PAPER

Inverse design of plasma metamaterial devices with realistic elements

To cite this article: Jesse A Rodríguez and Mark A Cappelli 2022 *J. Phys. D: Appl. Phys.* **55** 465203

View the article online for updates and enhancements.

Inverse design of plasma metamaterial devices with realistic elements

Jesse A Rodríguez[*] and **Mark A Cappelli**

Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, United States of America

E-mail: jrodrig@stanford.edu

CrossMark

Abstract

In an expansion of a previous study (Rodríguez *et al* 2021 *Phys. Rev. Appl.* **16** 014023), we apply inverse design methods to produce two-dimensional plasma metamaterial devices with realistic plasma elements which incorporate quartz envelopes, collisionality (loss), non-uniform density profiles, and resistance to experimental error/perturbation. Finite difference frequency domain simulations are used along with forward-mode differentiation to design waveguides and demultiplexers operating under the transverse magnetic polarization. Optimal devices with realistic elements are compared to previous devices with idealized elements, and several parameter initialization schemes for the optimization algorithm are explored, yielding a robust procedure for producing such devices. Demultiplexing and waveguiding are demonstrated for microwave-regime devices composed of plasma elements with reasonable space-averaged plasma frequencies $\sim$10 GHz and a collision frequency $\sim$1 GHz, allowing for future *in-situ* training and experimental realization of these designs.

Keywords: plasma, metamaterial, inverse design, machine learning

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

1. Introduction

Inverse design is a technique used to optimize device parameters given an *a priori* set of design criteria and metrics. When applied to electromagnetically-active systems [1–18], devices are conceived by solving Maxwell's equations for a particular material domain and set of sources, the properties of which (e.g. relative permittivity, permeability, source frequency and power) are mapped to trainable parameters. These parameters can be modified to alter the propagation of the source fields through the device in a manner which fulfills the performance criteria. Typically, these performance criteria are encoded as an objective function that depends on the input parameters which is minimized or maximized via the inverse design algorithm. As an example, in many scenarios, the objective function involves an inner product between the desired electromagnetic fields and those obtained via simulation in regions of interest. Automatic differentiation allows the calculation of a numerical gradient of such an objective with respect to the parameters that encode the permittivity structure of the training region. The inverse design algorithm then iteratively adjusts these parameters to minimize the distance between the simulated fields and the desired field in the chosen inner product space. The result is a novel, highly specialized device that meets the design criteria of the user [1–3].

Unfortunately, the set of optical devices which can be experimentally realized is a small subset of the set of devices which can be computationally modeled, and so inverse design often works within a constrained setup where various limitations on the device geometry are imposed. For example, allowing for a continuous range of permittivities throughout the simulation domain results in fast optimization with very effective designs, but achieving continuous variation of permittivity throughout the design region is essentially impossible using conventional manufacturing techniques. In prior work, inverse

---

[*] Author to whom any correspondence should be addressed.

design practitioners take advantage of more constrained systems such as photonic crystal-style devices [1, 4–7] and metasurfaces [8–11] where a simple map from the parameterized design region to a physically realizable device is guaranteed to exist. These maps typically encode restrictions on the range and spatial distribution of the domain permittivity (i.e. using only permittivity values and geometry corresponding to available materials and manufacturing processes). A common technique in the literature is the use of nonlinear projection to design binarized photonic devices [12] which can be printed or created via lithography as a single-use design. Each method for constraining inverse design algorithms introduces its own limitations to the device capabilities, but even simplistic devices can be optimized to perform tasks which would normally require expert-level human design to achieve the same or worse performance.

Rather than restricting permittivity values within our domain to some discrete set, we focus instead on limiting the spatial configuration by parameterizing the design region as a plasma metamaterial (PMM). Our PMM is a square 10-by-10 array of cylindrical gaseous plasma elements (rods) which can be experimentally realized via the use of discharge lamps or laser-generated plasmas. Because plasma density (and therefore permittivity) in sources such as these can be continuously tuned through a wide range of values, our PMM serves as a good candidate for inverse design as the continuous variability of the permittivity of the elements yields an infinite configuration space for training purposes [19] and, perhaps more importantly, allows a single device to serve multiple functions. By changing the discharge current or laser power supplied to each of the plasma elements within the device, one could quickly switch between different device configurations. Previous work with inverse design electromagnetic devices achieved reconfigurability via methods such as refractive index changes [8] in materials like liquid crystals [10], although the use of our PMM configuration is particularly unique as completely disparate functions can be performed with a single device as can be seen in our prior study where the PMM could function as both a demultiplexer and a photonic logic gate [20]. The inverse design process for our PMM geometry is summarized in figure 1(a). In our algorithm, there is a one-to-one correspondence between each plasma element and the $n_x n_y$ elements of the parameter vector ρ (figure 1(b)). In the end, the parameterization scheme results in a relative permittivity matrix ε_r (of size $N_x \times N_y$, dependent on the simulation resolution) containing permittivity values at each pixel in the simulation domain.

In addition to the complex refractive behavior which can be accessed in PMM devices, when the PMM is composed of elements that are small compared to the operating wavelength and the source is polarized properly, we can couple into localized surface plasmon (LSP) modes along the boundary between a positive permittivity background and a negative permittivity plasma element. Such surface modes can yield more complex and efficient device behavior [21–23], allowing for a particularly high degree of reconfigurability, but for devices which seek to maximize transmission, they can cause serious degradation of performance [20]. Prior experimental studies of
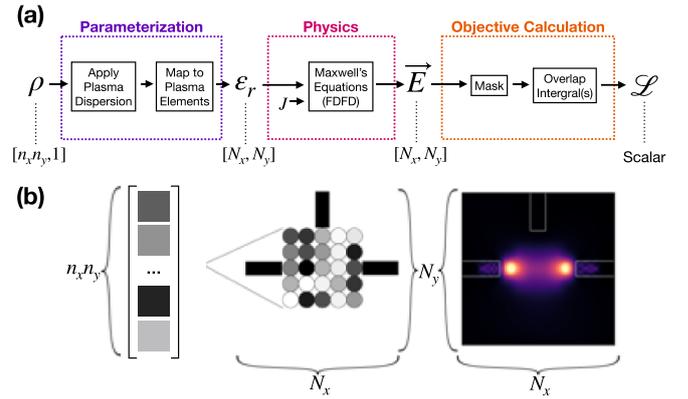

