac{H_W - H_0}{w}. \tag{5}$$

Hence

$$\int_{AB} \frac{1}{\epsilon} \nabla H \cdot \mathbf{n} ds \approx \left(\frac{H_W - H_0}{w} \right) \left(\frac{n}{2\epsilon_1} + \frac{s}{2\epsilon_2} \right) \tag{6}$$

and the other contribution to the contour integral in equation (4) are derived in a similar manner. The second integral in equation (4) can be approximated to:

$$\oint \oint_{ABCD} \mu H d\Omega \approx \frac{1}{4}(nw\mu_1 + ws\mu_2 + se\mu_3 + en\mu_4)H_0. \tag{7}$$

Hence the finite difference approximation can be written as:

$$\begin{aligned}
& \frac{1}{2w}\left(\frac{n}{\epsilon_1} + \frac{s}{\epsilon_2}\right) H_W + \frac{1}{2s}\left(\frac{w}{\epsilon_2} + \frac{e}{\epsilon_3}\right) H_S \\
+ \quad & \frac{1}{2e}\left(\frac{s}{\epsilon_3} + \frac{n}{\epsilon_4}\right) H_E + \frac{1}{2n}\left(\frac{e}{\epsilon_4} + \frac{w}{\epsilon_1}\right) H_N \\
- \quad & \frac{1}{2}\left[\frac{1}{w}\left(\frac{n}{\epsilon_1} + \frac{s}{\epsilon_2}\right) + \frac{1}{s}\left(\frac{w}{\epsilon_2} + \frac{e}{\epsilon_3}\right)\right] H_0 \\
- \quad & \frac{1}{2}\left[\frac{1}{e}\left(\frac{s}{\epsilon_3} + \frac{n}{\epsilon_4}\right) + \frac{1}{n}\left(\frac{e}{\epsilon_4} + \frac{w}{\epsilon_1}\right)\right] H_0 \\
+ \quad & \frac{\lambda}{4}\left(nw\mu_1 + ws\mu_2 + se\mu_3 + en\mu_4\right) H_0 = 0.
\end{aligned} \tag{8}$$

If $n = w = s = e = h$ and $\epsilon_j = \epsilon$, $\mu_j = \mu$, $j = 1, 2, 3, 4$ the equation can be rearranged to form the standard five point difference approximation.

The Bloch-Floquet theory is used to treat the periodicity [32] [33] and taking a rectangle as a unit cell with sides a, b the first Brillouin zone is given by

$$B = \mathbf{k} = s_1 \begin{pmatrix} \frac{1}{a} \\ 0 \end{pmatrix} + s_2 \begin{pmatrix} 0 \\ \frac{1}{b} \end{pmatrix}, -\pi < s_1, s_2 < \pi. \tag{9}$$

The solutions then take the form

$$H(\mathbf{x}) = exp(\imath\mathbf{k}.\mathbf{x})U(\mathbf{x}), \tag{10}$$

where $U(\mathbf{x})$ is a periodic function.

Substituting equation (10) into equation (8) gives the modified finite difference formula:

$$\begin{aligned}
& \frac{e^{-k_1 w}}{2w}\left(\frac{n}{\epsilon_1}+\frac{s}{\epsilon_2}\right)U_W + \frac{e^{-k_2 s}}{2s}\left(\frac{w}{\epsilon_2}+\frac{e}{\epsilon_3}\right)U_S \\
+ & \frac{e^{-k_1 e}}{2e}\left(\frac{s}{\epsilon_3}+\frac{n}{\epsilon_4}\right)U_E + \frac{e^{-k_2 n}}{2n}\left(\frac{e}{\epsilon_4}+\frac{w}{\epsilon_1}\right)U_N \\
- & \frac{1}{2}\left[\frac{1}{w}\left(\frac{n}{\epsilon_1}+\frac{s}{\epsilon_2}\right)+\frac{1}{s}\left(\frac{w}{\epsilon_2}+\frac{e}{\epsilon_3}\right)\right]U_0 \\
- & \frac{1}{2}\left[\frac{1}{e}\left(\frac{s}{\epsilon_3}+\frac{n}{\epsilon_4}\right)+\frac{1}{n}\left(\frac{e}{\epsilon_4}+\frac{w}{\epsilon_1}\right)\right]U_0 \\
+ & \frac{\lambda}{4}\left(nw\mu_1 + ws\mu_2 + se\mu_3 + en\mu_4\right)U_0 = 0.
\end{aligned} \tag{11}$$

