Multifunctional Volumetric Metaoptics

Thesis by
Conner Ballew

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended August 3, 2021

ACKNOWLEDGEMENTS

In 2016 I was in a Master's program at the University of Washington and had recently been accepted to a PhD program at Caltech. The UW Electrical Engineering department had a small news snippet on their website about what I planned to do after graduating, and I mentioned I would be going to Caltech for a PhD program. During the conversation I mentioned that I had an interest in both photonics and space technology, and that Caltech was a fortunate choice for me because it is both a great school *and* manages the Jet Propulsion Laboratory. When I actually moved to Caltech I had no concrete plans for working on space technology or anything like that — that would just be a bonus objective while my main goal was photonics. Shortly after I joined Andrei's lab I was given a collaborative project with a group at the JPL, which evolved into more projects and ultimately into a postdoctoral fellowship that I will be pursuing after my PhD. The bonus objective to my PhD plans simply fell into my lap! While this was extremely serendipitous, it simply would not have been possible without the help of many along the way.

Andrei Faraon deserves a lot of credit for the success of his students. He has a breadth of experience that has allowed him to pursue many different projects, which allows his students to pick exciting projects and freely focus on the scientific details. Andrei is excellent at enabling his students to be successful, and supporting them along the way. The lab group is full of happy and supportive individuals who I will dearly miss working closely with. Yu Horie, Ehsan Arbabi, Mahsa Kamali, Jon Kindem, Chuting Wang, Ioana Craiciu, Sadegh Faraji-Dana, Tian Zhong, John Bartholomew, Sarah Camayd-Munoz, Daniel Riedel, Hyounghan Kwon, Jake Rochman, Greg Roberts, Mi Lei, Andrei Ruskuc, Tianzhe Zheng, Tian Xie, Chun Ju Wu, Suki Gu, Ian Foo, Riku Fukumori, and Bilgehan Baspinar — thank you all for being wonderful labmates over the years! I especially want to thank Yu Horie, who mentored me in the critical first few months of my time in the lab before he graduated; Ehsan Arbabi and Mahsa Kamali, who are the ones who I consider foremost experts in metasurfaces and absolutely unparalleled in their ability to teach these topics to new students; and finally Greg Roberts, Tianzhe Zheng, and Sarah Camayd-Munoz, who are not only great friends but the ones who lived and breathed the same research as mine. Amir Arbabi and Oscar Bruno have consistently provided excellent feedback on project ideas and manuscripts. Thank you all!

Our focus and well-being as graduate students would not be attainable if not for the many staff members that keep the machine running. Cecilia Gamboa, Christy Jenstad, and Tanya Owen are but a few of the administrative experts that are the life blood of the Electrical Engineering and Applied Physics departments at Caltech. Without their diligence we would all just be spinning our wheels!

One of the greatest honors I have ever had was working with members of the Submillimeter-Wave Advanced Technology group at the JPL. I feel privileged to have worked with them because this collaboration amplified my research, gave it scientific purpose, and opened the door for a huge amount of future work. Maria Alonso-delPino, Goutam Chattopadhyay, Cecile Jung-Kubiak, and Sofia Rahiminejad — staggering intellect and ceaseless support. I owe you all a great deal, and I am excited to continue working with you as I transition to the JPL.

And lastly, but certainly not least, I wish to thank the family and friends that supported (and healthfully distracted me) along the way. Foremost among them is my partner Sullivan, my parents David and Diana, and my brothers Sean and Jason.

Conner Ballew
California Institute of Technology
August 3, 2021

ABSTRACT

The development of increasingly complex systems is often realized by adding increasingly more components to a system. Periodically, a new class of optical devices is discovered or engineered that reduces the required number of components in an optical system, which can lead to miniaturized systems with broader functionality. The ability to pattern materials at a scale comparable to the wavelength of the radiation being controlled has enabled unprecedented control of signals across the electromagnetic spectrum. At this scale, the properties of the device can be precisely engineered to behave in ways that are not observed in natural materials, thus lending these devices the name *optical metamaterials*.

A subset of optical metamaterials is the *metasurface*, which typically features an array of sub-wavelength elements spread over a single wavelength-thick sheet. Benefiting from available planar fabrication techniques and a deterministic design methodology, metasurfaces have started to be investigated in the early 1990s and now are at the point where they can enter into mainstream technology. Yet their thinness and the assumptions employed when designing them ultimately limit the bandwidth and complexity of the functions that metasurfaces can encode.

This thesis studies a new class of optical metamaterials that are patterned in all 3-dimensions, referred to as *volumetric metaoptics*. Both the design, which no longer benefits from simplifying assumptions employed in 2D metasurface design, and the fabrication are major challenges in the realization of these devices that the presented work addresses. Chapter 2 motivates and describes advanced inverse-design techniques that are suitable for finding the 3D refractive index distributions that optimally perform complex functionalities. The devices can be designed to respond in a desired way depending on all fundamental properties of light (intensity, frequency, polarization, and k-vector), and functionalities can be combined to enable extreme multifunctionality. Furthermore, the design methodology includes the ability to impose fabrication constraints, thereby finding solutions that can be fabricated with contemporary technology.

Chapter 3 showcases several examples of highly multifunctional metaoptics, including a scattering-based Bayer filter that outperforms traditional absorptive Bayer filters used in digital color imaging and a set of metaoptics that can be mechanically reconfigured to change optical functionality. Prototypes of these devices are experimentally verified at microwave frequencies. A passive device capable of classifying light according to its wavelength, polarization, and incident angle in an extremely small footprint is also designed with the potential to be fabricated with current CMOS processing fabrication tolerances. These devices allow massive amounts of information to be extracted with computational post-processing and can be tiled across a camera array to realize advanced computational imaging arrays.

Chapter 4 and Chapter 5 discuss 3D metaoptics at terahertz frequencies. The devices are made of stacked silicon wafers, each patterned with silicon micromachining and assembly technology capable of generating deeply subwavelength feature sizes and layer-to-layer alignments. The devices are well-suited for realizing dense focal plane arrays with large numbers of elements that are required for imaging planetary and astrophysical phenomena.

PUBLISHED CONTENT AND CONTRIBUTIONS

[1] C. Ballew, G. Roberts, S. Camayd-Muñoz, M. F. Debbas, and A. Faraon, "Mechanically reconfigurable multi-functional meta-optics studied at microwave frequencies", Scientific Reports **11**, 11145 (2021) `DOI`: `10.1038/s41598-021-88785-5`,
C.B. participated in the conception of the project, designed the devices, participated in constructing the measurement setup, and conducted the experiments.

[2] C. Ballew, G. Roberts, and A. Faraon, "Volumetric metaoptics for multifunctional applications", Manuscript in preparation (2021),
C.B. participated in the conception of the project, optimized the proposed device, and analyzed its performance.

[3] C. Ballew, G. Roberts, T. Zheng, and A. Faraon, "Constraining continuous topology optimizations to discrete solutions for photonic applications", (2021), arXiv:`2107.09468 [physics.app-ph]`,
C.B. participated in the conception of the project, implemented the proposed algorithm, and conducted the example simulations.

[4] P. Camayd-Muñoz, C. Ballew, G. Roberts, and A. Faraon, "Multifunctional volumetric meta-optics for color and polarization image sensors", Optica **7**, 280 (2020) `DOI`: `10.1364/optica.384228`,
C.B. conducted the measurements and analyzed the measurement results.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

ABBREVIATIONS

2D	two-dimensional.
3D	three-dimensional.
CMOS	complementary metal oxide transistor.
Si	silicon.
EM	electromagnetic.
FOM	figure of merit.
ASM	Angular Spectrum Method.
FFT	fast Fourier transform.
IFFT	inverse fast Fourier transform.
DFT	discrete Fourier transform.
RCWA	rigorous coupled-wave analysis.
FFF	fused filament fabrication.
RF	radio frequency.
THz	terahertz.
IR	infrared.
FDTD	finite-difference time-domain.
MPI	Message Passing Interface.
RGB	red-green-blue.
PLA	polylactide acid.
JPL	Jet Propulsion Laboratory.
AR	aspect ratio.
DRIE	deep reactive-ion etching.
HR-Si	high-resistivity silicon.
VNA	vector network analyzer.
CAD	computer-aided design.
SOI	silicon on insulator.
BOX	buried oxide.
ICP-RIE	inductively coupled plasma reactive ion etching.
SEM	scanning electron microscope.
TOX	thermally oxidized silicon.
HF	hydrofluoric acid.
LO	local oscillator.
***Q*-factor**	quality factor.
FSR	free spectral range.

DBR	distributed Bragg reflector.
FDFD	finite-difference frequency-domain.

Chapter 1

INTRODUCTION: MULTIFUNCTIONALITY IN OPTICAL COMPONENTS

The central idea of this thesis is *multifunctionality*, a term attributed to an optical device that responds in fundamentally different ways depending on the properties of the incident light. Basic forms of multifunctionality exist in single components: natural birefringent materials can treat polarizations differently, gratings separate wavelength, lenses separate light of different angles of incidence; but in general, a complex combination of these basic functions requires a combination of those basic optical components. This modular design paradigm follows a rough rule: *increasingly complex functionality requires increasingly more components.*

Optical engineers and graduate students have likely witnessed this in the lab. An optical setup starts with a laser whose beam size is not quite right, so we add one or two lenses to expand the beam and collimate it. Its alignment needs to be controlled precisely, so we add several mirrors to control that. Now, if we are for instance making a fluorescence microscope, we shine the laser light on a sample and observe the light that the specimen glows on a camera sensor, and quickly realize that the reflected laser light is saturating our camera, blinding us to the dimmer fluorescent glow of the sample. To fix this we add a dichroic mirror or spectral filter to separate the high-frequency laser light from the low-frequency fluorescent light. We can continue on, adding irises to filter out-of-focus light, adding combinations of polarizers and waveplates to ascertain polarization-dependent properties, and eventually the optical table will be full of components.

Rather than obtaining complex functionality through large numbers of basic components, we can instead obtain it through a very small number of complex components. These complex components can be multifunctional in nature: a lens for focusing light, a spectral splitter that routes different colors to different points, a polarizer that routes different polarizations to different points, an iris for filtering different angular components out — *all performed simultaneously using just one device*. Such complex components are, unfortunately, not naturally occurring (at least not to our current knowledge) and must be engineered instead, thus lending them the name *optical metamaterials*. This term alludes to the broader field of metamaterials which includes acoustic metamaterials, mechanical metamaterials, and generally components whose physical properties cannot be found in naturally occurring materials.

Metasurfaces are two-dimensional (2D) optical metamaterials. These components are typically operated as free-space components, like lenses and gratings, that can be placed in the path of some propagating beam. They are distinguished by their small feature sizes, which typically must be comparable to the wavelength of light that they are designed to control. In fact, asserting complex control of light *efficiently* in a small form factor device requires these small feature sizes. The metasurface features tend to vary spatially, thus imposing some spatially dependent function on the light input.

Optical metasurfaces (operating at near-infrared or visible wavelengths) were demonstrated in the 1990s. The spatial profile ingrained into the metasurface layer by varying the shape of low-loss dielectric elements provides a high-transmission way of manipulating the phase of light as it propagates through the metasurface, and in this way lenses (quadratic phase profiles) and beam deflectors (linear phase profiles) could be made in wavelength-thin layers. Initially most metasurfaces were made of long bars, bearing a resemblance to 1-dimensional diffraction gratings but with spatially varying periods or duty cycles [1, 2]. The importance of sub-wavelength feature size elements here becomes apparent when considering diffraction efficiency — the periodicity of the grating should be sufficiently small to suppress high-order diffraction components. 2D metasurfaces were demonstrated a few years later, replacing the long bars with small posts [3–5]

Major developments that enabled high efficiencies complex multifunctionality to be synthesized in single metasurface began by the early 2010s [6–8]. With improved fabrication techniques, the ability to reliably control the specific shape of the metasurface elements permitted more specific control of optical properties with improved efficiencies. For example, a metasurface consisting of circular elements exhibited polarization independent functionality, whereas a metasurface with elliptical elements could affect two orthogonal polarizations differently. Metasurface capabilities reached a critical milestone in 2015 with the demonstration of *complete phase and polarization control* [9]. Here it was demonstrated that metasurfaces could be designed to encode two arbitrary phase profiles into a single layer for two orthogonal polarizations. In other words, when illuminated with one polarization an arbitrary phase hologram can be projected, and when illuminated with the other polarization a completely different hologram can be projected. A further capability demonstrated here was the ability to encode different polarization rotations at different points along the metasurface, essentially creating a geometrically varying polarization waveplate. However, the arbitrary control of phase and the arbitrary control of polarization cannot be simultaneously controlled with a single metasurface (at least not without sacrificing efficiency).

The inability to simultaneously control phase and polarization efficiency is partially rooted in the design process of metasurface, but more fundamentally it is rooted in the thin nature of the metasurface. A single, thin layer metasurface simply does not have enough degrees of freedom to encode increasingly arbitrary functionalities [10]. The metasurface is also a diffractive device, owing to its quasi-periodic nature. The physics of diffraction is fundamentally wavelength-dependent, which introduces chromatic aberration and thus poses a problem for broadband behavior. Some of these problems can be overcome by cascading multiple metasurfaces [11, 12] or by applying post-processing algorithms to correct for certain aberrations [13].

These limitations on increasing multifunctionality, including operating across multiple wavelengths, is the primary problem that the work in this thesis seeks to solve. The subject of further increasing multifunctionality of metasurfaces has been the subject of intense research over the past several years. Cascading metasurfaces to improve performance falls back on the original, modular approach of enhancing functionality by adding more basic components rather than increasing the complexity of just a few components. Computational techniques can extract more functionality out of a single metasurface under certain assumption but cannot overcome the fundamental limitations inherent in the physics of thin metasurface [14, 15].

This thesis studies a different approach: add a third dimension to the metasurface. The term "metasurface" no longer applies when the device is three-dimensional (3D), so the term *volumetric metaoptics* or simply *metaoptics* is used instead. An extreme challenge in the realization of 3D metaoptics is the design procedure. Chapter 2 details one of the standard procedures for designing metasurface, with a particular emphasis on the assumptions that both render metasurface design tractable but also preclude its applicability to 3D devices. Tossing aside these assumptions, we are left with unsimplified electromagnetic problems that require efficient algorithms to solve. For this we turn to *inverse-design algorithms* that can take a large collection of arbitrary electromagnetic figure-of-merit functions, and find the optimal 3D material

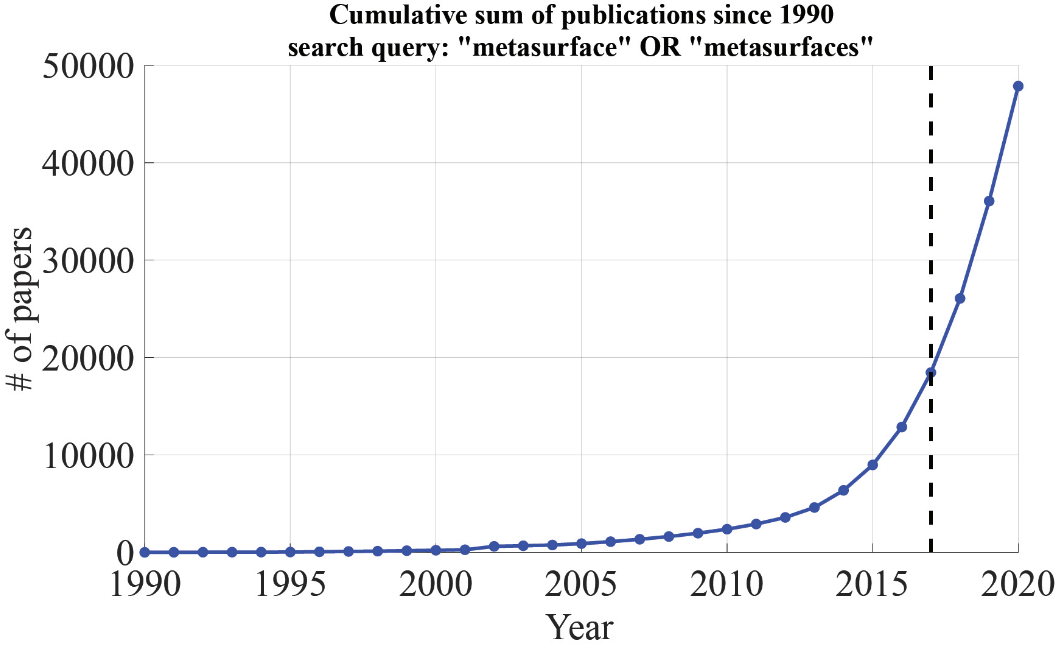

Figure 1.1: **Metasurface publications since 1990.** Data obtained by finding the number of results returned by Google Scholar with the search query *"metasurface" OR "metasurfaces"*. The black dashed line represents the year (2017) that the work in this thesis began.

distribution of a device to optimally perform those function. Furthermore, fabrication constraints can be directly incorporated into the inverse-design procedure, thereby ensuring that the design optimally converges to a fabricable structure.

This approach offers an unprecedented degree of multifunctionality which is showcased in Chapter 3. Inverse-designed spectral splitters that both outperform current color imaging filter and add an element of polarization control are designed and experimentally verified at microwave frequencies. This multifunctionality is then substantially expanded upon by incorporating mechanical reconfigurability into the structures, and again experimentally verifying at microwave frequencies. Finally, a 20-layer 3D device is designed with the capability of sorting light based on the fundamental properties of propagation direction, polarization, and frequency. The device is designed with fabrication tolerances corresponding to those used in complementary metal oxide transistor (CMOS) fabrication processes, which are not currently available for material systems consisting of low-loss optical materials but may be available in the near future as demand for miniaturized photonic components rises. This device has the potential to vastly improve upon current cameras and is particularly well-suited for enabling computational imaging algorithms to extract substantial amounts of information from the system.

The challenge of realizing 3D metaoptics is further complicated by the limitation of fabrication techniques. Some of this burden is eased by incorporating fabrication constraints into the inverse-design procedure as described, but further improvements can be made by employing state-of-the-art fabrication techniques. Available fabrication techniques depend on the frequency of operation of the device. Chapter 4 introduces both 2D metasurfaces and inverse-designed 3D metaoptics operating at THz frequencies. The devices are fabricated by patterning silicon (Si) with deeply subwavelength features, then stacking multiple Si wafers with deeply subwavelength layer-to-layer alignment. These devices have enormous potential for reducing the size, mass, and required power to operate relative to current state-of-the-art techniques. Chapter 5 takes a deeper dive into the theoretical prospects of these THz devices, introducing the concept of a high-resolution, high-bandwidth THz spectrometer capable of observing spectral lines critical to the understanding planetary studies and astrophysics.

Chapter 2

GUIDE TO SIMULATION AND DESIGN OF METAOPTICS

This chapter details the design process of metaoptics, and is intended to be a standalone guide that will aid in understanding subsequent chapters of this thesis. The guide will begin with a summary of textbook-level analyses of planar electromagnetic (EM) fields, which will prove essential for both design and analysis of metaoptics. The guide will then cover metasurface design using a "look-up table" approach to sample phase profiles with subwavelength-sized dielectric elements. This technique is commonly employed in metasurface design presently, and relies on the behavior of the individual elements of the metasurface being independent from its neighbors.

The next section describes the design of 3D metaoptics, which generally cannot be designed in the same way as metasurfaces due to their more complex electromagnetic response. Specifically, many scattering events typically take place in a large 3D volume, so the assumptions of locality, weak coupling of nearby elements, and angle independence that are assumed in metasurface design do not hold. The techniques used for 3D metaoptics here employ inverse-design, which involves making small perturbations to a design structure that incrementally increase its performance. The small perturbations are made in accordance with the computed gradient of one or more figures of merit (FoMs), which can be computed efficiently using a technique known as the *adjoint method*. Since the gradient is computed in a full-wave Maxwell solver, all of the electromagnetic scattering complexities are accounted for. Included in this section is a novel technique that is used to generally constrain continuous-permittivity optimizations to discrete solutions, which proved to be an essential technique for optimizing high-index contrast structures such as the Si-air metaoptics that are described in Ch. 4.

2.1 Planar analysis of electromagnetic fields

Many techniques exist for analyzing electromagnetic fields. The techniques described here are simply a subset of these techniques that are well-suited for the work done in this thesis. Most of these techniques will be simply discussed rather than proven rigorously, however I will specifically mention references that treat these topics rigorously.

Consider the scenario in which a beam of electromagnetic energy is propagating in 3D space, in the absence of any beams propagating in the reverse direction. This beam typically evolves as it moves along its axis of propagation, and thus it may seem like we must know everything about its 3D shape in order to fully understand the beam. However, it is sufficient to have just one 2D cross-section of the beam in which we know the amplitude and phase. This cross-section of the beam can be finite provided we are sure that the vast majority of the energy of the beam in the specific 2D plane is encapsulated within the finite bounds. If these conditions are satisfied, we can propagate the beam to any other plane in free-space, thus reconstructing the full 3D distribution of the beam at will.

Angular spectrum method

The primary method used in this thesis for propagating a 2D beam profile to another plane is known as the *Angular Spectrum Method (ASM)*. For a detailed discussion of this technique, see Chapter 3.10 of Introduction to Fourier Optics (Fourth Edition) by Joseph W. Goodman [16]. This method hinges on what is known as the *angular spectrum of plane waves*, which describes that the Fourier transform of a scalar electromagnetic field represents the decomposition of the field into a collection of plane waves propagating at different angles. Once the fields are decomposed into a collection of plane waves, the individual plane waves can be propagated analytically by applying an appropriate phase-shift, and in the case of evanescent fields an attenuation factor. The propagated plane waves can then be inverse Fourier transformed, yielding the propagated field in the original spatial coordinate system. A field propagated by a distance z_0 is evaluated using the ASM by the following:

$$E(x, y, z + z_0) = \mathcal{F}^{-1}\{\mathcal{F}\{E(x, y, z)e^{ik_z z_0}\}\} \tag{2.1}$$

$$k_z = \sqrt{k^2 - k_x^2 - k_y^2}$$

The quantities k_x, k_y, and k_z are the wave vectors in x, y, and z directions, and $k = \frac{2\pi}{\lambda}$ is the wavenumber. Note that k_z can be imaginary. This is the case in which fields are attenuated through free-space propagation. Any plane wave components existing outside of the circle defined by $k_x^2 + k_y^2 = |\boldsymbol{k}|^2$ will have an attenuation factor. This circle defines the *diffraction limit* and limits the resolution of any optical imaging system relying on free-space propagation. The fact that it is accounted for in the ASM ensures that we can use this method to analyze the diffraction effects of apertures without any special treatment of evanescent components. An example of applying the ASM to a Gaussian beam is shown in Fig. 2.1a-d.

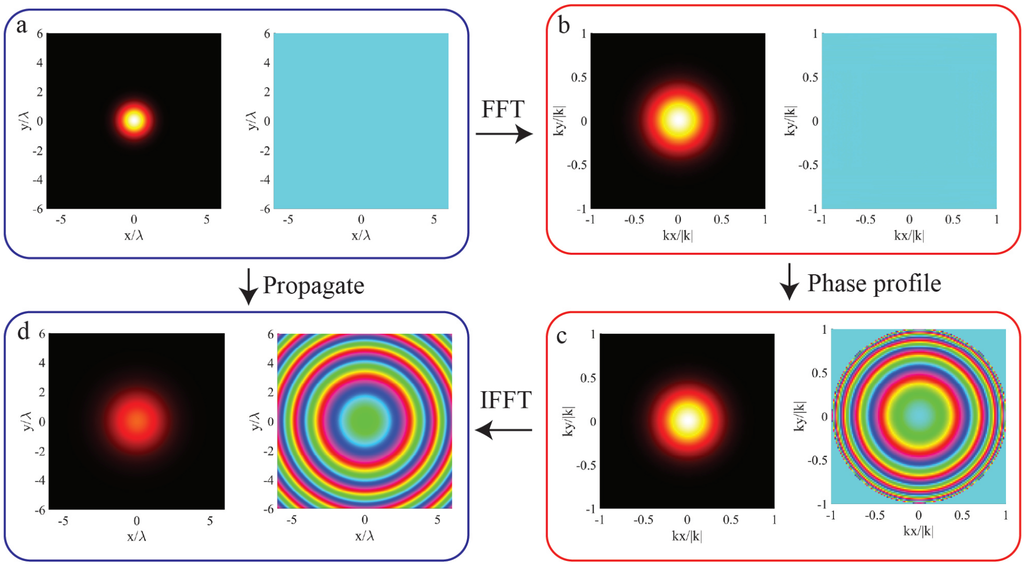

Figure 2.1: **Example: Angular spectrum method applied to a Gaussian beam.** A Gaussian beam of waist radius λ is propagated 5λ from its waist through vacuum ($n = 1$) using the ASM. For each box the amplitude (left) and phase (right) is shown. (a) The initial beam has a Gaussian amplitude profile and a constant phase profile, indicating it is located at the beam waist. (b) A fast Fourier transform (FFT) is applied to the complex beam profile. (c) The phase shift term $e^{ik_z z_0}$ is applied, as described by Eq. 2.1, leaving the amplitude unchanged. Interpreting the FFT as decomposing the beam into its individual planewave components, this phase term applies different propagation distances to the different planewave terms without altering their relative amplitudes. (d) The beam is converted back to spatial coordinates by applying an inverse fast Fourier transform (IFFT). The amplitude profile has expanded and the phase profile is no longer flat due to divergence of the beam.

The ASM does not give a sign to the k_z component of the fields. On one hand, this offers the advantage of being able to choose a sign, and thus choose whether to propagate the field forwards or backwards. (As an example, consider a Gaussian beam cross-section somewhere away from the beam waist. Then choosing one sign for k_z will cause the beam to diverge, and the other sign will cause the beam to converge.) On the other hand, this yields the disadvantage that the beam is assumed to be propagating in one direction only. In reality, it is possible to sample a field profile that contains counter-propagating waves. The work done in this thesis rarely encounters this situation.

Most of the work conducted in this thesis involves simulated and measured fields, so they are discretized and evaluated within a finite domain. Thus, fields are typically defined in matrix form and treated with numerical methods, rather than treated as functions and treated with analytical methods. Related to this is the topic of *sampling*, both in simulation and measurement. Unless otherwise mentioned, references to the Fourier transform refer to the FFT implementation of the discrete Fourier transform (DFT). The DFT requires sufficient sampling to avoid signal aliasing according to the Nyquist-Shannon sampling theorem. The required sampling spatial sampling rate of the signal is two times larger than the highest Fourier component of the signal. For collimated sources the max spatial component may be small. In general, the highest spatial frequency component of a propagating electromagnetic wave is $2\pi/\lambda$, corresponding to a spatial period of λ. Thus, sampling at $\lambda/2$ is sufficient. The finite window over which the DFT is evaluated is also important. As mentioned before, all of the electromagnetic energy should be encapsulated in this window. The fields may increase in spatial extent when being propagated with the ASM, such as is the case with a diverging beam. In this case, the window size must be sufficiently large to encapsulate all (in practice, *most of* but not all) the energy of the diverged beam, or else the Fourier-based methods will yield inaccurate results even if sampled above the Nyquist rate.

Mode overlap integral

It is often desired to examine the similarity of two electromagnetic fields quantitatively. For this, the *mode overlap integral* can be used.

$$\eta(E_1, E_2) = \frac{\left|\int \overline{E_1} E_2 dA\right|}{\int |E_1|^2 dA \int |E_2|^2 dA} \in [0, 1] \tag{2.2}$$

The function η evaluates how similar the fields E_1 and E_2 are, accounting for both amplitude and phase. In the absence of reflection, this quantity represents the fraction of power coupled from an input E_1 to an output E_2. As an example, suppose we are trying to couple power into a single-mode optical fiber. If we manage to create an electromagnetic field profile that exactly matches the conjugate of the fiber mode profile, then $\eta = 1$. However, there will likely be some reflection at the material interface at the input of the optical fiber, and therefore less than 100% of the power will transmitted.

The quantity is preserved under free-space propagation. This can be proven using Parseval's theorem (2.3).

$$\int_{-\infty}^{\infty} f(x)\overline{g(x)}dx = \int_{-\infty}^{\infty} \tilde{f}(x)\overline{\tilde{g}(x)} \tag{2.3}$$

Eq. 2.2 can be rewritten as

$$\eta(E_1, E_2) = \frac{\int \overline{E_1} E_2 dA \int E_1 \overline{E_2} dA}{\int E_1 \overline{E_1} dA \int E_2 \overline{E_2} dA} \tag{2.4}$$

Each integral in 2.4 is of the same form, so it is sufficient to show that each individual integral is preserved under free-space propagation.

$$\begin{aligned}
\iint \overline{E_1(x,y,z)} E_2(x,y,z)\,dx\,dy &= \iint \overline{\tilde{E}_1(x,y,z)}\tilde{E}_2(x,y,z)\,dk_x\,dk_y \\
&= \iint \overline{\tilde{E}_1(x,y,z)} e^{-ik_z z_0} \tilde{E}_2(x,y,z) e^{ik_z z_0}\,dk_x\,dk_y \\
&= \iint \overline{\tilde{E}_1(x,y,z) e^{ik_z z_0}} \tilde{E}_2(x,y,z) e^{ik_z z_0}\,dk_x\,dk_y \\
&= \iint \overline{\tilde{E}_1(x,y,z+z_0)} \tilde{E}_2(x,y,z+z_0)\,dk_x\,dk_y \\
&= \iint \overline{E_1(x,y,z+z_0)} E_2(x,y,z+z_0)\,dx\,dy
\end{aligned}$$

The mode overlap equation is particularly well-suited as a figure of merit since we can propagate fields to a convenient plane and quantitatively compare the fields to the desired fields.

2.2 Metasurface design

The techniques described in Section 2.1 can analyze and transform planar electromagnetic fields, and can be used to help design planar optical components. These techniques are useful tools in the design of dielectric metasurfaces, which are typically planar components made up of small elements that can alter the phase of incident light. This section is intended to provide a general overview of metasurface design, while 3 will delve into more detail on the various constraints, fabrication, and testing procedures for metasurfaces designed to operate at terahertz frequencies. Understanding the design process of metasurfaces helps in understanding the limitations of metasurfaces, both in their design and functionality, which consequentially motivates the use of 3D metaoptics described in this thesis.

The discussion here is limited to dielectric metasurfaces, which tend to be lossier than metallic-based plasmonic metasurfaces at non-RF and non-microwave frequencies. For a review of plasmonic metasurfaces, refer to [17–19]. The metasurfaces described here feature quasi-periodic arrays of sub-wavelength dielectric elements. The reason the array is considered quasi-periodic instead of periodic is because the elements are positioned on a periodic lattice, but are allowed to vary spatially. However, for purposes of design that will become clear shortly, it is beneficial to impose that the elements be at least quasi-periodic, meaning that the dissimilar elements populate a periodic lattice.

The goals of metasurface design done during the course of this thesis work are:

1. Identify a collection of elements that impose phase shifts with high transmission (or reflection, if desired)
2. Combine these elements to create a spatially varying phase mask that performs a desired functionality

The choice of individual elements that comprise a metasurface, sometimes called *meta-atoms*, imbues the metasurface with properties that are difficult to mimic with bulk, refractive components or basic diffractive components. For example, subwavelength elements without continuous or discrete cylindrical symmetry, such as non-circular ellipses and non-square rectangles, demonstrate a property called *local birefringence*. Light is coupled into the element, where the effective refractive index of cross-polarized modes is different. The response of these individual elements are therefore polarization-dependent, and can be varied across the metasurface to spatially control the phase and polarization of impinging light.

Depending on the period and geometry of the meta-atoms, neighboring elements may not interact strongly with one another. This simplification greatly simplifies the design procedure from metasurfaces because a desired transfer function can be "sampled" with a look-up table of meta-atoms, oftentimes without a severe reduction in performance due to coupling between neighboring elements. Backed by this assumption, a popular design methodology is to simulate the response of an element in a *periodic* lattice, where all elements are completely identical, and use its simulated response here to predict the response within a quasi-periodic lattice. This concept is illustrated in Fig. 2.2. The periodic simulation can be run extremely quickly because of the small (subwavelength) simulation region. A large batch of these simulations is run for different geometrical parameters, and from these simulations a set of optimal elements is extracted. Typically, an optimal element is chosen to be as close to the desired phase for each polarization as possible with as high of a power transmission as possible. For dual-polarization metasurfaces that control orthogonal polarizations independently, the goal is to find a set of metasurface elements that can yield 0 to 2π phase response for one polarization *at every 0 to* 2π *phase response for the other polarization*. With high index contrast elements, such as Si and air, this can be achieved with metasurface elements that are on the order of only a wavelength thick.

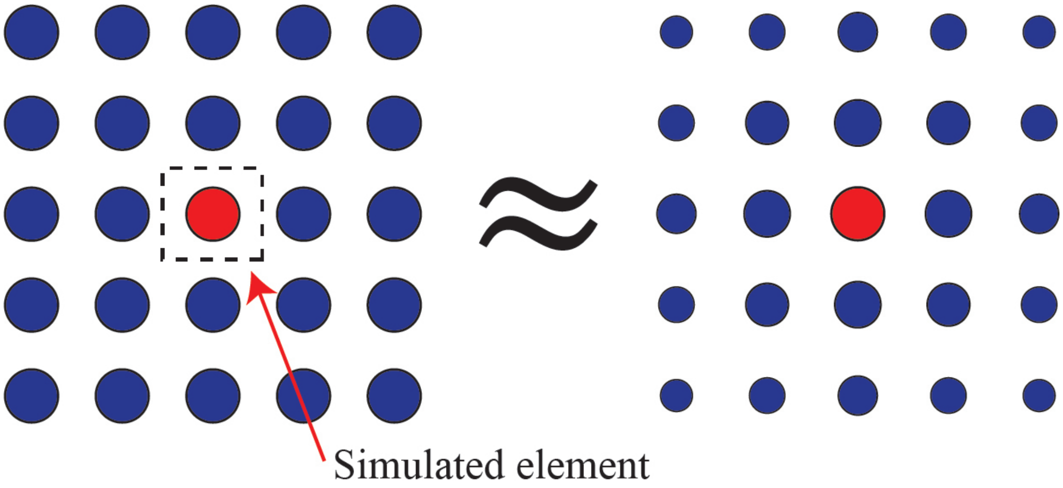

Figure 2.2: **Metasurface design assumption.** A typical assumption made in metasurface design is that the response of a single element (the red circle) is the same in a periodic lattice (left), which is computationally easy to simulate, as it is in a quasi-periodic lattice (right), which is more difficult to simulate.

Fig. 2.3 shows several examples of metasurfaces. A metasurface lookup table was simulated using rigorous coupled-wave analysis (RCWA). The individual elements are chosen to be elliptical air holes in a Si layer with a Si background. The reasons for this material system are described in detail in Chapter 4, but the details associated with this choice are not important for simply demonstrating how a metasurface is designed. The same design methodology would be applied to Si posts surrounded by air with a low index background material. The lookup table is then used to sample several phase profiles. The first is a polarization independent deflector that applies the same linear phase profile to both incident polarizations. The metasurface elements in this case are all circular. The second example is a polarizing beam splitter, which applies an equal-and-opposite linear phase profile to orthogonal linear polarizations. This would have the effect of routing x-polarized light to a −5°

angle, and routing y-polarized light to a +5° angle. The third example is a focusing phase profile, similar to a lens. The phase profile is chosen to be the reciprocal phase of a Gaussian beam (waist radius λ) that has been propagated 20λ in air using the ASM method described in the previous section. The final example combines the polarizing beam splitter profile with the focusing phase profile, yielding a device that both deflects the polarizations to equal-and-opposite angles and focuses them.

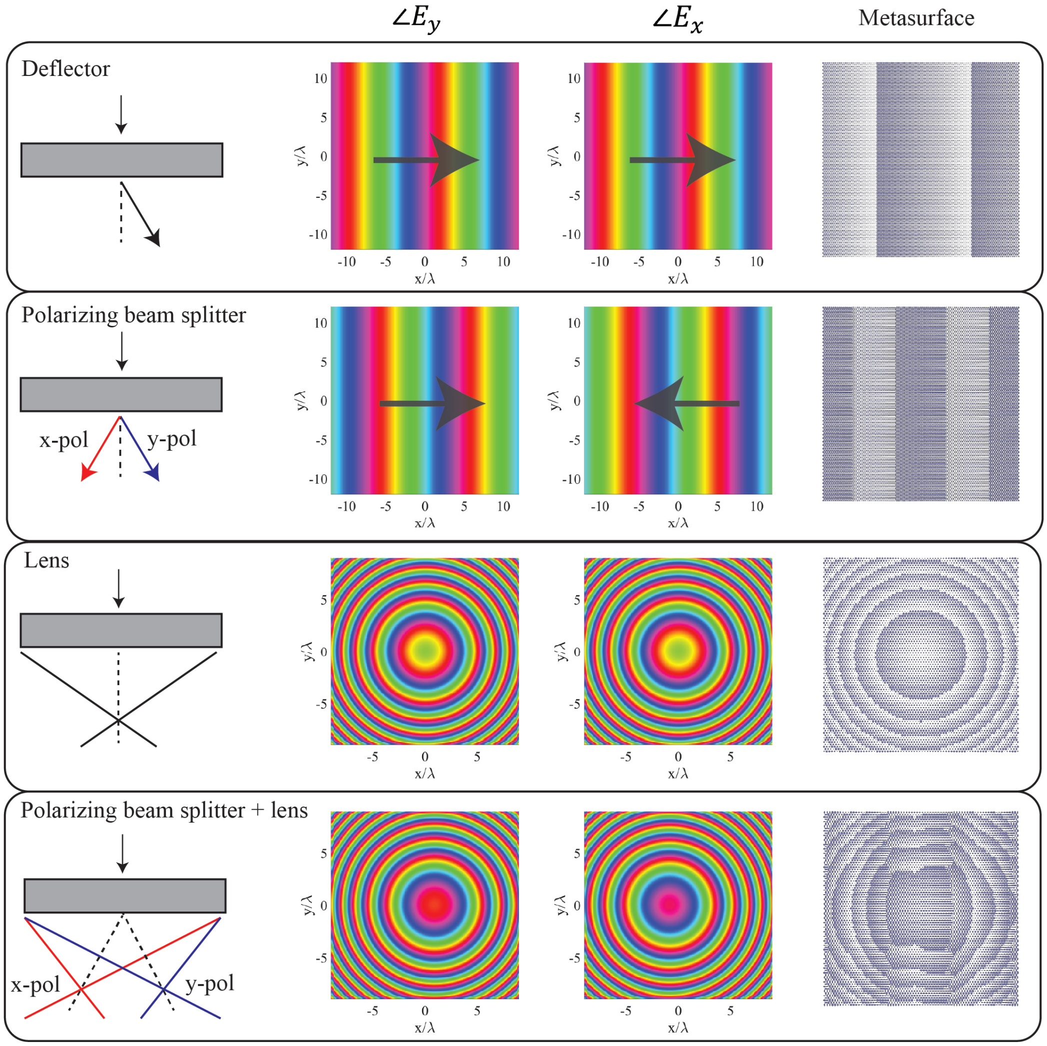

Figure 2.3: **Metasurface examples.**

With a basic overview of metasurface design as context the shortcomings of the methodology become more apparent. The phase profiles that were sampled can be done with very good accuracy, but typically only over a small bandwidth. The phase response of a meta-atom at one frequency is typically highly correlated with its response at other frequencies, thus making it difficult to independently engineer the response at different frequencies which manifest in strong chromatic aberration.

To some extent the function used to find the optimal elements (the ones that made it into the look-up table) can be made stricter to impose functionality at multiple frequencies. One promising effort controls the group delay of light within the meta-atoms [20]. These methods, which fundamentally rely on diffractive dispersion rather than the refractive dispersion utilized in bulk components like lenses, tend to fail as the size of the metasurface becomes increasingly large.

The metasurface is also limited in its multifunctional capabilities. For example, a metasurface can impose an arbitrary phase profile on two orthogonal polarization simultaneously, and it can convert any two arbitrary polarizations to any other two arbitrary polarization, but a single metasurface cannot perform both of these tasks. While this distinction perhaps seems esoteric, it gets at the nature of the limitation of metasurfaces — stated more simply, a metasurface can do A or B, but it cannot do A and B.

A single metasurface also cannot prescribe an amplitude profile to impinging light, at least not efficiently. The amplitude profile of the output light is important for coupling to a desired mode, in which case the amplitude, phase, and polarization must all match the target mode for max power transmission. Recall that in the example above the metasurface elements that were chosen for the lookup table were ones that exhibited high power transmission. Could we instead choose elements that also have some target amplitude transmission to better match to a mode? The answer is technically yes, but the physical mechanism by which this amplitude matching is achieved ends up hurting the power coupling efficiency in a different way. Since the metasurface is so thin (here we will say it is completely flat), power is not spatially redistributed within the metasurface layer. In fact, if it were then the assumption that nearby elements do not interact with each other would not be valid. Because of this the mechanisms by which amplitude is altered would be either reflection, absorption, or polarization rotation, all of which are undesired because that power is ultimately lost (not coupled to the desired mode).