Figure 1. (a) Flow chart describing the algorithmic design of our PMM array. *J* represents the modal source for the FDFD simulation. Examples of the training parameter vector ρ, permittivity matrix ε_r and simulated **E** fields for a simple PMM device are provided in (b). Reprinted (figure) with permission from [20], Copyright (2021) by the American Physical Society.

plasma photonic crystals (PPCs) have already highlighted the richness of this geometry in waveguide [20, 24] and bandgap [21, 25, 26] devices. The complicated array of electromagnetic phenomena present in PMM devices is well-matched to inverse design methods where more subtle physical phenomena can be leveraged by the optimization algorithm in a manner that would be impossible in a human-designed device.

2. Methods

Figure 2 provides a schematic of our PMM device in a directional waveguide configuration. A demultiplexer configuration, described later in the paper, utilizes the exact same array of tunable plasma elements but with different entrance/exit waveguides. The device is simply a 10×10 array of tunable plasma discharge tubes suspended in air and spaced according to the limits of existing experimental facilites [24, 25]. A detailed description of the plasma elements is found below. We chose a 10×10 array for this study for two reasons; first, this device size is about as large as we can reasonably produce experimentally in the near future, and second, stepping up to a larger array severely limits our ability to train devices in a reasonable amount of time. Although the raw simulation time simply scales as the square of the device dimension since the simulations are two-dimensional, the current design of the optimization scheme requires a large amount of memory, and so going any larger than 10×10 would necessitate the use of shared high-performance computing (HPC) resources for device training rather than our lab's own resources; requiring long wait times that are not conducive to rapid readjustments to the training parameters. We do believe that larger arrays would enable the device to perform more complex functions, but this will be left to future investigations.

The PMM device functions by redirecting modal sources introduced at the entrance waveguide(s) to the desired output waveguide(s). In this work, the electromagnetic field solutions are calculated via finite difference frequency domain (FDFD)
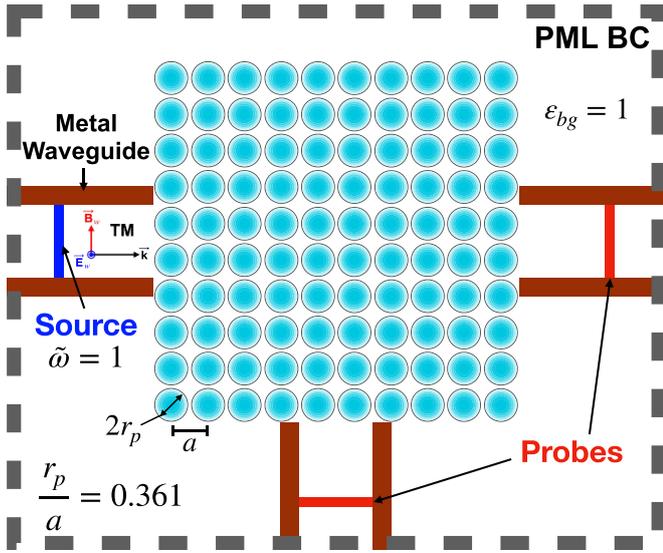
Figure 2. Schematic diagram for the PMM device in a waveguide configuration. The device consists of a 10×10 array of plasma discharge bulbs (cyan) suspended in air with metal waveguides (brown) functioning as the entrance and exit(s). The 'probe' slices (red) represent the domain over which the L_2 inner products that constitute the objective function are calculated. ε_{bg} is the background permittivity. $\tilde{\omega}$ is the nondimensionalized operating frequency $\omega/(2\pi c/a)$. Reprinted (figure) with permission from [20], Copyright (2021) by the American Physical Society.

simulations computed with Ceviche, an autograd-compliant electromagnetic simulation tool [2] that enables calculation of the gradient of the optimization objective with respect to the input parameters that encode the simulation domain. Note that the use of FDFD implies that all computed devices represent the steady state solution achieved after some characteristic time. The simulation domain was discretized using a uniform square mesh with a resolution of 75 pixels per lattice constant a due to the CPU memory and runtime limitations mentioned above, resulting in a domain size of $N_x = N_y = 1500$ pixels in each case. The 'optimal' devices were then run at higher resolutions up to 150 pixels per lattice constant on a HPC cluster to confirm that the device functionality was preserved; ensuring that the training resolution was high enough. A perfectly-matched boundary layer $2a$ in width was applied along the domain boundary.

The polarization of the input source has a strong effect in 2D devices such as these, either $\boldsymbol{E_z}$ (E out of the page), which we call the TM polarization, or $\boldsymbol{E_y}$ (H out of the page), which we call the TE polarization. Prior work with PPCs indicates that the response of the device to both TE and TM-polarized sources is highly tunable, with the former case benefiting (and suffering, depending on the desired device functionality) from the presence of LSP modes [21–23, 27], while the latter makes more direct use of dispersive and refractive effects [21, 24, 25]. In our previous study [20], we see that for devices which seek to preserve an input signal like the waveguides and demultiplexers we present below, the TE polarization leads to poor performance, so we choose to focus on the TM mode alone in this study. Thus, we will not expect to see any small, sub-wavelength-scale field structures around

our plasma elements or the large transmission losses which are indicative of the presence of LSP modes.

The simulated **E** fields are masked to compute the L_2-inner product integrals along planes of interest within the problem geometry. These integrals make up the optimization objective $\mathcal{L}(\rho)$. Our simulation tool, Ceviche, is then used to compute the numerical gradient of the objective function with respect to the training parameters via forward-mode differentiation [2]. Ceviche uses the Adam optimization algorithm [28] (gradient ascent with momentum, in essence) to iteratively adjust ρ and thereby maximize $\mathcal{L}$, or rather, minimize the L_2 distance between the simulated fields at the device output and the desired fields. Optimization was conducted with learning rates ranging from 0.0002 to 0.01, and with typical values for the Adam hyperparameters; $\beta_1 = 0.9$, and $\beta_2 = 0.999$.

