IOPscience iopscience.iop.org

Home Search Collections Journals About Contact us My IOPscience

Application-tailored optimisation of photonic crystal waveguides

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2016 J. Opt. 18 115005

(http://iopscience.iop.org/2040-8986/18/11/115005)

View the table of contents for this issue, or go to the journal homepage for more

Download details:

IP Address: 137.122.64.35
This content was downloaded on 18/10/2016 at 16:20

Please note that terms and conditions apply.

You may also be interested in:

Beyond the effective index method: improved accuracy for 2D simulations of photonic crystal waveguides
Sebastian A Schulz, Anthony H K Park, Israel De Leon et al.

Dispersion engineered slow light in photonic crystals: a comparison
S A Schulz, L O'Faolain, D M Beggs et al.

Low dispersion slow light waveguide with high coupling efficiency
Lei Dai and Chun Jiang

Silicon nitride membrane photonics
W H P Pernice, M Li, D F G Gallagher et al.

Multi-objective constrained design of nickel-base superalloys using data mining- and thermodynamics-driven genetic algorithms
Edern Menou, Gérard Ramstein, Emmanuel Bertrand et al.

Multi-objective optimization of two-dimensional porous phononic crystals
Hao-Wen Dong, Xiao-Xing Su and Yue-Sheng Wang

IOP Publishing Journal of Optics
J. Opt. 18 (2016) 115005 (8pp) doi:10.1088/2040-8978/18/11/115005

Application-tailored optimisation of photonic crystal waveguides

Sean Billings[1,2], Sebastian A Schulz[2], Jeremy Upham[2] and Robert W Boyd[2,3]

[1] School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, Canada
[2] Department of Physics, University of Ottawa, Ottawa, Canada
[3] Institute of Optics, University of Rochester, Rochester, NY, USA

E-mail: sschulz@uottawa.ca

Received 6 May 2016, revised 9 September 2016
Accepted for publication 27 September 2016
Published 17 October 2016

Abstract

Photonic crystal (PhC) waveguides are used for a wide range of applications with diverse performance metrics. A waveguide optimised for one application may not be suitable for others and no one-size-fits-all solution exists. Therefore each application requires a specialised waveguide design, a computationally and time intensive process. Here, we present a hybrid, multi-objective optimisation routine for PhC waveguides, to efficiently guide the device design. The algorithm can be configured to optimise for a wide range of performance metrics and applications. We demonstrate optimisations for three different applications: slow light performance, propagation loss due to fabrication disorder and delay line applications. For each optimisation target, our routine quickly finds practical waveguide designs (<48 h, on a laptop computer) that match or exceed the performance of state-of-the-art devices designed by the community over the last 10 years. This is also the first time that scattering loss from fabrication disorder has been incorporated into an optimisation algorithm, ensuring realistic predictions of a PhC waveguide design's practical performance.

Keywords: photonic crystals, optimisation algorithm, design

(Some figures may appear in colour only in the online journal)

Introduction

Slow light in photonic crystal (PhC) waveguides [1–3] has potential for a wide range of applications, e.g nonlinear optics [4, 5] or optical delay lines [6]. These applications are enabled through the processes of dispersion [7–10] and loss engineering [11, 12]. These engineering processes involve the tuning of the optical characteristics of a PhC waveguide through slight modifications to the lattice structure. While many successful PhC structures have been developed through systematic parameter sweeps and human intuition, the available design space is too large to fully explore manually. Therefore the computational optimisation of PhCs has revealed itself as a promising design method [13–21].

Previous work on the optimisation of PhC waveguides has been focused on the group index and the low dispersion bandwidth [15, 16], neglecting other figures of merit (FOMs). There are two big issues with this approach. First, the topic of propagation losses in PhC waveguides, often a limiting factor, has not been addressed in connection with optimisation algorithms. Second, the previous algorithms have only been demonstrated for the aforementioned FOMs (group index and bandwidth), neglecting the fact that different applications have different FOMs and optimial device design.

In this letter we demonstrate a flexible optimisation algorithm for the design of application-tailored PhC waveguides. The algorithm is capable of optimising for any definable FOMs, based on an application's particular needs, including propagation loss. A key issue for PhC optimisation is the often conflicting nature of the relevant FOMs [7], for example, the inherent trade-off between the group index and operating bandwidth of a waveguide. We therefore include a multi-objective optimisation algorithm, which can address a multitude of FOMs simultaneously.

Our algorithm consists of a hybrid optimisation scheme, combining a multi-objective evolutionary algorithm, for fast

 © 2016 IOP Publishing Ltd Printed in the UK

generation of a family of solutions, with a gradient descent algorithm, for optimisation of each individual solution towards a local optimum. The design space chosen for this paper includes hole position and radius variations. Viability of results is ensured through constraints on the design parameters, which ensure that it is feasible to fabricate the returned structures.

