Computation & Design for Nanophotonics

by

Ardavan Oskooi

S.M., Massachusetts Institute of Technology (2008) B.A.Sc., University of Toronto (2004)

Submitted to the Department of Materials Science and Engineering in partial fulfillment of the requirements for the degree of

Doctor of Science in Materials Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

© Massachusetts Institute of Technology 2010. All rights reserved.

| Author |
|---|
| Department of Materials Science and Engineering |
| May 10, 2010 |
| Way 10, 2010 |
| Certified by |
| Steven G. Johnson |
| Associate Professor of Mathematics |
| Thesis Supervisor |
| Certified by |
| Yoel Fink |
| Associate Professor of Materials Science |
| Thesis Supervisor |
| Accepted by |
| Christine Ortiz |
| Chair, Department Committee on Graduate Theses |

Computation & Design for Nanophotonics

by

Ardavan Oskooi

Submitted to the Department of Materials Science and Engineering on May 10, 2010, in partial fulfillment of the requirements for the degree of Doctor of Science in Materials Science and Engineering

Abstract

The versatility of computational design as an alternative to design by nanofabrication has made computers a reliable design tool in nanophotonics. Given that almost any 2d pattern can be fabricated at infrared length scales, there exists a large number of degrees of freedom in nanophotonic device design. However current designs are adhoc and could potentially benefit from optimization but there are several outstanding issues regarding PDE-based optimization for electromagnetism that must first be addressed: continuously and accurately deforming geometric objects represented on a discrete uniform grid while avoiding staircasing effects, reducing the computational expense of large simulations while improving accuracy, resolving the breakdown of standard absorbing boundary layers for important problems, finding robust designs that are impervious to small perturbations, and finally distinguishing global from local minima. We address each of these issues in turn by developing novel subpixel smoothing methods that markedly improve the accuracy of simulations, demonstrate the failure of perfectly matched layers (PML) in several important cases and propose a workaround, develop a simple procedure to determine the validity of any PML implementation and incorporate these and other enhancements into a flexible, free software package for electromagnetic simulations based on the finite-difference timedomain (FDTD) method. Next we investigate two classes of design problems in nanophotonics. The first involves finding cladding structures for holey photoniccrystal fibers at low-index contrasts that permit a larger class of materials to be used in the fabrication process. The second is the development of adiabatic tapers for coupling to slow-light modes of photonic-crystal waveguides that are insensitive to manufacturing and operational variability.

Thesis Supervisor: Steven G. Johnson Title: Associate Professor of Mathematics

Thesis Supervisor: Yoel Fink Title: Associate Professor of Materials Science Education is Not the Filling of a Pail but the Lighting of a Fire.

- William Butler Yeats

Acknowledgments

MIT is a special place. Little did I know during my orientation visit one frosty spring week six years ago how this incredible place would alter the course of my life. MIT is teeming with revolutionary ideas dreamt up by the world's most brilliant individuals.

I met the remarkable Steven G. Johnson during my first week at MIT who had recently started as an Assistant Professor. I was fortunate to become the first graduate student in his group and have benefited tremendously from the experience ever since. From the initial months of working together and throughout, I have continued to be awed by the depth and breadth of his scientific repertoire. His insatiable curiosity, meticulousness, tenacity to see things through to the end, ability to make connections among disparate fields and profound commitment to sharing the fruits of knowledge make him an exceptional researcher; one who has set the bar for my own professional career. He is always brimming with new research ideas and is never afraid to tackle the toughest of problems with his quintessential clarity of thought. His "Mathematical Methods in Nanophotonics" remains my favorite class of all time; the lectures are the stuff of legend. Steven cares deeply about his students. He spends an inordinate amount of time mentoring each of us and is never too tired to discuss science. Many times when I was stuck on research and feeling discouraged, a trip to Steven's office made all the difference to my rejuvenation. His mastery in problem solving is unparalleled and I affectionately refer to him as the "Oracle" for his singular ability to solve any problem, no matter how complex. Steven has consistently challenged me to become a better scientist. He has been a wonderful mentor, advisor and friend. Thanks for putting up with such a stubborn student, Steven.

There are three individuals who have left an indelible mark on my MIT experience. I am grateful to have gotten to know each of them well during my time here.

I met Wenhao Liu during the fateful summer before MIT at the IBM Almaden Research Center in San Jose, CA. Wen's charisma, wonder of the world around him and passion for technology are one of a kind. Happy memories abound: two years as roommates at the Sidney-Pacific Graduate Residence, trips to NYC, late nights in Chinatown, hiking in NH, exploring Boston's nightlife, and star-gazing on the MA coast. Wen has been a constant fixture of my life outside of work.

I met Alejandro Rodriguez-Wong in my first semester who was then a precocious and enthusiastic undergraduate. We have spent countless hours discussing physics, debating contemporary issues, and playing practical jokes on one another in CMT. Alejandro's strong work ethic and devotion to research set an excellent example for all. He is a diligent, gregarious and compassionate human being who relishes helping out those around him. His ebullient presence always livened our work environment. We share many unforgettable memories together.

I met Monika Schleier-Smith serendipitously on a warm summer evening in my third year while walking home from work. Monika is a supremely talented individual and she dedicates herself to everything she does with utmost zeal. She is a gifted scientist, musician and athlete with an endearing personality. I will always remember with fondness our bike rides, walks and runs around Cambridge and Boston, rowing on the Charles on the 4th of July, evenings sharing a meal, movies at the LSC, trips to the BSO, and the many hours in conversation. She has a special place in my heart.

I would also like to thank Professor John Joannopoulos for all his wisdom, support and guidance. JJ combines technical mastery, formidable communication skills, an effective management style with a vision for the future. He is the epitome of a great scientist. My sincere gratitude to Dr. Peter Bermel with whom I spent untold hours administering our group's computers. Peter's proficiency and affability made working together enjoyable. Thank you to all the talented members of the JDJ/SGJ/MS group, especially Dr. Aristeidis Karalis, Alexander McCauley and Hila Hashemi.

I dedicate this thesis to my parents, Fereydoun and Haydeh, who had the courage to leave our homeland during a very turbulent time in its history. They managed to raise two young children on their own in a foreign land and provided us with everything we needed. Their unyielding love and encouragement has been an everlasting source of inspiration. I am eternally grateful for all the sacrifices they have made.

I look back and smile. The best is yet to come.

Contents

| 1 | Introduction | | 25 |
|----------|--|--|-----------|
| 2 | Sub-pixel smoothing for dielectric media | | |
| | 2.1 | Summary | 31 |
| | 2.2 | Overview | 31 |
| | 2.3 | Designing subpixel smoothing algorithms with perturbation theory | 33 |
| | 2.4 | Analysis of smoothing perturbation | 36 |
| | 2.5 | Stable FDTD field-update implementation | 44 |
| | 2.6 | Numerical performance of methods for isotropic media | 45 |
| | 2.7 | Numerical performance of methods for anisotropic media | 49 |
| | 2.8 | Field singularities at sharp corners | 52 |
| | 2.9 | Conclusion | 53 |
| 3 | The | failure of perfectly matched layers | 55 |
| | 3.1 | Summary | 55 |
| | 3.2 | Overview | 56 |
| | | 3.2.1 Various PML Formulations | 57 |
| | | 3.2.2 PMLs in Photonic Crystals | 58 |
| | 3.3 | PMLs versus Adiabatic absorbers | 60 |
| | 3.4 | Brief review of PML | 61 |
| | | 3.4.1 Mathematical formulation | 61 |
| | | 3.4.2 Absorption profile | 63 |
| | 3.5 | Adiabatic theorems in electromagnetism | 64 |

| | 3.6 | Failure of PML | 65 |
|---|-----|---|-----|
| | | 3.6.1 Homogeneous & inhomogeneous media | 65 |
| | | 3.6.2 Backward-wave structures | 68 |
| | | 3.6.3 PMLs & adiabatic absorbers | 69 |
| | 3.7 | Smoothness & Reflection | 72 |
| | | 3.7.1 Numerical results | 73 |
| | | 3.7.2 Analysis | 78 |
| | | 3.7.3 Adiabatic theorems in discrete systems | 82 |
| | 3.8 | Towards Better Absorbers | 82 |
| | | 3.8.1 Smoothness & C_{∞} functions | 87 |
| | | 3.8.2 Balancing round-trip & transition reflections | 88 |
| | 3.9 | Conclusion | 90 |
| 4 | A s | imple validation scheme for perfectly matched layers | 93 |
| | 4.1 | Summary | 93 |
| | 4.2 | Overview | 93 |
| | 4.3 | PML formulation for lossless, non-dispersive, anisotropic media | 98 |
| | 4.4 | Non-PML materials in Meep | 100 |
| | 4.5 | PML formulation in frequency domain | 101 |
| | 4.6 | PML formulation in time domain | 103 |
| | 4.7 | Review of isotropic, nondispersive PML | 105 |
| | 4.8 | Concluding remarks | 106 |
| 5 | Me | ep: A flexible free-software package | |
| | for | electromagnetic simulations by the FDTD method | 107 |
| | 5.1 | Summary | 107 |
| | 5.2 | Overview | 109 |
| | 5.3 | Development History | 113 |
| | 5.4 | Grids and Boundary Conditions | 114 |
| | | 5.4.1 Coordinates and grids | 115 |
| | | 5.4.2 Grid chunks and owned points | 115 |
| | | | |

| | | 5.4.3 | Boundary conditions and symmetries | 117 | |
|---|---|---|---|---|--|
| | 5.5 | Interp | olation and the illusion of continuity | 119 | |
| | 5.6 | Mater | ials | 126 | |
| | | 5.6.1 | Nonlinear materials | 127 | |
| | | 5.6.2 | Absorbing boundary layers: PML, pseudo-PML, and quasi-PMI | 129 | |
| | 5.7 | Enabli | ng typical computations | 130 | |
| | | 5.7.1 | Computing flux spectra | 131 | |
| | | 5.7.2 | Analyzing resonant modes | 133 | |
| | | 5.7.3 | Frequency-domain solver | 134 | |
| | 5.8 | User in | nterface and scripting | 138 | |
| | 5.9 | Abstra | action versus performance | 142 | |
| | | 5.9.1 | Timestepping and cache tradeoffs | 143 | |
| | | 5.9.2 | The loop-in-chunks abstraction | 145 | |
| | 5.10 | Conclu | ıding remarks | 147 | |
| 6 | Zero-group-velocity modes in chalcogenide holey photonic-crystal | | | | |
| - | | 0 | | | |
| - | fibe | rs | | 149 | |
| - | fibe : 6.1 | rs Summ | ary | 149 149 | |
| - | fibe: 6.1 6.2 | rs Summ Introd | ary | 149 149 149 | |
| | fibe: 6.1 6.2 6.3 | rs Summ Introd Reviev | ary | 149 149 149 151 | |
| | fibe: 6.1 6.2 6.3 6.4 | rs Summ Introd Review Gaps a | ary | 149 149 149 151 155 | |
| | fibe: 6.1 6.2 6.3 6.4 6.5 | rs Summ Introd Review Gaps a Topolo | ary | 149 149 149 151 155 157 | |
| | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 | rs Summ Introd Review Gaps a Topolo Claddi | ary | 149 149 151 155 157 159 | |
| | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 | rs Summ Introd Review Gaps a Topolo Claddi Coupli | ary | 149 149 151 155 157 159 163 | |
| | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 | rs Summ Introd Review Gaps a Topolo Claddi Coupli Solid o | ary | 149 149 151 155 157 159 163 164 | |
| | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 | rs Summ Introd Review Gaps a Topolo Claddi Coupli Solid o Conclu | ary | 149 149 151 155 157 159 163 164 164 | |
| 7 | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 Rob | rs Summ Introd Review Gaps a Topold Claddi Coupli Solid o Conclu | ary | 149 149 151 155 157 159 163 164 164 165 | |
| 7 | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 Rob 7.1 | rs Summ Introd Review Gaps a Topolo Claddi Coupli Solid o Conclu oust de Summ | ary | 149 149 151 155 157 159 163 164 164 165 | |
| 7 | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 Rob 7.1 7.2 | rs Summ Introd Review Gaps a Topolo Claddi Coupli Solid o Conclu Dust de Summ Introd | ary | 149 149 149 151 155 157 159 163 164 164 165 165 166 | |
| 7 | fibe: 6.1 6.2 6.3 6.4 6.5 6.6 6.7 6.8 6.9 Rob 7.1 7.2 7.3 | rs Summ Introd Review Gaps a Topolo Claddi Coupli Solid o Conclu oust de Summ Introd | ary | 149 149 151 155 157 159 163 164 164 165 166 168 | |

| | | 1.5.1 | Taper snape and renection magnitude | 108 |
|----|-----------------|--------|--|-----|
| | | 7.3.2 | Parameter uncertainty and worst-case reflection magnitude | 170 |
| | | 7.3.3 | Nominal and robust taper shape problems $\ldots \ldots \ldots \ldots$ | 172 |
| | | 7.3.4 | Piecewise-linear taper shape parametrization | 173 |
| | | 7.3.5 | Successive refinement | 175 |
| | 7.4 | Comp | utation of reflection magnitude | 176 |
| | | 7.4.1 | Coupled-mode theory | 176 |
| | | 7.4.2 | Coupled-mode reflection gradient $\ldots \ldots \ldots \ldots \ldots \ldots$ | 178 |
| | | 7.4.3 | Brute-force verification | 179 |
| | | 7.4.4 | Worst-case reflection magnitude $\ldots \ldots \ldots \ldots \ldots \ldots \ldots$ | 180 |
| | 7.5 | Nume | rical results | 181 |
| | | 7.5.1 | Taper geometry and uncertainty model | 181 |
| | | 7.5.2 | Pessimizing method | 183 |
| | | 7.5.3 | Optimization | 184 |
| | | 7.5.4 | Results | 186 |
| | 7.6 | Conclu | usions | 187 |
| 8 | Cor | clusio | n | 189 |
| Bi | Bibliography 19 | | | 191 |

List of Figures

| 2-1 | Schematic of an interface perturbation: the interface between two | |
|-----|---|----|
| | materials $\boldsymbol{\varepsilon}^a$ and $\boldsymbol{\varepsilon}^b$ (possibly anisotropic) is shifted by some small | |
| | position-dependent displacement h | 34 |
| 2-2 | Schematic 2d Yee FDTD discretization near a dielectric interface, show- | |
| | ing the method [235] used to compute the part of E_x that comes from | |
| | \mathbf{D}_y and the locations where various ε^{-1} components are required | 35 |
| 2-3 | TE eigenfrequency error vs. resolution for a Bragg mirror of alternating | |
| | air and $\varepsilon = 12$ (inset). | 37 |
| 2-4 | TE eigenfrequency error vs. resolution for a square lattice of elliptical | |
| | air holes in $\varepsilon = 12$ (inset). | 39 |
| 2-5 | TM eigenfrequency error vs. resolution for a square lattice of elliptical | |
| | air holes in $\varepsilon = 12$ (inset). | 40 |
| 2-6 | Eigenfrequency error vs. resolution for a cubic lattice of $\varepsilon=12$ ellip- | |
| | soids in air (inset). | 41 |
| 2-7 | Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a square lat- | |
| | tice (period a) of 2d anisotropic ellipsoids (right inset) versus spatial | |
| | resolution (units of pixels per vacuum wavelength λ), for a variety of | |
| | subpixel smoothing techniques. Straight lines for perfect linear (black | |
| | dashed) and perfect quadratic (black solid) convergence are shown for | |
| | reference. Most curves are for the first eigenvalue band (left inset shows | |
| | E_z in unit cell), with vacuum wavelength $\lambda = 4.85a$. Hollow squares | |
| | show new method for band 15 (middle inset), with $\lambda = 1.7a$. Maximum | |
| | resolution for all curves is 100 pixels/ a | 43 |

| 2-8 | Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a cubic lattice | |
|------|--|----|
| | (period a) of 3d anisotropic ellipsoids (right inset) versus spatial resolu- | |
| | tion (units of pixels per vacuum wavelength λ), for a variety of subpixel | |
| | smoothing techniques. Straight lines for perfect linear (black dashed) | |
| | and perfect quadratic (black solid) convergence are shown for refer- | |
| | ence. Most curves are for the first eigenvalue band (left inset shows E_x | |
| | in xy cross-section of unit cell), with vacuum wavelength $\lambda = 5.15a$. | |
| | Hollow squares show new method for band 13 (middle inset), with | |
| | $\lambda = 2.52a$. New method for bands 1 and 13 is shown for resolution up | |
| | to 100 pixels/ a | 46 |
| 2-9 | Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a square lattice | |
| | (period a) of 2d anisotropic ellipses (green inset), versus spatial reso- | |
| | lution, for a variety of sub-pixel smoothing techniques. Straight lines | |
| | for perfect linear (black dashed) and perfect quadratic (black solid) | |
| | convergence are shown for reference | 47 |
| 2-10 | Relative error $\Delta \omega / \omega$ for an eigenmode calculation with cubic lattice | |
| | (period a) of 3d anisotropic ellipsoids (green inset), versus spatial res- | |
| | olution, for a variety of sub-pixel smoothing techniques. Straight lines | |
| | for perfect linear (black dashed) and perfect quadratic (black solid) | |
| | convergence are shown for reference | 48 |
| 2-11 | Degraded accuracy due to field singularities at sharp corners: TE eigen- | |
| | frequency error vs. resolution for square lattice of tilted-square air holes | |
| | in $\varepsilon = 12$ (inset) | 51 |
| | | |
| 3-1 | (a) PML is still reflectionless for inhomogeneous media such as waveg- | |
| | uides that are homogeneous in the direction perpendicular to the PML. | |
| | (b, c) PML is no longer reflectionless when the dielectric function is dis- | |
| | continuous (non-analytic) in the direction perpendicular to the PML, | |
| | as in a photonic crystal (b) or a waveguide entering the PML at an | |
| | angle (c) | 58 |

3-7 Field convergence factor [eq. (3.11)] (~ reflection/ L^2) vs. PML thickness L for 2d vacuum (inset) at a resolution of 20pixels/ λ for various shape functions s(u) ranging from linear [s(u) = u] to quintic [$s(u) = u^5$]. For reference, the corresponding asymptotic power laws are shown as dashed lines. Inset: $\Re[E_z]$ field pattern for the (point) source at the origin (blue/white/red = positive/zero/negative).

76

77

83

- 3-8 Field convergence factor [eq. (3.11)] (~ reflection/ L^2) vs. pPML thickness L for the discontinuous 2d periodic medium (left inset: square lattice of square air holes in $\varepsilon = 12$) with period a, at a resolution of 10pixels/a with a vacuum wavelength $\lambda = 0.6667a$ (not in a band gap) for various shape functions s(u) ranging from linear [s(u) = u] to quintic [$s(u) = u^5$]. For reference, the corresponding asymptotic power laws are shown as dashed lines. Right inset: $\Re[E_z]$ field pattern for the (point) source at the origin (blue/white/red = positive/zero/negative).
- 3-9 Reflectivity vs. PML thickness L for 1d vacuum (blue circles) at a resolution of 50pixels/ λ , and for pPML thickness L in the 1d periodic medium of Fig. 3-6 (red squares) with period a at a resolution of 50pixels/a with a wavevector $k_x = 0.9\pi/a$ (vacuum wavelength $\lambda = 0.9597a$. In both cases, a C_{∞} (infinitely differentiable) shape function $s(u) = e^{1-1/u}$ for u > 0 is used, leading to asymptotic convergence as $e^{-\alpha\sqrt{L}}$ for some constants α .
- 3-10 Reflectivity vs. PML thickness L for two different absorber profiles in the 1d uniform medium of Fig. 3-5 at a resolution of 50pixels/ λ . In both cases, a C_{∞} (infinitely differentiable) shape function is used, leading to asymptotic exponential convergence but the blue curve has a slightly lower constant factor since it approximates a quadratic taper profile for small taper lengths.

3-11 Reflectivity vs. pPML thickness L for two different absorber profiles in the 1d periodic medium of Fig. 3-6 with a wavevector $k_x = 0.9\pi/a$ (vacuum wavelength $\lambda = 0.9597a$ at a resolution of 50pixels/ λ . In both cases, a C_{∞} (infinitely differentiable) shape function is used, leading to asymptotic exponential convergence but the blue curve has a slightly lower constant factor since it approximates a quadratic taper profile for small taper lengths.

85

86

- 3-13 Reflectivity vs. PML thickness L for 1d vacuum (inset) at a resolution of 50pixels/ λ for $s(u) = u^2$, with the round-trip reflection either set to $R_0 = 10^{-16}$ (upper blue line) or set to match the estimated transition reflection from Fig. 3-5 (lower red line). By matching the round-trip reflection R_0 to the estimated transition reflection, one can obtain a substantial reduction in the constant factor of the total reflection, although the asymptotic power law is only changed by a ln L factor. . .

4-1 Results from 2d PWFD showing field convergence factor (a proxy for the reflection coefficient) versus resolution in pixels/ λ_v for both isotropic and anisotropic media with PML, Z-PML (from Ref. 253) and non-PML (conductivity absorber in electric field) absorbing boundaries. For the anisotropic medium, Z-PML fails to be reflectionless in the limit of high resolution. Inset: E_z field profile of a point source at the center of the 2d computational cell surrounded by absorbing material (blue/white/red = positive/zero/negative).

96

- 4-2 Results from 2d PWFD simulation showing field convergence factor (~ reflection/ L^2) vs. absorber thickness in units of vacuum wavelength (L/λ_v) for anisotropic media at a resolution of 9pixels/ λ_v for various polynomial absorber functions s(u) ranging from linear [s(u) = u] in blue to cubic $[s(u) = u^3]$ in green. As the absorber becomes thicker and the absorption is turned on more gradually, reflection goes to zero via the adiabatic theorem. For reference, the corresponding asymptotic power laws are shown as dashed lines. Fixing the round-trip reflection yields similar scaling relationships and values between the three types of absorbers. Inset: $\Re[E_z]$ field pattern for the (point) source at the center (blue/white/red = positive/zero/negative).

- 5-1 The computational cell is divided into chunks (left) that have a onepixel overlap (gray regions). Each chunk (right) represents a portion of the Yee grid, partitioned into *owned* points (chunk interior) and *not-owned* points (gray regions around the chunk edges) that are determined from other chunks and/or via boundary conditions. Every point in the interior of the computational cell is owned by exactly one chunk, the chunk responsible for timestepping that point. 116
- 5-3 A key principle of Meep is that continuously varying inputs yield continuously varying outputs. Here, an eigenfrequency of a photonic crystal varies continuously with the eccentricity of a dielectric rod, accomplished by subpixel smoothing of the material parameters, whereas the nonsmoothed result is "stairstepped." Specifically, the plot shows a TE eigenfrequency of 2d square lattice (period a of dielectric ellipses (ε=12) in air versus one semi-axis diameter of the ellipse (in gradations of 0.005a) for no smoothing (red squares, resolution of 20 pixels/a), subpixel smoothing (blue circles, resolution of 20 pixels/a) and "exact" results (black line, no smoothing at resolution of 200 pixels/a) . 120

- 5-8 The performance of a quasi-PML in the radial direction (cylindrical co-ordinates, left panel) at a resolution 20 pixels/ λ is nearly equivalent to that of a true PML (in Cartesian coordinates, right panel). The plot shows the difference in the electric field E_z (insets) from a point source between simulations with PML thickness L and L + 1, which is a simple proxy for the PML reflections [177]. The different curves are for PML conductivities that turn on as $(x/L)^d$ for d = 1, 2, 3 in the PML, leading to different rates of convergence of the reflection [177]. 127

| 5 - 12 | A simple Meep example showing the E_z field in a dielectric waveg- | |
|--------|---|--------|
| | uide ($\varepsilon = 12$) from a point source at a given frequency. A plot of | |
| | the resulting field (blue/white/red = $positive/zero/negative$) is in the | |
| | background, and in the foreground is the input file in the high-level | |
| | scripting interface (the Scheme language). | 139 |
| 6-1 | Projected band diagram (frequency ω vs. propagation constant β), for | |
| | a triangular lattice of holes (inset). Inset: optimized 2d ($\beta = 0$) gap | |
| | size vs. index contrast | 152 |
| 6-2 | Projected band diagram (frequency ω vs. propagation constant β), for | |
| | a triangular lattice of hexagonal-shaped holes (inset). Inset: optimized | |
| | 2d ($\beta = 0$) gap size vs. index contrast | 153 |
| 6-3 | Doubly-degenerate Γ_6 defect modes for a triangular lattice of hexagonal- | |
| | shaped holes with periodicity a obtained by varying inscribed defect | |
| | diameter of a hexagonal-shaped air core: a) $D = 1.6a$ (fundamental- | |
| | like) b) $D = 3.2a$ c) $D = 6.2a$ d) $D = 6.76a$ (blue/white/red = | |
| | negative/zero/positive). | 154 |
| 6-4 | Fraction of electric-field energy $\varepsilon \mathbf{E} ^2$ in the hexagonal-shaped air core | |
| | (as in Fig. 6-3) as a function of the core radius (radius of inscribed | |
| | circle). Inset: frequency ω at $\beta=0$ of guided mode vs. core radius. | 155 |
| 6-5 | Air-core guided mode in gap of Fig. 6-2, with insets showing electric- | |
| | field E_z and Poynting vector S_z (blue/white/red = negative/zero/positive | e).156 |
| 6-6 | Solid-core guided mode in gap of Fig. 6-2, with insets showing electric- | |
| | field E_z and Poynting vector S_z (blue/white/red = negative/zero/positive | e).157 |
| 6-7 | Left: the square lattice design with the largest fractional gap discovered | |
| | by the nonlinear optimization algorithm where every pixel in the unit | |
| | cell was a free parameter having refractive index in the range $\{1 \dots 3.4\}$. | |
| | Right: a simple two-parameter (radius and width) shape optimization | |
| | based on the adjacent design produces a complete 2d gap at an even | |
| | lower index contrast. | 160 |

- 7-1 Various tapers between uniform and periodic dielectric waveguides. (a) Periodic sequence of holes, where taper varies the radius and period of the holes, in 2d or 3d. (b) Periodic set of flanges, where taper varies the width of the flange, in 2d or 3d. (c) Periodic sequence of dielectric blocks, where taper varies the period Λ between the blocks. All three of these tapers, in 2d or 3d, can be efficiently optimized by the robust coupled-mode method, but this chapter focuses on (c) because it is also amenable to brute-force computation for verification purposes. 169

| Comparison of coupled-mode theory and brute force verification method | |
|---|---|
| (CAMFR [21, 22]) for a linear taper from length of 1 taper period | |
| through 100. The excellent agreement between the fast coupled-mode | |
| theory semi-analytical solver (blue circles) and the much slower brute- | |
| force method (red squares) to compute the objective function permits | |
| use of the former to quickly explore a large parameter space in the | |
| robust optimization. | 183 |
| Comparison of brute force computation (solid circles) and coupled- | |
| mode theory (hollow circles) of reflections from nominal taper designs | |
| optimized for each taper length. The performance of the nominal taper | |
| is clearly ruined by the slight pixellization effects introduced by the | |
| brute-force solver. | 184 |
| Brute force computation of reflections from linear (green), nominal | |
| (blue) and robust (red) taper designs for each length. The superior | |
| performance of the robust tapers, showing an exponential decrease of | |
| the reflection at shorter taper lengths before reaching a noise floor, | |
| is evident under the slight perturbation introduced by the brute-force | |
| solver's pixellization. | 185 |
| Taper profiles of linear (green), nominal (blue) and robust (red) designs | |
| for taper length of 20. The slow-light, periodic waveguide structure is | |
| at $u = 0$ on the left and the standard, strip waveguide is on the right | |
| of the axis. Note the delicate features of the nominal taper which arise | |
| from sensitive interference effects. The robust taper profile varies more | |
| gradually and has superior performance under the slight pixellizations | |
| effects of the brute-force solver | 187 |
| | Comparison of coupled-mode theory and brute force verification method (CAMFR [21, 22]) for a linear taper from length of 1 taper period through 100. The excellent agreement between the fast coupled-mode theory semi-analytical solver (blue circles) and the much slower brute- force method (red squares) to compute the objective function permits use of the former to quickly explore a large parameter space in the robust optimization |

Publications

<u>A.F. Oskooi</u>, D. Roundy, M. Ibanescu, P. Bermel, J.D. Joannopoulos and S.G. Johnson, "MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method," Computer Physics Communications, vol. 181, pp. 687-702 (2010).

<u>A.F. Oskooi</u>, C. Kottke, and S.G. Johnson, "Accurate finite-difference time-domain simulation of anisotropic media by subpixel smoothing," Optics Letters, vol. 34, pp. 2778-2780 (2009).

P.-R. Loh, <u>A.F. Oskooi</u>, M. Ibanescu, M. Skorobogatiy, and S.G. Johnson, "A fundamental relation between phase and group velocity, and application to failure of PML in backward-wave structures," Physical Review E, vol. 79, p. 065601 (2009).

<u>A.F. Oskooi</u>, J. D. Joannopoulos, and S.G. Johnson, "Zero–group-velocity modes in chalcogenide holey photonic-crystal fibers," Optics Express, vol. 17, pp. 10082-10090 (2009).

A. Mutapcic, S. Boyd, <u>A. Farjadpour</u>, S. Johnson, and Y. Avniel, "Robust design of slow-light tapers in periodic waveguides," Engineering Optimization, vol. 41, pp. 365-384 (2009).

<u>A.F. Oskooi</u>, L. Zhang, Y. Avniel, and S.G. Johnson, "The failure of perfectly matched layers, and towards their redemption by adiabatic absorbers," Optics Express, vol. 16, pp. 11376-11392 (2008).

L. Zhang, J.H. Lee, <u>A. Farjadpour</u>, J. White, and S. Johnson, "A novel boundary element method with surface conductive absorbers for 3-D analysis of nanophotonics," Microwave Symposium Digest, IEEE MTT-S International, pp. 523-526 (2008). C. Kottke, <u>A. Farjadpour</u>, and S.G. Johnson, "Perturbation theory for anisotropic dielectric interfaces, and application to sub-pixel smoothing of discretized numerical methods," Physical Review E, vol. 77, p. 036611 (2008).

K. K. Lee, <u>A. Farjadpour</u>, Y. Avniel, J.D. Joannopoulos, and S.G. Johnson, "A tale of two limits: fundamental properties of photonic-crystal fibers," Proc. SPIE, vol. 6901, p. 69010K, January 2008. Invited paper.

<u>A. Farjadpour</u>, D. Roundy, A. Rodriguez, M. Ibanescu, P. Bermel, J.D. Joannopoulos, S.G. Johnson, and G.W. Burr, "Improving accuracy by subpixel smoothing in the finite-difference time domain," Optics Letters, vol. 31, pp. 2972-2974 (2006) [also published in Proc. SPIE, vol. 6322, p. 63220G, (2006)].

G.W. Burr and <u>A. Farjadpour</u>, "Balancing accuracy against computation time: 3-D FDTD for nanophotonics device optimization," Proc. SPIE, vol. 5733, pp. 336-347, January 2005.

Chapter 1

Introduction

Progress in nanophotonics has been inextricably linked with the development of novel numerical modelling and design tools. These tools have been used to study photon modes whose solution are almost always non-analytical and thus necessitate a computational approach. Some common methods include finite-difference time domain (FDTD) [215], finite element (FE) [41], frequency domain planewave expansion eigensolvers [107], transfer matrix [22] and boundary element [245]. Each approach has its own advantages and disadvantages and is well suited for particular tasks: FDTD is commonly employed to compute scattering spectra, cavity resonances and for visualization of field patterns; FE is especially useful for problems involving metals where length scales vary greatly in different media; planewave expansion of Bloch modes for computing dispersion relations of photon modes; and transfer matrix and boundary element methods for scattering phenomena over large distances. These computational tools are increasingly being used to study light-matter interactions in new and novel regimes.

Photonic crystals (PhC) are periodic dielectric structures for which there exists a photonic bandgap when the wavelength of light is comparable to the lengthscale of the periodicity. These nanoscale structures were first studied in one dimension by Lord Rayleigh in 1887 who used wave scattering phenomenon to describe their properties. It was not until one hundred years later that analogues for higher dimensions were proposed for inhibiting spontaneous emission of atoms [241] and localization of light [100]. Among the many applications for PhCs are device components for integrated optics: waveguides [66], filters [67], switches [207] and optical buffers [243]. The vast and growing body of literature in the last twenty years is a testament to the versatility of PhCs in molding the flow of light with unprecedented control.

The ability to fabricate almost any 2d and increasingly 3d patterns at infrared length scales permits a huge number of degrees of freedom in nanophotonic device design. Researchers have hitherto made use of this in simple, primitive structures mainly by using apriori knowledge to guide initial design but more and more are turning to optimization as a means to explore the design space. Unfortunately optimization requires solving a large PDE that can take several hours for an objective function that is both nonlinear and highly non-convex. Furthermore there remains several outstanding issues related to this PDE-based optimization that need addressing: continuously deforming parameters represented on a discrete uniform grid, resolving the breakdown of conventional absorbing boundary layers for important problems, finding robust designs that are insensitive to fabrication and operational variabilities and finally discriminating between local and global minima.

One area of computational design that is becoming increasingly important given recent experimental advances in nanophotonics is that of robust optimization. Owing to delicate interference effects of scattered electromagnetic waves with dielectric structures, slight perturbations arising in device fabrication or operation conditions may significantly deteriorate nominal performance. Such sensitivity mandates robust designs that are insensitive to such irregularities. In robust optimization, one must simultaneously optimize potentially thousands of design parameters while pessimizing uncertainties (optimizing the worst case) and thus a single optimization-based design may require the PDE to be solved thousands or even tens of thousands of times.

Significant improvements can potentially be gained by exploiting intrinsic features of robust optimization: the holy grail would be to marry the optimization and the pessimization with the iterative PDE solver. This would involve updating the design parameters and pessimizing the uncertainties *during* iterations of the PDE solver based on the inexact solution; and would provide the benefit of finding an optimal solution at a cost of iteratively solving the PDE once.

A well-known bottleneck for non-convex problems is that local methods routinely get stuck in local optima; to address this, the optimization is usually repeated numerous times with different initial conditions and the best-performing design ultimately chosen. There are a variety of global optimum approaches but they become exponentially expensive in higher dimensions. A heuristic method we call successive refinement, however, seems to circumvent the issue of getting stuck in poor local minima, finding a simultaneously *more* robust and *more* global solution. Successive refinement consists of solving a sequence of design problems with successively finer resolution starting from the previous coarser design in each case.

When actually solving the PDEs, impediments to computational photonics can act as serious bottlenecks especially when optimizing over a large design space requiring numerous iterations. Two major aspects of computational electromagnetism relevant to nanophotonics that markedly increase computational cost are the fine resolution required to model objects with intricate geometries and the need for minimally reflective absorbing boundary layers for inhomogeneous media. The first is related to the question of how best to model arbitrary dielectric structures on a discrete grid and the second with how best to design absorbing boundary layers with the least numerical footprint.

A major challenge of computational electromagnetism with discrete, uniform grids is in the modelling of non-orthogonal dielectric interfaces not aligned with the grid. This issue is also prevalent in device optimization studies where some shape is continuously varied and accuracy is required to monitor a given optical property (e.g., cavity quality factor, bandgap, transmitted/reflected flux, etc.) without strange jumps from numerical artifacts in the data. Typically such "staircasing" effects often arise in nanophotonics and significantly degrade the accuracy of FDTD simulations. Chapters 2 presents a novel subpixel material averaging scheme for isotropic and anisotropic materials based on rigorous analytical arguments derived from perturbation theory that greatly improve the accuracy of FDTD simulations. We demonstrate its superior performance by comparing it to other previously published subpixel averaging methods. With this new method, researchers can now use modest computing resources to obtain high accuracy for large simulations.

A standard and popular approach to simulate open boundaries with minimal numerical artifacts in computational electromagnetism is the perfectly matched layer (PML). PMLs are absorbing boundary layers that surround the computational cell and are theoretically reflectionless. However, in Chapter 3 we show that PMLs irrecoverably fail when overlapping inhomogeneous media (e.g., photonic crystals) and lead to large reflections. We then demonstrate a simple replacement solution involving adiabatic absorbers that typically perform just as well as PMLs and establish the basic link between the reflections from these absorbers and their corresponding absorber profile. In Chapter 4, we introduce a simple procedure to validate the correctness of *any* PML formulation by providing analytical insights into key characteristics of PMLs. These findings will now make simulations of photonic crystals and other inhomogeneous media more accurate using FDTD and any other numerical methods that use PMLs.