To summarize, the main limitations of metasurface are a limited bandwidth and limited degrees of freedom to prescribe broad multifunctionality. One method to overcome these limitations is to add more metasurfaces. For instance, two metasurfaces can be used to match both the amplitude and phase of a target mode, rather than just the phase. Furthermore, many metasurfaces have been cascaded to create spectrometer systems operating over a relatively broad band compared to a single large-area metasurface. In general, moving to systems of metasurfaces is a popular choice for addressing the shortcomings of single metasurfaces, which falls back on the traditional design paradigm of adding more components to a system to enable increasingly complex functionality. Placing metasurfaces nearby one another invalidates the assumptions used in metasurface design, and thus the size of the systems tends to be dominated by the need for free-space propagation between cascaded metasurface elements.

2.3 Topology optimization: adjoint-based inverse design of 3D metaoptics

Metaoptic devices can be improved by adding thickness to them. While metasurfaces can benefit from increasing thickness of an individual layer (e.g. to increase phase coverage of the elements), in general it is desired to pattern the device in all three dimensions to take full advantage of the optical degrees of freedom available in the design volume. This patterning can be done in several ways: full 3D fabrication in which fabrication tolerance are similar in all directions; 2.5D fabrication, where individual layers are stacked directly on top of each other; or cascaded systems of metasurfaces, where individual metasurface are sufficiently separated to preserve their independent functionality. Cascaded systems are a topic of interest, but are not covered in detail here. Such systems benefit from an intuitive design process, in which individual elements are designed and combined to achieve complex functionality. However, designs rely on sufficiently large separation of the individual elements to preserve the design assumptions. Thus, these systems tend to be much larger than is truly necessary.

The availability of 3D fabrication or 2.5D fabrication is largely dependent on the frequency regime that the metaoptics device is designed for, and the corresponding fabrication techniques available to generate sub-wavelength feature sizes in large enough designs volumes. For example, standard fused filament fabrication (FFF) is capable of generating sub-wavelength 3D features with low-loss material at radio frequency (RF) and microwave frequencies. From terahertz (THz) to some infrared (IR) frequencies, some technology exists to print 3D features with direct-laser lithography, however materials choices are limited and material absorption creates barriers to realizing low-loss devices. From THz to optical frequencies, 2.5D fabrication techniques involving stacked layers can generate sub-wavelength features in a variety of material systems, including low-loss materials such as high-resistivity Si at THz and infrared frequencies, and TiO2 or SiO2 for optical frequencies.

Multiple scattering events within the 3D volume preclude the use of traditional look-up table techniques used for metasurfaces. Instead, full-wave solvers such as finite-difference time-domain (FDTD) must be used to simulate all of the electromagnetic phenomenon associated with the device. By relaxing assumptions the individual simulations take much longer, so the optimization technique must be efficient to minimize the number of simulations. Furthermore, the optimization technique must be able to guide the optimization of the device towards designs that respect various fabrication constraints.

To achieve this we turn to *inverse design* in this work, which involves defining an objective function (i.e. FOM) and iteratively perturbing a design region to maximize the FOM. The inverse design procedure is guided by the gradient of the FOM with respect to material permittivity, which is used to inform the optimization whether increasing or decreasing permittivity at a specific region will improve the device. To minimize the number of full-wave simulations required, the adjoint-method is used. This exploits adjoint symmetry in Maxwell's equations such that only two full-wave simulations are required to compute the desired gradient *at every point in the design region*. Knowing this gradient, the device can be adjusted to increase the FOM while simultaneously enforcing fabrication constraints to ultimately define the topology of an optimal device. This recipe for optimization is called *topology optimization*. Its theory and implementation is the primary focus of this section.

Adjoint state method and gradient descent

The key enabling feature of photonic topology optimization is the rapid retrieval of the gradient of a desired FOM. This can be done using the *adjoint state method*, referred to in this thesis more simply as the *adjoint method*. The adjoint method was first developed in mechanical applications, such as aerodynamics and fluid mechanics, to find optimal shapes for structures (see 1 for more historical context). The adjoint method relies on symmetry, and in the case of electromagnetic design stems from the reciprocity of Green's functions. For a detailed look at the principles of the adjoint method as it applies to electromagnetic problems, see Chapter 5 of Owen Miller's 2012 thesis from the University of California Berkeley [21].

Consider the case in which a plane wave is incident on a design region, and the FOM is the intensity of the electric field at a point in space x_0. By changing the permittivity of a voxel in the design volume at a point x a dipole moment is introduced, which in turn alters the FOM. To analyze the effect of this dipole moment, we could run a forward simulation, then run a second forward simulation that includes the dipole moment at x and study the resulting change in the field at x_0. Alternatively, the second simulation could instead feature a source located at x_0 in which we study the field at x, due to reciprocity. The key difference is the latter simulation can be used to analyze the effects of dielectric inclusions at *all voxels in the design region simultaneously*, since the information is derived from the simulation of the single dipole at x_0 rather than a series of simulations. That is, the effect of altering the dielectric constant at any point x_i in the design volume can be understood by running the simulation with a dipole at x_0 and studying the field at x_i, along with the related forward simulation. The gradient of the FOM can be found by simulating just the forward simulation (i.e. the incident plane wave), and a reverse simulation comprised of a dipole source at x with appropriate phase, which we call the *adjoint simulation*.

The adjoint source (the dipole in the above example) is dependent on the FOM and the phase of the forward simulation, and its electric and magnetic field terms are given by:

$$\mathbf{E}_{adj} = \frac{\partial f(\mathbf{E},\mathbf{H})}{\partial \mathbf{E}} \tag{2.5}$$

$$\mathbf{H}_{adj} = -\frac{1}{\mu_0}\frac{\partial f(\mathbf{E},\mathbf{H})}{\partial \mathbf{H}} \tag{2.6}$$

The FOM can be a collection of dipoles. This occurs when the FOM is evaluated over more than one point in space. In fact, any spatial function of electric and magnetic fields has an associated adjoint source as long as the FOM is differentiable.

Once the gradient of a FOM is efficiently calculated at all points in the design region using the adjoint method, the structure can be altered to improve the FOM. At this point, the physics of the problem has been used to efficiently compute the gradient and what remains is a gradient-based mathematical optimization problem. The functions being minimized are generally non-linear, thus a computed gradient is only valid in a locally linear region of f. Gradient descent is therefore an iterative optimization, comprised of computing ∇f, stepping f in the direction of $-\alpha\nabla f$ where α is a sufficiently small step size, and concluding the optimization when ∇f is sufficiently close to 0. A major enabling technique for gradient-based optimization is the *adjoint method* for efficiently computing a gradient. The technique has been applied to many fields (see [22] for a recent overview of the topic), including photonic design [23].

The gradient-descent procedure becomes increasingly complex as the dimensions of the parameter space increase and as constraints are imposed. For optimizing the topology of photonic devices, the parameter space includes the permittivity of every voxel in the design region, which can be arbitrarily large. For many photonic devices, the desired solution is constrained due to the limitations of available fabrication techniques that make only certain solutions physically realizable. Some common requirements of current fabrication techniques include:

- *Minimum feature size*: the size of the smallest uniform material piece. this is typically dictated by lithography, material etching, and material deposition capabilities.
- *Connectivity*: one or more materials are topologically connected. Ensuring the topology of the device is for mechanical considerations, which is useful in techniques such as 3D-printing.
- *Binary*: the device is made of only two materials.

There are many topology optimization algorithms and many reviews summarizing them [24]. They can be characterized as density-based optimization or shape optimization: density-based optimizations feature a density variable ρ that is evaluated across a grid of points, thus describing the device in an element-wise way. Density-based optimization can be further classified as discrete, where $\rho = 0$ or $\rho = 1$, or continuous, where $0 \leq \rho \leq 1$. Shape-based optimizations instead describe the boundary of the device Ω. One common technique for shape-based optimizations involves describing the boundary Ω as the level-set contour of a higher-dimensional function, called the level-set method. In the context of photonic inverse design, two methods have dominated the literature: continuous density-based optimizations, and shape optimizations using the level-set method.

Of the fabrication constraints listed above, the *binary* constraint is a substantial difference between the continuous density-based and the discrete level-set approaches to photonic optimization, with the former inherently describing a non-binary device and the latter inherently describing a binarized device. While the level-set method has the advantage of describing binarized devices, a known limitation is its ability to easily nucleate holes, bridge gaps, and generally converge to shapes that are substantially different from the starting topology. A technique that is now employed in photonic optimization is to begin with a density-based optimization, then switch to a level-set method to conclude the optimization [25]. Given the limitations of the level-set method, it is important that the density-based optimization converge to a sufficiently good starting point for the level-set method. In low contrast structures, the binarization process usually automatically finished as the optimization objective is optimized towards the gradient. However, in high contrast structures, the local minimum of the objective could be related to a grey-scale material. This represents a challenge for devices that have a high electric permittivity contrast between the two materials.

The goal of binarizing a continuous density-based optimization is not new, and is not limited to photonic inverse-design. Commonly used methods for this task include projection filters, which encourage the device to become more binary but fall short of explicitly requiring the device become binary. More advanced methods employing gradually strengthened sigmoidal or Heaviside filters have been used to guarantee convergence to a binary solution, but these also cannot fully guarantee that the device becomes binary enough that the final thresholding operation (in which all voxels are rounded to 0 or 1) does not incur a severe performance penalty that requires substantial level-set re-optimization to overcome.

2.4 Constraining continuous topology optimizations to discrete solutions

We present a method here that reformulates the gradient-based continuous density optimization in a manner that strictly enforces that the design become fully binary, thus preventing the option of settling at a non-binary solution. We have found this technique to be particularly useful in optimizing high index contrast devices. The method uses the gradient information at each iteration to step the permittivity in a direction that maximizes the performance of the device while constraining the permittivity step to binarize the device by some amount that is greater than zero. We refer to this procedure as "sub-optimization", since it occurs at every iteration of the overall photonic optimization. An important quality of this sub-optimization is that much of the computational burden of finding the ideal constrained permittivity is eased by instead solving the Lagrange dual-problem of the original sub-optimization, which converts the original N-dimensional linear optimization (where N is the number of voxels in the design region) into a 1-dimensional non-linear optimization which can be easily solved computationally.

The remainder of the section is composed as follows: first, we formulate the sub-optimization and derive the solution of the Lagrange dual problem; then, we provide an example of an optimized spectral demultiplexer. The examples are intended to illustrate the practical usage of the method, compare its performance to a traditional steepest-ascent optimization, integrate the technique with a subsequent level-set optimization, and describe a technique to control the overall fraction of material in the final device.

Lagrange dual problem

Setup

Consider a discretized 3D domain of interest $\mathcal{P}$. The domain contains a region of dielectric permittivity $\epsilon_i, i \in \mathcal{D} \subseteq \mathcal{P}$, where $\mathcal{D}$ is called *design region* or *device*. The objective function $f(\mathbf{E}, \mathbf{H}, \vec{\epsilon})$ depends on the electric and magnetic fields output from the device, and the permittivity of the design region $\vec{\epsilon}(x)$ with $\epsilon_i \in [\epsilon_{min}, \epsilon_{max}]\ \forall i \in \mathcal{D}$. The general goal of the inverse-design optimization is to find the optimal permittivity of the design region that maximizes $f(\mathbf{E}, \mathbf{H}, \vec{\epsilon})$, and here we wish to constrain the solutions to only permittivity distributions consisting of two materials. ϵ_{min} and ϵ_{max} are the lower and upper bounds of the permittivity, respectively, representing the two materials the final device is comprised of. The electric field $\mathbf{E}$ and magnetic field $\mathbf{H}$ inside the domain can be solved for using Maxwell's equation solvers such as the finite-difference-time-domain (FDTD). The adjoint method computes the gradient $\partial f / \partial \vec{\epsilon}$ at all points within the design region.

The quantities involved in 3D simulations, such as the device permittivity $\vec{\epsilon}$ and FoM gradient $\partial f / \partial \vec{\epsilon}$, are often computed and stored as 3D matrices. Here, we find it convenient to flatten these into N-dimensional column vectors, where N is the total number of voxels in the device region. For the remainder of this derivation, all multi-dimensional quantities are considered N-dimensional vectors.

We define a metric to quantify how binary the device is. In this case, we use a simple absolute value function to define the binarization of a single pixel, whose mean across all pixels quantifies the binarization of the device. The absolute value function is scaled and shifted such that it evaluates to 1 at $\epsilon = \epsilon_{min}$ and $\epsilon = \epsilon_{max}$, 0 at the permittivity midpoint, ϵ_{mid}.

$$B = \frac{1}{N} \sum_i \left| \frac{\epsilon_i - \epsilon_{mid}}{\epsilon_{max} - \epsilon_{mid}} \right|, i \in \mathcal{D} \tag{2.7}$$

$$\epsilon_{mid} = \frac{\epsilon_{min} + \epsilon_{max}}{2} \tag{2.8}$$

If we define $\vec{b} = dB/d\vec{\epsilon}$, then the approximate change in binarization of the device given a permittivity step $\vec{\Delta\epsilon}$ is $\vec{b}\cdot\vec{\Delta\epsilon}$. This term can thus be constrained to predictably control the change in device binarization.

Direct gradient ascent would shift the device permittivity in the direction of $\partial f/\partial\vec{\epsilon}$. However, the gradient may not point in a direction that increases the binarization of the device, so we reformulate the optimization to *improve the objective function as much as possible while enforcing that the binarization increases by some amount* β. This corresponds to solving the following constrained optimization problem, which we refer to as the sub-optimization problem from now on since it is solved at every iteration of the main photonic optimization.

$$\begin{aligned} \underset{\vec{\Delta\epsilon}}{\text{maximize}} \quad & \frac{\vec{\partial f}}{\partial\epsilon}\cdot\vec{\Delta\epsilon} \\ \text{subject to} \quad & \vec{b}\cdot\vec{\Delta\epsilon} \geq \beta, \\ & \Delta\epsilon_i \leq u_i \quad \forall i, \\ & \Delta\epsilon_i \geq l_i \quad \forall i \end{aligned} \tag{2.9}$$

The vectors $\vec{u}$ and $\vec{l}$ control the maximum step size that the permittivity can undergo in a single iteration. They are also used to enforce the maximum and minimum permittivity constraints (i.e. $\epsilon_i \in [\epsilon_{min}, \epsilon_{max}]\ \forall i$) since they can be set to 0 for a permittivity voxel at ϵ_{min} or ϵ_{max}.

We convert (2.9) to a canonical form.

$$\begin{aligned} \underset{\vec{x}}{\text{minimize}} \quad & \vec{c}\cdot\vec{x} \\ \text{subject to} \quad & \beta - \vec{b}\cdot\vec{x} \leq 0, \\ & x_i - u_i \leq 0 \quad \forall i, \\ & l_i - x_i \leq 0 \quad \forall i \end{aligned} \tag{2.10}$$

Since $\vec{x} \in \mathbb{R}^N$, there are $2N + 1$ inequality constraints in (2.10) in addition to the N dimensions of the non-constrained optimization problem. For large N, this problem is challenging to solve numerically. Fortunately, this multi-dimensional linear optimization can be reduced to a single-dimensional non-linear optimization, regardless of the value of N. This is done by instead solving the *Lagrange dual problem*, which in this case will have the same optimal point as (2.10) since the property of *strong duality* holds for linear optimizations [26].

Solution

To obtain the dual problem, we start with the *Lagrangian*. This is defined as:

$$\begin{aligned} L(\vec{x}, \vec{\lambda}_1, \vec{\lambda}_2, \nu) &= \vec{c} \cdot \vec{x} + \vec{\lambda}_1 \cdot (\vec{x} - \vec{u}) + \vec{\lambda}_2 \cdot (\vec{l} - \vec{x}) + \nu(\beta - \vec{b} \cdot \vec{x}) \\ &= \vec{x} \cdot (\vec{c} + \vec{\lambda}_1 - \vec{\lambda}_2 - \nu\vec{b}) - \vec{\lambda}_1 \cdot \vec{u} + \vec{\lambda}_2 \cdot \vec{l} + \nu\beta \end{aligned} \tag{2.11}$$

The *dual function* is defined as the infimum of the Lagrangian over $\vec{x}$:

$$g(\vec{\lambda}_1, \vec{\lambda}_2, \nu) = \inf_{\vec{x}} L(\vec{x}, \vec{\lambda}_1, \vec{\lambda}_2, \nu) \tag{2.12}$$

Since (2.11) is linear, the infimum is $-\infty$ if any slope is non-zero.

$$g(\vec{\lambda}_1, \vec{\lambda}_2, \nu) = \begin{cases} -\infty & \vec{c} + \vec{\lambda}_1 - \vec{\lambda}_2 - \nu\vec{b} \neq \vec{0} \\ -\vec{\lambda}_1 \cdot \vec{u} + \vec{\lambda}_2 \cdot \vec{l} + \nu\beta & \vec{c} + \vec{\lambda}_1 - \vec{\lambda}_2 - \nu\vec{b} = \vec{0} \end{cases} \tag{2.13}$$

The optimization we wish to solve is called the *Lagrange dual problem*:

$$\begin{aligned} \text{maximize} \quad & g(\vec{\lambda}_1, \vec{\lambda}_2, \nu) \\ \text{subject to} \quad & \vec{\lambda}_1 \geq 0, \\ & \vec{\lambda}_2 \geq 0, \\ & \nu \geq 0 \end{aligned} \tag{2.14}$$

Incorporating (2.13) into Eq. 2.14, we get:

$$\begin{aligned} \text{maximize} \quad & -\vec{\lambda}_1 \cdot \vec{u} + \vec{\lambda}_2 \cdot \vec{l} + \nu\beta \\ \text{subject to} \quad & \vec{\lambda}_1 \geq 0, \\ & \vec{\lambda}_2 \geq 0, \\ & \nu \geq 0, \\ & \vec{c} + \vec{\lambda}_1 - \vec{\lambda}_2 - \nu\vec{b} = 0 \end{aligned} \tag{2.15}$$

The fourth constraint of (2.15) can be used to write the objective function and the second constraint in terms of only ν and λ_1.

$$\begin{aligned} \text{maximize} \quad & -\vec{\lambda}_1 \cdot \vec{u} + (\vec{c} + \vec{\lambda}_1 - \nu\vec{b}) \cdot \vec{l} + \nu\beta \\ \text{subject to} \quad & \vec{\lambda}_1 \geq 0, \\ & \vec{\lambda}_1 \geq \nu\vec{b} - \vec{c}, \\ & \nu \geq 0 \end{aligned} \tag{2.16}$$

For any ν, the optimal elements of $\vec{\lambda}_1$ are as small as possible since $\vec{u} \geq 0$ and $\vec{l} \leq 0$. The solution to $\vec{\lambda}_1$ is obtained by satisfying the constraints while minimizing $\vec{\lambda}_1$. The solution for the i-th element of $\vec{\lambda}_1$ is thus:

$$\lambda_{1,i} = \begin{cases} 0 & \nu b_i - c_i \leq 0 \\ \nu b_i - c_i & \nu b_i - c_i > 0 \end{cases} = (\nu b_i - c_i)H(\nu b_i - c_i) \tag{2.17}$$

where $H(x)$ is the Heaviside step function. $\vec{\lambda}_2$ is solved using the fourth constraint of (2.15).

$$\vec{\lambda}_2 = \vec{c} + \vec{\lambda}_1 - \nu\vec{b} = \vec{c} + (\nu\vec{b} - \vec{c})H(\nu\vec{b} - \vec{c})^T - \nu\vec{b} \tag{2.18}$$

Here, $H(\vec{x})$ represents the Heaviside function applied element-wise to $\vec{x}$. Since the optimal $\vec{\lambda}_1$ and $\vec{\lambda}_2$ are dependent on ν, the dual problem is a function of only $\nu \in \mathbb{R}$. The optimization can be reduced to a single dimension as follows.

$$\begin{aligned} \underset{\nu}{\text{maximize}} \quad & \nu\vec{b} \cdot \vec{l} + (\vec{u} - \vec{l}) \cdot (\nu\vec{b} - \vec{c})H(\nu\vec{b} - \vec{c})^T - \nu\beta \\ \text{subject to} \quad & \nu \geq 0 \end{aligned} \tag{2.19}$$

This problem can be solved numerically using open-source functions. We denote the optimal dual variables as $(\nu^\star, \vec{\lambda}_1^\star, \vec{\lambda}_2^\star)$. The optimal solutions of the dual problem can be mapped back to original variable x using the *Karush–Kuhn–Tucker (KKT) conditions*, which are a set of conditions that hold under certain regularity conditions. Fortunately, linearity of the constraints is a sufficient condition for the KKT conditions to hold. We use the KKT condition of *complementary slackness* to translate $(\nu^\star, \vec{\lambda}_1^\star, \vec{\lambda}_2^\star)$ to the optimal $\vec{x}$ that we call $\vec{x}^\star$.

$$\vec{\lambda}_1^\star(\vec{x}^\star - \vec{u})^T = \vec{0} \tag{2.20}$$

$$\vec{\lambda}_2^\star(\vec{l} - \vec{x}^\star)^T = \vec{0} \tag{2.21}$$

The solution is

$$x_i^\star = \begin{cases} u_i & \lambda_{1,i} \neq 0 \\ l_i & \lambda_{1,i} = 0 \end{cases} \tag{2.22}$$

Example

This section provides example optimizations to illustrate the described techniques usage and elaborate on several key points. We study device performance relative to a basic implementation of regular gradient ascent in which the permittivity is stepped in the exact direction of the gradient. We study this for three different cases of material-to-void refractive index contrast: a low index contrast of 1.5:1, a medium index contrast of 2.5:1, and a high index contrast of 3.0:1. Next, we pair this optimization technique with a level-set optimization that thresholds the continous permittivity to a binary solution and re-optimizes the device with a level-set optimization to recover the lost performance from the threshold operation. Finally, we demonstrate how this technique can be used to control the overall ratio of material:void in the final device simply by shifting the binarization function in Eq. 2.7.

The example optimization we use here is a spectral demultiplexer, or "color splitter". Fig. **2.4**a illustrates the purpose of the device, which focuses three equally-sized frequency bands of a normally incident TE plane-wave to three distinct points in the focal plane. The FoM function is defined as the intensity at the desired point in the focal plane, which allows us to use a simple dipole source for the adjoint source. When quantifying the performance of the device we use the power transmission through apertures centered at the dipole source. These apertures are drawn in Fig. **2.4**b as red, green, and blue horizontal lines in the focal plane. The gradient of every FoM (there is one FoM for each simulated frequency) are combined using a weighted average to obtain a single gradient vector, which is input as $\partial f/\partial\vec{\epsilon}$ in Eq. 2.9. This weighting procedure is described in detail in [27].

The device is optimized over a 42% fractional bandwidth, which is comparable to fractional bandwidth of visible light. λ_0 and f_0 denote the center wavelength and frequency of the full bandwidth of the device. In Fig. **2.4**b the intensity evaluated at the various frequencies are overlaid using their equivalent hue in the visible regime. The quantitative performance of the device is shown in Fig. **2.4**c, which plots the power transmission into three monitors of size $1.56\lambda_0$, centered at the relevant focal points. Fabrication tolerance limits are not strictly enforced in this design, although the discretization of the geometry on a $\lambda_0/30$ grid and the discretization of the FDTD simulation on a $\lambda_0/10n_{max}$ grid preclude arbitrarily small features in the design. After all optimizations, we simulated the final device on a finer grid to ensure that the FDTD results had properly converged. We noticed very little change to device behavior and efficiency after halving the simulation grid size in all directions.

A video illustrating the evolution of the optimization is available on the Caltech-DATA repository [28]. The device begins as a uniform rectangle of material at the midpoint refractive index of the two material boundaries, which in this example had refractive indices of 1 and 1.5. The device quickly achieves the desired spectral demultiplexing functionality, and the optimization continues as the design region is gradually pushed towards a nearly binary solution.

In the optimization problem described by Eq. 2.9, the quantity β describes the amount that the device must binarize after stepping the permittivity. In our implementation we set this quantity to be a factor of the max possible binarization, denoted β_{max}, during the current iteration. β_{max} is not constant during the optimization, since points on the device eventually reach a material boundary. The maximum and minimum allowable step size for each permittivity point in the device is contained in the vectors $\vec{u}$ and $\vec{l}$, respectively, which can be used to find β_{max}. For this example, we use $\beta = 0.2\beta_{max}$. The maximum and minimum permittivity step sizes are chosen to be 0.02. The algorithm used in this example is summarized in Algorithm 1.

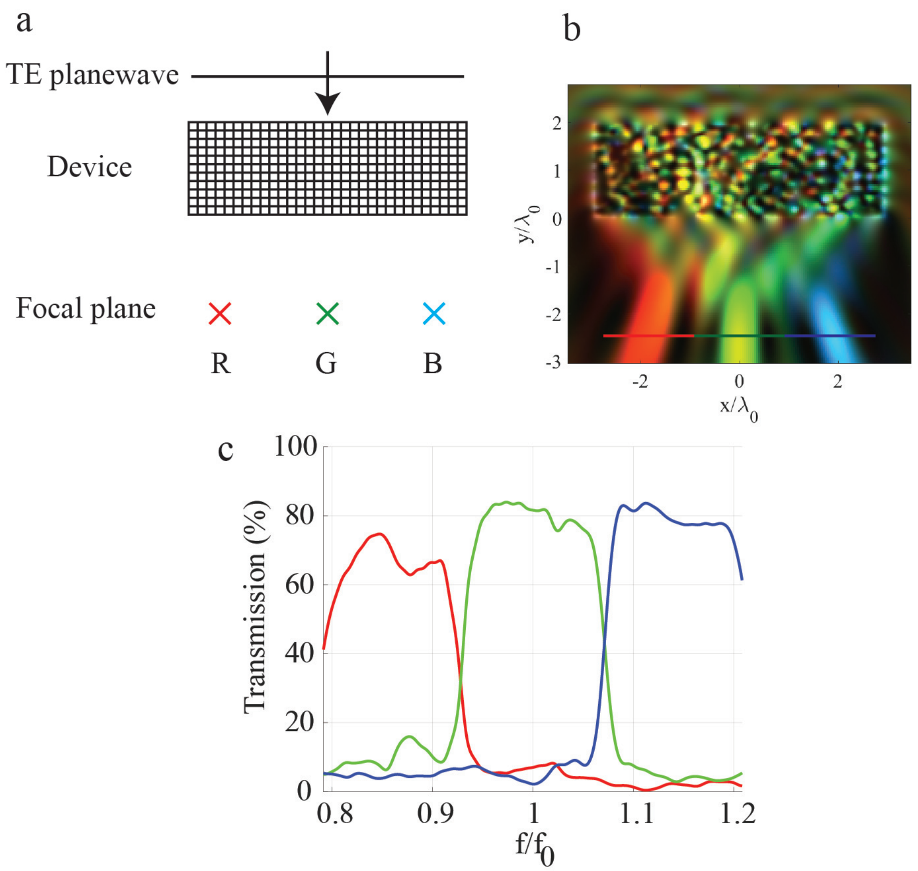

Figure 2.4: **Example device.** (a) A schematic of the device. The device is a spectral demultiplexer that focuses three equally spaced frequency bins to three distinct points marked in the focal plane. The input is a normally incident TE-polarized planewave. (b) The intensity $|E|^2$ of each frequency overlaid. Each frequency is drawn in its equivalent hue in the visible regime. (c) The power transmission of the output fields through the apertures. Each trace is drawn in the same color as the corresponding aperture drawn in (b).

Algorithm 1

```
initialize device permittivity;
while $B < B_{end}$ do
    compute gradient using adjoint method, averaging multiple FoM gradients as necessary for multi-functional performance;
    compute $\vec{u}$ and $\vec{l}$ based on the current permittivity and the maximum allowable step size;
    compute $\beta_{max}$ based on the values of $\vec{u}$ and $\vec{l}$, and choose $\beta_0$ such that $\beta = \beta_0\beta_{max}$;
    solve the optimization problem Eq. 2.19 using a numerical solver;
    use Eq. 2.17 to solve for $\lambda_1^\star$, then use Eq. 2.22 to solve for $x^\star$;
    update the permittivity of the device by $x^\star$;
end
```

We compare our method with a basic implementation of gradient ascent in order to demonstrate the differences in convergence properties and final device performance. In the basic implementation of gradient ascent, we step the device permittivity in the exact direction of the electromagnetic FoM gradient (i.e. a *steepest ascent* optimization). The permittivity update is computed with $\alpha(\partial f/\partial\vec{\epsilon})$, where α controls the specific step size. To ensure both stability and acceptable convergence, we dynamically alter α such that the maximum step in the absolute value of permittivity is 0.1. The reason for comparing these two methods is to observe their different behaviors, particularly with respect to converging to binary devices. We do not compare the speed of convergence. However, techniques for improving convergence speed (e.g. a *line search* for computing an optimal step size) apply to both methods.

The results of the traditional gradient ascent, which we call *direct gradient ascent*, and the technique that this paper describes, which we call *modified gradient ascent* here, are shown in Fig. **2.5** for different value of material:void refractive index contrast. The plotted average FoM is the transmission through the desired aperture averaged over all frequencies. In the low contrast case (1.5:1) in Fig. **2.5**a, the optimizations perform equally well in terms of FoM. In the medium contrast case (2.5:1) in Fig. **2.5**b, the direct gradient ascent optimization converges to a final FoM value of 84% while the binarization stagnates at 54%. This implies that optimization has found a local maximum of the FoM where the device is not binary. This type of solution is forbidden in the modified gradient ascent approach, which after 500 iterations is at a point where the FoM is 80% for both the FoM and binarization (the fact that FoM and binarization percentage are nearly identical is a coincidence). In fact, the FoM reaches a maximum at iteration 300, but continues to search for a solution that is binary in order to satisfy the binarization constraint. It is necessary that the maximum FoM of a binarized device is less-than-or-equal to the maximum FoM of a greyscale device, since the former is a subset of the latter (it is possible that a binarized device exists at a better *local* optimum than a seperately optimized greyscale device). The difference between the two methods becomes yet more apparent in the high index contrast case (3.0:1) in Fig. 2.5c. In this case, the direct gradient ascent optimization converges to a 48% binary device after 1000 iterations while the modified gradient ascent optimization reaches 80% binarization after 1000 iterations.

While the modified gradient ascent approach described above will eventually converge to a 100% binary device, this is not strictly required. Instead, the optimization can be concluded before the device is 100% binary, thresholded to the nearest material boundary, and then used as an initial seed to a level-set optimization. Level-set optimizations tend to be dependent on initial seed [24], and we study the dependence of the converged performance of a level-set optimization with respect to the binarization percentage of an initial seed obtained with the modified gradient ascent approach. For this study, we use an implementation of level-set optimization based on signed distance functions [29]. Fig. 2.6a shows the converged average FoM of a 2.5:1 medium index contrast device as a function of the binarization percentage of the initial seed and Fig. 2.6b shows the convergence curves for the individual level-set optimizations. The results suggest that the level-set optimization does indeed depend on initial seed, and that in general the more binarized starting seeds yield a better level-set optimized device. Beyond 40% binarization, the performance monotonically increased with starting binarization, however the diminishing gains suggest that terminating the density-based optimization around 80% will yield nearly optimal performance while minimizing the time spent in the density-based phase of the optimization.

The definition of the binarization function can be shifted with respect to permittivity in order to converge to devices with different overall material:void ratio. Eq. 2.7 uses the mid-point of the two material permittivities as the center of the absolute

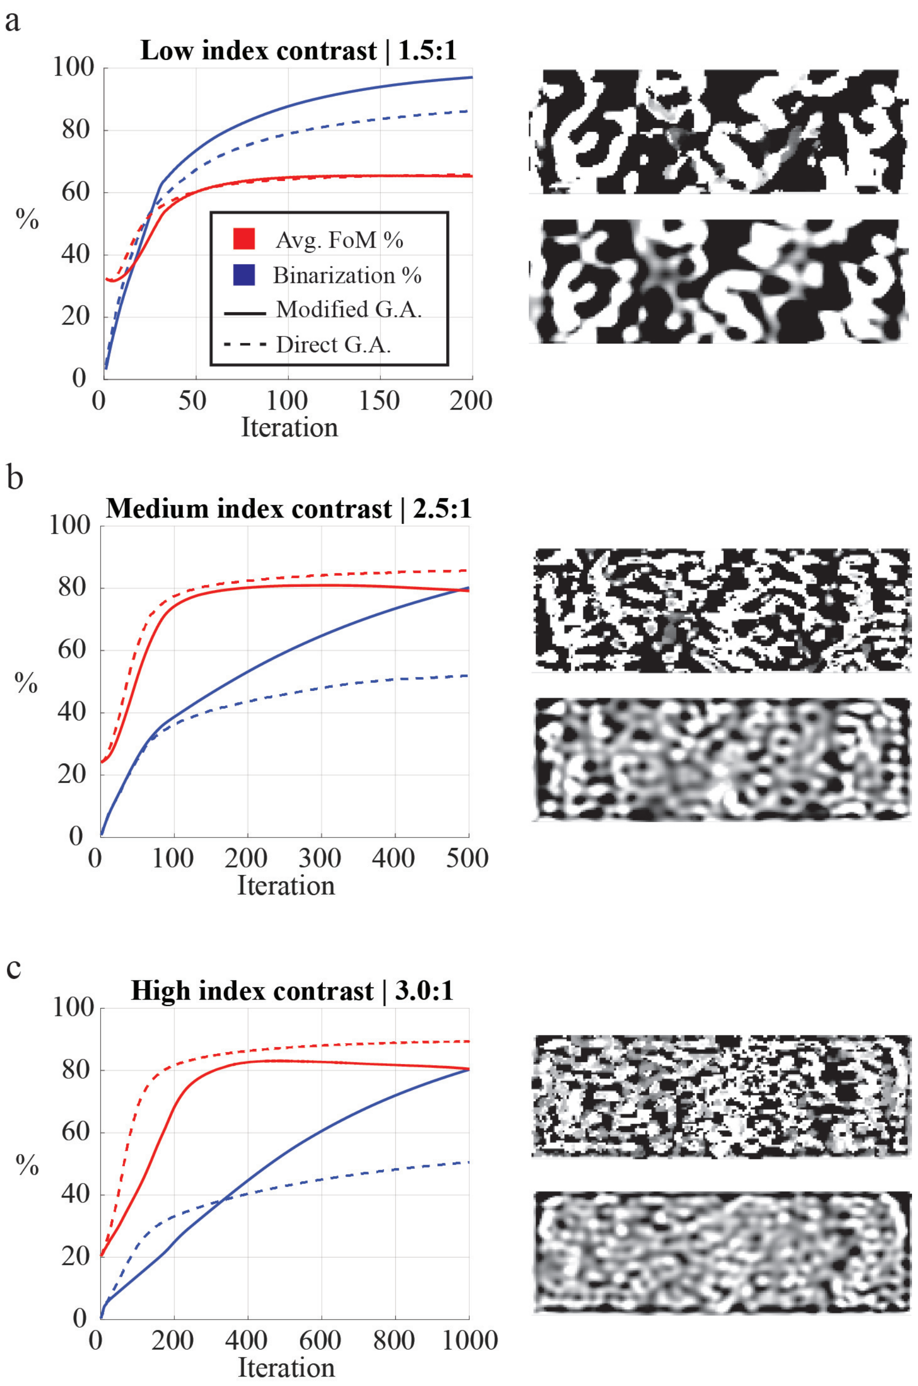

Figure 2.5: **FoM and binarization convergence curves for different index contrasts.** The FoM (averaged across all functionalities) and binarization are drawn as red and blue traces, respectively. The modified gradient ascent (G.A.) and direct gradient ascent methods are drawn as solid and dashed lines respectively. The final index profile of each device is shown as a greyscale colormap. In each subset, the device optimized with the modified gradient ascent procedure is the top device while the device optimized with the direct gradient ascent procedure is the bottom device. The index contrasts are: (a) 1.5:1, (b) 2.5:1, (c) 3.0:1. A video showing the evolution of the low index contrast example in (a) is available on the CaltechDATA repository [28].

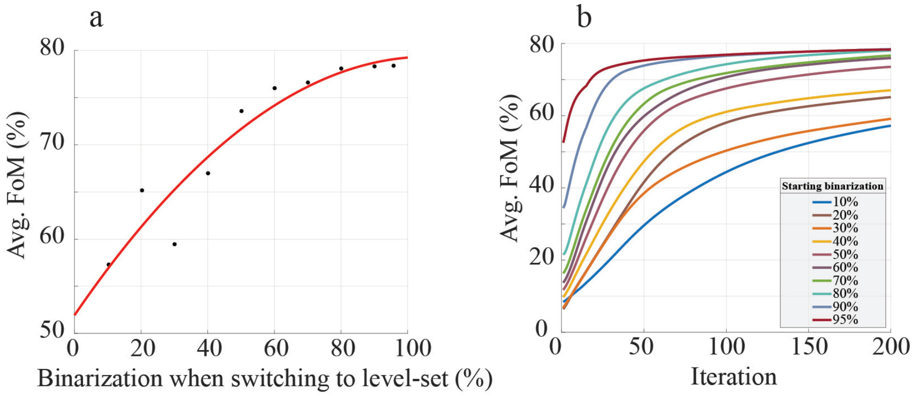

Figure 2.6: **Combining the continuous density-based optimization with a level-set optimization.** (a) The level-set optimization uses an initial seed given by the results of the continuous density-based optimization at different starting binarization values. The average FoM is plotted as a function of this starting binarization. The red curve is a quadratic fit. (b) The convergence curves of each level-set optimization.

value function. We can introduce a shift term, $\delta\epsilon$:

$$B = \frac{1}{N}\sum_{i}\left|\frac{\epsilon_i + \epsilon_{mid} - \delta\epsilon}{\epsilon_{max} - \epsilon_{mid}}\right|, i \in \mathcal{D} \tag{2.23}$$

Intuitively the binarization constraint in Eq. 2.9 can be thought of as defining whether a permittivity point should increase or decrease in order to binarize the device more, depending on its current value of permittivity. In Eq. 2.9 (with no shift about the midpoint), points that are above the midpoint are encouraged to increase further since this is the direction of binarization, and similarly points below the midpoint are encouraged to decrease further. Thus, shifting the cusp away from the midpoint affects how the permittivity is encouraged to become more binary. In general, shifting the cusp of B to the left (positive $\delta\epsilon$) yields a device with more material, and vice versa. This can be a useful feature for enabling shorter maskless lithography write times or ensuring mechanical robustness without incurring a substantial decrease in device performance.

This technique is demonstrated in a low index contrast (1.5:1) device. Fig. 2.7a shows the final material:void fraction as a function of $\delta\epsilon$, which is varied from -0.5 to +0.5 about the mid-point permittivity of 1.625. The resulting material distribution is also shown for each optimized device. Fig. 2.7b shows the evolution of the material:void fraction over the course of the optimization for the different $\delta\epsilon$ values. The largest change in material:void fraction tends to occur in the early stages of optimization in this example. Fig. 2.7c shows the evolution of the average FoM for different values of $\delta\epsilon$. The devices with the worst final performance are the ones with the most extreme $\delta\epsilon$, with the device consisting of primarily material being substantially worse than the rest. The best performing device was optimized with $\delta\epsilon = -0.25$.

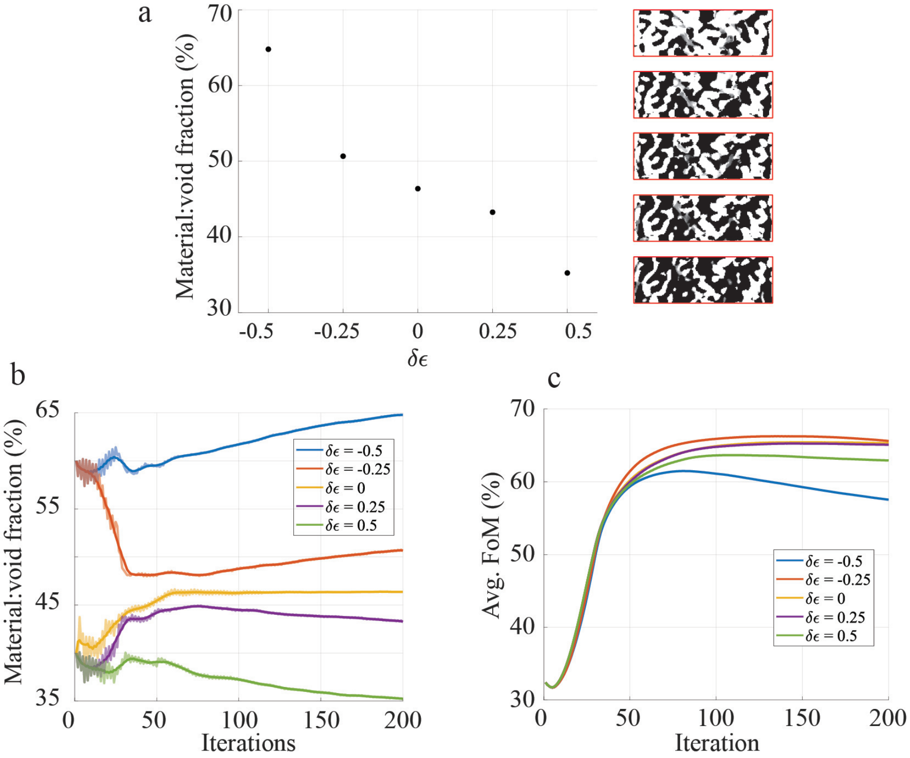

Figure 2.7: **Controlling the overall fraction of material.** (a) The material-to-void ratio for different shifts in the definition of binarization. To the right of the plots is a picture of the resulting optimized device, where white is material and black is void. From top to bottom the devices correspond to $\delta\epsilon = -0.5$, $\delta\epsilon = -0.25$, $\delta\epsilon = 0$, $\delta\epsilon = 0.25$, $\delta\epsilon = 0.5$. (b) The evolution of the material-to-void ratio during the optimization. Each traces shows the raw data and a 5-iteration rolling average. (c) The evolution of the FoM during the optimization.