To demonstrate the robustness and flexibility of our algorithm, we investigate waveguide designs for three different sets of FOMs. These sets are (1) the group index and bandwidth, for comparison with the state-of-the art dispersion engineered waveguides, (2) the group index and propagation loss, to further explore the loss engineering of PhC waveguides, and (3) we also consider the FOMs inherent to the design of an optical delay line. Finaly, we compare the results of our waveguide optimisation, which is performed using 2D simulations, with 3D simulations. The wide range of applications considered shows the feasibility of of our algorithm as a design tool through the favourable comparison of our results with the current state-of-the-art.

PhC waveguide optimisation

A good optimisation algorithm has to fulfil several requirements. It should be capable of optimising for multiple FOMs, even if there is an inherent trade-off between the FOMs under consideration. It should be flexible, allowing optimisation for different combinations of FOMs. It should provide a good coverage of the solution space, while also ensuring that each solution is at or close to the local optimum. And last but not least, the results should quickly converge and be good, i.e. comparable to or better than the state-of-the-art.

We address these requirements by combining two optimisation approaches into a hybrid optimisation algorithm, consisting of a multi-objective algorithm that quickly explores a large solution space for FOMs with inherent trade-off, with a single objective algorithm that ensures convergence towards local optima.

Objectives

Many FOMs can be used as objectives during PhC waveguide design, with the relevant choice of optimisation objectives depending on the target application. To address the general case, independent of the specific application under consideration, we represent an arbitrary set of objectives as an n-dimensional vector $[o_1, \ldots, o_n]$.

Design space

For our simulations, we consider W1 based PhC waveguides consisting of a Si membrane ($n = 3.48$) surrounded by air ($n = 1$) with a lattice period, a, of 400 nm, although our algorithm can be applied to other high-refractive-index-contrast material systems. Our chosen design space consists of hole shifts perpendicular [10] and parallel [22] to the waveguide, as well as variations of the hole radii [9]. These

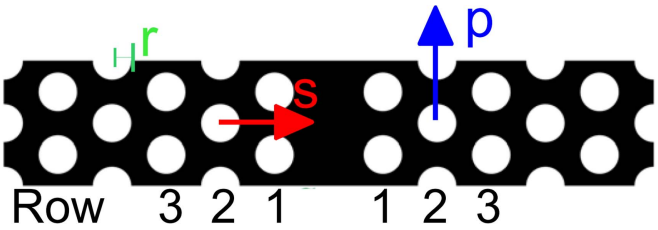

Figure 1. Sketch of a PhC outlining the design space parameters. A shift perpendicular to the waveguide is denoted by p, a shift parallel by s, and the hole radius by r. Rows are numbered going outwards from the waveguide and all variations are applied symmetrically to both sides of the waveguide.

variations are limited to the first three rows of the PhC lattice (see figure 1). We adopt the following conventions to describe the design parameters:

- p_n denotes a shift of the nth row parallel to the waveguide
- s_n denotes a shift of the nth row perpendicular to the waveguide
- r_n for $n \neq 0$ denotes the hole radius for the nth row
- r_0 denotes the hole radius for the remaining holes in the PhC.

We can represent an arbitrary solution in the design space with the vector $\overrightarrow{x} = [r_0, r_1, r_2, r_3, p_1, p_2, p_3, s_1, s_2, s_3]$.

Constraints

We introduce constraints on the design parameters range available to the optimisation algorithm, to account for fabrication limitations, for example the need for an interconnected membrane structure and minimum reproducible feature sizes. We enforce a minimum separation between holes of $0.1a$ (40 nm) and a minimum hole radius of $0.2a$ (80 nm), consistent with typical PhC fabrication [10, 11].

Parallel shifts p_n along the waveguide are periodic, with the period equal to the lattice constant a. In light of this periodicity, we reduce the design space under consideration by constraining the parallel shift of a row such that $-0.5a < p_n < 0.5a$.

Optimisation algorithm

Strength pareto evolutionary algorithm (*SPEA*)

The SPEA [23] is a genetic multi-objective optimisation algorithm, well suited to cases that involve a trade-off between diffeent objectives. Rather than finding a single optima, SPEA advances an initial set of solutions towards a front of well-distributed, non-dominated solutions in the objective space, using iterative randomisation, selection, and evolution.

A solution with objective o_n is non-dominated if there is no other solution (o_n'), such that ($o_i' > o_i$) for all objectives [23], i.e. no other solution outperforms this solution across all objectives. We use the convention that all objectives are to be maximised, and if necessary redefine our objectives

accordingly. For example, propagation loss, α, is undesired and should be minimised. The objective, o_i, corresponding to propagation loss is therefore evaluated as $1/\alpha$.

SPEA uses a fitness ('strength') evaluation metric to select the solutions from the population for evolution, and clustering to store a wide distribution of solutions [24]. This allows SPEA to generate a wide range of non-dominated solutions. This approach is particularly effective when there is no single 'best' solution in the objective space, as is typically the case when there is a trade-off between different objectives, e.g. group index and bandwidth. This broad coverage of the possible solution space addresses one of the key requirements for a waveguide optimisation algorithm.