3.2. Design Search

Starting with a very wide-ranging search and gradually homing in on promising ‘basins of attraction’ of the objective function landscape is a strategy with a good historical track-record in many fields and we have employed it here in the process of maximizing the gap-to-midgap ratio of a generic PhC geometry. The broad phase of the search is represented by a completely free ‘clean slate’ exploration of the design space defined by a 30×30 tile geometry (or, to be precise, over the 120-tile triangular section that defined the eight-fold symmetrical geometry), analyzed on a finite difference grid of the same resolution. We ran an evolutionary optimization process on a population of 10,000 individuals on a total computing budget of approximately 350

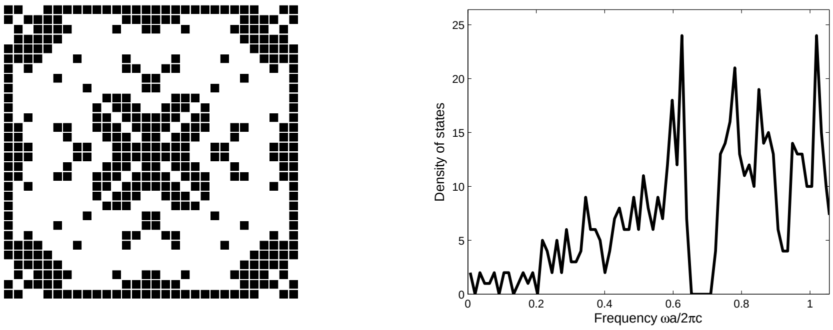

(a) Geometry. (b) Density of states.

Figure 2: The shape resulting from the ‘clean slate’ initial search.

CPU-hours, which, at a computational cost of around 10 seconds per fitness function evaluation[1], enabled the simulation of 32 generations. The resulting geometry is depicted in Figure 2(a), with Figure 2(b) showing its density of states diagram. The relatively high computational budget allocated to this part of the study reflects its importance, as it established the fundamental morphology of the design.

The design resulting from this study formed the basis of a greedy optimization heuristic (as described in the previous section), which yielded the result depicted in Figure 3(a) – once again, the corresponding density of states diagram is shown alongside the geometry, in Figure 3(b).

In spite of the optimizer having had the freedom to alter any part of the geometry, it can be seen to have preserved the general topology of the evolved design, which had served as its starting point. The convergence of the search

[1]Higher order analysis methods could, potentially, be considered for a more accurate assessment of the performance of a given design. However, their much higher computational cost would clearly render global searches of this breadth unfeasible.

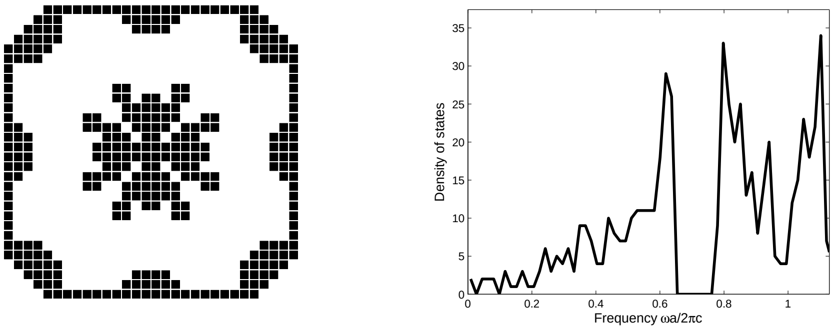

(a) Geometry (b) Density of states

Figure 3: A greedy search over the full geometry generated by the original evolutionary process resulted in this geometry.

on this class of shapes indicated that this was the appropriate stage for an increase of the geometry definition resolution to a 60×60 grid for a finer tuning of the unit cell layout. The combined cost-increasing effect of the requirement for a higher resolution finite difference analysis (more expensive calculation of the fitness function) and the quadrupled shape complexity necessitated a restriction of the search space to the finer detail around the (scaled up version of the) layout found earlier: the next evolutionary search was restricted to the two layers of elements around the edges of the high dielectric regions of the current design. The shape depicted in Figure 4(a) resulted from this search. As before, a greedy search was used to fine-tune the shape at its current resolution – its outcome can be seen in Figure 5.