To summarize, the PMM optimization problem can be expressed as

$$\max_{\rho} \quad \mathcal{L}(\mathbf{E})$$

$$\text{given} \quad \nabla \times \frac{1}{\mu_0} \nabla \times \mathbf{E} - \omega^2 \varepsilon(\omega, \rho)\mathbf{E} = -i\omega J.$$

where $\mathcal{L}$ is the objective composed of a set of L_2 inner product integrals of the simulated field with the desired propagation mode, **E** is the electric field, ρ is an $n_x n_y$-dimensional vector that contains the parameters that set the electron density of each of the PMM elements, μ_0 is the vacuum permeability, ω is the field frequency, and $\varepsilon(\omega, \rho)$ is the spatially-dependent permittivity that is encoded by ρ, and J is a current density used to define a fundamental modal source at the input waveguide. In practice, the permittivity distribution among the plasma elements is controlled by varying the discharge current (which therein alters the plasma density) in each of the PMM elements according to the Drude model, which is where the dependence on ω arises. The non-dimensionalized plasma frequencies in ρ are mapped to element permittivities via the Drude model;

$$\varepsilon = 1 - \frac{\omega_p^2}{\omega^2 + i\omega\gamma}$$

where γ is the collision/damping rate, $\omega_p^2 = \frac{n_e e^2}{\varepsilon_0 m_e}$ is the plasma frequency squared, n_e is the electron density, e is the electron charge, m_e is the electron mass, and ε_0 is the free-space permittivity.

The devices that follow have been explored previously in [20], but mostly as a proof of concept with very simplistic plasma elements that were collisionless, uniform-density rods with no quartz envelope. In this study, we attempt to model all major non-ideal factors in order to assess whether or not such devices will be readily experimentally realized. Real gas discharge plasmas will be affected by a range of non-ideal factors in addition to the damping rate included in the Drude model above.

Collisional damping in discharges can be brought to very low levels in practice by reducing the discharge pressure, but this will come at the expense of plasma density (and correspondingly, plasma frequency), which in turn restricts the range of values of the plasma dielectric constant that we can access.
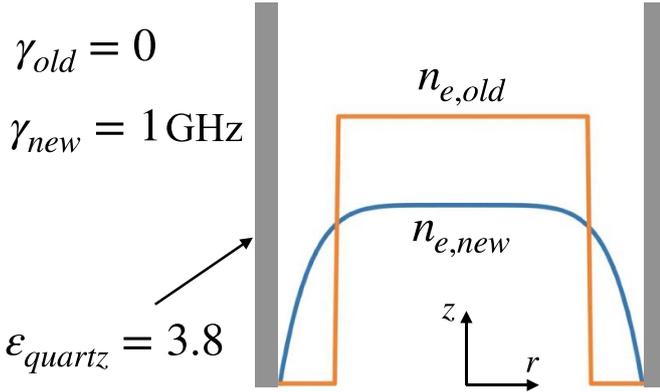
Figure 3. Illustration comparing our model (labeled 'new') and the idealized model (labeled 'old') from [20] for the plasma discharge tubes. The idealized model does not include the quartz bulb.

Furthermore, the plasma density profile within the gaseous elements is not uniform and dependent on the methods used to generate the plasma. For example, within a discharge tube at nominal operation conditions, the plasma has a core of relatively uniform density plasma and then a sheath region near the bulb where the density drops to negligible values at the inner bulb surface. As one would expect, the plasma density profile can have a strong effect on device performance in both the TE mode by affecting LSP coupling [22, 23] and in the TM mode by altering the refraction of impinging waves. In the past, our group has approximated this with a uniform plasma column of radius $r_i/\sqrt{2}$ where r_i is the inner radius of the bulb's quartz envelope. This leads to decent agreement with experiment in photonic crystal configurations [21, 24, 25, 27], but may not be appropriate for simulating devices with EM wave propagation structures as complicated as are found in our inversely designed PMMs. To best account for the non-uniform density profile, we choose as our model a 6th-order polynomial as a function of radius, where the trainable parameters ρ map to the space-averaged plasma frequency of each of the rods. Although a simple analysis considering diffusion and recombination at the plasma-quartz interface would suggest a Bessel function solution which would be better approximated by a 2nd-order polynomial, the 6th-order polynomial profile has been shown to lead to the closest match between simulated and observed transmission spectra for our PPC devices [29]. In addition, we include the $\varepsilon = 3.8$ quartz bulb around the plasma elements that was omitted in [20]. Figure 3 illustrates the difference between the simplistic plasma elements in [20] and our more realistic ones.

In addition to all of this, while an accurate and precise control of element permittivity by varying discharge current is possible in theory, in practice a number of factors (such as thermal variation, mechanical/electrical error, etc) can lead to small unknown perturbations in discharge current ∼1%– 5%, which leads to density perturbations of about the same order of magnitude relative to the nominal density. Thus, if one wishes to translate these designs to a physical device,

such perturbations must be accounted for. This can be done in one of three ways: (a) random noise can be programmed into the learning algorithm itself at greater computational cost as training would take longer, (b) the device can be trained *in-situ* which would automatically take experimental noise within the apparatus into account, or (c) one can optimize the device without taking any noise into account and then do a sensitivity test on the optimized parameters to determine how resistant the resultant designs are to perturbations in plasma density. In this study, we carry out the third option.