Relatively steep (RS) gradient descent

An approach to solving multi-objective optimisation problems is to combine the relevant objectives into a single scoring function S, for example through a weighted sum, and then apply single-objective optimisation routines [23], such as gradient descent [25].

An approximated gradient $\hat{\nabla}S$ is used to iteratively minimise S. This iterative minimisation works by updating an initial solution vector $\overrightarrow{x_n}$ via $x_{n+1} = x_n - c\hat{\nabla}S$. Repeated iterations of this transformation ensure convergence of a solution to a local minimum, where $\hat{\nabla}S \approx 0$, in the neighbourhood of the initial solution [25].

The RS gradient descent algorithm uses a variation of backtracking [25] to scale the descent direction $c\hat{\nabla}S$. Backtracking is a method of scaling changes to $\overrightarrow{x_n}$, which ensures that the descent of a solution is convergent. Backtracking iteratively reduces the magnitude of descent in order to ensure that changes to $\overrightarrow{x_n}$ are sufficiently small. In contrast, the RS gradient descent variant alternates between maximising and minimising the magnitude of the descent direction. The variations of a solution $\overrightarrow{x}$ in the design space can correspond to much larger fluctuations in the scored objectives. Therefore, it is more efficient to assume that the magnitude of the optimal descent direction is proportionately smaller than the full gradient in most cases. However, when a solution approaches a local minima, the magnitude of the steepest descent direction converges to the gradient, and in this case we maximise the magnitude of the descent direction relative to the gradient.

The gradient descent approach allows one to use any single objective scoring function S, and thus it is very flexible.

SPEARS

SPEA, like other Pareto algorithms, leads to a structurally diverse set of solutions while consistently advancing the Pareto front during optimisation. These solutions serve as excellent starting points for RS gradient descent. Applying gradient descent ensures faster convergence of solutions to nearby local minima compared to SPEA alone. Furthermore, the RS variant of gradient descent is particularly adapted to the sensitive behaviour of PhC FOMs. Our hybrid algorithm, the SPEARS optimisation routine, does exactly the step described above. It sequentially combines SPEA, to gain effective coverage of solutions in the objective space, with RS gradient descent, to optimise these solutions locally. A general case is shown in figure 2. Initial solutions are generated in the n-dimensional design space and their results for two FOMs, objective 1 and 2, are calculated (white dots). First the SPEA algorithm optimises these initial results to a collection of non-dominated solutions (black dots). It also ensures that our final solution set has good coverage across the objective space and provides a 'front' of solutions rather than isolated clusters. The subsequent RS algorithm allows each these diverse solutions to efficiently converge to their respective local maxima (arrows). This routine is highly customisable and can be easily extended to higher dimension objective spaces and to account for larger numbers of FOMs.

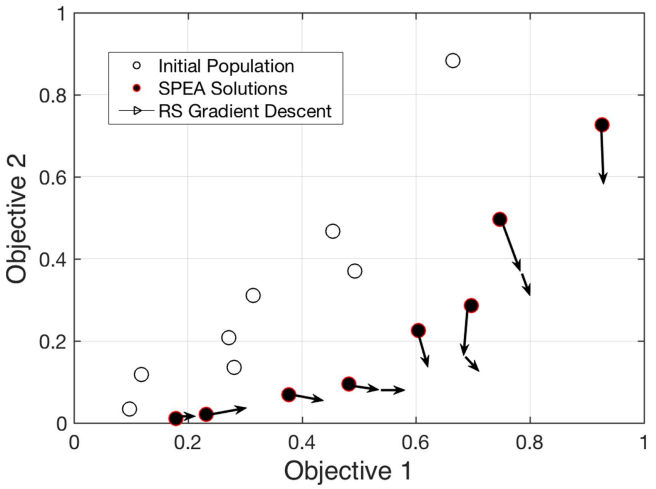

Figure 2. Schematic illustrating how the SPEARS algorithm uses an initial population (white dots) to find a family of optimised solutions (black dots) that span the objective space and quickly converge to their optimal solutions (black arrows). Multiple arrows versus a single arrow for a solution indicate that different solutions can take a different number of iterations during the gradient descent step before convergence is achieved. Here the initial population is optimised to maximise objective 1 and minimise objective 2.