The resolution of the grid was increased once more at this point to 120×120 tiles, through a factor two scaling of the best 60×60 layout found in the previous stage the search. The next stage of the evolutionary optimization was then applied to the two rows of tiles around the edges of this baseline

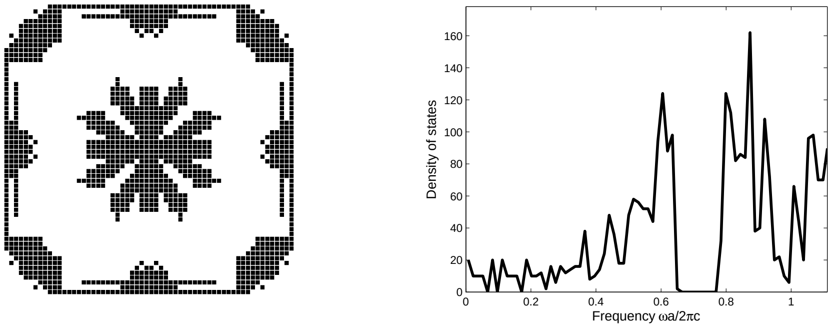

(a) Geometry (b) Density of states

Figure 4: The result of the evolutionary search at the 60×60 resolution.

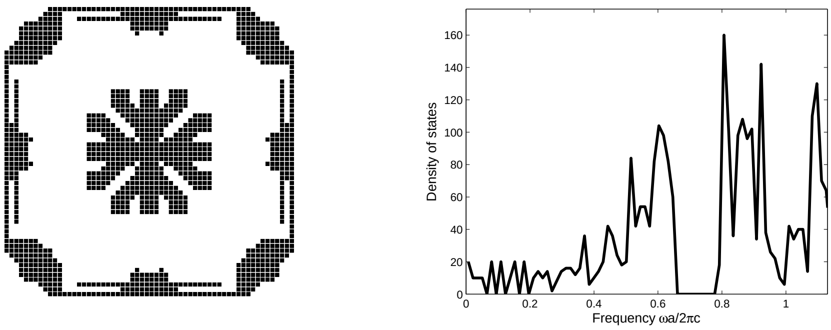

(a) Geometry (b) Density of states

Figure 5: Shape resulting from the greedy fine-tuning of the design evolved over the 60×60 grid.

(a) Geometry

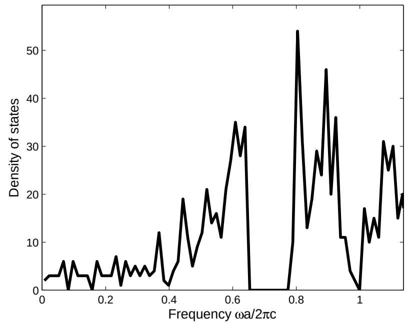

(b) Density of states

Figure 6: Result of evolutionary search over the edges of the shape represented on the 120×120 FD grid.

design – the resulting structure is shown in Figure 6. Breaking the layout down into three major constituents – the structures at the center, at the corner and on the side of the unit cell – it was apparent that the central part and the corners were now relatively stable. We therefore smoothed and froze these and concentrated the remainder of the evolutionary search time budget on the sides of the unit cell. The outcome of this further exploitation of the band gap landscape is shown in Figure 7, with an ultimate maximum gap-to-midgap ratio of 20.7% (Figure 7(c) depicts a larger section of the crystal structure comprising 3x3 unit cells).