Our developments in subpixel averaging and PML in addition to a number of other improvements have been incorporated into our flexible, free-software package for electromagnetic simulations by the FDTD method known as Meep (an acronym for MIT Electromagnetic Equation Propagation) detailed in Chapter 5. The current range of Meep's functionality permits simulations in 1d, 2d, 3d or cylindrical co-ordinates of a large class of electromagnetic phenomena involving arbitrary anisotropic, nonlinear, dispersive, active and conductive materials. Meep is also fully parallelized so that it can run on supercomputers for large applications.

Next, we demonstrate the utility of optimization in nanophotonics by focusing on two important design problems.

Chapter 6 describes a new class of holey photonic-crystal fibers. These are fibers with a 2d photonic crystal in the cladding where the complete 2d bandgap extended over a range of axial propagation wavevectors acts to confine slow light to the waveguide core which is either solid or air. The principal aim of this research was to determine how low an index contrast could be found for various 2d photonic crystal geometries so as to enable a larger class of materials for the fabrication of such fibers for slow-light applications. The challenge here is that the smaller the index contrast, the more difficult the problem becomes of finding a structure with a complete bandgap so we decided to use an optimization strategy to guide the design. We were able to find designs for an index contrast as low as 2.6:1 which now permits for the first time the use of an important class of materials known as chalcogenide glasses to be used in the fiber drawing process.

Chapter 7 focuses on designing waveguide tapers that can be used to couple an optical mode from a standard dielectric (or strip) waveguide into the slow-light mode of a photonic-crystal waveguide. Slow-light optical modes are important for the investigation of a number of interesting physical phenomena, such as nonlinearities, gain, tunability and magneto-optics. The principal aim of this project was to design waveguide couplers with maximal transmission of photons over a narrow range of frequencies near the band edge that are also insensitive to fabrication imperfections. The latter point related to robustness is key as slow-light optical modes are strongly affected by any irregularities or disorder in the device. This work was conducted in collaboration with the optimization research group of Professor Stephen Boyd at Stanford University where we combined our group's nanophotonic design expertise and semi-analytical numerical tools for electromagnetism with their advanced optimization toolbox. We have been successful in designing waveguide tapers for a number of different 2d structures that operate in the challenging slow-light regime, were designed with rigorous optimization methods and had very small reflections (less than 1%) suitable for telecommunication applications.

Chapter 2

Sub-pixel smoothing for dielectric media

2.1 Summary

Finite-difference time-domain (FDTD) methods suffer from reduced accuracy when modelling discontinuous dielectric materials, due to the inherent discretization ("pixellization"). We show that accuracy can be significantly improved by using a sub-pixel smoothing of both the isotropic and anisotropic dielectric function, but only if the smoothing scheme is properly designed. We develop such a scheme based on a criterion taken from perturbation theory, and compare it to other published FDTD smoothing methods. In addition to consistently achieving the smallest errors, our scheme is the only one that attains quadratic convergence with resolution for arbitrarily sloped interfaces. Finally, we discuss additional difficulties that arise for sharp dielectric corners.

2.2 Overview

A popular numerical tool for photonics is the finite-difference time-domain (FDTD) method, which discretizes Maxwell's equations on a grid in space and time [215]. Here, we address difficulties in representing a discontinuous permittivity (ε) on such a grid,

by reviewing previously proposed anisotropic sub-pixel ε smoothing schemes adapted from spectral methods [107, 124, 153]. The work in this chapter provides a clear justification for the second-order convergence of prior proposed schemes for isotropic media in terms of perturbation theory. It enables extension to anisotropic media and clarifies the role of sharp corners. We show that our method consistently achieves the smallest errors compared to previous smoothing schemes for FDTD [56, 110, 160]. Subpixel smoothing has an additional benefit: it allows the simulation to respond continuously to changes in the geometry, such as during optimization or parameter studies, rather than changing in discontinuous jumps as interfaces cross pixel boundaries. This technique additionally yields much smoother convergence of the error with resolution, which makes it easier to evaluate the accuracy and enables the possibility of extrapolation to gain another order of accuracy [235]. Unlike methods that require modified field-update equations [57] or larger stencils and complicated position-dependent difference equations for higher order accuracy [255], our method uses the standard center-difference expressions and is easy to implement requiring only preprocessing of the materials (free code is available [176]).

The presence of material discontinuities in degrading the order of accuracy of underlying finite differences has been prevalent in a number of other computational schemes involving solid, fluid and heat equations [78, 134, 143]. Here we investigate the effects of such discontinuities in computational electromagnetism. When ε is represented by "pixels" on a grid (or "voxels" in 3d), two difficulties arise. First, a uniform grid makes it more difficult to model small features or to optimize device performance by continuous variation of geometric parameters. Second, the pixellized ε may be a poor representation of the dielectric function: diagonal interfaces produce "staircasing," and even interfaces aligned with the grid may be shifted by as much as a pixel. This increases the computational errors, and can even degrade the rate of convergence with the grid resolution—as was pointed out in Ref. 57, ε interfaces actually reduce the *order* of convergence from the nominal quadratic (error ~ Δx^2) of standard FDTD to only linear (error ~ Δx). We address both of these difficulties in this chapter.

2.3 Designing subpixel smoothing algorithms with perturbation theory

Our basic approach is to smooth the structure to eliminate the discontinuity before discretizing, but because the smoothing itself changes the geometry we use first-order perturbation theory to select a smoothing with zero first-order effect. For isotropic materials, this approach makes rigorous a smoothing scheme that had previously been proposed heuristically [132, 153] and we can now explain its second-order accuracy. Advances in perturbation theory have enabled us to extend this scheme to interfaces between anisotropic materials, initially for a planewave method [124]. Here, we adapt the technique to FDTD, combined with a recent FDTD scheme with improved stability for anisotropic media [235]. Although this chapter focuses on the case of isotropic and anisotropic electric permittivity ε , exactly the same smoothing and discretization schemes apply to magnetic permeabilities μ due to the equivalence in Maxwell's equations under interchange of ε/μ and \mathbf{E}/\mathbf{H} .

There are many ways to formulate perturbation techniques in electromagnetism. One common formulation, analogous to "time-independent perturbation theory" in quantum mechanics [48], is to express Maxwell's equations as a generalized Hermitian eigenproblem $\nabla \times \nabla \times \mathbf{E} = \omega^2 \varepsilon \mathbf{E}$ in the frequency ω and electric field \mathbf{E} (or equivalent formulations in terms of the magnetic field \mathbf{H}) [99], and then to consider the firstorder change $\Delta \omega$ in the frequency from a small change $\Delta \varepsilon$ in the dielectric function $\varepsilon(\mathbf{x})$ (assumed real and positive), which turns out to be [99]:

$$\frac{\Delta\omega}{\omega} = -\frac{\int \mathbf{E}^* \cdot \Delta\varepsilon \mathbf{E} \, d^3 \mathbf{x}}{2 \int \mathbf{E}^* \cdot \varepsilon \mathbf{E} \, d^3 \mathbf{x}} + O(\Delta\varepsilon^2), \qquad (2.1)$$

where **E** and ω are the electric field and eigenfrequency of the *unperturbed* structure ε , respectively, and * denotes complex conjugation.

As was shown by Johnson et al. [105], eq. (2.1) is not valid when $\Delta \varepsilon$ is due to a small change in the position of a boundary between two dielectric materials (except in the limit of low dielectric contrast), but a simple correction is possible. In particular,



Figure 2-1: Schematic of an interface perturbation: the interface between two materials ε^a and ε^b (possibly anisotropic) is shifted by some small position-dependent displacement h.

let us consider situations like the one shown in Fig. 2-1, where the dielectric boundary between two materials ε^a and ε^b is shifted by some small displacement h (which may be a function of position). Directly applying eq. (2.1), with $\Delta \varepsilon = \pm (\varepsilon^a - \varepsilon^b)$ in the regions where the material has changed, gives an incorrect result, and in particular $\Delta \omega / h$ (which should ideally go to the exact derivative $d\omega / dh$) is incorrect even for $h \to 0$. The problem turns out to be not so much that $\Delta \varepsilon$ is not small, but rather that **E** is discontinuous at the boundary, and the standard method in the limit $h \to 0$ leads to an ill-defined surface integral of **E** over the interfaces. For *isotropic* materials, corresponding to *scalar* $\varepsilon^{a,b}$, the correct numerator instead turns out to be the following surface integral over the boundary as shown by Johnson et al. [105]:

$$\int \mathbf{E}^* \cdot \Delta \varepsilon \mathbf{E} \, d^3 \mathbf{x} \longrightarrow$$

$$\iint \left[\left(\varepsilon^a - \varepsilon^b \right) \left| \mathbf{E}_{\parallel} \right|^2 - \left(\frac{1}{\varepsilon^a} - \frac{1}{\varepsilon^b} \right) \left| D_{\perp} \right|^2 \right] \mathbf{h} \cdot d\mathbf{A}, \quad (2.2)$$

where \mathbf{E}_{\parallel} and D_{\perp} are the (continuous) components of \mathbf{E} and $\mathbf{D} = \varepsilon \mathbf{E}$ parallel and perpendicular to the boundary, respectively, $d\mathbf{A}$ points towards ε^{b} , and \mathbf{h} is the displacement of the interface from ε^{a} towards ε^{b} .

In our previous work spearheaded by Chris Kottke [124] which is briefly reviewed here, we generalized eq. (2.2) to handle the case where the two materials are *anisotropic*, corresponding to arbitrary 3×3 tensors ε^a and ε^b (assumed Hermitian



Figure 2-2: Schematic 2d Yee FDTD discretization near a dielectric interface, showing the method [235] used to compute the part of E_x that comes from D_y and the locations where various ε^{-1} components are required.

and positive-definite to obtain a well-behaved Hermitian eigenproblem). In the generalized case, it is convenient to define a local coordinate frame (x_1, x_2, x_3) at each point on the surface, where the x_1 direction is orthogonal to the surface and the (x_2, x_3) directions are parallel. We also define a continuous field "vector" $\mathbf{F} = (D_1, E_2, E_3)$ so that $F_1 = D_{\perp}$ and $\mathbf{F}_{2,3} = \mathbf{E}_{\parallel}$. The resulting numerator of eq. (2.1), generalizing eq. (2.2), Kottke et al. showed to be:

$$\iint \mathbf{F}^* \cdot \left[\boldsymbol{\tau} \left(\boldsymbol{\varepsilon}^a \right) - \boldsymbol{\tau} \left(\boldsymbol{\varepsilon}^b \right) \right] \cdot \mathbf{F} \, \mathbf{h} \cdot d\mathbf{A}, \tag{2.3}$$

where $\boldsymbol{\tau}(\boldsymbol{\varepsilon})$ is the 3 × 3 matrix in eq. (2.6) which reduces to eq. (2.2) when $\boldsymbol{\varepsilon}$ is a scalar multiple ε of the identity matrix. (Our assumption that $\boldsymbol{\varepsilon}$ is positive-definite guarantees that $\varepsilon_{11} > 0$).

We define an interface-relative coordinate frame as in Fig. 2-2, so that the first component "1" is the direction normal to the interface. Previously, for an interface between two *isotropic* materials ε^a and ε^b , Meade et al. [153] showed (without rigor-

ous analytical arguments) that the proper smoothed permittivity (in this coordinate frame) at each point is:

$$\tilde{\boldsymbol{\varepsilon}} = \begin{pmatrix} \langle \varepsilon^{-1} \rangle^{-1} & 0 & 0 \\ 0 & \langle \varepsilon \rangle & 0 \\ 0 & 0 & \langle \varepsilon \rangle \end{pmatrix}, \qquad (2.4)$$

where $\langle \cdots \rangle$ denotes an average over one pixel. Equation (2.4) uses the mean $\langle \varepsilon \rangle$ for the surface-parallel **E** components and the harmonic mean $\langle \varepsilon^{-1} \rangle^{-1}$ for the surface-perpendicular component. For an interface between *anisotropic* materials, Kottke et al. [124] showed that the following subpixel smoothing scheme is the appropriate choice (having zero first-order perturbation) [124]:

$$\tilde{\boldsymbol{\varepsilon}} = \tau^{-1} \left[\langle \tau(\boldsymbol{\varepsilon}) \rangle \right], \tag{2.5}$$

where $\tau(\varepsilon)$ and its inverse are defined by

$$\tau\left(\varepsilon\right) = \begin{pmatrix} -\frac{1}{\varepsilon_{11}} & \frac{\varepsilon_{12}}{\varepsilon_{11}} & \frac{\varepsilon_{13}}{\varepsilon_{11}} \\ \frac{\varepsilon_{21}}{\varepsilon_{11}} & \varepsilon_{22} - \frac{\varepsilon_{21}\varepsilon_{12}}{\varepsilon_{11}} & \varepsilon_{23} - \frac{\varepsilon_{21}\varepsilon_{13}}{\varepsilon_{11}} \\ \frac{\varepsilon_{31}}{\varepsilon_{11}} & \varepsilon_{32} - \frac{\varepsilon_{31}\varepsilon_{12}}{\varepsilon_{11}} & \varepsilon_{33} - \frac{\varepsilon_{31}\varepsilon_{13}}{\varepsilon_{11}} \end{pmatrix}, \qquad (2.6)$$

$$\tau^{-1}[\tau] = \begin{pmatrix} -\frac{1}{\tau_{11}} & -\frac{\tau_{12}}{\tau_{11}} & -\frac{\tau_{13}}{\tau_{11}} \\ -\frac{\tau_{21}}{\tau_{11}} & \tau_{22} - \frac{\tau_{21}\tau_{12}}{\tau_{11}} & \tau_{23} - \frac{\tau_{21}\tau_{13}}{\tau_{11}} \\ -\frac{\tau_{31}}{\tau_{11}} & \tau_{32} - \frac{\tau_{31}\tau_{12}}{\tau_{11}} & \tau_{33} - \frac{\tau_{31}\tau_{13}}{\tau_{11}} \end{pmatrix}.$$
 (2.7)

The derivation of this result is nontrivial and is explained in Kottke et al. [124] and we will not repeat it here, but we point out that eq. (2.4) is now obtained as the special case for isotropic ε .

2.4 Analysis of smoothing perturbation

Here we review work first outlined by us for isotropic media [69] and later generalized by Kottke et al. to anisotropic media [124]. In any numerical method involving the


Figure 2-3: TE eigenfrequency error vs. resolution for a Bragg mirror of alternating air and $\varepsilon = 12$ (inset).

solution of the full-vector Maxwell's equations on a discrete grid or its equivalent, such as the planewave method above [107] or the finite-difference time-domain (FDTD) method [215], discontinuities in the non-discretized dielectric function ε (and the corresponding field discontinuities) generally degrade the accuracy of the method, typically reducing it to only linear convergence with resolution [57, 107]. Unfortunately, piecewise-continuous ε is the most common situation for computational simulations, so a technique to improve the accuracy (without switching to an entirely different computational method) is desirable. One simple approach that has been proposed by several authors is to smooth the dielectric function, or equivalently to set the ε of each "pixel" to be some average of ε within the pixel, rather than merely sampling ε in a "staircase" fashion [56,107,110,132,153,160,166]. Unfortunately, this smoothing itself changes the structure, and therefore introduces errors. The problem is closely related to perturbation theory: one desires a smoothing of ε that has zero first-order effect, to minimize the error introduced by smoothing and so that the underlying secondorder accuracy can potentially be preserved. At an interface between two isotropic dielectric materials, the first-order perturbation is given by eq. (2.2), and this leads to an *anisotropic* smoothing: one averages ε^{-1} for field components perpendicular to the interface, and averages ε for field components parallel to the interface, a result that had previously been proposed heuristically by several authors [107, 132, 153].

In this section, we generalize that result to interfaces between anisotropic materials, and illustrate numerically in the following sections that it leads to both dramatic improvements in the absolute magnitude and the convergence rate of the discretization error. In the anisotropic-interface case, a heuristic subpixel smoothing scheme was previously proposed [107], but Kottke et al. [124] showed that this method was suboptimal: although it is better than other smoothing schemes, it does not set the first-order perturbation to zero and therefore does not minimize the error or permit the possibility of second-order accuracy. Specifically, as discussed more explicitly below, a second-order smoothing is obtained by averaging $\tau(\varepsilon)$ and then inverting $\tau(\varepsilon)$ to obtain the smoothed "effective" dielectric tensor. Because this scheme is analytically guaranteed to eliminate the first-order error otherwise introduced by smoothing,



Figure 2-4: TE eigenfrequency error vs. resolution for a square lattice of elliptical air holes in $\varepsilon = 12$ (inset).



Figure 2-5: TM eigenfrequency error vs. resolution for a square lattice of elliptical air holes in $\varepsilon = 12$ (inset).

we expect it to generally lead to the smallest numerical error compared to competing smoothing schemes, and there is the hope that the overall convergence rate may be quadratic with resolution.

First, let us analyze how perturbation theory leads to a smoothing scheme. Suppose that we smooth the underlying dielectric tensor $\boldsymbol{\varepsilon}(\mathbf{x})$ into some locally averaged tensor $\bar{\boldsymbol{\varepsilon}}(\mathbf{x})$, by some method to be determined below. This involves a change $\Delta \boldsymbol{\varepsilon} = \bar{\boldsymbol{\varepsilon}} - \boldsymbol{\varepsilon}$, which is likely to be large near points where $\boldsymbol{\varepsilon}$ is discontinuous (and, conversely, is zero well inside regions where $\boldsymbol{\varepsilon}$ is constant). In particular, suppose that we employ a smoothing radius (defined more precisely below) proportional to the spatial resolution Δx of our numerical method, so that $\Delta \boldsymbol{\varepsilon}$ is zero [or at most $O(\Delta x^2)$] except within a distance $\sim \Delta x$ of discontinuous interfaces. To evaluate the effect of this large perturbation near an interface, we must employ an equivalent reformulation of eq. (2.3):



Figure 2-6: Eigenfrequency error vs. resolution for a cubic lattice of $\varepsilon = 12$ ellipsoids in air (inset).

$$\Delta \omega \sim \int \mathbf{F}^* \cdot \Delta \boldsymbol{\tau} \cdot \mathbf{F} \, d^3 \mathbf{x}, \qquad (2.8)$$

where $\Delta \boldsymbol{\tau} = \boldsymbol{\tau}(\bar{\boldsymbol{\varepsilon}}) - \boldsymbol{\tau}(\boldsymbol{\varepsilon})$. It is sufficient to look at the perturbation in ω , since the same integral appears in the perturbation theory for many other quantities (such as scattered power, etc.). If we let x_1 denote the (local) coordinate orthogonal to the boundary, then the x_1 integral is simply proportional to $\sim \int \Delta \boldsymbol{\tau} \, dx_1 + O(\Delta x^2)$: since \mathbf{F} is continuous and $\Delta \boldsymbol{\tau} = 0$ except near the interface, we can pull \mathbf{F} out of the x_1 integral to lowest order. That means, in order to make the first-order perturbation zero for all fields \mathbf{F} , it is sufficient to have $\int \Delta \boldsymbol{\tau} \, dx_1 = 0$. This is achieved by averaging $\boldsymbol{\tau}$ as follows.

The most straightforward interpretation of "smoothing" would be to convolve ε with some localized kernel $s(\mathbf{x})$, where $\int s(\mathbf{x}) d^3 \mathbf{x} = 1$ and $s(\mathbf{x}) = 0$ for $|\mathbf{x}|$ greater than some smoothing radius (the support radius) proportional to the resolution $\sim \Delta x$. That is, $\bar{\boldsymbol{\varepsilon}}(\mathbf{x}) = \boldsymbol{\varepsilon} * s = \int \boldsymbol{\varepsilon}(\mathbf{y}) s(\mathbf{x} - \mathbf{y}) d^3 \mathbf{y}$. For example, the simplest subpixel smoothing, simply computing the average of $\boldsymbol{\varepsilon}$ over each pixel, corresponds to s = 1 inside a pixel at the origin and s = 0 elsewhere. However, this will not lead to the desired $\int \Delta \boldsymbol{\tau} = 0$ to obtain second order accuracy. Instead, we employ:

$$\bar{\boldsymbol{\varepsilon}}(\mathbf{x}) = \boldsymbol{\tau}^{-1}[\boldsymbol{\tau}(\boldsymbol{\varepsilon}) * s] = \boldsymbol{\tau}^{-1} \left\{ \int \boldsymbol{\tau} \left[\boldsymbol{\varepsilon}(\mathbf{y}) \right] s(\mathbf{x} - \mathbf{y}) d^3 \mathbf{y} \right\}.$$
 (2.9)

where τ^{-1} is the inverse of the $\tau(\varepsilon)$ mapping, given by eqs. (2.7) and (2.6) respectively.

The reason why eq. (2.9) works, regardless of the smoothing kernel $s(\mathbf{x})$, is that

$$\int \Delta \boldsymbol{\tau} \, d^3 \mathbf{x} = \int d^3 \mathbf{x} \left\{ \int \boldsymbol{\tau} \left[\boldsymbol{\varepsilon}(\mathbf{y}) \right] \, s(\mathbf{x} - \mathbf{y}) \, d^3 \mathbf{y} - \boldsymbol{\varepsilon}(\mathbf{x}) \right\}$$
$$= \int d^3 \mathbf{y} \, \boldsymbol{\tau} \left[\boldsymbol{\varepsilon}(\mathbf{y}) \right] \left\{ \int \, s(\mathbf{x} - \mathbf{y}) \, d^3 \mathbf{y} - 1 \right\}$$
$$= 0. \tag{2.10}$$

This guarantees that the integral of $\Delta \tau$ is zero over all space, but above we required what appears to be a stronger condition, that the local, interface-perpendicular integral $\int \Delta \tau \, dx_1$ be zero (at least to first order). However, in a small region where the interface is locally flat (to first order in the smoothing radius), $\Delta \tau$ must be a function of x_1 only by translational symmetry, and therefore eq. (2.10) implies that $\int \Delta \tau \, dx_1 = 0$ by itself. Although the above convolution formulas may look complicated, for the simplest smoothing kernel $s(\mathbf{x})$ the procedure is quite simple: in each pixel, average $\tau(\varepsilon)$ in the pixel and then apply τ^{-1} to the result. (This is not any more difficult to apply than the procedure implemented in Ref. 107, for example.)

Strictly speaking, the use of this smoothing does not guarantee second-order accuracy, even if the underlying numerical method is nominally second-order accurate or better. For one thing, although we have canceled the first-order error due to smoothing, it may be that the next-order correction is not second-order. Precisely



Figure 2-7: Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a square lattice (period *a*) of 2d anisotropic ellipsoids (right inset) versus spatial resolution (units of pixels per vacuum wavelength λ), for a variety of subpixel smoothing techniques. Straight lines for perfect linear (black dashed) and perfect quadratic (black solid) convergence are shown for reference. Most curves are for the first eigenvalue band (left inset shows E_z in unit cell), with vacuum wavelength $\lambda = 4.85a$. Hollow squares show new method for band 15 (middle inset), with $\lambda = 1.7a$. Maximum resolution for all curves is 100 pixels/a.

this situation occurs if one has a structure with sharp dielectric corners, edges, or cusps, as discussed in Ref. 105: in this case, smoothing leads to a convergence rate between first order (what would be obtained with no smoothing) and second order, with the exponent determined by the nature of the field singularity that occurs at the corner. This is discussed in more detail in Ch. 2.8 below.

2.5 Stable FDTD field-update implementation

An additional difficulty for anisotropic material tensors occurs in FDTD: to accurately discretize the spatial derivatives, each field component is discretized on a different grid. In the standard Yee discretization for grid coordinates [i, j, k] = $(i\Delta x, j\Delta y, k\Delta z)$, the E_x and D_x components are discretized at [i + 0.5, j, k] while E_y/D_y are at [i, j + 0.5, k] and E_z/D_z are at [i, j, k + 0.5] [215]. At each time step, $\mathbf{E} = \varepsilon^{-1} \mathbf{D}$ must be computed, but any off-diagonal parts of ε couple components stored at different locations. For example, a nonzero $(\varepsilon_{xy})^{-1}$ means that the computation of E_x requires D_y , but the value of D_y is not available at the same grid point as E_x , as depicted in Fig. 2-2. One approach is to average the four adjacent D_y values and use them in updating E_x , along with $(\varepsilon_{xy})^{-1}$ at the E_x point [69,235]. This approach, however, is theoretically unstable and leads to divergences for a long simulation [235]. Instead, a modified technique was recently shown to satisfy a necessary condition for stability with Hermitian ε [235]: as depicted in Fig. 2-2, one first averages D_y at $[i, j \pm 0.5, k]$ and multiplies by $(\varepsilon_{xy})^{-1}$ at [i, j, k], and then averages the two results at [i, j, k] and [i + 1, j, k] to update E_x at [i + 0.5, j, k]. (Although Ref. 235 derives no *sufficient* condition for stability with inhomogeneous media once the Yee time discretization is included, this method has been stable in all numerical experiments to date.) We use this scheme here, and find that it greatly improves stability compared to the simpler scheme from the previous chapter [69]. The subpixel averaging is performed as follows. At the E_x point [i+0.5, j, k] (orange dot in Fig. 2-2), the smoothed $\tilde{\varepsilon}$ is computed by eq. (2.5), averaging over the pixel centered at that point. Then, $\tilde{\boldsymbol{\varepsilon}}$ is inverted to obtain $(\tilde{\boldsymbol{\varepsilon}}^{-1})_{xx}$, which is stored at the E_x point. The subpixel averaging $\tilde{\varepsilon}$ is also performed for a pixel centered at the [i, j, k] point (blue dot) halfway between two D_y points (red dot), and $(\tilde{\varepsilon}^{-1})_{xy}$ is computed and stored at that point. Similarly for other components. (Note that the $\tilde{\varepsilon}$ tensor from eq. (2.5) must be rotated from the interface-normal to Cartesian coordinates at each point.) Thus, for each Yee cell in three dimensions, the subpixel averaging is performed four times, obtaining $(\tilde{\varepsilon}^{-1})_{xx}$ at [i+0.5, j, k], $(\tilde{\varepsilon}^{-1})_{yy}$ at [i, j+0.5, k], $(\tilde{\varepsilon}^{-1})_{zz}$ at [i, j, k+0.5], and all off diagonal components are at [i, j, k]—in other words, we apply the same averaging procedure eq. (2.5) to pixels centered around different points/corners in the Yee cell, and then for each point we store only the components of ε^{-1} necessary for that point. Each component of ε^{-1} need only be stored at most once per Yee cell, so no additional storage is required compared to other anisotropic FDTD schemes. After this smoothing, the anisotropic FDTD scheme proceeds without modification.

2.6 Numerical performance of methods for isotropic media

To evaluate the discretization error, we compute an eigenfrequency ω of a periodic (square or cubic, period a) lattice of dielectric shapes with 12:1 ε contrast, a photonic crystal [99]. In particular, we compute the smallest ω for an arbitrarily chosen Bloch wavevector **k** (not aligned with the grid), so that the wavelength is comparable to the feature sizes. We perform an FDTD simulation with Bloch-periodic boundaries and a Gaussian pulse source, analyzing the response with a filter-diagonalization method [145] to obtain the eigenfrequency ω . This is compared to the "exact" ω_0 from a planewave calculation [107] at a very high resolution, plotting the relative error $|\omega - \omega_0|/\omega_0$ versus FDTD resolution. ω is a good proxy for other common computations, because both the change in the frequency and the scattered power for a small $\Delta \varepsilon$ go as $\Delta \varepsilon |\mathbf{E}|^2$ to lowest order [108]. We compare to three other smoothings. The simplest is to use the scalar mean $\langle \varepsilon \rangle$ for all components, [56] which is incorrect for the surface-normal fields. Kaneda [110] proposed an anisotropic smoothing that leads



Figure 2-8: Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a cubic lattice (period a) of 3d anisotropic ellipsoids (right inset) versus spatial resolution (units of pixels per vacuum wavelength λ), for a variety of subpixel smoothing techniques. Straight lines for perfect linear (black dashed) and perfect quadratic (black solid) convergence are shown for reference. Most curves are for the first eigenvalue band (left inset shows E_x in xy cross-section of unit cell), with vacuum wavelength $\lambda = 5.15a$. Hollow squares show new method for band 13 (middle inset), with $\lambda = 2.52a$. New method for bands 1 and 13 is shown for resolution up to 100 pixels/a.



Figure 2-9: Relative error $\Delta \omega / \omega$ for an eigenmode calculation with a square lattice (period *a*) of 2d anisotropic ellipses (green inset), versus spatial resolution, for a variety of sub-pixel smoothing techniques. Straight lines for perfect linear (black dashed) and perfect quadratic (black solid) convergence are shown for reference.

to diagonal $\tilde{\varepsilon}^{-1}$ tensors. We also consider the "VP-EP" scheme, [160] which is exactly the diagonal part of eq. (2.4) for s = 1. Both Kaneda and VP-EP are equivalent to eq. (2.4) for flat interfaces oriented along the grid (xyz) directions, but they do not satisfy the perturbation criterion for diagonal interfaces. Yet another method [166] was found to be numerically unstable for our test cases, which prevented us from evaluating it; however, it is equivalent to eq. (2.4) only for flat x/y/z interfaces. Other schemes, not considered here, were developed for perfect conductors [215, 249] or for non-Yee lattices in 2d. [164]

To start with, we look at a 1d case in Fig. 2-3 where Kaneda, VP-EP, and Nadobny are equivalent to our method: a distributed bragg reflector (DBR) along the x direction, with a **k** vector in the xy plane so that the eigenfield **E** has components both parallel and perpendicular to the interfaces. We find that both the no-smoothing and simple mean- ε cases both have only linear convergence, whereas the new method



Figure 2-10: Relative error $\Delta \omega / \omega$ for an eigenmode calculation with cubic lattice (period *a*) of 3d anisotropic ellipsoids (green inset), versus spatial resolution, for a variety of sub-pixel smoothing techniques. Straight lines for perfect linear (black dashed) and perfect quadratic (black solid) convergence are shown for reference.

(and Kaneda and Nadobny) have quadratic convergence.

Since Kaneda, VP-EP, and our method are equivalent for grid-parallel interfaces (and we obtain quadratic convergence for all these methods), we focus instead on a more complicated case: a square lattice of elliptical air holes shown in the inset of Fig. 2-4, for the TE polarization (**E** in the 2d plane). Our new method (hollow squares) has the smallest errors by large margin, while the Kaneda and VP-EP methods are actually worse than no smoothing. As mentioned above, all methods except ours converge linearly, whereas we expect our method to be asymptotically quadratic. As a trick to make the quadratic convergence of our method more apparent, we double the smoothing diameter to s = 2 (filled squares), at the expense of increasing the absolute error.

The TM polarization (**E** out of the plane) is shown in Fig. 2-5, but is less interesting: all the smoothing methods are equivalent to the simple mean ε , all decrease the error compared to no smoothing, and *all* methods (including no smoothing) exhibit quadratic convergence. Since **E** is everywhere continuous, TM is the "easy" case for numerical computation (and perturbative methods [105, 108]).

In three dimensions, we used a cubic lattice of $\varepsilon = 12$ ellipsoids, with an arbitrary orientation, in air. The results in Fig. 2-6 again show that the new method has the smallest error, and is again quadratic. Notice that the ordering of the other methods has changed, and in general we observe them to yield erratic accuracy.

2.7 Numerical performance of methods for anisotropic media

To illustrate the discretization error for the anisotropic smoothing algorithm of eqs. (2.6) and (2.7), we repeat similar numerical experiments where we compute an eigenfrequency ω of a periodic (square in 2d or cubic in 3d, period *a*) lattice of dielectric ellipsoids made of ε^a surrounded by ε^b , a photonic crystal [99]. We choose $\varepsilon^{a,b}$ to be random positive-definite symmetric matrices with eigenvalues in the interval [1, 5] (1.45, 2.81, and 4.98) for ε^a and in [9, 12] (8.49, 8.78, and 11.52) for ε^b . We compute the lowest ω for an arbitrary Bloch wavevector $\mathbf{k} = (0.4, 0.2, 0.3)\frac{2\pi}{a}$, giving wavelengths comparable to the feature sizes. In an FDTD simulation with Bloch-periodic boundaries and a Gaussian pulse source, we analyze the response with a filter-diagonalization method [145] to obtain the eigenfrequency ω , obtaining the relative error $|\omega - \omega_0|/\omega_0$ by comparison with the "exact" ω_0 from a planewave calculation [107] at a high resolution. We looked at eigenvalue bands 1 and 15 (in 2d) or 1 and 13 (in 3d), where the higher band is clearly nonplanewave-like (see inset fields), to counter suggestions that subpixel averaging may perform poorly for higher bands [235].

We compare the new smoothing technique of eq. (2.5) to the nonsmoothed case as well as to two simple smoothing techniques: using the mean $\langle \varepsilon \rangle$ [56] and also the harmonic mean $\langle \varepsilon^{-1} \rangle^{-1}$. We do not compare to a previous heuristic that we had proposed without the benefit of perturbation theory [107], since the previous chapter already demonstrated that this heuristic (which does not yield zero firstorder perturbation) is much less accurate than the new method [124], and first-order FDTD accuracy for that heuristic was also shown in Ref. 235 [who did not examine the isotropic case where eq. (2.4) remains correct].

Results from 2d and 3d simulations are shown in Fig. 2-7 and Fig. 2-8, respectively. In both cases, similar to our previous results for isotropic materials [69], the new smoothing algorithm has the lowest error, often by an order of magnitude or more, and is the only technique that appears to give second-order accuracy in the limit of high resolution. (The simple mean $\langle \varepsilon \rangle$ does better than the harmonic mean $\langle \varepsilon^{-1} \rangle^{-1}$, probably because it treats roughly two of the three field components correctly [69].) Similar accuracy is obtained for both lower and higher (non-planewave-like) bands at comparable resolutions per wavelength (although higher bands require greater absolute resolution per *a*, of course, because their wavelengths are smaller). As we have noted, apparent quadratic convergence obtained in a single structure [107] can sometimes be fortuitous [124], but we have confidence in these results (obtained now in multiple settings) because they are backed by a clear theory rather than an *ad hoc* heuristic.

We repeat a similar experiment this time using a planewave method consisting of a preconditioned conjugate-gradient minimization of the block Rayleigh quotient from a free-software package [107]. As before, we first consider a two-dimensional example problem: a square lattice (period *a*) of ellipses made of ε^a surrounded by ε^b , where we will find the lowest- ω Bloch eigenmode. As above, we choose the dielectric tensors to be random positive-definite symmetric matrices with random eigenvalues in [2, 12] for ε^a and in [1, 5] for ε^b , and the ellipses are oriented at an arbitrary angle, at an arbitrary Bloch wavevector $\mathbf{k}a/2\pi = (0.1, 0.2, 0.3)$, to avoid fortuitous symmetry effects. (The vacuum wavelength λ corresponding to the eigenfrequency ω is $\lambda = 5.03a$.) For each resolution Δx , we assign an $\overline{\varepsilon}$ to each pixel by computing τ^{-1} of the average of $\tau(\varepsilon)$ within that pixel. Then, we compute the relative error $\Delta \omega/\omega$ (compared to a calculation at a much higher resolution) as a function of resolution. For comparison, we also consider four other smoothing techniques: no smoothing,



Figure 2-11: Degraded accuracy due to field singularities at sharp corners: TE eigenfrequency error vs. resolution for square lattice of tilted-square air holes in $\varepsilon = 12$ (inset).

averaging ε in each pixel [56], averaging ε^{-1} in each pixel, and a heuristic anisotropic averaging proposed by Ref. 107 in analogy to the scalar case. The results are shown in Fig. 2-9 and show that the new smoothing technique clearly leads to the lowest errors $\Delta \omega / \omega$. Also, whereas the other methods yield clearly first-order convergence, the new method seems to exhibit roughly second-order convergence. The no-smoothing case has extremely erratic errors, as is typical for stair-casing phenomena.

In Fig. 2-10, we also show results from a similar calculation in three dimensions. Here, we look at the lowest eigenmode of a cubic lattice (period a) of 3d ellipsoids (oriented at a random angle) made of ε^a surrounded by ε^b , both random positivedefinite symmetric matrices as above. The frequency ω , at an arbitrarily chosen wavevector $\mathbf{k}a/2\pi = (0.4, 0.3, 0.1)$, corresponds to a vacuum wavelength $\lambda = 3.14a$. Again, the new method almost always has the lowest error by a wide margin, especially if the unpredictable dips of the no-smoothing case are excluded, and is the only one to exhibit (apparently) better than linear convergence.