Simulation implementation

We use the MIT photonics band plane wave solver (MPB) [26] to solve the PhC waveguides. As 3D simulations would be too resource-consuming for an optimisation algorithm without migrating to large scale computing systems, we use the effective period approximation [27] to solve a 2D equivalent of each waveguide (a reduction in the computation time by a factor of 100–1000). Our MPB base code [27, 28] not only solves for the PhC band structure, but also for the propagation loss [11, 27]. All results presented within this paper are obtained through simulation runs on a laptop computer equiped with a 2.2 GHz Intel Core i7 processor, running 6 threads in parallel. The total memory requirement is less than 6 GB and an optimisation run is completed in less than 48 h. For all applications that we consider the standard W1 waveguide was used as a seed solution, i.e. no prior user knowledge of PhC waveguide design is assumed, unless

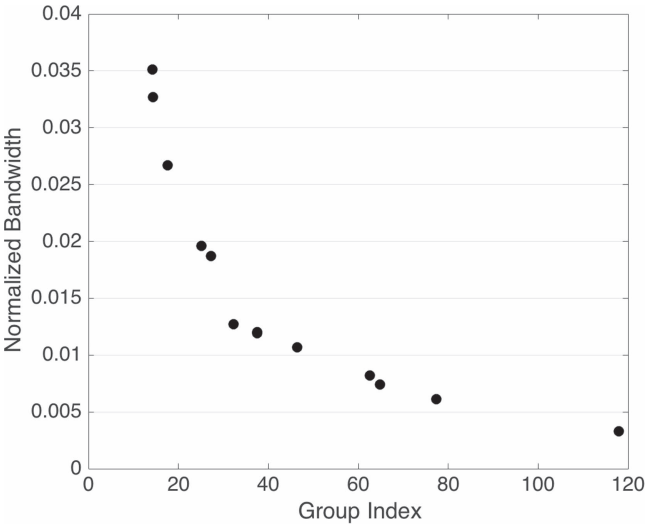

Figure 3. Distribution of solutions in the group index—bandwidth space. Note that the solutions, all of which are local maxima of the group index bandwidth product, span both solutions with high group indices and those with high bandwidths.

stated otherwise. The waveguide is assumed to be in a 220 nm thick silicon slab.

Band structure optimisation

The group index, bandwidth and the group-index-bandwidth product (GBP) are the standard FOMs for slow light waveguides [7] and have been used extensively to compare different dispersion engineering approaches.

We use the standard definitions of the group index (n_g), normalised bandwidth ($\Delta\omega$) and GBP:

$$n_g = \frac{d\omega}{dk}, \tag{1}$$

$$\Delta\omega = \frac{\omega_p - \omega_m}{\omega_0}, \tag{2}$$

$$\mathrm{GBP} = n_g \Delta\omega. \tag{3}$$

Here, ω_p and ω_m are the frequencies at which the group index reaches plus or minus 10% compared to the group index at the central frequency, ω_0. The bandwidth and the group index of a PhC waveguide are competing properties, i.e as the group index increases the bandwidth decreases and vice versa. Therefore the GBP has been used as the standard objective for assessing waveguide optimisation aproaches [16]. Typically waveguides have group indices in the range of $10 < n_g < 100$, with GBP values up to 0.3 (experiment) and 0.5 (theory) [7, 29].

We utilise the SPEARS routine by applying SPEA over both the group index and bandwidth FOMs to generate a set of non-dominated solutions, and then apply RS gradient descent using GBP as a scoring function to further optimise the generated solutions. Figure 3 shows the distribution of optimisation results from a single execution of the SPEARS optimisation routine. Table 1 presents a selection of 10 designs from this solution front. Several solutions (6–8 from table 1) are state-of-the-art PhCs in terms of the GBP. The remaining solutions demonstrate designs with high GBP for the associated group index. These results demonstrate that the SPEARS routine can quickly generate optimised slow light waveguide designs that are on-par with the state-of-the-art.

Propagation loss minimisation

While the design of high GBP waveguides is important, waveguide designs that exhibit high propagation loss are often impractical and not suited for many applications. Furthermore, there is a strong dependence of propagation loss on the group index [11], another key FOM for many applications, once again necessitating a trade-off during device design.

We optimise for these two FOMs, the group index, n_{g0}, and the loss computed at this group index α_0. It is worth noting that our design space for this application includes only radius variations and hole shifts perpendicular to the waveguide. We have previously shown [30] that parallel hole shifts, p_n, do generally not result in a reduction of the propagation loss and initial optimisation runs over the full design space confirmed this. Runs over the reduced design space yield a family of solutions, see figure 4, that once again meet or exceed the current state-of-the-art. For example, a waveguide with the design parameters [$r_0 = 0.214$, $r_1 = 0.2$, $r_2 = 0.285$, $r_3 = 0.291$, $s_3 = -0.0047$, $s_2 = -0.00041$, $s_1 = -0.0489$], which has a predicted propagation loss of $\alpha_0 = 46\,\mathrm{dB\,cm^{-1}}$ at $n_{g0} = 37$, matching the best reported propagation loss in this group index range ($\alpha_0 = 50\,\mathrm{dB\,cm^{-1}}$ at $n_{g0} = 36$ [11]).

PhC delay line design

Optical delay lines, especially those with tunable delay values, are a key application for PhC waveguides [2, 3], and have been demonstrated experimentally [2, 6, 31–33]. The working principle of a tunable delay line is simple: a PhC waveguide supports both fast (n_g of approximately 5) and slow light regions. An external tuning mechanism, such as thermal tuning [32, 34], electrical tuning [34] or ultrafast, adiabatic frequency tuning [33] can then be used to select which of these regions a signal pulse will experience and therefore the time taken to travel through the device.