Discussion

This section presented a sub-optimization technique that is constrained to converge to fully binary solutions. The technique involves solving a sub-optimization at each iteration that maximizes the change in device performance while forcing the step in dielectric permittivity to increase the overall binarization of the device by a certain amount. This technique is computationally infeasible to solve directly with a numerical solver, but the optimization problem is reduced to a one-dimensional non-linear optimization problem by instead solving the Lagrange dual problem.

The usefulness of the technique becomes apparent when optimizing devices with a large refractive index contrast between the two material boundaries, where traditional gradient-based techniques tend to stagnate at non-binary solutions. It is shown that our level-set optimization depends on the initial seed, as has been discussed in literature over the years, and that beginning the level-set optimization with more binary seeds leads to better final device performance. Lastly, we showed an example where the material-to-void fraction of the optimized device can be controlled by simply shifting the definition of binarization.

The specific sub-optimization presented here was intended to constrain the device only to binary solutions. However, there were no assumptions made on the binarization function B other than the existence of a derivative with respect to ϵ. Thus, this function can describe other properties of the device and the optimization procedure can be used to constrain other device properties in future works.

2.5 Imposing a minimum feature size

The electromagnetic optimization techniques that have been presented can be used at any frequency in the electromagnetic spectrum. Unfortunately, fabrication techniques are not so universal, and careful consideration of the capabilities of available fabrication techniques must be considered to ensure a device can be reliably fabricated. The way this challenge is addressed at present is to impose a general *minimum feature size* constraint that restricts that places a lower bound on the allowed gap size and radius of curvature of the many features within the design region. In general, this ensures the device can be fabricated using techniques that scan some patterning element over a grid. Examples of this very general technique that were used for the work in this thesis include include: water jet cutting in Section 3.1, fused-filament-fabrication 3D printing in Section 3.2, and photolithography for the THz devices in Ch. 4. Further examples include two-photon polymerization and electron-beam lithography, which are techniques capable of fabricating subwavelength feature sizes at infrared and optical frequencies, respectively.

The density-based technique only had to impose an rough feature size constraint, since the optimization will ultimately be converted a level-set optimization. The most basic way to do this is to interpolate the geometry of the structure onto a grid that is the same size as the minimum feature size. With this technique, the simulation grid should be kept at whatever value is required to obtain accurate simulation results, and the result interpolated from that grid onto a different grid used to define the geometry. One drawback to this method is that it unnecessarily forces the positions of the geometry to be limited to the geometry grid size, which is overly restrictive. For example, a geometric grid size of 100 nm could be used to limit the minimum feature size to 100 nm, but it will also force features to be integer multiple of 100 nm apart, when in reality having two feature 120 nm may be well within the capabilities of the tool.

A more advanced technique is to incorporate a Gaussian blur to the gradient. With this procedure, the gradient can be interpolated onto a finer geometric grid, then blurred with a Gaussian kernel to enforce a minimum feature size. This has the benefit of retaining accurate positioning resolution. However, this technique is not guaranteed to converge to a device that has a minimum feature size imposed for both material and void features. For example, two blurred material regions can still form a narrow void gap, which is oftentimes as difficult to fabricate as a narrow material gap.

The intricacies of imposing a minimum feature size during the density-based optimization become a secondary concern if the optimization is concluded with a level-set optimization. The minimum feature size constraint imposed during this phase can be corrected for (or ruined) to some extent during the level-set phase. This thesis borrows a clever technique from Ref. [30] which, at each iteration, analytically computes a fabrication penalty function that can be minimized. The penalty function consists of two terms, one based on a minimum gap size, and one based on a minimum radius of curvature. The first forbids the creation of features such as long narrow bridges (of either material or void), and the second forbids the creation of small holes (again, of either material or void). The penalty functions are analytically computed using the first and second derivatives of the level-set-function, which are computationally easy quantities to compute. Thus, their gradients can be computed relatively quickly compared to the electromagnetic gradients, and this gradient can be incorporated into the gradient-based optimization to minimize the fabrication penalty.

Comments on fabrication robustness

Fabrication techniques can be further characterized by the range of sizes that a feature may deviate from its nominal value. For example, over- or under-exposure of resist can cause dilation or erosion of the patterned feature. So even if a photonic device has been optimized such that all features have achievable sizes, it is not necessarily the case that the device will perform well when these feature sizes are deviated within their fabrication tolerance ranges. To address this, the device can be explicitly optimized for fabrication *robustness*.

For the work presented in this thesis, fabrication robustness was not considered because the tolerances for the microwave and THz devices measured during this thesis are deeply subwavelength. In general, it was found that optimizing a device to be non-resonant (either by optimizing for broadband performance, or by adding a small bandwidth to devices that are intended to be used at just a single frequency) was enough to ensure the device performed as expected in measurement. However, examples from literature also suggest that fabrication robustness can be optimized for by simultaneously optimizing eroded and dilated versions of the device [31]. For this technique, 3× as many simulations are required to optimize the device, which does not pose a substantial burden on optimization time if parallelized properly.

2.6 The complete optimization: logistics and parallelization

The optimization of a highly multifunctional photonic device is, above all else, a logistical and computational challenge. This section describes how optimizations can be organized to minimize the required number of simulations and parallelize computations as much as possible.

The basic unit of the optimization is a pairing of a forward source to an adjoint source. Since the adjoint-method combines the radiated fields from a pair of sources to compute the gradient of an associated FOM, each pairing is representative of a unique FOM and thus the term FOM is often used in this thesis to refer to such a pairing. To distill this unit down to its most basic form, we can consider the pair of sources to be monochromatic with a defined polarization. Thus, an example of a basic unit of the optimization may be a TE-polarized 532 nm planewave as the forward source, and a TE-polarized 532 nm dipole as the adjoint source. This pairing could be used to optimize a device that focuses the forward source to the point of the dipole, seeking to maximize a FOM defined as the intensity of the electric field at the location of the dipole source.

If we wish for the device to focus more wavelengths than just 532 nm, we can add more FoMs to the optimization that are evaluated at different wavelengths. To make the device focus TM-polarization as well as TE-polarization, we can similarly add more FoMs that pair off TM-polarized sources at the desired wavelengths. By considering more FoMs, we can begin to introduce multifunctionality to the optimization.

Weighting

A choice that must be made is how the results of these simulations are combined. Consider a device that we wish to focus both 532 nm light and 620 nm light to a single point. We can compute the gradient of each of the FoMs using the adjoint method, but we must then combine these gradients in some way to encourage both FoMs to be optimized. This can be as simple as averaging the two gradients, allowing the permittivity to step in a direction that hopefully increases both FoMs. The simplicity of this scheme can be problematic in certain cases, and the plausibility of these cases becomes more common for optimizations with a very large number of FoMs. As an example, consider a case of two FoMs where FOM 1 is performing well and still in the process of being optimized, but FOM 2 is performing so poorly that it is nearly at a local minimum. The gradient of FOM 2 is small compared to FOM 1, since it is near a local minimum. Thus, when the gradients are averaged the gradient of FOM 1 will be dominant, which causes FOM 1 to continue to optimize and FOM 2 to be practically ignored.

We can assert more control in the averaging scheme by performing a weighted average instead. The weights can be chosen based on the current performance of each FOM, such that poorly performing functionalities are given more weight than better performing functionalities (Eq. 2.24). The gradients are then combined into a single gradient (still evaluated at every point in the design region) using Eq. 2.25.

$$w_i = \max\left(\frac{2}{N} - \frac{FoM_i^2}{\sum FoM_j^2}, 0\right) \tag{2.24}$$

$$g(x, y) = \sum w_i \frac{\partial FoM_i}{\partial \epsilon} \tag{2.25}$$

In these equations, FoM represents the current value of a figure of merit, N is the total number of figures of merit, and w_i represents the weight applied to its respective merit function's gradient. The maximum operator is used to ensure the weights are never negative, thus ignoring the gradient of high-performing figures of merit rather than forcing the figure of merit to decrease. The $2/N$ factor is used to ensure all weights conveniently sum to 1 unless some weights were negative before the maximum operation. Although the maximum operator is non-differentiable, this function is used only to apply the gradient rather than to compute it. Therefore, it does not affect the applicability of the adjoint method.

Parallelization

In general both the individual electromagenetic simulations and the overall optimization procedure can be highly parallelized. Commercial Maxwell's solvers are generally capable of utilizing multi-core processors on a single computer, and often can utilize Message Passing Interface (MPI) to deploy large simulations across multi-node clusters. Increasing the number of processes used to conduct individual the simulations is one simple way to speed up optimizations. The other method to increase speed is to parallelize the many different simulations that are required to simulate multiple FoMs.

The basic concept of parallelizing the optimization is to, for each iteration, simulate all of the unique sources from all FoMs, then combine the results of all simulations in as parallel a manner as possible. In general every FOM represents a unique pair of sources, but FoMs may share one source. As an example, a power-splitter device would be optimized using two adjoint sources (representing the two output ports) that share an identical input source. Thus, simulating only the unique sources among the collection of all FOM sources is important.

A special case for parallelization occurs when the sources for different FoMs differ only in frequency, since these simulations can be batched very efficiency using time-domain techniques such as FDTD. FDTD simulations model the response of a system to a time-domain pulse, and since very short pulses contain a large range of frequency components FDTD simulation times tend to scale very well with increasing frequency points (it is primarily computation time and RAM requirements that increase with the number of frequency points being analyzed). Thus, simply increasing the frequency range over which a particular functionality operates does not add substantially to the optimization time as it does not increase the required number of FDTD simulations. However, it is important to note that the weighting scheme discussed in the previous section still requires FoMs of this type to be treated with distinct weights.

Remaining cases (such as simulating orthogonal polarizations, different input angles, or generally different amplitude/phase profiles) require independent simulations. Paying attention to simulate only the unique sources among all of the FoMs, these simulations can be deployed across a cluster. Once all simulations are complete, the gradient of each FOM must then be computed. This process can be parallelized by creating a job for each FOM that retrieves and combines the results its forward/adjoint source pair. Once this is done, the results from all FoMs are combined with a weighted average and the primary computational task of that iteration is complete. The remaining update procedure and enforcement of fabrication constraints typically need not be parallelized beyond the standard multi-core parallel computations employed by standard programming languages.

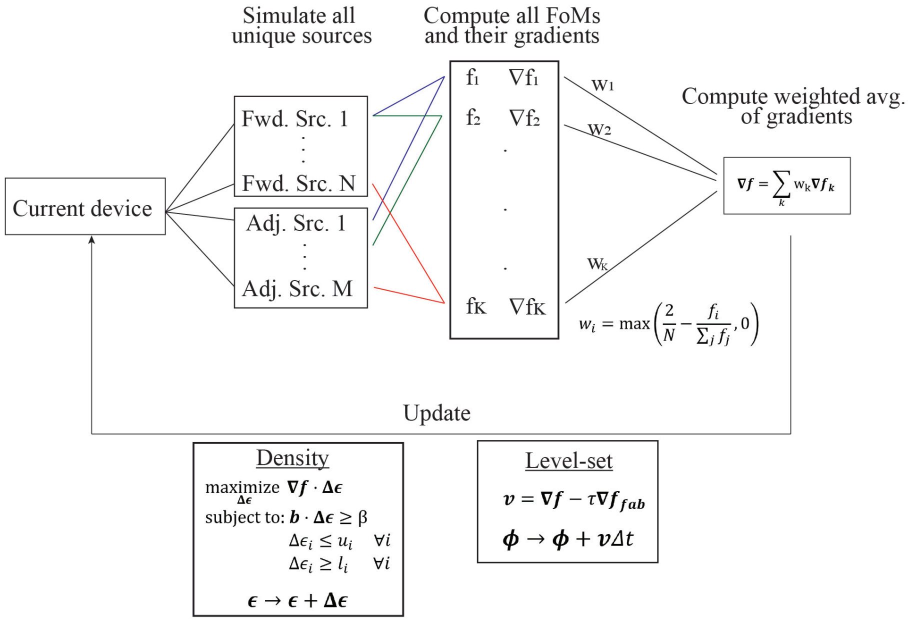

Figure 2.8: An overview of the design methodology.

The distinction between *distributed computing* and *concurrent computing* is important, particularly because both types of computing are used extensively in photonic design. Distributed computing describes how a large problem can be broken into smaller pieces that are solved on different nodes of a cluster. In the context of photonic simulation, distributed computing is typically used to solve a very large simulation that requires more RAM and compute than a typical single computer has. On the other hand, concurrent computing solves multiple problems in parallel by assigning each problem to their own node. Concurrent computing is typically used to solve parametric problems, such as sweeping a geometrical parameter of a geometry, where each node of the cluster is given the task of running a simulation for one parameter value.

For these iterative inverse design optimization, we are primarily interested in concurrent computing. Since many simulations need to be run each iteration, and since many iterations need to be performed to converge to a final device, the time required to run a single simulation is typically short. There is, however, a large volume of simulations that must be run each iteration (one for every forward and adjoint source), so it is optimal to distribute each simulation to its own node on a cluster.

This distinction between distributed computing and concurrent computing is particularly important when setting up jobs on a cluster. Oftentimes a cluster is set up for distributed computing by default, allowing all nodes on the cluster to communicate with one another. Unless specified, a photonic simulation deployed to the cluster may therefore utilize processes that are localized to different nodes on the cluster, which can create a bottleneck in speed as intercommunication between nodes tends to be slower than local communication on one node. It is faster to give each node exclusive access to a simulation, which only requires intercommunication between nodes to report the final results of the simulation.

Update procedure

The device permittivity is updated according to the weighted average of all gradients, computed by 2.25. This can be done using either a density-based or level-set update procedure. Every optimization conducted in this thesis utilized the sub-optimization technique described in Section 2.4. For the level-set update procedure, a level-set function based on signed-distance functions and velocity extensions was adapted from Ref. [29]. The ability to enforce minimum feature size constraints, described in greater detail in Section 2.5, involves analytically minimizing a fabrication penalty function based on gap-size and radius-of-curvature penalties, which was adapted from Ref. [30].

Chapter 3

MULTIFUNCTIONAL VOLUMETRIC METAOPTICS

The inverse design techniques described in Chapter 2 allow us to find optimal shapes of devices to perform arbitrary electromagnetic functions. The device shapes are subject to fabrication constraints, but otherwise able to assume arbitrary shapes. Numerous functions can be co-optimized to create devices that respond in some desired way to different wavelengths, polarizations, or incident angles. The work done during the course of this thesis has leveraged these capabilities to design fabricable 3D metaoptics that vastly expand on the multifunctional capabilities of metasurfaces, and this chapter is dedicated to describing these devices in detail.

The chapter begins with Section 3.1 introducing the work that first demonstrated these capabilities in a 3D platform, referred to as volumetric metaoptics. This work designed and experimentally verified devices capable of efficient 3-channel spectral demultiplexing red-green-blue (RGB) bins that exceed the theoretical efficiency of traditional absorptive Bayer filters. Furthermore, an additional capability of sorting orthogonal polarizations of the green bin to different locations is included. Given 2D metasurfaces are typically limited in bandwidth, the demonstration of both spectral demultiplexing across the entire visible spectrum with added polarization sorting capabilities far exceeds any studied behavior from 2D metasurfaces.

This work is then expanded upon by introducing designs of further increased multifunctional behavior. Section 3.2 studies various platforms in which the metaoptics device is mechanically reconfigured to perform different functionalities. For example, by simply translating alternating layers of an 8-layer device by only two center wavelengths, the devices can switch between one of three functionalities: a spectral demultiplexer, a broadband focuser, and a polarization sorter.

The device described in Section 3.3 expands on the multifunctional capabilities even further by designing a device that can classify properties related to the wavelength, polarization state, and incident angle without the need for mechanical reconfiguration. It can be thought of as a direct extension to the device in Section 3.1, except it is placed over a larger number of CMOS pixels in order to classify more properties of the incident light. Specifically, when excited with a monochromatic planewave the device sorts light to a combination of nine different locations in the focal plane array, where a 3×3 grid of CMOS pixels could be placed. The ratio of power transmission to two of the nine pixels is used to quantify the wavelength of the light, the same ratio applied to two different pixels is used to quantify the amplitude difference of two orthogonal polarization states, and the ratio of power transmissions to five of the nine pixels are used to identify the angle of incidence with a 5° cone of incidence. This device could be tiled to make an array, and could directly replace the microlens arrays in wavefront sensors (Shack-Hartmann sensors) and plenoptic sensors. Since the properties of light are determined by the ratio of power transmission to a small number of pixels rather than the spatial location of a focal point, the device can have an extremely compact form factor compared to the microlens arrays often employed in these systems. These devices can enable a new generation of computation imaging cameras that are extremely compact and capable of imaging all of the fundamental properties of incident light.

3.1 Introduction: A scattering-based Bayer filter for improved color imaging

This section primarily serves to introduce the concept of inverse-designed multifunctional metaoptics by summarizing the work from Ref. [32]. My contribution to this work involved conducting and analyzing the measurements (see Fig. 3.4), so I eleborate on these techniques beyond what is in the original scope of Ref. [32], bearing in mind that this measurement setup is used again for the work described in Ch. 4 for measuring THz metaoptics.

This work sought to design a spectral demultiplexer capable of replacing traditional absorptive Bayer filters that are used in digital color imaging, which are limited in maximum efficiency due to their reliance on absorption. CMOS pixels tend to have broad absorption spectra, and information regarding the color of the imaged light is lost after absorption. To enable color imaging, a tessellated array of absorptive filters is placed above the CMOS array (one common arrangement features one red pixel, one blue pixel, and two green pixels, which is shown in Fig. 3.1a). Considering a planewave incident on this array, at most 25% of the power of red and blue light (and 50% for green light) is sensed by the CMOS array. This efficiency limit can be overcome by replacing the absorptive filters with a low-loss dielectric device that scatters light to the different pixels, which is shown as the 3D structured filter array in Fig. 3.1a.

Metasurfaces have demonstrated multifunctional capabilities by controlling both phase and polarization of incident wavefronts, but obtaining broadband performance from metasurfaces remains difficult, particularly for multifunctional devices. Increasing the thickness of metasurfaces (which makes it no longer a *surface*, but rather a *volume* hence the name *volumetric metaoptics*) allows for more degrees of freedom, and thus more complex functionality without loss in efficiency. The challenge of finding the optimal refractive index for these devices is a computational difficult problem, but with the techniques described in detail in 2 it is a tractable one. Furthermore, these design approaches offer substantial versatility with the ability to optimize arbitrary electromagnetic functions. To demonstrate this, the devices designed in this work feature polarization control for green light, focusing x-polarized green light to one pixel and y-polarized green light to another. The intensity profiles of red (650 nm) and blue (450 nm) are shown in Fig. **3.1**b. These functionalities are polarization independent. The intensity profiles for x-polarized and y-polarized green light (532 nm) are shown in Fig. **3.1**c.

The device in Fig. **3.1**a is fully 3D, meaning its fabrication resolution is identical in all directions. These sorts of devices are possible to fabricate with 3D printing, even at the nanoscale with techniques such as two-photon direct-laser lithography (although this technique has a worse resolution along the axial direction of the lithography laser beam). However, obtaining the resolutions required for devices operating in the visible spectrum is difficult with current technology, and material choices are limited. A more suitable fabrication process for complex dielectric components, featuring a much wider range of low-loss dielectric materials, is sequential layered fabrication which yield a device comprised of layers that are typically thicker than the resolution obtainable with lithography. Thus, these layered devices are often referred to as 2.5D devices, rather than 3D. A typical process for creating such a structure involves depositing a thin layer of a dielectric, patterning it, backfilling the void features with a new dielectric, and repeating. It is often necessary to add a planarizing step as well, such that the next layer is deposited on a flat surface. A proposed device comprised of five $2\mu m \times 2\mu m$ layers, with each layer made of 400 nm thick SiO2/TiO2, is shown in Fig. **3.2**a, and a cross-section of one layer is shown in Fig. **3.2**b.

The 2.5D device performs the same function as the 3D device from **3.1**, and it is characterized in simulation by exciting it with a linearly polarized planewave at all frequencies (Fig. **3.2**c). The sorting efficiency, defined as the fraction of power that transmits to one of four $1\mu m$ square pixels in the focal plane, is shown for each pixel in Fig. **3.4**d. The results of approximately 60% for red and blue vastly outperform

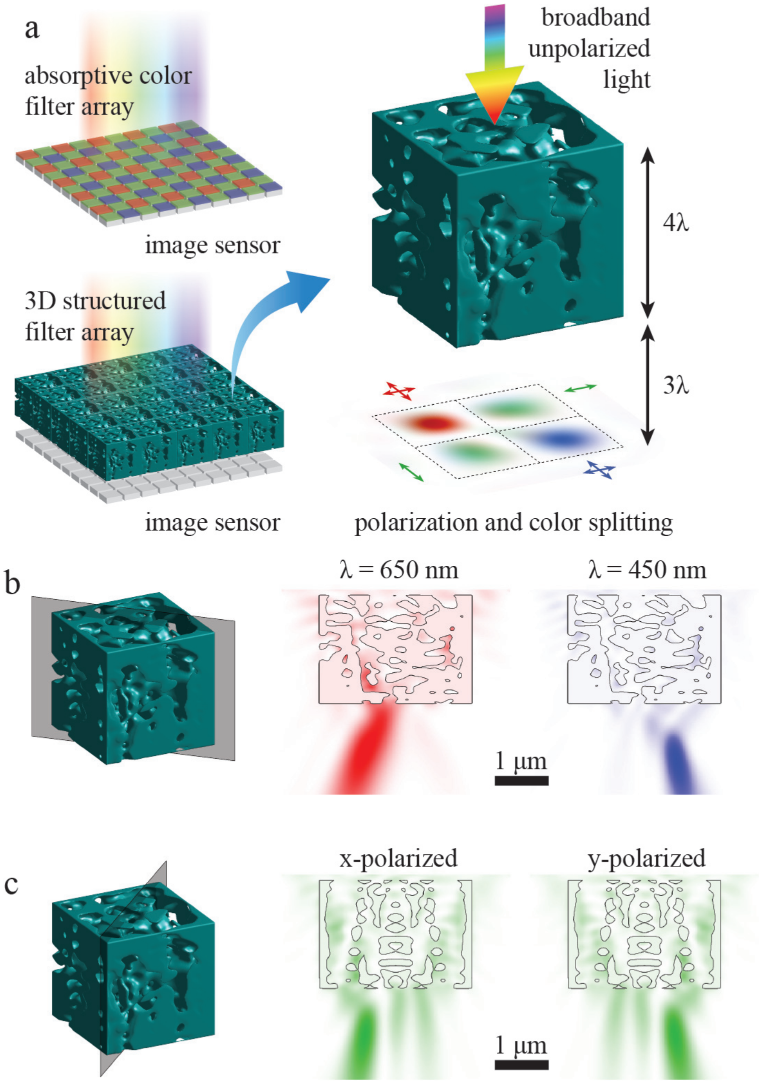

Figure 3.1: **Inverse-designed color and polarization sorter.** (a) Spectral filter consisting of a 3D nanostructured polymer cube. Incident light is efficiently focused into four sub-pixels on a focal plane below the device. (b) Red (600–700 nm) and blue (400–500 nm) spectral bands are sorted into opposite quadrants. The color indicates field intensity. The brightness of the output field intensity demonstrates the focusing nature of the devices. (c) Similarly, the green (500–600 nm) band is further split according to linear polarization. Red and blue quadrants are polarization independent.

traditional absorptive Bayer filters. The green performance of 50% also outperforms Bayer filters in practice, nearly matching their theoretical max sorting efficiency of 50%. The fourth pixel is dark since that corresponds to the y-polarization of green light, which is not excited here. Due to diagonal symmetry of the device, which was imposed during the optimization, the sorting efficiencies are identical for the y-polarization, except the illuminated green pixel is switched.

We investigate the trade-off between multifunctionality and device thickness by designing a series of devices with different numbers of layers. Each device follows the same design algorithm using 400 nm layers. Fig. 3.2e shows the sorting efficiency of each device, averaged across the visible spectrum. While the single-layer metasurface performs marginally better than empty space, efficiencies grow steadily with device thickness. In addition, thicker devices exhibit improved color and polarization contrast. Contrast is defined here as the difference in normalized power between the two brightest quadrants, and it therefore reflects the capacity to distinguish incident colors and polarizations. With five layers, the volumetric scattering element outperforms an absorptive filter with respect to sorting efficiency (57%), color contrast (29%), and polarization contrast (41%).

The chosen dimensions can be fabricated in a semiconductor cleanroom with significant investment in process development, but they are challenging to achieve in a university cleanroom with limited resources. Thus, we experimentally validate the fabrication-constrained design using a centimeter-scale analog (Fig. 3.3) operating at microwave frequencies in the Ka band (26– 40GHz). This device is designed to maximize the sorting efficiency at 18 wavelengths, equally spaced across the Ka band, and for both orthogonal linear polarizations. The device is constructed from a stack of 20 polypropylene (PP) sheets (index = 1.5), each 1.6 mm thick and patterned with 1 mm minimum feature size. A metallic boundary on the sides of the cube is incorporated to limit interference from the measurement apparatus. The microwave device occupies a 35 mm× 35 mm footprint, the same as its optical analog relative to the operating wavelength. The index contrast of the scattering element is also similar in both cases: $n_{TiO2}/n_{SiO2} = 1.67$, $n_{PP}/n_{air} = 1.5$. Fig. 3.3b shows the assembled device.

We characterize the performance of the spectral filter by measuring the complex microwave field scattered by the device when illuminated by a collimated Gaussian beam. We sample the scattered microwave near field in the Ka band using a WR-28 waveguide flange. An incident Gaussian beam (FWHM = 25 mm) is generated by

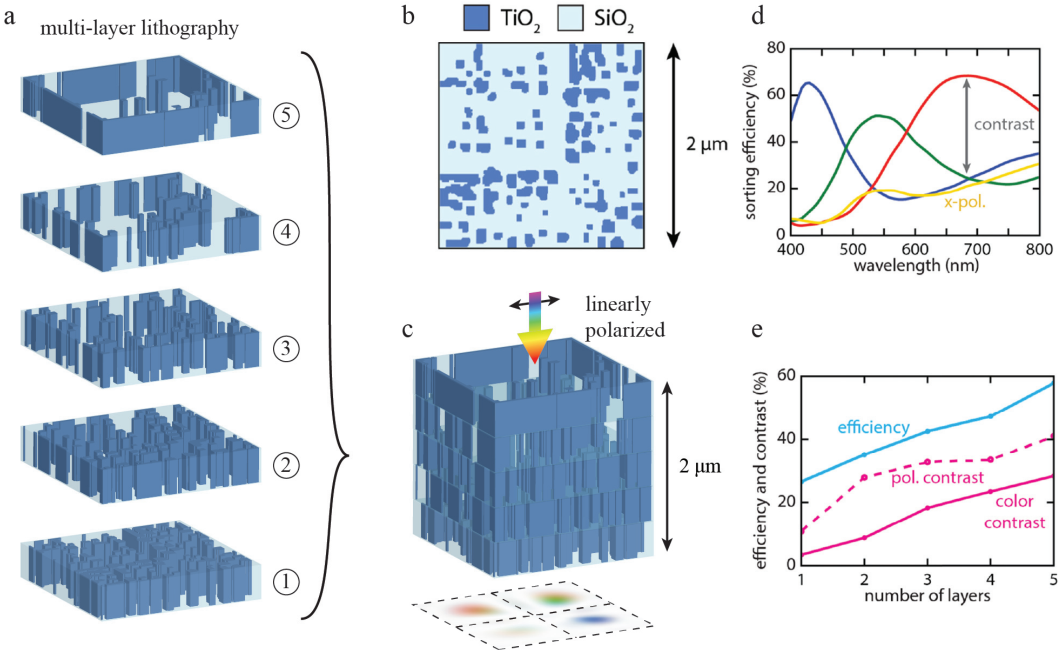

Figure 3.2: **Spectral filter designed for multilayer lithography.** (a) Three-dimensional scattering element constructed by stacking multiple two-dimensional layers. (b) Each layer consists of a series of TiO2 mesas, which can be fabricated using standard lithography and material deposition. Posts are infiltrated with an SiO2 matrix, forming a flat substrate for subsequent layers. (c) When combined, the five-layer stack performs the same function as the freeform design shown in Fig. 3.1. Linearly polarized light is focused into distinct sensor regions depending on frequency and polarization. (d) Sorting efficiency within the visible spectrum, defined as the fraction of incident power reaching the target quadrant as indicated in (c). The yellow curve indicates light present at the cross-polarized quadrant. (e) Spectrum-averaged sorting efficiency, polarization contrast, and color contrast for devices with increasing number of layers, each 400 nm thick.

a vector network analyzer coupled to free space via a microwave horn antenna and focusing mirror (Fig. S6). We sample the S_{21} scattering parameter at a measurement plane 62 mm beyond the output aperture of the device in order to recover the local electric field. By scanning the position of the probe antenna in the measurement plane, we can measure the local electric field profile. We apply a deconvolution filter to account for the anisotropy of the probe and the finite aperture of the device. Specifically, we deconvolve the measured results with the far-field radiation pattern of the waveguide flange, then computationally backpropagate the complex fields to the focal plane of the device, where they can be analyzed in the same way as the simulated device (see Fig. 3.4c-f). After conducting the procedure for one orientation of the device, the device is rotated 90° and the process is repeated to measure the response of the device to the orthogonal linear polarization. Finally, the measurement is repeated without any device present to measure the total power of the input beam. The output fields intensities of the the device are normalized to this quantity to compute efficiency values (see Fig. 3.4g,h).

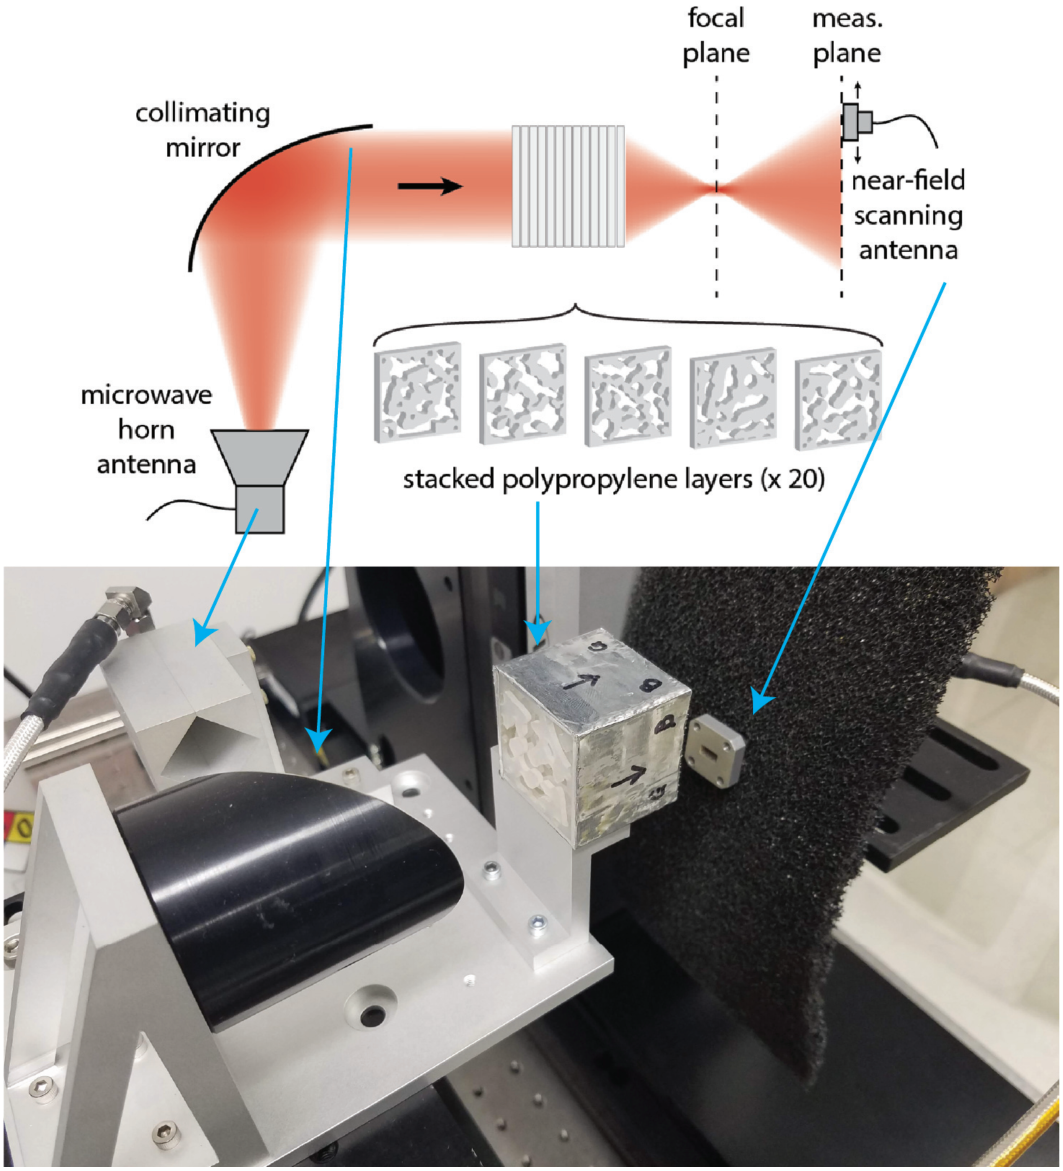

Figure 3.3: **Experimental characterization at microwave frequencies.** Microwave device consists of 20 patterned polypropylene sheets, assembled into a cube. The cube is illuminated by a collimated microwave source operating in the Ka band (26.5 – 40 GHz). Local electric fields are collected by a WR-28 waveguide flange (labelled near-field scanning antenna), and back-propagated to the focal plane. The black foam surrounding the probe is an absorbing foam used to reduce interference from reflections, and isolate the system from other microwave fields present in the laboratory.

Figs. 3.4c–3.4f show the simulated and measured intensity of the microwave fields at the focal plane of the device, summed over all frequencies within the measurement bandwidth. The color corresponds to the observed hue of the analogous optical fields when scaled to visible frequencies. These field profiles show efficient color and polarization sorting into each target quadrant, with good agreement between experiment and simulation. Figures 3(g) and 3(h) show the relative sorting efficiency of the spectral filter across the measurement spectrum. These efficiencies are defined as power transmitted through each target quadrant, normalized to the total power at the focal plane. We observe close agreement between experimental and simulated efficiencies, shown as solid and dashed curves, respectively. Each band shows efficient sorting with low crosstalk from out-of-band light, roughly 10%. The individual peaks in within each band correspond to the optimized frequencies. Sharp transitions between spectral bands highlight the improved color discrimination over typical dispersive scattering elements. We observe similar discrimination between orthogonal polarizations, shown in green and yellow curves. The remaining intensity is lost to reflection and oblique scattering outside of the target area.

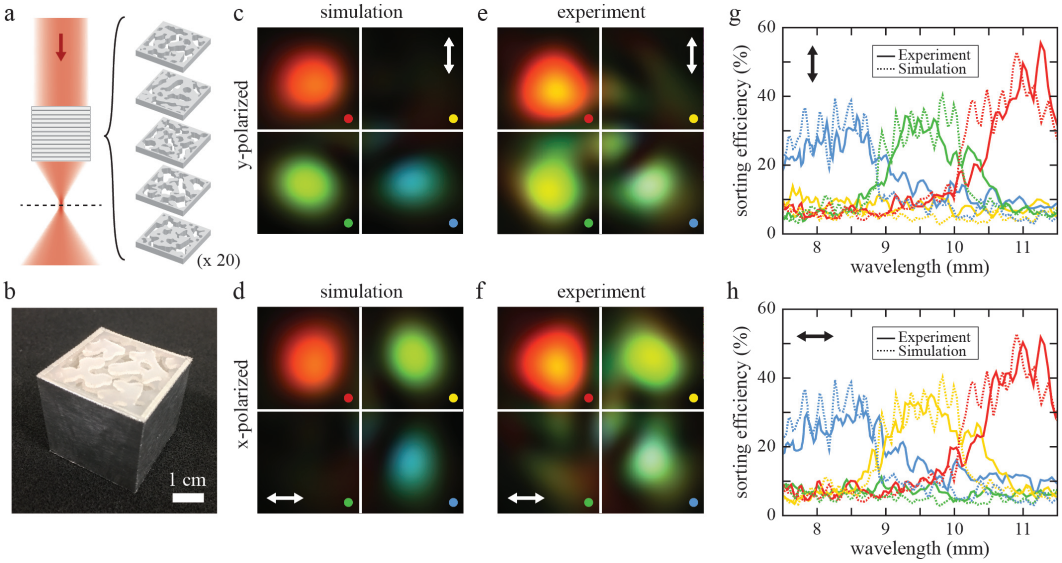

Figure 3.4: **Microwave measurement and simulation results.** (a) and (b) Microwave filter composed of patterned polypropylene sheets, assembled into a cube. The cube is illuminated by a collimated microwave source operating in the Ka band (26–40 GHz). (c) and (d) Field intensity at the focal plane of the microwave filter for linearly polarized illumination, as obtained by full wave simulation. Colors represent the equivalent hue of visible light, scaling the microwave frequency by a factor of ×17,500; (c) and (d) correspond to vertical and horizontal polarization, respectively. (e) and (f) Measured microwave field intensity at the focal plane. (g) and (h) Measured and simulated microwave sorting efficiency for each polarization. Dashed lines represent simulated efficiency, while solid lines represent experimental values. Each color corresponds to a quadrant as indicated by the dots in (c)–(f).

3.2 Mechanically reconfigurable metaoptics

Metasurfaces advanced the field of optics by reducing the thickness of optical components and merging multiple functionalities into a single layer device. However, this generally comes with a reduction in performance, especially for multi-functional and broadband applications. Three-dimensional metastructures can provide the necessary degrees of freedom for advanced applications, while maintaining minimal thickness. This work explores mechanically reconfigurable devices that perform focusing, spectral demultiplexing, and polarization sorting based on mechanical configuration. As proof of concept, a rotatable device, a device based on rotating squares, and a shearing-based device are designed with adjoint-based topology optimization, 3D-printed, and measured at microwave frequencies (7.6-11.6 GHz) in an anechoic chamber.

Introduction

Traditional optical design combines independent bulk elements to achieve complex functionality. Recent advances in nano-fabrication technologies enabled the miniaturization of bulk components by synthesizing multiple optical functionalities into single, subwavelength-thick layers called metasurfaces [7, 33]. However, the performance of metasurfaces is limited by the number of optical modes that can be controlled, which scales with the volume of the device and the maximum refractive index contrast [10, 34].

Systems of metasurfaces have been used to perform functions that are beyond the capabilities of a single metasurface. A recent novel platform for this technique involves fabricating two metasurfaces on opposite sides of a transparent wafer [35, 36]. A “folded-optics” platform uses reflective metasurfaces on both sides of a wafer to enable many interactions with different metasurfaces in a small, wafer-thick region [12]. This technique still uses a modular design approach, requiring the different metasurfaces be spatially separated to mitigate the unaccounted-for effects of multiple scattering events between elements that effectively invalidate the approximations commonly used in metasurface design [9].

To expand the functionality of metaoptics, adjoint-based topology optimization is capable of designing high-performing dielectric structures with nonintuitive index distributions and objective functions that are challenging to achieve with traditional methods [23, 37]. While much of the work has been applied to 2D platforms such as silicon photonics [25, 38–40] and flat optics [41–43], free-space 2.5D and 3D devices have also been explored recently [32, 44–46]. In a previous work, a passive 3D device was designed that functions as a RGB color splitter and polarizer, scaled up to microwave frequencies [32]. An interesting prospect for this type of device is to expand its functionality while minimizing any corresponding reduction in efficiency. For a passive device, multiple dependent functionalities will necessarily compete for efficiency. For example, the efficiency of demultiplexing spectral bins will be reduced if the device simultaneously sorts polarization. One way to maintain high efficiencies amidst multiple competing functionalities is to design active devices that can switch functionalities in time to perform the different tasks.

Previous work has explored the feasibility of coupling mechanical reconfiguration with metaoptics to enhance multi-functionality. Examples include Moiré-lens systems [47], MEMS-based Alvarez lenses [48, 49], and computational imaging with shiftable diffractive optics [50]. In this work, we extend this concept by using inverse design techniques and topology optimization to design devices capable of switching between optical functions through mechanical reconfiguration. We explore three different methods of reconfigurability: a rotatable device that switches between three-band spectral splitting and broadband focusing when rotated 180°; a device based on rotating squares that performs the same functions as the rotatable device; and a device based on shearing adjacent layers that can switch between three-band spectral splitting, broadband focusing, and polarization sorting.