Our previous heuristic proposal from Ref. 107, while better than the other smoothing schemes (and less erratic than no smoothing), is clearly inferior to the new method. Previously, we had observed what seemed to have been quadratic convergence from the heuristic scheme [107], but this result seems to have been fortuitous—as we demonstrated recently, even non-second-order schemes can sometimes appear to have second-order convergence over some range of resolutions for a particular geometry [69]. The key distinction of the new scheme, that lends us greater confidence in it than one or two examples can convey, is that it is no longer heuristic. The new smoothing scheme is based on a clear analytical criterion—setting the first-order perturbative effect of the smoothing to zero—that explains why it should be an accurate choice in a wide variety of circumstances.

2.8 Field singularities at sharp corners

Finally, we consider a qualitatively different case, in which *none* of the methods satisfy our zero-perturbation criterion: the presence of a sharp corner leads to a new field singularity. Figure 2-11 shows the error for a square lattice of tilted air squares in $\varepsilon = 12$ (inset). Because our new method at least handles the flat edges properly, it still has lower error than other smoothing schemes, although suboptimal handling of the corner limits the differences. Fits of this data indicate that our method seems to be converging as $\Delta x^{1.4}$, and in fact this can be predicted analytically. Quite generally, any corner leads to a singularity where **E** diverges as r^{p-1} for a radius r from the corner, with p given by a transcendental equation in the corner angle and ε 's (here, $p \approx 0.702$) [4]. This leads to a perturbation in the frequency $\sim \int_{\Delta} |\mathbf{E}|^2 r dr \sim \Delta r^{2p} \approx \Delta r^{1.404}$, where Δr is the size of the perturbation (the pixel). Other smoothing schemes, in contrast, are limited by the linear error from the flat interfaces.

2.9 Conclusion

We have described in this chapter a method for designing subpixel smoothing algorithms for dielectric media having zero first-order effect. We then used this method to develop smoothing algorithms for isotropic and anisotropic media and verified their property of restoring the second-order accuracy of standard FDTD simulations. Because the new smoothing scheme greatly improves the accuracy of FDTD simulation for isotropic and anisotropic materials, without increasing the computational/storage cost (other than a one-time preprocessing step), it should be an attractive technique. A remaining challenge is to accurately handle objects with sharp corners, where the resulting field singularities are known to degrade the accuracy to between first- and second-order once the smoothing eliminates the first-order error. We are hopeful that an accurate smoothing can be developed for corners once the corresponding perturbation theory is derived.

Chapter 3

The failure of perfectly matched layers

3.1 Summary

Although perfectly matched layers (PMLs) have been widely used to truncate numerical simulations of electromagnetism and other wave equations, we point out important cases in which a PML fails to be reflectionless even in the limit of infinite resolution. In particular, the underlying coordinate-stretching idea behind PML breaks down in photonic crystals and in other structures where the material is not an analytic function in the direction perpendicular to the boundary, leading to substantial reflections. The alternative is an adiabatic absorber, in which reflections are made negligible by gradually increasing the material absorption at the boundaries, similar to a common strategy to combat discretization reflections in PMLs. We demonstrate the fundamental connection between such reflections and the smoothness of the absorption profile via coupled-mode theory, and show how to obtain higher-order and even exponential vanishing of the reflection with absorber thickness.

3.2 Overview

A perfectly matched layer (PML) is an artificial absorbing medium that is commonly used to truncate computational grids for simulating wave equations (e.g. Maxwell's equations), and is designed to have the property that interfaces between the PML and adjacent media are reflectionless in the exact wave equation [17, 215]. We describe important cases in which PML *fails* to be reflectionless, even in the exact (non-discretized) Maxwell equations, most notably in the case of periodic media (photonic crystals [99])—contrary to previous suggestions of photonic-crystal "PML" absorbers [120, 122, 123, 224, 234]. In these cases (similar to PML reflections due to discretization error [39, 215]), the remaining approach to reduce reflections is to "turn on" the absorption gradually, asymptotically approaching an "adiabatic" limit of zero reflections [103] regardless of whether the medium forms a true PML—here, we provide a deeper understanding of all such adiabatic absorbers by showing that the reflection's dependence on the thickness of the absorbing layer is determined by the smoothness of the absorption profile, and can be predicted by coupled-mode theory approximations. For a fixed absorption profile (typically quadratic or cubic in previous work [215]), the reflection decreases with absorber thickness L proportional to some characteristic power law determined by the smoothness (e.g. $1/L^6$ for quadratic absorption). (The same smoothness/reflection relation can be applied to adiabatic absorbers in boundary element methods where a true PML is much less practical to terminate infinite surfaces like waveguides even in cases where it is theoretically possible [251].) As the absorber becomes thicker, smoother absorptions become favorable, and we show that it is even possible to obtain exponential decrease of the reflection with L by new choices of the absorption profile. The role of PML (when it works), compared to ordinary absorbing materials, is to improve the constant factor in this reflection convergence, rather than the functional form. For homogeneous materials as in most previous analyses, although some attempts have been made to optimize the PML profile among various polynomial functions [36, 109, 149, 190], a quadratic or cubic profile works so well [215] that further attempts at optimization are arguably superfluous. On the other hand, for periodic media—especially when operating in modes with low group velocity—the required absorber thickness can become so large that the choice of absorption profile becomes critical. We also discuss the possibility of other optimizations, such as balancing the "transition" reflection from the absorber interface with the "round-trip" reflection due to the finite absorption, but these optimizations depend more sensitively on the incident-wave medium.

3.2.1 Various PML Formulations

There are several nearly equivalent formulations of PML. Berenger's original formulation [17] split the wave solution into the sum of two new artificial field components. A more common "UPML" (uniaxial-PML) formulation expresses the PML region as the ordinary wave equation with a combination of artificial anisotropic absorbing materials [197]. Both of these formulations were originally derived by laboriously computing the solution for a planewave incident on the absorber interface at an arbitrary angle and polarization, and then solving for the conditions in which the reflection is always zero. Both formulations, however, can also be derived by a complex "stretchedcoordinate" approach [40, 189, 218]—this much simpler and more elegant derivation of PML reveals its underlying meaning and generalizes more easily to inhomogeneous media, other wave equations, and other coordinate systems. In particular, the coordinate-stretching approach derives PML by an analytic continuation of Maxwell's equations into complex spatial coordinates, where the oscillating fields become exponentially decaying [40,49,189,218]. (This description can then be converted back into a change of materials via a complex coordinate transformation [218, 230]. A real coordinate transformation would be suitable only for waves that are already decaying, but not propagating, as it would merely act to shorten the distance over which the wave has decayed to some negligible amount.) By viewing PML as an analytic continuation, it can be shown to be reflectionless even for *inhomogeneous* media such as in Fig. 3-1(a) [88]: for a waveguide entering the PML perpendicularly, complex coordinate stretching is still possible because the material parameters (and hence Maxwell's equations) are analytic functions (constants) in that direction. The same derivation



Figure 3-1: (a) PML is still reflectionless for inhomogeneous media such as waveguides that are homogeneous in the direction perpendicular to the PML. (b, c) PML is no longer reflectionless when the dielectric function is discontinuous (non-analytic) in the direction perpendicular to the PML, as in a photonic crystal (b) or a waveguide entering the PML at an angle (c).

of PML, however, also immediately points to situations where PML is *inapplicable*: in any problem where the material parameters are *not* described by analytic functions in the direction perpendicular to the boundary, a reflectionless absorber cannot be designed by complex coordinate stretching. As discussed in more detail below, this means that "PML" is not reflectionless for photonic crystals as in Fig. 3-1(b) where the dielectric function varies discontinuously in the direction perpendicular to the boundary, or even in cases where a dielectric waveguide hits the PML obliquely [Fig. 3-1(c)]. (In fact, even for rare cases in which an oscillating dielectric function is analytic in the PML direction, we will explain that the analytic-continuation idea still does not yield a useful PML absorber in the discretized equations.)

3.2.2 PMLs in Photonic Crystals

Previous suggestions to apply PML to photonic crystals by simply overlapping a "PML" anisotropic absorber with the periodic dielectric function [120, 122, 123, 224, 234] (including a similar suggestion for integral-equation methods [183]) were therefore not "true" PML media in the sense that the reflection will not go to zero even in the limit of infinite resolution. In this thesis, we will refer to such an absorbing layer as a *pseudo-PML* (pPML). (In the special case of an effectively one-dimensional medium where there is only a single propagating mode, such as a single-mode waveguide surrounded by a complete-bandgap medium, it is possible to arrange an "impedancematched" absorber to approximately cancel that one mode [155], or alternatively to specify analytical boundary conditions of zero reflection for that one mode [162]. More generally, in a transfer-matrix or scattering-matrix method where one explicitly computes all propagating modes and expands the fields in that basis, it is possible to impose analytically reflectionless boundary conditions [135, 182], but such methods become very expensive in three dimensions.) These previous authors were nevertheless able to observe small reflections in a pPML only because they overlapped the pPML with many periods of the crystal and thereby turn on the pPML very gradually. As we explain below, such absorbing layers are more properly understood as adiabatic absorbers rather than PML media, and indeed the "PML" property only improves the constant factor in the long-wavelength limit of an effective homogeneous medium, or in any case where there are large homogeneous-material regions compared to the wavelength. Moreover, as we describe, the reflections worsen rapidly as the group velocity decreases (e.g. as a band edge is approached).

Even in the case of a homogeneous medium (or one uniform in the direction perpendicular to the boundary), where true PML applies, there are well-known numerical reflections due to the finite discretization [39,215]. It is sometimes claimed that the solutions for a PML converge exponentially to the solution of the open problem as the PML thickness is increased [130,259]. This is true, but only in the limit where the discretization error is negligible. Once the discretization reflections dominate, we show in Ch. 3.7 that the convergence rate with PML thickness depends on the smoothness of the PML profile in the same way as for any other adiabatic absorber, and the rate is only polynomial for a fixed polynomial profile. That is, there is a universal relationship between smoothness and reflectivity for all adiabatic absorbers, whether discrete or continuous and whether PML or non-PML. Other authors have remarked that the numerical reflection seems to be dominated by the discontinuity in the profile or its derivatives at the PML boundary [215], but have not presented a precise analysis of the relationship between convergence and smoothness.

3.3 PMLs versus Adiabatic absorbers

The notion of an adiabatic absorber is an old idea. Anechoic chambers have been used to provide minimally reflecting walls for echo suppression in acoustic and electromagnetic wave experiments [46]. The chamber walls are typically made of pyramidshaped pieces of lossy materials where the graded-geometry profile is a stand-in for the continuously-varying conductivity profile of numerical absorbing boundary layers. Given that the wavelength of sound waves in air is commensurate with radio waves in vacuum (i.e. centimeters), the same chamber design can be used for both types of waves (albeit constructed with different radiation absorbent materials). Anechoic chambers continue to be widely used for a number of applications including acoustic speaker testing, RF antenna design, and measurements of noise radiation from industrial machinery.

PMLs have also been used in a number of different areas outside of computational electromagnetism: modelling acoustic waves in uniform media [14, 113, 140], fluids [203], periodic media (phononic crystals) [128, 159, 181, 213] and even piezoelectric crystals [33] as well as in simulating the Schrodinger equation [3, 256]. However, similar to electromagnetism, claims of PML for periodic acoustic media (phononic crystals) [128, 159, 181, 213] appear to be erroneous for the same reason as for photonic crystals in this chapter and are really just another example of adiabatic absorbers.

The following chapter is structured as follows. We begin, in Ch. 3.4, with a very brief review of the derivation of PML in the simple case of one and two dimensions, and define the key quantities. Then, in Ch. 3.6, we explain and demonstrate the failure of PML for periodic media, even in the simplest case of one-dimensional structures where only normal-incident, non-evanescent waves are present, and even when the dielectric function varies analytically (sinusoidally). In fact, in this case, pPML may do no better than an ordinary absorbing medium (e.g., a scalar electric conductivity). Next, in Ch. 3.7, we analyze the relationship of the reflection to the smoothness of the absorption profile, and show via both 1d and 2d numerical calculations that the asymptotic behavior is predicted by coupled-mode theory, as well as the effect of group velocity. In Ch. 3.8, we describe how the coupled-mode understanding of this transition reflection points the way towards improved absorbing layers—ideally, layers whose reflection decreases exponentially with thickness (*not* the case even for true PML with a conventional quadratic profile, as mentioned above). Finally, we conclude with some remarks about future directions in Ch. 3.9.

3.4 Brief review of PML

3.4.1 Mathematical formulation

Consider Maxwell's equations in two dimensions (xy) for the TM polarization, in which the electric field (**E**) is in the z direction and the magnetic field (**H**) is in the xy plane, for a current source $\mathbf{J}_{\mathbf{z}}$ and a dielectric function $\varepsilon(x, y)$ in natural units $(\varepsilon_0 = \mu_0 = 1)$, with time-harmonic fields (time-dependence $\sim e^{-i\omega t}$) are:

$$\nabla \times \mathbf{H} = \frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} = -i\omega\varepsilon E_z \tag{3.1}$$

$$\frac{\partial E_z}{\partial y} = i\omega H_x \tag{3.2}$$

$$\frac{\partial E_z}{\partial x} = -i\omega H_y \tag{3.3}$$

One can now derive a PML absorbing boundary in the x direction, assuming for now that ε is a function of y only (e.g., the medium is homogeneous, or a waveguide in the x direction, near the computational cell boundary). In this case, one performs an analytic continuation to complex x coordinates by the transformation:

$$\frac{\partial}{\partial x} \to \frac{1}{1+i\frac{\sigma(x)}{\omega}}\frac{\partial}{\partial x},$$
(3.4)

in terms of a PML profile $\sigma(x)$, which plays the role of a conductivity or absorption strength. The profile $\sigma(x)$ can also be a complex function, where the imaginary part corresponds to a real coordinate stretching and enhances the attenuation of purely evanescent waves [68, 215], but in this thesis we focus on the case of real σ and the absorption of propagating waves. Maxwell's equations then become:

$$\frac{\partial H_y}{\partial x} - \left(1 + \frac{i\sigma}{\omega}\right)\frac{\partial H_x}{\partial y} = -i\omega\varepsilon E_z + \sigma\varepsilon E_z \tag{3.5}$$

$$\frac{\partial E_z}{\partial y} = i\omega H_x \tag{3.6}$$

$$\frac{\partial E_z}{\partial x} = -i\omega H_y + \sigma H_y \tag{3.7}$$

Note the $\sigma \varepsilon E_z$ and σH_y terms, which have the form of electric and magnetic conductivities, respectively. The remaining $i\sigma/\omega$ term becomes an integral or convolution in time-domain and is typically handled by integrating an auxiliary differential equation [215], but is trivial in frequency domain. The extension to PMLs in other directions is straightforward and is not reviewed here.

In a medium independent of x, the wave solutions can be decomposed into normal modes with x dependence $\exp(ik_x x)$ and $k_x > 0$ for right-going waves in a righthanded [58] medium (e.g. planewaves in a homogeneous medium or waveguide modes in a waveguide). The point of this transformation (3.4) is that these normal modes are thereby analytically continued to decaying solutions $\exp[ik_x x - \frac{k_x}{\omega}\int^x \sigma(x')dx']$ wherever $\sigma > 0$. The $1/\omega$ factor is desirable because, at least in a homogeneous dispersionless medium, the attenuation factor k_x/ω is independent of frequency (but not of incidence angle).

Outside the PML regions, where $\sigma = 0$, the wave equation and thus the solution are unchanged, and it is only inside the PML ($\sigma > 0$) that the oscillating solution becomes exponentially decaying with no reflections (in theory) no matter how fast σ changes, even if σ changes discontinuously. After a short distance L in the PML, the computational cell can then be truncated (e.g. with Dirichlet boundaries), with an exponentially small round-trip reflection

$$R_{\text{round-trip}} \sim e^{-4\frac{k_x}{\omega} \int_0^L \sigma(x') dx'},\tag{3.8}$$

where we have started the PML at x = 0, and the factor of 4 is because the reflection is proportional to the round-trip (2L) field squared. In the exact Maxwell equations, the PML could be made arbitrarily thin by making σ very large, but this is not feasible in practice because, once Maxwell's equations are discretized (in a finite-difference or finite-element scheme) the reflectionless property disappears. That is, it is not meaningful to analytically continue the discretized equations, and thus in the discretized system there are *numerical reflections* from the PML boundary that disappear in the limit of high resolution. To reduce these numerical reflections, most authors suggest that the PML be turned on gradually, i.e. that $\sigma(x)$ be a continuous function starting at zero, typically chosen to grow quadratically or cubically [215].

3.4.2 Absorption profile

More precisely, let us define $\sigma(x)$ in the PML $(x \in [0, L])$ by a shape function $s(u) \in [0, 1]$:

$$\sigma(x) = \sigma_0 \, s(x/L) \tag{3.9}$$

where the argument of s(u) is a rescaled coordinate $u = x/L \in [0, 1]$ and σ_0 is an overall amplitude set to achieve some theoretical maximum round-trip absorption R_0 for normal-incident waves in a medium of index n ($k_x = \omega n$). Using eq. (3.8) for R_0 , we define:

$$\sigma_0 = \frac{-\ln R_0}{4Ln \int_0^1 s(u')du'}.$$
(3.10)

For x < 0, outside the PML, $\sigma = 0$, i.e. s(u < 0) = 0. As L is made longer and longer for a fixed s(u), the PML profile σ turns on more and more gradually [both because s(u) is stretched out and because σ_0 decreases], and the numerical reflections decrease. Several authors have suggested $s(u) = u^2$ (quadratic) or $s(u) = u^3$ (cubic) turn-on of the PML, which have discontinuities at u = 0 in the second and third derivatives respectively [215]. In Ch. 3.7, we show that there is a simple correspondence between the smoothness of s(u) and the rate of decrease of absorption with L, as a consequence of the adiabatic theorem and coupled-mode theory. Note that the smoothness of s(u)is still relevant in a discretized system—with a fixed resolution and wavelength, as Lis increased one samples s(x/L) more and more finely and a discrete version of the adiabatic theorem applies [61].

In fact, we will see that the same adiabatic theorem and the same rate of decrease apply for *any* absorption, whether or not the absorbing material forms a PML. For example, if we only include σ on the right-hand-side of eq. (3.5), and neither on the left-hand-side nor in eqs. (3.6) and (3.7), it corresponds to an ordinary scalar electric conductivity. As we see in Ch. 3.6, the advantage of PML over this ordinary conductivity is not that the reflection decreases faster with L, but that this decrease is multiplied by a much smaller constant factor (which decreases with increasing resolution) in the case of PML. This advantage mostly disappears for periodic media where analytic continuation fails, but the same relationship between the rate of decrease and the smoothness of s(u) applies.

In general, therefore, we will divide the reflections from PML into two categories: the exponentially small round-trip reflections (above), and transition reflections from the boundary between $\sigma \neq 0$ and $\sigma = 0$ (which can arise either from numerical discretization or from other failures of PML as described in the next section). It is possible to obtain exactly zero reflection by balancing the round-trip and transition reflections so that they destructively interfere, but this cancellation can only occur for isolated wavelengths (and incident angles) [109] and hence is not useful in general. Instead, we will begin by setting the estimated round-trip reflection R_0 to be negligibly small (10⁻²⁵) and focus on the transition reflection; we return to the question of balancing round-trip and transition reflections in Ch. 3.8.

3.5 Adiabatic theorems in electromagnetism

Adiabatic theorems have been widely used in quantum mechanics involving timedependent Hamiltonians where at each temporal cross section the solution is expanded as an eigenmode series [8,43,112,157,169,247]. Such theorems have also been applied to electromagnetism for studying ordinary waveguides [148]. Johnson et al. [102] extended the adiabatic theorem to strongly-grated waveguides (photonic crystals) with arbitrary index modulation by generalizing coupled-mode theory to handle arbitrary nonuniform gratings using an instantaneous Bloch-mode basis. The analysis yields a continuous set of differential equations for the basis coefficients and the general principal is that as the system is varied more and more gradually, these coupling coefficients converge rapidly to constants (see Appendix A of Ref. 102). We will be returning to similar coupled-mode equations in later sections where we compute the reflection coefficient from the absorbing boundary layer.

3.6 Failure of PML

3.6.1 Homogeneous & inhomogeneous media

To illustrate the failure of PML in periodic media, we consider a finite-difference frequency-domain simulation (FDFD, with a second-order-accurate Yee grid) [44] of the simplest possible case: a periodic dielectric function $\varepsilon(x)$ in one dimension [so that we only have the E_z and H_y fields in eqs. (3.5) and (3.7)]. Given a point dipole source at some position (outside the absorber), we then compute the reflection coefficient from a pPML of thickness L as a function of both L and resolution.

Here, pPML (pseudo-PML) is defined by using eqs. (3.5) and (3.7): exactly the same equations as for an ordinary PML, but with an inhomogeneous ε function overlapping the "PML" as in Refs. 120, 122, 123, 224, 234. For comparison, we also show a *non-PML* absorber in which σ is included only in eq. (3.5) but not in eq. (3.7), i.e. an ordinary electric conductivity only. We consider two dielectric functions: vacuum ($\varepsilon = 1$) for comparison, and a periodic dielectric function $\varepsilon(x) = 6 + 5 \sin(2\pi x/a)$ that varies from 1 to 11 with period *a*. Like all one-dimensional periodic structures, this $\varepsilon(x)$ has photonic band gaps that prohibit propagation in certain frequency ranges [99], but we operate at a vacuum wavelength $\approx a$ slightly below the first bandgap (at a wavevector $k_x = 0.9\pi/a$ and vacuum wavelength $\lambda = 0.9597a$). The reflection is computed as the squared amplitude of the reflected Bloch wave, given by the total field minus the incident Bloch wave (computed by numerically solving for the Bloch waves of the discretized unit cell). Of course, there are two boundaries, at +x and -x, but we make the latter reflection negligible by using an absorber of thickness 5L on the left (and verified that further increasing the left-absorber thickness does not change the result). In this section, we use a quadratic shape function $s(u) = u^2$ for the absorber profile σ as defined above.

The absorber, here, is a pPML because it is not derived by analytic continuation of the dielectric function, and is instead formed by simply applying the homogeneous-PML equations on top of the inhomogeneous medium, leading to intrinsic reflections. However, in this case the periodic $\varepsilon(x)$ function is actually analytic in x, so in principle one *could* have derived a *true* PML by using eq. (3.5) with the analytically continued dielectric function $\varepsilon[x + \frac{i}{\omega} \int^x \sigma(x') dx']$. Unfortunately, this introduces new problems: the sine of a complex argument has an exponentially growing real part, causing the solutions to oscillate exponentially rapidly and leading to a breakdown of the discretization as the oscillation exceeds the Nyquist frequency. In practice therefore, we find that such a "true" PML with exponentially increasing $\Re[\varepsilon]$ leads to large reflections that (at best) decrease extremely slowly with resolution. So, one still cannot use a true PML in practice for the discretized problem (and the same is true any analytical periodicity, via Fourier expansion of ε). In any case, this possibility is not applicable in the vast majority of practical periodic structures, which more commonly involve a discontinuous (non-analytic) ε , so we do not consider analytically continuing $\varepsilon(x)$ further here and focus only on the pPML case.

Figure 3-2 shows the results of these one-dimensional FDFD simulations, and the difference between the uniform medium (where PML works) and the periodic medium (where it does not) is stark. In the uniform medium, the reflection from PML rapidly goes to zero as resolution is increased (and in fact, goes to zero quadratically with resolution because FDFD's center-difference discretization is second-order accurate), whereas the non-PML absorber in the uniform medium goes to a *constant* nonzero reflection (the Fresnel reflection coefficient from the exact Maxwell equations). For the periodic medium, both the pPML and non-PML absorbers behave roughly the same, going to a constant nonzero-reflection in the high-resolution limit: the pPML is *not* reflectionless for the exact Maxwell equations.



Figure 3-2: Reflection coefficient as a function of discretization resolution for both a uniform medium and a periodic medium with PML and non-PML absorbing boundaries (insets). For the periodic medium, PML fails to be reflectionless even in the limit of high resolution, and does no better than a non-PML absorber. Inset: reflection as a function of absorber thickness L for fixed resolution ~ 50pixels/ λ : as the absorber becomes thicker and the absorption is turned on more gradually, reflection goes to zero via the adiabatic theorem; PML for the uniform medium only improves the constant factor.

3.6.2 Backward-wave structures

Here we identify and explain a fundamental failure of PMLs for backward-wave structures. In particular, the stretched-coordinate derivation of PML suggests that such waves should be exponentially *growing* in the PML, and we explain this physically by pointing out that PML is an anisotropic "absorber" with *gain* in the longitudinal direction, which dominates for backward-wave modes.

PML is reflectionless because it corresponds merely to a complex coordinate stretching $z \rightarrow (1 + \frac{i\sigma}{\omega})z$, so that propagating waves $e^{i\beta z}$ are transformed into exponentially decaying waves $e^{i\beta z - \sigma z/v_p}$ for some PML strength σ . From this perspective, an obvious problem occurs for backward waves: if $v_p < 0$ for $v_g > 0$, then a +z-propagating wave ($v_g > 0$) will undergo exponential growth for $\sigma > 0$. (This is entirely distinct from the failure of PML in medium periodic in the z direction as described in the previous section, which in that case is due to the non-analyticity of Maxwell's equations and lead to reflections but not instability.)

We demonstrate this PML failure in the case of a backward-wave waveguide made of positive-index materials. For example, a hollow metallic waveguide containing a concentric dielectric cylinder was shown to support backward-wave modes [47]. More recently, the same phenomenon was demonstrated in all-dielectric (positiveindex, non-dispersive) photonic-crystal Bragg and holey fibers, and in general can be explained as an avoided eigenvalue crossing from a forced degeneracy at $\beta = 0$ [90]. An example of such a structure is shown in the inset of Fig. 3-3, which shows the crosssection of a Bragg fiber formed by alternating layers of refractive indices $n_{\rm hi} = 4.6$ (thickness 0.25*a*) and $n_{\rm lo} = 1.4$ (thickness 0.75*a*) with period *a*. The central highindex core has radius 0.45*a* and the first low-index ring has thickness 0.32*a*. For this geometry, one of the guided modes (with angular dependence $e^{im\phi}$ and m = 1) has the dispersion relation $\omega(\beta)$ shown in the inset of Fig. 3-3: at $\beta = 0$, $d^2\omega/d\beta^2 < 0$, resulting in a downward-sloping backward-wave region with $v_g v_p < 0$. As the index of the core cylinder is varied, this curvature can be changed from negative to positive in order to eliminate the backward-wave region. The exponential growth of fields within the PML region is observed for this structure as shown in Fig. 3-3. We simulated the backward-wave structure of Fig. 3-3(inset) with a finite-difference time-domain (FDTD) simulation in cylindrical coordinates [176,215], terminated in the z direction with PML layers. Both the forwardand backward-wave modes were excited with a short-pulse current source, and the fields in the PML region after a long time were fit to an exponential in order to determine the decay rate. Figure 3-3 plots this decay rate as a function of the curvature $\partial^2 \omega / \partial \beta^2|_{\beta=0}$ as the core-cylinder index is varied from 2.6 to 5.0. The appearance of negative curvature, which indicates the appearance of a backward-wave region, precisely coincides with the decay rate changing sign to exponential growth.

In the usual case in which group and phase velocities are oriented in the same direction, the overall rate constant is negative and this causes absorptive loss in the PML. In the case of backward waves, however, the ratio v_g/v_p is negative, and thus the overall rate constant is positive, i.e., PML produces gain.

In a homogeneous backward-wave medium, this problem can be solved merely by making $\sigma < 0$ in the negative-index frequency ranges [53, 59]. This solution is impossible in the case of Fig. 3-3(inset), however, because at the *same* ω one has *both* forward and backward waves—no matter what sign is chosen for σ , one of these waves will experience exponential growth in the PML.

3.6.3 PMLs & adiabatic absorbers

One way of understanding why pPML is not reflectionless for a periodic medium was described in the previous section: the equations with "PML" absorption are no longer derived via analytic continuation of Maxwell's equations, and so the fundamental justification for PML disappears. This has nothing to do with either evanescent waves or glancing-angle waves, neither of which are present in one dimension, nor is it a numerical reflection from discretization (since it does not vanish as resolution is increased). Another way of understanding this is that the propagating waves in a periodic medium are Bloch waves [99], and consist of a superposition of reflections from all interfaces (all places where ε changes) in the medium—when we absorb



Figure 3-3: Field decay rate within the PML vs. curvature of the dispersion relation at $\beta = 0$, showing onset of gain for $v_g < 0$ and loss for $v_g > 0$. Inset. Dispersion relation (of the first TE band) with $v_g < 0$ region at $\beta = 0$, and cross-section of the Bragg fiber.

waves reflected from interfaces within the "PML," we have effectively terminated the periodicity and hence see reflections from this termination. (Similar but even stronger reflections are observed if one terminates the periodicity *before* it enters the absorbing region [155].)

However, the inset of Fig. 3-2 shows a way in which the reflections can still be made small for the periodic medium: by increasing the thickness L of the absorbing layer. As L is increased, we see that the reflections in *all* four cases (PML and non-PML, uniform and periodic) go to zero as $1/L^6$ asymptotically (although the periodic media take longer to attain this asymptotic power law). The true PML in the uniform medium is only different in that it has a better constant factor (which depends on resolution). The reason for this, as described in the next section, is that all transition reflections can be understood via the same coupled-mode mechanism, and



Figure 3-4: Field convergence (~ reflection/L²) vs. absorber length for various σ ranging from linear $[\sigma(z/L) \sim (z/L)]$ to quintic $[\sigma(z/L) \sim (z/L)^5]$. For reference, the corresponding asymptotic power laws are shown as dashed lines. *Inset:* Bragg-fiber structure in cylindrical computational cell with absorbing regions used in simulation.

the $1/L^6$ rate is a consequence of the second-derivative discontinuity in $s(u) = u^2$. This reduction of reflection with L is *adiabatic absorption*, distinct from the PML concept, and it is such adiabatic absorption that one must better understand in order to efficiently truncate periodic media.

With the understanding that the standard formulation of PML fails for backward waves, we now turn to a discussion of what can be done instead. As pointed out above, previous corrections for left-handed media [53,59] are inapplicable here because one has forward and backward waves at the same ω . Since the reflectionless property of PML fundamentally arises from the coordinate-stretching viewpoint, and gain is predicted by coordinate-stretching above, we are led to the conclusion that PML must be abandoned entirely for such backward-wave structures. The alternative is to use a

scalar absorbing material, e.g. a scalar conductivity σ , which is absorbing for all field orientations and therefore cannot lead to gain (unlike in the previous example for periodic media where an anisotropic "pseudo-PML" could still be employed). At the interface of such a material, however, there will be reflections. Such reflections can be made arbitrarily small, however, by turning on the absorption by a sufficiently gradual taper transition similar to our approach for an unrelated failure of PML as described in the next section. Even for PML, numerical reflections due to discretization require a similar gradual σ taper. In both cases, the reflection R(L) goes to zero as the absorber thickness L is made longer (and more gradual), and the impact of PML (when it works) is merely to multiply R(L) by a smaller constant coefficient. Even without PML, the rate at which R(L) goes to zero can be made more rapid by reducing the discontinuity in σ : for example, if $\sigma \sim (z/L)^2$ (for z > 0) then its second derivative is discontinuous at the transition z = 0 and R(L) consequently scales as $1/L^4$, while if $\sigma \sim (z/L)^3$ then $R(L) \sim 1/L^6$. Figure 3-4 shows how a scalar conductivity σ can be used as a last-resort replacement for PML in the backward-wave structure of Fig. 3-3(inset). The plot shows the difference-squared of the magnetic field at a test point for absorber lengths L and L + 1 (which scales as $R(L)/L^2$) versus L for various conductivity profiles σ . Even with both forward and backward waves excited, the reflection can indeed be made small for a sufficiently thick absorber (albeit thicker than a PML for purely forward-wave modes) and displays the expected scaling $1/L^{2d+2}$ for $\sigma \sim (z/L)^d$ discussed next.

3.7 Smoothness & Reflection

In this section, we demonstrate and explain the relationship between the smoothness of the absorber profile's shape function s(u) and the dependence of reflection on absorber thickness L. The basic principle is that, as L increases, the rate of change of the absorption (PML or otherwise) becomes more and more gradual—as it approaches a perfectly uniform (or perfectly periodic) limit, there is an *adiabatic theorem* stating that the reflections must go to zero. Such an adiabatic theorem is the well-understood


Figure 3-5: Reflectivity vs. PML thickness L for 1d vacuum (inset) at a resolution of 50pixels/ λ for various shape functions s(u) ranging from linear [s(u) = u] to quintic $[s(u) = u^5]$. For reference, the corresponding asymptotic power laws are shown as dashed lines.

mechanism behind gradual waveguide tapers [186], and adiabatic theorems also hold in periodic media with slowly varying unit cells [103], and there is also an adiabatic theorem for slowly-varying discretized systems [61]. Moreover, as we discuss in the next section, the rate at which the adiabatic (zero-reflection) limit is approached is determined by the smoothness of the transition s(u).

3.7.1 Numerical results

First, however, let us present the results of numerical experiments using second-order FDFD discretization for four structures: uniform and periodic media in one and



Figure 3-6: Reflectivity vs. pPML thickness L for the 1d periodic medium (inset) with period a, as in Fig. 3-2, at a resolution of 50pixels/a with a wavevector $k_x = 0.9\pi/a$ (vacuum wavelength $\lambda = 0.9597a$, just below the first gap) for various shape functions s(u) ranging from linear [s(u) = u] to quintic $[s(u) = u^5]$. For reference, the corresponding asymptotic power laws are shown as dashed lines.

two dimensions (with continuous and discontinuous ε , respectively). The reflection versus PML/pPML absorber length L in one dimension is shown for uniform media in Fig. 3-5 and for a periodic medium (the same medium as for Fig. 3-2) in Fig. 3-6, for a variety of shape functions $s(u) = u^d$ for exponents $d \in \{1, 2, 3, 4, 5\}$. In both cases, there is a striking pattern: the reflection asymptotically follows a power law $1/L^{2d+2}$, which we will explain analytically below in terms of the smoothness of s(u).

In two dimensions, we looked at the boundary reflection from a point source at the center of the cell. In this case, defining a single "reflection" coefficient is more difficult because the point source emits waves at multiple angles. Instead, we look at the convergence of the electric field as L is increased, and defined a *field convergence factor*

$$\frac{|E_z^{(L+1)}(x,y) - E_z^{(L)}(x,y)|^2}{|E_z^{(L)}(x,y)|^2}$$
(3.11)

in terms of the electric field E_z at a point (x, y) (chosen roughly halfway between the point source and the absorbing layer) for two PML/pPML thicknesses L and L + 1. This difference should go to zero as $L \to \infty$, assuming that the reflection goes to zero in this limit (and hence the field converges to the solution for open boundaries). Indeed, this adiabatic limit is observed for both the uniform medium (vacuum) in Fig. 3-7 and for a periodic medium (a square lattice of width-0.7*a* square air holes in $\varepsilon = 12$) in Fig. 3-8. Again, there is a simple power-law relationship evident in both plots: when $s(u) = u^d$, the field convergence factor goes as $1/L^{2d+4}$.

In 1d, we found that the reflection went as $1/L^{2d+2}$ for $s(u) = u^d$, and in 2d we found that the corresponding field convergence factor went as $1/L^{2d+4}$. These two results are mathematically equivalent, for the following reason. Suppose that the reflection coefficient (for waves at any angle) goes asymptotically as $1/L^{2\alpha}$ for some exponent α ; it follows that the reflected electric field goes as $1/L^{\alpha}$, and hence $E^{(L)}(x, y) = E^{(\infty)}(x, y) + O(1/L^{\alpha})$. Substituting this expression into eq. (3.11) and expanding in powers of 1/L, one finds that the field convergence factor goes as $1/L^{2\alpha+2}$, exactly the difference of $1/L^2$ that we observed above. [There is a subtlety in this derivation: it implicitly assumes that the *phase* of the $O(1/L^{\alpha})$ term, i.e. the reflected



Figure 3-7: Field convergence factor [eq. (3.11)] (~ reflection/ L^2) vs. PML thickness L for 2d vacuum (inset) at a resolution of 20pixels/ λ for various shape functions s(u) ranging from linear [s(u) = u] to quintic [$s(u) = u^5$]. For reference, the corresponding asymptotic power laws are shown as dashed lines. Inset: $\Re[E_z]$ field pattern for the (point) source at the origin (blue/white/red = positive/zero/negative).