The manner in which we conduct the sensitivity test is as follows: Since the elements of the device would be individually controlled in practice, we add mean-zero independently distributed Gaussian noise to the density of each element with a standard deviation $\sigma = p\sqrt{n_e}$ where p represents a 'perturbation factor' that sets the variance of the random noise. If $p = 0.05$, we can expect variation in plasma density amongst the columns on the order of ∼5% of the plasma density in each individual column. This is the same as reassigning the plasma density of each element according to

$$n_{e,perturbed} = \left| n_{e,opt} + \mathcal{N}(0, p^2) n_{e,opt} \right|,$$

where the absolute value is applied so we do not end up with non-physical negative densities when the Gaussian values are spuriously large and negative. For each device, we run ten different random perturbations with the same value for p and increase p in increments of 0.01 until we see a significant failure mode in at least three of the ten runs. For a failure mode to be deemed 'significant', we require a drop in transmission in the correct output port of more than 10%. For the demultiplexer, since a drop in transmission for one frequency can be accompanied with an increase in transmission in the correct output for another frequency, we require a drop in transmission in the correct output of 10% or greater for all frequencies. The dominant failure mode amongst the three or more dysfunctional perturbed devices along with the average drop in transmission in the correct port across all ten simulations is presented below alongside the original design to illustrate how robust each device is and how the device functionality will deteriorate if the plasma density is perturbed to a high enough degree.

Once the optimization algorithm converges on a device, we proceed by conducting a quantitative evaluation to complement the E-field results. To do this, we calculate the S-parameters, insertion loss, and loss to the plasma elements via the time-averaged Poynting flux through specific flux surfaces in the domain. The time-averaged Poynting vector was calculated using the phasor field amplitude results from the FDFD calculation as $\vec{S}_p = \text{Re}\{\vec{E} \times \vec{H}^*\}$. After determining the total energy flux out of the source, the energy flux out of the input waveguide into the domain is calculated to yield the insertion loss, and the net energy flux into the plasma elements is calculated to yield the loss to the plasma elements. The locations of the source and the probes where the objective is calculated are used as the ports for the calculation of the relevant
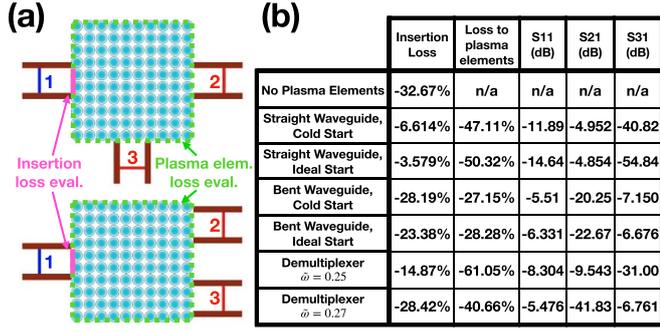
Figure 4. (a) Illustration denoting the loss-evaluation surfaces and ports for the waveguide (top) and demultiplexer (bottom) configurations and (b) insertion loss, loss to the plasma elements, and S-parameters for each of the optimal devices presented in this study.

S-parameters, which we define in the following way for this study:

$$S_{11} = 10\log_{10}\left[\frac{S_{\text{tot}} - \int \vec{S}_{\text{p}} \cdot d\vec{A}_1}{S_{\text{tot}}}\right]$$

$$S_{21} = 10\log_{10}\left[\frac{\int \vec{S}_{\text{p}} \cdot d\vec{A}_2}{S_{\text{tot}}}\right]$$

$$S_{31} = 10\log_{10}\left[\frac{\int \vec{S}_{\text{p}} \cdot d\vec{A}_3}{S_{\text{tot}}}\right],$$

where S_{tot} is the total energy flux out of the source divided by 2 (i.e. the positive-x directed energy flux out of the source). A diagram illustrating the location of the ports and loss evaluation surfaces is given in figure 4(a).

3. Results and discussion

3.1. S-parameters

In figure 4, the quantitative measures of performance for each of the device functions is given, with the 'cold' and 'ideal' start designations referring to whether the initial parameters before optimization were from the idealized study (ideal) or just had all the plasma elements inactive (cold). One important aspect to note before discussing each individual device is that since we did not give any consideration to the coupling between the input and exit waveguides and the device, even when the plasma elements are not present about 1/3 of the input power is reflected at the waveguide exit interface. This insertion loss is improved in every case by the optimal design, however one can imagine that the use of a horn or some other apparatus to reduce the insertion loss would improve the overall performance of the device in all cases. As such, the relatively high $S11$ values in each case are not a major concern.

In all cases, the transmission through the correct port is well over an order of magnitude higher than that of the incorrect port. The worst comparison comes in the 'cold start' bent
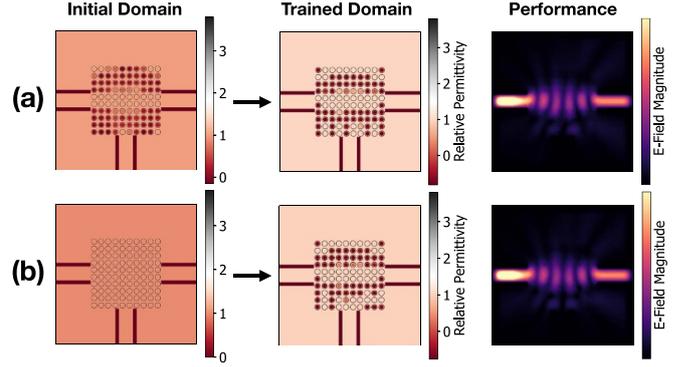

Figure 5. Initial and trained relative permittivity domains and FDFD-simulated field magnitude $|\mathbf{E_z}|$ ($\hat{\mathbf{z}}$ out of the page) showing the resultant performance for the straight waveguide initialized with (a) the optimal parameters from [20] and (b) with all plasma elements deactivated, where the maximum plasma frequency among the plasma elements is about $\tilde{\omega}_{\text{p}} = 0.35$ which translates to $f_{\text{p}} = 7\,\text{GHz}$.

waveguide case where there is a 13 dB difference in transmission between the two ports, and the best comes in the 'ideal start' straight waveguide case where there is a 50 dB difference in transmission between the correct and incorrect ports. In all the waveguide cases, about half of the signal power is lost due to insertion losses and dissipation within the plasma, while in the demultiplexer case these losses are a bit higher at around 70% of the input power. Again, we stress that these losses would likely be partially mitigated by the introduction of a better coupling scheme between the inputs/outputs and the device elements. The values given in figure 4 give us a strong quantitative evaluation of each of the PMM functions, showing that the device consistently directs much more power to the correct output than the incorrect output.