Measurement setup and design methods

Although 3D devices can be fabricated at the micro-scale with methods such as two-photon direct laser writing and greyscale lithography, it remains challenging to fabricate fully 3D structures with the required features sizes and alignment accuracy for optical and near-infrared applications [51]. For this work, we test our ideas in the microwave range from 7.6 to 11.6 GHz where Fused Filament Fabrication 3D-printing with polylactide acid (PLA) can generate subwavelength feature sizes with minimal material absorption. The 42% fractional bandwidth of this frequency range is the same fractional bandwidth as the range of 450 nm to 690 nm light, which is nearly the full visible spectrum.

We use a homemade anechoic microwave chamber to measure the electric fields of the devices, shown in Fig. 3.5. The design and construction of this setup was inspired by the work in Ref. [52]. The measurement system consists of a parallel plate waveguide supporting only the fundamental TEM mode in the frequency range

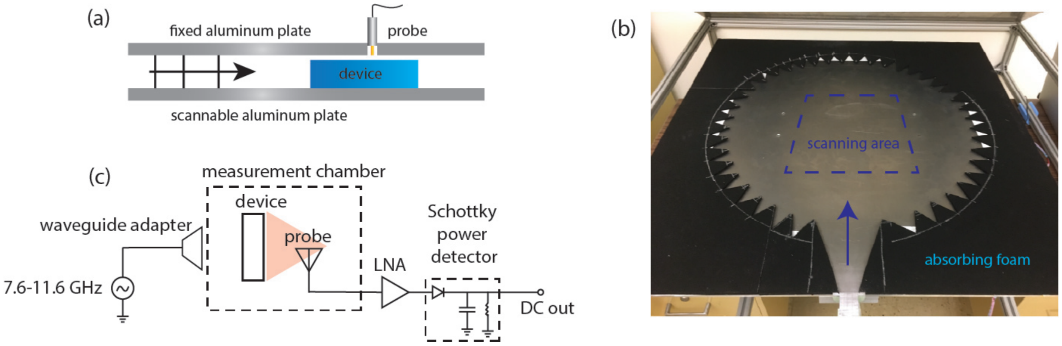

Figure 3.5: **The anechoic chamber used to measure the 3D printed devices.** (a) A sideview of the chamber. A TEM mode is injected into the anechoic chamber, exciting the device at approximately normal incidence. The fields in the cavity, including fields above the device, are measured with a small antenna probe extending through the top aluminum plate. The bottom aluminum plate is scanned in two directions to make a full 2D scan. (b) A picture of the anechoic chamber with the top aluminum plate removed. The scanning area is surrounded by absorbing foam. Triangles were cut out of the absorbing foam to reduce reflections. (c) A schematic representation of the measurement setup. The microwave source (Windfreak SynthHD) is scanned from 7.6 GHz to 11.6 GHz and is injected through a waveguide adapter. The signal from the measurement probe is first amplified with two low-noise amplifiers, and the signal is then passed through a Schottky power detector (Keysight 8470B). The output is proportional to the field amplitude in the cavity.

of interest when propagating through air or PLA. A sideview schematic of the system is shown in Fig. 1(a), and the constructed system with the top plate removed is shown in Fig. 1(b). A circuit view of the system is shown in Fig. 1(c). The output of a microwave synthesizer (Windfreak SynthHD) is injected into the chamber through a waveguide feed. This source is scanned from 7.6 GHz to 11.6 GHz, which is the designed bandwidth of operation for all presented devices. A small probe antenna (SMA Connector Receptacle) is extended through the top plate of the chamber to probe the fields inside the cavity. The probe antenna signal is amplified with low noise amplifiers and propagated through a Schottky diode detector (Keysight 8470B). The resulting signal is proportional to the amplitude of the fields in the cavity. The bottom plate is moved with stepper motors to make a 2D scan of the field amplitude inside the chamber. The probe antenna does not extend into the cavity, so the field intensity in the small air gap above the device can be measured. These fields are approximately proportional to the fields within the device material [52].

An important characteristic of this test setup is it measures 2D devices under TE-polarized illumination. The response of the device in the microwave chamber is nearly identical to the response of a device placed in a free space with the same index distribution, assuming the index distribution is extruded infinitely in the direction separating the waveguide plates [52]. This simplification from 3D to 2D implies we can optimize the device in a 2D coordinate system, thus drastically reducing the overall time of the optimization. Although these devices can be fully described by a 2D coordinate system, they are analogous to 3D free-space or 2D integrated photonics devices since patterning occurs in the direction of light propagation.

We design our devices assuming a non-dispersive refractive index of 1.5. While the actual refractive index of PLA is closer to 1.65 [53], for simplicity of measurement the devices only fill about 80% of the gap in the microwave test chamber. The height of the cavity is chosen such that only this fundamental mode propagates, even when the cavity is fully loaded with PLA, so it is sufficient to evaluate only the effective index of the fundamental mode. This can be done using eigenmode solvers, and in this case we use a mode source in Lumerical FDTD to compute the frequency-dependent effective index of the fundamental mode of the cavity. The effective index as a function of the percentage of the cavity filled with PLA is plotted in Fig. 3.6. Individual lines are plotted for frequencies ranging from 7.6 GHz to 11.6 GHz, which shows some dispersion. This dispersion was not considered when optimizing the devices.

There are some sources of error in our choice of refractive index which likely contributes to discrepancies between simulated and measured device performance, in addition to the simulated assumption that it is non-dispersive. One source of the uncertainty is predicted to originate from the layered nature of the 3D print which has previously been reported to cause a 7% anisotropy in PLA permittivity, with the permittivity ranging from 2.75 in the direction of normal to the printing surface and 2.96 in the direction parallel to the printing surface [53]. Furthermore, if the permittivity depends on the layered nature of the print, then it is likely affected by the 3D printer quality and layer height, which may be different from the references used to estimate the PLA permittivity.

Topology optimization

The three designs presented are all optimized with the same strategy. The optimization technique is similar to the techniques presented in Chapter 2, with some differences to account for the mechanical reconfiguration of each device, which will be briefly summarized here. The adjoint variable method is a technique to efficiently compute the gradient of a figure of merit with respect to a design region permit-

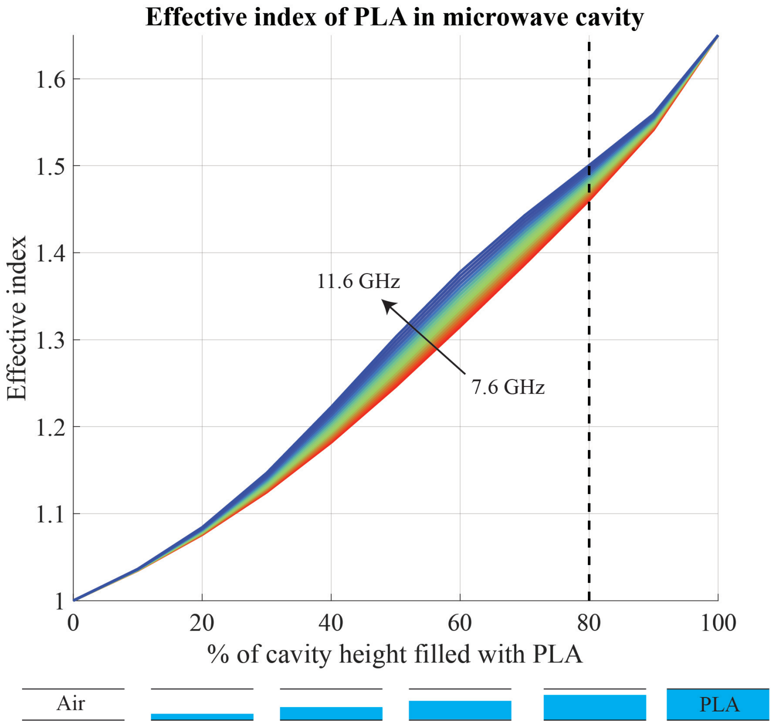

Figure 3.6: **Effective index of a PLA slab in the microwave cavity** The colored lines represent the frequency, with red at 7.6 GHz and blue at 11.6 GHz. The approximate height of the designed PLA devices is 80% the cavity height, which is shown with the dashed black line.

tivity by combining the results of just two simulations, a "forward" simulation and a backward, or "adjoint", simulation. The figure of merit used here, shown in Eq. 3.1, is the electric field intensity at the center of the desired pixel, denoted (x_0, y_0), so the adjoint simulation features a dipole excitation at this point. $\mathbf{E}$ denotes the electric field vector, and λ is the wavelength. We use a time-domain Maxwell's equations solver (Lumerical FDTD), so each simulation contains results for the entire bandwidth of interest.

$$FoM(\mathbf{E}, f) = |\mathbf{E}(x_0, y_0; f)|^2 \tag{3.1}$$

The optimization begins with each device modelled as a block of material with a permittivity between that of air and PLA. The gradient of every figure of merit is computed at every point in the design region for each of its potential configurations, and combined with a weighted average according to the discussion in Section 2.6.

At each iteration a sub-optimization is performed to find a permittivity step that maximizes the change in the overall device performance (the product of the gradient and the permittivity step) while constrained such that the overall binarization of the device, which is a measure of how close the device is to being comprised of either air or PLA, increases with each step. This sub-optimization has the effect of increasing binarization while maximally increasing, or minimally decreasing, the device performance and is described in detail in Chapter 2 Section 2.4. We conclude the optimization by forcing each pixel to be either PLA or air which has a negligible effect on the performance of the device since the device is nearly binary by the end of the optimization. Plots showing the convergence of the figure of merit and binarization of each device are shown in Fig. 3.7. Once the design is finalized, the resulting permittivity distribution is exported to an .STL file and 3D printed in PLA using an Ultimaker S3.

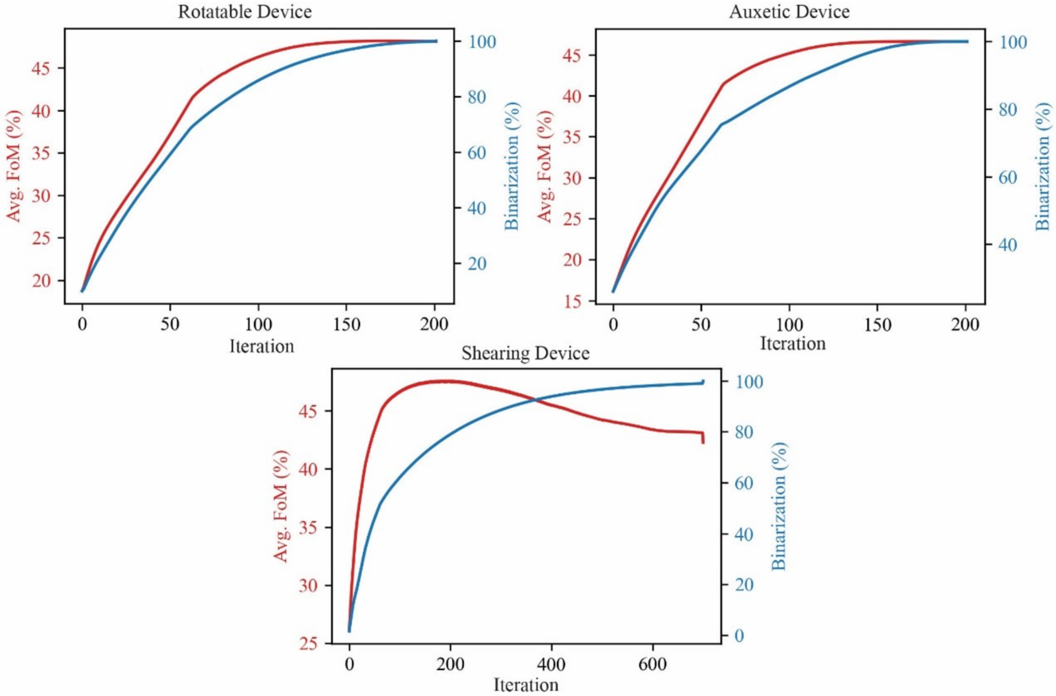

Figure 3.7: **Convergence plots of the different mechanically reconfigurable devices.** The red lines indicate the average figure of merit, and the blue lines represent the overall binarization of the device, On the final iteration, each device undergoes a thresholding operation that forces each voxel to its nearest material boundary. The performance is not substantially affected since the devices are nearly binary prior to the operation.

Results

Here we describe the results from the three different devices. Measurement results feature a 2D scan of the microwave cavity chamber for each configuration, with the red-green-blue colors representing the equivalent hue of visible light when the microwave frequencies are scaled by a factor of ×59,618. The microwave components have varying scattering parameters over the bandwidth of interest, a scan of the empty cavity is used to normalize device measurements to the total power injected into the cavity. Each device focuses to one or more of three separate pixels, depending on the desired functionality of the configuration. The pixels are depicted as either red, green, or blue, indicating the assigned color in the spectral splitting configuration with red as the lowest frequency bin (7.6-8.9 GHz), green as the middle frequency bin (8.9-10.2 GHz), and blue as the high frequency bin (10.2-11.6 GHz). These frequency bins are referred to simply as red, green, or blue from now on.

The focal plane is analyzed in simulation to determine the sorting efficiency, defined as the fraction of incident power transmitting through the target pixel. To quantitatively corroborate the measured intensity profile with the simulated intensity profile, the measured devices feature a series of plots that compare the normalized intensity integrated over the frequency bands of interest for each functionality.

Rotatable device

The first device presented here changes its function by undergoing a 180° rotation. The device is designed only for TE light. It focuses broadband light to the center pixel when illuminated by a normally-incident planewave from one side, and focuses red, green, and blue light to the respectively colored pixels when illuminated from the other side, as shown in Fig. 2(a-d). The footprint of the device is 6.2 cm × 18.6 cm, which is $2\lambda \times 6\lambda$ at the center wavelength of 3.1 cm, shown in Fig. 2(e). To ensure structural stability the device has a thin frame of PLA, and a connectivity constraint is enforced every ten iterations of the optimization so that the index distribution converges to a fully connected structure. Connectivity is enforced by first labelling disconnected regions, then connecting them with least-cost paths (i.e. a minimum spanning tree). The cost is considered the required permittivity change to flip a voxel from low to high permittivity. In this work, we reflect the permittivity about the midpoint of air and PLA. The figure of merit does not decrease substantially due to this operation according to the convergence plots in Fig. 3.7.

This device is uniquely simple among the devices presented, since a 180° rotation does not alter the scattering matrix beyond a transposition due to reciprocity. A reasonable concern is that this device will be limited in performance because of this. However, reciprocity does not strongly affect this device, since inputs from both sides are assumed to be normally incident. Thus, only scattering components mapping a normally incident input to a normally incident output are coupled by reciprocity, while the desired output fields are comprised of many more uncoupled planewave components. The sorting efficiency as a function of frequency is shown in Fig. 2(f).

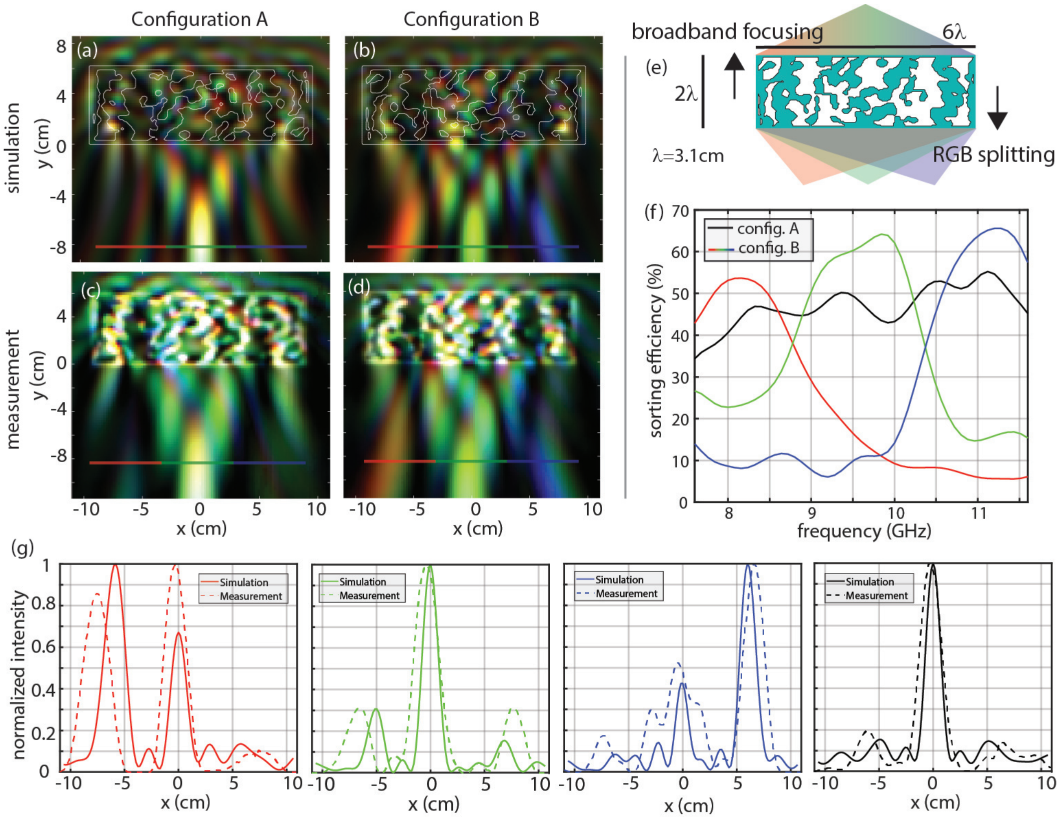

Figure 3.8: **Rotatable device performing broadband focusing in one configuration and spectral splitting in the other configuration.** (a) Broadband focusing simulation, (b) spectral splitting simulation, (c) broadband focusing measurement, and (d) spectral splitting measurement. (e) Schematic of the device, where blue represents PLA and white represents air. The footprint of the device is 6.2 cm × 18.6 cm. (f) Analysis of the simulated fields at the focal plane showing the sorting efficiency for each function, defined as the fraction of incident power reaching the target pixel. Red, green, and blue light are focused to their respectively colored pixels, and sorting efficiency is drawn in red, green, and blue, respectively. The broadband focusing function focuses all light to the middle pixel and is drawn in black. (g) Comparison of simulated and measurement normalized intensity profiles at the focal plane. Configuration B intensities are integrated over (left) 7.6 to 8.9 GHz, (center-left) 8.9 GHz to 10.2 GHz, (center-right) 10.2 to 11.6 GHz. (right) Configuration A intensity integrated from 7.6 to 11.6 GHz.

The sorting efficiency averaged across the relevant frequencies for each function is 50.6%. This outperforms traditional three-pixel absorptive Bayer filter arrays, which have a maximum theoretical sorting efficiency of 33% for each frequency band. To quantitatively compare the intensity profiles of measured and simulated results, normalized intensity profiles are shown in Fig. 2(g) for each function.

Rotating squares

The second device we studied is comprised of rotatable squares. The first configuration features a 0° rotation of all squares, while the second configuration features a ±90° rotation of each square with each square having an opposite angular rotation as its neighboring squares. This transformation can be achieved using the scheme depicted in Fig. 3(e), where the squares are hinged appropriately as described in previous works [54–56].

The two functions chosen here are broadband focusing and spectral splitting for TE polarization — the same as the rotatable device. The device is fully connected, with a frame around each square element to ensure structural stability.

A total of three devices were fabricated. The first device demonstrates the transformation. To fabricate this, the 3D-print was paused at half the thickness of the device, a nylon mesh was manually inserted, and the 3D-print was then resumed to completion. After the print, the nylon mesh was cut to leave only the required flexible hinges between each square. The other two devices are the two different configurations printed directly, which were then measured.

This design and rotatable design have the same device size and focal length, so the performance of the two devices can be directly compared. The simulated sorting efficiency is shown in Fig. 3(f). Since the current device has a frame around each square element, less of the design area is available for optimization. This detracts from the degrees of freedom of the device, which may explain why the device has a 48.0% average sorting efficiency, 2.6% less than the rotatable device. The difference in average efficiency may also arise from the different natures of the mechanical reconfiguration. The simulated and measured intensities within the full test chamber are shown in Fig. 3(a-d). To quantitatively compare the intensity profiles of measured and simulated results, normalized intensity profiles are shown in Fig. 3(g) for each function.

Shearing device

The final device features a transformation based on shearing alternate layers. This action is achievable at the microscale through MEMS electrostatic actuation. Unlike

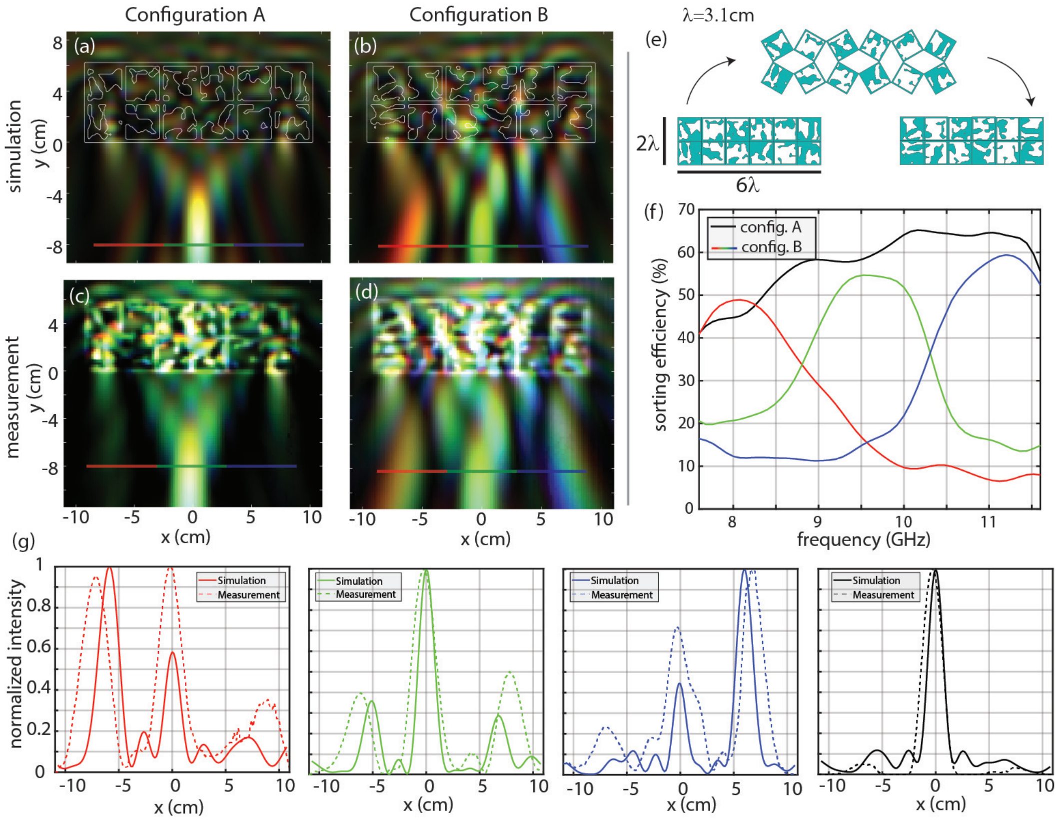

Figure 3.9: **The device performing broadband focusing in one configuration and spectral splitting in the other configuration.** (a) Broadband focusing simulation, (b) spectral splitting simulation, (c) broadband focusing measurement, and (d) spectral splitting measurement. (e) Demonstration of the transformation and device footprint, where blue represents PLA and white represents air. The footprint of the device is 6.2 cm × 18.6 cm. (f) Analysis of the focal plane of this device, showing the sorting efficiency for each function, defined as the fraction of incident power reaching the target pixel. Red, green, and blue light are focused to their respectively colored pixels, and sorting efficiency is drawn in red, green, and blue, respectively. The broadband focusing function focuses all light to the middle pixel and is drawn in black. (g) Comparison of simulated and measured normalized intensity profiles at the focal plane. Configuration B intensities are integrated over (left) 7.6 to 8.9 GHz, (center-left) 8.9 GHz to 10.2 GHz, (center-right) 10.2 to 11.6 GHz. (right) Configuration A intensity integrated from 7.6 to 11.6 GHz.

the previous devices, this device switches between three different functionalities: spectral splitting shown in Fig. 4(a), broadband focusing shown in Fig. 4(b), and polarization splitting shown in Fig. 4(c). All functionalities are designed for both TE and TM polarizations.

The added polarization control demands more degrees of freedom than the previous devices, and it was found that the device needed to be nearly 3× thicker than the previous devices to achieve satisfactory performance of 59.2% average sorting efficiency. This device does not have a supporting frame and does not enforce connectivity like the previous devices.

This device was analyzed only in simulation due to limitations within the measurement system: the measurement chamber only supports a TE-polarized TEM mode, and the scannable region is too small to measure this device. A 4-layer device was designed and tested with similar agreement between simulation and measurement to the previous devices, but the device could not achieve all objective functions. Data for this device is available on request.

The configurations and simulation results for the 8-layer simulated device are summarized in Fig. 4. The device is a stack of eight 8 cm × 2 cm layers. The mechanical actuation displaces adjacent layers in equal and opposite directions by 3 cm, which is approximately one wavelength. The first configuration, shown in Fig. 4(a), shows spectral splitting behavior with 42% average sorting efficiency for TE and 40% average sorting efficiency for TM. The crosstalk between the different spectral bins is worse than the previous devices, with the worst case occurring for the TM-polarized green input which focuses only 1.2× more power to the desired green pixel than the undesired blue and red pixels.

The neutral position of the device, in which all layers are aligned as shown in Fig. 4(b), features broadband focusing with high efficiency. The aperture size of this configuration is smaller than the other two configurations, and the sorting efficiency is normalized to the power incident on this smaller aperture when analyzing this configuration. The sorting efficiency shown in Fig. 4(e) is expected to be uniformly high across all frequencies and both polarizations.

The final configuration sorts TE polarization to the leftmost pixel and TM polarization to the center pixel. The transmission through each pixel, averaged across the entire spectrum, is summarized in matrix form in Fig. 4(f): a TE input focuses 47% power to the correct pixel and 28% power to the incorrect pixel, the TM input focuses 62% power to the correct pixel and 24% power to the incorrect pixel. More power is coupled to the desired pixel than the undesired pixel at all frequencies except for the case of TE input frequencies below approximately 9 GHz. In this case more power is directed towards the incorrect green pixel than the correct red pixel. It is possible that this could be fixed by sacrificing performance in the other functionalities, such as by tuning the weighting scheme described in Eq. 2.24, or by increasing the thickness or index contrast of the device.

Discussion

We have studied dielectric scattering elements capable of substantially expanding their functionality through mechanical reconfiguration. The work here is a step towards answering an important question in optics —- how much performance and functionality can be squeezed into a certain volume? The miniaturization of optical systems has led to new applications in lasers [8, 57], biomedical optics [58, 59], space instrumentation [60], and generally applications where strict size requirements exist. Yet there is still much progress that can be made, since systems of cascaded metasurfaces still generally require some free-space propagation to preserve their independence and the validity of the techniques often employed in metasurface design. Furthermore, complex functionality such as combining polarization control, spectral splitting, and imaging into single elements has been elusive.

The approach shown here showcases the advantages of coupling mechanical reconfiguration with adjoint-based electromagnetic optimization. Miniaturization is achieved by employing adjoint-based inverse design to find a locally optimal shape of a small dielectric volume, while broad multi-functional performance is achieved through mechanical reconfiguration. The designed structures are difficult to fabricate at the nanoscale for optical applications, but the design process is flexible enough to incorporate nanoscale fabrication requirements as shown in other works [61]. In particular, some areas of fabrication that can fabricate the required subwavelength features for these devices include: multilayer fabrication common in CMOS or MEMS processes [62]; direct-write two-photon laser lithography capable of designing subwavelength 3D elements in the infrared [51]; and closely aligned stacks of Silicon wafers with subwavelength features at terahertz frequencies [63, 64]. All of these techniques have demonstrated active mechanical control that could be useful for multiplexing functions like what was demonstrated in this work.

The devices here were studied at microwave frequencies for ease of fabrication,

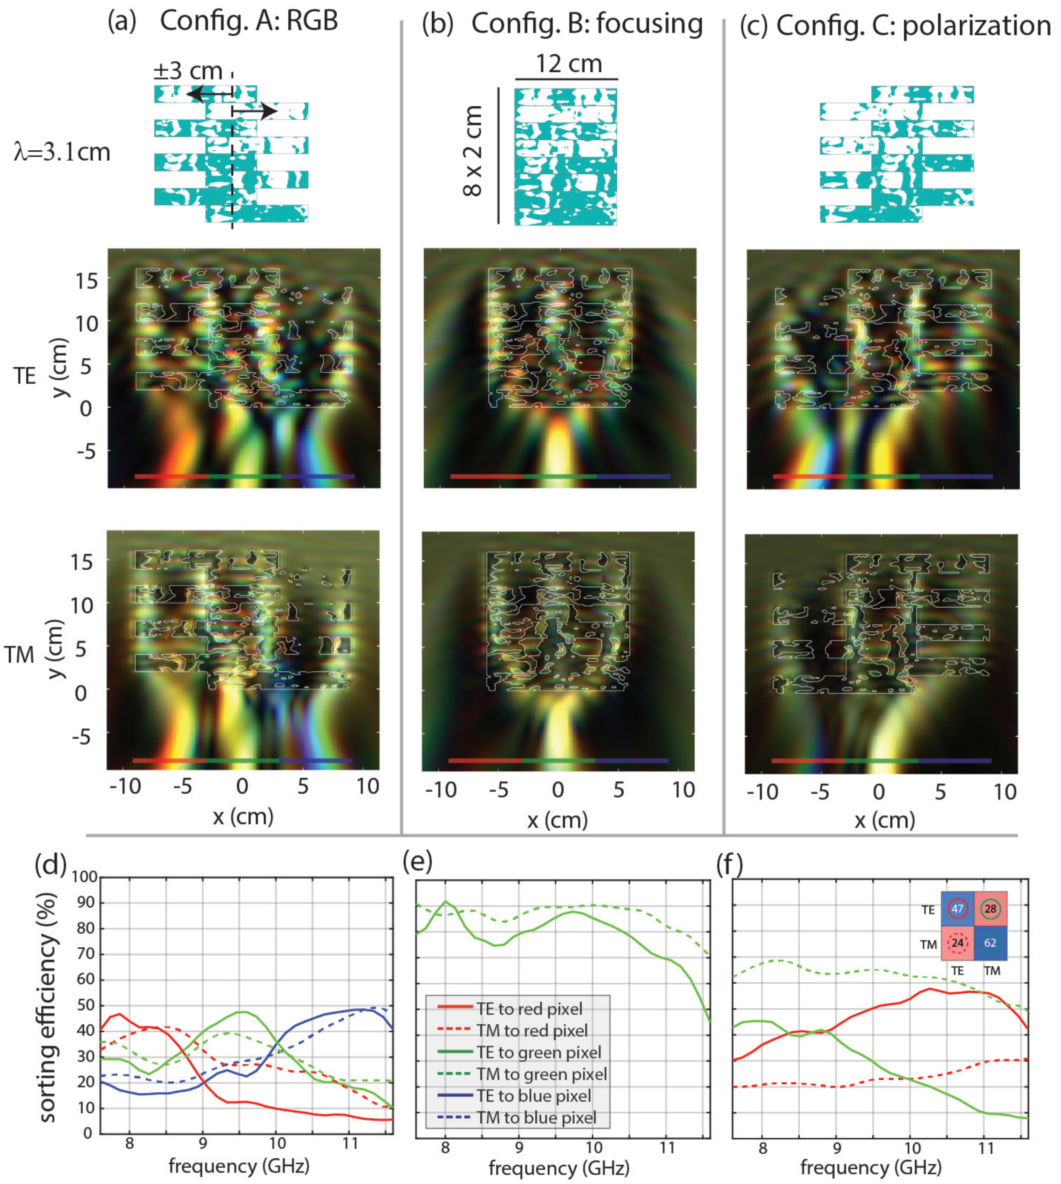

Figure 3.10: **TE and TM fields for a device based on a net shearing movement of 6 cm.** (a) The spectral splitting configuration splits red, green, and blue light to the red, green, and blue pixels, respectively. This is analyzed for both TE and TM polarizations. (b) The neutral state of the device performs broadband focusing to the center pixel for both TE and TM polarizations. (c) The polarization sorting configuration sorts broadband TE light to the left (red) pixel and focuses broadband TM light to the center (green) pixel. (d-f) Sorting efficiency of the different configurations. The line color represents the sorting efficiency to the similarly colored pixel as depicted in the color plots. Solid lines represent the TE response, and dashed lines represent the TM response. (d) Sorting efficiency in the spectral splitting configuration. (e) Sorting efficiency for the broadband focusing configuration. (f) Sorting efficiency for the polarization splitting configuration. (Inset) A confusion matrix representation of the sorting efficiency, with true input on the vertical axis and predicted input on the horizontal axis. Each matrix entry is determined by averaging the relevant trace over the full bandwidth.

yet the applications are particularly useful for optics. The scalability of Maxwell's equations ensures that the simulations used to compute the gradient are frequency agnostic. However, the design may need to be changed to account for different deformation and actuation mechanisms available for microscopic devices. For example, the device based on a shearing actuation presented here is analogous to electrostatically actuated devices seen in MEMS, yet considerations such as the pull-in effect would need to be carefully considered for such a device operating at the microscale [65, 66]. Some of these challenges have been addressed in recent works that combine metaoptics with MEMS [49, 67, 68].

Adjoint-optimization is well-suited for optimizing the optical properties of mechanically reconfigurable devices. In this work, the mechanical reconfiguration scheme was predetermined, and the material was patterned to enable multi-functional performance. The performance of the rotatable device and the device based on rotating squares were comparable, suggesting that the specific choice of mechanical transformation may not drastically affect optical performance. This offers the advantage that a certain transformation can be chosen for mechanical purposes and can be subsequently patterned to imbue the device with desirable optical properties. Alternatively, the mechanical and optical properties could be co-optimized. Future work could combine existing optimization techniques for designing mechanical metamaterials[69, 70] with the techniques presented here, yielding devices that have co-optimized mechanical and optical performance.

3.3 Volumetric metaoptics for simultaneous imaging, spectrum demultiplixing, and polarization analysis

For this work, we target a platform in which a periodic array of the designed 3D devices can be placed atop a CMOS imaging array, which allows for the small 3D devices to be tiled to cover as large of an area as desired. The devices are designed for three functionalities: polarization splitting of two linearly polarized states, wavelength splitting of two distinct wavelengths, and a basic form of imaging in which five planewave components are focused to five different locations. This totals to nine states we seek to classify, so in practice the device can be placed above a 3×3 grid of CMOS pixels similar to absorptive arrays used in color imaging.

The polarization and wavelength splitting functionalities are designed to work for all of the assumed incident angles, which include a normally incident case and non-normal incidences that are 5 degrees tilted from normal (polar angle $\theta = 5°$) with azimuthal angles $\phi \in \{0°, 90°, 180°, 270°\}$. The two design wavelengths are 532 nm and 620 nm. To prevent the structure from being highly resonant, frequencies within a ±5 THz range (approximately ±4 nm) of these design frequencies are also optimized. The two polarizations are linear and orthogonal, with one polarized in the xz-plane and the other in the yz-plane. The device is a 3µm × 3µm × 4µm stack of 20 layers. Each 200 nm layer is comprised of titanium dioxide (TiO2, n=2.4) and silicon dioxide (SiO2, n=1.5), with a minimum feature size of 50 nm. The background material is assumed to be SiO_2. A schematic of the device, CMOS sensor, and examples of the fields at the focal plane under a 620 nm xz-polarized planewave excitation angled at $(\theta, \phi) = (5°, 90°)$ is shown in Fig. 3.11. Note that since this 5° polar angle is in SiO_2 with $n = 1.5$, these planewave components couple to 7.5° planewave components in air.

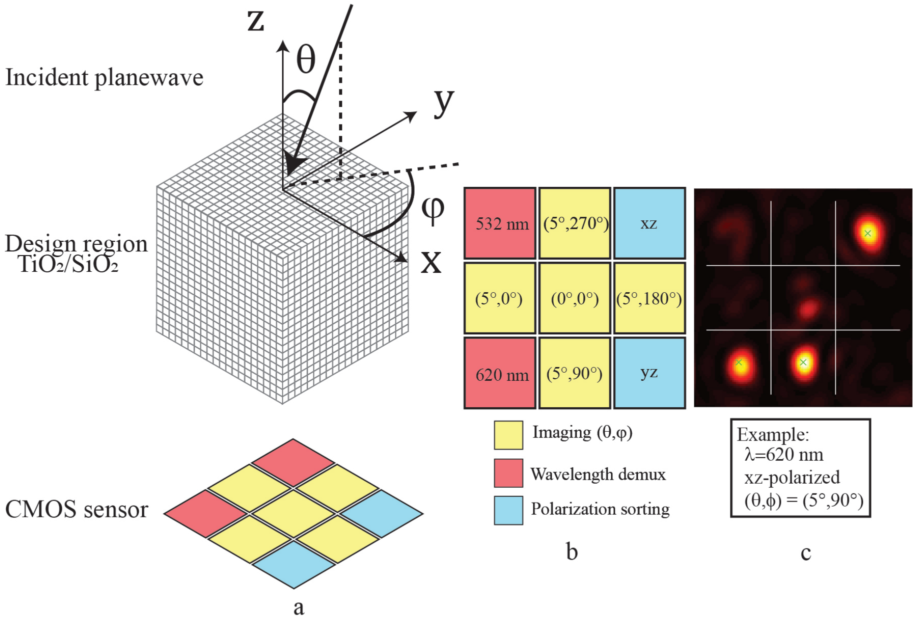

Figure 3.11: **The layout of the system, which features a device region placed above an array of CMOS pixels.** (a) An incident planewave is input to a device comprised of SiO_2 and TiO_2 at an angle (θ, ϕ). The bottom of the device is 1.5 μm above a CMOS array. The background material is assumed to be SiO_2. (b) The distribution of functionalities across the pixel array. The device focuses light to the different pixels depending on the state of the input light. (c) The output of 620 nm, xz-polarized planewave input incident at an angle $(\theta, \phi) = (5°, 90°)$

This section begins by describing the inverse design process used to optimize the devices. Next, the performance of the device under the assumed input conditions are studied. We then observe and comment on two important properties of the optimized device. We first observe that the independence of functionalities is preserved and predictable when excited at states in between the states the device was optimized for. For example, as the input light wavelength is continuously shifted from 532 nm to 620 nm, the ratio of the 620 nm pixel transmission to the 532 nm pixel transmission monotonically increases, while the imaging and polarization functions remain efficient. This occurs despite the device only being optimized in a narrow wavelength range around 532 nm and 620 nm. Similar behavior is observed in the imaging functionality when excited at input angles that were not explicitly optimized within a 5° cone. The read-out from these pixels can thus be used to infer the state of the light incident on the device so long as the input is assumed to be a monochromatic planewave. We find that this behavior does not depend on the chosen mapping of functionality to specific pixels in the focal plane by analyzing the behavior of devices optimized for 20 unique pixel distributions.

Methods

The goal of photonic topology optimization is to find a refractive index distribution that maximizes an electromagnetic figure-of-merit (FoM). Since the device presented here is highly multi-functional, its optimization is multi-objective. Each objective is a mapping of an input to a FoM: the input is a planewave with a specific angle, polarization, and wavelength; the FoM is the power transmission through the desired pixel. The general procedure for this optimization is shown in Fig. 2.8. It consists of three main steps: first, the FoMs and their associated gradients are computed [23]; second, the gradients are combined with a weighted average [27]; third, the device is updated in accordance with the averaged gradient using either a density-based optimization of a continuous permittivity or a level-set optimization of a discrete permittivity [30].

The individual FoM gradients are evaluated at every point in the design region using the adjoint method, which entails combining the electric fields in the design region for a "forward" and an "adjoint" simulation to compute the desired gradient. In this case, the forward case simulates the device under the assumed planewave excitation, and the adjoint case simulates a dipole (with a particular phase and amplitude based on the forward simulation) placed at the center of the desired pixel. Choosing a dipole as an adjoint source means that the device is optimized to focus light to the point the dipole is located at. However, we record the performance of the device as power transmission through the desired pixel rather than intensity at a point, since power transmission better represents the signal a CMOS sensor pixel would record. The remainder of this section elaborates on this procedure.

All unique forward and adjoint simulations are first simulated in parallel on a high-performance cluster. Next, the associated FoMs are recorded from the forward simulation results, and the FoM gradients are computed by combining the results of appropriate forward and adjoint source pairs. The individual gradients are spatially averaged in the z-direction for each layer, yielding a 2D gradient for each layer of the device. All FoM gradients are combined using a weighted average into a single gradient, still evaluated at every point in the design region. Information on this weighting procedure is described in detail in Section 2.5. This weighted gradient is then interpolated from the simulation mesh grid to the grid on which the geometry is defined, which we refer to here as the geometric grid. Interpolating on this geometric grid decouples the geometry from the specific choice of simulation accuracy while forcing the design to obey a rough minimum feature size constraint, which in this case is 50 nm. Finally, this interpolated gradient is used to update all voxels in the device structure. The optimization is done in two phases: a density-based phase, and a level-set phase. Each phase has a unique update procedure.