Figure 3-8: Field convergence factor [eq. (3.11)] (~ reflection/ L^2) vs. pPML thickness L for the discontinuous 2d periodic medium (left inset: square lattice of square air holes in $\varepsilon = 12$) with period a, at a resolution of 10pixels/a with a vacuum wavelength $\lambda = 0.6667a$ (not in a band gap) for various shape functions s(u) ranging from linear [s(u) = u] to quintic $[s(u) = u^5]$. For reference, the corresponding asymptotic power laws are shown as dashed lines. Right inset: $\Re[E_z]$ field pattern for the (point) source at the origin (blue/white/red = positive/zero/negative).

phase, goes to a constant as $L \to \infty$ in order to expand in powers of 1/L. This assumption is confirmed by our numerical results, but it is also predicted analytically by the coupled-mode theory result eq. (3.13) in the next section.]

3.7.2 Analysis

The natural way to analyze waves propagating along a medium that is slowly varying in the propagation direction (say x) is coupled-mode theory (or coupled-wave theory) [103, 148]: at each x, one expands the fields in the basis of the eigenmodes (indexed by ℓ) of a uniform structure with that cross-section in terms of expansion coefficients $c_{\ell}(x)$. (The eigenmodes have x-dependence $e^{i\beta_{\ell}x}$ for some propagation constants β_{ℓ} .) The expansion coefficients c_{ℓ} in this basis are then determined by a set of ordinary differential equations for dc_{ℓ}/dx coupling the different modes, where the coupling coefficient is proportional to the rate of change [here, the derivative s'(x/L)]. In the limit where the structure varies more and more slowly, the solution approaches an "adiabatic" limit in which the c_{ℓ} are nearly constant (i.e. no scattering between modes). Although coupled-mode theory was originally developed for media that are slowly varying in the propagation direction [148], it has been generalized to periodic media with a slowly varying unit cell |103|, in which very similar coupled-mode equations appear. A similar adiabatic limit has also been derived for slowly varying discrete systems. Using coupled-mode theory, one can derive a universal relationship between the smoothness of the rate of change [s'(u)] and the asymptotic rate of convergence to the adiabatic limit. This relationship, derived below, analytically predicts the convergence rates of the reflection with absorber length observed in the previous section.

Coupled Mode Theory

We omit the derivation of the coupled-mode equations here; their general form is considered in detail elsewhere [103, 148]. We simply quote the result: in the limit of slow variation (large L), the equations can be solved to lowest order in 1/L in terms of a simple integral. In particular, if the structure is smoothly parameterized by a shape function s(x/L) (e.g. the absorption profile as given here), then the amplitude c_r (corresponding to a reflected power $|c_r|^2$) of a reflected mode is given to lowest-order (for large L) by [103]:

$$c_r(L) = \int_0^1 s'(u) \frac{M[s(u)]}{\Delta\beta[s(u)]} e^{iL \int_0^u \Delta\beta[s(u')]du'} du.$$
(3.12)

Here, M is a coupling coefficient depending on the mode overlap between the incident and reflected field (in the changing part of the structure) and $\Delta\beta \neq 0$ is the difference $\beta_i - \beta_r$ between the propagation constants of the incident and reflected modes. Both of these are some analytic functions of the shape s(u). In general, there may be more than one reflected mode, and in a periodic structure the coefficient even for a single reflected mode is a sum of contributions of above form from the different Brillouin zones [103], but it suffices to analyze the rate of convergence of a single such integral with L. The basic reason for the adiabatic limit is that, as L grows, the phase term oscillates faster and faster and the integral of this oscillating quantity goes to zero.

Convergence Analysis

There are many standard methods to analyze the asymptotic (large L) properties of such an integral. In particular, we apply a technique that is commonly used to analyze the convergence rate of Fourier series: one simply integrates by parts repeatedly until a nonzero boundary term is obtained [26, 152]. Each integration by parts integrates the $e^{iL\int\Delta\beta}$ term, dividing the integrand by $iL\Delta\beta(u)$, and differentiates the $s'M/\Delta\beta$ term. (If $\Delta\beta$ is real as in the case of waveguides but not absorbers then we can turn this expression into a Fourier transform, otherwise we have to evaluate the expression explicitly as shown next.) After integrating by parts d times, the boundary term at u = 0 is zero if the corresponding derivative $s^{(d)}(0^+)$ is zero, whereas the boundary term at u = 1 is always negligible because of the absorption (leading to a complex $\Delta\beta$ and exponential decay), assuming a small round-trip reflection R_0 . The dominant asymptotic term is the *first* (lowest-d) u = 0 boundary term that is nonzero, since all subsequent integrations by parts have an additional factor of 1/L. [Here, we have assumed that s is a smooth function in (0, 1) so that there are never deltafunction contributions from the interior. In systems with purely decaying solutions (e.g. elliptic equations), mapping the domain $[0,\infty]$ to [0,1] requires some care since if the wave oscillations vary too rapidly (exceeding the discrete grid's Nyquist limit), large reflections will arise [26].] The result is the following asymptotic form for $c_r(L)$, independent of the particular details of the geometry or the modes:

$$c_r(L) = s^{(d)}(0^+) \frac{M(0^+)}{\Delta\beta(0^+)} [-iL\Delta\beta(0^+)]^{-d} + O(L^{-(d+1)}), \qquad (3.13)$$

where $s^{(d)}(0^+)$ is the first nonzero derivative of s(u) at $u = 0^+$, and integrating by parts d times yielded a division by $(-iL\Delta\beta)^d$ (flipping sign each time). This result corresponds to what is sometimes called "Darboux's principle:" the convergence is dominated by the lowest-order singularity [26], which here is the first discontinuity in the rate of change s'(u) at u = 0. A similar result applies, for example, to the convergence rate of a Fourier series: a function that has a discontinuity in the dth derivative has a Fourier series whose coefficients c_n decrease asymptotically as $1/n^{(d+1)}$ [26,152] (the d+1 instead of d is due to the fact that our integral starts with s').

Equation (3.13) would seem to imply that the reflection $\sim |c_r|^2$ is $O(L^{-2d})$, but this is not the case because there is a hidden 1/L factor in the coupling coefficient M, thanks to the 1/L dependence of σ_0 in eq. (3.10). The coupling coefficient Mis a matrix element proportional to the rate of change of the materials [103], which in this case is $\frac{\partial \sigma}{\partial u} = s'(u)\sigma_0 \sim 1/L$. Therefore, the reflection scales as $|M|^2/L^{2d} = O(L^{-(2d+2)})$, exactly corresponding to our numerical results above.

Other useful results can be obtained from eq. (3.13), and in particular one can show that the reflections due to nonuniformity worsen in a periodic structure as a flat band edge (β_0, ω_0) is approached [186]. As a quadratic-shaped band-edge $\omega - \omega_0 \sim (\beta - \beta_0)^2$ is approached, the group velocity $v_g = \frac{d\omega}{d\beta}$ scales proportional to $\beta - \beta_0$, while the $\Delta\beta$ between the forward and reflected modes is $2(\beta - \beta_0) \sim v_g$. Also, the coupling coefficient M is proportional to $1/v_g$ because of the constant-power normalizations of the incident and reflected modes [103, 186]. Hence, by inspection of eq. (3.13), the reflection $|c_r|^2 = O(v_g^{-(2d+4)})$. For example, the reflection is $O(v_g^{-6})$ for a linear taper s(u) = u [186]. Because of this unfavorable scaling, an imperfect absorbing layer such as a pPML is most challenging in periodic structures when operating close to a band edge where there are slow-light modes (in the same way that other taper transitions are challenging in this regime [186]).

Derivation of reflection coefficient

Here we show how eq. (3.13) is derived from eq. (3.12) using the methods of Fourier analysis, specifically the integration by parts method. We start with

$$c_r(L) = \int_0^1 s'(u) \frac{M[s(u)]}{\Delta\beta[s(u)]} e^{iL \int_0^u \Delta\beta[s(u')]du'} du$$
(3.14)

where the upper bound of the integral can be extended to $+\infty$ since $\Delta\beta[s(u')] = 2k + 2i\sigma(u)$ is itself complex owing to the absorption profile $\sigma = \sigma_0 s(u)$ within the "PML" region. Using this fact with slight rearrangement of the terms, we arrive at:

$$c_r(L) \approx \int_0^\infty \frac{s'(u) \frac{M[s(u)]}{\Delta\beta[s(u)]}}{iL\Delta\beta[s(u)]} \frac{d}{du} \left[e^{iL\int_0^u \Delta\beta[s(u')]du} \right] du.$$
(3.15)

Equation (3.15) is now in the standard form of $\int u dv = uv - \int v du$ and after one iteration of integration by parts becomes:

$$c_r(L) = \frac{s'(u)\frac{M[s(u)]}{\Delta\beta[s(u)]}}{iL\Delta\beta[s(u')]} e^{iL\int_0^u \Delta\beta[s(u')]du} \Big|_{0^+}^\infty - \int_0^\infty \left[\frac{s'(u)\frac{M[s(u)]}{\Delta\beta[s(u)]}}{iL\Delta\beta[s(u')]}\right]' e^{iL\int_0^u \Delta\beta[s(u')]du}.$$
 (3.16)

The first term with integral bounds is zero as $s'(0^+) = 0$ on the left boundary and exponentially approaches zero on the right boundary. This iteration repeats until a discontinuity is reached when $s^{(d)}(0^+) \neq 0$ resulting in:

$$c_r(L) = \frac{s^{(d)}(0^+)\frac{M(0^+)}{\Delta\beta(0^+)}}{\left[-iL\Delta\beta(0^+)\right]^d} + O\left(\frac{1}{L^{(d+1)}}\right).$$
(3.17)

3.7.3 Adiabatic theorems in discrete systems

There is one thing missing from the above analysis, and that is the discretized-space adiabatic case. In a slowly varying discrete system [i.e, sampling some slow change $s_n = s(n\Delta x/L)$ as L grows larger], there is still a proof of the adiabatic theorem $(c_r \to 0)$, but the only published proof is currently for the lossless case (unitary evolution) [61]. Also, an analogous integral form of the lowest-order reflection has not been presented, nor has the rate of convergence to the adiabatic limit been analyzed in the discrete case. So, our prediction of the asymptotic convergence rate is rigorously proven only for the case of the continuous-space wave propagation. However, our numerical results demonstrate that a slowly-changing discretized system exhibits *exactly* the same scaling (e.g. in the PML case for uniform media, where the only reflections are due to discretization). (This seems analogous to the fact that the discretization error of a discrete Fourier transform converges at the same rate as the decay of the coefficients of the continuous-space Fourier series [26].) In future work, we hope to further validate our numerical result for the convergence rate in discretized space with a proper generalization of the coupled-mode analysis.

3.8 Towards Better Absorbers

From the previous section, there is a close relationship between the smoothness of the absorption profile and the asymptotic convergence rate of the reflections R(L)as a function of absorber thickness L: if the profile s(u) has a discontinuity in the d-th derivative (e.g. for $s = u^d$), then the reflection coefficient goes as $1/L^{2d+2}$ for a fixed round-trip reflection. This result raises several interesting questions. Can one do better than polynomial convergence? What is the optimal shape s(u)? And what if the round-trip reflection is not fixed?



Figure 3-9: Reflectivity vs. PML thickness L for 1d vacuum (blue circles) at a resolution of 50pixels/ λ , and for pPML thickness L in the 1d periodic medium of Fig. 3-6 (red squares) with period a at a resolution of 50pixels/a with a wavevector $k_x = 0.9\pi/a$ (vacuum wavelength $\lambda = 0.9597a$. In both cases, a C_{∞} (infinitely differentiable) shape function $s(u) = e^{1-1/u}$ for u > 0 is used, leading to asymptotic convergence as $e^{-\alpha\sqrt{L}}$ for some constants α .



Figure 3-10: Reflectivity vs. PML thickness L for two different absorber profiles in the 1d uniform medium of Fig. 3-5 at a resolution of 50pixels/ λ . In both cases, a C_{∞} (infinitely differentiable) shape function is used, leading to asymptotic exponential convergence but the blue curve has a slightly lower constant factor since it approximates a quadratic taper profile for small taper lengths.



Figure 3-11: Reflectivity vs. pPML thickness L for two different absorber profiles in the 1d periodic medium of Fig. 3-6 with a wavevector $k_x = 0.9\pi/a$ (vacuum wavelength $\lambda = 0.9597a$ at a resolution of 50pixels/ λ . In both cases, a C_{∞} (infinitely differentiable) shape function is used, leading to asymptotic exponential convergence but the blue curve has a slightly lower constant factor since it approximates a quadratic taper profile for small taper lengths.



Figure 3-12: Reflectivity vs. pPML thickness L for two different absorber profiles in the 2d periodic medium of Fig. 3-8 consisting of a square lattice of square air holes in $\varepsilon = 12$ with period a, at a resolution of 10pixels/a with a vacuum wavelength $\lambda = 0.6667a$ (not in a band gap). Right inset: $\Re[E_z]$ field pattern for the (point) source at the origin (blue/white/red = positive/zero/negative). In both cases, a C_{∞} (infinitely differentiable) shape function is used, leading to asymptotic exponential convergence but the blue curve has a slightly lower constant factor since it approximates a quadratic taper profile for small taper lengths.

3.8.1 Smoothness & C_{∞} functions

The above result relating smoothness and convergence has a natural corollary: if s(u)is C_{∞} , i.e. all of its derivatives are continuous, then the reflection goes to zero faster than any polynomial in 1/L. This is similar to a well-known result for the convergence of Fourier series of C_{∞} functions [26]; the exact rate of faster-than-polynomial convergence again depends on the strongest singularity in s(u). For example, for $s(u) = (\tanh(u) + 1)/2$, which goes exponentially to zero as $u \to -\infty$ and to one as $u \to +\infty$, the reflection should decrease exponentially with L, as determined by contour integration from the residue of the pole at $u = \pm i\pi/2$ that is closest to the real axis (similar to the analysis for the convergence of a Fourier series for an analytic function [26,63]). However, such an absorption taper would require an infinitely thick absorber in order to avoid discontinuously truncating the exponential tail of tanh(u). To have a C_{∞} function with a finite absorber, with s(u) = 0 for $u \leq 0$, the s(u) function must be non-analytic; a standard example of such a function is $s(u) = e^{1-1/u}$ for u > 0 (all of whose derivatives go to zero as $u \to 0^+$, where there is an essential singularity). Because $s(u) = e^{1-1/u}$ is C_{∞} , its reflection R(L) must decrease faster than any polynomial. Exactly how much faster than polynomial is determined by asymptotically evaluating the integral of eq. (3.12) by a saddle-point method [27, 37]: the result is that R(L) decays asymptotically as $e^{-\alpha\sqrt{L}}$ for some constant $\alpha > 0$ [27]. This is confirmed by Fig. 3-9, which plots the PML/pPML reflection for the 1d uniform and periodic cases on a semilog scale versus \sqrt{L} , and results clearly approach a straight line as expected.

Although $s(u) = e^{1-1/u}$ yields an exponential convergence of the absorption in Fig. 3-9, the constant factor and the exponential rate are almost certainly suboptimal for this arbitrary choice of C_{∞} function. If we compare Fig. 3-9 to Fig. 3-5 for the uniform case and Fig. 3-6 for the periodic case, we see that this $C_{\infty} s(u)$ is superior to the polynomial s(u) for the periodic case where PML is not perfect, but inferior for the uniform case until the reflection becomes inconsequential (~ 10^{-20}). This is still a useful result in the sense that one mainly needs to improve pPML for the periodic case, whereas PML is already good enough for uniform media. However, one would ideally prefer a shape function that is consistently better than the polynomial s(u), regardless of the dielectric function, so further exploration of the space of possible absorption profiles seems warranted. A possibility here would be to design a custom absorber profile with optimal performance that combines the superior constant factor of quadratic absorbers at short taper lengths with the exponential convergence of C_{∞} absorber at long taper lengths. The function $s(u) = u(e^{\beta u} - 1)/(e^{\beta} - 1)$ where $\beta = 2L^{0.5}$ is such a function and its performance is demonstrated for 1d uniform (see Fig. 3-10), 1d periodic (see Fig. 3-11) and 2d periodic media (see Fig. 3-12) comparing to the simple C_{∞} profile $s(u) = u^{1-1/u}$. In all three examples, the constant factor of the more complicated custom absorber profile is lower for small taper lengths (since it approximates a simple quadratic profile in this regime) while both absorber profiles clearly demonstrate the exponential convergence as the taper length is increased. This is just a simple demonstration of the utility of custom absorber profiles for adiabatic tapers and we hope further research will continue to improve their properties.

3.8.2 Balancing round-trip & transition reflections

Finally, in the above analysis we fixed the round-trip reflection R_0 , via the estimate of eq. (3.8), to approximately 10^{-25} in order for our calculations to isolate the effect of the transition reflection. Obviously, in a real application, one is unlikely to require such low reflections and one will set R_0 to a larger value, corresponding to a larger $\sigma_0 \sim \ln R_0$ in eq. (3.10). This will also reduce the transition reflection [as seen from eq. (3.13)], but only by a logarithmic constant factor. The best choice to minimize reflection for a given absorber length, in principle, is to set R_0 to be roughly equal to the transition reflection for that length. (Another reason to make them equal is the possibility of destructive interference between the round-trip and transition reflection [109], but such destructive interference is inherently restricted to narrow bandwidths and ranges of incident angles and so we do not concern ourselves with this possibility.) In order to make them roughly equal, one needs an estimate of the transition reflection; for example, one could simply numerically fit the power law of



Figure 3-13: Reflectivity vs. PML thickness L for 1d vacuum (inset) at a resolution of 50pixels/ λ for $s(u) = u^2$, with the round-trip reflection either set to $R_0 = 10^{-16}$ (upper blue line) or set to match the estimated transition reflection from Fig. 3-5 (lower red line). By matching the round-trip reflection R_0 to the estimated transition reflection, one can obtain a substantial reduction in the constant factor of the total reflection, although the asymptotic power law is only changed by a ln L factor.

eq. (3.13). The result of such matching is shown in Fig. 3-13 for a quadratic profile $s(u) = u^2$ in 1d uniform media, and the overall reflection is reduced by a factor of 3-400 compared to a fixed $R_0 = 10^{-16}$. This is a significant reduction, but is not overwhelming (especially for smaller L) and changes the asymptotic convergence rate eq. (3.13) only by a factor of $\ln R_0 \sim \ln L$. The drawback of this optimization is that it is difficult to determine the transition reflection analytically for inhomogeneous media, and so one is generally forced to make a conservative estimate of R_0 , which reduces the advantage gained.

3.9 Conclusion

Perfectly matched layers are an extraordinarily powerful technique to absorb waves incident on the boundaries of wave-equation simulation, but they are not a panacea. In particular, for cases such as photonic crystals where the medium is not analytic in the direction perpendicular to the boundary, the fundamental coordinate-stretching idea behind PML breaks down, and the interface has intrinsic reflections (even for simple 1d cases with only normal-incident non-evanescent waves). However, one can still obtain small reflections by gradually ramping up the "pseudo-PML" (pPML) absorption, similar to the idea behind the quadratic PML profiles commonly used to circumvent discretization-based reflections in uniform media, forming an adiabatic absorber. In fact, for both cases (pPML in periodic media and PML in discretized uniform media), we show that the basic mechanism behind the reflection is determined in the same way by the smoothness of the absorption profile, which can be predicted analytically by coupled-mode theory. More generally, an adiabatic absorber is applicable in any situation where a true PML is inconvenient or impossible to implement.

The same theory then predicts that an *exponential absorber*, one whose reflections decrease exponentially with some power of the absorber thickness L, is possible, for example by using an infinitely differentiable absorption profile. (In contrast, ordinary PML in a uniform medium with a quadratic profile is *not* an exponential absorber: its

numerical-discretization reflections decrease as $1/L^6$.) We gave a simple C_{∞} example profile that led to such exponential absorption, but much future work remains to be done in identifying profiles with both exponential absorption and good constant factors. In particular, one possibility that we will examine in a subsequent manuscript is an absorption profile whose smoothness increases with L, so that it matches simple quadratic profiles for small L but becomes exponentially smoother with large L. (Such L-varying profiles require a more careful convergence analysis, however, in order to ensure that they approach the adiabatic zero-reflection limit. A closely related mathematical idea is explored in Ref. 27.)

Chapter 4

A simple validation scheme for perfectly matched layers

4.1 Summary

We introduce a numerical method to verify the correctness of perfectly matched layer (PML) absorbing boundaries. Our method is straightforward and can be applied to any PML regardless of implementation. To demonstrate its utility, we derive a correct uniaxial-PML formulation for lossless, non-dispersive, anisotropic media for the finite-difference time-domain (FDTD) method. Our formulation consists of at most four auxiliary variables and is as computationally efficient in storage as previous implementations based on the split-field approach. We numerically verify the validity of our method and also demonstrate that certain previously reported formulations are incorrect.

4.2 Overview

We introduce an efficient numerical method to verify the correctness of the perfectly matched layer (PML) irrespective of the details of its implementation. Our method is both simple and intuitive, consisting of two parts: the first requires that reflections from the PML boundaries reduce to zero in the limit of increasing resolution and the second shows the reflection decreasing with absorber thickness L proportional to some characteristic power law determined by the smoothness (e.g. $1/L^6$ for quadratic absorption) when the round-trip reflection is fixed [177]. This method is then used to demonstrate the failure of at least one previously published proposal for PML in anisotropic media. We then derive our own efficient formulation for an anisotropic absorber for anisotropic media and verify its correctness with this approach.

A perfectly matched layer (PML) is an artificial absorbing medium that is commonly used to truncate computational grids for simulating wave equations (e.g. Maxwell's equations), and is designed to have the property that interfaces between the PML and adjacent media are reflectionless in the exact wave equation [17, 215]. There are several nearly equivalent formulations of PML. Berenger's original formulation [17] split the wave solution into the sum of two new artificial field components, which while effective as a numerical method for absorbing incident waves nevertheless does not reveal physical insights into its operation or connection with Maxwell's equations. Moreover implementation of split-field PMLs typically requires laboriously computing the solution for a planewave incident on the absorber interface at an arbitrary angle and polarization, and then solving for the conditions in which the reflection is always zero [74, 179].

A more common "UPML" (uniaxial-PML) formulation expresses the PML region as the ordinary wave equation with a combination of artificial anisotropic absorbing materials [75, 197, 215, 221]. This modification of the permittivity and permeability material parameters, as Ward and Pendry clarified [230], is equivalent to a co-ordinate transformation of Maxwell's equations. In the case of the UPML, this "stretched-coordinate" corresponds to an analytic continuation of real space into the complex plane where oscillating fields become exponentially decaying [40, 189, 220]. UPML's principal advantage over split-field PML is that it is based on the unmodified Maxwell's equations (the absorption appearing as anisotropic material tensors or fictitious conductivity absorbers) which makes implementation in numerical code straightforward. Several attempts have been made at UPML formulations for anisotropic materials [141,219,220,222,253] yet some, we will argue, are incorrect as we demonstrate analytically and confirm with a numerical experiment while others lack the full details of an FDTD implementation. Certain authors have proposed "materialindependent PMLs" containing conductivities σ_D and σ_B (instead of σ_E and σ_H) in the split-field [252, 254] and UPML [228, 253] formulations while failing to recognize in the derivation equations that the Jacobian stretch matrices and full anisotropic material (ε or mu) tensors do not commute.

To illustrate the failure of incorrect PML formulations in anisotropic media, we first consider a frequency-domain simulation with a planewave expansion method (PWFD) solved with an iterative bi-conjugate gradient algorithm [13] in two dimensions (2d). Given a point dipole source at the center of the cell, we then wish to compute the reflection coefficient from a surrounding absorber of thickness L as a function of both L and resolution. Defining a single "reflection" coefficient is difficult because the point source emits waves at multiple angles, so instead we look at the convergence of the electric field by defining a *field convergence factor*

$$\frac{|E_z^{(L+1)}(x,y) - E_z^{(L)}(x,y)|^2}{|E_z^{(L)}(x,y)|^2}$$
(4.1)

in terms of the electric field E_z at a point (x, y) (chosen directly adjacent to the absorbing layer) for two PML thicknesses L + 1 and L (L = 1 in this example). This difference should go to zero as both resolution increases and $L \to \infty$ for a true PML, assuming that the reflection goes to zero in this limit (and hence the field converges to the solution for open boundaries).

Figure 4-1 shows the results of these two-dimensional PWFD simulations for uniform isotropic ($\varepsilon = 10$) and anisotropic (ε chosen to be a random positive-definite symmetric matrix with eigenvalues [8.4896, 8.7820, 11.5210])) media; the difference between a correct PML and an incorrect Z-PML (from Ref. 253) or non-PML (conductivity absorber for the electric field only) is stark. In the correct PML formulation for isotropic and anisotropic media, the reflection from the boundaries rapidly goes to zero as resolution is increased, whereas the Z- and non- PML absorbers asymptote to a *constant* nonzero reflection.



Figure 4-1: Results from 2d PWFD showing field convergence factor (a proxy for the reflection coefficient) versus resolution in pixels/ λ_v for both isotropic and anisotropic media with PML, Z-PML (from Ref. 253) and non-PML (conductivity absorber in electric field) absorbing boundaries. For the anisotropic medium, Z-PML fails to be reflectionless in the limit of high resolution. Inset: E_z field profile of a point source at the center of the 2d computational cell surrounded by absorbing material (blue/white/red = positive/zero/negative).

The inability to detect the failure of the Z-PML is a result of a common testing procedure used by many researchers in the past who proposed various PML formulations. The dominant verification method typically consisted of computing low reflections from the PML absorbing regions as a function of the wave incident angle or fixed position in the non-absorbing region of the computational cell. Unfortunately, this method is not a sufficient condition to verify a true PML. The key numerical test to determine whether any absorber is a true PML, as just demonstrated, is to show that reflections from the absorbing region decrease to zero with increasing resolution. In other words, in the limit as the discretized wave equation approaches a continuum, the reflections must reduce to zero (corresponding to that of a true PML in a non-discretized system).



Figure 4-2: Results from 2d PWFD simulation showing field convergence factor (~ reflection/ L^2) vs. absorber thickness in units of vacuum wavelength (L/λ_v) for anisotropic media at a resolution of 9pixels/ λ_v for various polynomial absorber functions s(u) ranging from linear [s(u) = u] in blue to cubic $[s(u) = u^3]$ in green. As the absorber becomes thicker and the absorption is turned on more gradually, reflection goes to zero via the adiabatic theorem. For reference, the corresponding asymptotic power laws are shown as dashed lines. Fixing the round-trip reflection yields similar scaling relationships and values between the three types of absorbers. Inset: $\Re[E_z]$ field pattern for the (point) source at the center (blue/white/red = positive/zero/negative).

Additionally, any absorber whether a PML or conductivity absorber for one field having fixed round-trip reflection, shows a characteristic scaling of the transition reflections from the absorbing region with absorber length set by the smoothness of the polynomial absorption profile (i.e., $1/L^{2d+4}$ for a $s(u) = u^d$ profile using the field convergence factor of eq. (4.1) [177]). The reason for this is that all transition reflections can be understood via the same coupled-mode mechanism, and the $1/L^8$ rate is a consequence of the second-derivative discontinuity in $s(u) = u^2$ [177]. Any proposed PML formulation that satisfies both conditions is then guaranteed to be a true PML.

4.3 PML formulation for lossless, non-dispersive, anisotropic media

We now present our implementation of UPML absorbing boundary layers for FDTD simulations. It is assumed that the reader is already familiar with the basic stretched-coordinate derivation of PML, as well as the general technique by which a coordinate transformation can be expressed as a transformation of ε and μ (see Appendix of Ref. 124), leading to the so-called "UPML" formulation. The formulation presented below has been implemented in our popular free software package for FDTD [176].

The subtlety is that the transformations/materials of PML are frequency-dependent, and so to express them in time domain involves the evolution of appropriate auxiliary differential equations. (Equivalently, multiplication by a frequency-dependent susceptibility corresponds in the time domain to a convolution, leading to so-called "convolutional PML" formulations [192].) The emphasis is on keeping the number of auxiliary differential equations (and the resulting memory and computational costs) to a minimum, while not making the PML region too complicated compared to the non-PML regions.

Because the treatment of ε and μ are identical except for interchange of **D** with **B** and **E** with **H** (and a sign flip from Ampere's to Faraday's law), we only describe the ε and $d\mathbf{D}/dt$ equations (Ampere's law) here.

We proceed slightly differently from the UPML as derived in e.g. the Taflove and Hagness FDTD textbook [215]. As reviewed in the appendix, this "standard" UPML performs a matrix factorization that relies on ε commuting with the PML Jacobian, and this is not the case for an arbitrary anisotropic ε . Our factorization, instead, works for anisotropic ε , and turns out to have a nice property: the PML just adds two auxiliary fields **U** and **W** with *diagonal* relationships to **D** and **E**, and allow us to use the non-PML timestepping operations *unchanged* except for the addition of these diagonal ODE updates. In particular, we get a true PML for anisotropic, dispersive, media "for free" in the sense that the code for those portions is unchanged (although the case of conducting media necessitates an extra auxiliary field [139]).



Figure 4-3: Results from 2d FDTD simulations showing field convergence factor eq. (4.1) vs. absorber thickness in units of vacuum wavelength (L/λ_v) for anisotropic media at a resolution of 20pixels/ λ_v for various polynomial absorber functions s(u) ranging from linear [s(u) = u] in blue to $[s(u) = u^3]$ in green. Left inset: field convergence factor versus resolution in pixels/ λ_v showing correct PML scaling relationship. Top right inset: $\Re[E_z]$ field pattern snapshot in time for the (point) source at the center (blue/white/red = positive/zero/negative).

This is nice because the timestepping for anisotropic media requires special care for stability, and dispersive media may have complicated polarization-update equations, and this way we don't need to modify any of that in the slightest for the PML. Our approach also makes a clean separation between the **D** update equation (from $\nabla \times \mathbf{H}$) and terms that affect the **E** update equation (from the constitutive equations), and correspondingly in the Meep code these are handled separately in step_db and update_eh, respectively.

4.4 Non-PML materials in Meep

In Meep [176], we support dispersive media $\varepsilon(\omega)$ of the form:

$$\varepsilon(\omega) = \left(1 + \frac{i\sigma_D}{\omega}\right) \left[\varepsilon_{\infty} + \chi_E(\omega)\right],$$

where σ_D is a conductivity (which may be anisotropic, but must be diagonal in Cartesian coordinates and hence commutes with \mathcal{J}), ε_{∞} is a non-dispersive part of the permittivity (which may be non-diagonal anisotropic) and $\chi_E(\omega)$ is some additional dispersive part (possibly anisotropic) implemented by auxiliary ODEs that solve $\mathbf{P} = \chi(\omega) \mathbf{E}$ (currently, a sum of Lorentzians). This corresponds, in time-domain, to Ampere's law of the form:

$$\mathbf{K} = \nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \sigma_D \mathbf{D}, \qquad (4.2)$$

denoting $\nabla \times \mathbf{H}$ by \mathbf{K} for later convenience, and a constitutive equation

$$\mathbf{D} = \varepsilon_{\infty} \mathbf{E} + \mathbf{P},$$

where **P** is time-evolved via a system of ODEs derived from $\mathbf{P} = \chi(\omega)\mathbf{E}$. (We don't include free currents **J** here, since they have no impact on the PML equations, and indeed one rarely puts free currents inside the PML anyway.)

In FDTD, we discretize these in space and time. Let us denote the time discretization by a superscript: \mathbf{D}^n denotes $\mathbf{D}(n\Delta t)$, and similarly for \mathbf{P}^n and \mathbf{E}^n . The magnetic fields are offset in time by half a time step, giving $\mathbf{H}^{n+0.5}$ and $\mathbf{K}^{n+0.5}$. To time-step these fields in FDTD, we first compute \mathbf{D}^{n+1} from \mathbf{D}^n and $\mathbf{K}^{n+0.5}$ by the curl equation, then compute \mathbf{E}^{n+1} from \mathbf{D}^{n+1} and \mathbf{P}^{n+1} by the constitutive equation, and finally compute \mathbf{P}^{n+2} . (Note that, with Lorentzian dispersion, we can compute \mathbf{P}^{n+2} from \mathbf{P}^{n+1} and \mathbf{P}^n given \mathbf{E}^{n+1} .

Using the standard second-order center-difference approximations, the equation

for **D** becomes:

$$\mathbf{K}^{n+0.5} = \frac{\mathbf{D}^{n+1} - \mathbf{D}^n}{\Delta t} + \sigma_D \frac{\mathbf{D}^{n+1} + \mathbf{D}^n}{2},$$

giving

$$\mathbf{D}^{n+1} = \left(1 + \frac{\sigma_D \Delta t}{2}\right)^{-1} \left[\left(1 - \frac{\sigma_D \Delta t}{2}\right) \mathbf{D}^n + \mathbf{K}^{n+0.5} \right].$$
(4.3)

The constitutive equation is simply:

$$\mathbf{E}^{n+1} = \varepsilon_{\infty}^{-1} \left(\mathbf{D}^{n+1} - \mathbf{P}^{n+1} \right).$$
(4.4)

(This is nontrivial to implement correctly for anisotropic media [235], and the details have been covered in Ch. 2.5.)

4.5 PML formulation in frequency domain

PML is simplest to derive in frequency domain, where the fields all have timedependence $e^{-i\omega t}$. An ordinary PML in Cartesian coordinates is derived by a complex coordinate stretching, where each coordinate is stretched by a factor

$$s_{x,y,z} = 1 + \frac{i\sigma_{x,y,z}}{\omega},$$

where ω is the frequency and σ is the PML "conductivity." For example, to terminate the cell in the *x* direction, only σ_x is nonzero. These coordinate stretchings can be absorbed into Maxwell's equations as a change in ε and μ . The original permittivity ε in the PML region is replaced by an effective tensor $\tilde{\varepsilon}$ given by

$$\tilde{\varepsilon} = \frac{\mathcal{J}\varepsilon\mathcal{J}^{T}}{\det\mathcal{J}},$$

where $\mathcal{J} = \text{diag}(s_x^{-1}, s_y^{-1}, s_z^{-1})$ is the Jacobian matrix of the coordinate stretching.

In frequency domain, replacing $\varepsilon(\omega)$ with $\tilde{\varepsilon}$, eq. (4.2) becomes:

$$\mathbf{K} = \nabla \times \mathbf{H} = -i\omega s_x s_y s_z \left(1 + \frac{i\sigma_D}{\omega} \right) \begin{pmatrix} s_x^{-1} & & \\ & s_y^{-1} \\ & & s_z^{-1} \end{pmatrix} [\varepsilon_\infty + \chi_E(\omega)] \begin{pmatrix} s_x^{-1} & & \\ & s_y^{-1} \\ & & s_z^{-1} \\ & & (4.5) \end{pmatrix} \mathbf{E} \cdot \mathbf{K} = \mathbf{E} \cdot \mathbf{E$$

Although this is straightforward to implement in frequency domain, where one $\varepsilon(\omega)$ is as good as another, in the time domain a frequency-dependent term requires care. In time-domain, a frequency-dependent term can be thought of as a convolution with a filter (hence the viewpoint of a "convolutional PML" adopted by some authors), where in the language of signal-processing we wish to find a stable recursive filter to implement this convolution with as few taps as possible (to minimize the memory and computational burden). Equivalently, the frequency dependence may be implemented by auxiliary ordinary differential equations (discretized ODE = recursive filter). The most convenient ODEs to discretize are first-order ODEs. For example, $a = sb = (1 + i\sigma/\omega)b$ gives $-i\omega a = -i\omega b + \sigma b$, which corresponds to the first-order ODE $\frac{da}{dt} = \frac{db}{dt} + \sigma b$.