3.2. Directional waveguide

Electromagnetic waveguides are fairly common in inverse design devices either as the primary function or as important building blocks to more complicated devices [5, 18]. Presented in figure 5 are field simulations and initial/final (initial and final here meaning before and after training) domain permittivities of optimal straight waveguide configurations for our PMM. Operating at $\omega = 0.25 \times (2\pi c/a)$, or a non-dimensionalized source frequency $\tilde{\omega} = \omega/(2\pi c/a) = 0.25$, this represents the simplest type of optimization problem for our PMM. Because of the simplicity of this geometry, we limit the plasma frequency to $f_p = 7\,\text{GHz}$, a modest operating condition that could be easily reached and run at long time scales (~hours) without damaging the discharge bulbs. In this and all other simulations, the relative permittivity of the input and output waveguide walls was set to $\varepsilon = -1000$ to serve as a lossless metal. The waveguide objective is the L_2 inner product of the simulated field with an $m = 1$ propagation mode (in order to preserve the input mode) evaluated at the probe in the desired exit waveguide minus the integrated

J. Phys. D: Appl. Phys. **55** (2022) 465203

J A Rodríguez and M A Cappelli

field intensity at the incorrect exit to discourage leakage into the wrong output:

$$\mathcal{L}_{wvg} = \int E \cdot E_{m=1}^* \mathrm{d}l_{\text{desired exit}} - \int |E|^2 \mathrm{d}l_{\text{incorrect exit}}.$$

One could likely improve performance even further by discouraging all leakage; i.e. subtracting the field intensity integrated over all space outside the device other than within the correct output, but for the devices in this study the improvement is marginal.

In the prior idealized study [20], the initial parameters were set to be uniform among all the plasma elements since the initial conditions were found to be unimportant; i.e. the optimal parameters were very similar between runs with different intial parameters. Upon adding all the non-ideal factors described above, the optimization objective function became significantly more complicated with many more local minima than before, and so the initial conditions became closely linked with the device performance once the algorithm converged. In figure 5, the results for both initializations look similar in both optimal parameters and field distributions, but when one considers the S-parameters of the devices given in figure 4 we can see that when initializing the domain with all the plasma elements turned off (figure 5(b), the 'cold' start), the functionality of the device suffers, whereas if the device is initialized with the optimal parameters from [20] (figure 5(a), the 'ideal' start), the final device performs much better even though the parameters change significantly during training. The trained domain for the naively initialized device, while very similar to the device initialized with the idealized results, lacks activation of several bulbs near the incorrect exit, resulting in a signal in that port more than an order of magnitude higher than the 'ideal' start case. The 'ideal' case also exhibits much better coupling into the device from the waveguide, reducing S11.

Next, we modified the objective to produce a waveguide with a 90° bend by simply switching the role of the desired and undesired exit probes in the straight waveguide objective. The results for this case are found in figure 6 where the difference in performance between the two cases is slightly easier to observe in the field plots. Again we see that, while both cases result in decent performance, the device initialized with the optimal parameters from [20] performs better with less loss to the surroundings and lower insertion loss. As this objective is obviously more difficult than the straight waveguide, there is a much smaller difference in transmission between the two waveguides, but it is still very strong. Should the plasma frequency limit for this device be slightly increased, the transmission difference between ports 2 and 3 would likely increase at the expense of long-term performance of the plasma elements or the time scales upon which the device could be active. With the limit as it is at $f_\mathrm{p} = 7$ GHz, this device could be operated at very long time scales (∼hours) for many months as this is essentially the nominal operating condition specified by the bulb manufacturer. Driving the elements slightly harder would reduce the amount of time they could be active before runaway
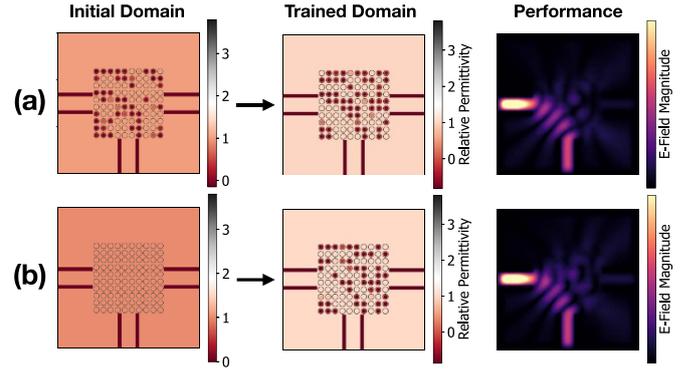

Figure 6. (a) Initial and trained relative permittivity domains and FDFD simulated field magnitude $|\mathbf{E_z}|$ ($\hat{\mathbf{z}}$ out of the page) showing the resultant performance for the bent waveguide initialized with (a) the optimal parameters from [20] and (b) with all plasma elements deactivated, where the maximum plasma frequency among the plasma elements is about $\tilde{\omega}_\mathrm{p} = 0.35$ which translates to $f_\mathrm{p} = 7$ GHz.

thermal effects begin to degrade performance on both the large and small time scales; though in order to reach catastrophic degradation the discharges have to be pushed quite far beyond the 7 GHz limit.

The final configurations shown in figures 5 and 6, much like in the previous study [20], appear to utilize the plasma elements in a much more nuanced way than a human-designed device would. Human-designed devices for a configuration like this would resemble a simple photonic crystal waveguide like that of Wang and Cappelli [24], where all of the discharges within the device are tuned to their maximum plasma density except for those along the path that the user would like the source to propagate, seeking to mimic a metallic waveguide. For any finite device with non-ideal elements, this more naive approach results in much more loss and very little signal at the correct exit, particularly in the bent waveguide case. The plasma density distribution obtained via inverse design, despite its strange structure which, in regions, mirrors the seemingly random structures that often arise from inverse design schemes [12, 15], has some plasma elements that are activated along the path of propagation of the source, even partially covering the entrance and exit which is entirely unintuitive. We suspect that the algorithm is finding an optimal balance between utilizing the reflective behavior of over-dense plasma and the refractive behavior of under-dense plasma to steer the light in a more effective way than simply attempting to mimic a metallic waveguide. It is also interesting to note how the devices create a gradually expanding and then contracting path within the device, similar to two horns directed at one another. This geometry is likely part of what reduces the insertion loss compared to the case when the plasma elements are not present. This demonstrates that inverse design allows us to arrive at novel, highly effective configurations that would not be formulated through conventional means.