In the density-based optimization the permittivity of the device is modelled as a grid of greyscale permittivity values between the permittivity of the two material boundaries. This permittivity representation is effectively fictitious (unless an effective index material can be reliably fabricated), and the goal is to converge to a device that performs well and is comprised of only two materials which we refer to as binary. In past works density-based techniques have been prone to converging to non-binary solutions. An improved method is required that forbids voxels from being non-binary, and that can force the design to optimally converge to a binary solution. We present such a method in Ref. [71].

While the density-based optimization can converge to a fully binary solution, it is faster to terminate the optimization early by forcing each device voxel to its nearest material boundary. This thresholding step reduces the device performance, which we recover by further optimizing the device with a level-set optimization. Level-set optimization models the device boundaries as the zero-level contour of a level-set function ($\phi(x, y) = 0$), and thus benefits from describing inherently binary structures [29]. Empirically the final device performance is dependent on initial seed, hence the need for the improved density-based optimization that converges to a near-binary solution. Here we use level-set techniques to simultaneously optimize device performance and ensure the final device obeys fabrication constraints. This technique is described in detail in Ref. [30].

The convergence plots showing the average device FOM and device binarization are shown in Fig. 3.12. At various points the binarization is forced to increase by passing the permittivity through a sigmoidal function and changing the device permittivity to the output of this function. The FOM is then allowed to recover before repeating this discrete push in binarization. The device loses performance as the device becomes more binary. At iteration 512 the device is fully binarized and the optimization switches to a level-set optimization, which recovers some of the performance that was lost during the final phases of the density optimization. During the level-set optimization, the binarization is recorded as approximately 90% because the boundaries of the device are smoothed out yielding a continuous permittivity value. This can intuitively be thought of as the level-set function passing through a simulation mesh voxel, which is modelled in FDTD as a dielectric volume average of the two materials. Therefore the density optimization is overly restrictive at binarization values above 90%, since it does not allow for this type of border smoothing, instead modelling the device as a discretized grid of binary voxels. This explains why the level-set function is able to recover substantially more performance than what is lost at the initial conversion from density-based optimization to a level-set optimization, recovering to approximately the same average FOM value as the density-based optimization when it was 90% binary.

Results

We quantify the performance of the device in two ways. First, we check the performance of the device under the excitation beams that were assumed when optimizing the device, which we refer to as *training modes*. Second, we study the device at different input angles, wavelengths, and polarization states to analyze its ability to classify states. We refer to these as *validation modes*.

The design methodology co-optimizes 20 different input excitations featuring unique combinations of wavelength, polarization, and angle of incidence. The transmission to each pixel for each input state is shown in Fig. 3.13. For each state the transmission to the three correct pixels, the transmission to the six incorrect pixels, and the transmission elsewhere (oblique scattering, back-scattering, etc.) is shown. On average, the transmission to the correct pixels totals 47.7%, the transmission to incorrect pixels totals 16.8%, and the transmission elsewhere totals 35.5%.

An important metric is the contrast between pixels of the same function because the measured contrast of the pixels can be used to classify the input state. The

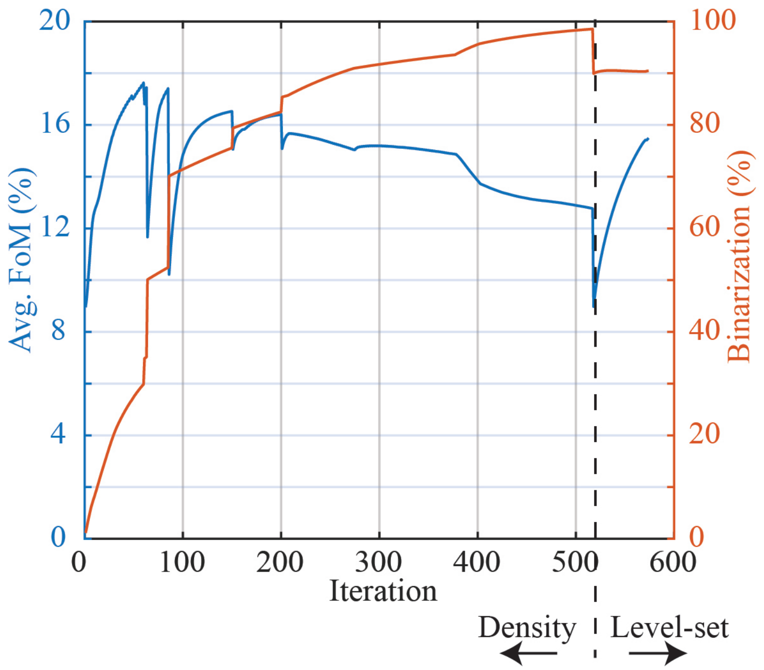

Figure 3.12: **Convergence plots of the 3D metaoptics device.** The initial density-based optimization features numerous points at which the binarization of the device is forced to increase by passing the current permittivity through a sigmoidal function. At iteration 512 the optimization is converted to a level-set optimization, which recovers the performance lost from that transition. Furthermore, the level-set optimization improves the performance of the device to approximately the same point as when the density-based optimization was 90% binary, which is the approximate binarization of the when modelled by a level-set curve due to the continuous smoothing of the permittivity at simulation voxel that are intersected by the level-set curve.

average contrast for the various imaging, wavelength, and polarization functionalities is 5.9, 6.2, and 11.2, respectively (equivalently 7.7 dB, 7.9 dB, and 10.5 dB). Early on during the optimization the polarization functionality was exhibiting poor contrast, so we chose to introduce an additional figure of merit that sought to minimize the transmission to the undesired polarization pixel, and this may be why the optimization ultimately converged to a solution where the polarization contrast is superior to other functionalities.

The design process optimizes the performance of the device only for the training modes, and it is unclear what the output of the device will be if the excitation parameters are continuously varied. Here we quantify the performance of the device at input states between that of the designed input states. We study the performance at wavelengths in between 532 nm and 620 nm, at a large set of angles within a 10° input angle cone, and at arbitrary polarization states.

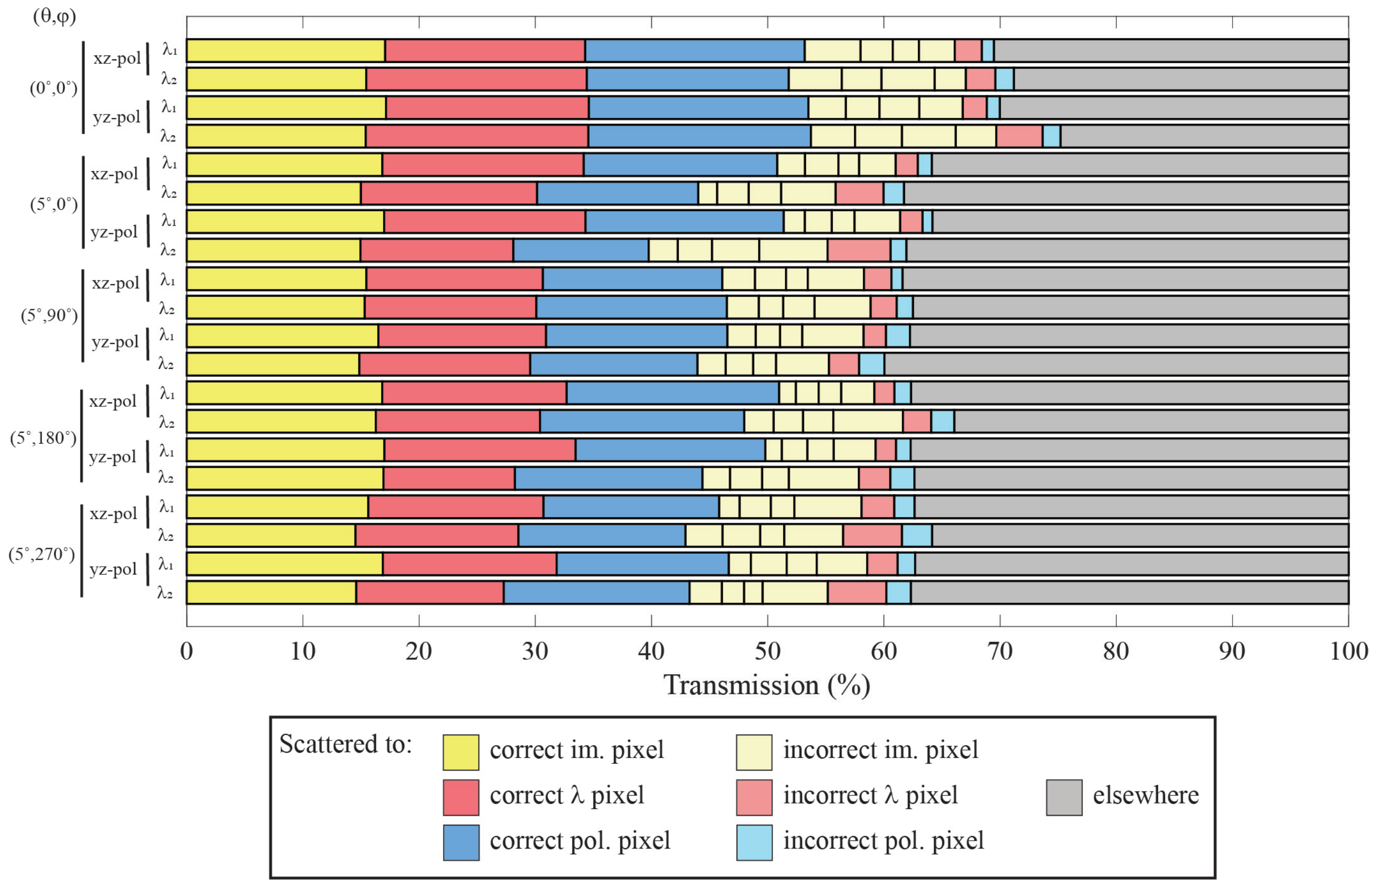

Figure 3.13: The performance of the device under excitation of training modes. The input angle, polarization state, and wavelength are represented by the y-axis. The x-axis represents the fraction of power scattered to different regions. The bold colored bars quantify transmission to the correct pixel, while the pale colored bars quantify transmission to the incorrect pixel. The grey bar quantifies the power that is back-scattered or obliquely scattered, thus not reaching any of the pixels in the focal plane.

As the wavelength shifts from 532 nm to 620 nm, the transmission through the 532 nm pixel smoothly decreases, and the transmission through the 620 nm pixel smoothly increases. The imaging functionality varies smoothly with respect to incident angles (θ, ϕ), but in a less intuitive way. As an example, when $(\theta, \phi) = (2.5°, 45°)$, the light is primarily distributed across three of the imaging pixels: the $(0°, 0°)$, $(5°, 0°)$, $(5°, 90°)$ pixels. Furthermore, in this case 532 nm is not focused to the 620 nm pixel which is the pixel that the incident planewave is oriented towards and as such is where light would be focused if the device were replaced with a microlens. Instead, the wavelength sorting and polarization sorting functionalities are mostly preserved under all excitation angles within an acceptance cone. Videos illustrating some examples as these parameters are swept are available in the CaltechDATA repository [72].

The full angular response of each individual pixel in a 10° cone is shown in Fig. 3.14a-i. The wavelength and polarization demultiplexing functionalities show a more uniform response than the imaging functionality, which means these functions are not highly dependent on incident angle. The imaging functionality is, by design, sensitive to the incident angle. To obtain these plots, the transmission values are averaged across wavelength, polarization, or both depending on the functionality: for the wavelength sorting functionality the values are averaged across polarization; for the polarization sorting functionality the values are averaged across wavelength; and for the imaging functionality the values are averaged across both wavelength and polarization.

Fig. 3.14j shows 1D traces of the transmissions, plotted across the horizontal dashed lines of Fig. 3.14a-i. The red line corresponds to the 620 nm pixel transmission, the green light to the 532 nm pixel transmission, and the blue solid and dashed lines to xz- and yz-polarized light, respectively. The solid yellow lines reference the imaging pixels with a high signal, which in this case are the ones in Fig. 3.14d-f. The yellow dashed lines refer to Fig. 3.14b,h in this case. Fig. 3.14k plots the transmission along the vertical dashed lines of Fig. 3.14a-i. Here the solid yellow lines refer to the imaging pixels in 3.14b,e,h, while the dashed yellow lines refer to 3.14d,f.

Fig. 3.14l shows the polarization and wavelength sorting functionalities as a function of wavelength. Each transmission trace is averaged over all simulated $\phi \in [0, 360°]$ and $\theta \in [0, 5°]$. As expected, the transmission to the 532 nm pixel is maximized for 532 nm wavelengths, and likewise for the 620 nm pixel. The traces smoothly vary

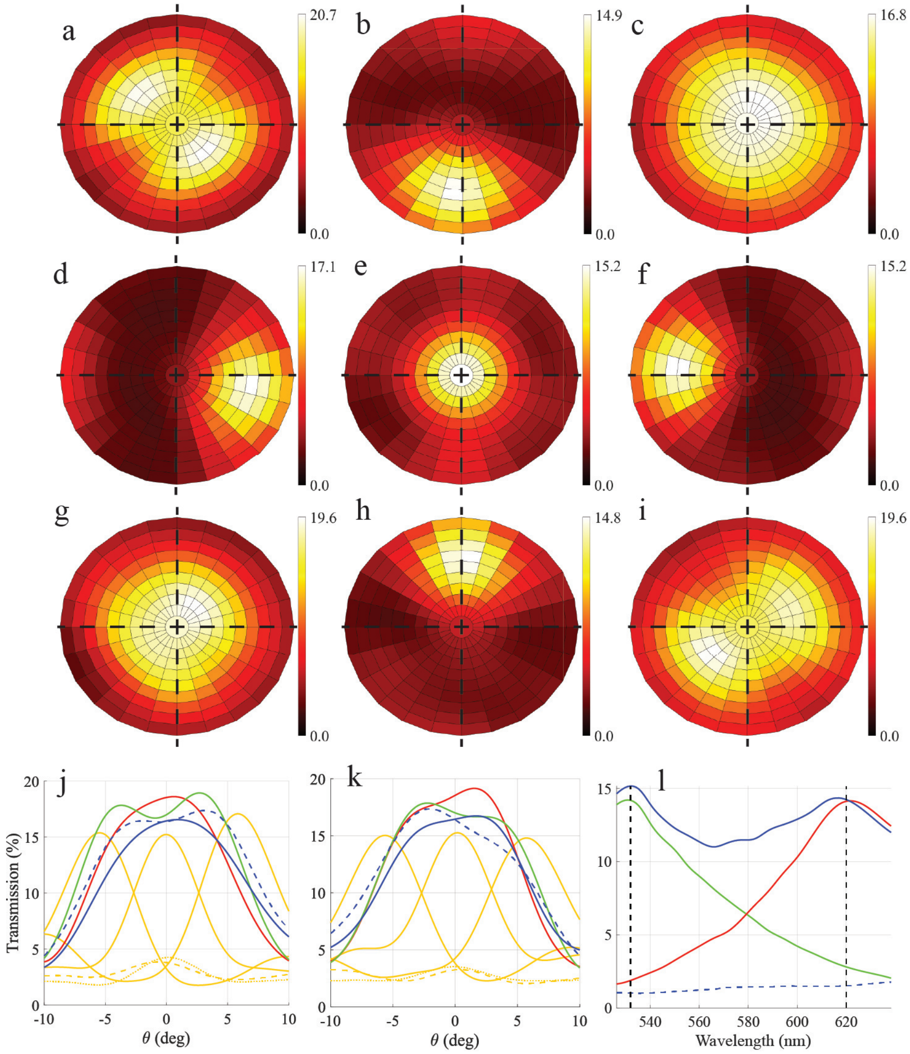

Figure 3.14: The angular sensitivity of the device. (a)-(i) plot the transmission to the different pixels in the focal plane. The top-down left-right ordering of these plots are matched to the ordering of the pixels depicted in Fig. 3.11b. The wavelength demultiplexing plots (a) and (g) are averaged in polarization. The polarization sorting plots (c) and (i) are averaged in wavelength. The imaging plots (b), (d-f), and (h) are averaged in both polarization and wavelength. (j,k) are horizontal (j) and vertical (k) traces of the surface plots (a)-(i) along the black dashed lines. For (j) and (k), the line colors represent transmission to different pixels: the green line is the 532 nm pixel, the red line is the 620 nm pixel, the solid blue is the xz-polarization pixel, the dashed blue is the yz-polarization pixel, the solid yellow lines are the on-axis imaging pixels, and the dashed lines are the off-axis imaging pixels. For (l), the green and red lines represent the transmission to the 532 and 620 nm pixels respectively. The solid blue line is the transmission to the correct polarization pixel given the polarization of the input source, and the dashed blue line is the transmission to the incorrect polarization pixel. For (l), all traces are averaged over all evaluated input angles within a 5° cone.

between these two wavelengths, indicating that the power is predictably redistributed between the pixels in these cases. The blue line represents the average transmission to the *correct* polarization pixel (for example, xz-polarized input being focused to the xz-polarization pixel), and the blue dashed line represents the average transmission to the *incorrect* polarization pixel (for example, xz-polarized input being focused to the yz-polarization pixel). While there is a drop in efficiency at the non-optimized wavelengths, the polarization is still accurately classifiable due to the high contrast between the solid and dashed blue lines at all wavelengths.

The effect of altering the polarization state is predictable since the polarization state of any input excitation can be described by two orthogonal linearly polarized states with a relative amplitude and phase shift between the orthogonal components. As an example, a simple experiment involving a linearly polarized source input into a Wollaston prism would show that as the source is rotated, the output of the Wollaston prism fluidly shifts from one linearly polarized output beam to the other. This behavior is observed in this device in the polarizing functionality, where the ratio between the polarizer pixels transmission can be used to infer the relative power of the two orthogonal linear polarization states. However, we would hope that the device performance for the wavelength and imaging functionalities is not adversely effected by the polarization states that were not explicitly optimized for. The efficiency of these functionalities may be altered as the cross-polarized output of one beam interferes with the parallel-polarized output of the orthogonally polarized input beam. We define the amount of cross-polarization by integrating the ratio $|E_x|^2/|E_y|^2$ over the output plane for all incident angles. The highest amount of cross-polarization is -11.9 dB at 620 nm and -13.6 dB at 532 nm. To verify that the various functionalities are not strongly affected, we analyze the results of all functionalities at all possible input polarization states by sweeping the relative amplitude and phase of the orthogonal input components, and we observe that interference effects can alter the transmission to the various pixels by a few percentages. Data and further commentary on this matter is available in Appendix A. Note that if a particular design seeks to further minimize cross-polarization, then cross-polarization can be explicitly minimized during the optimization. Doing so will consume some design degrees of freedom, which may detract from the efficiency of the various functionalities, but the optimization will not require substantially more time to optimize since the number of electromagnetic simulations per iteration will be unchanged.

Discussion

Based on these results the device could be used to classify the state of incident light at wavelengths between 532 nm and 620 nm and inputs with incident angles within an approximately 5° cone. Additionally, the relative amplitude of the xz-polarized and yz-polarized component can be classified. The classification is done by measuring the ratio of intensity between pixels and using a look-up table to then approximate the state of light. In the case of classifying wavelength and imaging angle, the behavior of smoothly transitioning between states is non-trivial and is not exhibited in typical optical devices such as gratings and lenses. These components tend to spatially shift a beam in the focal plane as the wavelength is varied (grating) or the incident angle is varied (lens), whereas our device only alters the ratio of transmission to the various pixels with very little spatial shift of the beams. In the case of classifying polarization this behavior tends to occur naturally in other optical devices such as birefringent prisms because any polarization state can be decomposed into a coherent sum of two orthogonally polarized states.

While classifying the wavelength and input angle states without ambiguity is not significantly affected by the incident polarization state, as was discussed in the Results section, it does require that the incident light be a monochromatic planewave. The planewave assumption is a common assumption employed in Shack-Hartman sensors and plenoptic sensors, and the assumption of monochromaticity can be satisfied with color filters, or if the light fundamentally comes from a narrow-band source such as a laser. Thus, there are numerous applications in which this device can be useful. The device can be tiled across a CMOS array to enhance the functionality of an image sensor [32, 73, 74]. Such an array could be used to accurately classify the properties of a laser beam, including all fundamental properties of wavelength, polarization, and incident angle within the device's acceptance cone. Note that this would require reoptimizing the device with periodic boundary conditions instead of PML boundary conditions. The angle-dependent nature of the device hearkens to angle-sensitive CMOS pixels[75], and could be used for lightfield imaging since the incident wavefront angle can be computationally determined using the relative intensity of the imaging pixels [14, 76–78]. If coupled with a device such as a tunable bandpass filter, then a wavelength-dependent light field image can be obtained, with the added functionality of measuring the relative intensity of the two orthogonal linear polarization components for basic polarimetry. In general, we believe this type of device will open up new degrees of freedom in computational imaging applications.

We investigated whether the qualitative behavior of the device, and its corresponding ability to interpolate input states effectively, depends on the specific assignment of functionalities to pixels by studying the behavior of devices that were optimized for different pixel distributions. There are a large number of pixel permutations, and currently it is not easy to design a device for all of them. Instead, we investigated twenty different pixel distributions. Pixel distributions were chosen randomly, but were chosen to be sufficiently different from the original device by satisfying two criteria: 1) all pixels had to be moved from their original location, 2) no pixel could be rotated by 90°, 180°, or 270° degrees, thus ensuring that devices were not similar by any rotational symmetry. Each device was optimized until the device was more than 15% binary. We compare the angular response of each functionality (the plots in Fig. **3.13**c-k) to one another using the overlap integral in Eq. 3.2, integrated over all simulated (θ, ϕ) points. This equation quantifies the similarity in the angular transmission profile, evaluating to zero if the responses are completely dissimilar and one if the responses are the same.

$$\eta(T_1, T_2) = \left[\frac{\left| \iint T_1 T_2 d\theta d\phi \right|^2}{\iint |T_1|^2 d\theta d\phi \iint |T_2|^2 d\theta d\phi} \right]^{1/2} \tag{3.2}$$

For each functionality, we evaluate the overlap integral for all combinations of devices. Each result is plotted as a scatter point in Fig. **3.15**, where the x-axis represents the nine different functionalities and the y-axis represents the computed overlap integral. The results indicate that all transmission profiles are very similar, suggesting that the qualitative behavior of the device is not strongly dependent on the mapping of the different functionalities to pixels. We withhold making any stronger claims than this qualitative similarity in behavior due to the limited number of pixel distributions that were studied.

The fabrication of this device would be ambitious, but achievable with state-of-the-art current or near-future technology. The optimization procedure ensures that the device obeys a minimum feature size fabrication constraint of 50 nm, which is obtainable with standard electron-beam lithography. A twenty-layer device with good layer-to-layer alignment is achievable in CMOS fabrication, although such processes are primarily used for electronics rather than optics and the materials used are not suitable for low-loss optical devices. A promising recent advancement involves designing a photonic device in a CMOS platform, then removing the metal layers with a wet etch to leave only materials made of the dielectric layers and air

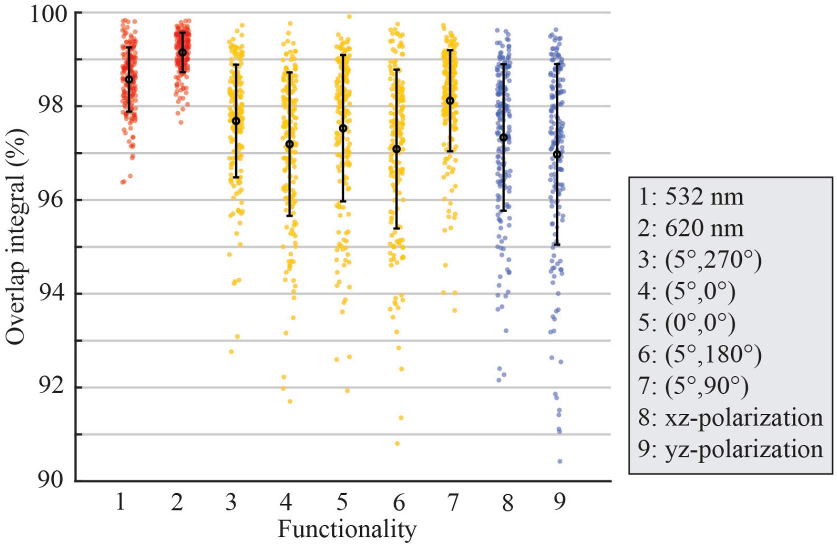

Figure 3.15: The dependence of device performance on pixel distribution. For each functionality the overlap integral in Eq. 3.2 is evaluated on all combinations of pixel distributions. The average overlap integral is greater than 96.9% for all functionalities, suggesting that the arbitrary choice of mapping functionalities to specific pixels has little effect on the behavior of the device.

[79]. This "subtractive photonics platform" could potentially be used to fabricate the devices presented here, particularly if the air was backfilled with TiO2. Note that this procedure would require that all layers be connected to maintain structural integrity when the metallic layers are etched away, and such connectivity constraints have been enforced in past works.

Conclusion

We presented an extremely multi-functional 20-layer device made of SiO2 and TiO2 capable of classifying the fundamental properties of an incident light wave: polarization, wavelength, and incident angle. The ability to classify wavelength and angle is contingent upon assumptions of the incident light. Monochromaticity is required for uniquely classifying the wavelength, which is a common assumption made in laser applications. A flat phase front corresponding to an incident angle of approximately 5° or less is required for classifying the incident angle. This an assumption that is commonly employed in Shack-Hartman and plenoptics sensors for applications such as wavefront sensing and computation imaging, and is valid when the device is laterally small. Microlenses are often used for these applications, which tend to be much larger than the device presented here.

The device was designed using topology optimization, aided by adjoint-based inverse design. The multi-objective optimization used a novel technique to approach a nearly binary solution during the continuous permittivity design phase, before concluding with a level-set approach that respects fabrication constraints. It was found that the device functions at many more input conditions than those that were explicitly optimized for, which is an extremely promising feature that eases the computational burden of optimizing many different merit functions.

The fabrication of this device, while ambitious, could be achieved with nanofabrication techniques used in CMOS fabrication facilities. However, there does not yet exist foundry processes for TiO2/SiO2 material systems with fabrication tolerances akin to CMOS devices, and this fact precludes the ability to fabricate these device at the present. In the mean time, a scaled-up version of this device could possibly be fabricated and tested at infrared or mm-wave frequencies to enable new imaging technologies. The remainder of this thesis covers similarly designed devices operating at higher frequencies.

Chapter 4

TERAHERTZ METAOPTICS

The THz, or *sub-millimeter*, portion of the electromagnetic spectrum is the range extending from about 300 GHz to 3 THz (these bounds depend on the source of information). There is immense scientific value in this spectrum, particularly due to the existence of many vibrational and rotational spectral lines of important molecules, including water and oxygen [80, 81]. Furthermore, over 98% of all photons *in the universe* exist within the THz regime [82]. With this in mind it is no wonder that the THz spectrum has garnered immense scientific interest. Despite this interest, THz technology remains difficult to develop. The technological gap between microwave electrons and optical/infrared photonics has been deemed *the THz gap*. This chapter will explore the scientific value of the THz spectrum, the challenges in working with it, and how metaoptics are uniquely suited for enabling compact, multi-pixel THz imaging systems to help bridge the technological gap. Chapter 5 continues by introducing advanced topics that metaoptics are well-suited to advance such as spectroscopy.

4.1 The science in THz radiation

Interest in terahertz research is primarily motivated by astrophysical, planetary, and Earth science-based applications [81]. Substantial research potential comes from the existence of many vibrational and rotational spectral lines within the THz band, defined here from 300 GHz to 3 THz, which fueled the vast growth in terahertz imaging and spectroscopy technologies over the past few decades. A recent example is the potential detection of phosphine on Venus using the Atacama Large Millimeter Array and James Clerk Maxwell Telescope at 267 GHz, fueling the speculation of existence of life in the Venusian atmosphere [83]. This research could be pursued further by analyzing the even stronger spectral line of phosphine at 533 GHz, particularly with a space-based Venus limb viewing spectrometer instrument, to discern whether the phosphine detection is accurate, image its spatial distribution, and help determine if it has a biological origin. Recent results of the detection of water on the sunlit surface of the moon [84] also opens up the possibility of further studies using high-resolution spectrometers at 557 GHz, which happens to be the strongest water line. A highly compact 3-D metasurface based spectrometer proposed here will be highly attractive to confirm the results with high-fidelity. Other applications of THz imaging and radiometry include the observation of molecular components in Earth's atmosphere, and the observation of the critical CI line at 1.9 THz that is considered key to understanding the evolution of young galaxies.

Considerable investment has been placed in terahertz technology, including spaceborne instruments that operate above the relatively opaque Earth atmosphere. For example, the Herschel Space Observatory active from 2009 to 2013, housed the Heterodyne Instrument for the Far Infrared (HIFI) which featured both direct-detector based imaging and high spectral resolution heterodyne imaging. The past success of THz imaging observatories like Herschel and ALMA, along with the present excitement of remote sensing molecular species at THz frequencies, testifies to the immense scientific interest still ongoing for THz applications.

THz cameras have historically been limited to single-pixel detection schemes. Direct-detector systems and heterodyne detector systems operate near their theoretical noise limits, which led to an intense interest to expand to multi-pixel cameras. This interest led to the successful integration of multi-pixel focal plane arrays by leveraging Si micromachining practices that can pattern many devices across the flat area of low-loss Si wafers [85]. With this promising platform for highly integrable arrays, a subsequent push to improve focal plane array efficiencies and enable complex functionality, such as polarization and spectral sorting, is underway.

4.2 Fabrication and assembly of Silicon THz devices

Efficient metaoptics require sub-wavelength fabrication tolerances. This includes the ability to pattern subwavelength features, as well as aligning individual components with high accuracy. Metal machining and Si micromachining are two fabrication techniques capable of patterning features on the scale of terahertz wavelengths. Metals tend to be low-loss at microwave frequencies, high-loss at optical frequencies, and somewhere in between in the THz regime. For some applications, such as air waveguides with metallic boundaries, the loss is acceptable over short distances. However, both the absorption loss of metal and fabrication tolerances of metal machining are inferior to that of Si micromachining. For this reason, substantial effort has been focused on creating low-loss dielectric components with micromachining instead of metal machining. This decision can intuitively be thought of as focusing on *scaling up* thin film fabrication techniques used in CMOS electronics, MEMS, and optical metaoptics, instead of *scaling down* metal machining dominant in RF and microwave technologies.

Substantial work has been performed at the Jet Propulsion Laboratory (JPL) to create complex THz components using stacks of Si wafers that have been patterned with microfabrication-based technology [63]. In general a single layer benefits from: (1) extremely accurate lithography resolution with respect to THz wavelengths; (2) the ability to fabricate large batches across wafers; (3) the ability to etch high aspect ratio (AR) structures using mature deep reactive-ion etching (DRIE) Bosch processing. Furthermore, the layers can be coated with metal to create metallic components with fabrication tolerances superior to that of metal machining.

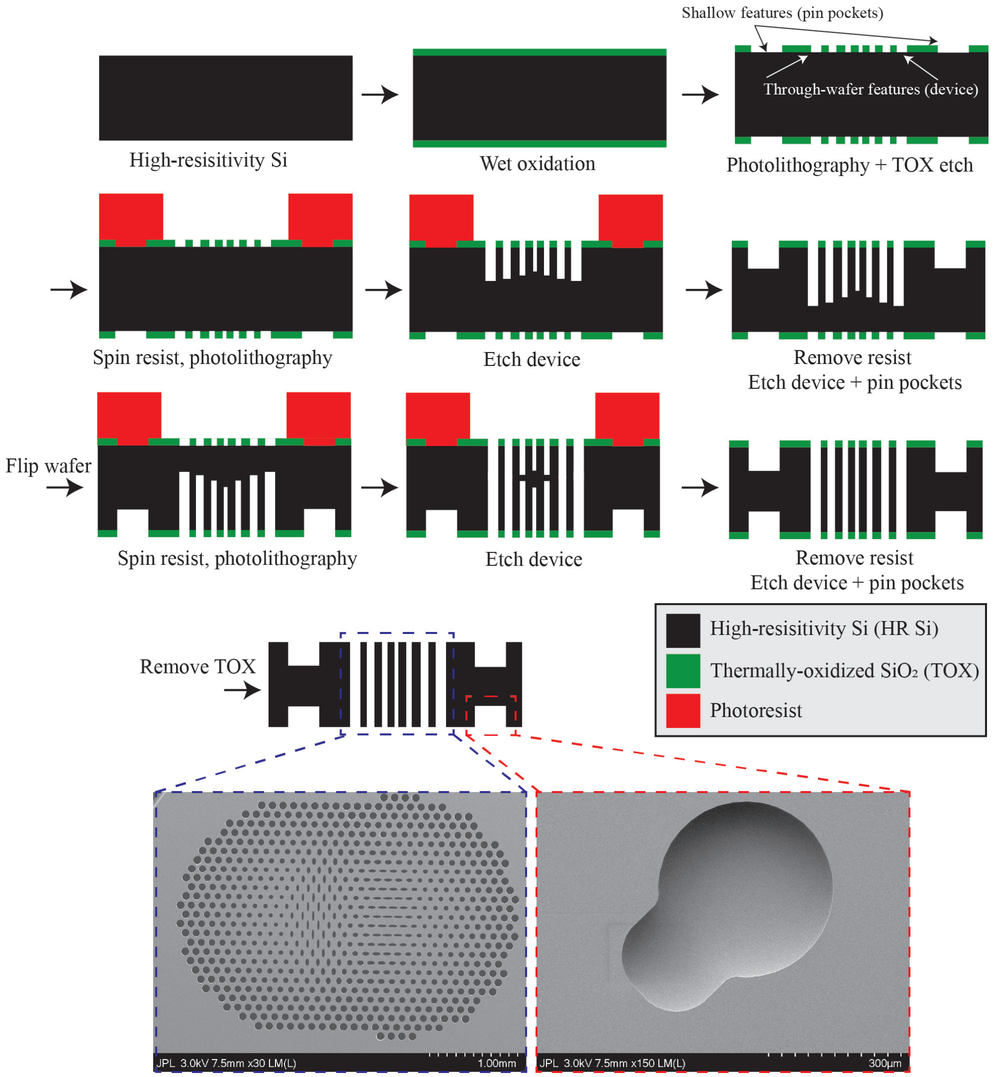

Figure 4.1: **Procedure for fabricating THz devices.** The fabrication procedure is a two-step DRIE etch on both sides of a wafer. One set of features is etched entirely through the device, the other set of features is only partially etched into each side.

Individual layers can be stacked to create complex 3D devices with extremely accurate layer-to-layer alignment. This alignment is achieved by patterning shallow depression at several locations across a device layer, then placing microfabricated Si compression pins in these pockets. A separate layer with identical "pin pockets" can then be placed on the first layer and manually aligned so that the compression pins lock the two layers into place laterally. The layer-to-layer alignment has been shown to be better than 2 μm [85].

4.3 Dual-polarization metasurface for 500-600 GHz focal plane arrays

The basic optical system required for a THz telescope is a large collector, such as a lens or curved mirror, that focuses radiation of different incident angles to different points in the focal plane. One or more sensor elements is placed in the focal plane to sense this radiation. If only one element is used, then the formation of an image requires scanning some feature of the telescope. Examples include mechanical scanning of the full telescope system, mechanically scanning of the sensor in the focal plane (referred to as dithering), or electrically tuning the system such as a phase array system. To reduce the amount of scanning necessary, multi-pixel arrays can be used to image different input angles simultaneously. Multi-pixel arrays are extremely common at optical and infrared frequencies, since densely integrated CMOS arrays patterened on flat substrates are readily available at these frequencies. However, large-scale integrated focal plane arrays remain difficult at THz primarily due to strict power, size, and efficiency requirements, as well as limited fabrication choices capable of densely packing pixels across a large area.

The fabrication techniques described in the previous section permit the patterning of well-aligned high-resistivity silicon (HR-Si) structures. Since HR-Si is an extremely low-loss material at THz frequencies, and since features can be patterned across large areas with mature fabrication techniques, it is a natural choice for creating high-efficiency focal plane array elements. A recent work utilized Si microlenses for coupling radiation to waveguides, which are then routed to highly sensitive detectors such as hot electron bolometers. The intention of the current work is to replace these microlenses with metasurface lenses, which offer several advantages: (1) metasurfaces are planar, and thus easier to integrate over large areas; (2) metasurfaces feature custom phase profiles dictated by the shape of lithography-defined elements, rather than the curvature of a material, so certain optical precise phase-matching and reduction of critical abberations can be achieved; and (3) metasurfaces offer polarization control, which offers several advantages that we elaborate on below.

The polarization control capabilities of metasurfaces allow for either polarimetric imaging, if that is desired for a particular scientific study, or for improving the received power and imaging resolution of the system. When coupling to waveguides, single mode waveguides are typically used to minimize losses. With this limitation, polarization-insensitive microlenses focus both polarization to a single-mode waveguide, and thus half the power is lost assuming the incident radiation contained equal power for the orthogonal polarizations. It is possible to place two orthogonally polarized waveguides in close proximity to another, but this necessarily reduces the efficiency of the system since at least one waveguide must be spatially separated from the focal spot of the lens. It is also possible to focus both polarizations to a square waveguide, capable of supporting modes from both polarizations, then separating the polarization immediately using an element known as an orthomode transducer. This approach has been demonstrated before and shown excellent polarization contrast and overall transmission [63].

Metasurfaces, like the orthomode transducer, separate the polarization components but they do so by focusing two orthogonal polarizations to distinct locations in the focal plane. The ability to encode arbitrary phase profiles for both polarizations independently ensures that excellent coupling to the specific modes of the waveguide can be achieved. A benefit to seperating the polarizations in free-space rather than within the waveguides layers is an increase in the resolution of the imaging system when viewing a polarization independent scene. When tiling waveguides in the focal plane of the lens element (microlens or metasurface), parallel-polarized waveguides must be spatially separated to reduce cross-coupling. However, cross-polarized waveguides do not couple substantially, and thus can be placed much closer to each other. Thus, a grid of parallel-polarized and cross-polarized waveguides can be interleaved to make a dense imaging array.

System design and fabrication

The metasurface is designed to couple an incident planewave to two orthogonally polarized waveguides. To simplify the design process, we can invoke reciprocity and equivalently design a device that collimates the output of the waveguide to a planewave. A prototype containing two cross-polarized waveguides coupled to the dual-polarization metasurface is shown in Fig. 4.2, and to describe it in detail we can follow the path of a signal being injected into the waveguides from the bottom, radiating to the metasurface, then coupling out to a free-space planewave.

At the bottom of the stack is a metal-machined waveguide adapter that can connect to the WR-1.5 waveguides of a vector network analyzer (VNA). The remaining elements in the stack are all fabricated in HR-Si using a process similar to the one described in 4.1. The variations to the fabrication procedure for individual layers are illustrated with their own figures below. The bottom two wafers are waveguide layers that route the signal from the VNA to the location of the metasurface focal point. The next wafer is a two-step taper that transitions from the single mode waveguide to a square-waveguide. Next, a SOI wafer is patterned with an iris that matches the square waveguide to a Fabry-Perot cavity, which itself is made by placing a HR-Si wafer with an air cavity etched into it above the iris. The purpose of this cavity is to expand the beam emitted by the waveguide, and thus increase its directivity. The radiation pattern of a small waveguide is nearly isotropic, which would require the metasurface element to be extremely large to capture all radiated energy. The beam output from the Fabry-Perot cavity, on the other hand, is well-collimated, reducing the required size of the metasurface substantially. After the Fabry-Perot cavity, numerous blank HR-Si wafers are stacked. These wafers are "blank" only in the sense that there are no patterned features within the path of the emitted beam. There are, however, patterned pin slots to ensure good alignment between the metasurface and the waveguides. The metasurface is then added to the stack. Finally, the metasurface is capped with another blank Si wafer that has an SU-8 anti-reflection coating patterned on it. The motivation for incorporating this will be made clear once the design of the metasurface element is discussed.