So, before we proceed to time domain, we want to *factorize* eq. (4.17) into terms with only *one* factor of s (or σ/ω) each (or ratios of single s factors). In doing so, we are free to change the definition of **D** and introduce new auxiliary fields as desired, since the fields in the PML region are not physical. A key trick is the factorization:

$$s_{x}s_{y}s_{z}\begin{pmatrix} s_{x}^{-1} & & \\ & s_{y}^{-1} & \\ & & s_{z}^{-1} \\ & & & s_{z}^{-1} \end{pmatrix} = \begin{pmatrix} s_{y} & & \\ & s_{z} & \\ & & s_{x} \end{pmatrix} \begin{pmatrix} s_{z} & & \\ & s_{x} & \\ & & s_{y} \end{pmatrix}.$$

Using this factorization, and defining new auxiliary fields, \mathbf{C} , \mathbf{U} , and \mathbf{W} , we can factorize eq. (4.17) into the following equivalent form (giving the same relationship between \mathbf{H} and \mathbf{E}):

$$\mathbf{K} = \nabla \times \mathbf{H} = -i\omega \left(1 + \frac{i\sigma_D}{\omega} \right) \mathbf{C}$$
(4.6)

$$\mathbf{U} = \begin{pmatrix} s_y^{-1} & & \\ & s_z^{-1} & \\ & & s_x^{-1} \end{pmatrix} \mathbf{C}$$
(4.7)

$$\mathbf{D} = \begin{pmatrix} s_z^{-1} & & \\ & s_x^{-1} & \\ & & s_y^{-1} \end{pmatrix} \mathbf{U}$$
(4.8)

$$\mathbf{W} = \varepsilon_{\infty}^{-1} \left(\mathbf{D} - \mathbf{P} \right) \tag{4.9}$$

$$\mathbf{P} = \chi_E(\omega) \mathbf{W} \tag{4.10}$$

$$\mathbf{E} = \begin{pmatrix} s_x & & \\ & s_y & \\ & & s_z \end{pmatrix} \mathbf{W}$$
(4.11)

4.6 PML formulation in time domain

Given eqs. (4.6–4.11), a time-domain formulation is now easy because each equation only includes first-order factors in ω (corresponding to first-order time derivatives). Moreover, all of the nontrivial equations, namely (4.6), (4.9), and (4.10), are *exactly* the same as in the non-PML case, meaning that we can use *exactly* the same code except passing new fields **C** and **W** instead of **D** and **E**. The other equations are diagonal ODEs that are trivial to implement. It may seem wasteful to have three new auxiliary fields, but in many cases they can be omitted: except in corners of the computational cell, one has PML only in one direction that only one component of $s_{x,y,z}$ is $\neq 1$, and in these cases one need only store one component of **U** and one component of **W** (with the other components replaced by **D** and **E**, respectively); **C** can be omitted in non-conducting materials ($\sigma_D = 0$).

For completeness, we write out the equations here, in the order that they would be evaluated. The **C** update from eq. (4.6) is identical to the **D** timestep from eq. (4.12)

$$\mathbf{C}^{n+1} = \left(1 + \frac{\sigma_D \Delta t}{2}\right)^{-1} \left[\left(1 - \frac{\sigma_D \Delta t}{2}\right) \mathbf{C}^n + \mathbf{K}^{n+0.5} \right].$$
(4.12)

Each component U_k of **U** is updated from the corresponding component of **C** by the ODE $dU_k/dt + \sigma_{k+1}U_k = dC_k/dt$ (where $\sigma_{k\pm 1}$ is interpreted as a cyclic shift, e.g. $\sigma_{y+1} = \sigma_z, \sigma_{z+1} = \sigma_x$), giving:

$$U_{k}^{n+1}\left(1+\frac{\sigma_{k+1}\Delta t}{2}\right)^{-1}\left[\left(1-\frac{\sigma_{k+1}\Delta t}{2}\right)U_{k}^{n}+C_{k}^{n}-C_{k}^{n-1}\right].$$
 (4.13)

Each component of **D** is then updated from the corresponding component of **U** by the ODE $dD_k/dt + \sigma_{k-1}D_k = dU_k/dt$, giving:

$$D_{k}^{n+1}\left(1+\frac{\sigma_{k-1}\Delta t}{2}\right)^{-1}\left[\left(1-\frac{\sigma_{k-1}\Delta t}{2}\right)D_{k}^{n}+U_{k}^{n}-U_{k}^{n-1}\right].$$
 (4.14)

W is then updated from $\mathbf{D} - \mathbf{P}$ exactly as in eq. (4.15):

$$\mathbf{W}^{n+1} = \varepsilon_{\infty}^{-1} \left(\mathbf{D}^{n+1} - \mathbf{P}^{n+1} \right).$$
(4.15)

Then, each component of **E** is updated from each component of **W** by the ODE $dE_k/dt = dW_k/dt + \sigma_k W_k$, giving:

$$E_k^{n+1} = E_k^n + \left(1 + \frac{\sigma_k \Delta t}{2}\right) W_k^{n+1} - \left(1 - \frac{\sigma_k \Delta t}{2}\right) W_k^n.$$

$$(4.16)$$

Finally, \mathbf{P}^{n+2} is computed using $\mathbf{P} = \chi_E(\omega) \mathbf{W}$, exactly as for the non-PML case (but with \mathbf{W} replacing \mathbf{E}).

Note that, in order to avoid having to save the fields from the previous timestep in yet more auxiliary arrays, the **C**, **U**, and **D** updates [eqs. (4.12–4.14)] have to be performed in a single loop body, while the **W** and **E** updates [eqs. (4.15–4.16)] must be performed in another single loop body. [One cannot easily merge the two loops because the offdiagonal anisotropic terms in $\varepsilon_{\infty}^{-1}$ combined with the staggered Yee grid mean that eq. (4.15) is effectively nonlocal in space, requiring **D**^{*n*+1} components at several spatial points to determine **W**^{*n*+1}. Separating the two loops has the additional advantage that it reduces the combinatorial explosion of the number of material cases that must be handled.]

4.7 Review of isotropic, nondispersive PML

In this appendix, we review the standard "textbook" UPML for isotropic nondispersive media as found in, for example, Taflove and Hagness [215]. It is important to distinguish mere differences in notation (the derivation in T&H is far less compact than the one here) from substantive differences, and in particular the usual UPML corresponds to a different factorization of eq. (4.17) than ours. This textbook formulation only works for isotropic (or at least diagonal) ε , because it assumes that an arbitrary diagonal matrix commutes with ε .

In particular, dropping the σ_D and χ_E terms (assuming nondispersive media), the textbook PML formulation corresponds to the following refactorization of eq. (4.17):

$$\mathbf{K} = \nabla \times \mathbf{H} = -i\omega \begin{pmatrix} s_y & & \\ & s_z & \\ & & s_x \end{pmatrix} \varepsilon_{\infty} \begin{pmatrix} \frac{s_z}{s_x} & & \\ & \frac{s_x}{s_y} & \\ & & \frac{s_y}{s_z} \end{pmatrix} \mathbf{E}.$$
 (4.17)

This then factorizes into two equations:

$$\begin{split} \mathbf{K} &= \nabla \times \mathbf{H} = -i\omega \begin{pmatrix} s_y & & \\ & s_z & \\ & & s_x \end{pmatrix} \mathbf{D}, \\ \mathbf{E} &= \varepsilon_{\infty}^{-1} \begin{pmatrix} \frac{s_x}{s_z} & & \\ & \frac{s_y}{s_x} & \\ & & \frac{s_z}{s_y} \end{pmatrix} \mathbf{D}. \end{split}$$

The former discretizes similarly to eq. (4.3), with the PML $\sigma_{y,z,x}$ taking the place of σ_D , and the latter turns into the ODE $dE_k/dt + \sigma_{k-1}E_k = \varepsilon_{\infty}^{-1}[dD_k/dt + \sigma_k D_k]$ which is easily discretized.

Only two fields need be stored: **D** and **E**. If a conductivity σ_D (or similar) is included, however, one needs an additional auxiliary field **C** (or similar). It turns out that the case of a dispersive $\mathbf{P} = \chi_E(\omega)\mathbf{E}$ can also be handled with no additional storage compared to the non-PML case, at least for a Lorentzian χ_E that requires both \mathbf{P}^{n+1} and \mathbf{P}^n to be stored anyway. So, the main cost of the Meep formulation is that two additional auxiliary fields **U** and **W** must be stored in the PML; however, this is actually only a slight additional burden in storage since for most PML regions only one component of **U** and **W** actually must be stored (see above). The main problem with the textbook implementation is that it fails for anisotropic media (even if one naively plugs a tensor $\varepsilon_{\infty}^{-1}$ into the **E** update, the reflection will not go to zero with increasing resolution). This is especially important given the fact that, even for nominally isotropic media, accurate subpixel averaging requires the discretization to use effective anisotropic media at material boundaries.

4.8 Concluding remarks

We have demonstrated a straightforward method to verify the correctness of any PML formulation. Our approach consists of two parts. The first test involves computing the quantity known as the field convergence factor, related to the reflection from an absorbing boundary, as a function of the resolution of the discrete computational grid. A correct PML, equivalent to a complex co-ordinate stretching of real space, has reflections that decrease to zero with increasing resolution. Any violation of this property indicates a non-PML absorber. The second test consists of analyzing the transition reflections that arise for waves entering the absorbing regions when the absorber's round-trip reflection is fixed. All absorbers with the round-trip reflection properly set will show identical scaling relationships of the reflection with absorber length based on the smoothness of the absorber profile. We demonstrated the utility of this method in showing that at least one previously proposed PML formulation for anisotropic media was incorrect. We then derived our own PML formulation and verified its correctness using this scheme. The deeper analytical insights into the PML that these findings provide should now enable researchers to correctly develop PML for a number of different media and co-ordinate systems with ease.

Chapter 5

Meep: A flexible free-software package for electromagnetic simulations by the FDTD method

5.1 Summary

This chapter describes Meep, a popular free implementation of the finite-difference time-domain (FDTD) method for simulating electromagnetism. In particular, we focus on aspects of implementing a full-featured FDTD package that go beyond standard textbook descriptions of the algorithm, or ways in which Meep differs from typical FDTD implementations. These include pervasive interpolation and accurate modelling of subpixel features, advanced signal processing, support for nonlinear materials via Padé approximants, and flexible scripting capabilities.

Program Summary

Program title: Meep Program summary URL: http://ab-initio.mit.edu/meep Licensing provisions: GNU GPL No. of lines in distributed program, including test data, etc: 58000 No. of bytes in distributed program, including test data, etc: 734K

Distribution format: tar.gz

Programming language: C++

Computer: any computer with a Unix-like system and a C++ compiler; optionally exploits additional free software packages: GNU Guile [1], libctl interface library [2], HDF5 [3], MPI message-passing interface [4], and Harminv filter-diagonalization [5]. Developed on 2.8 GHz Intel Core 2 Duo.

Operating system: any Unix-like system; developed under Debian GNU/Linux 5.0.2

RAM: problem dependent (roughly 100 bytes per pixel/voxel)

Classification: Electrostatics and Electromagnetics

External routines/libraries: optionally exploits additional free software packages: GNU Guile [1], libctl interface library [2], HDF5 [3], MPI message-passing interface [4], and Harminv filter-diagonalization [5] (which requires LAPACK and BLAS linear-algebra software [6]).

Nature of problem: classical electrodynamics

Solution method: finite-difference time-domain (FDTD) method

Running time: problem dependent (typically about 10 ns per pixel per timestep)

References:

- 1. GNU Guile, http://www.gnu.org/software/guile
- 2. Libctl, http://ab-initio.mit.edu/libctl
- M. Folk, R.E. McGrath, N. Yeager, HDF: An update and future directions, in: Proc. 1999 Geoscience and Remote Sensing Symposium (IGARSS), Hamburg, Germany, vol. 1, IEEE Press, 273–275, 1999.
- T.M. Forum, MPI: A Message Passing Interface, in: Supercomputing '93, Portland, OR, 878–883, 1993.
- 5. Harminv, http://ab-initio.mit.edu/harminv
- 6. LAPACK, http://www.netlib.org/lapack/lug
5.2 Overview

One of the most common computational tools in classical electromagnetism is the finite-difference time-domain (FDTD) algorithm, which divides space and time into a regular grid and simulates the time evolution of Maxwell's equations [64, 127, 212, 215, 248]. This chapter describes our free, open-source implementation of the FDTD algorithm: Meep (an acronym for MIT Electromagnetic Equation Propagation), available online at http://ab-initio.mit.edu/meep. Meep is full-featured, including, for example: arbitrary anisotropic, nonlinear, and dispersive electric and magnetic media; a variety of boundary conditions including symmetries and perfectly matched layers (PML); distributed-memory parallelism; Cartesian (1d/2d/3d) and cylindrical coordinates; and flexible output and field computations. It also includes some unusual features, such as advanced signal processing to analyze resonant modes, accurate subpixel averaging, a frequency-domain solver that exploits the time-domain code, complete scriptability, and integrated optimization facilities. Here, rather than review the well-known FDTD algorithm itself (which is thoroughly covered elsewhere), we focus on the particular design decisions that went into the development of Meep whose motivation may not be apparent from textbook FDTD descriptions, the tension between abstraction and performance in FDTD implementations, and the unique or unusual features of our software.

Why implement yet another FDTD program? Literally dozens of commercial FDTD software packages are available for purchase, but the needs of research often demand the flexibility provided by access to the source code (and relaxed licensing constraints to speed porting to new clusters and supercomputers). Our interactions with other photonics researchers suggest that many groups end up developing their own FDTD code to serve their needs (our own groups have used at least three distinct in-house FDTD implementations over the past 15 years), a duplication of effort that seems wasteful. Most of these are not released to the public, and the handful of other free-software FDTD programs that could be downloaded when Meep was first released in 2006 were not nearly full-featured enough for our purposes. Since then,

Meep has been cited in over 100 journal publications and has been downloaded over 10,000 times, reaffirming the demand for such a package.

FDTD algorithms are, of course, only one of many numerical tools that have been developed in computational electromagnetism, and may perhaps seem primitive in light of other sophisticated techniques, such as finite-element methods (FEMs) with high-order accuracy and/or adaptive unstructured meshes [98, 198, 245], or even radically different approaches such as boundary-element methods (BEMs) that discretize only interfaces between homogeneous materials rather than volumes [25,41,188,225]. Each tool, of course, has its strengths and weaknesses, and we do not believe that any single one is a panacea. The nonuniform unstructured grids of FEMs, for example, have compelling advantages for metallic structures where micrometer wavelengths may be paired with nanometer skin depths. On the other hand, this flexibility comes at a price of substantial software complexity, which may not be worthwhile for dielectric devices at infrared wavelengths (such as in integrated optics or fibers) where the refractive index (and hence the typical resolution required) varies by less than a factor of four between materials, while small features such as surface roughness can be accurately handled by perturbative techniques [105]. BEMs, based on integral-equation formulations of electromagnetism, are especially powerful for scattering problems involving small objects in a large volume, since the volume need not be discretized and no artificial "absorbing boundaries" are needed. On the other hand, BEMs have a number of limitations: they may still require artificial absorbers for interfaces extending to infinity (such as input/output waveguides) [250]; any change to the Green's function (such as introduction of anisotropic materials, imposition of periodic or symmetry boundary conditions, or a switch from three to two dimensions) requires reimplementation of large portions of the software (e.g. singular panel integrations and fast solvers) rather than purely local changes as in FDTD or FEM; continuously varying (as opposed to piecewise-constant) materials are inefficient; and solution in the time domain (rather than frequency domain, which is inadequate for nonlinear or active systems in which frequency is not conserved) with BEM requires an expensive solver that is nonlocal in time as well as in space [25]. And then, of course, there are specialized tools that solve only a particular type of electromagnetic problem, such as our own MPB software that only computes eigenmodes (e.g. waveguide modes) [107], which are powerful and robust within their domain but are not a substitute for a general-purpose Maxwell simulation. FDTD has the advantages of simplicity, generality, and robustness: it is straightforward to implement the full time-dependent Maxwell equations for nearly arbitrary materials (including nonlinear, anisotropic, dispersive, and time-varying materials) and a wide variety of boundary conditions, one can quickly experiment with new physics coupled to Maxwell's equations (such as populations of excited atoms for lasing [19,35,89,167,257]), and the algorithm is easily parallelized to run on clusters or supercomputers. This simplicity is especially attractive to researchers whose primary concern is investigating new interactions of physical processes, and for whom programmer time and the training of new students is far more expensive than computer time.

The starting point for any FDTD solver is the time-derivative parts of Maxwell's equations, which in their simplest form can be written:

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} - \mathbf{J}_B \tag{5.1}$$

$$\frac{\partial \mathbf{D}}{\partial t} = +\nabla \times \mathbf{H} - \mathbf{J},\tag{5.2}$$

where (respectively) **E** and **H** are the macroscopic electric and magnetic fields, **D** and **B** are the electric displacement and magnetic induction fields [94], **J** is the electriccharge current density, and \mathbf{J}_B is a fictitious magnetic-charge current density (sometimes convenient in calculations, e.g. for magnetic-dipole sources). In time-domain calculations, one typically solves the initial-value problem where the fields and currents are zero for t < 0, and then nonzero values evolve in response to some currents $\mathbf{J}(\mathbf{x}, t)$ and/or $\mathbf{J}_B(\mathbf{x}, t)$. (In contrast, a *frequency-domain* solver assumes a time dependence of $e^{-i\omega t}$ for all currents and fields, and solves the resulting linear equations for the steady-state response or eigenmodes [99, app. D].) We prefer to use dimensionless units $\varepsilon_0 = \mu_0 = c = 1$. From our perspective, this choice emphasizes both the scale invariance of Maxwell's equations [99, chap. 2] and also the fact that the most meaningful quantities to calculate are almost always dimensionless ratios (such as scattered power over incident power, or wavelength over some characteristic lengthscale). The user can pick any desired unit of distance a (either an SI unit such as $a = 1 \ \mu \text{m}$ or some typical lengthscale of a given problem), and all distances are given in units of a, all times in units of a/c, and all frequencies in units of c/a. In a linear dispersionless medium, the constituent relations are $\mathbf{D} = \varepsilon \mathbf{E}$ and $\mathbf{B} = \mu \mathbf{H}$, where ε and μ are the relative permittivity and permeability (possibly tensors); the case of nonlinear and/or dispersive media (including conductivities) is discussed further in Ch. 5.6.

The remaining chapter is organized as follows. Chapter 5.3 gives a brief history of Meep's development. In Ch. 5.4, we discuss the discretization and coordinate system; in addition to the standard Yee discretization [215], this raises the question of how exactly the grid is described and divided into "chunks" for parallelization, PML, and other purposes. Chapter 5.5 describes a central principle of Meep's design, *pervasive interpolation* providing (as much as possible) the illusion of continuity in the specification of sources, materials, outputs, and so on. This led to the development of several techniques unique to Meep, such as a scheme for subpixel material averaging designed to eliminate the first-order error usually associated with averaging techniques or stairstepping of interfaces. In Ch. 5.6, we describe and motivate our techniques for implementing nonlinear and dispersive materials, including a slightly unusual method to implement nonlinear materials using a Padé approximant that eliminates the need to solve cubic equations for every pixel. Chapter 5.7 describes how typical computations are performed in Meep, such as memory-efficient transmission spectra or sophisticated analysis of resonant modes via harmonic inversion. This section also describes how we have adapted the time-domain code, almost without modification, to solve frequency-domain problems with much faster convergence to the steadystate response than merely time-stepping. The user interface of Meep is discussed in Ch. 5.8, explaining the considerations that led us to a scripting interface (rather than a GUI or CAD interface). Chapter 5.9 describes some of the tradeoffs between performance and generality in this type of code and the specific compromises chosen in Meep. Finally, we make some concluding remarks in Ch. 5.10.

5.3 Development History

Prior to Meep, the Ab Initio physics group of Professor John Joannopoulos had been using a custom-built FDTD software program written in Fortran by Shanhui Fan (with contributions from other group members including Chiyan Luo) known as "TD3D". This code, as its name implies supported only 3d simulations, was used extensively in the early years of the group's photonic crystal research throughout the 1990s and early 2000s. However its inflexible user interface and unsuitability for simulating other co-ordinate systems, particularly cylindrical, was becoming increasingly evident as the research frontier evolved towards more exotic material systems as well as larger and more complicated geometries. In early 2003, David Roundy, then a postdoctoral researcher in the Joannopoulos group, initiated a new FDTD software program written from scratch in C++ known as "Dactyl" as it was originally designed to perform simulations in only 2d and cylindrical co-ordinates. David was soon joined by graduate students Mihai Ibanescu and Peter Bermel working together to expand its features to include polarization fields to model dispersive dielectrics, third-order Kerr nonlinearities, split-field perfectly matched layers, various symmetry and boundary conditions with a C++ object-oriented interface. At the time, David, Mihai and Peter were primarily interested in simulating fiber structures that required cylindrical co-ordinates and so this was where the initial emphasis was placed. The notion of pervasive interpolation to create an illusion of continuity for the end user was first conceived by David Roundy and made into a central design philosophy of the code. Steven Johnson became involved in this nascent project late in 2003 and provided a loop-in-chunks routine that greatly simplified calculations involving fields that were spread out arbitrarily over the computational cell, parallelization of the code using chunks and a Scheme front-end interface that was bolted on from another electromagnetics software he had developed (MPB). I started working on Meep soon after joining the group in Fall 2004 following FDTD research at the IBM Almaden Research

Center under the guidance of Dr. Geoffrey W. Burr and have been responsible for: implementing the subpixel averaging algorithm, implementing anisotropic permittivity and permeability, implementing second-order Pockel's nonlinearity, rewriting the entire time stepping portion of the code for improved stability and to replace the original unwieldy Haskell generated files, implementing uniaxial-PML to replace the split-field PML, implementing the frequency-domain solver among many other fixes and improvements. The code, now renamed as Meep, gradually evolved to include more functionality and was released as open-source software in April 2006. The latest release of Meep is currently version 1.1 and the project continues to mature with an expanded feature set that now includes more than just classical electromagnetism but also quantum Casimir phenomena through the work of Alejandro Rodriguez-Wong and Alexander McCauley. Meep has now been made into standard Debian and Ubuntu packages and several researchers at Moscow State University and Ghent University have developed their own freely available Python interface. I am confident that Meep will continue to be enhanced well into the future to entail an even broader set of physical phenomena in optics-related research.

5.4 Grids and Boundary Conditions

The starting point for the FDTD algorithm is the discretization of space and time into a grid. In particular, Meep uses the standard *Yee grid* discretization which staggers the electric and magnetic fields in time and in space, with each field component sampled at different spatial locations offset by half a pixel, allowing the time and space derivatives to be formulated as center-difference approximations [246]. This much is common to nearly every FDTD implementation and is described in detail elsewhere [215]. In order to parallelize Meep, efficiently support simulations with symmetries, and to efficiently store auxiliary fields only in certain regions (for PML absorbing layers), Meep further divides the grid into *chunks* that are joined together into an arbitrary topology via boundary conditions. (In the future, different chunks may have different resolutions to implement a nonuniform grid [18, 117, 173, 258]). Furthermore, we distinguish two coordinate systems: one consisting of integer coordinates on the Yee grid, and one of continuous coordinates in "physical" space that are interpolated as necessary onto the grid (see Ch. 5.5). This section describes those concepts as they are implemented in Meep, as they form a foundation for the remaining sections and the overall design of the Meep software.

5.4.1 Coordinates and grids

The two spatial coordinate systems in Meep are described by the vec, a continuous vector in \mathbb{R}^d (in *d* dimensions), and the **ivec**, an integer-valued vector in \mathbb{Z}^d describing locations on the Yee grid. If **n** is an ivec, the corresponding vec is given by $0.5\Delta x$ **n**, where Δx is the spatial resolution (the same along x, y, and z)—that is, the integer coordinates in an ivec correspond to *half*-pixels, as shown in the right panel of Fig. 5-1. This is to represent locations on the spatial Yee grid, which offsets different field components in space by half a pixel as shown (in 2d) in the right panel of Fig. 5-1. In 3d, the E_x and D_x components are sampled at ivecs $(2\ell + 1, 2m, 2n)$, E_y and D_y are sampled at ivecs $(2\ell, 2m+1, 2n)$, and so on; H_x and B_x are sampled at ivecs $(2\ell, 2m+1, 2n+1)$, H_y and B_y are sampled at ivecs $(2\ell+1, 2m, 2n+1)$, and so on. In addition to these grids for the different field components, we also occasionally refer to the *centered* grid, at odd ivecs $(2\ell + 1, 2m + 1, 2n + 1)$ corresponding to the "center" of each pixel. (The origin of the coordinate systems is an arbitrary ivec that can be set by the user, but is typically the center of the computational volume.) The philosophy of Meep, as described in Ch. 5.5, is that as much as possible the user should be concerned only with continuous physical coordinates (vecs), and the interpolation/discretization onto ivecs occurs internally as transparently as possible.

5.4.2 Grid chunks and owned points

An FDTD simulation must occur within a finite volume of space, the *computational cell*, terminated with some boundary conditions and possibly by absorbing PML regions as described below. This (rectilinear) computational cell, however, is further



Figure 5-1: The computational cell is divided into chunks (left) that have a onepixel overlap (gray regions). Each chunk (right) represents a portion of the Yee grid, partitioned into *owned* points (chunk interior) and *not-owned* points (gray regions around the chunk edges) that are determined from other chunks and/or via boundary conditions. Every point in the interior of the computational cell is owned by exactly one chunk, the chunk responsible for timestepping that point.

subdivided into convex rectilinear *chunks*. On a parallel computer, for example, different chunks may be stored at different processors. In order to simplify the calculations for each chunk, we employ the common technique of padding each chunk with extra "boundary" pixels that store the boundary values [137] (shown as gray regions in Fig. 5-1)—this means that the chunks are *overlapping* in the interior of the computational cell, where the overlaps require communication to synchronize the values.

More precisely, the grid points in each chunk are partitioned into *owned* and *not-owned* points. The *not-owned* points are determined by communication with other chunks and/or by boundary conditions. The *owned* points are time-stepped within the chunk, independently of the other chunks (and possibly in parallel), and *every grid point inside the computational cell is owned by exactly one chunk*.

The question then arises: how do we decide which points within the chunk are owned? In order for a grid point to be owned, the chunk must contain all the information necessary for timestepping that point (once the not-owned points have been communicated). For example, for a D_y point $(2\ell, 2m + 1, 2n)$ to be owned, the H_z points at $(2\ell \pm 1, 2m + 1, 2n)$ must both be in the chunk in order to compute $\nabla \times \mathbf{H}$ for timestepping \mathbf{D} at that point. This means that the D_y points along the left (minimum-x) edge of the chunk (as shown in the right panel of Fig. 5-1) cannot be owned: there is no H_z point to the left of it. An additional dependency is imposed by the case of anisotropic media: if there is an ε_{xy} coupling E_x to D_y , then updating E_x at $(2\ell + 1, 2m, 2n)$ requires the four D_y values at $(2\ell + 1 \pm 1, 2m \pm 1, 2n)$ (these are the surrounding D_y values, as seen in the right panel of Fig. 5-1). This means that the E_x (and D_x) points along the *right* (maximum-x) edge of the chunk (as shown in the right panel of Fig. 5-1) cannot be owned either: there is no D_y point to the right of it. Similarly for $\nabla \times \mathbf{D}$ and anisotropic μ .

All of these considerations result in the shaded-gray region of Fig. 5-1(right) being not-owned. That is, if the chunk intersects k+1 pixels along a given direction starting at an **ivec** coordinate of 0 (e.g. k = 5 in Fig. 5-1), the endpoint **ivec** coordinates 0 and 2k + 1 are not-owned and the interior coordinates from 1 to 2k (inclusive) are owned.

5.4.3 Boundary conditions and symmetries

All of the not-owned points in a chunk must be determined by boundary conditions of some sort. The simplest boundary conditions are when the not-owned points are owned by some other chunk, in which case the values are simply copied from that chunk (possibly requiring communication on a multiprocessor system) each time they are updated. In order to minimize communications overhead, all communications between two chunks are batched into a single message (by copying the relevant notowned points to/from a contiguous buffer) rather than sending one message per point to be copied.

At the edges of the computational cell, some user-selected boundary condition must be imposed. For example, one can use perfect electric or magnetic conductors where the relevant electric/magnetic-field components are set to zero at the boundaries. One can also use Bloch-periodic boundary conditions, where the fields on one side of the computational cell are copied from the other side of the computational cell, optionally multiplied by a complex phase factor $e^{ik_i\Lambda_i}$ where k_i is the propagation constant in the ith direction, and Λ_i is the length of the computational cell in the same direction. Meep does *not* implement any absorbing boundary conditions—absorbing boundaries are, instead, handled by an artificial material, perfectly matched layers (PML), placed adjacent to the boundaries [215].

Bloch-periodic boundary conditions are useful in periodic systems [99], but this is only one example of a useful symmetry that may be exploited via boundary conditions. One may also have mirror and rotational symmetries. For example, if the materials and the field sources have a mirror symmetry, one can cut the computational costs in two by storing chunks only in half the computational cell and applying mirror boundary conditions to obtain the not-owned pixels adjacent to the mirror plane. As a more unusual example, consider an S-shaped structure as in Fig. 5-2, which has no mirror symmetry but is symmetric under 180-degree rotation, called C_2 symmetry [93]. Meep can exploit this case as well (assuming the current sources have the same symmetry), storing only half of the computational cell as in Fig. 5-2 and inferring the not-owned values along the dashed line by a 180-degree rotation. (In the simple case where the stored region is a single chunk, this means that the not-owned points are determined by owned points in the same chunk, requiring copies, possibly with sign flips.) Depending on the sources, of course, the fields can be even or odd under mirror flips or C_2 rotations [99], so the user can specify an additional sign flip for the transformation of the vector fields (and pseudovector \mathbf{H} and \mathbf{B} fields, which incur an additional sign flip under mirror reflections [94, 99]). Meep also supports fourfold rotation symmetry (C_4) , where the field can be multiplied by factors of 1, *i*, -1, or -i under each 90-degree rotation [93]. (Other rotations, such as threefold or sixfold, are not supported because they do not preserve the Cartesian Yee grid.) In 2d, the xy plane is itself a mirror plane (unless in the presence of anisotropic materials) and the symmetry decouples TE modes (with fields E_x , E_y , and H_z) from TM modes $(H_x, H_y, and E_z)$ [99]; in this case Meep only allocates those fields for which the corresponding sources are present.



Figure 5-2: Meep can exploit mirror and rotational symmetries, such as the 180-degree (C_2) rotational symmetry of the S-shaped structure in this schematic example. Although Meep maintains the illusion that the entire structure is stored and simulated, internally only half of the structure is stored (as shown at right), and the other half is inferred by rotation. The rotation gives a boundary condition for the not-owned grid points along the dashed line.

A central principle of Meep is that symmetry optimizations be transparent to the user once the desired symmetries are specified. Meep maintains the illusion that the entire computational cell is computed—for example, the fields in the entire computational cell can still be queried or exported to a file, flux planes and similar computations can still extend anywhere within the computational cell, and so on. The fields in the non-stored regions are simply computed behind the scenes (without ever allocating memory for them) by transforming the stored chunks as needed. A key enabling factor for maintaining this illusion efficiently is the *loop-in-chunks* abstraction employed by the Meep code, described in Ch. 5.9.

Meep also supports continuous rotational symmetry around a given axis, where the structure is invariant under rotations and the fields transform as $e^{im\phi}$ for some m [99], but this is implemented separately by providing the option to simulate Maxwell's equations in the (r, z) plane with cylindrical coordinates, for which operators like $\nabla \times$ change form.

5.5 Interpolation and the illusion of continuity

A core design philosophy of Meep is to provide the illusion of continuous space and time, masking the underlying discretization from the user as much as possible. There



Figure 5-3: A key principle of Meep is that continuously varying inputs yield continuously varying outputs. Here, an eigenfrequency of a photonic crystal varies continuously with the eccentricity of a dielectric rod, accomplished by subpixel smoothing of the material parameters, whereas the nonsmoothed result is "stairstepped." Specifically, the plot shows a TE eigenfrequency of 2d square lattice (period *a* of dielectric ellipses (ε =12) in air versus one semi-axis diameter of the ellipse (in gradations of 0.005*a*) for no smoothing (red squares, resolution of 20 pixels/*a*), subpixel smoothing (blue circles, resolution of 20 pixels/*a*) and "exact" results (black line, no smoothing at resolution of 200 pixels/*a*)

are two components to this approach: the input and the outputs. Continuously varying inputs, such as the geometry, materials, and the source currents, lead to continuously varying outputs, as in the example of Fig. 5-3. Similarly, the value of any field (or any function of the fields) can be output at any point in space or integrated over any region. Furthermore, the effects of these inputs and the resulting outputs must converge as quickly as possible to the exact solution as the resolution increases. In this section, we discuss how this illusion of continuity is implemented for field outputs, current inputs, and geometry/materials.

Any field component (or any combinations such as flux, energy, and user-defined functions) can be evaluated at any point in space. In general, this requires interpolation from the Yee grid. Since the underlying FDTD center-difference algorithm has second-order accuracy, we linearly interpolate fields as needed (which also has second-order accuracy for smooth functions). Similarly, we provide an interface to integrate any function of the fields over any convex rectilinear region (boxes, planes, or lines), and the integral is computed by integrating the linear interpolation of the fields within the integration region. This is straightforward, but there are two subtleties due to the staggered Yee grid. First, computation of quantities like $\mathbf{E} \times \mathbf{H}$ that mix different field components requires an additional interpolation: first, the fields are interpolated onto the centered grid (Ch. 5.4), then the integrand is computed, and then the linear interpolation of the integrand is integrated over the specified region. Second, the computation of quantities like $\mathbf{E} \times \mathbf{H}$ mixes two fields that are stored at different times: **H** is stored at times $(n-0.5)\Delta t$, while **E** is stored at times $n\Delta t$ [215]. Simply using these time-offset fields together is only first-order accurate. If secondorder accuracy is desired, Meep provides the option to temporarily synchronize the electric and magnetic fields: the magnetic fields are saved to a backup array, stepped by Δt , and they are averaged with the backup array to obtain the magnetic fields at $n\Delta t$ with $O(\Delta t^2)$ accuracy. (The fields are restored from backup before resuming timestepping.) This restores second-order accuracy at the expense of an extra half a timestep's computation, which is usually negligible because such field computations are rarely required at every timestep of a simulation—see Ch. 5.7 for how Meep performs typical transmission simulations and other calculations efficiently.

The conceptually reversed process is required for specifying sources: the current density is specified at some point (for dipole sources) or in some region (for distributed current sources) in continuous space, and then must be *restricted* to a corresponding current source on the Yee grid. Meep performs this restriction using exactly the same code (the loop-in-chunks abstraction of Ch. 5.9) and the same weights as the interpolation procedure above. Mathematically, we are exploiting a well-known concept (originating in multigrid methods) that restriction can be defined as the *transpose* of interpolation [223]. This is illustrated by a 2d example in Fig. 5-4. Suppose that the bilinear interpolation f (blue) of four grid points (red) is $f = 0.32f_1 + 0.48f_2 + 0.08f_3 + 0.12f_4$, which can be viewed as multiplying a vector of



Figure 5-4: Left: a bilinear interpolation of values $f_{1,2,3,4}$ on the grid (red) to the value f at an arbitrary point. Right: the reverse process is *restriction*, taking a value J at an arbitrary point (e.g. a current source) and converting into values on the grid. Restriction can be viewed as the transpose of interpolation and uses the same coefficients.

those fields by the row-vector [0.32, 0.48, 0.08, 0.12]. Conversely, if we place a pointdipole current source J (blue) at the same point, it is restricted on the grid (red) to values $J_1 = 0.32J$, $J_2 = 0.48J$, $J_3 = 0.08J$, and $J_4 = 0.12J$ as shown in Fig. 5-4, corresponding to multiplying J by the column vector $[0.32, 0.48, 0.08, 0.12]^{T.1}$ Such a restriction has the property of preserving the sum (integral) of the currents, and typically leads to second-order convergence of the resulting fields as the resolution increases (see below). An example of the utility of this continuous restriction process is shown in Fig. 5-5 via the phenomenon of Cerenkov radiation [129]: a point charge qmoving at a constant velocity \mathbf{v} with a magnitude 1.05c/n exceeding the phase velocity c/n in the medium emits a shockwave-like radiation pattern, and this can be directly modelled in Meep by a continuously moving current source $\mathbf{J} = -\mathbf{v}q\delta(\mathbf{x} - \mathbf{v}t)$ [142]. In contrast, pixelizing the motion into discrete jumps to the nearest grid point leads to visible numerical artifacts in the radiation, as seen in the right panel of Fig. 5-5.

All of the second-order accuracy of FDTD and the above interpolations is generally spoiled to only first-order, however, if one directly discretizes a discontinuous material boundary [57,69]. Moreover, directly discretizing a discontinuity in ε or μ leads to "stairstepped" interfaces that can only be varied in discrete jumps of one

¹Technically, for a dipole-current source given by a delta function with amplitude I, the corresponding current density is $J = I/\Delta x^d$ in d dimensions.