4. Sensitivity Analysis and Robustness

When the resolution of the design space across which a shape optimization process operates is comparable with the resolution of the fabrication technologies expected to be employed to manufacture the resulting shapes, the question of *performance robustness* arises. In mathematical terms this

(a) Geometry

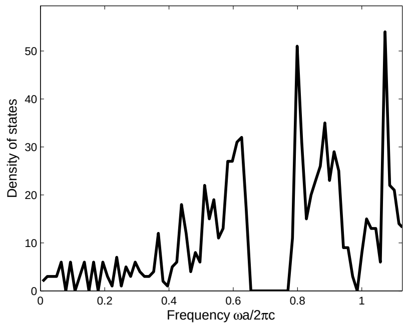

(b) Density of states

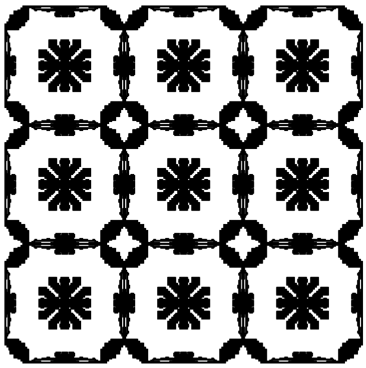

(c) Crystal structure

Figure 7: The final design: a GaAs structure with a gap-to-midgap ratio of 20.7%.

amounts to the ensemble of sensitivities of the performance metric (the total band gap ratio in this case) with respect to geometrical deviations from the nominal shape yielded by the design search. The aim of a robustness study then is to create a link between performance uncertainty and the shape uncertainty arising from manufacturing errors.

A full model of this link can only be constructed if the precise error margins of the manufacturing process, as well as of the analysis capability are known, but a qualitative assessment can be gained even with partial information. We endeavour to accomplish this here by considering the impact of losing GaAs tiles from the nominal design or accidentally failing to remove tiles from what constitutes the 'air' region of the nominal design.

Let us first consider the case when the fabrication process accidentally results in high dielectric material in an 'air' region. The impact of this type of scenario on the total gap-to-midgap ratio of the structure is shown in Figure 8(a). This is a depiction of the following experiment. The final structure generated by the optimization (Figure 7(a)) is taken as the baseline geometry for a series of band gap computations, where a square block of nine GaAs tiles is added to the structure, centred around each of the blank (air) tiles[2]. The colour of these tiles is a representation of the total gap-to-midgap ratio of the 'flawed' structure, with the colour range being determined by the interdecile interval of the range of total gap-to-midgap ratios obtained during the series of computational experiments.

The baseline structure is represented by black dots in Figure 8(a), as well

[2]The added tiles are not copied across the structure as part of its symmetrical folding process.

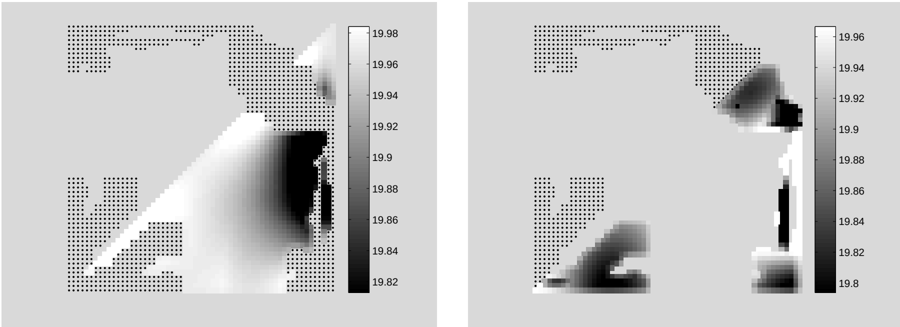

(a) Nine extra tiles. (b) Nine missing tiles.

Figure 8: The impact on the gap-to-midgap ratio (in %) of accidentally adding or failing to remove a square block of nine tiles during the fabrication process.

as in Figure 8(b), the latter showing the results of the opposite experiment, that is, when for each GaAs tile we consider the effect of missing out a square block of nine tiles.