In figure 7 we see the results of the sensitivity analysis for the best straight and bent waveguide configurations. In both
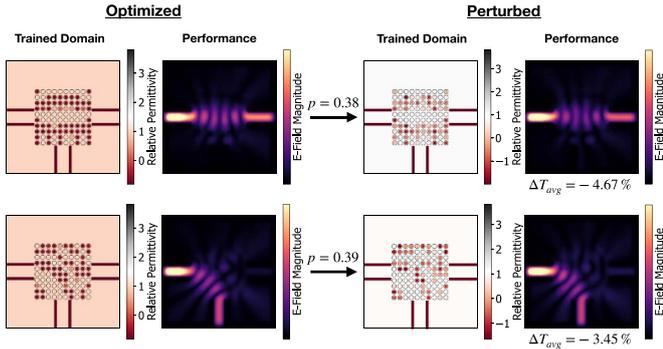
Figure 7. Result for the waveguide stress test including the perturbed domain and field results along with average transmission drop in the correct output port across ten perturbations with the denoted value of p.

cases we see that the optimal design is very robust to random perturbations of the device parameters as p needs to be increased to $\geqslant 0.38$ to see an appreciable loss in functionality. Some of the rods in this case are so greatly altered that at the peak plasma density at the center of the plasma rods the real part of the permittivity drops as low as -1.5 where before the perturbations it was -0.7, and likewise many of the rods are nearly completely inactivated. The field plots show a very minor increase in loss to the surroundings and a drop in transmission at the correct port, but even when three of the ten perturbed devices failed, the average drop in transmission among the perturbed devices was only -4.67% in the straight waveguide case and -3.45% in the bent waveguide case. Keeping experimental error within $\sim 38\%$ is well within current capabilities and so we can assume that this particular non-ideal factor will not hinder experimental realization of these designs.

The results demonstrated in figures 5–7 are reassuring. As mentioned before, simple waveguides such as these can serve as building blocks for more complicated devices, and strong performance like what we see here with very conservative operating conditions opens up many possibilities. It is also important here to remind the reader that unlike most devices produced via inverse design, the PMM geometry allows a single device to perform both functions as the plasma elements can be tuned electronically. As an example, quickly switching between the straight and bent propagation mode could enable the design of transistor-like optical switches that are experimentally feasible. The square lattice PMM configuration is capable of more than just waveguiding, however, as we will explore in the next subsection.

3.3. Demultiplexer

Next, we present a demultiplexer designed to discriminate between frequencies of $\tilde{\omega}_1 = 0.25$ and $\tilde{\omega}_2 = 0.27$. Inverse design methods have been utilized in the past to create frequency demultiplexers [9, 12], and the demultiplexer is a key component in many optical computing schemes [30]. The optimization objective for the demultiplexer is quite different

from the waveguide case since the objective must take into account two separate simulations as the permittivity of the elements is frequency-dependent. With this in mind, the demultiplexer objective is:

$$\mathcal{L}_{mp} = \left(\int E_{\omega_1} \cdot E^*_{m=1} \mathrm{d}l_{\omega_1 \mathrm{exit}} \right) \left(\int E_{\omega_2} \cdot E^*_{m=1} \mathrm{d}l_{\omega_2 \mathrm{exit}} \right)$$
$$- \int |E_{\omega_1}|^2 \mathrm{d}l_{\omega_2 \mathrm{exit}} - \int |E_{\omega_2}|^2 \mathrm{d}l_{\omega_1 \mathrm{exit}},$$

where E_{ω_1} is the simulated field for the $\tilde{\omega} = 0.25$ source and E_{ω_2} is the simulated field for the $\tilde{\omega} = 0.27$ source. Since this objective is significantly more complex than the waveguide case, we allow the plasma frequency of the rods to increase up to $f_\mathrm{p} = 13\,\mathrm{GHz}$, an operating condition which can be maintained for time scales $\leqslant 10\,\mathrm{s}$ at a time without damaging or otherwise altering the operation of the discharge bulbs. If operated in short ($\sim$1–10 ms) pulses with a duty cycle that limits the proportion of time that the bulbs are active, this higher plasma density will not damage the discharges over much larger time scales (e.g. the pulsed operation affects the elements in a similar manner to the nominal operating conditions) and thus represents a very reasonable operating condition that could be sustained over a long device lifetime.

As an additional result of the significant jump in complexity from the waveguide case, it is not sufficient to simply take the optimal parameters from the idealized study [20] and take those as our initial parameters. Instead, we borrow from this idea and train the device parameters in several steps, starting with all the plasmas off and training a domain where the only non-ideal factor is the presence of the quartz envelopes which has an objective function with fewer local minima, allowing the algorithm to converge quickly on a configuration with strong performance. Then, once an optimal design was obtained for this simpler case, we introduce the non-uniform plasma density profile and train using the optimal parameters from the previous run. Finally, we add loss to the rods, train once again and arrive at a final trained domain. The parameter evolution along with the device performance is shown below in figure 8. While this design exhibits much more insertion loss and dissipation within the plasma elements, the difference in transmission between the correct and incorrect output ports is very large and the device performance is very similar to that achieved in [20].