Figure 5-5: Cerenkov radiation emitted by a point charge moving at a speed v = 1.05c/n exceeding the phase velocity of light in a homogeneous medium of index n=1.5. Thanks to Meep's interpolation (or technically *restriction*), the smooth motion of the source current (left panel) can be expressed as continuously varying currents on the grid, whereas the non-smooth pixelized motion (no interpolation) (right panel) reveals high-frequency numerical artifacts of the discretization (counterpropagating wavefronts behind the moving charge).

pixel at a time. Both of these problems are solved in Meep by using an appropriate subpixel smoothing of ε and μ : before discretizing, discontinuities are smoothed into continuous transitions over a distance of one pixel Δx , using a carefully designed averaging procedure. Any subpixel smoothing technique will achieve the goal of continuously varying results as the geometry is continuously varied. In the case of Meep this is illustrated by Fig. 5-3: in a 2d photonic crystal (square lattice of dielectric rods), the lowest TE-polarization eigenfrequency (computed as in Ch. 5.7) varies continuously with the eccentricity of the elliptical rods for subpixel averaging, whereas the nonaveraged discontinuous discretization produces a stairstepped discontinuous eigenfrequency. On the other hand, most subpixel smoothing techniques will not increase the accuracy of FDTD—on the contrary, smoothing discontinuous interfaces changes the structure, and generally introduces *additional* error into the simulation [69]. In order to design an accurate smoothing technique, we exploited recent results in perturbation theory that show how a particular subpixel smooth-



Figure 5-6: Appropriate subpixel averaging can *increase* the accuracy of FDTD with discontinuous materials [69,175]. Here, relative error $\Delta\omega/\omega$ (comparing to the "exact" ω_0 from a planewave calculation [107]) for an eigenmode calculation (as in Ch. 5.7) for a cubic lattice (period a) of 3d anisotropic- ε ellipsoids (right inset) versus spatial resolution (units of pixels per vacuum wavelength λ), for a variety of subpixel smoothing techniques. Straight lines for perfect linear (black dashed) and perfect quadratic (black solid) convergence are shown for reference. Most curves are for the first eigenvalue band (left inset shows E_x in xy cross-section of unit cell), with vacuum wavelength $\lambda = 5.15a$. Hollow squares show Meep's method for band 13 (middle inset), with $\lambda = 2.52a$. Meep's method for bands 1 and 13 is shown for resolutions up to 100 pixels/a.

ing can be chosen to yield zero first-order error [69, 105, 124, 175]. The results are shown in Fig. 5-6 and Fig. 5-7: for both computation of the eigenfrequencies (of an anisotropic photonic crystal) in Fig. 5-6 and the scattering loss from a bump on a strip waveguide in Fig. 5-7, the errors in Meep's results decrease quadratically $[O(\Delta x^2)]$, whereas doing no averaging leads to erratic linear convergence $[O(\Delta x)]$. Furthermore, Fig. 5-6 compares to other subpixel-averaging schemes, including the obvious strategy of simply averaging ε within each pixel [56], and shows that they lead to first-order convergence no better than no averaging at all.

The subpixel averaging is discussed in more detail elsewhere [69, 124, 175], so we



Figure 5-7: The relative error in the scattered power from a small semicircular bump in a dielectric waveguide ($\varepsilon = 12$), excited by a point-dipole source in the waveguide (geometry and fields shown in inset), as a function of the computational resolution. Appropriate subpixel smoothing of the dielectric interfaces leads to roughly secondorder $[O(\Delta x^2)]$ convergence (red squares), whereas the unsmoothed structure has only first-order convergence (blue circles).

only briefly summarize it here. In order for the smoothing to yield zero first-order perturbation, the smoothing scheme must be anisotropic. Even if the initial interface is between isotropic materials, one obtains a tensor ε (or μ) which uses the mean ε for fields parallel to the interface and the harmonic mean (inverse of mean of ε^{-1}) for fields perpendicular to the interface—this was initially proposed heuristically [153] and later shown to be justified via perturbation theory [69, 105]. (If the initial materials are anisotropic, a more complicated formula is needed [124, 175].) The key point is that, even if the physical structure consists entirely of isotropic materials, the discretized structure will use anisotropic materials. Stable simulation of anisotropic media requires an FDTD variant recently proposed in Ref. 235.

There are a few limitations to this subpixel averaging. First, the case of perfect metals requires a different approach [6, 158] that is not yet implemented in Meep. Although Meep does not yet implement subpixel averaging for dispersive materials, there is numerical evidence that similar accuracy improvements are obtained in that case by the same technique [54], and we suspect that a similar derivation can be applied (using the unconjugated form of perturbation theory for the complex-symmetric Maxwell equations in reciprocal media with losses [133]. Second, once the smoothing eliminates the first-order error, the presence of sharp corners (associated with field singularities) introduce an error intermediate between first- and second-order [69], which we hope to address in future work. Third, the fields directly *on* the interface are still at best first-order accurate even with subpixel smoothing—however, these localized errors are equivalent to currents that radiate zero power to first order [108, 124]. The improved accuracy from smoothing is therefore obtained for fields evaluated off of the interface as in scattered flux integrated over a surface away from the interface (Fig. 5-7), for nonlocal properties like resonant frequencies and eigenfrequencies (Fig. 5-6), and for overall integrals of fields and energies [to which the interface contributes only $O(\Delta x)$ of the integration domain and hence first-order errors on the interface have a second-order effect].

5.6 Materials

Time-dependent methods for electromagnetism, given their generality, allow for the simulation of a broad range of material systems. Certain classes of materials, particularly active and nonlinear materials which do not conserve frequency, are ideally suited for modelling by such methods. Materials are represented in Maxwell's equations (5.1) and (5.2) via the relative permittivity $\varepsilon(\mathbf{x})$ and permeability $\mu(\mathbf{x})$ which in general depend on position, frequency (material dispersion) and the fields themselves (nonlinearities). Meep currently supports arbitrary anisotropic material tensors, anisotropic dispersive materials (Lorentz–Drude models and conductivities, both magnetic and electric), and nonlinear materials (both second- and third-order nonlinearities), which taken together permit investigations of a wide range of physical phenomena. The implementation of these materials in Meep is mostly based on standard techniques [215], so we will focus here on two places where Meep differs from



Figure 5-8: The performance of a quasi-PML in the radial direction (cylindrical coordinates, left panel) at a resolution 20 pixels/ λ is nearly equivalent to that of a true PML (in Cartesian coordinates, right panel). The plot shows the difference in the electric field E_z (insets) from a point source between simulations with PML thickness L and L + 1, which is a simple proxy for the PML reflections [177]. The different curves are for PML conductivities that turn on as $(x/L)^d$ for d = 1, 2, 3 in the PML, leading to different rates of convergence of the reflection [177].

the usual approach. For nonlinearities, we use a Padé approximant to avoid solving cubic equations at each step. For PML absorbing media in cylindrical coordinates, we only use a "quasi-PML" [138] based on a Cartesian PML, but explain why its performance is comparable to a true PML while requiring less computational effort.

5.6.1 Nonlinear materials

Optical nonlinearities arise when large field intensities induce changes in the local ε or μ to produce a number of interesting effects: temporal and spatial soliton propagation, optical bistability, self-focusing of optical beams, second- and third-harmonic generation, and many other effects [2,24]. Such materials are usually described by a power-series expansion of **D** in terms of **E** and various susceptibilities. In many common materials, or when considering phenomena in a sufficiently narrow bandwidth (such as the resonantly enhanced nonlinear effects [194] well-suited to FDTD calculations), these nonlinear susceptibilities can be accurately approximated via nondispersive (instantaneous) effects [28]. Meep supports instantaneous isotropic (or diagonal anisotropic) nonlinear effects of the form:

$$D_i - P_i = \varepsilon^{(1)} E_i + \chi_i^{(2)} E_i^2 + \chi_i^{(3)} |\mathbf{E}|^2 E_i, \qquad (5.3)$$

where $\varepsilon^{(1)}$ represents all the linear nondispersive terms and P_i is a dispersive polarization $\mathbf{P} = \chi^{(1)}_{\text{dispersive}}(\omega)\mathbf{E}$ from dispersive materials such as Lorentz media [215]. (A similar equation relates **B** and **H**.) Implementing this equation directly, however, would require one to solve a cubic equation at each time step [215, sec. 9.6], since **D** is updated from $\nabla \times \mathbf{H}$ before updating **E** from **D**.

However, eq. (5.3) is merely a power series approximation for the true material response, valid for sufficiently small field intensities, so it is not necessary to insist that it be solved exactly. Instead, we approximate the solution of eq. (5.3) by a Padé approximant [11], which matches the "exact" cubic solution to high-order accuracy by the rational function:

$$E_{i} = \left[\frac{1 + \left(\frac{\chi^{(2)}}{[\varepsilon^{(1)}]^{2}}\tilde{D}_{i}\right) + 2\left(\frac{\chi^{(3)}}{[\varepsilon^{(1)}]^{3}}\|\tilde{\mathbf{D}}\|^{2}\right)}{1 + 2\left(\frac{\chi^{(2)}}{[\varepsilon^{(1)}]^{2}}\tilde{D}_{i}\right) + 3\left(\frac{\chi^{(3)}}{[\varepsilon^{(1)}]^{3}}\|\tilde{\mathbf{D}}\|^{2}\right)}\right] \left[\varepsilon^{(1)}\right]^{-1}\tilde{D}_{i},$$
(5.4)

where $\tilde{D}_i = D_i - P_i$. For the case of isotropic $\varepsilon^{(1)}$ and $\chi^{(2)} = 0$, so that we have a purely Kerr $(\chi^{(3)})$ material, this matches the "exact" cubic E to $O(D^7)$ error. With $\chi^{(2)} \neq 0$, the error is $O(D^4)$.

For more complicated dispersive nonlinear media or for arbitrary anisotropy in $\chi^{(2)}$ or $\chi^{(3)}$, one approach that Meep may implement in the future is to incorporate the nonlinear terms in the auxiliary differential equations for a Lorentz medium [215].

5.6.2 Absorbing boundary layers: PML, pseudo-PML, and quasi-PML

A perfectly matched layer (PML) is an artificial absorbing medium that is commonly used to truncate computational grids for simulating wave equations (e.g. Maxwell's equations), and is designed to have the property that interfaces between the PML and adjacent media are reflectionless in the exact wave equation [215]. There are various interchangeable formulations of PML for FDTD methods [215], which are all equivalent to a coordinate stretching of Maxwell's equations into complex spatial coordinates; Meep implements a version of the uniaxial PML (UPML), expressing the PML as an effective dispersive anisotropic ε and μ [215]. Meep provides support for arbitrary user-specified PML absorption profiles (which have an important influence on reflections due to discretization error and other effects) for a given round-trip reflection (describing the strength of the PML in terms of the amplitude of light passing through the PML, reflecting off the edge of the computational cell, and propagating back) [177]. For the case of periodic media such as photonic crystals, the medium is not analytic and the premise of PML's reflectionless property is violated; in this case, a "PML" material overlapped with the photonic crystal is only a "pseudo-PML" that is reflectionless only in the limit of a sufficiently thick and gradual absorber, and control over the absorption profile is important [177].

For the radial direction in cylindrical coordinates, a true PML can be derived by coordinate-stretching, but it requires more storage and computational effort than the Cartesian UPML [84, 217], as well as increasing code complexity. Instead, we chose to implement a *quasi-PML* [138], which simply consists of using the Cartesian UPML materials as an approximation for the true radial PML. This approximation becomes more and more accurate as the outer radius of the computational cell increases, because the implicit curvature of the PML region decreases with radius and approaches the Cartesian case. Furthermore, one must recall that *every* PML has reflections once space is discretized [215], which can be mitigated by gradually turning on the PML absorption over a finite-thickness PML layer. The quasi-PML approximation

is likewise mitigated by the same gradual absorption profile, and the only question is that of the constant factor in the reflection convergence: how thick does the quasi-PML need to be to achieve low reflections, compared to a true PML? Figure 5-8 shows that, for a typical calculation, the performance of the quasi-PML in cylindrical coordinates (left) is comparable to that of a true PML in Cartesian coordinates (right). Here, we plot a measure of the reflection from the PML as a function of the PML absorber length L, for a fixed round-trip reflection [177], using as a measure of the reflection the "field convergence" factor: the difference between the \mathbf{E} field at a given point for simulations with PML absorber lengths L and L + 1. The PML conductivity $\sigma(x)$ is turned on gradually as $(x/L)^d$ for d = 1, 2, 3, and it can be shown that this leads to reflections that decrease as $1/L^{2d+2}$ and field-convergence factors that decrease as $1/L^{2d+4}$ [177]. Precisely these decay rates are observed in Fig. 5-8, with similar constant coefficients. As the resolution is increased (approaching the exact wave equations), the constant coefficient in the Cartesian PML plot will decrease (approaching zero reflection), while the quasi-PML's constant coefficient will saturate at some minimum (corresponding to its finite reflectivity in the exact wave equation for a fixed L). This difference seems of little practical concern, however, because the reflection from a one-wavelength thick quasi-PML at a moderate resolution $(20 \text{ pixels}/\lambda)$ is already so low.

5.7 Enabling typical computations

Simulating Maxwell's equations in the time domain enables the investigation of problems inherently involving multiple frequencies, such as nonlinearities and active media. However, it is also well adapted to solving frequency domain problems since it can solve large bandwidths at once, for example analyzing resonant modes or computing transmission/reflection spectra. In this section, we describe techniques Meep uses to efficiently compute scattering spectra and resonant modes in the time domain. Furthermore, we describe how the time domain method can be adapted to a purely frequency domain solver while sharing almost all of the underlying code.

5.7.1 Computing flux spectra

A principle task of computational time-domain tools are investigations of transmission or scattering spectra from arbitrary structures, where one wants to compute the transmitted or scattered power in a particular direction as a function of the frequency of incident light. One can solve for the power at many frequencies in a single timedomain simulation by Fourier transforming the response to a short pulse. Specifically, for a given surface S, one wishes to compute the integral of the Poynting flux:

$$P(\omega) = \Re \oint_{S} \mathbf{E}_{\omega} \left(\mathbf{x} \right)^{*} \times \mathbf{H}_{\omega} \left(\mathbf{x} \right) d\mathbf{A},$$
(5.5)

where \mathbf{E}_{ω} and \mathbf{H}_{ω} are the fields produced by a source at frequency ω , and \Re denotes the real part of the expression. The basic idea, in time-domain, is to use a short-pulse source (covering a wide bandwidth including all frequencies of interest), and compute \mathbf{E}_{ω} and \mathbf{H}_{ω} from the Fourier transforms of $\mathbf{E}(t)$ and $\mathbf{H}(t)$. There are several different ways to compute these Fourier transforms. For example, one could store the electric and magnetic fields throughout S over all times and at the end of the simulation perform a discrete-time Fourier transform (DTFT) of the fields:

$$\mathbf{E}_{\omega} = \sum_{n} e^{i\omega n\Delta t} \mathbf{E}(n\Delta t) \Delta t, \qquad (5.6)$$

for all frequencies (ω) of interest, possibly exploiting a fast Fourier transform (FFT) algorithm. Such an approach has the following computational cost: for a simulation having T timesteps, $F \ll T$ frequencies to compute, N_S fields in the flux region and N pixels in the entire computational cell this approach requires $\Theta(N + N_S T)$ storage and $\Theta(NT + T \log T)$ time (using a FFT-based chirp-z algorithm [10]).² The difficulty with this approach is that if a long simulation (large T) is required to obtain a high frequency resolution by the usual uncertainty relation [174], then the $\Theta(N_S T)$ storage requirements for the fields $\mathbf{E}(t)$ and $\mathbf{H}(t)$ at each point in S become excessive. Instead, Meep accumulates the DTFT summation of the fields at every point in

²Here, Θ has the usual meaning of an asymptotic tight bound [50].



Figure 5-9: Relative error in the quality factor Q for a photonic-crystal resonant cavity (inset, period a) with $Q \sim 10^6$, versus simulation time in units of optical periods of the resonance. Blue circles: filter-diagonalization method. Red squares: least-squares fit of energy in cavity to a decaying exponential. Filter-diagonalization requires many fewer optical periods than the decay time Q, whereas curve fitting requires a simulation long enough for the fields to decay significantly.

S as the simulation progresses; once the time stepping has terminated, eq. (5.5) can be evaluated using these Fourier-transformed fields.³ The computational cost of this approach is $\Theta(N + N_S F)$ storage [much less than $\Theta(N_S T)$ if $F \ll T$] and $\Theta(NT + N_S FT)$ time. Although our current approach works well, another possible approach that we have been considering is to use Padé approximation: one stores the fields at every timestep on S, but instead of using the DTFT one constructs a Padé approximant to extrapolate the infinite-time DTFT from a short time series [79]. This requires $\Theta(N + N_S T)$ storage (but T is potentially much smaller) and $O(NT + T \log^2 T)$ time [31].

 $^{^{3}}$ It is tempting to instead accumulate the Fourier transform of the Poynting flux at each time, but this is not correct since the flux is not a linear function of the fields.

5.7.2 Analyzing resonant modes

Another major goal of time-domain simulations is analysis of resonant phenomena, specifically by determining the resonant frequency ω_0 and the quality factors Q (i.e., the number of optical cycles $2\pi/\omega_0$ for the field to decay by $e^{-2\pi}$) of one or more resonant modes. One straightforward and common approach to compute ω_0 and Qis by computing the DTFT of the field at some point in the cavity in response to a short pulse [215]: ω_0 is then the center of a peak in the DTFT and 1/Q is the fractional width of the peak at half maximum. The problem with this approach is that the Fourier uncertainty relation (equivalently, spectral leakage from the finite time window [174]) means that resolving the peak in this way requires a simulation much longer than Q/ω_0 (problematic for structures that may have very high Q, even 10^9 or higher [193]). Alternatively, one can perform a least squares fit of the field time-series within the cavity to an exponentially decaying sinusoid, but this leads to an ill-conditioned, non-convex, nonlinear fitting problem (and is especially difficult if more than one resonant mode may be present). If only a single resonant mode is present, one can perform a least-squares fit of the energy in the cavity to a decaying exponential in order to determine Q, but a long simulation is still required to accurately resolve a large Q (as shown below). A more accurate and efficient approach, requiring only a short simulation even for very large Q values, is the technique of *fil*ter diagonalization originally developed for NMR spectroscopy, which transforms the time-series data into a small eigenproblem that is solved for all resonant frequencies and quality factors at once (even for multiple overlapping resonances) [145]. Chapter 16 of Ref. 215 compared the DFT peak-finding method with filter-diagonalization by attempting to resolve two near-degenerate modes in a microcavity, and demonstrated the latter's ability to accurately resolve closely-spaced peaks with as much as a factor of five times fewer timesteps. Rodriguez et al. [193] have used filter diagonalization to compute quality factors of 10^8 or more using simulations only a few hundred optical cycles in length. We quantify the ability of filter diagonalization to resolve a large $Q \sim 10^6$ in Fig. 5-9, comparing the relative error in Q versus simulation time for filter diagonalization and the least-squares energy-fit method above. (The specific cavity is formed by a missing rod in a two-dimensional photonic crystal consisting of a square lattice of dielectric rods in air with period a, radius 0.2a, and $\varepsilon = 12$ [99].) Figure 5-9 demonstrates that filter diagonalization is able to identify the quality factor using almost an order of magnitude fewer time steps than the curve fitting method. (Another possible technique to identify resonant modes uses Padé approximants, which can also achieve high accuracy from a short simulation [55,79].)

5.7.3 Frequency-domain solver

A common electromagnetic problem is to find the fields that are produced in a geometry in response to a source at a single frequency ω . In principle, the solution of such problems need not involve time at all, but involve solving a linear equation purely in the frequency domain [99, appendix D]; this can be achieved by many methods, such as finite-element methods [98, 198, 245], boundary-element methods [25, 41, 188, 225], or finite-difference frequency-domain methods [44]. However, if one already has a fullfeatured parallel FDTD solver, it is attractive to exploit that solver for frequencydomain problems when they arise. The most straightforward approach is to simply run a simulation with a constant-frequency source—after a long time, when all transient effects from the source turn-on have disappeared, the result is the desired frequency-domain response. The difficulty with this approach is that a very long simulation may be required, especially if long-lived resonant modes are present at nearby frequencies (in which case a time $\gg Q/\omega$ is required to reach steady state). Instead, we show how the FDTD time-step can be used to directly plug a frequency-domain problem into an iterative linear solver, finding the frequency-domain response in the equivalent of many fewer timesteps while exploiting the FDTD code almost without modification.

The central component of any FDTD algorithm is the time step: an operation that advances the field by Δt in time. In order to extract a frequency-domain problem from this operation, we first express the timestep as an abstract linear operation: if \mathbf{f}^n represents all of the fields (electric and magnetic) at time step n, then (in a linear time-invariant structure) the time step operation can be expressed in the form:

$$\mathbf{f}^{n+1} = \hat{T}_0 \mathbf{f}^n + \mathbf{s}^n, \tag{5.7}$$

where \hat{T}_0 is the timestep operator with no sources and \mathbf{s}^n are the source terms (currents) from that time step. Now, suppose that one has a time-harmonic source $\mathbf{s}^n = e^{-i\omega n\Delta t}\mathbf{s}$ and wish to solve for the resulting time-harmonic (steady state) fields $\mathbf{f}^n = e^{-i\omega n\Delta t}\mathbf{f}$. Substituting these into eq. (5.7), we obtain the following linear equation for the field amplitudes \mathbf{f} :

$$\left(\hat{T}_0 - e^{-i\omega\Delta t}\right)\mathbf{f} = -\mathbf{s}.$$
(5.8)

This can then be solved by an iterative method, and the key fact is that iterative methods for Ax = b only require one to supply a function that multiplies the linear operator A by a vector [13]. Here, A is represented by $\hat{T}_0 - e^{-i\omega\Delta t}$ and hence one can simply use a standard iterative method by calling the unmodified timestep function from FDTD to provide the linear operator. To obtain the proper right-hand side \mathbf{s} , one merely needs to execute a single timestep (5.7), with sources, starting from zero field $\mathbf{f} = 0$. Since in general this linear operator is not Hermitian (especially in the presence of PML absorbing regions), we employ the BiCGSTAB-*L* algorithm (a generalization of the stabilized biconjugate gradient algorithm, where increasing the integer parameter *L* trades off increased storage for faster convergence) [204, 205].

This technique means that all of the features implemented in our time-domain solver (not only arbitrary materials, subpixel averaging, and other physical features, but also parallelization, visualization, and user-interface features) are immediately available as a frequency-domain solver. To demonstrate the performance of this frequency-domain solver over the straightforward approach of simply running a long simulation until transients have disappeared, we computed the root-mean-square error in the field as a function of the number of time steps (or evaluations of \hat{T}_0 by BiCGSTAB-L) for two typical simulations. The first simulation, shown in Fig. 5-10, consists of a point source in vacuum surrounded by PML (inset). The frequency-



Figure 5-10: Root-mean-square error in fields in response to a constant-frequency point source in vacuum (inset), for frequency-domain solver (red squares, adapted from Meep time-stepping code) vs. time-domain method (blue circles, running until transients decay away).

domain solver (red squares) shows rapid, near-exponential convergence, while the error in the time-domain method (blue circles) decreases far more gradually (in fact, only polynomially). A much more challenging problem is to obtain the frequency-domain response of a cavity (ring resonator) with multiple long-lived resonant modes: in the time domain, these modes require a long simulation ($\sim Q$) to reach steady state, whereas in the frequency domain the resonances correspond to poles (near-zero eigenvalues of A) that increase the condition number and hence slow convergence [13]. Figure 5-11 shows the results for a ring resonator cavity with multiple closely-spaced resonant modes, excited at one of the resonant frequencies (inset)—although both frequency- and time-domain methods take longer to converge than for the non-resonant case of Fig. 5-10, the advantage of the frequency-domain's exponential convergence is even more clear. The convergence is accelerated in frequency domain by using L = 10 (green diamonds) rather than L = 2 (at the expense of more



Figure 5-11: Root-mean-square error in fields in response to a constant-frequency point source exciting one of several resonant modes of a dielectric ring resonator (inset, $\varepsilon = 11.56$), for frequency-domain solver (red squares, adapted from Meep timestepping code) vs. time-domain method (magenta triangles, running until transients decay away). Green diamonds show frequency-domain BiCGSTAB-*L* solver for five times more storage, accelerating convergence. Blue circles show time-domain method for a more gradual turn-on of source, which avoids exciting long-lived resonances at other frequencies.

storage). In time domain, the convergence is limited by the decay of high-Q modes at other frequencies, and the impact of these modes can be reduced by turning on the constant-frequency source more gradually (magenta triangles, hyperbolic-tangent turn-on of the source over 175 optical periods).

This is by no means the most sophisticated possible frequency-domain solver. For example, we currently do not use any preconditioner for the iterative scheme [13]. In two dimensions, a sparse-direct solver may be far more efficient than an iterative scheme [13]. The key point, however, is that programmer time is much more expensive than computer time, and this technique allows us to obtain substantial improvements in solving frequency-domain problems with only minimal changes to an existing FDTD program.

5.8 User interface and scripting

In designing the style of user interaction in Meep, we were guided by two principles. First, in research or design one hardly ever needs just *one* simulation—one almost always performs a whole series of simulations for a class of related problems, exploring the parameter dependencies of the results, optimizing some output as a function of the input parameters, or looking at the same geometry under a sequence of different stimuli. Second, there is the Unix philosophy: "Write programs that do one thing and do it well" [200]—Meep should perform electromagnetic simulations, while for additional functionality it should be combined with other programs and libraries via standard interfaces like files and scripts.

Both of these principles argue against the graphical CAD-style interface common in commercial FDTD software. First, while graphical interfaces provide a quick and attractive route to setting up a single simulation, they are not so convenient for a series of related simulations. One commonly encounters problems where the size/position of certain objects is determined by the size/position of other objects, where the number of objects is itself a parameter (such as a photonic-crystal cavity surrounded by a variable number of periods [99]), where the length of the simulation is controlled by a complicated function of the fields, where one output is optimized as a function of some parameter, and many other situations that become increasingly cumbersome to express via a set of graphical tools and dialog boxes. Second, we don't *want* to write a mediocre CAD program—if we wanted to use a CAD program, we would use a professional-quality one, export the design to a standard interchange format, and write a conversion program to turn this format into what Meep expects. The most flexible and self-contained interface is, instead, to allow the user to control the simulation via an arbitrary program. Meep allows this style of interaction at two levels: via a low-level C++ interface, and via a standard high-level scripting language (Scheme) implemented by an external library (GNU Guile). The potential slowness



Figure 5-12: A simple Meep example showing the E_z field in a dielectric waveguide ($\varepsilon = 12$) from a point source at a given frequency. A plot of the resulting field (blue/white/red = positive/zero/negative) is in the background, and in the foreground is the input file in the high-level scripting interface (the Scheme language).

of the scripting language is irrelevant because all of the expensive parts of the FDTD calculation are implemented in C/C++.

The high-level scripting interface to Meep is documented in detail, with several tutorials, on the Meep web page (http://ab-initio.mit.edu/meep), so we restrict ourselves to a single short example in order to convey the basic flavor. This example, in Fig. 5-12, computes the (2d) fields in response to a point source located within a dielectric waveguide. We first set the size of the computational cell to 16×8 (via geometry-lattice, so-called because it determines the lattice vectors in the periodic case)—recall that the interpretation of the unit of distance is arbitrary and up to the user (it could be $16 \,\mu\text{m} \times 8 \,\mu\text{m}$, in which case the frequency units are $c/\mu\text{m}$, or $16 \,\text{mm} \times 8 \,\text{mm}$ with frequency units of c/mm, or any other convenient distance unit). Let us call this arbitrary unit of distance a. Then we specify the geometry within the cell as a list of geometric objects like blocks, cylinders, etcetera—in this case by a single block defining the waveguide with $\varepsilon = 12$ —or optionally by an arbitrary user-defined function $\varepsilon(x, y)$ (and μ , etcetera). A layer of PML is then specified around the

boundaries with thickness 1; this layer lies *inside* the computational cell and overlaps the waveguide, which is necessary in order to absorb waveguide modes when they reach the edge of the cell. We add a point source, in this case an electric-current source **J** in the z direction (sources of arbitrary spatial profile can also be specified). The time-dependence of the source is a sharp turn-on to a continuous-wave source $\cos(\omega t)$ at the beginning of the simulation; gradual turn-ons, Gaussian pulses, or arbitrary user-specified functions of time can also be specified. The frequency is 0.15 in units of c/a, corresponding to a vacuum wavelength $\lambda = a/0.15$ (e.g. $\lambda \approx 6.67 \,\mu\text{m}$ if $a = 1 \,\mu\text{m}$). We set the resolution to 10 pixels per unit distance (10 pixels/a), so that the entire computational cell is 160×80 pixels, and then run for 200 time units (units of a/c), corresponding to $200 \times 0.15 = 30$ optical periods. We output the dielectric function at the beginning, and the E_z field at the end.

In keeping with the Unix philosophy, Meep is not a plotting program; instead, it outputs fields and related data to the standard HDF5 format for scientific datasets [71], which can be read by many other programs and visualized in various ways. (We also provide a way to effectively "pipe" the HDF5 output to an external program within Meep: for example, to output the HDF5 file, convert it immediately to an image with a plotting program, and then delete the HDF5 file; this is especially useful for producing animations consisting of hundreds of frames.)

Another important technique to maintain flexibility is that of higher-order functions [1]: wherever it is practical, our functions take functions as arguments instead of (or in addition to) numbers. Thus, for example, instead of specifying special input codes for all possible source distributions in space and time, we simply allow a user-defined function to be used. More subtly, the arguments **output-epsilon** and **output-efield-z** to the **run-until** function in Fig. 5-12 are actually functions themselves: we allow the user to pass arbitrary "step functions" to **run-until** that are called after every FDTD timestep and which can perform arbitrary computations on the fields as desired (or halt the computation if a desired condition is reached). The **output-efield-z** is simply a predefined step function that outputs E_z . These step-functions can be modified by transformation functions like **at-end**, which take step functions as arguments and return a new step function that only calls the original step functions at specified times (at the end of the simulation, or the beginning, or at certain intervals, for example). In this way, great flexibility in the output and computations is achieved. One can, for example, output a given field component only at certain time intervals after a given time, and only within a certain subvolume or slice of the computational cell, simply by composing several of these transformations. One can even output an arbitrary user-defined function of the fields instead of predetermined components.

There is an additional subtlety when it comes to field output, because of the Yee lattice in which different field components are stored at different points; presented in this way to the user, it would be difficult to perform post-processing involving multiple field components, or even to compare plots of different field components. As mentioned in Ch. 5.5 and again in Ch. 5.9.2, therefore, the field components are automatically interpolated from the Yee grid onto a fixed "centered" grid in each pixel when exported to a file.

Although at a simplistic level the input format can just be considered as a file format with a lot of parentheses, because Scheme is a full-fledged programming language one can control the simulation in essentially arbitrary ways. Not only can one write loops and use arithmetic to define the geometry and the relationships between the objects or perform parameter sweeps, but we also expose external libraries for multivariable optimization, integration, root-finding, and other tasks in order that they can be coupled with simulations.

Parallelism is completely transparent to the user: exactly the same input script is fed to the parallel version of Meep (written with the MPI message-passing standard for distributed-memory parallelism [72]) as to the serial version, and the distribution of the data across processors and the collection of results is handled automatically.

5.9 Abstraction versus performance

In an FDTD simulation, essentially just one thing has to be fast: inner loops over all the grid points or some large fraction thereof. Everything else is negligible in terms of computation time (but not programmer time!), so it can use high-level abstractions without penalty—for example, the use of a Scheme interpreter as the user interface has no performance consequences for a typical computation, because the inner loops are not written in Scheme.⁴ For these inner loops, however, there is a distinct tension between abstraction (or simplicity) and performance, and in this section we discuss some of the tradeoffs that result from this tension and the choices that have been made in Meep.

The primacy of inner loops means that some popular principles of abstraction must be discarded. A few years ago, a colleague of ours attempted to write a new FDTD program in textbook object-oriented C++ style: every pixel in the grid would be an object, and every type of material would be a subclass overriding the necessary timestepping and field-access operations. Timestepping would consist of looping over the grid, calling some "step" method of each object, so that objects of different materials (magnetic, dielectric, nonlinear etcetera) would dynamically apply the corresponding field-update procedures. The result of this noble experiment was a working program but a performance failure, many times slower than the aging Fortran software it was intended to replace: the performance overhead of object dereferencing, virtual method dispatch, and function calls in the inner loop overwhelmed all other considerations. In Meep, each field's components are stored as simple linear arrays of floating-point numbers in row-major (C) order (parallel-array data structures worthy of Fortran 66), and there are separate inner loops for each type of material (more on this below). In a simple experiment on a 2.8 GHz Intel Core 2 CPU, merely moving the if statements for the different material types into these inner loops decreased

⁴The exception to this rule is when the user supplies a Scheme function and asks that it be evaluated for every grid point, for example to integrate some function of the fields. If this is done frequently during the simulation, it is slow; in these circumstances, however, the user can replace the Scheme function with one written in C/C++ if needed. This is rare because most such functions that might be used frequently during a simulation, such as energy or flux, are already supplied in C/C++ within Meep.

Meep's performance by a factor of two in a typical 3d calculation and and by a factor of six in 2d (where the calculations are simpler and hence the overhead of the conditionals is more significant). The cost of the conditionals, including the cost of mispredicted branches and subsequent pipeline stalls [85] along with the frustration of compiler unrolling and vectorization, easily overwhelmed the small cost of computing, e.g., $\nabla \times \mathbf{H}$ at a single point.

5.9.1 Timestepping and cache tradeoffs

One of the dominant factors in performance on modern computer systems is not arithmetic, but memory: random memory access is far slower than arithmetic, and the organization of memory into a hierarchy of caches is designed to favor locality of access [85]. That is, one should organize the computation so that as much work as possible is done with a given datum once it is fetched (temporal locality) and so that subsequent data that are read or written are located nearby in memory (spatial locality). The optimal strategies to exploit both kinds of locality, however, appear to lead to sacrifices of abstraction and code simplicity so severe that we have chosen instead to sacrifice some potential performance in the name of simplicity.

As it is typically described, the FDTD algorithm has very little temporal locality: the field at each point is advanced in time by Δt , and then is not modified again until *all* the fields at *every* other point in the computational cell have been advanced. In order to gain temporal locality, one must employ *asynchronous timestepping*: essentially, points in small regions of space are advanced several steps in time before advancing points far away, since over a short time interval the effects of far-away points cannot cannot be felt. A detailed analysis of the characteristics of this problem, as well as a beautiful "cache-oblivious" algorithm that automatically exploits a cache of any size for grids of any dimensionality, is described in Ref. 73. On the other hand, an important part of Meep's usability is the abstraction that the user can perform arbitrary computations or output using the fields in any spatial region at any time, which seems incompatible with the fields at different points in space being outof-sync until a predetermined end of the computation. The bookkeeping difficulty of reconciling these two viewpoints led us to reject the asynchronous approach, despite its potential benefits.

However, there may appear to be at least a small amount of temporal locality in the synchronous FDTD algorithm: first **B** is advanced from $\nabla \times \mathbf{E}$, then **H** is computed from **B** and μ , then **D** is advanced from $\nabla \times \mathbf{H}$, then **E** is computed from **D** and ε . Since most fields are used at least once after they are advanced, surely the updates of the different fields can be merged into a single loop, for example advancing **D** at a point and then immediately computing **E** at the same point—the **D** field need not even be stored. Furthermore, since by merging the updates one is accessing several fields around the same point at the same time, perhaps one can gain spatial locality by interleaving the data, say by storing an array of $(\mathbf{E}, \mathbf{H}, \varepsilon, \mu)$ tuples instead of separate arrays. Meep does not do either of these things, however, for two reasons, the first of which is more fundamental. As is well-known, one cannot easily merge the **B** and **H** updates with the **D** and **E** updates at the same point, because the discretized $\nabla \times$ operation is nonlocal (involves multiple grid points)—this is why one normally updates **H** everywhere in space before updating **D** from $\nabla \times \mathbf{H}$, because in computing $\nabla \times \mathbf{H}$ one uses the values of \mathbf{H} at different grid points and all of them must be in sync. A similar reasoning, however, applies to updating E from D and **H** from **B**, once the possibility of anisotropic materials is included—because the Yee grid stores different field components at different locations, any accurate handling of off-diagonal susceptibilities must also inevitably involve fields at multiple points (as in Ref. 235). To handle this, **D** must be stored explicitly and the update of **E** from **D** must take place after **D** has been updated everywhere, in a separate loop. And since each field is updated in a separate loop, the spatial-locality motivation to merge the field data structures rather than using parallel arrays is largely removed.