Both figures highlight the areas where fabrication errors would be most critical. In terms of the accidental fabrication of a structure with nine supranumeral tiles, the most sensitive region appears to be near the vertical side of the triangle. The impact of missing out GaAs tiles is also the greatest near the vertical side of the triangle. Both figures indicate, however, that the structure is relatively robust to manufacturing errors.

5. Conclusions and Future Work

Design optimization based on numerical solution of partial differential equations always presents a delicate trade-off. On the one hand, the analyst has the option of deploying higher order, high fidelity, but correspondingly expensive codes to drive a limited scope optimization process. The alternative is to opt for lower cost, lower fidelity analysis, to support a wide ranging, global scope search process. For the illustrative example presented here we opted for the latter in the interest of demonstrating a broad exploration of a wide variety of shapes, as opposed to the limited fine-tuning of existing topologies that more expensive codes would have restricted us to. This choice resulted in the example optimization process yielding a structure with a high gap-to-midgap ratio (20.7%) with a topology that is radically different from previously published PhC structures.

This result also confirms that an iterative, multi-resolution search process comprising an evolutionary heuristic and a greedy exploitation phase is a viable PhC topology optimization tool in the search for structures with specific optical effects. The chosen geometry representation method has also proven effective in supporting this multi-level approach, while being sufficiently flexible and generic to allow design searches over different types of domains.

More problems

Future studies could include the investigation of the impact of the choice of dielectric material on the optimal shape – this would essentially involve repeating the study described here for different dielectric constants. Hexagonal lattices could also be considered, as they show promising performance for simple geometries – the methodology described here could be adapted very readily to the optimization of the layouts of such structures. In terms of different objective functions, the selective pressure of the evolutionary algorithm could be biased in favour of structures featuring flat frequency bands, as well as in favour of PhC layouts with multiple total band gaps.

6. Acknowledgements

We gratefully acknowledge Microsoft for their support of the Microsoft Institute for High Performance Computing, University of Southampton. The work of the second author was funded by the Royal Academy of Engineering and the UK Engineering and Physical Sciences Research Council through their Research Fellowship Scheme. The authors would like to thank Dr. Darren Bagnall and an anonymous referee for their useful suggestions.

References

[1] T. Krauss, R. D. L. Rue, S. Brand, Two-dimensional photonic-bandgap structures operating at near-infrared wavelengths, Nature 383 (1996) 699 – 702. doi:10.1038/383699a0.

[2] C. M. Anderson, K. P. Giapis, Larger two-dimensional photonic band gaps, Physics Review Letters 77 (14) (1996) 2949–2952. doi:10.1103/PhysRevLett.77.2949.

[3] S. J. Cox, D. C. Dobson, Band structure optimization of two-dimensional photonic crystals in H-polarization, Journal of Computational Physics 158 (2000) 214–224.

[4] M. E. Zoorob, M. D. B. Charlton, G. J. Parker, J. J. Baumberg, M. C. Netti, Complete photonic bandgaps in 12-fold symmetric quasicrystals, Nature 404 (2000) 740–743. doi:10.1038/35008023.

[5] P. Vukusic, J. R. Sambles, Photonic structures in biology, Nature 424 (2003) 852–855.

[6] G. Parker, Biomimetically-inspired photonic nanomaterials, Journal of Materials Science: Materials in Electronics 21 (10) (2010) 965–979.

[7] C. Y. Kao, S. Osher, E. Yablonovitch, Maximising band gaps in two-dimensional photonic crystals by using level set methods, Journal of Applied Physics B: Lasers and Optics 81 (2005) 235–244.

[8] J. M. Geremia, J. Williams, H. Mabuchi, Inverse-problem approach to designing photonic crystals for cavity QED experiments, Physical Review E 66 (2002) 066606.

[9] D. Englund, I. Fushman, Vačković, General recipe for designing photonic crystal cavities, Jounal of the Optical Society of America 13(16) (2005) 5961–5975.

[10] S. J. Cox, D. C. Dobson, Maximizing band gaps in two-dimensional photonic crystals, SIAM Journal on Applied Mathematics 59 (6) (1999) 2108–2120.