Ideally, a delay line should provide the largest possible delay, over the largest possibly bandwidth, with the smallest footprint, the least amount of group velocity dispersion and the lowest propagation loss. Here, we optimise a waveguide for operation as a static element for a $100\,\mathrm{Gbit\,s^{-1}}$ delay line (assuming 10 ps pulses).

For a given delay, the footprint is dependent on the group velocity (and therefore the group index) of the PhC waveguide, with a larger group index leading to a smaller footprint. Therefore, the group index will be the first objective for our optimisation routine. The second optimisation objective is the normalised bandwidth $\Delta\omega$, as defined in equation (2). As the last objective, we choose the maximum achievable delay τ_m

Table 1. SPEARS generated GBP designs. The bandwidth is normalised to the central frequency, see equation (2), and the radii and position shift parameters are normalised to the lattice constant.

	GPB	n_g	Bandwidth	r_0	r_1	r_2	r_3	p_1	p_2	p_3	s_1	s_2	s_3
1	0.467	64.8	0.0074	0.233	0.200	0.384	0.200	−0.136	−0.0422	0.0059	−0.0693	−0.0756	−0.0654
2	0.447	37.5	0.0119	0.228	0.238	0.4	0.2	0.125	0.179	0.0309	0.0639	−0.0978	−0.0761
3	0.392	118	0.0033	0.233	0.2	0.384	0.2	−0.136	0.0051	0.0045	−0.0697	−0.0757	−0.0656
4	0.484	25.1	0.0196	0.2	0.2	0.311	0.2	0.125	0.211	−0.0427	−0.133	−0.0954	−0.078
5	0.465	77.4	0.0061	0.237	0.203	0.384	0.2	−0.142	0.0289	0.0283	−0.0625	−0.102	−0.0467
6	0.493	62.5	0.0082	0.237	0.2	0.384	0.2	−0.135	0.005	0.0066	−0.625	−0.0948	−0.0467
7	0.455	14.4	0.0327	0.207	0.2	0.278	0.227	0.226	0.306	−0.0012	0.198	0.161	0.0254
8	0.491	46.4	0.0107	0.237	0.2	0.384	0.2	−0.113	0.005	0.0066	−0.0624	−0.0989	−0.0466
9	0.491	27.2	0.0187	0.2	0.261	0.356	0.2	−0.142	−0.183	−0.0505	−0.1349	−0.0922	0.0808
10	0.484	14.2	0.0351	0.2	0.2	0.264	0.232	−0.162	−0.211	−0.0427	−0.122	−0.153	−0.0704

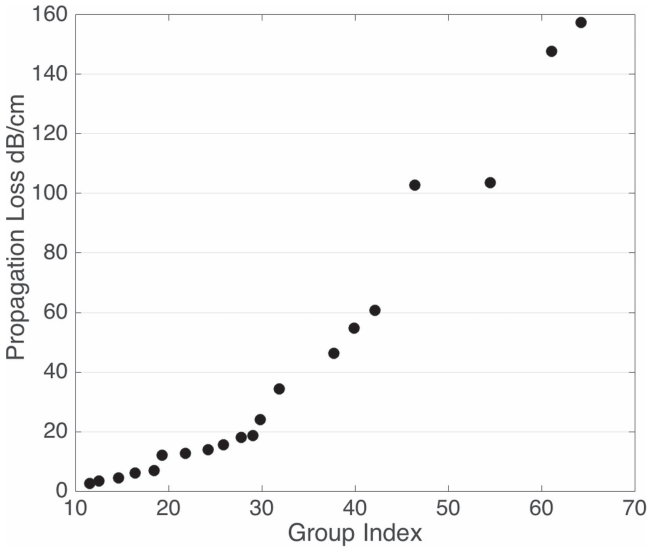

Figure 4. Family of solutions for an optimisation run for propagation loss and group index, over a reduced design space (hole shifts perpendicular to the waveguide and radius reduction only).

for a given design. This delay is limited either by the group velocity dispersion or the propagation loss and our routine is configured to automatically assess which factor is most limiting for a given design.

As in previous work, we set the total acceptable propagation loss for a device to be $IL_{\mathrm{m}} = 10$ dB [7]. Using this constraint, we derive the loss-limited device length:

$$L_{\alpha} = IL_{\mathrm{m}}/\alpha. \tag{4}$$

When an optical pulse encounters group velocity dispersion the pulse width is increased, leading to a reduced signal to noise ratio, increased bit error rate and, in the worst case, an overlap with following pulses. We use the acceptable pulse width broadening length L_{D} from [35], which is defined as follows, assuming an initial pulse width of $\tau_0 = 10$ ps:

$$L_{\mathrm{D}} = \frac{\tau_0^2}{(4\ln 2)\beta_2}. \tag{5}$$

In the above, $\beta_2 = \delta^2 k/\delta\omega^2$ is the group velocity dispersion coefficient. The loss- and dispersion-limited delay are then related to the limiting device length by

$$\tau_{\alpha} = \frac{n_{\mathrm{gs}} - n_{\mathrm{g0}}}{c} L_{\alpha}, \tag{6}$$

and

$$\tau_{\beta_2} = \frac{n_{\mathrm{gs}} - n_{\mathrm{g0}}}{c} L_{\mathrm{D}}, \tag{7}$$

respectively. Here, n_{gs} is the group index in the slow light region and n_{g0} is the group index in the fast light region. For simplicity and without loss of generality we assume $n_{\mathrm{g0}} = 5$ throughout our analysis.