To conclude our analysis of the demultiplexer, we run our stress test procedure on the final design and find that in order to see three failures out of ten perturbed cases, the perturbation parameter p must be increased all the way to 0.31, which is very close to the values found in the waveguide sensitivity test above. The optimized and perturbed domain along with the degraded performance can be found below in figure 9. We also note that since the degraded performance of 10% or greater in three perturbations out of ten is required for both frequencies, the average transmission drop across the ten perturbations is higher. Even with this considered, the high p needed to produce this level of degradation indicates that the plasma
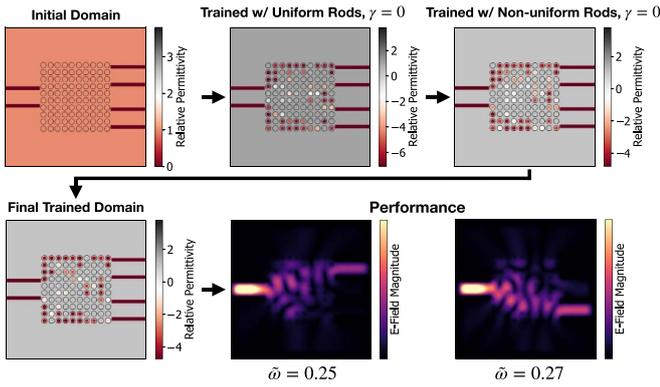
Figure 8. Evolution of the relative permittivity of demultiplexer simulation domain throughout the training stages as the plasma model complexity is increased, along with FDFD simulations showing final device performance at each operating frequency. The maximum plasma frequency among the elements in this case is about $\tilde{\omega}_p = 0.8$ which translates to $f_p = 13$ GHz.
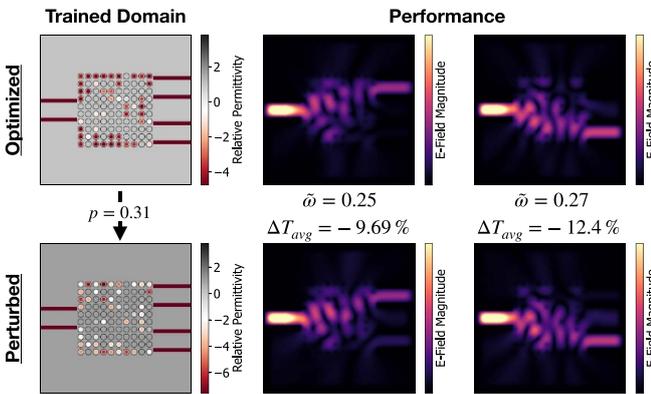

Figure 9. Result for the demultiplexer stress test including the perturbed domain and field results along with average transmission drop in the correct output ports across 10 perturbations, showing a significant degradation of performance when the perturbation parameter p is increased to 0.31.

density distribution that we arrived at in this case is very robust to experimental noise/error.

The results shown here are very compelling. For both source frequencies, the field intensity at the correct output waveguide is substantially higher than at the incorrect output (2–3 orders of magnitude higher as shown in figure 4(b)), and the designs are extremely robust to random perturbation of the parameters. In addition, the insight gained from the optimization procedure for the demultiplexer objective in particular paves the way for both experimental realization of these and more complex devices, such as the logic gates presented in [20]. If one wants to achieve a complicated device function, we have shown here that a good way of doing so is to optimize in stages while adding one non-ideal factor at a time, and then perhaps even perform *in-situ* training of the device using the optimal parameters of the most realistic optimized device as the initial parameters. In the end, these results suggest that the outlook for experimental realization is very bright.

4. Conclusion

This work opens up many opportunities for future work, both in simulation and experiment. On the simulation front, with the improved plasma element model, new device factors could be introduced, such as different geometries (hexagonal lattices, different experimental parameters/bulb designs, larger arrays, etc), variable operating frequencies, externally applied magnetic fields, and time-dependent behavior. For magnetized plasma in particular, the use of circularly-polarized waves gives the user access to a resonance that makes the plasma capable of achieving $\varepsilon > 1$, making much more complex operations possible. Since Ceviche also allows for time-domain simulations, the non-linear response of the plasma elements to a high-power source could also be explored, eventually enabling the design of highly sophisticated devices like all-optical recurrent neural networks. In addition to all of this, the model for our plasma elements could be further improved via the use of kinetic particle-in-cell simulations for each bulb given an operating voltage, current, and fill pressure/composition. The particle density functions obtained from these codes could be used to provide a more accurate density profile than the polynomial profile used in this study, and more detailed collision modeling can be incorporated. With such an accurate model of our elements, one could likely easily take optimal parameters straight from the *in-silico* inverse design algorithm and achieve high levels of performance in the lab.

However, should experimental realization via the direct use of simulation results prove challenging, the nature of the PMM device, specifically the fact that its elements can be tuned continuously in real-time, allows for the training to be carried out fully *in-situ* with plasma element currents adjusted while the gradient of the objective (transmission at a certain point, for example) is determined in real-time. This could be achieved by either carrying out the Ceviche forward-mode differentiation scheme using the experimental device in place of the simulation, or even a more complicated field measurement-based approach that would allow one to calculate the gradients through back-propagation [31]. Doing the device training *in-situ* would automatically take all non-ideal factors into account. The design and construction of such a system will not be straightforward as we require programmatic control of ~100 individual plasma discharge systems that can be smoothly varied in plasma density, but work has already begun on this front and is proceeding nicely.