Of course, not all simulations involve anisotropic materials—although they appear even in many simulations with nominally isotropic materials thanks to the subpixel averaging discussed in Ch. 5.5—but this leads to the second practical problem with merging the **E** and **D** (or **H** and **B**) update loops: the combinatorial explosion of the possible material cases. The update of **D** from $\nabla \times \mathbf{H}$ must handle 16 possible
cases, each of which is a separate loop (see above for the cost of putting conditionals inside the loops): with or without PML (4 cases, depending upon the number of PML conductivities and their orientation relative to the field), with or without conductivity, and with the derivative of two **H** components (3d) or only one **H** component (2d TE polarization). The update of **E** from **D** involves 12 cases: with or without PML (2 cases, distinct from those in the **D** update), the number of off-diagonal ε^{-1} components (3 cases: 0, 1, or 2), and with or without nonlinearity (2 cases). If we attempted to join these into a single loop, we would have $16 \times 12 = 192$ cases, a code-maintenance headache. (Note that the multiplicity of PML cases comes from the fact that, including the corners of the computational cell, we might have 0 to 3 directions of PML, and the orientation of the PML directions relative to a given field component matters greatly.)

The performance penalty of separate **E** and **D** (or **H** and **B**) updates appears to be modest. Even if, by somehow merging the loops, one assumes that the time to compute $\mathbf{E} = \varepsilon^{-1} \mathbf{D}$ could become *zero*, benchmarking the relative time spent in this operation indicates that a typical 3d transmission calculation would be accelerated by only around 30% (and less in 2d).

5.9.2 The loop-in-chunks abstraction

Finally, let us briefly mention a central abstraction that, while not directly visible to end-users of Meep, is key to the efficiency and maintainability of large portions of the software (field output, current sources, flux/energy computations and other field integrals, and so on). The purpose of this abstraction is to mask the complexity of the partitioning of the computational cell into overlapping chunks connected by symmetries, communication, and other boundary conditions as described in Ch. 5.4.

Consider the output of the fields at a given timestep to an HDF5 datafile. Meep provides a routine get-field-pt that, given a point in space, interpolates it onto the Yee grid and returns a desired field component at that point. In addition to interpolation, this routine must also transform the point onto a chunk that is actually stored (using rotations, periodicity, etcetera) and communicate the data from

another processor if necessary. If the point is on a boundary between two chunks, the interpolation process may involve multiple chunks, multiple rotations etcetera, and communications from multiple processors. Because this process involves only a single point, it is not easily parallelizable. Now, to output the fields everywhere in some region to a file, one approach is to simply call get-field-pt for every point in a grid for that region and output the results, but this turns out to be tremendously slow because of the repeated transformations and communications for every single point. We nevertheless want to interpolate fields for output rather than dumping the raw Yee grid, because it is much easier for post-processing if the different field components are interpolated onto the same grid; also, to maintain transparency of features like symmetry one would like to be able to output the whole computational cell (or an arbitrary subset) even if only a part of it is stored. Almost exactly the same problems arise for integrating things like flux $\mathbf{E} \times \mathbf{H}$ or energy or user-defined functions of the fields (noting that functions combining multiple field components require interpolation), and also for implementing volume (or line, or surface) sources which must be projected onto the grid in some arbitrary volume.

One key to solving this difficulty is to realize that, when the field in some volume V is needed (for output, integration, and so on), the rotations, communications, etcetera for points in V are identical for all the points in the intersection of V with some chunk (or one of its rotations/translations). The second is to realize that, when interpolation is needed, there is a particular grid for which interpolation is easy: for *owned* points of the *centered* grid (Ch. 5.4) lying at the center of each pixel, it is always possible to interpolate from fields on any Yee grid without any inter-chunk communication and by a simple equal-weight averaging of at most 2^d points in d dimensions.

The combination of these two observations leads to the *loop-in-chunks* abstraction. Given a (convex rectilinear) volume V and a given grid (either centered, or one of the Yee-field grids), it computes the intersection of all the chunks and their rotations/translations with V. For each intersection it invokes a caller-specified function, passing the portion of the chunk, the necessary rotations (etc.) of the fields, and interpolation weights (if needed, for the boundary of V). That function then processes the specified portion of the chunk (for example, outputting it to the corresponding portion of a file, or integrating the desired fields). All of this can proceed in parallel (with each processor considering only those chunks stored locally). This is (relatively) fast because the rotations, interpolations, and so on are computed only once per chunk intersection, while the inner loop over all grid points in each chunk can be as tight as necessary. Moreover, all of the rather complicated and error-prone logic involved in computing V's intersection with the chunks (e.g., special care is required to ensure that each conceptual grid point is processed exactly once despite chunk overlaps and symmetries) is localized to one place in the source code; field output, integration, sources, and other functions of the fields are isolated from this complexity.

5.10 Concluding remarks

We have reviewed in this chapter a number of the unusual implementation details of Meep that distinguish our software package from standard textbook FDTD methods. Beginning with a discussion of the fundamental structural unit of chunks that constitute the Yee grid and enable parallelization: we provided an overview of Meep's core design philosophy of creating an illusion of continuous space and time for inputs and outputs; we explained and motivated the somewhat unusual design intricacies of nonlinear materials and PMLs; we discussed important aspects of Meep's computational methods for flux spectra and resonant modes; we demonstrated the formulation of a frequency-domain solver requiring only minimal modifications to the underlying time-stepping algorithm. In addition to the inner workings of Meep's internal structure, we reviewed how such features are accessible to users via an external scripting interface.

We believe that a free/open-source, full-featured FDTD package like Meep can play a vital role in enabling new research in electromagnetic phenomena. Not only does it provide a low barrier to entry for standard FDTD simulations, but the simplicity of the FDTD algorithm combined with access to the source code offers an easy route to investigate new physical phenomena coupled with electromagnetism. For example, we have colleagues working on coupling multi-level atoms to electromagnetism within Meep for modelling lasing and saturable absorption, adapting published techniques from our and other groups [19,35,89,167,257], but also including new physics such as diffusion of excited gases. Other colleagues have modified Meep for modelling gyromagnetic media in order to design new classes of "one-way" waveguides [229]. Meep is even being used to simulate the quantum phenomena of Casimir forces (from quantum vacuum fluctuations, which can be computed from classical Green's functions) [150,195]—in fact, this was possible without any modifications of the Meep code due to the flexibility of Meep's scripting interface. We hope that other researchers, with the help of the understanding of Meep's architecture that this chapter provides, will be able to adapt Meep to future phenomena that we have not yet envisioned.

Chapter 6

Zero–group-velocity modes in chalcogenide holey photonic-crystal fibers

6.1 Summary

We demonstrate that a holey photonic-crystal fiber with chalcogenide-glass index contrast can be designed to have a complete gap at a propagation constant $\beta = 0$ that also extends into the non-zero β region. This type of bandgap (previously identified only at index contrasts unattainable in glasses) opens up a regime for guiding zero– group-velocity modes not possible in holey fibers with the more common finger-like gaps originating from $\beta \to \infty$. Such modes could be used to enhance nonlinear and other material interactions, such as for hollow-core fibers in gas-sensor applications.

6.2 Introduction

Photonic-crystal holey fibers have been of great interest for a variety of different applications, mainly using silica or polymers with low index contrasts ($\sim 1.5 : 1$) [196]. Researchers have also studied photonic-crystal fiber-like geometries with high index contrast materials (eg. Si or GaAs, index ~ 3.4) [70,80,90,147,154,184,237] and shown

that they support interesting zero-group-velocity modes [90], but to our knowledge such modes have not been described for fibers made of easily drawable materials. In this work, we demonstrate the possibility of obtaining zero-group-velocity modes in uniform fiber geometries using chalcogenide glasses (index \sim 2.8 [16]), which have proven amenable to drawn microstructured fibers [30, 115, 126, 161, 170, 180]. Holey fibers, formed by a lattice of air holes in the fiber cross section, are best known for supporting "finger-like" band gaps opening towards the high-frequency regime, which can open even for arbitrarily low index contrasts [7,99,118,184,196]. However, these gaps close before reaching a zero propagation constant β , and the guided modes that they support have all been found to have nonzero group velocity. If the index contrast is high enough to support a complete band gap for all polarizations in *two* dimensions, however, then the resulting *three* dimensional holey fiber has a gap extending from $\beta = 0$ to some nonzero β . Although such gaps appear in some earlier work for very high index contrasts (3.3-3.5:1), [70, 80, 90, 147, 154, 184, 237] here we point out that they are attainable in lower-contrast glassy materials (chalcogenides). Moreover, we argue that the key advantage of these gaps is that they can support guided fiber modes that have a zero group velocity at $\beta = 0$. The slow-light modes close to the zero-velocity band edge should enhance a wide variety of nonlinear phenomena and material interactions [99, 209], such as fiber-based sensors [87, 95, 121, 191], fiber lasers [136,227], or Raman scattering [15,151], and the band edge should also support gap solitons [210]. Numerous experiments have demonstrated slow-light effects in planar optical devices [9,151,172,202,226]. One simple structure that has a complete two-dimensional (2d) gap for chalcogenide/air index contrast is a triangular lattice of circular air holes. In this chapter, we employ a modification of this structure that is optimized to have a slightly larger gap, but either structure (and any future complete-gap 2d designs) creates well-localized zero-group-velocity fiber modes.

6.3 Review of fiber properties

Let us first review the basic terminology and characteristics of holey fibers [99, 196], and the origin of the gaps and zero–group-velocity modes in this chapter. The propagating modes of a fiber with a constant permittivity cross-section $\varepsilon(x,y)$ can be described as some xy electric field pattern $\mathbf{E}(x, y)$ multiplied by $e^{i(\beta z - \omega t)}$, where β is the propagation constant and ω is the frequency. A holey fiber consists of a periodic cladding (usually a triangular lattice of air holes), as well as a core (solid or hollow) that breaks the periodicity and supports guided modes. The dispersion relation, the plot of $\omega(\beta)$ for all solutions, can be divided into several regions (as in Figs. 6-1) and 6-2). First, there is a continuous (shaded) region, the light cone, consisting of all cladding (non-guided) modes that can propagate in the cladding far from the core. There are also regions of (β, ω) that have no cladding modes: band gaps within the light cone, which can confine gap-guided modes, and also an empty space below the light cone that can confine index-guided modes. The guided modes, exponentially localized to the vicinity of the core, appear as discrete bands $\omega_n(\beta)$ within the gaps and/or under the light cone. (Technically, in a finite-size fiber the gap-guided modes are leaky, but as this leakage rate decreases exponentially with the periodic cladding thickness it can be made negligible in practical contexts.) In order to confine light in an air core, the gaps and guided modes must lie above the light line $\omega = c\beta$ of air (since modes below the light line of air are evanescent in air regions). Normally, these guided bands are monotonically increasing, corresponding to a positive group velocity $d\omega/d\beta$ (and there is a proof that this is always the case for index-guided modes with a homogeneous cladding [12]). Zero group velocity (standing-wave modes) typically occurs only at values of β that have $z \rightarrow -z$ reflection symmetry (with rare exceptions [90, 96]), which in a uniform–cross-section fiber only occurs at $\beta = 0$. Index-guided modes are not possible at $\beta = 0$ (they become rapidly more weakly confined as $\beta \to 0$), so one must consider bandgap-guided modes. Unfortunately, the typical gaps that arise in holey fibers have their origin in the $\beta \to \infty$ limit (where the field patterns approximate those of a 2d metallic system [23,99]), and are observed to



Figure 6-1: Projected band diagram (frequency ω vs. propagation constant β), for a triangular lattice of holes (inset). Inset: optimized 2d ($\beta = 0$) gap size vs. index contrast.

close well before $\beta = 0$ is reached. The $\beta = 0$ point corresponds to a *two*-dimensional photonic crystal with in-plane propagation, whose modes can be decomposed into TE (E in xy plane) and TM (E in z direction) polarizations [99]. Typically, low-contrast materials such as silica/air will have a gap only for one of these polarizations (e.g. TE for air holes) [99]. Such a single-polarization 2d gap is not useful for guiding modes in a fiber, because the TE/TM distinction disappears for $\beta \neq 0$ and hence a single-polarization gap disappears [99]. On the other hand, if one can obtain an overlapping TE/TM gap at $\beta = 0$, which typically requires higher index contrasts, then it should be expected to persist for a nonzero range of β , even after the TE/TM distinction disappears [70,80,90,147,154,184,237]. We demonstrate that this, in fact, occurs, for index contrasts attainable in chalcogenide glasses that have been used for fiber drawing [30, 115, 126, 161, 170, 180], contrary to some previous suggestions [184]. The resulting gap around $\beta = 0$ therefore supports guided modes that attain zero group velocity as $\beta \to 0$. In practice, one does not operate at the zero-velocity point itself, but rather at nearby frequencies, so that by operating closer and closer to the zero-velocity band edge one can make the group velocity of light arbitrarily small in principle (at the expense of bandwidth and greater sensitivity to absorption and other loss, as discussed below).



Figure 6-2: Projected band diagram (frequency ω vs. propagation constant β), for a triangular lattice of hexagonal-shaped holes (inset). Inset: optimized 2d ($\beta = 0$) gap size vs. index contrast.

Several other mechanisms have been proposed for creating zero–group-velocity modes in fibers. Bragg fibers, consisting of concentric rings of two or more materials forming a one-dimensional photonic crystal, have a gap originating at $\beta = 0$ [99, 196] and consequently their guided modes attain zero group velocity at this point. (Although Bragg fibers do not have a complete 2d gap, this is compensated for by the rotational symmetry which eliminates modes propagating in the azimuthal direction at large radii [99].) These fiber modes resemble those of hollow metallic waveguides [91], which also have zero group velocity at their cutoff frequencies. However, Bragg fibers require two solid materials in the cladding, which makes fabrication more challenging, while metallic waveguides become lossy at infrared frequencies. With a traditional core-clad fiber or with a holey fiber, zero group velocity can instead be attained by periodic modulation of the structure along the axial direction. For example, a fiber Bragg grating is formed by a weak modulation of the refractive index "burned" in by a photorefractive effect. Because this index modulation is typically much less than 1%, however, the low group-velocity bandwidth is small in fiber Bragg gratings. Furthermore, one can only modulate the index of a solid material, greatly reducing the grating effect for modes confined in an air core. It has been proposed that spherical particles could be introduced into an air core in order to create a periodic modula-



Figure 6-3: Doubly-degenerate Γ_6 defect modes for a triangular lattice of hexagonalshaped holes with periodicity *a* obtained by varying inscribed defect diameter of a hexagonal-shaped air core: a) D = 1.6a (fundamental-like) b) D = 3.2a c) D = 6.2ad) D = 6.76a (blue/white/red = negative/zero/positive).

tion [208], but such structures seem challenging to produce on large scales compared to drawn fibers. Previous work showed that semiconductor (silicon) index contrasts (3.5:1) could support zero-group velocity modes in fiber-like geometries [90], and here we underline the existence and utility of analogous modes with conventional fiber materials. Furthermore, our previous work demonstrated that such zero group-velocity modes can even be converted into backwards-wave and ultra-flat bands by careful tuning of the waveguide core [90], and we expect that similar phenomena should be possible in chalcogenide fibers and other lower-contrast materials.



Figure 6-4: Fraction of electric-field energy $\varepsilon |\mathbf{E}|^2$ in the hexagonal-shaped air core (as in Fig. 6-3) as a function of the core radius (radius of inscribed circle). Inset: frequency ω at $\beta = 0$ of guided mode vs. core radius.

6.4 Gaps and defect modes

One 2d photonic crystal structure that is well known to have a complete gap for sufficiently large index contrast is a triangular lattice (period a) of cylindrical air holes (radius r) in dielectric [99], similar to the geometry of many fabricated holey fibers [196]. This geometry with r = 0.45a turns out to have a 4.4% complete gap at $\beta = 0$ for a refractive index of 2.8, chosen to correspond to that of a typical chalcogenide glass (e.g., As₂Se₃ at $\lambda = 1.5 \mu m$ [16]). We also considered a slightly modified 2d photonic crystal consisting of a triangular lattice of dielectric rods in air connected by thin veins (resembling hexagonal-shaped holes) [38,242]. The gap size was optimized over two parameters, rod radius and vein thickness, yielding a 5.4% gap-to-midgap ratio for a rod radius of 0.16a and a vein thickness of 0.2a. The gap in this structure persisted for index contrasts as low as 2.6:1 (as shown in the inset of Fig. 6-2). The Maxwell eigenproblem was solved with an iterative (conjugate gradient) method in a planewave basis [107]. The resulting band diagrams, with gaps that extend over a range of nonzero β , are shown in Figs. 6-1 and 6-2. Since the modified structure of Fig. 6-2 has a slightly larger gap, we focus on this structure for the remainder of the chapter; similar results can be obtained for the cylindrical-hole



Figure 6-5: Air-core guided mode in gap of Fig. 6-2, with insets showing electric-field E_z and Poynting vector S_z (blue/white/red = negative/zero/positive).

structure.

An air core is formed by removing some dielectric material, and here we do so by a hexagonal-shaped air core with an inscribed-circle diameter D in a 15a by 15asupercell. (This supercell is large enough that, for all guided modes considered here, the guided-mode field has decayed to negligible values by the edge of the supercell and hence the finite supercell size is irrelevant.) Depending on the core diameter D, different types of modes with varying symmetry and degrees of confinement can be localized [99]. We chose D to satisfy two criteria. First, the confined mode should be of the right symmetry to be excited by an incident planewave source—technically, this means that the mode is doubly degenerate and belongs to the Γ_6 representation of the sixfold (C_{6v}) symmetry group [92] of the hexagonal core. As D is varied, we obtain a variety of different Γ_6 defect modes, as shown in Fig. 6-3. For small D, we obtain fundamental-like fields patterns as in Fig. 6-3(a), whereas for larger D we obtain more complicated field patterns that are, however, better confined in the air core as in Fig. 6-3(d). For a given mode with strong air-core confinement, we then chose Dto maximize the fraction of the electric-field energy ($\varepsilon |\mathbf{E}|^2$) in the air core at $\beta = 0$ (see Fig. 6-4) while also eliminating the influence of surface states [116,199,236]. This is desirable in air-core fiber applications to reduce absorption loss from the cladding and increase light-gas interactions. In particular, we chose the mode from Fig. 6-



Figure 6-6: Solid-core guided mode in gap of Fig. 6-2, with insets showing electric-field E_z and Poynting vector S_z (blue/white/red = negative/zero/positive).

3(d) (D = 6.76a) for specificity, and the resulting structure is shown along with its dispersion relation in Fig. 6-5. The field profile (which is TM at $\beta = 0$) is still strongly confined at a non-zero axial wavevector $(\beta a/2\pi = 0.14)$, as shown by the inset.

6.5 Topology optimization of cladding structure

Another approach to finding a cladding structure with low index contrast instead of starting with a pre-determined lattice and unit cell would be to make *every* pixel in the unit cell a free parameter. The ε of each pixel could then be constrained to lie in some interval (between say 1 and 12) and we could proceed to maximize the bandgap between any set of two adjacent bands. Cox and Dobson [51, 52] optimized the bandgap of square lattice designs of 2d photonic crystals using evolutionary algorithms while Yablonovitch et al. [111] employed level-set methods in a similar approach. Both groups used the absolute size of the gap as their objective function and only considered one type of unit cell arrangement. Maximizing a gap's absolute size is misleading since it is by definition relative to an arbitrary frequency of one where the lattice constant is unity. In practical circumstances, the relevant figure of merit is the fractional gap as this quantity is normalized by the operating frequency.

We thus choose the fractional gap as our objective function. A potential objective

function could then be,

$$\max_{\varepsilon} \left\{ 2 \frac{[\min_{\mathbf{k}} \omega_{n+1}(\mathbf{k})] - [\max_{\mathbf{k}} \omega_{n}(\mathbf{k})]}{[\min_{\mathbf{k}} \omega_{n+1}(\mathbf{k})] + [\max_{\mathbf{k}} \omega_{n}(\mathbf{k})]} \right\}.$$
(6.1)

However the problem with expressing the problem this way is that the objective function in eq. (6.1) is not differentiable at accidental degeneracies in the dispersion relation corresponding to two different wavevector points in the irreducible Brillouin zone with the same frequency. These accidental degeneracies are bound to arise as the optimization algorithm seeks to align the frequencies at two wavevectors and then proceeds to "push" them up together to maximize the gap. A non-gradient based algorithm for these types of problems is computationally intractable given that the number of iterations would be at least several times the number of dimensions (for example, in *d* dimensions without gradient information, d+1 evaluations are required just to explore the variation along each dimension in order to make the *first* update of the design variables). Our only recourse then is to use gradient-based optimization algorithms but any such method would encounter obstacles in exploring the large design space as just described. Fortunately, we can reformulate the minimax problem of eq. (6.1) as an equivalent problem with nonlinear constraints [231],

$$\max_{\varepsilon, f_1, f_2} \qquad 2\frac{f_2 - f_1}{f_2 + f_1} \\ \text{subject to} \qquad f_1 \ge \omega_n(\mathbf{k}), \ f_2 \le \omega_{n+1}(\mathbf{k}).$$
(6.2)

Equation (6.2) and its associated gradient can now be solved with suitable inequalityconstrained nonlinear algorithms. In this case we chose a recent refinement of the method of moving asymptotes (MMA) [214] implemented in a free-software package of nonlinear-optimization algorithms [101]. The key property of the MMA algorithm that makes it suitable for such problems is that it produces a feasible solution at every iteration enabling the algorithm to be terminated at any time (most methods typically satisfy constraints only when they are close to converging). This is particularly useful as we are not interested in computing the fractional bandgap to a very high degree of accuracy and can thus quickly explore the design space with many different initial designs. Another useful attribute of the MMA algorithm is that it is globally convergent (that is, it provably converges to *some* local optimum from any feasible starting point). We combined the MMA algorithm with our frequency-domain planewave expansion Maxwell solver [107] for a 2d computational cell consisting of 64 by 64 pixels (4096 degrees of freedom). The unit cell symmetry was fixed to be either square (symmetry group C4v) or triangular (C6v) since if a unique global optimum exists it must be symmetric (on intuitive grounds we would normally expect the optimal unit cell to be symmetric in order that it have the same overlapping gap in different directions of the Brillouin zone). The MMA algorithm found optimal structures having a triangular lattice at high-index contrasts but unexpectedly a square-lattice arrangement at low-index contrasts having an 8% complete 2d gap at an index contrast of 2.6:1 (see Fig. 6-7). We then used this design to formulate a simple two-parameter, non-gradient optimization involving the inner and outer radius of the ring and width of the connected rod. This design produced a 4.5% gap at even smaller index contrast of 2.4:1.

Such an approach combining the results from a nonlinear optimization having inequality constraints with low-parameter topology optimization presents intriguing possibilities for future explorations of cladding designs. The example presented here demonstrates the potential utility of these hybrid methods to find non-conventional cladding designs.

6.6 Cladding losses in hollow-core fibers

One source of loss is the material absorption in the cladding, which for bulk As₂Se₃ is about 36 dB/m at $\lambda = 1.5 \ \mu m$ [16]. For a guided mode in the hollow core, this absorption loss is suppressed by a factor of fc/v_gn , where f is the fraction of the electric-field energy in the cladding, v_g is the group velocity, and n is the cladding refractive index [99,106]. For the mode of Fig. 6-5 at $\beta a/2\pi = 0.14$, where $v_g = 0.22c$ and f = 0.19, the absorption loss of the mode is therefore 11.1 dB/m, which is sufficient for short-distance fiber devices. Lower loss could be obtained by operating



Figure 6-7: Left: the square lattice design with the largest fractional gap discovered by the nonlinear optimization algorithm where every pixel in the unit cell was a free parameter having refractive index in the range $\{1...3.4\}$. Right: a simple two-parameter (radius and width) shape optimization based on the adjacent design produces a complete 2d gap at an even lower index contrast.

at a longer wavelength such as 3 or 10 μ m, where the losses of chalcogenide are much lower while the index of refraction remains larger than 2.7 [16].

Another general strategy to counteract such cladding-based losses is to increase the diameter of the air core [106], which means that a smaller fraction of the airguided mode will reside in the cladding. Unfortunately, increasing the core diameter leads to other problems, such as increased bending losses and other inter-modal coupling [106], very similar to the tradeoffs that were faced decades ago in designing low-loss microwave transmission tubes (where the fields in the cladding gave Ohmic losses) [232], and so it is desirable to increase the core diameter as little as possible. Johnson et al. [106] previously showed the scaling losses for the case of a cylindrical omniguide fiber and here we show that for a holey photonic crystal fiber having a cladding structure made of a much more complicated triangular lattice. This is a review of the general argument, based on the scalar limit, that the cladding-related losses should (asymptotically) scale inversely with the *cube* of the core diameter, similar to the well-known result for metallic tubes [232], and we demonstrate this scaling with numerical results. All cladding-based losses in a dielectric waveguide scale with the fraction of the electric-field energy in the cladding. Material absorption losses, in fact, can be shown from perturbation theory to be exactly equal to the fraction of $\int \varepsilon |\mathbf{E}|^2$ residing in the cladding material, multiplied by a constant factor proportional to the material's extinction coefficient and divided by the mode's group velocity [106,233]. The precise analysis of surface-roughness losses is complicated [108], but nevertheless is proportional to the mean $|\mathbf{E}|^2$ at the surfaces (with a complicated proportionality factor depending on the roughness shapes and correlations [108]) and hence also scales as the fraction of the field energy in the cladding [106]. The field leakage through the finite number of crystal layers can easily be reduced exponentially by adding more layers [99, 106], so it is less of a concern, but it too scales with the squared field amplitude in the cladding [106]. Therefore it is sufficient to consider the scaling of the absorption loss, or of the fraction of field energy in the cladding, with the air-core radius R.

For any given mode, in the limit of large R the mode becomes more and more similar to a plane wave propagating along the z axis. Its dispersion relation approaches the air light line, and its penetration depth into the cladding becomes negligible compared to the scale of the transverse oscillations. This condition, of the penetration into the inhomogeneous materials becoming small compared to the transverse wavelength $(2\pi/k_t)$, was precisely the condition in which the scalar limit applies. In this limit, we can describe the mode as a linear polarization multiplied by a scalar amplitude $\psi(x, y)$ that is zero in the cladding. In reality, there is some small nonzero amplitude in the cladding, but because of the approximate zero boundary condition at r = R, the amplitude of the cladding field goes as 1/R just as we explained for te₀₁. Thus, all modes approach a $1/R^3$ scaling.

This is true in two-dimensionally periodic photonic-crystal fibers, such as the hollow-core holey structures described in this chapter. Overall, the same asymptotic $1/R^3$ scaling applies: the core interface/area ratio goes as 1/R and there is an additional $1/R^2$ factor from the cladding field amplitude in the scalar limit. However, an additional wrinkle is provided by the proliferation of surface states. Unless a crystal



Figure 6-8: Scaling of the absorption suppression factor α/α_0 versus core radius R, at mid-gap, for the fundamental mode of a hollow-core holey fiber (blue circles/lines); this factor tends to a $1/R^3$ scaling (black line, for reference). Insets show the intensity pattern (time-average Poynting flux) of the fundamental mode for two core radii, R = 0.83a and R = 12.1a. The dielectric interfaces are shown as black lines; the air core is hexagonal and terminates the crystal in such a way as to remove the possibility of surface states [236].

termination is chosen that eliminates surface states [116, 199, 236], as the core size is increased we will get more and more surface states. These surface states cross the guided band and chop up its usable bandwidth [236]. Precisely such a phenomenon was observed experimentally when Ref. 146 replaced the air core of an earlier holeyfiber experiment by Ref. 206 with one of about 2.2 times the diameter: the losses were reduced by a factor of eight (from 13 dB/km to 1.6 dB/km), but the bandwidth was reduced by a factor of five because the surface states were not eliminated. (The surface states below the light line do *not* have absorption/leakage/scattering losses that decrease with R, because they remain localized at the cladding surface regardless of R.)

To see the $1/R^3$ scaling more convincingly in a holey fiber, however, one must look at a larger range of core diameters, in a computer simulation where all other things can be kept rigorously equal (as opposed to two experiments by different groups that may or may not be comparable), and choose terminations so that the influence of surface states is eliminated. We have done so, looking at the fundamental air-guided mode of a hollow-core holey fiber in $\varepsilon = 2.1025$ silica with radius 0.47*a* air holes in a triangular lattice with period *a*, similar to Fig. 6-1. The core is a hexagonal-shaped air region carved out of the crystal with "radius" *R*, where 2*R* is the distance between two parallel sides of the hexagon. This structure, along with the intensity pattern of the fundamental mode, is shown for two sample core radii in the insets of Fig. 6-8. The core termination is chosen to roughly bisect a layer of holes, as in Ref. 236, to eliminate surface states. Then, we varied the core radius *R* (in steps of $a\sqrt{3}/2$ to preserve the termination) and computed the fraction of the fundamental mode was computed near mid-gap, at $\beta a/2\pi = 1.65$. The result is shown in Fig. 6-8, and indeed quickly approaches the expected $1/R^3$ asymptotic scaling.

6.7 Coupling to slow-light modes

Another practical challenge in all slow-light structures is coupling from a non-slow source; one very general technique is a gradual "taper" transition to a higher-velocity waveguide [102,165,177,186], for example by gradually scaling the structure [131,144] to a larger diameter to shift the band edge down to increase the group velocity at the operating ω . (Alternatively, rather than rescaling the whole structure, gradually *decreasing* the core diameter while keeping the cladding unchanged turns out to shift the band edge down in this geometry.) (Theoretically, a gradual enough transition can couple any pair of waveguides, no matter how different, with arbitrarily low reflection limited only by fabrication capabilities [102].) Minimization of reflections by proper design of couplers between very different modes of dielectric and photoniccrystal waveguides, including slow-light modes, has been studied elsewhere [102, 156, 163, 165, 186, 216], and a specific design for this fiber lies beyond the scope of this manuscript.

6.8 Solid cores

In contrast to air cores, solid (dielectric-filled) cores can be used to enhance interactions and nonlinearities with solid materials, such as for fiber lasers [136,227]. Here, we form a small solid core by filling a hexagonal-shaped core (D = 1.62a) with dielectric. This confines a doubly-degenerate mode with an extremely flat dispersion relation, in addition to its zero-group-velocity point at $\beta = 0$, as shown in Fig. 6-6. This extreme flatness could potentially be transformed into a higher-order (e.g. quartic) band edge or even a concave backward-wave band-edge, via proper tuning of the solid core geometry [90].

6.9 Concluding remarks

Any holey photonic-crystal geometry with a complete gap for both polarizations in two dimensions can be used to obtain zero–group-velocity modes in a fiber geometry our triangular lattice structure of hexagonal holes, here, is only one such example. An opportunity for future designs is to find complete gap structures with even lower index contrasts, in order that a wider range of materials become available for the fabrication of such slow-light devices. The ideal result would be a structure that has a complete 2d gap at silica/air index contrasts (1.5:1), but we are not currently aware of any geometry with this property. We have also highlighted the importance of the scalar short-wavelength limit, which has many informative consequences even at moderate wavelengths, by analyzing the scaling of the cladding losses with core radius.

Chapter 7

Robust design of slow-light tapers in periodic waveguides

7.1 Summary

This chapter describes the design of tapers for coupling power between uniform and slow-light periodic waveguides. New optimization methods are reviewed for designing robust tapers that had previously been proposed, which not only perform well under nominal conditions, but also over a given set of parameter variations. When the set of parameter variations models the inevitable variation typical in the manufacture or operation of the coupler, a robust design is one that will have a high yield, despite these parameter variations.

We review the notion of successive refinement, and robust optimization based on multi-scenario optimization with iterative sampling of uncertain parameters, using a fast method for approximately evaluating the reflection coefficient. Robust tapers designed over a range of different lengths are compared to a linear taper and to optimized tapers that do not take parameter variation into account. Finally, robust performance of the resulting designs is verified using an accurate, but much more expensive, method for evaluating the reflection coefficient.

7.2 Introduction

This chapter begins by reviewing optimization methods first presented by Mutapcic et al. [165] for designing *robust* tapers, which not only perform well under nominal conditions, but also over a given set of parameter variations. In contrast, the nominal optimum produced by straightforward optimization of this problem relies on delicate interference effects that are destroyed by any deviation from the design. Here we present results that extend our previous work in Ref. 165 (spearheaded by Almir Mutapcic) which involved optimizing taper designs at a single taper length to a range of lengths. In this work, we incorporate uncertainties in the taper shape and length, different from Mutapcic et al. which considered the taper length and operating frequency as we have now discovered that our semi-analytical computational method is inaccurate for calculating taper reflections over a set of different frequencies. We now also provide insights into the few key parameters among the many degrees of control in the optimization algorithm that strongly govern robust taper designs. Furthermore, we study the trade off between reflection and taper length by optimizing each taper length separately. This analysis will demonstrate the impact of robustness on device design more clearly. We also focus on more realistic designs having shorter taper lengths (1-20 lattice constants) amenable to fabrication whereas Mutapcic et al. [165] considered a relatively large range of taper lengths (1-100) as part of our initial work demonstrating the proof-of-concept of this approach.

The methods outlined optimize over an arbitrary variable taper rate, described by hundreds (or thousands) of degrees of freedom, in order to find a design with performance orders of magnitude better than that of a simple linear (constant-rate) taper. Accurate techniques from coupled-mode theory [102] are used to quickly explore different shapes; the results are validated against a direct numerical solution of Maxwell's equations [21,22]. Because the set of parameter variations models the inevitable variations typical in the manufacture or operation of the coupler, and is explicitly accounted for in the optimization, a robust design will have a high yield despite these parameter variations.

A standard component of optical and microwave devices is a waveguide taper which couples light from one waveguide to another by means of a gradual transition. Although a sufficiently gradual taper approaches an adiabatic limit of 100% transmission, in a practical setting the challenge is to design a taper as short as possible, or with as low a loss as possible for a given length. Perhaps the most challenging case is to design a short taper between an ordinary uniform waveguide and a periodic waveguide [62], a special case of a general class of periodic optical structures known as 'photonic crystals' [99]. Periodic waveguides are both useful and challenging for the same reason: a periodic waveguide has a 'slow-light' band edge for which the group velocity of light slows down as it approaches a certain frequency. Operating in this slow-light region is useful because it increases the interaction of light with the material, enhancing nonlinearities [209, 239], tunability [186], gain [244], and other effects. However, as the group velocity decreases, the 'impedance mismatch' between the periodic and uniform waveguide increases, and a longer taper is generally required to achieve the same coupling loss [186]. If the waveguides are simply butt-coupled without a taper, the transmission goes to zero as the zero-velocity band-edge is approached [201].

A variety of techniques have been employed to select a taper shape for coupling to periodic waveguides. Most of this previous work examines cases operating far from any band edge (so the group velocity is not small) and focuses on simple linear (constant-rate) tapers [20, 42, 83, 104, 156, 178, 185, 216, 240] or families of quadratic shapes [60, 114]. Genetic algorithms have also been employed to design couplers using arbitrarily placed scattering cylinders [82, 97]. Non-taper-based couplers from free space or parallel waveguides have also been considered [125, 187]. Although this previous work did not explicitly account for uncertainties in the model parameters, the mostly small number of design parameters combined with the moderate group velocities help avoid non-manufacturable designs. As soon as the design involves optimization over a large number of free parameters, the nominal optimum tends to be a non-robust design that relies on delicate interference effects. (A similar result was observed as a strong frequency sensitivity in genetic optimization over many degrees of freedom [81]). Previously, Povinelli et al. [186] considered a slow-light periodicwaveguide coupler with higher-degree polynomial taper shapes, and used a simple regularization technique to avoid non-robust solutions.

There are several general models of parameter uncertainty, as well as general approaches for dealing with uncertain parameters. The approach taken in this design is *worst-case robust optimization* or (minimax optimization), as first described in our previous work spearheaded by Almir Mutapcic [165]. Here the parameters are modelled as lying in some given set of possible values, but without any known distribution; a taper design is chosen, whose worst-case objective value, over the given set of possible uncertainties, is minimized. In this model, one does not rely on any knowledge of the distribution of uncertain parameters (which indeed, need not be stochastic). There is no claim that worst-case robust optimization is superior to other approaches; but it is generally found that worst-case robust designs produced by the methods of this project also perform well when analyzed under a stochastic model of parameter variation.