Both τ_{β_2} and τ_{α} constrain the maximum achievable delay of a device. To ensure that both constraints are met, we evaluate the maximum achievable, tunable delay τ_{m} as the minimum of the two limiting values, i.e.:

$$\tau_{\mathrm{m}} = \min(\tau_{\alpha}, \tau_{\beta_2}). \tag{8}$$

Our optimisation routine quickly produces designs with delays exceeding 700 ps, once again meeting or exceeding the state of the art (217 ps [36]). A collection of these solutions is seen in table 2. In fact, these delay values are comparable to those of other types of intricate delay line structures such as ring-resonator-based delay lines.

Figure 5 shows the maximum tunable delay τ_{m} versus the group index n_{g}. In many devices, e.g. optical resonators, the bandwidth and achievable delay share a trade off similar to that of group index and bandwidth, unless the system is dyamically tuned [37] and it is generally assumed that this also applies to PhC waveguides. However this behaviour is not observed here. Instead, devices with an increased delay show a larger operating bandwidth. We attribute this behaviour to the fact that our devices are mainly loss limited. As seen in figure 5, a large delay is correlated with a lower operating group index, and therefore a larger bandwidth.

Comparison with 3D simulations

As mentioned earlier, the optimisation routine is implemented using 2D simulations, with the effective period method [27], as 3D simulations are too time consuming for a full optimisation run. In this section we assess the accuracy of the optimisation results when compared to 3D simulations of the structure with the same design parameters. We chose a set of the most interesting results that match or exceed the current state-of-the-art.

Performing 3D simulations of designs generated by the SPEARS algorithm (in 2D) shows that the 2D results generally underestimate the group index. This is consistent with the results in [27]. It is worth reiterating here that other 2D approximation methods, for example the effective index method, also underestimate the group index and have in fact a bigger discrepancy to the 3D results [27]. The effect of the underestimation of the group index on the results of optimised waveguides depends on the FOMs used for optimisation. As far as the GBP is concerned, a change in the group index will also lead to a change in the bandwidth and depending on the exact dispersion curve this means that the 2D simulations might under- or overestimate the GBP. For example, in our simulation run most of the devices had an overengineered dispersion curve, with multiple peaks in the group index spectrum. Here an underestimation of the group index also underestimates the group velocity dispersion and hence overestimates the bandwidth. Therefore, in our case all of the GBP results are an overestimation of the true 3D results. For example a 2D GBP of 0.4843 (result 10 in table 1) reduces to 0.2688 in the 3D simulations. For other FOMs the change is more systematic and clearer. When modelling propagation loss, an underestimation of the group index implies that in the real 3D structure the low scattering region is at a slightly higher group index, making a comparison with the state of the

Table 2. SPEARS generated delay designs.

	Tunable delay (ps)	Group index	Bandwidth	r_0	r_1	r_2	r_3	s_1	s_2	s_3
1	775	14.1	0.0216	0.204	0.2	0.2	0.269	−0.0701	−0.0035	−0.0061
2	858	12.8	0.0137	0.249	0.2	0.205	0.249	−0.0358	−0.0004	−0.0046
3	258	40.8	0.0044	0.2	0.2	0.278	0.280	−0.0329	0.0013	0.0019
4	159	50.1	0.003	0.249	0.2	0.307	0.249	−0.049	−0.0013	0.0009
5	388	27.2	0.0064	0.237	0.2	0.268	0.261	−0.0649	0.0044	−0.0067
6	627	17.6	0.0188	0.221	0.2	0.223	0.297	−0.0299	0.005	−0.0024
7	515	21.6	0.0081	0.233	0.2	0.244	0.261	−0.0457	0.0048	−0.0047
8	339	35.0	0.0088	0.2	0.2	0.271	0.277	−0.0459	−0.0004	−0.0046
9	837	13.4	0.0169	0.225	0.2	0.2	0.261	−0.046	0.0037	−0.0067
10	105	76.2	0.0032	0.212	0.2	0.308	0.249	−0.046	−0.0021	0.0009