In conclusion, we apply inverse design methods to create highly optimized and realistic two-dimensional PMM devices including directional waveguides and a demultiplexer for the TM polarization. Despite the consideration of several non-ideal factors such as the presence of quartz envelopes, collisionality/loss, non-uniform density profiles, and experimental error/noise, the devices all perform well and exhibit nuanced, complicated designs that would not be possible without inverse design methodology. A robust procedure for producing effective, experimentally realizable PMMs consisting of a gradual increase in model complexity is presented. The optimal parameters of each of the devices are experimentally feasible with existing equipment, enabling imminent

J. Phys. D: Appl. Phys. **55** (2022) 465203

J A Rodríguez and M A Cappelli

realization of these designs. Further work is needed in both simulation and experiments to enable development of devices with more exotic objectives like boolean logic [20] as well as others with complicated EM wave propagation structures.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

ORCID iDs

Jesse A Rodríguez ⓘ https://orcid.org/0000-0001-7892-6207
Mark A Cappelli ⓘ https://orcid.org/0000-0003-3093-3357

References

[1] Minkov M, Williamson I A, Andreani L C, Gerace D, Lou B, Song A Y, Hughes T W and Fan S 2020 Inverse design of photonic crystals through automatic differentiation *ACS Photon.* **7** 1729

[2] Hughes T W, Williamson I A D, Minkov M and Fan S 2019 Forward-mode differentiation of Maxwell's equations *ACS Photon.* **6** 3010

[3] Su L, Vercruysse D, Skarda J, Sapra N V, Petykiewicz J A and Vučković J 2020 Nanophotonic inverse design with SPINS: software architecture and practical considerations *Appl. Phys. Rev.* **7** 011407

[4] Burger M, Osher S J and Yablonovitch E 2004 Inverse problem techniques for the design of photonic crystals *IEICE Trans. Electron.* **87** 258

[5] Borel P I, Harpøth A, Frandsen L H, Kristensen M, Shi P, Jensen J S and Sigmund O 2004 Topology optimization and fabrication of photonic crystal structures *Opt. Express* **12** 1996

[6] Lin Z, Pick A, Lončar M and Rodriguez A W 2016 Enhanced spontaneous emission at third-order dirac exceptional points in inverse-designed photonic crystals *Phys. Rev. Lett.* **117** 107402

[7] Miller O D 2013 Photonic design: from fundamental solar cell physics to computational inverse design (arXiv:1308.0212)

[8] Christiansen R E, Lin Z, Roques-Carmes C, Salamin Y, Kooi S E, Joannopoulos J D, Soljačić M and Johnson S G 2020 Fullwave maxwell inverse design of axisymmetric, tunable and multi-scale multi-wavelength metalenses *Opt. Express* **28** 33854

[9] Pestourie R, Pérez-Arancibia C, Lin Z, Shin W, Capasso F and Johnson S G 2018 Inverse design of large-area metasurfaces *Opt. Express* **26** 33732

[10] Chung H and Miller O D 2020 Tunable metasurface inverse design for 80% switching efficiencies and 144° angular deflection *ACS Photon.* **7** 2236

[11] Meem M, Banerji S, Pies C, Oberbiermann T, Majumder A, Sensale-Rodriguez B and Menon R 2020 Large-area, high-numerical-aperture multi-level diffractive lens via inverse design *Optica* **7** 252

[12] Piggott A Y, Lu J, Lagoudakis K G, Petykiewicz J, Babinec T M and Vučković J 2015 Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer *Nat. Photon.* **9** 374

[13] Liu D, Gabrielli L H, Lipson M and Johnson S G 2013 Transformation inverse design *Opt. Express* **21** 14223

[14] Hughes T W, Minkov M, Williamson I A and Fan S 2018 Adjoint method and inverse design for nonlinear nanophotonic devices *ACS Photon.* **5** 4781

[15] Molesky S, Lin Z, Piggott A Y, Jin W, Vucković J and Rodriguez A W 2018 Inverse design in nanophotonics *Nat. Photon.* **12** 659

[16] Christiansen R E and Sigmund O 2021 Inverse design in photonics by topology optimization: tutorial *J. Opt. Soc. Am.*B **38** 496

[17] Jiang J, Chen M-K and Fan J A 2020 Deep neural networks for the evaluation and design of photonic devices (arXiv:2007.00084 [eess.IV])

[18] Andrade N M, Hooten S, Fortuna S A, Han K, Yablonovitch E and Wu M C 2019 Inverse design optimization for efficient coupling of an electrically injected optical antenna-led to a single-mode waveguide *Opt. Express* **27** 19802

[19] Iwai A, Righetti F, Wang B, Sakai O and Cappelli M A 2020 A tunable double negative device consisting of a plasma array and a negative-permeability metamaterial *Phys. Plasmas* **27** 023511

[20] Rodríguez J A, Abdalla A I, Wang B, Lou B, Fan S and Cappelli M A 2021 Inverse design of plasma metamaterial devices for optical computing *Phys. Rev. Appl.* **16** 014023

[21] Wang B, Rodríguez J A and Cappelli M A 2019 3D woodpile structure tunable plasma photonic crystal *Plasma Sources Sci. Technol.* **28** 02LT01

[22] Sakai O, Iwai A, Omura Y, Iio S and Naito T 2018 Wave propagation in and around negative-dielectric-constant discharge plasma *Phys. Plasmas* **25** 031901

[23] Sakai O, Maeda J, Shimomura T and Urabe K 2013 Functional composites of plasmas and metamaterials: flexible waveguides and variable attenuators with controllable phase shift *Phys. Plasmas* **20** 073506

[24] Wang B and Cappelli M A 2016 Waveguiding and bending modes in a plasma photonic crystal bandgap device *AIP Adv.* **6** 065015

[25] Wang B and Cappelli M A 2016 A plasma photonic crystal bandgap device *Appl. Phys. Lett.* **108** 161101

[26] Wang B, Rodríguez J A, Miller O and Cappelli M A 2021 Reconfigurable plasma-dielectric hybrid photonic crystal as a platform for electromagnetic wave manipulation and computing *Phys. Plasmas* **28** 043502

[27] Righetti F, Wang B and Cappelli M 2018 Enhanced attenuation due to lattice resonances in a two-dimensional plasma photonic crystal *Phys. Plasmas* **25** 124502

[28] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)

[29] Thomas W 2019 Development and validation of computational tools for the characterization and analysis of non-linear plasma photonic crystals *Master's Thesis* University of Washington Research Works Archive (available at: https://digital.lib.washington.edu/researchworks/handle/1773/43281)

[30] Xu X *et al* 2021 11 tops photonic convolutional accelerator for optical neural networks *Nature* **589** 44

[31] Hughes T W, Minkov M, Shi Y and Fan S 2018 Training of photonic neural networks through *in situ* backpropagation and gradient measurement *Optica* **5** 864