[11] N. Kiziltas, N. Kikuchi, J. L. Volakis, J. Halloran, Topology optimization of dielectric substrates for filters and antennas using SIMP, Archives of Computational Methods in Engineering 11 (4) (2004) 355–388.

[12] J. S. Jensen, O. Sigmund, Topology optimization of photonic crystal structures: a high-bandwidth low-loss t-junction waveguide, Jounal of the Optical Society of America B 22(6) (2005) 1191–1198.

[13] J. S. Jensen, Topology optimization of nonlinear optical devices, Structural and Multidisciplinary Optimization DOI 10.1007/s00158-011-0640-0 (to appear) (2010) 1–13.

[14] Y. Jiao, S. Fan, D. A. B. Miller, Demonstration of systematic photonic crystal device design and optimization by low-rank adjustments: an extremely compact mode separator, Optics Letters 30(2) (2005) 141–143.

[15] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, R. D. Meade, Photonic crystals moulding the flow of light, Princeton University Press, 2008.

[16] H. Men, N. C. Nguyen, R. M. Freund, P. A. Parrilo, J. Peraire, Bandgap optimization of two-dimensional photonic crystals using semidefinite programming and subspace methods, Journal of Computational Physics (229) (2010) 3706–3725.

[17] O. Sigmund, K. Hougaard, Geometric properties of optimal photonic crystals, Physical Review Letters 100 (153904).

[18] M. Qiu, S. He, Optimal design of a two-dimensional photonic crystal of square lattice with a large complete two-dimensional bandgap, Journal of the Optical Society of America B 17 (6) (2000) 1027–1030.

[19] J. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, Ann Arbor, MI, 1975.

[20] M. P. Bendsøe, O. Sigmund, Topology Optimization: Theory, Methods and Applications, Springer, 2003.

[21] G. Rozvany, M. Bendsøe, U. Kirsch, Layout optimization of structures, Applied Mechanics Reviews 48 (1995) 41–119.

[22] C. D. Chapman, K. Saitou, M. J. Jakiela, Genetic algorithms as an approach to configuration and topology design, ASME Journal of Mechanical Design 116 (4) (1994) 1005–1012.

[23] L. Shen, Z. Ye, S. He, Design of two-dimensional photonic crystals with large absolute band gaps using a genetic algorithm, Physical Review B 68 (2003) 035109.

[24] R. P. Drupp, J. A. Bossard, D. H. Werner, T. S. Mayer, Single-layer multiband infrared metallodielectric photonic crystals designed by genetic algorithm optimization, Applied Physics Letters 86 (2005) 081102.

[25] S. Preble, M. Lipson, H. Lipson, Two-dimensional photonic crystals designed by evolutionary algorithms, Applied physics letters 86 (2005) 061111.

[26] P. R. Villeneuve, M. Piché, Photonic bandgaps: What is the best numerical representation of periodic structures?, Journal of Modern Optics 41 (2) (1994) 241–256.

[27] P. R. Villeneuve, M. Piché, Photonic band gaps of transverse-electric modes in two-dimensionally periodic media, Journal of the Optical Society of American A 8 (8) (1991) 1296–1305.

[28] P. R. Villeneuve, M. Piché, Photonic band gaps in periodic dielectric structures, Progress in Quantum Electronics 18 (1994) 153–200.

[29] H. Y. D. Yang, Finite difference analysis of 2-D photonic crystals, IEEE Transactions on Microwave Theory and Techiques 44 (12(2)) (1996) 2688–2695.

[30] K. Bierwirth, N. Schulz, F. Arndt, Finite-difference analysis of rectangular dielectric waveguide structures, IEEE Transactions on Microwave Theory and Techiques 34(11) (1986) 1104–1114.

[31] R. S. Varga, Matrix Iterative Analysis, Prentice Hall, 1962.

[32] F. Bloch, Über die quantenmechanik der elektronen in kristallgittern, Zeitschrift für Physik A Hadrons and Nuclei 52 (1929) 555–600.

[33] G. Floquet, Sur les équations différentieller linéaries à coefficients périodiques., Annales Scientifiques de l'École Normale Supérieure 12 (1883) 47–88.