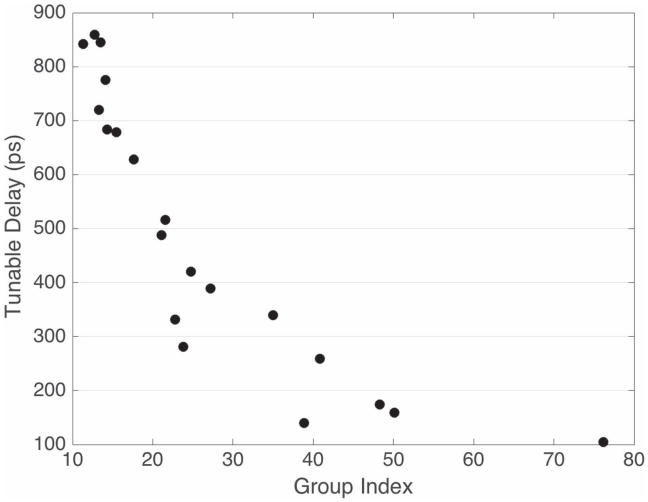

Figure 5. Distribution of delay solutions in terms of group index.

art even more favourable. In our chosen example, the 2D structure had a loss of $\alpha_0 = 46$ dB cm^{-1} at $n_g = 37$, while the 3D simulations showed an even lower propagation loss ($\alpha_0 = 36$ dB cm^{-1}) at a slightly higher group index ($n_g = 41.5$). For the delay-line application, both the propagation loss and bandwidth are of importance. However, since most devices are loss-limited, and the propagation loss is well approximated by the 2D simulations, most of the delay-line results are also in reasonable agreement with 3D simulations. For example, a design with a predicted delay of 676 ps from 2D simulations has a predicted delay of 520 ps in the 3D simulations, still comparing favourably with the current state-of-the-art

Conclusion

We demonstrate that the SPEARS optimisation routine is well suited for PhC waveguide design. Starting from a standard W1 waveguide it very quickly generates state-of-the-art solutions for diverse optimisation objectives, with low computational requirements (less than 48 h on a quad core processor, less than 6 GB RAM). Longer simulation runs or runs with high performance computers, both beyond the scope of this paper, should result in further improvements in performance, far in excess of the current state-of-the-art. We demonstrate the flexibility of this approach by optimising for different applications, including a tunable delay line and the first demonstration of optimisation for low propagation loss at high group indices, achieving better than state-of-the art solutions in all cases. The SPEARS routine can easily be configured for different applications, through a modification of the optimisation objectives.

All resources necessary for the SPEARS routine will be available online [38] and we invite other researchers to make use of this tool in their own work.

We want to reiterate that our optimisation was performed using 2D simulations and that before fabrication of devices all results from 2D simulations should always be corroborated using 3D modelling. Our own comparison with 3D models showed that the quality of results is dependent on the FOMs under consideration, although in all cases the results were of sufficient quality to serve as a starting point for 3D finalisation of the device design, which can then be performed over a much more limited design space.

Acknowledgments

The authors thank M J Huttunen for fruitful discussion. This work was funded by the Canada Excellence Research Chair (CERC) program.

References

[1] Krauss T F 2007 Slow light in photonic crystal waveguides *J. Phys. D: Appl. Phys.* **40** 2666–70
[2] Baba T 2008 Slow light in photonic crystals *Nat. Photon.* **2** 465–73
[3] Krauss T 2008 Why do we need slow light? *Nat. Photon.* **2** 448–50
[4] Corcoran B, Monat C, Grillet C, Moss D J, Eggleton B J, White T P, O'Faolain L and Krauss T F 2009 Green light emission in silicon through slow-light enhanced third-harmonic generation in photonic-crystal waveguides *Nat. Photon.* **3** 206
[5] Monat C, De Sterke M and Eggleton B J 2010 Slow light enhanced nonlinear optics in periodic strucutres *J. Opt.* **12** 104003