The metasurface is fabricated using the techniques in Section 4.2, exactly as described in Fig. 4.1. The waveguide wafers are fabricated using a similar 2-step per side DRIE etch, but there are enough differences to warrant further discussion. Taking the direction that the wafers are stacked in to be "vertical, the bottom two waveguide wafers route the signal vertically, then transition to a horizontal propagation, then back to vertical propagation. The transition from horizontal to vertical propagation and vice versa can cause large reflections in the signal, which can be mitigated by incorporating a step feature that serves to discretely taper the geometry. The size of this step is constrained to the depth of the pin pocket, which is nominally 120 μm. This concept is illustrated in Fig. 4.3, which shows a cross-section of the waveguides in Fig. 4.3a, a fabricated gold-coated wafer in Fig. 4.3b, and a CAD drawing of the bottom waveguide layer in Fig. 4.3c. Note that there may be flexibility in this quantity by fabricating a different set of Si compression pins with different thicknesses, but this was not tested. There is also a tolerance on the depth

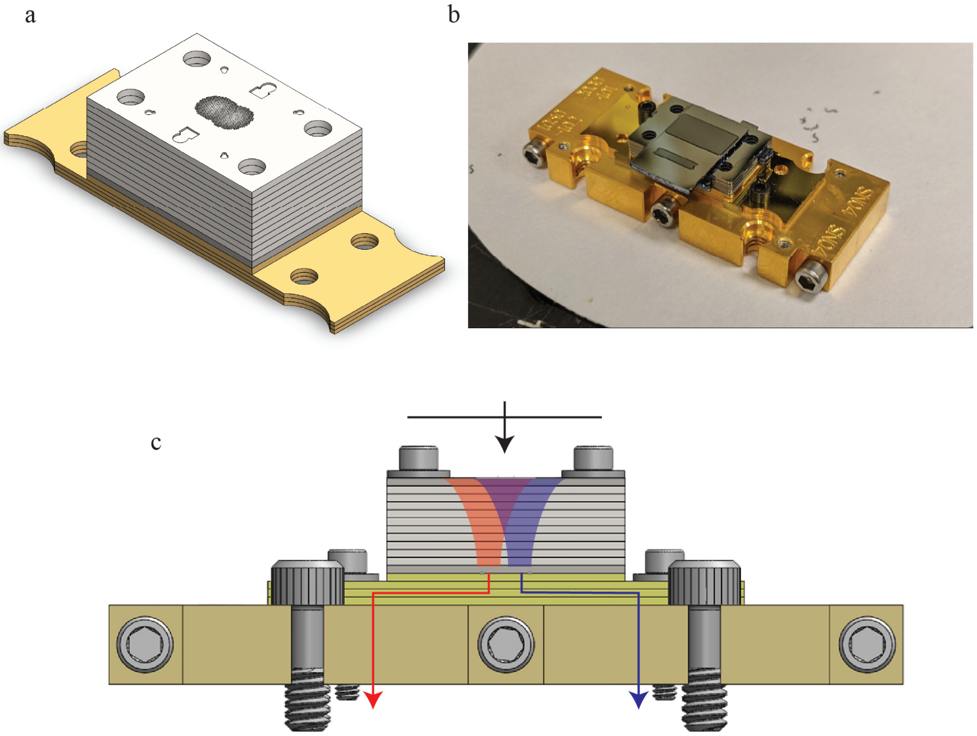

Figure 4.2: **Wafer stack for THz focal plane array.** (a) A computer-aided design (CAD) drawing of the assembled device without an AR coated top layer. (b) A picture of an assembled device, with an AR coated top layer made of patterned SU-8. (c) A schematic demonstrating the working principle in reception mode. The two orthogonal linear polarizations of a planewave incident from the top is focused into two cross-polarized waveguides, routed horizontally, then down through the adapter block to a VNA.

of the pin pockets: if too shallow, then the Si compression pin will protrude from the bottom wafer too far and lift up the top wafer; if too deep, then the Si compression will not protrude enough to permit a stable lock between the wafers. Thus, an etch depth of 112 μm was used which was both adequate for the compression pins to work as intended, as well as a locally optimal value for minimizing reflections at 550 GHz. After etching the waveguides 2 μm of gold is sputtered on both sides of the device.

The design of the iris, which is detailed in [85], also has a slightly different fabrication process. The iris requires a very specific thickness, the precision of which is difficult to achieve with the process described thus far. The difficulty stems from a variable DRIE etch rate across the wafer, and a variable etch rate across a given feature which causes the center of the feature to etch faster than the edges. To overcome this, an silicon on insulator (SOI) wafer is used whose buried oxide (BOX) layer is used as

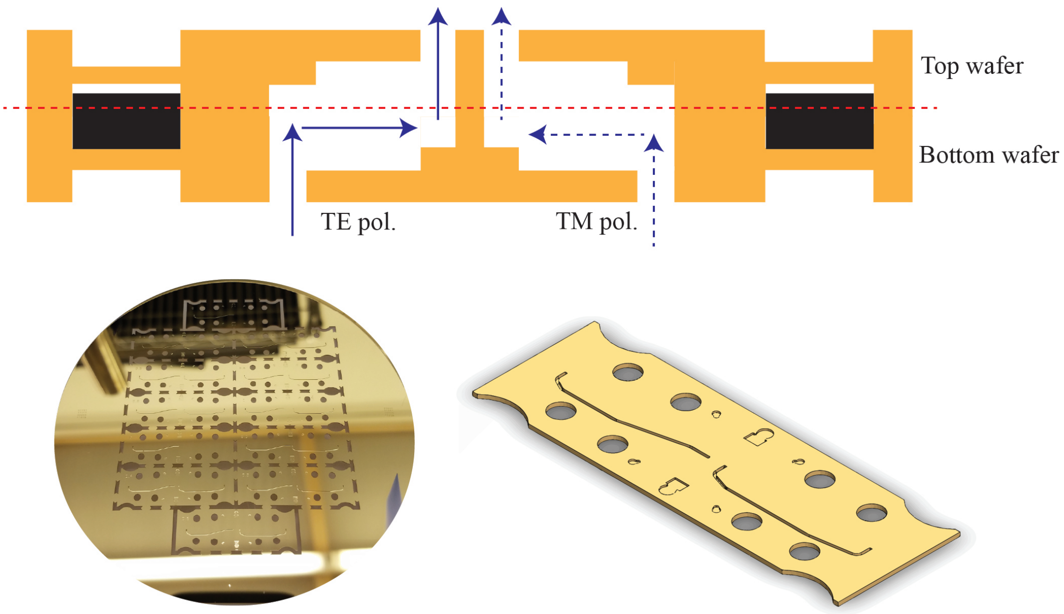

Figure 4.3: **Waveguide layers.** The signal from two cross-polarized waveguides is routed using single-mode waveguides. The rectangular waveguides, consisting of air surrounded by gold, are made by aligning two independently fabricated gold-coated Si wafers. (Top) A cross-section of the waveguide stack. (Bottom-left) A fabricated wafer containing multiple waveguide devices corresponding to different waveguide spacings. (Bottom-right) A CAD drawing of the bottom waveguide layer.

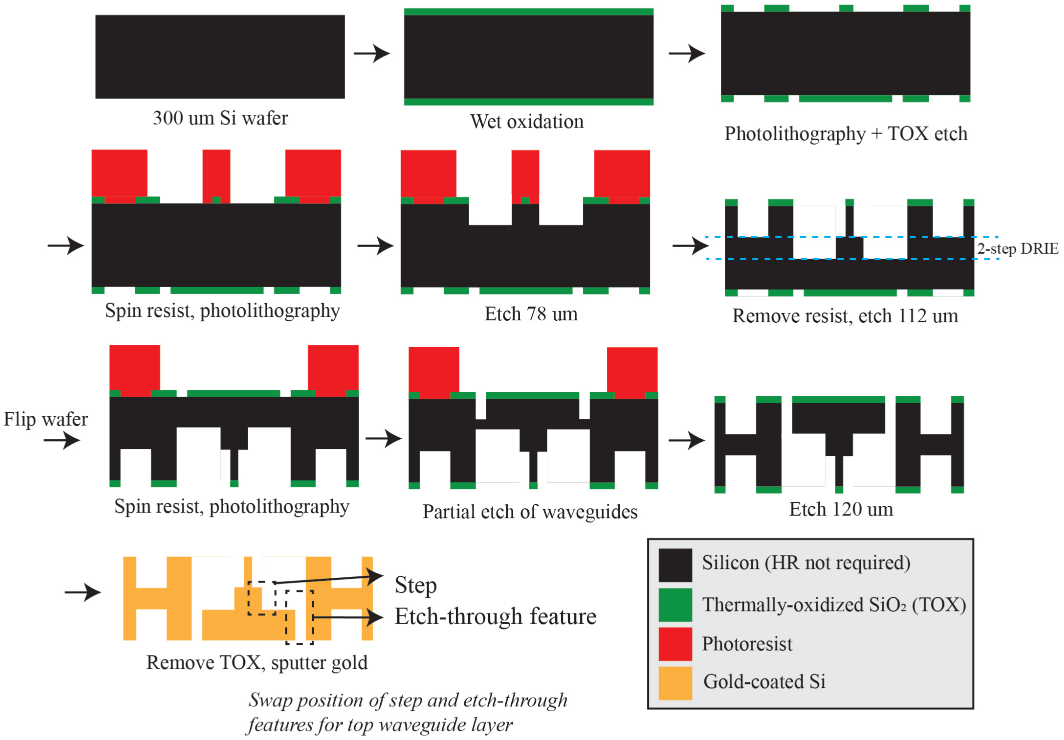

Figure 4.4: **Waveguide fabrication procedure.**

an etch-stop. Fig. 4.5 shows the fabrication process. The most notable change is that three separate etches are required to create the top-side pin pockets. First, the pockets are etched through the device layer to the BOX layer using DRIE. Then, the BOX layer is etched using a fluorine-based inductively coupled plasma reactive ion etching (ICP-RIE) etch. The DRIE etch is then resumed to complete the etch. An important note is that spin-coating resist becomes prohibitively difficult after the first DRIE etch, and thus the exact same photoresist mask should be used for all three etches. For this process, a very thick resist layer of approximately 15 μm was used, which was obtained by spin-coating AZ-5214 at 750 RPM, baking at 110 C°, then repeating this procedure one more time.

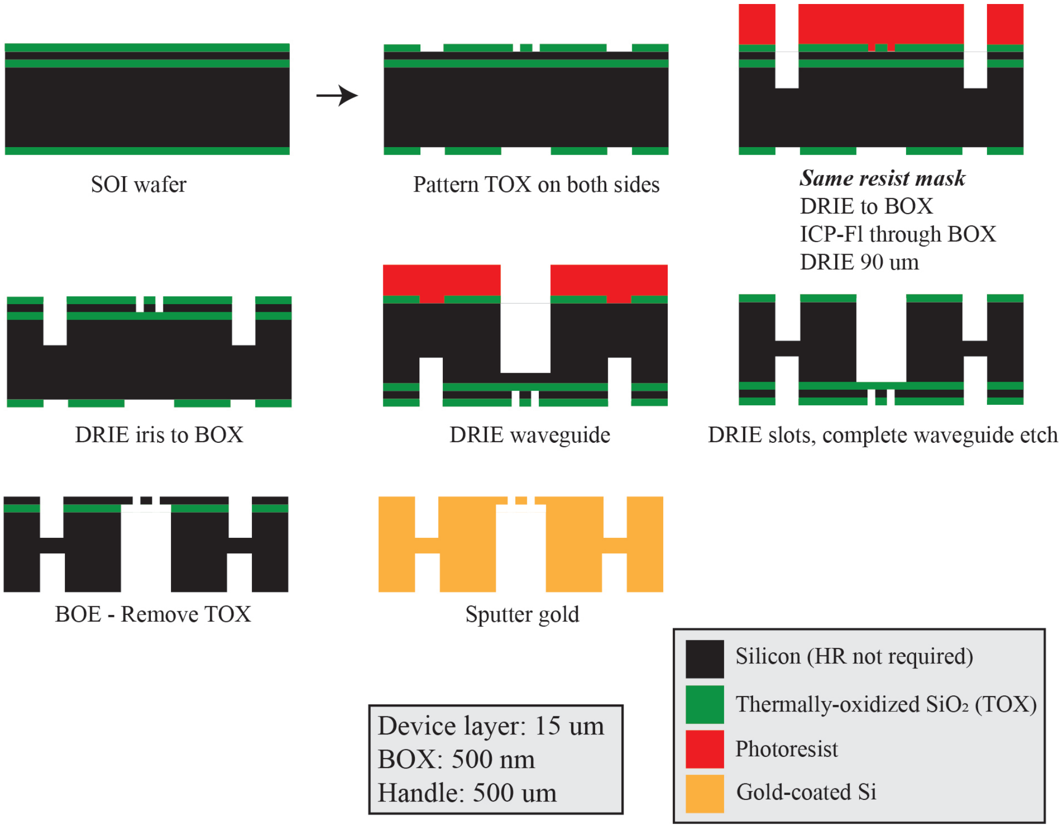

Figure 4.5: **Iris fabrication procedure.**

Metasurface design methods

The first step in designing the metasurface is finding a suitable collection of sub-wavelength elements that can alter phase from 0 to 2π with high transmission. In the case of dual-polarization metasurfaces, the demand is that the collection of elements can prescribe a 0 to 2π phase shift to the incident light of one polarization *for all phases of the orthogonal polarization in 0 to* 2π. For this work we choose the metasurface elements to be elliptical holes rather than posts that are commonly used in metasurface design. This allows the entire substrate to be fully connected, which is critical since we are etching the metasurface elements entirely through the HR-Si wafer.

The lattice period of metasurface is chosen based on the numerical aperture of the lens, which sets an upper bound on the largest angular component that we hope to control with the metasurface and thus a lower bound on the period required to suppress high order diffraction components at this maximum angle of incidence. This condition is analytically solved for in the supplementary documentation of Ref. [86], and the relevant equation is repeated below. Λ is the metasurface period, λ is the free-space wavelength, n is the refractive index of the substrate, and NA is the numerical aperture of the lens.

$$\Lambda < \frac{\lambda}{n + NA} \tag{4.1}$$

The height of the metasurface layer is chosen to be 625 μm, which was the smallest height at which satisfactory phase coverage was achieved. A large number of 625 μm elliptical holes varying semi-major and semi-minor axis lengths are simulated using a periodic RCWA solver (Stanford Stratified Structure Solver described in Ref. [87]), and the complex transmission coefficient to the zeroth diffraction order is extracted. Since the periodicity is sub-wavelength, this is the only diffraction order that contains energy. Compiling these results, the complex transmission coefficient as a function of radius along the x-direction, r_x, and the y-direction, r_y, is known and denoted as $t(r_x, r_y)$. This is simulated for both x-polarized and y-polarized light, so that we can decompose the complex transmission function into two components: $t_x(r_x, r_y)$ and $t_y(r_x, r_y)$. The output of an x-polarized input field E_x is thus $t_x E_x$, and is similarly $t_y E_y$ for a y-polarized input field E_y.

Taking ϕ to represent the phase of t, and noting that any subscript applied to ϕ represents the polarization to which that phase applies, we wish to find the optimal post for every pair $(\phi_x, \phi_y$. An optimal post is one whose transmitted phase is as close as possible to the target values, and whose power transmission is as high as possible. To do this we can find post that minimize an error function defined in Eq. 4.2.

$$\text{Error} = \left|T_x e^{i\theta_x} - e^{i\phi_x}\right|^2 + \left|T_y e^{i\theta_y} - e^{i\phi_y}\right|^2 \tag{4.2}$$

Where T is the transmitted power for the post ϕ is the desired phase shift, and θ is the actual transmitted phase. For every phase (4.2) is minimized with respect to hole radius and the resulting collection of metasurface holes are used as a look-up table for sampling phase functions.

An important feature of metasurface is their ability to maintain high transmission. This can be thought of as adequately impedance matching the medium above and below the metasurface. Many near-infrared and optical metasurfaces feature patterned Si nanoposts with air on one side and SiO2 on the other, and these can feature very high transmission values because the high-index Si elements can adequately impedance match the relatively low-index air and SiO2 layers. For these THz metasurfaces, the system is made of only Si and air which have a large index contrast. We found that the metasurface cannot impedance match air to Si very well, with typical transmission values in the range of 70%-80%. To obtain high transmission, we surround the metasurface entirely with Si. This is achieved in assembly by placing blank Si wafers above and below the metasurface layer. It is still important that high tranmission is maintained when coupling the radiation between the homogeneous Si medium and the homogeneous air medium that contains whatever scene the device is imaging. To do this, a $\lambda/4$ anti-reflection coating made of SU-8 is placed on the topmost Si wafer in the wafer stack. SU-8 has an index of approximately 1.73 at 550 GHz [88], which is close to the ideal value of $\sqrt{n_{Si}} = 1.84$ for $n_{Si} = 3.41$. A $\lambda/(4n_{SU8}$ at 550 GHz is 79 μm, which we obtain by spin-coating SU-8 (MicroChem SU-8 2050) at 2000 RPM. The SU-8 was patterned via photolithography and the height of the layer was measured with a profilometer for validation.

To find the phase profile that the metasurface should encode we sample the conjugate phase of the radiating antenna feed. The feed, consisting of the waveguide loaded by an iris and Fabry-Perot cavity, is simulated in Ansys HFSS and the resulting field profile in a plane slightly above the Fabry-Perot cavity is extracted. These fields are then computationally propagated using the ASM (see Ch. 2 for details) to the plane where the metasurface is placed. The phase is then conjugated, and the phases at the location of metasurface posts is interpolated. At each point, the appropriate metasurface element is placed to achieve the desired phase shift.

A choice must be made for how large we make the metasurface. If the metasurface is too small, then most of the radiated power will not be collimated as intended. On the other hand making the metasurface much larger than the beam is unnecessary, since very little power will be coupled to the power at high angular components.

Furthermore, reducing the size of the metasurface improves the resolution of an array of metasurfaces. Thus, the metasurface size should be chosen to be as small as possible while coupling as much power as possible. Here we use a common choice used in antenna theory that maximizes a term called the *illumination efficiency*, denoted η_i. The term is computed using the mode overlap integral (described in Section 2.1) comparing the radiated antenna field with an apertured uniform amplitude beam that shares the same phase as the radiating beam, representing an ideal phase mask.

$$\eta_i = \frac{\left| \int_{-\infty}^{\infty} E_1(x, y) E_2(x, y)) dA \right|^2}{\int_{-\infty}^{\infty} |E_1(x, y)|^2 \, dA \int_{-\infty}^{\infty} |E_2(x, y)|^2 \, dA} \tag{4.3}$$

$$E_2(x, y) = \begin{cases} 0 & (x, y) \notin \mathcal{A} \\ -\angle E_1 & (x, y) \in \mathcal{A} \end{cases} \tag{4.4}$$

where $\mathcal{A}$ denotes the aperture and E_1 represents the radiating field profile of the antenna. Using (4.4) to simplify (4.3) we get:

$$\eta_i = \frac{\int_{\mathcal{A}} |E_1(x, y)|^2 \, dA}{\int_{-\infty}^{\infty} |E_1(x, y)|^2 \, dA \int_{\mathcal{A}} dA} \tag{4.5}$$

A plot showing η_i as a function of circular aperture diameter is shown in Fig. **4.6**. The optimal aperture size is found to be a diameter of approximately 5λ, which corresponds closely to the contour line at which the normalized electric field power falls to -9 dB of its peak value.

Multiple metasurfaces are designed corresponding to different spacing of the cross-polarized waveguide feeds. The feeds are spaced out from 0, in which they are directly overlapping, to being a full metasurface diameter away in which case the two antenna apertures do not have any overlap. The -9 dB aperture size is used for both of the waveguide feeds, and the overlap of their individual apertures vary with the separation of the waveguide.

The error function (4.2) typically evaluates to very small errors for optimal posts. However, the process utilizes only information from the periodic simulations used to originally ascertain the behavior of the metasurface elements, which may change when the neighboring elements are altered, thus breaking the periodicity. The

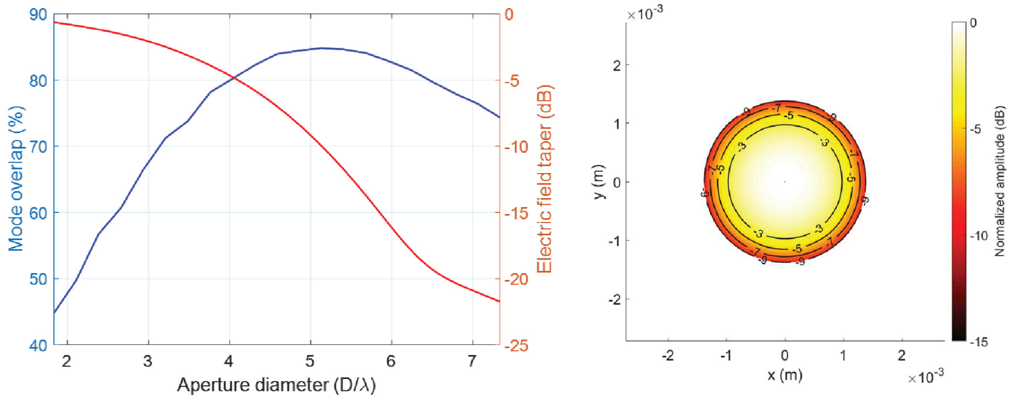

Figure 4.6: **Illumination efficiency calculation for determining metasurface aperture size.** (Left) The aperture efficiency, computed as a mode overlap integral in Eq. 4.5, is drawn in blue as a function of aperture diameter. The electric field taper, which is defined as the ratio of the electric field intensity at the edge of the aperture relative to the center of the aperture, is drawn in red. The maximum illumination efficiency occurs at 5D/λ, corresponding to an electric field taper of -9 dB. (Right) The -9 aperture drawn of the electric field intensity of the beam profile.

assumption that this *does not* happen has helped make metasurface design efficient and scalable. However, perhaps because the metasurface elements here are more densely packed than typical metasurfaces because of the high-index Si substrate, the efficiency of the metasurface is lower than expected when evaluated with full Maxwell's equations solvers. Here, we import the geometry from the designed metasurface into Lumerical FDTD for further evaluation, and observe a drop in efficiency relative to what is expected from a perfect phase mask.

The metasurface is fabricated exactly as described in Fig. 4.1. The etched holes of the metasurface were inspected in detail by cleaving a metasurface (Fig. 4.8a) along its horizontal axis and imaging the cleave line in an scanning electron microscope (SEM). The sample metasurface used here is etched into a 600 μm wafer and contains only circular features ranging in diameter from 27 μm (AR=22) to 77 μm (AR=7.8). Fig. 4.8b shows the cleaved metasurface from an angle, Fig. 4.8c shows the largest holes, and Fig. 4.8d shows the smallest holes. The large holes have highly vertical sidewalls, while the smaller features begin to taper inwards slightly. The small features are not etched entirely through the wafer, but this was a deliberate choice that sheds light on the exact etch rate of the smallest features, and in subsequent etch process the thermally oxidized silicon (TOX) etch mask was thick enough to support fully etching the smallest features. The walls of the features are left slightly roughened by the etch. This roughness is deeply subwavelength at 550 GHz, and

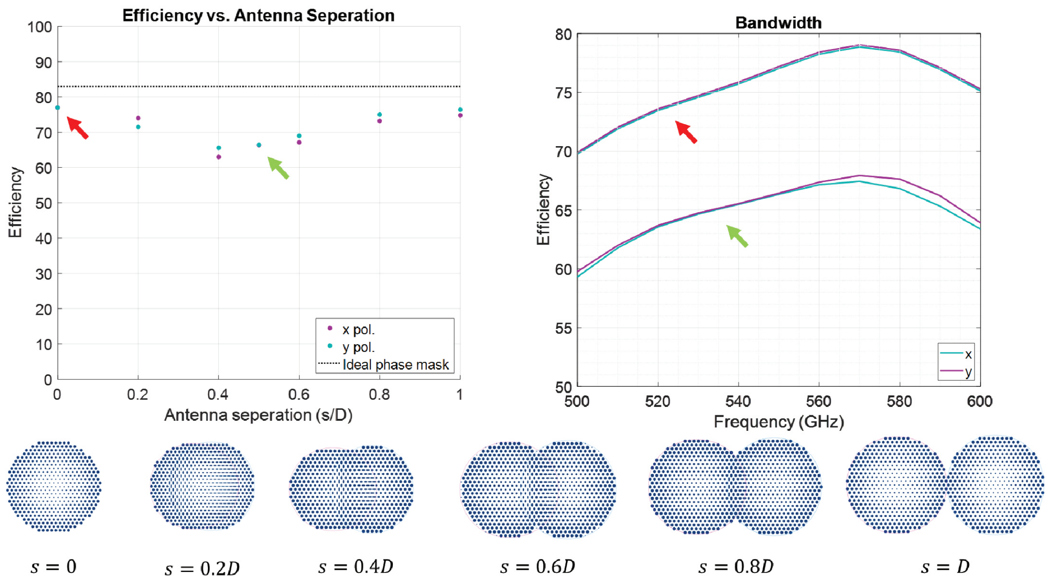

Figure 4.7: **THz metasurface efficiency and bandwidth.** (Left) The efficiency of metasurfaces designed for different distances between the cross-polarized waveguide feeds, evaluated at 550 GHz. (Right) The efficiency over a bandwidth of 500-600 GHz for two different antenna spacings.

thus likely has little effect on metasurface performance. At higher frequencies of operation, the roughness may become more of an issue, and in this case the wafer can be treated by growing TOX, then removing it with an hydrofluoric acid (HF) wet etch. This effectively smooths the surface, and has been demonstrated before for improving the performance of devices operating at 1.9 THz [63].

4.4 Improving the performance with 3D metaoptics

The metasurface is limited in performance in two ways, one theoretical and one based on limitations in the design process. The theoretical limitation exists due to the thinness of the metasurface, which fundamentally limits it to efficiently function only as a phase mask. To maximize the coupled power, both amplitude and phase need to be matched, hence the limit on the maximum efficiency. One may question why a metasurface cannot also act as an amplitude mask. After all, some simulated metasurface elements have lower transmissions than others, so why not utilize this to vary the amplitude of the transmitted wave? The issue lies in *how* the metasurface would control amplitude in that case — through reflection, or perhaps if the metasurface is lossy, then through absorption. In both cases, the power is lost. So while this technique would create a better-matched mode profile, it would do so by sacrificing power, and ultimately the efficiency would not be improved.

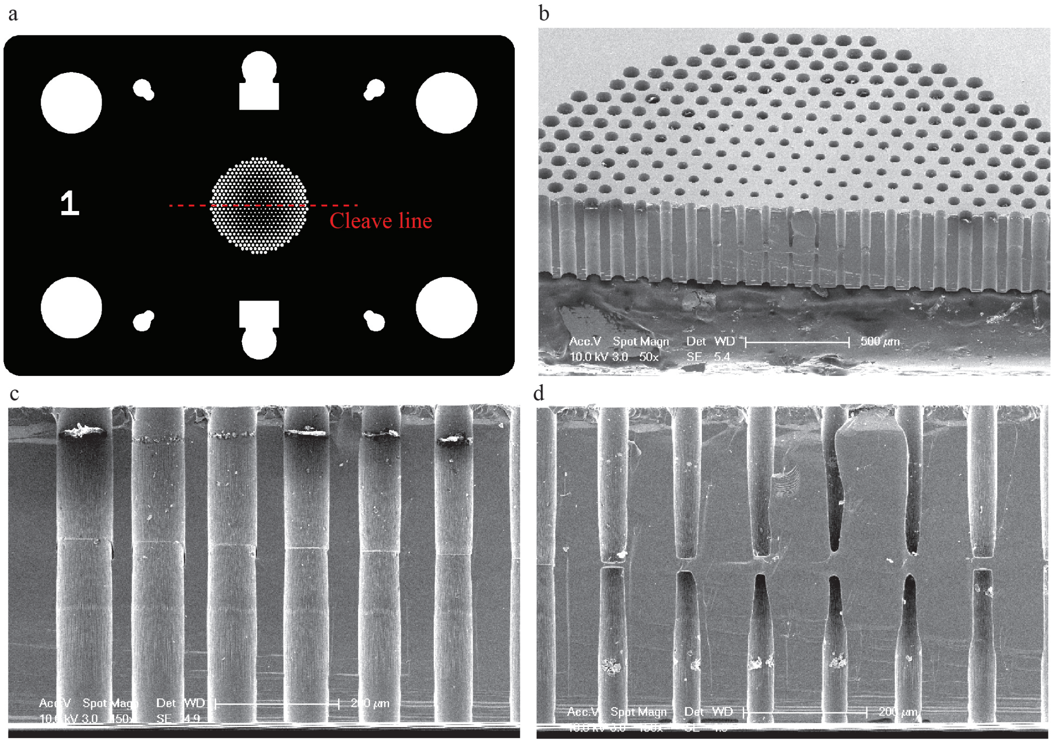

Figure 4.8: **SEM of a cleaved metasurface.** (a) The lithography mask for a metasurface sample. A 600 μm thick metasurface sample is carefully cleaved along the red dashed line using a Dynatex GST-150 scriber-breaker. (b) An angled view of the cleaved metasurface. (c) The larger circular features, ranging in diameter from 77 μm (AR=7.8) on the left to 50 μm (AR=12) on the right, have nearly vertical sidewalls. (d) The sidewalls of smaller features taper inwards slightly. These features were deliberately not etched fully through the wafer so that an accurate measurement of the etch rate could be found from this image. The smallest feature is 27 μm, which corresponds to an AR of 22.

This limitation can be overcome by adding more thickness to the metasurface, which allows the incoming fields to reorient themselves spatially while propagating through the device. This leads us to the next issue that prohibits the metasurface from being highly efficient — the design process itself. The design process relies on simulating a metasurface element in a periodic lattice, which introduces a major challenge in designing a specific spatial routing of light from one point in the actual quasi-periodic lattice of the metasurface. To design this, we would require knowledge of how non-identical elements couple radiation to one another, which violates the assumption in the design process that metasurface elements do not couple substantial amounts of energy.

The goal then is to thicken the metasurface and alter the way it is designed to more accurately account for the complexities of the system. Increasing the thickness beyond 625 μm is difficult for fabrication reasons, so we can pattern the blank

Si layers beneath the metasurface layer to extend the height of the device. Such a device closely resembles the 2.5D layered device described in Section 3.3, and we can use the same inverse-design procedures to optimize this 2.5D device. One major issue in optimizing these devices is the extreme index contrast between Si (n=3.41) and air (n=1). This makes it difficult for the optimization to converge to a binary device during first stage of optimization when the device is described by a continuous, grey-scale refractive index. However, this is overcome using the procedure developed in Section 2.4. In fact, it was this work on THz Si-air material systems that drove the development of that technique.

One difference between this inverse-design procedure and the ones describe in Ch. 3 is the difference in adjoint source. The devices in Ch. 3 were designed to focus all light to a point, which made the adjoint source a dipole at that point. If we used this FOM to optimize the present devices, then the amplitude of the transmitted radiation would focus to the device well, but the phase and polarization would not be guaranteed to match that of the waveguide. Thus, we alter the FOM to be a mode-overlap integral. The adjoint source, which is derived in Ref. [23], is described by equation 4.6. Intuitively, this adjoint source can be thought of as launching the antenna waveguide mode into the device with a complex amplitude that is dependent on the forward simulation.

$$E_{adj}(x,y) = \overline{E_m(x,y)} \iint \overline{E_{fwd}(x,y)} E_m(x,y)\,dxdy \tag{4.6}$$

All field quantities are evaluated in the in the xy-plane at which the mode overlap integral FOM is calculated. $E_{adj}(x,y)$ is the adjoint source in this plane. $E_{fwd}(x,y)$ is the electric field output from the device when excited by the forward source, computed via the forward simulation. $E_m(x,y)$ is the desired mode profile, which in this case is the complex conjugate of the radiating antenna feed profile.

The convergence plots showing the average device FOM and device binarization are shown in Fig. **4.9**. This optimization is very similar to the optimization of the sugarcube device in Chapter 3 (see Fig. **3.12**), but the details are repeated here for convenience. At various points the binarization is forced to increase by passing the permittivity through a sigmoidal function and changing the device permittivity to the output of this function. The FOM is then allowed to recover before repeating this discrete push in binarization. The device loses performance as the device becomes more binary. At iteration 2739 the optimization switches to a level-set optimization, which recovers the the performance that was lost during the thresholding operation that binarizes the device fully at the end of the density optimization. During the level-set optimization, the binarization is recorded as approximately 95% because the boundaries of the device are smoothed out yielding a continuous permittivity value. This can intuitively be thought of as the level-set function passing through a simulation mesh voxel, which is modelled in FDTD as a dielectric volume average of the two materials.

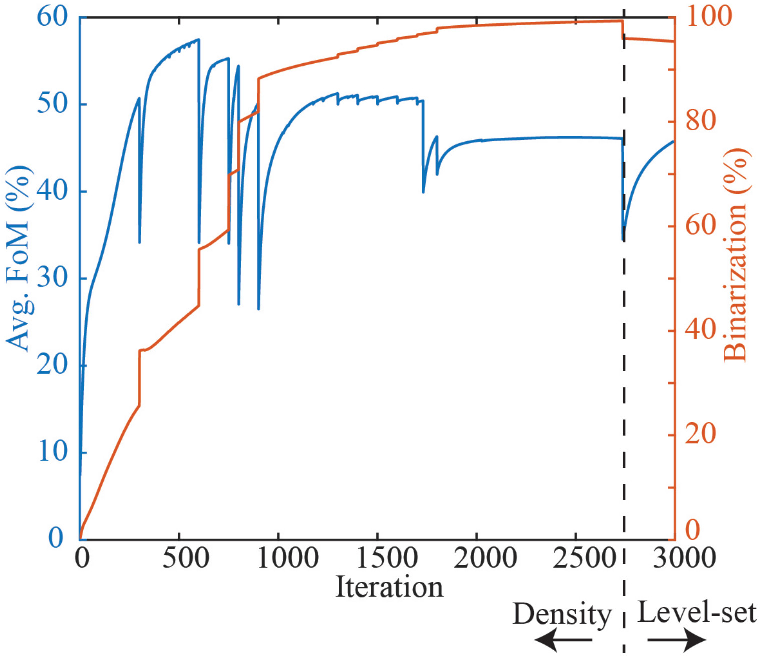

Figure 4.9: **Convergence plot for the inverse-design of the THz 3D metaoptics element.** The initial density-based optimization features numerous points at which the binarization of the device is forced to increase by passing the current permittivity through a sigmoidal function. At iteration 2739 the optimization is converted to a level-set optimization, which recovers the performance lost from that transition.

The designed device has an efficiency that is approximately 1.8× better than the metasurface (evaluated at 550 GHz, and averaging the efficiencies of the two polarizations) designed for the waveguide spacing and focal length at 550 GHz. The efficiency over the full bandwidth from 500-600 GHz is shown in Fig. 4.10. The efficiency in this case is calculated as the ratio of the power coupled into the waveguide (evaluated using the mode overlap integral in Eq. 2.2) normalized to the power incident on the full rectangular aperture of the 3D device, whereas the metasurface efficiencies shown in Fig. 4.7 were normalized to the -9 dB circular apertures for each waveguide.

The efficiency of the 3D device also exceeds that of an ideal phase mask, denoted by the horizontal black line in Fig. 4.10. This is possible only by tailoring both the *phase and amplitude* to the waveguide mode, which is impossible to do with high efficiency with thin metasurfaces [10]. For a detailed discussion, see Section 2.3 and Ref. [10].

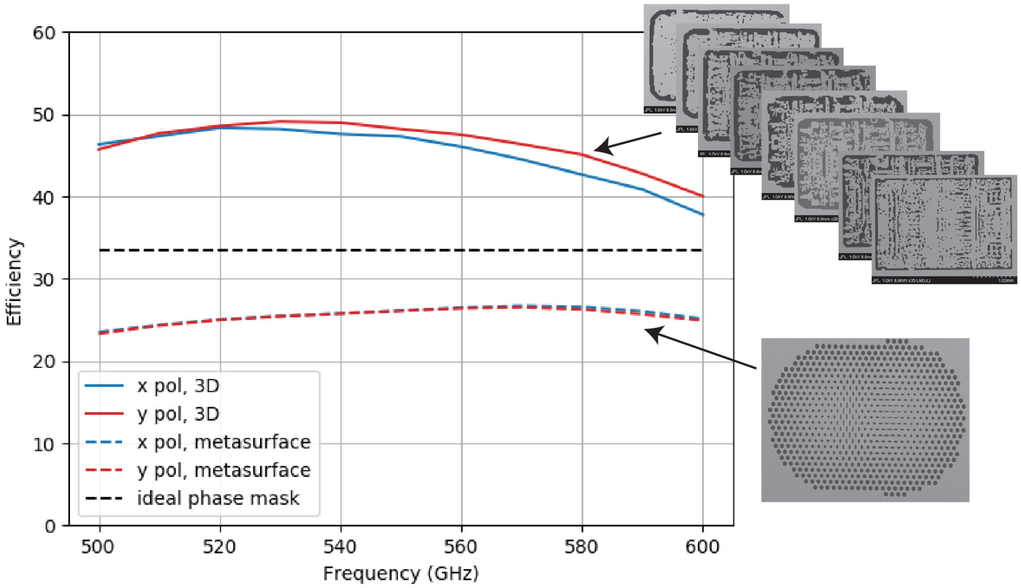

Figure 4.10: **Efficiency of the 3D THz metaoptics over a 500-600 GHz bandwidth.** The efficiency of the 3D THz device is 1.8× better than the metasurface (evaluated at 550 GHz, and averaging the efficiencies of the two polarizations), and 1.4× better than an ideal phase mask. Note that in Fig. 4.7 the metasurface efficiency is defined with respect to the circular aperture over which the phase is sampled, whereas here it is defined over the full rectangular aperture of the 3D device.

Like the metasurface, the Si of each layer of this multi-layer device must be fully connected. This was imposed in a different way than a minimum feature size constraint described in Section 2.5, which incorporated the optimization of minimum feature size into the overall optimization loop alongside the electromagnetic FOM. Instead, connectivity is enforced at discrete points along the optimization, typically every 10-20 iterations, during the continuous phase. To connect a single layer, first a copy of the device is made whose greyscale permittivity is thresholded to its nearest material boundary. Then, all disconnected islands are identified and labelled including the surrounding substrate. Returning to the original greyscale permittivity device, a network is created where each node represents a permittivity point in the design and the bidirectional edges that connect these nodes represent the "cost" of flipping that permittivity from low to high about the mid-point permittivity. Since we do not wish to flip Si to air at any point, then we set the cost function to be zero when moving from a Si feature to another Si. The action of flipping certain permittivity points is what will be used to create paths that link the islands, and quantifying the cost of flipping these points allows us to numerically compute a set of least-cost paths using graph theory algorithms. A *minimum spanning tree* is computed using a bi-directional variant of Djikstra's algorithm, and the nodes of the computed paths are flipped to Si.

The device is fabricated with the same procedure as the metasurface, and the same procedure outlined in Fig. 4.1. The wafer thickness is reduced from 625 μm to 300 μm. All devices of the 8-layer device are fabricated on a single wafer in parallel, and once complete the wafer is broken into the individual pieces for stacking. SEM images of the fabricated device is shown in Fig. 4.12. The connectivity algorithm described above yielded devices that were fully connected, but the mechanical robustness of the connected structures was not considered during the design. Fortunately, all devices were fabricated without any mechanical issues. Devices 1 and 8 featured thick tabs that were place in four locations in the event that mechanical issues arose, however these were deemed to be unnecessary after devices 2-7 were fabricated without issue.

4.5 Measurement and analysis

The measurements of the metasurface and 3D metaoptics devices are, at present, ongoing. This section serves as an update on the efforts and a look at the problems facing the measurement of these devices. The measurement setup is similar to the one shown in Chapter 3 Fig. 3.5, except multipliers are attached to the ports of the VNA to extend its range to the desired frequency range of 500-600 GHz. The multipliers, which are the black rectangular prisms in Fig. 4.13, have WR-1.5 waveguides which the metal-machined bottom block of the stack shown in Fig. 4.2 is designed to attach to. The second multiplier is placed on the scanning stage, and a diagonal horn (25 dB gain, Virginia Diodes) is attached to its WR 1.5 waveguide to make the radiation pattern more smooth and predictable.

Our observations thus far show that one polarization behaves substantially different than expected. Measurements performed with and without the metaoptics incorporated in the stack suggest that the current issues do not stem from the metaoptics themselves, but rather an issue with the radiated beam from the waveguide antenna

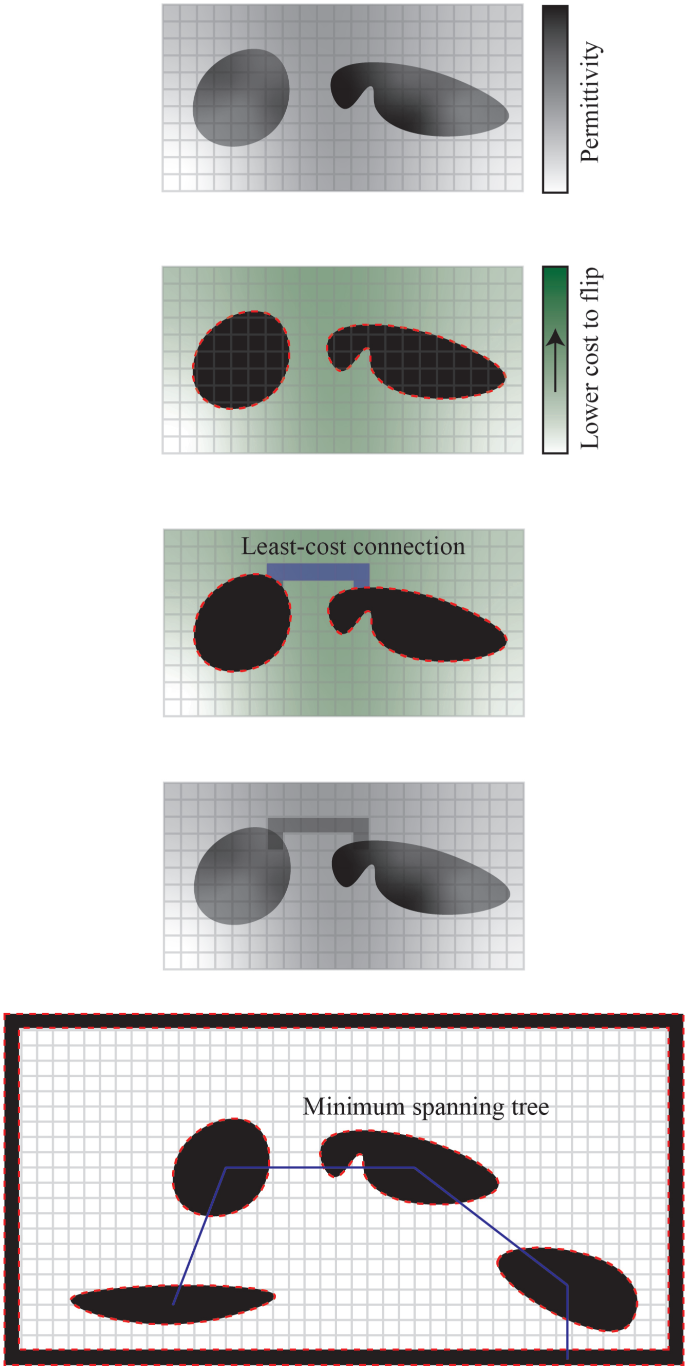

Figure 4.11: **Algorithm for fully connecting a structure during optimizations.** The device is forced to be connected at various points during the optimization, including when the permittivity is continuous. The connection procedure begins by rounding permittivity to their nearest material boundary. All connected regions are then identified (denoted by red dashed borders). The continuous permittivity value of the remaining pixels are used to define a cost function. with pixels closer to the mid-point permittivity considered low-cost. A least-cost path can be calculated using Djikstra's algorithm. When many disconnected regions are present, a minimum spanning tree can be used to connect all regions with the minimum possible cost.

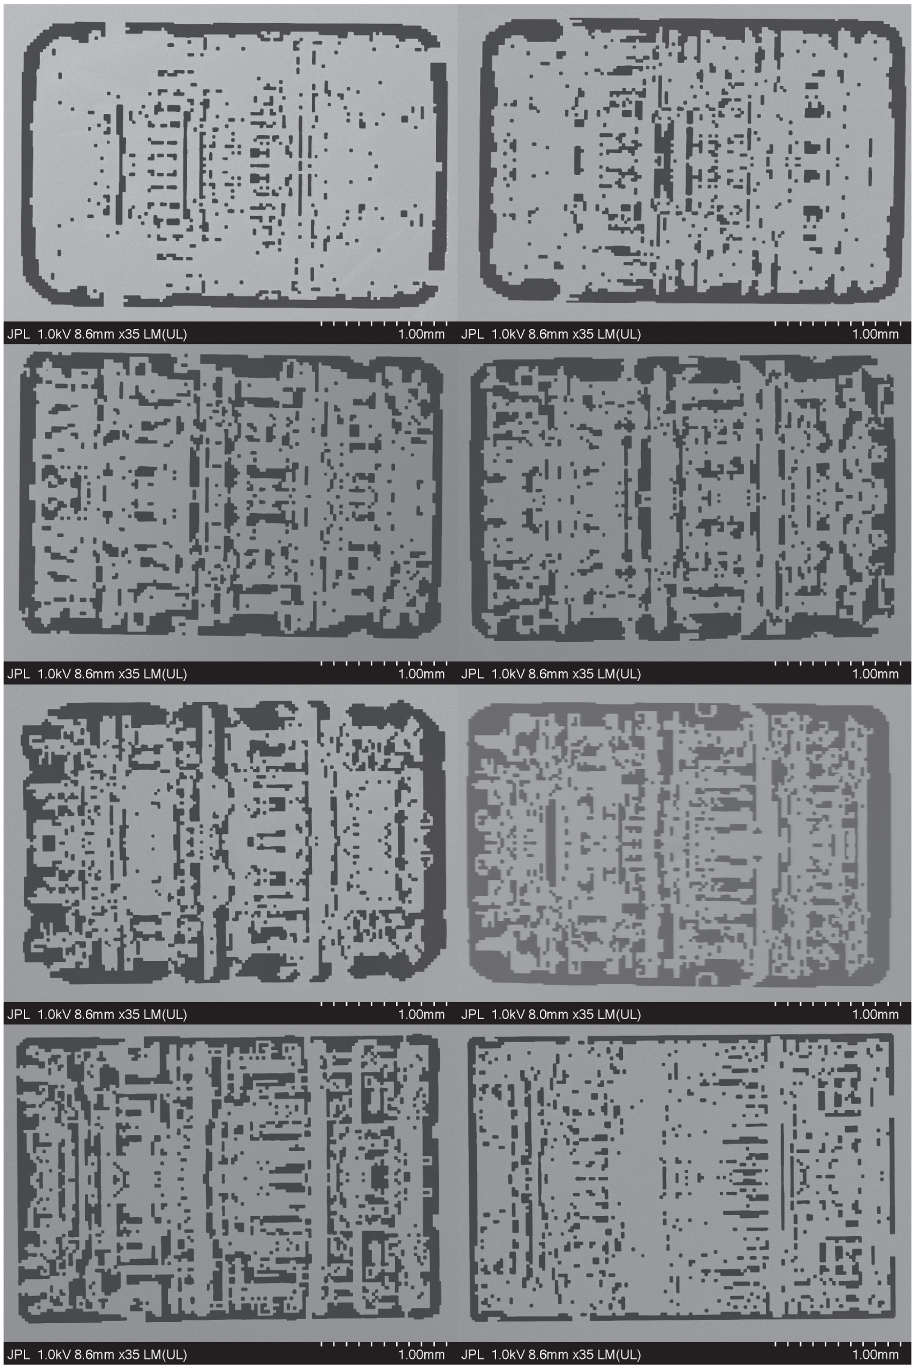

Figure 4.12: **Fabricated 3D THz devices.** The devices are fully connected and etched entirely through a 300 μm wafer. Although some of the layers are connected by thin bridges, there were no instances of mechanical failure in any samples.

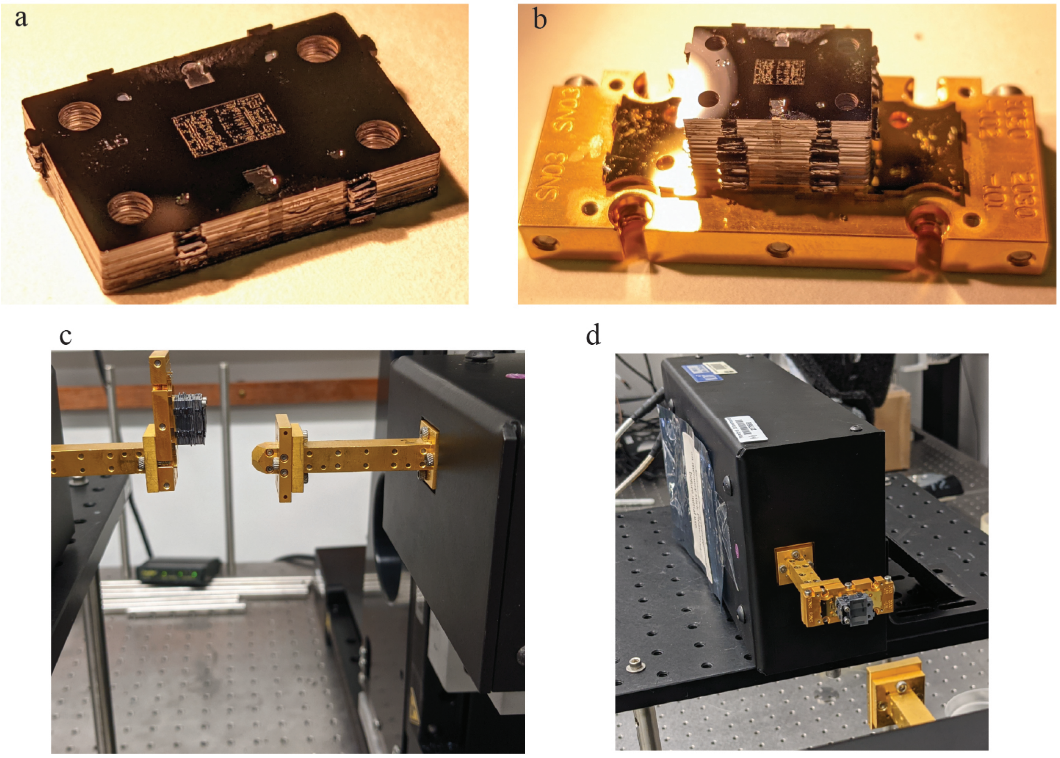

Figure 4.13: **Assembled THz device and measurement setup.** (a) The 8-layer 3D THz device fully assembled. (b) The THz device assembled with the rest of the wafer stack, including a waveguide adapter block used to attach the device to WR1.5 waveguides of a VNA. (c) The device (left) is attached to one THz multiplier, and a diagonal horn (right) is attached to the other THz multiplier. (d) One polarization component is rotated in the waveguide layers of the assembly. To measure this polarization, the multiplier block is rotated.

feeds. The primary issue is the spatial distribution of the beam is different than expected, and the light for one polarization is radiated at a large angle of approximately 21° (6.2° in Si when refraction is considered). This cannot stem from the single-mode waveguide layers, which act as a single-mode spatial filter. The problem also likely does not lie in the waveguide taper layer, which contains only 2 spatial modes with similarly restrictive spatial mode profiles. Therefore, it is most likely that an issue has arisen in the iris layer, Fabry-Perot cavity layer, or the intermediate blank Si layers. Since the spatial distribution was repeatable across multiple across multiple independently stacked assemblies, it is unlikely that the issue arose due to factors like debris accruing on the layers during the assembly, which would have introduced more randomness across samples than was observed here.

The measured fields for a metasurface device (waveguide spacing of 0.6D in Fig. 4.7) are shown in Fig. 4.14. Simulated field values are also shown for comparison.

These are obtained by simulating the propagation of the field profiles, then altering the phase of the fields by the ideal phase that the metasurface is designed to impose, then propagating the fields through the remaining homogeneous Si and air to the approximate location of the measurement probe. The location to the measurement plane is fine-tuned by maximizing the mode-overlap integral between the simulated and measured fields as a function of the simulated propagation distance in air.

The measured fields for the 8-layer 3D metaoptics device is shown in Fig. **4.15**, along with the fields that were measured when the 3D device was removed from the wafer stack. The x-polarized case (the problematic polarization) shows that there is very little difference between the fields output from the device and the fields from the incident beam on the device, which may indicate that the angled beam does not impinge on the device area substantially. The y-polarization is not substantially angled, but its mode profile differs substantially from the expected beam. Nonetheless the 3D device does indeed collimate the y-polarized fields to some extent, which is indicated by the higher intensity fields output by the 3D device relative to the incident beam.

It is difficult to quantify the efficiency of the metaoptics devices because evidently the beams incident upon them do not have the expected profile, and further work is required to diagnose why this occurs.

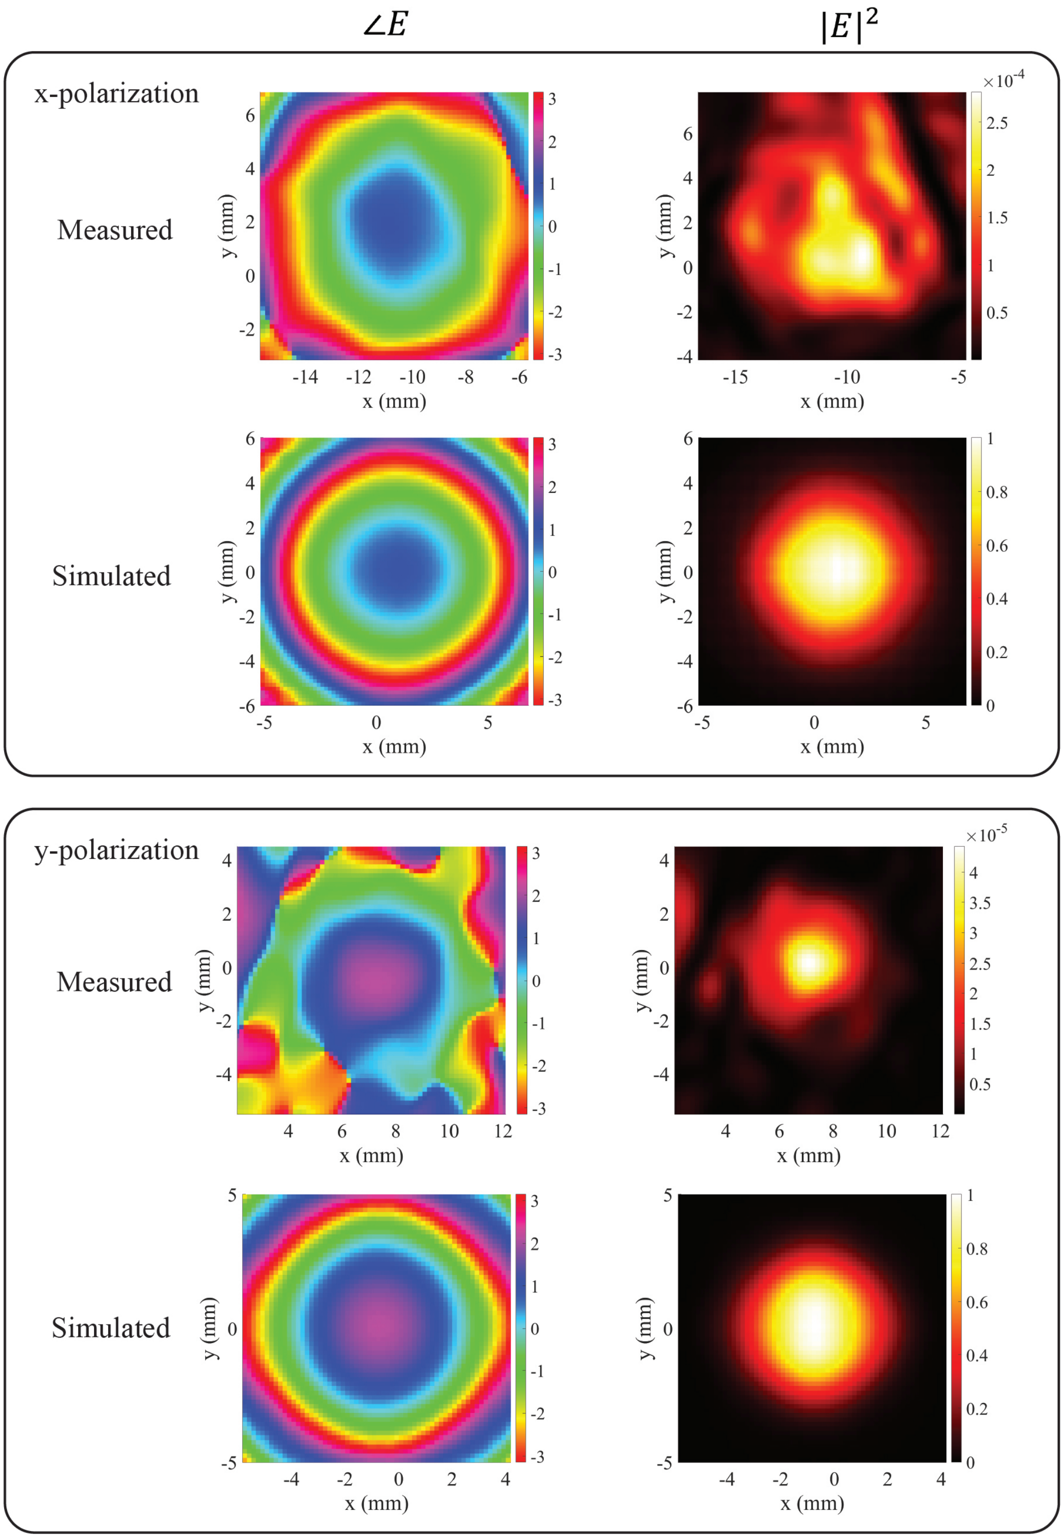

Figure 4.14: **THz measurement results for a metasurface device.** The metasurface is designed for focusing to cross-polarized waveguides separated by 0.6D (see Fig. 4.7). Simulated and measured fields are compared for both polarizations.

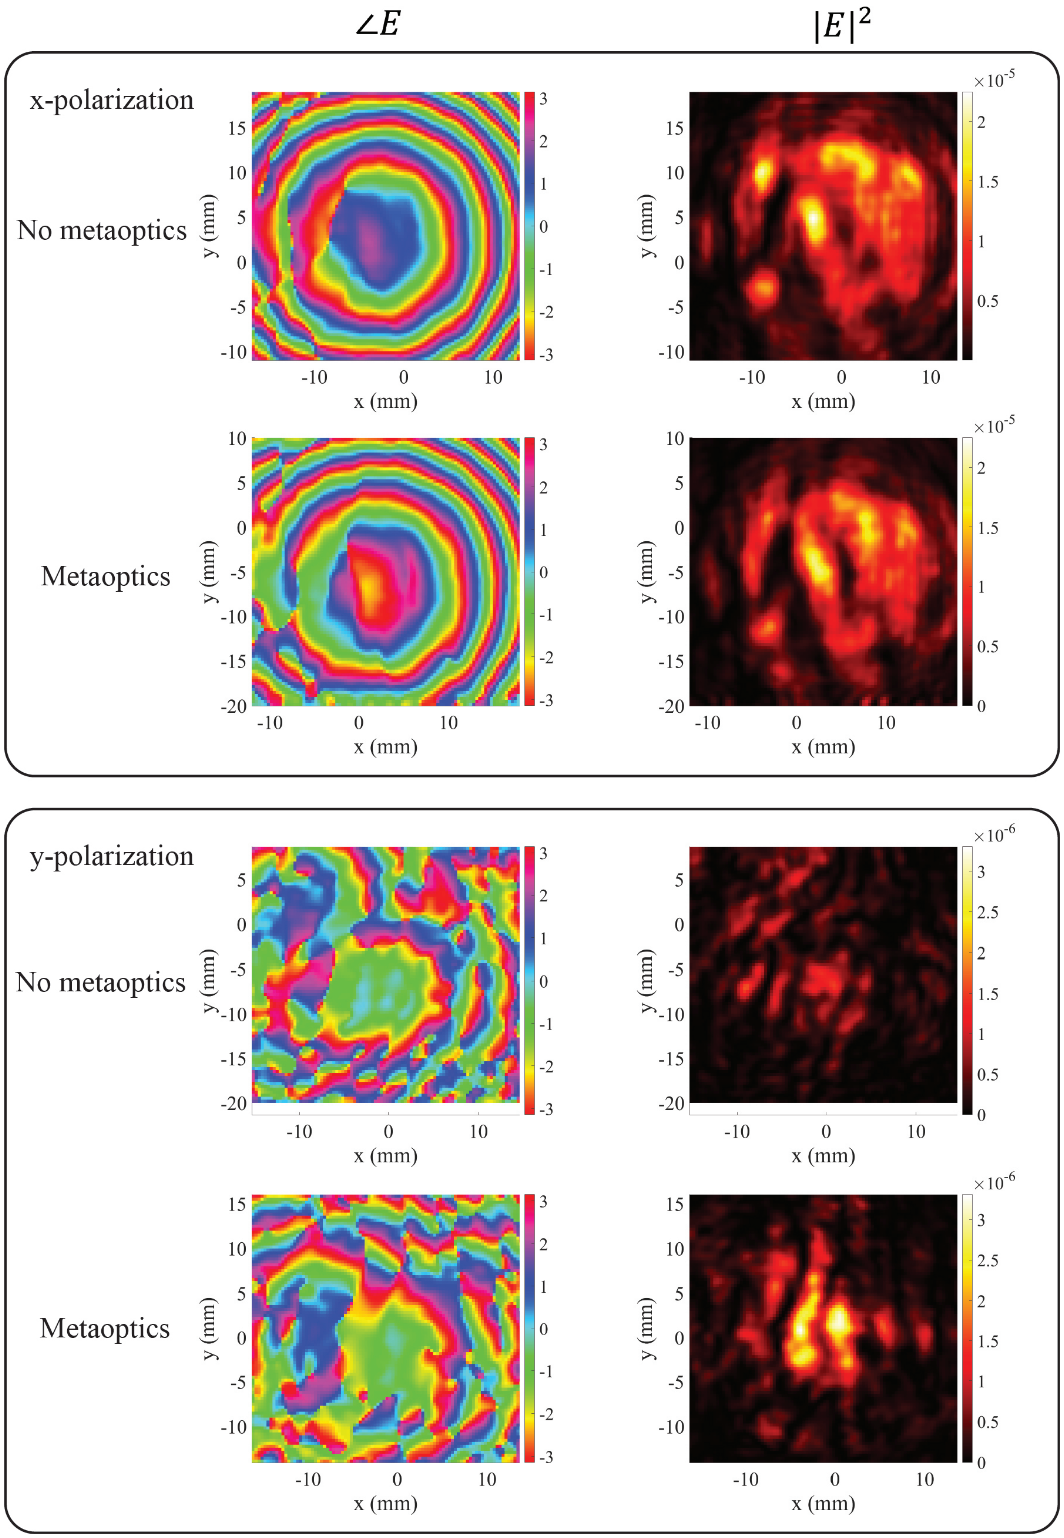

Figure 4.15: **THz measurement results for a metasurface device.** The fields that are compared for each polarization correspond to measurements made with and without the metaoptics in the wafer stack.

Chapter 5

ADVANCED APPLICATIONS IN THZ METAOPTICS

The 3D device described in Chapter 4 offers a large efficiency improvement over a metasurface, and even a large improvement over ideal phase masks which typically represent the best-case scenario for thin metasurface performance. While this comparison illustrates one advantage that is obtained from moving to 3D devices, there remains one important question: *what can 3D metastructure do that 2D metasurfaces cannot do?* This chapter describes several applications which have never been achieved with metasurfaces or compact systems of bulk components.

5.1 Direct-detector based spectral imaging

Technology background

The dominant technique for THz spectral imaging is coherent heterodyning-based detection. This is the process of mixing a broadband incoming signal, commonly referred to as the RF signal, with an extremely narrow-band local oscillator (LO) signal. The output of the mixer, denoted IF for *intermediate frequency*, can then be used to infer the intensity of the RF signal within the narrow spectral range of the LO signal. With decades of improvements steering this technology, LO signals can have excellent spectral resolutions ($> 10^5$) over moderate bandwidths (20% fractional bandwidth). However, the LO requires substantial power to generate and the overall systems tend to be fairly large compared to the 3D metaoptics described

The alternative technology to heterodyne detector systems are ones based on incoherent "direct" detectors. These detectors can be extremely sensitive, but do not distinguish between different incoming frequencies. They can be used to measure the intensity of an incoming signal integrated over a large spectral band, or can be coupled with a different element that ensures the signal being passed to the direct detector contains only frequencies within a desired range for spectroscopy. An example would include coupling high quality factor (Q-factor) outputs from a resonator to a direct detector, and possibly sweeping the resonant frequency to scan over some spectral range. These systems can be extremely low power, since the resonator element does not require power beyond whatever mechanism is used to scan its resonant frequency. However, the resonators are often seriously limited in bandwidth. There is a typical, but *not* fundamental, trade-off that exists for these systems in which improving the Q-factor of the resonances, and thus the spectral resolution of the system, reduces the free spectral range (FSR) of the resonator. The FSR is the quantity that measures the spectral distance between the resonant peaks, which determines the range over which a measured signal can be unambiguously assigned to one spectral peak. Taking a Fabry-Perot cavity as an example resonator, the Q-factor of a resonance increases with mirror reflectivity and with cavity length. The FSR decreases with cavity length. Thus, increasing cavity length introduces the trade-off between spectral resolution and bandwidth.

Ultimately, heterodyning tends to hold the advantage for current space-based instrumentation which typically demand both high spectral resolution and high bandwidth. Some direct-detector spectrometers have obtained extremely good spectral sensitivity ($Q > 10^6$) over a small bandwidth of about 311 MHz [89]. The fact that high Q-factors can be obtained is promising, and the challenge of expanding the bandwidths of these types of device is a task particular well-suited for the 3D inverse-designed spectral splitters that were described in Chapter 3.

Proposed solution

The proposed solution involves placing a 3D metaoptics device at the output of a Fabry-Perot resonator. High Q-factor resonances transmitted through the cavity and input into an inverse-design 3D metaoptics device. The metaoptics are designed to focus spectral bins that are equal in size to the FSR of the cavity to different points in its focal plane, which contains an array of sensitive direct detectors. Thus, the bandwidth of the resonator is expanded from one FSR to however many FSR bins are focused by the metaoptics device. As a prototype, the 3-bin 42% fractional bandwidth devices that were demonstrated in both Ch. 2 and Ch. 3 is used, and are scaled to operate from 440 GHz to 660 GHz.

The resonator is made of two distributed Bragg reflector (DBR) mirrors made of Si and air. DBR mirrors are broadband, dielectric mirrors made by cascading multiple period of $\lambda/4$ layers. The bandwidth over which the mirror reflectivity is high is determined by the refractive index contrast between the two materials, and in the case of Si and air an extremely large bandwidth extending from 340 GHz to 740 GHz. With only 5 Si-air period, the reflectivity of a normally incident planewave exceed 99.998% at 550 GHz.

The DBR mirrors can be fabricated using the procedures outline in Ch. 4. As was done in the iris fabrication shown in Fig. 4.5, an SOI wafer can be used to ensure the

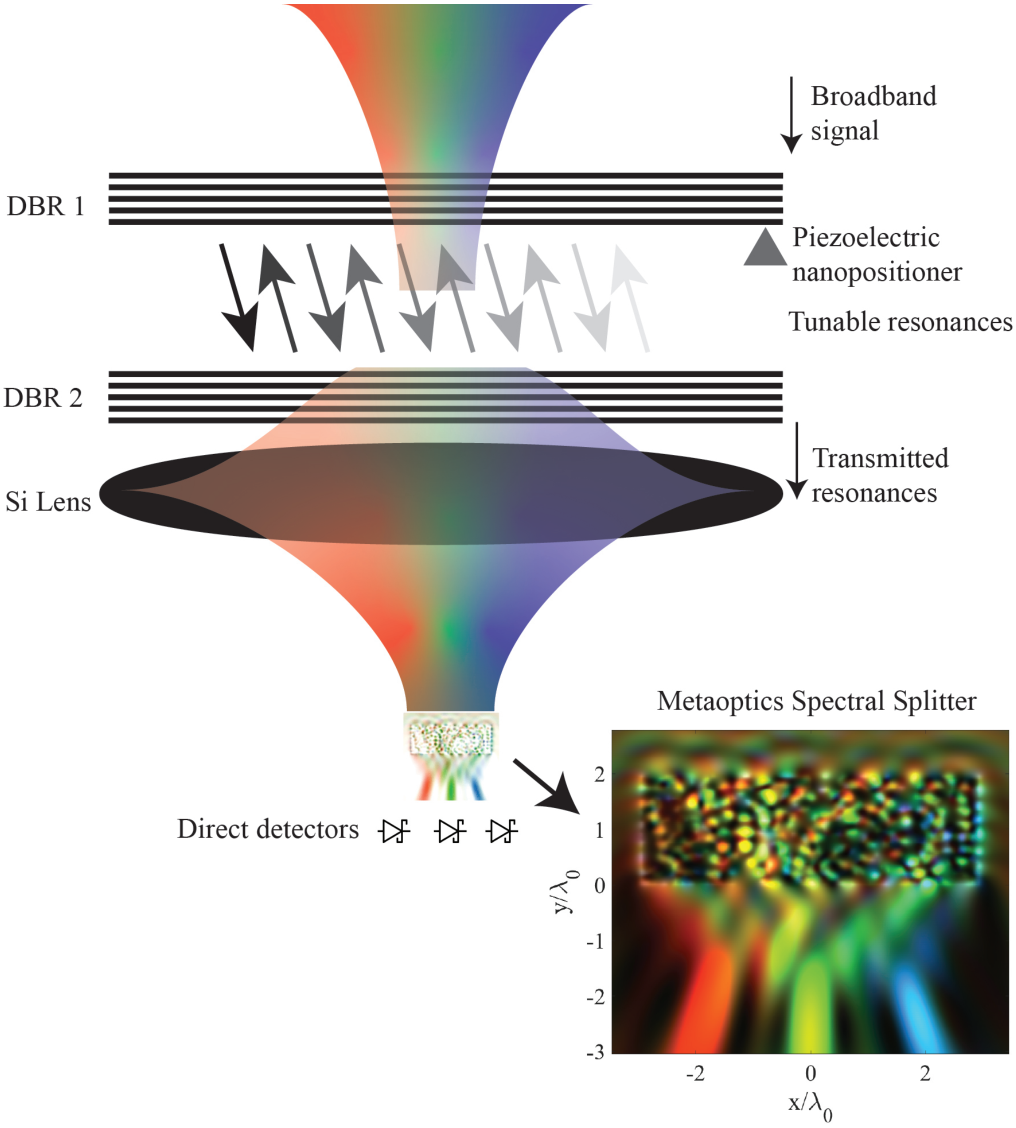

Figure 5.1: **Illustration of the proposed THz spectrometer working principles.** A broadband signal is passed through a Fabry-Perot resonator, and the output is spatially larger beam that contains only the resonant frequencies. The output is focused through a Si lens onto a metaoptics device that sorts the different resonant frequencies to direct detectors. The spatial separation of the resonant frequencies is key to eliminating the need for coherent detection through mixing the signal with an LO signal.

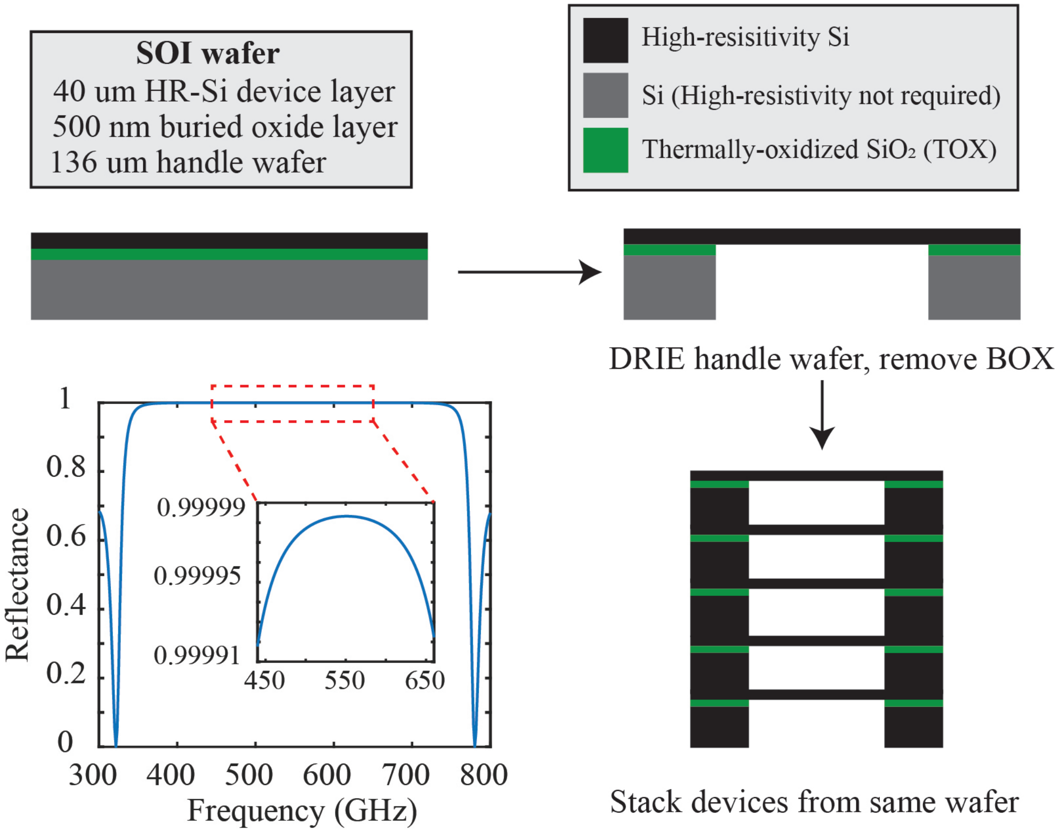

Figure 5.2: **Design and fabrication of Si-air DBR mirrors.** The fabrication is done by patterning an SOI wafer. The air portion of the DBR is etched through the handle wafer, and the BOX. The device are then stacked. A 5-period stack designed for a center frequency of 550 GHz has a reflectance of nearly 99.999% at 550 GHz in theory.

the fabricated layer heights are extremely accurate. A suitable fabrication procedure would involve procuring a SOI wafer with a 39 μm HR-Si device layer ($\lambda/4$ at 550 GHz in Si) and a 136 μm handle wafer ($\lambda/4$ at 550 GHz in air). The handle wafer can then be etched through to the BOX layer. Pin pockets can be fabricated on both sides to enable accurate stacking if integrated into a larger system.

The Fabry-Perot cavity is made by cascading two well-aligned DBR mirrors with a small air gap in between. The transmission through the cavity is maximized for frequencies that satisfy the resonance condition. In the absence of absorption, the max transmission is 1 if both mirrors have the same reflectivity. The reflectivity of the mirrors can be well-matched by fabricating both mirrors on the same wafer. To scan the Fabry-Perot cavity, highly sensitive and low-power piezo-electric motors (nanopositioners) can be used. These nanopositioners claim to feature sub-nm positioning resolution, as well as positioning feedback to ensure convergence to position set-points. A resonance can be scanned repeatedly using these nanopositioners to continuously scan a desired spectral line.

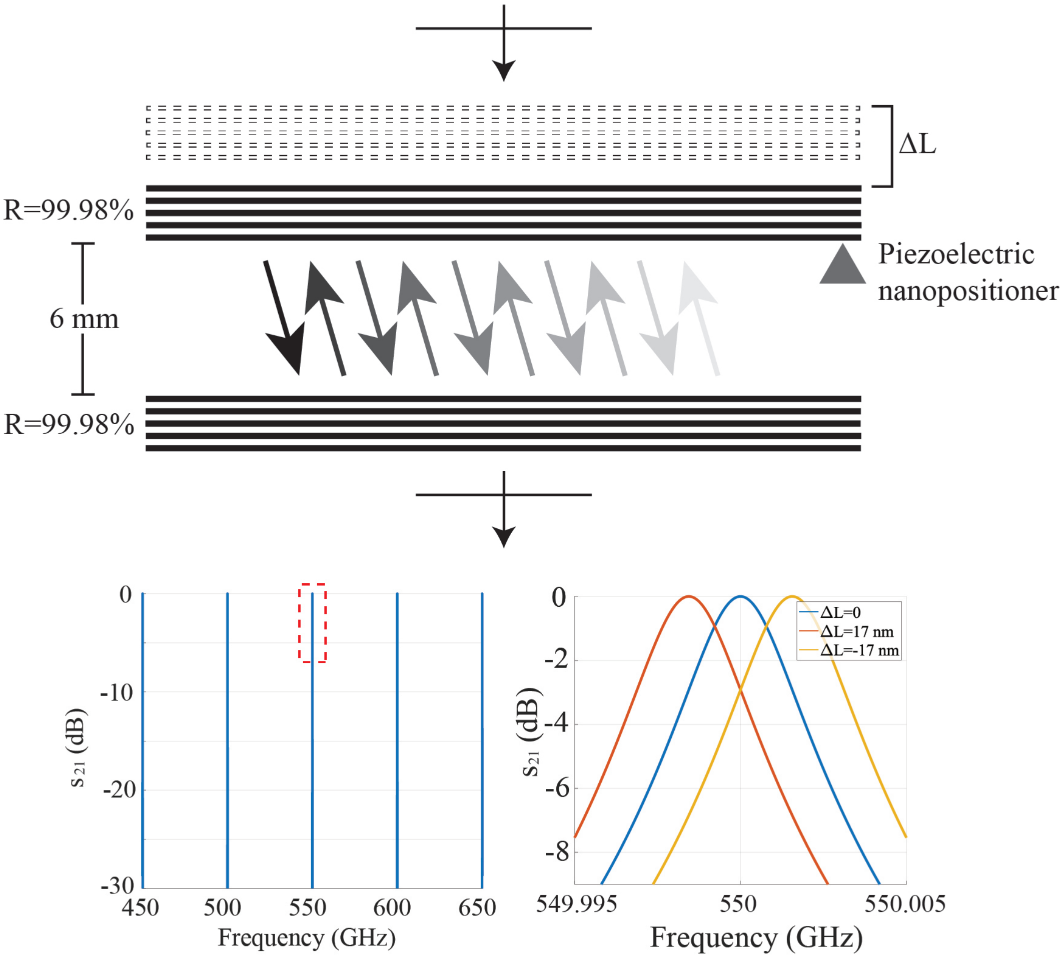

Figure 5.3: **Design and fabrication of DBR cavity.** (Top) The cavity is comprised of a 6 mm air gap between two 99.98% relective DBR mirrors. (Bottom-left) The transmission of resonant frequencies is 1, and the frequency separation (FSR) between peaks is 50 GHz. (Bottom-right) By tuning the length of the cavity with a piezoelectric nanopositioner the resonance can be shifted to scan the full spectrum.

One drawback to this design is that the mirrors are planar. In general, cavity Q-factors suffer without some guiding mechanism. This can be done by coupling power into a waveguide, like was done in [89], or by curving one of the mirrors of the Fabry-Perot cavity such that the beam is re-collimated with every round trip. Alternatively, the beam waist and Fabry-Perot cavity can be chosen to simply be large enough to minimize its divergence angle. FDTD simulations of Gaussian beam propagation through a Fabry-Perot cavity (5-layer DBR mirrors, 6-mm air cavity length) indicate that Q-factors of over 10^5 can be achieved for a 3-mm beam waist radius Gaussian beam.

To decouple the lateral size of the beam and cavity from the size of the metaoptics device, a lens can be placed at the output of the cavity that focuses the output fields onto the metaoptics device. The exact spatial distribution of the fields output by the cavity can be measured and incorporated into the inverse-design of the metaoptics device. The metaoptics device can then be designed using procedures detailed throughout this thesis, particularly like the design shown in Section 4.4. The device in Ch. 2 Fig. **2.5**c is very similar to the device proposed here.

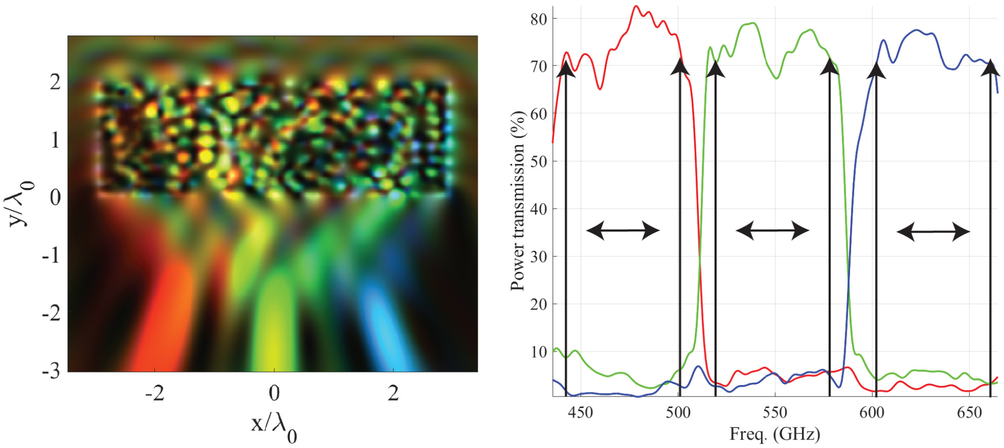

Figure 5.4: **Design of metaoptics device for expanding Fabry-Perot spectrometer bandwidth.** The metaoptics device focuses different spectral bands to different locations in a focal plane. The size of the bands can be chosen to be equal to the FSR of the Fabry-Perot cavity, thus sorting individual resonant frequencies to their own detectors. This effectively multiplies the usable bandwidth of a traditional Fabry-Perot spectrometer by the number of individual bins the metaoptics can be optimized for.

Is the resonant cavity required?

The DBR mirrors, cavity, and metaoptics are all comprised entirely of Si and air. The entire system is therefore a subset of the solutions that can be found with inverse-design, and thus there is a possibility that the inverse-design process may be able to find a better if we simply take the entire region containing all elements as the design region. However, this approach is current infeasible given the current limitations on computing power. The scale of such a problem is vastly larger than any structures that were optimized in this thesis. The geometric scale is rather large already, but the overall structure is required to be highly resonant which makes the FDTD simulations extremely long. This can potentially be overcome by switching to finite-difference frequency-domain (FDFD) simulations, but since the system must be optimized over a very large bandwidth many FDFD simulations would need to be conducted over that full bandwidth.

Thus, applying inverse-design to the entire proposed structure remains difficult. For now the creation of complex structures such as this still requires the integration of modular components. However, the 3D metaoptics offers a substantial miniaturization of these types of systems. The ability to optimize the device under exactly the electromagnetic field profiles that are output from the cavity is an advantage that is extremely difficult to achieve with other design methodologies, and the optimization over a large bandwidth is a task that 2D metasurface design has historically struggled to do.