[6] Mori D and Baba T 2004 Dispersion-controlled optical group delay device by chirped photonic crystal waveguides *Appl. Phys. Lett.* **85** 1101–3
[7] Schulz S A, O'Faolain L, Beggs D M, White T P, Melloni A and Krauss T F 2010 Dispersion engineered slow light in photonic crystals: a comparison *J. Opt.* **12** 104004
[8] Petrov A Y and Eich M 2004 Zero dispersion at small group velocities in photonic crystal waveguides *Appl. Phys. Lett.* **85** 4866
[9] Frandsen L, Lavrinenko A V, Fage-Pederson J and Borel P 2006 Photonic crystal waveguides with semi-slow light and tailored dispersion properties *Opt. Express* **14** 9444–50
[10] Li J, White T, O'Faolain L, Gomez-Iglesias A and Krauss T 2008 Systematic design of flat band slow light in photonic crystal waveguides *Opt. Express* **16** 6227–32
[11] O'Faolain L, Schulz S A, Beggs D M, White T P, Spasenovic M, Kuipers L, Morichetti F, Melloni A, Mazoyer S, Hugonin J P, Lalanne P and Krauss T F 2010 Loss engineered slow light waveguides *Opt. Express* **18** 27627–38
[12] Li L, O'Faolain L, Schulz S A and Krauss T F 2012 Low loss propagation in slow light photonic crystal waveguides at group indices up to 60 *Photon. Nanostruct.* **10** 589–93
[13] Goh J, Fushman I, Englund D, Goh J V, Fushman I, Englund D and Vučković J 2007 Genetic optimization of photonic bandgap structures *Opt. Express* **15** 8218
[14] Jensen J S and Sigmund O 2011 Topology optimization for nano-photonics *Laser Photon. Rev.* **5** 308–21
[15] Mirjalili S, Abedi K and Mirjalili S 2013 Optical buffer performance enhancement using particle swarm optimization in ring-shape-hole photonic crystal waveguide *Optik* **124** 5989
[16] Mirjalili S, Mirjalili S, Lewis A and Abedi K 2014 A tri-objective particle swarm optimizer for designing line defect photonic crystal waveguides *Photon. Nanostruct.* **12** 152
[17] Mirjalili S, Mirjalili S and Lewis A 2013 A novel multi-objective optimization framework for designing photonic crystal waveguides *IEEE Photon. Technol.* **26** 146–9
[18] Preble S, Lipson M and Lipson H 2005 Two-dimnesional photonic crysta; designed by evolutionary algorithms *Appl. Phys. Lett.* **86** 061111
[19] Dripp R, Bossard J, Werner D and Mayer T 2005 Single-layer multiband infrared metallodielectric photonic crystals designed byb genetic algorithm optimization *Appl. Phys. Lett.* **86** 081102
[20] Minkov M and Savona V 2014 Automated optimization of photonic crystal slab cavities *Sci. Rep.* **4** 5124
[21] Lai Y, Pirotta S, Urbinati G, Gerace D, Minkov M, Savonna V, Badolato A and Galli M 2014 Genetically designed L3 photonic crystal nanocavities with measured quality factor exceeding one million *Appl. Phys. Lett.* **104** 241101
[22] Hamachi Y, Kubo S and Baba T 2009 Slow light with low dispersion and nonlinear enhancement in a lattice shifted photonic crystal waveguide *Opt. Lett.* **34** 1072–4
[23] Marler R T and Arora J S 2004 Survery of multi-objective optimization methods for engineering *Struct. Multidisc. Optim.* **26** 369–95
[24] Zitzler E and Thiele L 1998 An evolutionary algorithm for multiobjective optimization: the strength pareto approach *TIK-Report* No. 43
[25] Nocedal J and Wright S J 2006 *Numerical Optization* 2nd edn (Berlin: Springer)
[26] Johnson S and Joannopolous J 2001 Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis *Opt. Express* **8** 173–90
[27] Schulz S A, Park A H K, De Leon I, Upham J and Boyd R 2015 Beyond the effective index method: improved accuracy for 2D simulations of photonic crystal waveguides *J. Opt.* **17** 075006
[28] S A Schulz and H K Park, https://github.com/sschulz365/PhC_simulations/tree/2D-simulations or https://researchgate.net/publication/288993468_2D_PhC_waveguides
[29] Zhao Y, Zhang Y-N, Wang Q and Hu H 2015 Review on the optimization methods of slow light in photonic crystal waveguide *IEEE Trans. Nanotechnol.* **14** 407
[30] Schulz S A, Whelan-Curtin W, Rey I H and Krauss T F 2011 Understanding propagation loss in slow light waveguides *Advanced Photonics, OSA Technical Digest (CD)* (Optical Society of America) paper SLTuB1
[31] Baba T, Kawasaki T, Sasaki H, Adachi J and Mori D 2008 Large delay-bandwidth product and delay tuning of slow light pulse in photonic crystal coupled waveguide *Opt. Express* **16** 9245–53
[32] Adachi J, Ishikura N, Sasaki H and Baba T 2010 Widerange tuning of slow light pulse in SOI photonic crystal coupled waveguide via folded chirping *Sel. Top. Quantum Electron.* **16** 192–9
[33] Beggs D M, Rey I H, Kampfrath T, Rotenberg N, Kuipers L and Krauss T F 2011 Ultrafast tunable optical delay line based on indirect photonic transistions *Phys. Rev. Lett.* **108** 213901
[34] Hayakawa R, Ishikura N, Nguyen H C and Baba T 2013 High-speed delay tuning of slow light in pin-diode-incorporated photonic crystal waveguide *Opt. Lett.* **38** 2680–2
[35] Melloni A, Morichetti F, Ferrari C and Martinelli M 2008 Continuously tunable 1 byte delay in coupled-resonator optical waveguides *Opt. Lett.* **33** 2389–91
[36] Lin C-Y, Subbaraman H, Hosseini A, Wang A X, Zhu L and Chen R T 2012 Silicon nanomembrane based photonic crystal waveguide array for wavelength-tunable true-time-delay lines *Appl. Phys. Lett.* **101** 051101
[37] Upham J, Fujita Y, Kawamoto Y, Tanaka Y, Song B S, Asano T and Noda S 2013 The capture, hold and forward release of an optical pulse from a dynamic photonic crystal nanocavity *Opt. Express* **21** 3809–17
[38] https://github.com/sschulz365/PhC_Optimization