Chapter 6

A FORWARD-LOOKING CONCLUSION

The devices described in this thesis are small dielectric devices capable of controlling all of the fundamental properties of light: wavelength, polarization, and k-vector. Their extreme versatility stems from their thickness, which compared to metasurfaces allows for more complex electromagnetic scattering within the device volume, and an inverse-design procedure that can identify a 3D refractive index distribution that adequately achieves the prescribed optical functionalities. The extreme complexity of the problem at hand requires that the devices be relatively small, around a few wavelengths per side, and thus the relevant applications tend to focus on arrays of these devices placed over imaging arrays.

Chapter 2 offered an in-depth discussion of how these devices are designed. The procedure uses gradient-based techniques to slowly perturb the dielectric structure to optimize a set of arbitrary electromagnetic FOM. While gradient-based optimization strategies are employed across practically all fields of engineering, physics, and mathematics, the optimizations here are highly multi-dimensional and thus it is impractical to compute the required gradients with "brute force" techniques. Instead, a technique referred to as the adjoint-state method is used to efficiently compute the gradient.

Chapter 2 also describes a technique that was devised to make the optimization of high-index structures more tractable. A common strategy for optimizing photonic structures begins with a fictitious grey-scale permittivity at the midpoint of two material permittivities. Guided by the computed gradients, this continuous permittivity is perturbed to improve performance. When the device is sufficiently binary, the continuous permittivity is rounded to its nearest material boundary, and the optimization switches over to a level-set based optimization that perturbs the shape of the material boundaries. A problem that occurs in this approach is that the continuous phase may not converge to a solution that closely resembles a binary structure, which is often required for the subsequent level-set optimization to converge to a good solution. This is particularly problematic when the two materials have a high-index contrast. The procedure laid out in Section 2.4 addresses this by solving a sub-optimization problem at every iteration of the overall photonic optimization.

The sub-optimization uses the gradient information to find the permittivity step that maximizes the change in performance of the device while constraining the solutions to only be ones that drive the device to a more binary solutions.

An intriguing prospect for future work is to replace the binarization constraint with different constraints, such that at every iteration of the optimization some parameter is forced to be increased or decreased. One example that could be tested is minimizing a fabrication penalty term during the continuous phase of the optimization. Another direction could be to treat some electromagnetic FoMs differently with this constraint — for example, the spectral splitter example of Section 2.4 could include a FOM for the overall transmission of the entire device which is forced to increase at every step of the optimization. Such a technique could prove advantageous for explicitly prioritizing one set of FoMs over another set.

During the course of this thesis multiple microwave prototypes were measured, which are described in Ch. 3. The devices measured in Section 3.1 and Section 3.2 certainly have the potential to improve optical technology, but the forward-looking lesson from these measurements is that the complex devices that were designed with extensive simulation behave as expected in reality. This allows us to design more ambitious devices with the confidence that accurate fabrication procedures will ensure the device behaves as expected in reality. Thus, Ch. 3 further describes a device that was studied only in simulation that assumes our fabrication capabilities mimic that of modern day CMOS processes used in microelectronics (see Section 3.3). This device is capable of sorting light based on its incident angle, wavelength, and polarization. Furthermore, it can be tiled across a sensor array (similar to current color filter arrays) to make highly multifunctional image sensors. In the context of the history of wavefront-sensing arrays, this device is a suitable successor to the microlens arrays, which in turn were successors to arrays of slits placed over an image sensor.

While the simulated device's fabrication tolerances were chosen to emulate that of current CMOS capabilities, the material system does not and this is potentially the largest obstacle in development. CMOS processes are used in microelectronics for encapsulating metallic wires in a dielectric, and decades of research and developments have pushed feature sizes down. Unfortunately, metals are not a good material for optical devices due to their high loss, so for the proposed device the metal is replaced with TiO2. The transition to a new material system is likely not limited by physics, but rather a matter of allocating time and resources to creating all-dielectric foundry processes with <100 nm resolutions.

In the meantime, we can move to larger wavelengths to realize novel devices at infrared and THz frequencies. In particular, THz devices can be fabricated by stacking individually patterned silicon wafers, which is described in Ch. 4. The stack has extremely accurate layer-to-layer alignment achieved by locking the layers together with compression pins, which are lithographically defined with excellent resolution. Furthermore, the individual layers of the device are patterned with deeply subwavelength features thanks to highly developed Si micromachining techniques. These fabrication techniques were used to realize highly efficient polarization splitters that can outperform ideal metasurfaces.

Ch. 5 details one of the exciting prospects of 3D THz devices — spectroscopy. THz spectrometers are critical in the study of planets and astrophysics due to the presence of key spectral lines. THz spectrometers tend to utilize coherent heterodyne-based detection, which are bulky and require LO sources that are limited in bandwidth. The inverse-designed devices described in this thesis are perfectly suited for improving THz spectrometers since they are small and passive, yet provide a huge amount of functionality and versatility including spectral demultiplexing.

One basic spectrometer consists of a resonator, such as a Fabry-Perot cavity, whose resonances are scanned to trace out the spectrum of the incident light. These resonators can have extremely high resolutions, but are limited in bandwidth by the cavity FSR. The novel spectral demultiplexing elements described in Ch. 2 and Ch. 3 can be cascaded with a high-Q-factor resonator to increase the bandwidth by integer multiples of the cavity FSR. Such an application, if successful in its implementation, would yield THz spectrometers that are smaller, consume less power, and easily integrated into multi-pixel arrays by distributing devices across a wafer.

Multifunctional volumentric metaoptics have the capability of vastly improving technology from microwave to optical frequencies. As more time and resources are invested in this exciting new area of photonics, the fabrication of multi-layer structures with subwavelength feature sizes and alignments in new materials systems will ultimately yield devices operating at smaller and smaller wavelengths. Currently, the THz regime is an exciting area with both scientific value and available fabrication techniques. Exciting applications exist for infrared and optical regimes which this thesis only began to touch on, such as highly multifunctional camera sensors capable of lightfield imaging with simultaneous spectral and polarization characterization. The realization of devices with the requisite feature sizes is hopefully a matter of time and resource allocation, rather than a limitation imposed by physics, as evidenced by the advancements in CMOS fabrication techniques over several decades.

BIBLIOGRAPHY

[1] W. Stork, N. Streibl, H. Haidner, and P. Kipfer, “Artificial distributed-index media fabricated by zero-order gratings”, Opt. Lett. **16**, 1921–1923 (1991) `DOI`: `10.1364/OL.16.001921` (cited on p. 2).

[2] M. E. Warren, R. E. Smith, G. A. Vawter, and J. R. Wendt, “High-efficiency subwavelength diffractive optical element in GaAs for 975 nm”, Opt. Lett. **20**, 1441–1443 (1995) `DOI`: `10.1364/OL.20.001441` (cited on p. 2).

[3] F. T. Chen and H. G. Craighead, “Diffractive phase elements based on two-dimensional artificial dielectrics”, Opt. Lett. **20**, 121–123 (1995) `DOI`: `10.1364/OL.20.000121` (cited on p. 2).

[4] F. T. Chen and H. G. Craighead, “Diffractive lens fabricated with mostly zeroth-order gratings”, Opt. Lett. **21**, 177–179 (1996) `DOI`: `10.1364/OL.21.000177` (cited on p. 2).

[5] P. Lalanne, S. Astilean, P. Chavel, E. Cambril, and H. Launois, “Blazed binary subwavelength gratings with efficiencies larger than those of conventional échelette gratings”, Optics Letters **23**, 1081 (1998) `DOI`: `10.1364/ol.23.001081` (cited on p. 2).

[6] N. Yu, P. Genevet, M. A. Kats, F. Aieta, J.-P. Tetienne, F. Capasso, and Z. Gaburro, “Light Propagation with Phase Discontinuities: Generalized Laws of Reflection and Refraction”, Science **334**, 333–337 (2011) `DOI`: `10.1126/science.1210713` (cited on p. 2).

[7] N. Yu and F. Capasso, “Flat optics with designer metasurfaces”, **13**, 139–150 (2014) `DOI`: `10.1038/nmat3839` (cited on pp. 2, 47).

[8] A. Arbabi, R. M. Briggs, Y. Horie, M. Bagheri, and A. Faraon, “Efficient dielectric metasurface collimating lenses for mid-infrared quantum cascade lasers”, Optics Express **23**, 33310 (2015) `DOI`: `10.1364/oe.23.033310`, arXiv:`1511.08127` (cited on pp. 2, 58).

[9] A. Arbabi, Y. Horie, M. Bagheri, and A. Faraon, “Dielectric metasurfaces for complete control of phase and polarization with subwavelength spatial resolution and high transmission”, Nature Nanotechnology **10**, 937–943 (2015) `DOI`: `10.1038/nnano.2015.186`, arXiv:`1411.1494` (cited on pp. 2, 47).

[10] A. Arbabi and A. Faraon, “Fundamental limits of ultrathin metasurfaces”, Scientific Reports **7**, 1–9 (2017) `DOI`: `10.1038/srep43722`, arXiv:`1411.2537` (cited on pp. 3, 47, 91).

[11] A. Arbabi, E. Arbabi, Y. Horie, S. M. Kamali, and A. Faraon, “Planar metasurface retroreflector”, Nature Photonics **11**, 415–420 (2017) `DOI`: `10.1038/nphoton.2017.96` (cited on p. 3).

[12] M. Faraji-Dana, E. Arbabi, H. Kwon, S. M. Kamali, A. Arbabi, J. G. Bartholomew, and A. Faraon, "Hyperspectral Imager with Folded Metasurface Optics", ACS Photonics **6**, 2161–2167 (2019) DOI: `10.1021/acsphotonics.9b00744` (cited on pp. 3, 47).

[13] S. Colburn, A. Zhan, and A. Majumdar, "Metasurface optics for full-color computational imaging", Science Advances **4**, DOI: 10.1126/sciadv.aar2114 (2018) DOI: `10.1126/sciadv.aar2114` (cited on p. 3).

[14] J. N. Mait, G. W. Euliss, and R. A. Athale, "Computational imaging", Advances in Optics and Photonics **10**, 409 (2018) DOI: `10.1364/aop.10.000409` (cited on pp. 3, 70).

[15] Z. Lin, C. Roques-Carmes, R. Pestourie, M. Soljačić, A. Majumdar, and S. G. Johnson, "End-to-end nanophotonic inverse design for imaging and polarimetry", Nanophotonics **10**, 1177–1187 (2021) DOI: `doi:10.1515/nanoph-2020-0579` (cited on p. 3).

[16] J. W. Goodman, "Introduction to fourier optics", Introduction to Fourier optics, 3rd ed., by JW Goodman. Englewood, CO: Roberts & Co. Publishers, 2005 **1** (2005) (cited on p. 7).

[17] Y. Zhao and A. Alù, "Manipulating light polarization with ultrathin plasmonic metasurfaces", Phys. Rev. B **84**, 205428 (2011) DOI: `10.1103/PhysRevB.84.205428` (cited on p. 11).

[18] F. Ding, A. Pors, and S. I. Bozhevolnyi, "Gradient metasurfaces: a review of fundamentals and applications", Reports on Progress in Physics **81**, 26401 (2017) DOI: `10.1088/1361-6633/aa8732` (cited on p. 11).

[19] F. Ding, Y. Yang, R. A. Deshpande, and S. I. Bozhevolnyi, "A review of gap-surface plasmon metasurfaces: fundamentals and applications", Nanophotonics **7**, 1129–1156 (2018) DOI: `doi:10.1515/nanoph-2017-0125` (cited on p. 11).

[20] E. Arbabi, A. Arbabi, S. M. Kamali, Y. Horie, and A. Faraon, "Controlling the sign of chromatic dispersion in diffractive optics with dielectric metasurfaces", Optica **4**, 625–632 (2017) DOI: `10.1364/OPTICA.4.000625` (cited on p. 14).

[21] O. Miller, "Photonic design: from fundamental solar cell physics to computational inverse design", PhD thesis (EECS Department, University of California, Berkeley, 2012) (cited on p. 16).

[22] D. Givoli, "A tutorial on the adjoint method for inverse problems", Computer Methods in Applied Mechanics and Engineering **380**, 113810 (2021) DOI: `10.1016/j.cma.2021.113810` (cited on p. 17).

[23] C. M. Lalau-Keraly, S. Bhargava, O. D. Miller, and E. Yablonovitch, "Adjoint shape optimization applied to electromagnetic design", Optics Express **21**, 21693 (2013) DOI: `10.1364/oe.21.021693` (cited on pp. 17, 47, 62, 89).

[24] O. Sigmund and K. Maute, "Topology optimization approaches: A comparative review", Structural and Multidisciplinary Optimization **48**, 1031–1055 (2013) DOI: `10.1007/s00158-013-0978-6` (cited on pp. 18, 28).

[25] A. Y. Piggott, J. Lu, K. G. Lagoudakis, J. Petykiewicz, T. M. Babinec, and J. Vucković, "Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer", Nature Photonics **9**, 374–377 (2015) DOI: `10.1038/nphoton.2015.69`, arXiv:`1504.00095` (cited on pp. 19, 47).

[26] S. Boyd and L. Vandenberghe, *Convex optimization* (Cambridge University Press, 2004), DOI: `10.1017/CBO9780511804441` (cited on p. 21).

[27] C. Ballew, G. Roberts, S. Camayd-Muñoz, M. F. Debbas, and A. Faraon, "Mechanically reconfigurable multi-functional meta-optics studied at microwave frequencies", Scientific Reports **11**, 11145 (2021) DOI: `10.1038/s41598-021-88785-5`, (cited on pp. 24, 62).

[28] C. Ballew, "Thesis - multifunctional volumetric metaoptics - section 2.4 supplementary video", `DOI: 10.22002/D1.2082` (2021) DOI: `10.22002/D1.2082` (cited on pp. 25, 29).

[29] N. Lebbe, C. Dapogny, E. Oudet, K. Hassan, and A. Gliere, "Robust shape and topology optimization of nanophotonic devices using the level set method", Journal of Computational Physics **395**, 710–746 (2019) DOI: `10.1016/j.jcp.2019.06.057` (cited on pp. 28, 38, 63).

[30] D. Vercruysse, N. V. Sapra, L. Su, R. Trivedi, and J. Vučković, "Analytical level set fabrication constraints for inverse design", Scientific Reports **9**, 1–7 (2019) DOI: `10.1038/s41598-019-45026-0` (cited on pp. 33, 38, 62, 63).

[31] E. W. Wang, D. Sell, T. Phan, and J. A. Fan, "Robust design of topology-optimized metasurfaces", Opt. Mater. Express **9**, 469–482 (2019) DOI: `10.1364/OME.9.000469` (cited on p. 34).

[32] P. Camayd-Muñoz, C. Ballew, G. Roberts, and A. Faraon, "Multifunctional volumetric meta-optics for color and polarization image sensors", Optica **7**, 280 (2020) DOI: `10.1364/optica.384228`, (cited on pp. 40, 48, 70).

[33] S. M. Kamali, E. Arbabi, A. Arbabi, and A. Faraon, *A review of dielectric optical metasurfaces for wavefront control*, 2018, DOI: `10.1515/nanoph-2017-0129`, arXiv:`1804.09802` (cited on p. 47).

[34] D. A. B. Miller, "Fundamental limit for optical components", Journal of the Optical Society of America B **24**, A1 (2007) DOI: `10.1364/josab.24.0000a1` (cited on p. 47).

[35] A. Arbabi, E. Arbabi, Y. Horie, S. M. Kamali, and A. Faraon, "Planar metasurface retroreflector", Nature Photonics **11**, 415–420 (2017) DOI: `10.1038/nphoton.2017.96` (cited on p. 47).

[36] B. Groever, W. T. Chen, and F. Capasso, "Meta-lens doublet in the visible region", Nano Letters **17**, 4902–4907 (2017) DOI: `10.1021/acs.nanolett.7b01888` (cited on p. 47).

[37] J. S. Jensen and O. Sigmund, *Topology optimization for nano-photonics*, 2011, DOI: `10.1002/lpor.201000014` (cited on p. 47).

[38] J. S. Jensen and O. Sigmund, "Systematic design of photonic crystal structures using topology optimization: Low-loss waveguide bends", Applied Physics Letters **84**, 2022–2024 (2004) DOI: `10.1063/1.1688450` (cited on p. 47).

[39] P. I. Borel, B. Bilenberg, L. H. Frandsen, T. Nielsen, J. Fage-Pedersen, A. V. Lavrinenko, J. S. Jensen, O. Sigmund, and A. Kristensen, "Imprinted silicon-based nanophotonics", Optics Express **15**, 1261 (2007) DOI: `10.1364/oe.15.001261` (cited on p. 47).

[40] T. W. Hughes, M. Minkov, I. A. Williamson, and S. Fan, "Adjoint Method and Inverse Design for Nonlinear Nanophotonic Devices", ACS Photonics **5**, 4781–4787 (2018) DOI: `10.1021/acsphotonics.8b01522`, arXiv:`1811.01255` (cited on p. 47).

[41] D. Sell, J. Yang, S. Doshay, R. Yang, and J. A. Fan, "Large-Angle, Multi-functional Metagratings Based on Freeform Multimode Geometries", Nano Letters **17**, 3752–3757 (2017) DOI: `10.1021/acs.nanolett.7b01082` (cited on p. 47).

[42] R. Pestourie, C. Pérez-Arancibia, Z. Lin, W. Shin, F. Capasso, and S. G. Johnson, "Inverse design of large-area metasurfaces", Optics Express **26**, 33732 (2018) DOI: `10.1364/oe.26.033732`, arXiv:`1808.04215` (cited on p. 47).

[43] A. Zhan, R. Gibson, J. Whitehead, E. Smith, J. R. Hendrickson, and A. Majumdar, "Controlling three-dimensional optical fields via inverse Mie scattering", Science Advances **5**, eaax4769 (2019) DOI: `10.1126/sciadv.aax4769` (cited on p. 47).

[44] Y. Deng and J. G. Korvink, "Topology optimization for three-dimensional electromagnetic waves using an edge element-based finite-element method", Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **472**, 20150835 (2016) DOI: `10.1098/rspa.2015.0835` (cited on p. 48).

[45] M. Mansouree, H. Kwon, E. Arbabi, A. McClung, A. Faraon, and A. Arbabi, "Multifunctional 25D metastructures enabled by adjoint optimization", Optica **7**, 77 (2020) DOI: `10.1364/optica.374787` (cited on p. 48).

[46] H. Chung and O. D. Miller, "High-NA achromatic metalenses by inverse design", Optics Express **28**, 6945 (2020) DOI: `10.1364/oe.385440`, arXiv:`1905.09213` (cited on p. 48).

[47] K. Iwami, C. Ogawa, T. Nagase, and S. Ikezawa, “Demonstration of focal length tuning by rotational varifocal moiré metalens in an ir-A wavelength”, Optics Express **28**, 35602 (2020) DOI: `10.1364/oe.411054` (cited on p. 48).

[48] A. Zhan, S. Colburn, C. M. Dodson, and A. Majumdar, “Metasurface Freeform Nanophotonics”, Scientific Reports **7**, 1–9 (2017) DOI: `10.1038/s41598-017-01908-9` (cited on p. 48).

[49] Z. Han, S. Colburn, A. Majumdar, and K. F. Böhringer, “MEMS-actuated metasurface Alvarez lens”, Microsystems and Nanoengineering **6**, 1–11 (2020) DOI: `10.1038/s41378-020-00190-6` (cited on pp. 48, 60).

[50] F. Heide, Q. Fu, Y. Peng, and W. Heidrich, “Encoded diffractive optics for full-spectrum computational imaging”, Scientific Reports **6**, 1–10 (2016) DOI: `10.1038/srep33543` (cited on p. 48).

[51] H. Wei, F. Callewaert, W. Hadibrata, V. Velev, Z. Liu, P. Kumar, K. Aydin, and S. Krishnaswamy, “Two-Photon Direct Laser Writing of Inverse-Designed Free-Form Near-Infrared Polarization Beamsplitter”, Advanced Optical Materials **7**, 1900513 (2019) DOI: `10.1002/adom.201900513` (cited on pp. 48, 58).

[52] B. J. Justice, J. J. Mock, L. Guo, A. Degiron, D. Schurig, and D. R. Smith, “Spatial mapping of the internal and external electromagnetic fields of negative index metamaterials”, Optics Express **14**, 8694 (2006) DOI: `10.1364/oe.14.008694` (cited on pp. 48–50).

[53] J. M. Felício, C. A. Fernandes, and J. R. Costa, “Complex permittivity and anisotropy measurement of 3D-printed PLA at microwaves and millimeter-waves”, in Icecom 2016 - conference proceedings, 22nd international conference on applied electromagnetics and communications (2017), DOI: `10.1109/ICECom.2016.7843900` (cited on p. 50).

[54] J. N. Grima and K. E. Evans, “Auxetic behavior from rotating squares”, Journal of Materials Science Letters **19**, 1563–1565 (2000) DOI: `10.1023/A:1006781224002` (cited on p. 55).

[55] J. N. Grima, V. Zammit, R. Gatt, A. Alderson, and K. E. Evans, “Auxetic behaviour from rotating semi-rigid units”, physica status solidi (b) **244**, 866–882 (2007) DOI: `10.1002/pssb.200572706` (cited on p. 55).

[56] H. M. Kolken and A. A. Zadpoor, *Auxetic mechanical metamaterials*, 2017, DOI: `10.1039/c6ra27333e` (cited on p. 55).

[57] Y. Y. Xie, P. N. Ni, Q. H. Wang, Q. Kan, G. Briere, P. P. Chen, Z. Z. Zhao, A. Delga, H. R. Ren, H. D. Chen, C. Xu, and P. Genevet, “Metasurface-integrated vertical cavity surface-emitting lasers for programmable directional lasing emissions”, Nature Nanotechnology **15**, 125–130 (2020) DOI: `10.1038/s41565-019-0611-y` (cited on p. 58).

[58] F. Yesilkoy, E. R. Arvelo, Y. Jahani, M. Liu, A. Tittl, V. Cevher, Y. Kivshar, and H. Altug, *Ultrasensitive hyperspectral imaging and biodetection enabled by dielectric metasurfaces*, 2019, DOI: `10.1038/s41566-019-0394-6` (cited on p. 58).

[59] H. Kwon, E. Arbabi, S. M. Kamali, M. S. Faraji-Dana, and A. Faraon, "Single-shot quantitative phase gradient microscopy using a system of multifunctional metasurfaces", Nature Photonics **14**, 109–114 (2020) DOI: `10.1038/s41566-019-0536-x` (cited on p. 58).

[60] Z. Wu, L. Li, Y. Li, and X. Chen, "Metasurface Superstrate Antenna with Wideband Circular Polarization for Satellite Communication Application", IEEE Antennas and Wireless Propagation Letters **15**, 374–377 (2016) DOI: `10.1109/LAWP.2015.2446505` (cited on p. 58).

[61] A. Y. Piggott, J. Petykiewicz, L. Su, and J. Vučković, "Fabrication-constrained nanophotonic inverse design", Scientific Reports **7**, 1–7 (2017) DOI: `10.1038/s41598-017-01939-2`, arXiv:`1612.03222` (cited on p. 58).

[62] A. Y. Piggott, E. Y. Ma, L. Su, G. H. Ahn, N. V. Sapra, D. Vercruysse, A. M. Netherton, A. S. Khope, J. E. Bowers, and J. Vučković, "Inverse-Designed Photonics for Semiconductor Foundries", ACS Photonics **7**, 569–575 (2020) DOI: `10.1021/acsphotonics.9b01540`, arXiv:`1911.03535` (cited on p. 58).

[63] C. Jung-Kubiak, T. J. Reck, J. V. Siles, R. Lin, C. Lee, J. Gill, K. Cooper, I. Mehdi, and G. Chattopadhyay, "A Multistep DRIE Process for Complex Terahertz Waveguide Components", IEEE Transactions on Terahertz Science and Technology **6**, 690–695 (2016) DOI: `10.1109/TTHZ.2016.2593793` (cited on pp. 58, 76, 78, 87).

[64] S. Rahiminejad, M. Alonso-del Pino, T. J. Reck, A. Peralta, R. Lin, C. Jung-Kubiak, and G. Chattopadhyay, "A Low-Loss Silicon MEMS Phase Shifter Operating in the 550 GHz Band", IEEE Transactions on Terahertz Science and Technology, 1–1 (2021) DOI: `10.1109/TTHZ.2021.3085123` (cited on p. 58).

[65] H. C. Nathanson, W. E. Newell, R. A. Wickstrom, and J. R. Davis, "The Resonant Gate Transistor", IEEE Transactions on Electron Devices **ED-14**, 117–133 (1967) DOI: `10.1109/T-ED.1967.15912` (cited on p. 60).

[66] G. I. Taylor, "The coalescence of closely spaced drops when they are at different electric potentials", Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences **306**, 423–434 (1968) DOI: `10.1098/rspa.1968.0159` (cited on p. 60).

[67] T. Roy, S. Zhang, I. W. Jung, M. Troccoli, F. Capasso, and D. Lopez, "Dynamic metasurface lens based on MEMS technology", APL Photonics **3**, 021302 (2018) DOI: `10.1063/1.5018865`, arXiv:`1712.03616` (cited on p. 60).

[68] E. Arbabi, A. Arbabi, S. M. Kamali, Y. Horie, M. S. Faraji-Dana, and A. Faraon, "MEMS-tunable dielectric metasurface lens", Nature Communications **9**, 1–9 (2018) DOI: `10.1038/s41467-018-03155-6`, arXiv:`1712.06548` (cited on p. 60).

[69] J. Schwerdtfeger, F. Wein, G. Leugering, R. F. Singer, C. Körner, M. Stingl, and F. Schury, "Design of auxetic structures via mathematical optimization", Advanced Materials **23**, 2650–2654 (2011) DOI: `10.1002/adma.201004090` (cited on p. 60).

[70] N. T. Kaminakis and G. E. Stavroulakis, "Topology optimization for compliant mechanisms, using evolutionary-hybrid algorithms and application to the design of auxetic materials", in Composites part b: engineering, Vol. 43, 6 (2012), pp. 2655–2668, DOI: `10.1016/j.compositesb.2012.03.018` (cited on p. 60).

[71] C. Ballew, G. Roberts, T. Zheng, and A. Faraon, "Constraining continuous topology optimizations to discrete solutions for photonic applications", (2021), arXiv:`2107.09468 [physics.app-ph]`, (cited on p. 63).

[72] C. Ballew, "Thesis - multifunctional volumetric metaoptics - section 3.3 supplementary videos", DOI: `10.22002/D1.2081` (2021) DOI: `10.22002/D1.2081` (cited on p. 67).

[73] J. Hong, H. Son, C. Kim, S.-E. Mun, J. Sung, and B. Lee, "Absorptive metasurface color filters based on hyperbolic metamaterials for a CMOS image sensor", Opt. Express **29**, 3643–3658 (2021) DOI: `10.1364/OE.415874` (cited on p. 70).

[74] N. Zhao, P. B. Catrysse, and S. Fan, "Perfect RGB-IR Color Routers for Sub-Wavelength Size CMOS Image Sensor Pixels", Advanced Photonics Research **2**, 2000048 (2021) DOI: `https://doi.org/10.1002/adpr.202000048` (cited on p. 70).

[75] A. Wang, P. Gill, and A. Molnar, "Light field image sensors based on the Talbot effect", Appl. Opt. **48**, 5897–5905 (2009) DOI: `10.1364/AO.48.005897` (cited on p. 70).

[76] E. H. Adelson and J. Y. A. Wang, "Single lens stereo with a plenoptic camera", IEEE Transactions on Pattern Analysis and Machine Intelligence **14**, 99–106 (1992) DOI: `10.1109/34.121783` (cited on p. 70).

[77] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz, "Light Field Microscopy", ACM Trans. Graph. **25**, 924–934 (2006) DOI: `10.1145/1141911.1141976` (cited on p. 70).

[78] M. Levoy, Z. Zhang, and I. McDowall, "Recording and controlling the 4D light field in a microscope using microlens arrays", Journal of Microscopy **235**, 144–162 (2009) DOI: `https://doi.org/10.1111/j.1365-2818.2009.03195.x` (cited on p. 70).

[79] R. Fatemi, C. Ives, A. Khachaturian, and A. Hajimiri, “Subtractive photonics”, Opt. Express **29**, 877–893 (2021) `DOI`: `10.1364/OE.410139` (cited on p. 72).

[80] P. H. Siegel, *Terahertz technology*, 2002, `DOI`: `10.1109/22.989974` (cited on p. 74).

[81] P. H. Siegel, *THz instruments for space*, 2007, `DOI`: `10.1109/TAP.2007.908557` (cited on p. 74).

[82] D. T. Leisawitz, W. C. Danchi, M. J. DiPirro, L. D. Feinberg, D. Y. Gezari, M. Hagopian, W. D. Langer, J. C. Mather, S. H. Moseley, Jr., M. Shao, R. F. Silverberg, J. G. Staguhn, M. R. Swain, H. W. Yorke, and X. Zhang, “Scientific motivation and technology requirements for the SPIRIT and SPECS far-infrared/submillimeter space interferometers”, in Uv, optical, and ir space telescopes and instruments, Vol. 4013, 28 (2000), pp. 36–46, `DOI`: `10.1117/12.393957` (cited on p. 74).

[83] J. S. Greaves, A. M. Richards, W. Bains, P. B. Rimmer, H. Sagawa, D. L. Clements, S. Seager, J. J. Petkowski, C. Sousa-Silva, S. Ranjan, E. Drabek-Maunder, H. J. Fraser, A. Cartwright, I. Mueller-Wodarg, Z. Zhan, P. Friberg, I. Coulson, E. Lee, and J. Hoge, “Phosphine gas in the cloud decks of Venus”, Nature Astronomy, 1–10 (2020) `DOI`: `10.1038/s41550-020-1174-4`, arXiv:`2009.06593` (cited on p. 74).

[84] C. I. Honniball, P. G. Lucey, S. Li, S. Shenoy, T. M. Orlando, C. A. Hibbitts, D. M. Hurley, and W. M. Farrell, “Molecular water detected on the sunlit Moon by SOFIA”, Nature Astronomy, 1–7 (2020) `DOI`: `10.1038/s41550-020-01222-x` (cited on p. 74).

[85] M. Alonso-Delpino, T. Reck, C. Jung-Kubiak, C. Lee, and G. Chattopadhyay, “Development of Silicon Micromachined Microlens Antennas at 1.9 THz”, IEEE Transactions on Terahertz Science and Technology **7**, 191–198 (2017) `DOI`: `10.1109/TTHZ.2017.2655340` (cited on pp. 75, 77, 80).

[86] S. M. Kamali, E. Arbabi, A. Arbabi, Y. Horie, and A. Faraon, “Highly tunable elastic dielectric metasurface lenses”, Laser & Photonics Reviews **10**, 1002–1008 (2016) `DOI`: `https://doi.org/10.1002/lpor.201600144` (cited on p. 83).

[87] V. Liu and S. Fan, “S^4 : A free electromagnetic solver for layered periodic structures”, Computer Physics Communications **183**, 2233–2244 (2012) `DOI`: `10.1016/j.cpc.2012.04.026` (cited on p. 83).

[88] N. Ghalichechian and K. Sertel, “Permittivity and Loss Characterization of SU-8 Films for mmW and Terahertz Applications”, IEEE Antennas and Wireless Propagation Letters **14**, 723–726 (2015) `DOI`: `10.1109/LAWP.2014.2380813` (cited on p. 84).

[89] F. Hindle, R. Bocquet, A. Pienkina, A. Cuisset, and G. Mouret, “Terahertz gas phase spectroscopy using a high-finesse Fabry–Pérot cavity”, Optica **6**, 1449 (2019) `DOI`: `10.1364/optica.6.001449` (cited on pp. 100, 103).

Appendix A

EFFECTS OF CROSS-POLARIZATION ON MULTIFUNCTIONAL SUGARCUBE

The device discussed in the main manuscript sorts incident light based on its polarization, wavelength, and incident angle. The device was designed explicitly for two orthogonal linearly-polarized sources, one polarized in the xz plane and the other in the yz plane. Here, we wish to determine the effect on device performance when excited with planewaves of arbitrary polarization. The response of the device for any polarization state can be determined by coherently summing the electric fields, applying relative amplitude and phase shifts (i.e. a Jones vector) to the two linearly-polarized components to obtain arbitrary elliptical polarizations.

The Jones vector is defined as:

$$\begin{pmatrix} E_{0x} \\ E_{0y}e^{i\phi_J} \end{pmatrix}$$

To sweep the Jones vector over all possibilities, we introduce a variable $\theta_J \in [0, \pi]$ and set $E_{0x} = \cos(\theta_J)$ and $E_{0y} = \sin(\theta_J)$. We then plot the transmission to all pixels as function of (θ_J, ϕ_J) with $\phi_J \in [0, \pi]$. Since we wish to observe interference effects here, we do not average over any properties as was done in quantifying performance in the main manuscript. Instead, the results for all training modes are plotted independently. The transmission to the pixels do exhibit a dependence on the incident, but it is not substantial enough to alter the behavior of the device. Thus, the device will exhibit the behavior reported in the main manuscript under any polarization input.

Figures A.1-A.5 below show these results. θ_J, the relative amplitude, is the x-axis and ϕ_J, the relative phase, is the y-axis for each subplot. Each figure contains the results for all pixels at a particular incident angle (θ, ϕ). The specific pixels to which transmission is *not* maximized are faded. The ordering of the pixels is identical to that of Fig. 3.11 in Chapter 3.

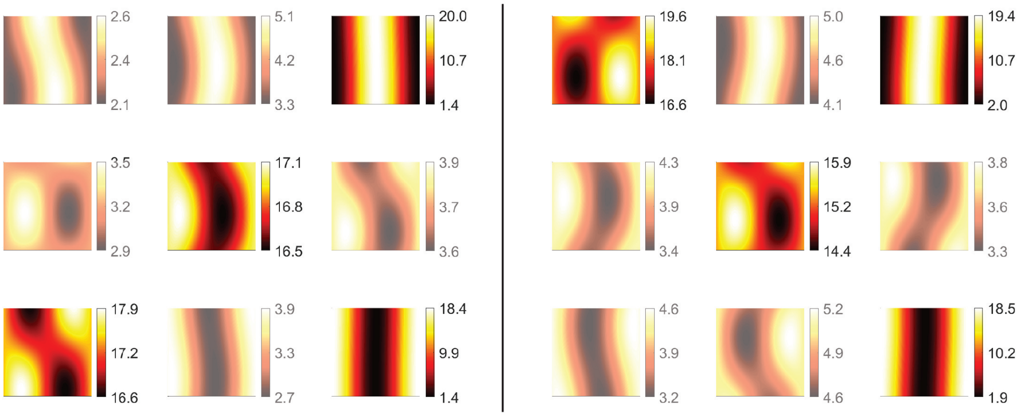

Figure A.1: $(\theta, \phi) = (0°, 0°)$. Left set of figures correspond to 620 nm, the right set to 532 nm. The non-greyed out subset of figures represent the pixels that receive the most power flux under the current excitation state.

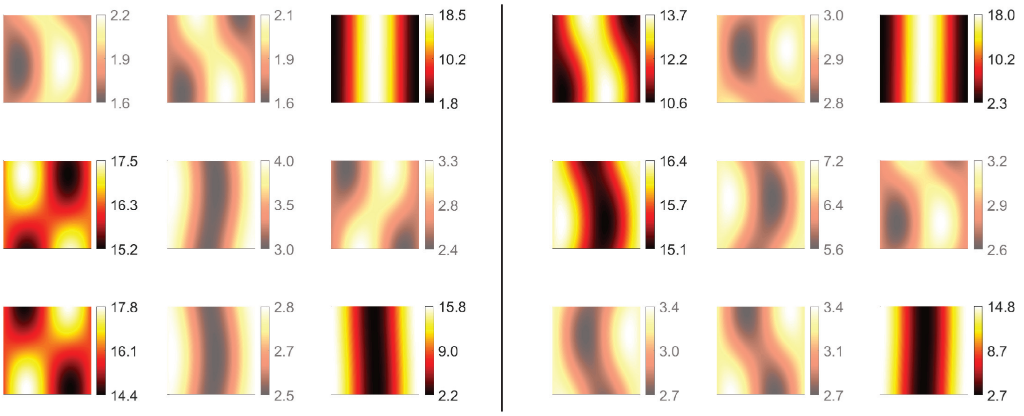

Figure A.2: $(\theta, \phi) = (5°, 0°)$. Left set of figures correspond to 620 nm, the right set to 532 nm. The non-greyed out subset of figures represent the pixels that receive the most power flux under the current excitation state.

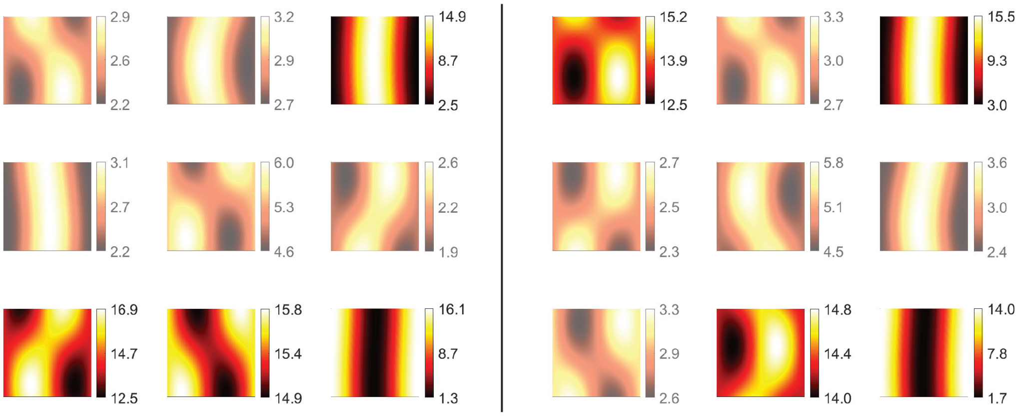

Figure A.3: $(\theta, \phi) = (5°, 90°)$. Left set of figures correspond to 620 nm, the right set to 532 nm. The non-greyed out subset of figures represent the pixels that receive the most power flux under the current excitation state.

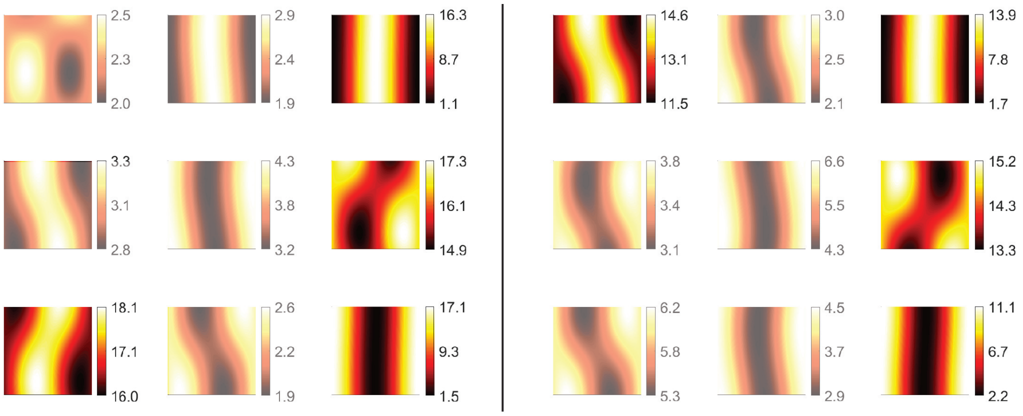

Figure A.4: $(\theta, \phi) = (5°, 180°)$. Left set of figures correspond to 620 nm, the right set to 532 nm. The non-greyed out subset of figures represent the pixels that receive the most power flux under the current excitation state.

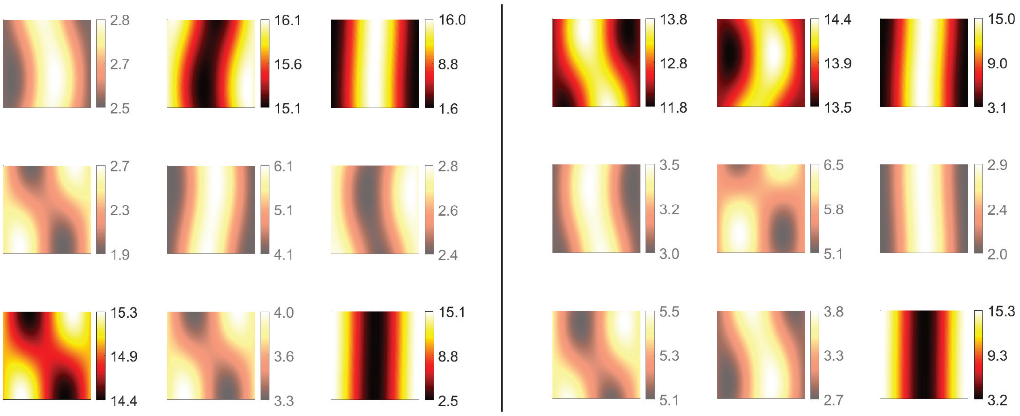

Figure A.5: $(\theta, \phi) = (5°, 270°)$. Left set of figures correspond to 620 nm, the right set to 532 nm. The non-greyed out subset of figures represent the pixels that receive the most power flux under the current excitation state.