7.3 Nominal and robust taper design problems

This section reviews material presented in Ref. 165.

7.3.1 Taper shape and reflection magnitude

Consider a taper with length L that couples a uniform and a slow-light waveguide structure with period Λ . The taper is a quasi-periodic structure that is parametrized by the *taper shape function* $s : [0, 1] \rightarrow \mathbf{R}_+$. The argument of the taper shape function is the normalized length variable u = z/L, where z is the physical coordinate along the taper. Each value of s corresponds to an intermediate periodic structure between the taper endpoints, for example in Fig. 7-1, s could correspond to the width of the flanges, the radius of the holes, or the separation of the blocks. The varying periodic structure described by s(u) defines a taper as described in Ref. 102; essentially, the taper matches the cross-section of the periodic structure s(u) at z = Lu. The taper



Figure 7-1: Various tapers between uniform and periodic dielectric waveguides. (a) Periodic sequence of holes, where taper varies the radius and period of the holes, in 2d or 3d. (b) Periodic set of flanges, where taper varies the width of the flange, in 2d or 3d. (c) Periodic sequence of dielectric blocks, where taper varies the period Λ between the blocks. All three of these tapers, in 2d or 3d, can be efficiently optimized by the robust coupled-mode method, but this chapter focuses on (c) because it is also amenable to brute-force computation for verification purposes.

shape function is constrained at its starting and its final point, with s(0) = 0 denoting the starting uniform structure, and s(1) = 1 denoting the final periodic structure. Figure 7-2 illustrates a sample taper and its shape function, where in this case s is simply proportional to the continuously varying width of the flanges.

Given a taper shape function, one can evaluate the magnitude of the reflection from an incoming light wave coupled from the uniform into the slow-light waveguide, for example by numerical simulation of the wave equation. The reflection magnitude is denoted R; it depends on the taper shape function s, as well as various parameters such as the refractive index (which might, indeed, vary spatially), the wavelength, and so on. These parameters are denoted by a vector $v \in \mathbf{R}^m$; to emphasize that R is a function (or, since s is a function, a functional) of the taper shape s and the parameter vector v, it will sometimes be written as R(s, v).

Let v_{nom} be the nominal value of v, i.e. a typical (or expected) value of the



Figure 7-2: Top. A taper coupling uniform and slow-light waveguide structures. Bottom. Its taper shape function s.

parameter vector. The nominal reflection magnitude is defined as

$$R_{nom}(s) = R(s, v_{nom}). \tag{7.1}$$

The nominal reflection magnitude is a functional of the taper shape function s, and gives the magnitude of the reflection when the parameter vector is equal to its nominal value.

7.3.2 Parameter uncertainty and worst-case reflection magnitude

Parameter uncertainty, which can be caused by manufacturing imperfections, wavelength variation, model parameter errors, etc., is modelled by a set $\mathcal{V} \subseteq \mathbf{R}^m$. The set \mathcal{V} can be thought of as the set of possible values of the parameter vector. It will be assumed that $v_{nom} \in \mathcal{V}$. As a simple (but important) example, \mathcal{V} can be a finite set $\mathcal{V} = v_1, \ldots, v_K$. In this case the index *i* is referred to as a *scenario*, with associated parameter vector v_i . As another common example, \mathcal{V} can be a box in \mathbb{R}^m , for example, centered at the nominal parameter value,

$$\mathcal{V} = v||v_i - v_{nom,i}| \le \xi_i, i = 1, \dots, m, \tag{7.2}$$

where ξ_i gives the radius or half-range of the variation in parameter *i*. (This type of parameter variation can be described as $v_i = v_{nom,i} \pm \xi_i$.)

The performance of a taper design, in the presence of parameter uncertainty, is judged by the worst-case (largest possible) reflection magnitude over all possible $v \in \mathcal{V}$. The worst-case reflection magnitude is defined as

$$R_{wc}(s) = \sup_{\upsilon \in \mathcal{V}} R(s, \upsilon).$$
(7.3)

The worst-case magnitude reflection R_{wc} is a functional of the taper function s. It is always the case that $R_{wc}(s) \ge R_{nom}(s)$ for any s; indeed the ratio $R_{wc}(s)/R_{nom}(s)$ gives a measure of (worst-case) performance degradation of the taper, due to parameter variation.

For a scenario model of parameter uncertainty, i.e. when $\mathcal{V} = v_1, \ldots, v_K$, the worst-case reflection magnitude has the form

$$R_{wc}(s) = \max_{i=1,\dots,K} R(s, v_i),$$
(7.4)

the maximum reflection magnitude over the K scenarios. But in most cases, $R_{wc}(s)$ cannot be computed exactly, since this involves solving a non-convex optimization problem. It can be approximately computed, however, using several methods below.

7.3.3 Nominal and robust taper shape problems

In the nominal taper shape problem, a taper shape function s is found that minimizes the nominal reflection magnitude R_{nom} , subject to some constraints:

minimize
$$R_{nom}(s)$$

subject to $s(0) = 0, s(1) = 1$ (7.5)
 $0 \le s(u) \le S^{max}, |s'(u)| \le D^{max} \text{ for } 0 \le u \le 1.$

The optimization variable is the taper shape function $s : [0, 1] \to \mathbf{R}_+$ The problem parameters are the maximum allowed shape value S^{max} , the maximum allowed taper slope D^{max} , and of course, the objective function R_{nom} . A solution of this problem is called a *nominal optimal taper*.

In the robust taper shape problem, the goal is to find a taper shape function s that minimizes the worst-case reflection magnitude R_{wc} , subject to some constraints:

minimize
$$R_{wc}(s)$$

subject to $s(0) = 0, s(1) = 1$ (7.6)
 $0 \le s(u) \le S^{max}, |s'(u)| \le D^{max} \text{ for } 0 \le u \le 1.$

A solution of this problem is called a *robust optimal taper*. The main goal of this chapter is to present a tractable way to (approximately) solve the robust taper shape problem eq. (7.6).

Both the nominal and robust taper shape problems eqs. (7.5) and (7.6) are infinitedimensional optimization problems, since the optimization variable is a function [5], and they include semi-infinite constraints [86], i.e. an infinite set of constraints indexed by a continuous variable (u). Both of these issues will be (approximately) handled by searching over a finite-dimensional set of shape functions, for which the semi-infinite constraints can be expressed in a simple way. The complexity of the algorithm grows linearly with the dimension of the finite-dimensional parametrization, and easily scales to dimensions large enough (e.g. thousands) that errors due to the finite-dimensional parametrization are negligible.

A more fundamental issue is that the problems eqs. (7.5) and (7.6) are not



Figure 7-3: A piecewise-linear taper shape s with n = 4, with grid points u_1, \ldots, u_4 . The taper shape satisfies s(0) = 0, $s(u_1) = x_1, \ldots, s(u_4) = x_4$, and s(1) = 1.

convex (since the objectives are, in general, not convex), which makes it unlikely that the global solutions can be found efficiently. So one must settle for locally optimal solutions of the problems, which need not be globally optimal. In Ch. 7.3.5, a successive refinement approach is described, which appears to be quite resistant to getting trapped in poor local minima.

7.3.4 Piecewise-linear taper shape parametrization

It is assumed that the taper shape functions are piecewise-linear, parametrized by their values at *n* fixed *control* or *break* or *grid* points u_1, \ldots, u_n with $0 < u_1 < \ldots < u_n < 1$, and passing through the endpoints s(0) = 0 and s(1) = 1. This is illustrated for n = 4 grid points in Fig. 7-3. This can be expressed as

$$s(u) = x_1 f_1(u) + \ldots + x_n f_n(u) + f_{n+1}(u),$$
(7.7)

with

$$f_{i}(u) = \begin{cases} (u - u_{i-1})/(u_{i} - u_{i-1}) & u_{i-1} \leq u \leq u_{i}, \\ (u_{i+1} - u)/(u_{i+1} - u_{i}) & u_{i} \leq u \leq u_{i+1}, \\ 0 & \text{otherwise}, \end{cases}$$

where $u_0 = 0$ and $u_{n+1} = 1$, and

$$f_{n+1}(u) = \begin{cases} (u-1)/(u_n-1) & u_n \le u \le 1, \\ 0 & \text{otherwise.} \end{cases}$$

The vector $x \in \mathbf{R}^n$ is referred to as the *taper shape vector*. Evidently $s(u_i) = x_i$.

With this parametrization, the endpoint constraints s(0) = 0 and s(1) = 1 hold automatically, for any shape vector x. Moreover, the semi-infinite constraints

$$0 \le s(u) \le S^{max}, |s'(u)| \le D^{max} \text{ for } 0 \le u \le 1,$$

hold if and only if

$$0 \le x_i \le S^{max}, \qquad i = 1, \dots, n,$$

$$|x_{i+1} - x_i| \le D^{max}(u_{i+1} - u_i), \quad i = 1, \dots, n - 1,$$

$$|x_1| \le D^{max}u_1, \qquad |1 - x_n| \le D^{max}(1 - u_n).$$
(7.8)

These are a set of 4n linear inequalities on the shape vector x. The notation $x \in S$ will be used to denote this, where S is the (polyhedral) set of x for which eq. (7.8) holds.

With some abuse of the notation, $R_{nom}(x)$ and $R_{wc}(x)$ will be used to denote the values of $R_{nom}(s)$ and $R_{wc}(s)$, for the shape function s associated with the shape vector x. With piecewise-linear parametrization of taper shapes, the nominal taper design problem can be expressed as

minimize
$$R_{nom}(x)$$

subject to $x \in \mathcal{S}$, (7.9)

and the robust taper design problem as

minimize
$$R_{wc}(x)$$

subject to $x \in \mathcal{S}$. (7.10)

These are finite-dimensional optimization problems, with optimization variable $x \in \mathbf{R}^n$, and 4n linear inequality constraints.

7.3.5 Successive refinement

The taper design problems described in eqs. (7.9) and (7.10) are non-convex and local methods which can (and do) get stuck in poor locally optimal points. A common method to fix this problem is to run the algorithm multiple times, starting with different initial taper designs, picking the best design obtained among the runs of the algorithm. In previous work lead by Mutapcic et al. [165], a method called *successive refinement* was presented, however, which seems to avoid the problem of getting caught in poor local minima, and eliminates the need for multiple runs from different starting points.

In successive refinement, a sequence of design problems with successively finer piecewise-linear taper shape functions is solved, in each case starting from the previous design. One starts with a single grid point, i.e. n = 1, and runs a global search of the optimal robust taper, which is tractable only for this single dimensional problem. One then adds two more grid points, in between 0 and the first grid point, and the first grid point and 1, so that n = 3, and runs the robust taper shape (RTS) algorithm (as described in [165]), starting from the previous design. This is repeated until some maximum value of n is reached. This is illustrated in Fig. 7-4.

In numerical experiments we started with the initial grid point at 1/2, and in each successive refinement step, new grid points are added halfway in-between the old ones (and 0 and 1). At the Mth refinement step there will be $n = 2^M - 1$ grid points, with values

$$u_i^M = i2^{-M}, \ i = 1, \dots, 2^M - 1.$$
 (7.11)

This approach is related in spirit to the multigrid methods [29], where the latter uses both successive refinements and coarsenings in order to speed up convergence of a linear solver rather than to avoid local minima. Successive refinement have been successfully applied in circuit design [34], in motion estimation for video coding [45], etc.



Figure 7-4: Top left. Linear taper with a single grid point. Top right. Full search performed to obtain a global optimum taper with a single grid point. Bottom left. Two new grid points added and taper values interpolated at $u_1^{(2)}$ and $u_3^{(2)}$. Optimization algorithm is run starting from this taper. Bottom right. New local optimum.

7.4 Computation of reflection magnitude

This section reviews material presented in Ref. 165. To optimize the taper shape function s, one needs a rapid method to compute the reflected power fraction R and its gradient for light incident on a particular taper structure. This project employs two such methods, described below: a fast approximate method for the optimization (including the computation of the gradient), and a slower brute-force method for verification of the final design.

7.4.1 Coupled-mode theory

In general, computing the reflection from an arbitrary structure could require an expensive solution of the complete Maxwell equations, evaluated to high accuracy in order to distinguish the tiny reflected field in a well-designed gradual taper. In the present case, however, the fact that the structure is *nearly* periodic (slowly-varying) and the reflection is consequently small, can be exploited to utilize a fast semi-analytical method based on *coupled-mode theory*.

Coupled-mode theory, also known as *coupled-wave theory* or the *slowly-varying envelope approximate* (SVEA), involves an expansion of the electromagnetic field along the waveguide taper in terms of the eigenmodes (indexed by k) of a uniform periodic waveguide matching the cross-section at each point. The expansion coefficients c_k in this basis are then determined by a set of ordinary differential equations for dc_k/dz along the taper direction (z), where the different modes are coupled by terms proportional to the rate of change of the structure. Because the structure is slowly varying, the expansion coefficients approach an 'adiabatic' limit in which the c_k are nearly constant. In this limit, the equations can be integrated approximately, to first-order in the taper rate, to yield a simple integral for the reflection coefficient. (Reflection dominates the loss in slow-light tapers.)

The most common form of coupled-mode theory was developed for nearly uniform waveguides [148] but has recently been generalized to strongly periodic waveguides of the type considered in this chapter [102]. The results of a simple first-order calculation were found to be nearly exact as long as the reflections were under 10%, making them ideal for the present case where the taper designs all have reflections well under 1%.

In particular, coupled-mode theory of a taper shape s(u) with length L leads to a first-order reflection amplitude c_r , where the fraction of reflected power is $R = |c_r|^2$, given by an integral of the form:

$$c\{s(u)\} = \int_0^1 du \frac{ds}{du} \sum_k \frac{M_k[s(u)]}{\Delta\beta_k[s(u)]} e^{iL\int_0^u \Delta\beta_k[s(u')]du'}.$$
(7.12)

Here, M_k and $\Delta\beta_k$ are given functions of the taper parametrization s. That is, each s denotes a given intermediate periodic structure, $M_k(s)$ is a complex-valued coupling coefficient determined from the eigenfields of that structure, and $\Delta\beta_k$ is a real phase-mismatch factor. The summation must in principle run over all integers k, but in practice only a handful of terms are required because the contributions decrease rapidly with k. (In particular, $\Delta\beta_k(s) = \Delta\beta(s) + 2\pi k/\Lambda(s)$, where $\Lambda(s)$ is the variable period along the taper.)

The derivation of these coupled-mode equations is rather complicated and will not

be reproduced here ¹. The key point, however, is that the full Maxwell equations need only be solved once: a set of small calculations for the eigenmodes of the periodic structures at each s, by a spectral method [107], yields the functions $M_k(s)$ and $\Delta\beta(s)$. One can then re-use these functions to compute the reflection for any taper shape s(u) and any length L by a single integral, which allows quick exploration and optimization over many different shapes.

The equations are the same regardless of the dimensionality of the problem, and have previously been used by Povinelli et al. to compute taper reflections and perform simple optimizations in large three-dimensional structures where direct simulation was not possible [186].

7.4.2 Coupled-mode reflection gradient

To carry out taper shape optimization we need to evaluate not only the reflection magnitude R but also its functional derivative (gradient) $\partial R/\partial s$. In general, such gradients can be computed by an adjoint method [32], but in this case the problem is simple enough that one can derive the same thing without resorting to such cumbersome techniques.

In particular, since $R = |c_r|^2$ and c_r is a summation over k, it suffices to compute the gradient of each k term in the summation equation for c_r above. Dropping the ksubscript for simplicity, each k term corresponds to the functional:

$$c\{s(u)\} = \int_0^1 dus'(u)F[s(u)]e^{\int_0^u f[s(u')]du'},$$
(7.13)

where $F(s) = M_k(s)/\Delta\beta_k(s)$ and $f(s) = iL\Delta\beta_k(s)$. The gradient g(u) of this functional is defined by the first-order change of $c\{s(u)\}$ under a small change $\delta s(u)$ (where $\delta s(0) = \delta s(1) = 0$ to preserve the boundary conditions):

$$\delta c = c\{s + \delta s\} - c\{s\} = \int_0^1 g(u)\delta s(u)du.$$
(7.14)

¹The original derivation [102] did not include an explicit shape function s(u). However, it was noted that the coupling matrix elements were simply proportional to the taper rate, and this is what allows us to pull out the taper-rate dependence as an s'(u) term in the integral

The explicit gradient g can be derived by substituting $s + \delta s$ into c, dropping terms higher than first-order in δs , and integrating by parts to eliminate the $\delta s'$ term. After some algebra, one obtains:

$$g(u) = -F[s(u)]f[s(u)]e^{\int_0^u f[s(u')]du'} + f'[s(u)]\int_u^1 d\tilde{u}s'(\tilde{u})F[s(\tilde{u})]e^{\int_0^{\tilde{u}} f[s(\tilde{u})]du'}, \quad (7.15)$$

which is a single integral in terms of s(u) and the known functions F and fand their derivatives, which means that the gradient can be evaluated with roughly the same cost as evaluating $c\{s(u)\}$ (similar to what one would expect for adjoint methods).

In practice, of course, infinitely many degrees of freedom are not present in s(u). As explained in Ch. 7.3.4, a piecewise-linear parametrization $s(u) = \sum_i x_i f_i(u)$, for 'tent' functions $f_i(u)$ and parameters x_i , is employed. One therefore needs only the finite-dimensional gradient with respect to the x_i :

$$\frac{\partial c}{\partial x_i} = \int_0^1 g(u) f_i(u) du. \tag{7.16}$$

The gradient of the reflection R is then found by first summing $\partial c/\partial x_i$ over k to obtain $\partial c_r/\partial x_i$, and then $\partial R/\partial x_i$ is the real part of $2c_r^*\partial c_r/\partial x_i$.

7.4.3 Brute-force verification

Because coupled-mode theory involves some approximations, it is also desirable to directly solve the Maxwell equations, with no assumptions, in order to verify the correctness of the solutions. Such a direct solution allows one to consider the effect of imperfections that violate the slow-taper assumption underlying coupled-mode theory; in particular, one can include rapid small variations in the structure corresponding to fabrication imperfections (e.g., surface roughness). We use two different computational methods to validate the coupled-mode theory: the first is an eigenmode-expansion, or transfer-matrix, method that is implemented in a free software package called CAMFR [21, 22], the second is our Meep software package for FDTD simulations described in Chapter 4.

CAMFR works by expanding the fields at every z in terms of the eigenmodes of that cross-section, with perfectly-matched layer (PML) absorbing boundaries in the lateral directions [17]. In this sense, it is related to the classic coupled-mode method mentioned above [148]. Unlike the first-order integration above, however, CAMFR makes no assumption of small scattering or slow variation, and computes a complete transfer matrix at each point where the cross-section changes that couples all the modes according to the continuity conditions on the electromagnetic field. In this sense, it is a 'brute-force' method: it solves the complete Maxwell equations with no assumptions, to an arbitrary accuracy given enough computational time and memory (i.e. a large enough eigenmode basis).

Moreover, CAMFR imposes the incident-wave boundary conditions (at z = 0 and z = L) analytically, thanks to its eigenmode basis, and hence can distinguish even a tiny reflection coefficient with high accuracy. It is most effective, however, when the two ends of the simulation are terminated by semi-infinite uniform waveguide, and so the CAMFR simulations are performed using a double taper, which tapers from uniform to periodic, then five periods in the periodic structure, and then tapers back from periodic to uniform.

7.4.4 Worst-case reflection magnitude

The problem of finding the worst, or at least a bad, value of the parameter $v \in \mathcal{V}$, for a given taper shape s, is called *pessimizing*, since the goal is to find the least favorable value of the parameter for the given shape. When \mathcal{V} is finite, exact pessimizing can be carried out by evaluating R under each scenario and taking the largest value found.

When \mathcal{V} is infinite it is difficult to compute the exact value of the worst-case reflection magnitude

$$R_{wc}(s) = \sup_{v \in \mathcal{V}} R(s, v)$$
along with a (worst-case) parameter v^* that achieves this supremum, since in general R(s, v) is not concave in v (and \mathcal{V} need not be convex). Options for pessimizing include direct search methods [119, 168, 238], or any standard local optimization method such as sequential quadratic programming methods [76, 171]. With any of these methods, the algorithm is run from a number of starting points in \mathcal{V} ; the largest value of R found is then an estimate of R_{wc} .

When \mathcal{V} is a box (1), one can easily guess a value of v that often leads to large (if not largest) R. The gradient of R with respect to v is evaluated at v_{nom} ; the approximate pessimizer is then

$$\upsilon_i^* \approx \begin{cases} \upsilon_{nom,i} + \xi_i, & \partial R / \partial \upsilon_i > 0\\ \upsilon_{nom,i} - \xi_i, & \partial R / \partial \upsilon_i < 0 \end{cases}$$

(This is the maximizer of the first-order approximation of R over \mathcal{V} .) This point can, of course, be used as the starting point for a local optimization method.

7.5 Numerical results

In this section some numerical results for a particular structure are presented. The details of the optimization algorithms are discussed in our previous publication elsewhere [165].

7.5.1 Taper geometry and uncertainty model

The two dimensional taper depicted in Fig. 7-1 (c), similar to the one considered in Ref. 102 and identical to the structure considered by Mutapcic et al. [165], is optimized, in order to have a structure where the brute-force CAMFR method is efficient (and thus can be used to validate the coupled-mode theory for a large number of values of the parameters). The periodic structure is a sequence of dielectric blocks with period Λ_0 , size $0.4\Lambda_0 \times 0.4\Lambda_0$, and dielectric constant $\varepsilon = 12$. The blocks are surrounded by air ($\varepsilon = 1$). The electric field is polarized perpendicular to the 2d plane ('TM' polarization). As described in Ref. 102, this structure supports true localized guided modes by the mechanism of index-guiding [65], and has a zero–groupvelocity band edge at a frequency of $\omega \Lambda_0/2\pi = 0.2434$. The operating frequency is $\omega \Lambda_0/2\pi c = 0.23$, ² which is slightly below the band edge, where the group velocity is under c/4 and the waveguide is single-mode at every point along the taper.

A uniform waveguide width of $0.4\Lambda_0$, which can be treated as a sequence of (touching) blocks with period $0.4\Lambda_0$, is tapered to the periodic structure by gradually spreading the blocks apart. That is, their period varies as $\Lambda(s) = \Lambda_0[s + 0.4(1 - s)]$, so that s = 0 corresponds to the uniform structure with pitch $0.4\Lambda_0$ and s = 1 corresponds to the periodic structure with pitch Λ_0 . The problem is then to determine the function s describing how fast the period (pitch) varies along the taper. The taper will be optimized individually at each length from 1 to 100 with maximum shape value $S^{max} = 1$, and maximum slope $D^{max} = 5$.

On physical grounds, one expects the optimal taper to be more rapid at the u = 0 corresponding to the uniform waveguide where the group velocity is larger, and to be more gradual at the u = 1 end corresponding to the periodic waveguide where the group velocity is low (and thus the structure is more sensitive to small changes [186]). This is precisely what is found, below, although the exact taper rate is difficult to predict *a priori*.

The following parameter uncertainty model is used. The taper length varies 1% around its nominal value (Mutapcic et al. [165] had previously considered a 1% variation of the operating frequency while using the same coupling coefficients in the coupled-mode theory for all frequencies which we have determined to be inaccurate and thus avoided); variation in the taper shape function is bounded at each grid point by ± 0.001 around the current value, with the perturbed shape within the bounds 0 and 1. The shape variation is meant to model, for example, manufacturing variation.

²It is convenient to use dimensionless frequency units of $2\pi c/\Lambda_0$, where c is the speed of light in vacuum, due to the scale invariance of Maxwell's equations [99].



Figure 7-5: Comparison of coupled-mode theory and brute force verification method (CAMFR [21, 22]) for a linear taper from length of 1 taper period through 100. The excellent agreement between the fast coupled-mode theory semi-analytical solver (blue circles) and the much slower brute-force method (red squares) to compute the objective function permits use of the former to quickly explore a large parameter space in the robust optimization.

7.5.2 Pessimizing method

The following method is used to carry out approximate worst-case analysis. At each of 11 values of the taper length, uniformly spaced over the interval of 1% around its nominal value, the approximate worst-case shape perturbation at the current point s is found using the derivative heuristic as described in Ch. 7.4.4, i.e.

$$s^{*}(u) = \begin{cases} \min\{s(u) + 0.001, 1\}, & \partial R/\partial s(u) > 0\\ \max\{s(u) - 0.001, 0\}, & \partial R/\partial s(u) < 0 \end{cases}$$

(The worst-case shape perturbation depends on the taper length L.) The reflection magnitude is evaluated for each taper length L with its associated approximate worst-case taper shape. The result is the approximate worst-case reflection over the shape uncertainty; maximizing over the 11 values of the taper length yields the approximation of R_{wc} .



Figure 7-6: Comparison of brute force computation (solid circles) and coupled-mode theory (hollow circles) of reflections from nominal taper designs optimized for each taper length. The performance of the nominal taper is clearly ruined by the slight pixellization effects introduced by the brute-force solver.

We cannot claim that this pessimization heuristic gives the true worst-case value. However it has been tested extensively, by attempting to find worse parameter values using other methods, such as derivative-free optimization, SQP, and simply sampling random parameter values in \mathcal{V} . In no case was a significantly worse value of the parameter found.

7.5.3 Optimization

Tapers were found using the nominal taper shape (NTS) and robust taper shape (RTS) algorithms as described in Ref. 165, with the following parameters: initial $\rho = 0.1S^{max}$, $\rho^{min} = 0.001S^{max}$, $\rho^{max} = 0.5S^{max}$, $\alpha^{decr} = 0.75$, $\alpha^{incr} = 1.25$, and $N^{max} = 150$, terminating also if no improvement is made. 10 iterations of successive refinement are used, with evenly-spaced grid points, which results in a final taper with n = 1023 grid points. Global optimization is carried out during the first step of the successive refinement, after which the obtained shape is used to construct initial



Figure 7-7: Brute force computation of reflections from linear (green), nominal (blue) and robust (red) taper designs for each length. The superior performance of the robust tapers, showing an exponential decrease of the reflection at shorter taper lengths before reaching a noise floor, is evident under the slight perturbation introduced by the brute-force solver's pixellization.

points for the subsequent steps.

The NTS algorithm, and the multi-scenario taper shape (MSTS) algorithms carried out in each iteration of the RTS algorithm, usually terminate in 50-70 steps, due to no improvement in objective value. The RTS algorithm converged around 30-40 steps (each of which consisted of an approximate worst-case analysis and a multi-scenario optimization). For the highest level of refinement, the RTS algorithm required a total of around 2000 basic iterations (each requiring an approximate worstcase analysis, the solution of an LP, etc.).

The algorithms were implemented in Matlab, solving the update step subproblems using CVX [77], which calls the SeDuMi solver [211]. The subproblem calculation for the NTS algorithm with n = 1023 variables (the last step in successive refinement) takes about a second, while the subproblem calculation for the MSTS algorithm with n = 1023 and K = 50 scenarios takes about ten seconds (on a personal computer). Solving the NTS problem required a total of around 40 seconds, and solving the RTS problem required a total of around 20 minutes. Had algorithms been implemented in C, using a custom LP solver for the particular structure that arises in these problems, these times would likely have been far smaller, by a factor exceeding 10.

7.5.4 Results

Figure 7-5 shows the reflections from a linear taper structure spanning taper lengths 1 to 100 computed using two different methods: the slow, exact (apart from discretizations), brute-force CAMFR method and the much faster, approximate, semi-analytical coupled-mode theory. (In order to minimize numerical errors from the boundaries in the simulations the double taper setup as described in [102] Ch. 7.4 is used.) The excellent agreement between the two methods confirms the validity of the coupled-mode theory which, given its speed, is used in the optimization algorithms. Coupled-mode theory's ability to quickly and accurately evaluate the taper reflection objective function thus permits exploration of a large design parameter space.

The nominal taper designs, on the other hand, show very different performance as illustrated in Fig. 7-6. The delicate interference effects give rise to taper designs that have low reflections using coupled-mode theory, but in the presence of slight perturbations (e.g. the pixellization effects of the brute-force solver) their performance is degraded by as much as four orders of magnitude at certain taper lengths.

The robust taper designs outperform the linear and nominal taper designs at all taper lengths as demonstrated in Fig. 7-7. The optimal robust design has a reflection coefficient that is nearly two orders of magnitude lower than the nominal design, and almost three orders of magnitude lower than the simple linear design under the slight pixellization effects introduced by the discrete brute-force solver. Among all the parameters used in the robust optimization, the taper-slope trust region was found to be dominant for producing smooth taper shapes without strange features. This was true regardless of the inclusion of taper slope constraints and modifications to the number of iterations of the multi-scenario taper shape algorithm (N^{max}) and robust taper algorithm.

An investigation of the taper profiles produced by the nominal and robust taper



Figure 7-8: Taper profiles of linear (green), nominal (blue) and robust (red) designs for taper length of 20. The slow-light, periodic waveguide structure is at u = 0 on the left and the standard, strip waveguide is on the right of the axis. Note the delicate features of the nominal taper which arise from sensitive interference effects. The robust taper profile varies more gradually and has superior performance under the slight pixellizations effects of the brute-force solver.

shape optimization algorithms provides insight into their performance. Figure 7-8 shows the linear, nominal and robust taper profiles designed at a taper length of 20. The nominal taper design has very fine, non-robust features that arise from delicate interference effects required to produce small reflections. The robust taper design varies more gradually and thus should be expected to tolerate slight perturbations as its performance confirms.

7.6 Conclusions

In this chapter, an approach to non-convex robust optimization has been reviewed that was first presented by Mutapcic et al. [165], which is applied to the challenging problem of designing robust taper transitions to 'slow-light' periodic waveguides. The robust optimization algorithm is based on multi-scenario optimization with iterative sampling of uncertain parameters, and uses fast and accurate coupled-mode computations in order to quickly explore different taper designs. The approach also uses the idea of successive refinement in order to avoid poor locally optimal points and to improve design robustness to taper shape uncertainty.

Numerical experiments verify that the obtained robust tapers perform well under the slight pixellization effects introduced by the brute-force solver, while the optimized tapers that do not take parameter variation into account perform quite poorly. An inspection of the taper shape profiles provides clear insight into their performance as only smooth designs lacking fine features are insensitive to slight perturbations. The parameters having the dominant effect on producing suitable robust taper shapes have also been identified.

In the future, we aim to further extend the techniques reviewed in this chapter to more complicated 3d taper structures where brute-force calculation of the objective function would be intractable. Our hope is to design robust tapers for complicated geometries that can ultimately be fabricated, tested and deployed in real-world experiments.

Chapter 8

Conclusion

This thesis has explored a number of areas relevant in the computation and design of nanophotonics. The first half was devoted to resolving key issues limiting the use of large, bulky simulations in computational electromagnetism for device optimization. We developed a novel subpixel smoothing algorithm, for both isotropic and anisotropic media, based on rigorous analytical arguments and demonstrated its superior performance over that of previously published results. Our method restores the quadratic accuracy of the underlying second-order finite differences and thereby permits simulations with high accuracy at low resolutions and thus computational cost. Next we demonstrated the irrecoverable failure of perfectly matched layer (PML) absorbing boundaries for a number of important problems and proposed a workaround involving adiabatic absorbers. We demonstrated the fundamental connection between reflections from any type of absorber, PML or non-PML, by making a link with the smoothness of the absorption profile. From these fundamental properties of PMLs, we then introduced a simple method to verify the correctness of any PML formulation and demonstrated its utility by proposing and validating a uniaxial-PML (UPML) formulation for anisotropic media in FDTD simulations. We then incorporated these and many other enhancements into a flexible, free-software package for electromagnetism, Meep, that is becoming increasingly popular in the optics community. Meep is an implementation of the FDTD method that has a rich set of features and functionalities and continues to evolve. The second half of this thesis was devoted to the design of nanophotonic devices. We used topology optimization to design cladding structures for holey-photonic crystal fibers having a low-index contrast which now permits a new class of materials known as chalcogenide glasses to be used to fabricate such fibers. Finally, we developed fast, semi-analytical tools to design adiabatic tapers for coupling to slow-light modes of a photonic-crystal waveguide. The tapers were designed with careful performance metrics to withstand small manufacturing and operational perturbations. These robust tapers were shown to be a significant improvement over nominal tapers that were not designed with these principals in mind. Several enhancements relating to PDE-constrained optimization were introduced to overcome technical challenges in the design of such taper structures. The tools and design methodologies described in this thesis can be readily extended to other problems in electromagnetism. It is our hope that the work presented here has opened new lines of inquiry and will continue to be developed and improved upon by future generations of researchers.

Bibliography

- Harold Abelson and Gerald Jay Sussman. Structure and Interpretation of Computer Programs. MIT Press, Cambridge, MA, 1985.
- [2] Govind P. Agrawal. Nonlinear Fiber Optics. Academic Press, San Diego, 3rd edition, 2001.
- [3] A. Ahland, D. Schulz, and E. Voges. Accurate mesh truncation for schrodinger equations by a perfectly matched layer absorber: Application to the calculation of optical spectra. *Phys. Rev. B*, 60(8):R5109, 1999.
- [4] J. Bach Andersen and V.V. Solodukhov. Field behavior near a dielectric wedge. IEEE Trans. Anten. Prop., 26(4):598–602, 1978.
- [5] E.J. Anderson and P. Nash. *Linear programming in infinite-dimensional spaces*. New York: John Wiley, 1987.
- [6] J. Anderson, M. Okoniewski, and S. S. Stuchly. Practical 3-d contour/staircase treatment of metals in fdtd. *IEEE Microwave Guided Wave Lett.*, 6(3):146–148, 1996.
- [7] A. Argyros, T. Birks, S. Leon-Saval, C. M. Cordeiro, F. Luan, and P. S. J. Russell. Photonic bandgap with an index step of one percent. *Opt. Express*, 13(1):309–314, 2005.
- [8] Joseph E. Avron and Alexander Elgart. Adiabatic theorem without a gap condition. *Commun. Math. Phys.*, 203:445–463, 1999.
- [9] T. Baba. Slow light in photonic crystals. *Nature Photonics*, 2:465–473, 2008.
- [10] David H. Bailey and Paul N. Swartztrauber. A fast method for the numerical evaluation of continuous fourier and laplace transforms. SIAM J. Sci. Comput., 15(5):1105–1110, 1994.
- [11] Jr. Baker, George A. and Peter Graves-Morris. Padé Approximants. Cambridge University Press, 2nd edition, 1996.
- [12] A. Bamberger and A. S. Bonnet. Mathematical analysis of the guided modes of an optical fiber. J. Math. Anal., 21(6):1487–1510, 1990.

- [13] R. Barrett, M. Berry, T.F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the Solution* of Linear Systems: Building Blocks for Iterative Methods. SIAM, Philadelphia, PA, 1994.
- [14] U. Basu and A.K. Chopra. Perfectly matched layers for time-harmonic elastodynamics of unbounded domains: theory and finite-element implementation. *Compt. Meth. in Appl. Mech. and Engr.*, 192(11–12):1337–1375, 2003.
- [15] F. Benabid, J. C. Knight, G. Antonopoulos, and P. St. J. Russell. Stimulated raman scattering in hydrogen-filled hollow-core photonic crystal fiber. *Science*, 298:399–402, 2002.
- [16] G. Benoit and Y. Fink. Mit electronic database of optical properties. http: //mit-pbg.mit.edu/Pages/DataBase.html, 2006.
- [17] Jean-Paul Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. J. Comput. Phys., 114(1):185–200, 1994.
- [18] Marsha J. Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. J. Comp. Phys., 53:484–512, 1984.
- [19] P. Bermel, E. Lidorikis, Y. Fink, and J.D. Joannopoulos. Active materials embedded in photonic crystals and coupled to electromagnetic radiation. *Phys. Rev. B*, 73(165125), 2006.
- [20] P. Bienstman, S. Assefa, S.G. Johnson, J.D. Joannopoulos, G.S. Petrich, and L.A. Kolodziejski. Taper structures for coupling into photonic crystal slab waveguides. J. Opt. Soc. Am. B, 20(9):1817–1821, 2003.
- [21] Peter Bienstman. CAMFR: CAvity Modeling FRamework. Software at http://camfr.sf.net.
- [22] Peter Bienstman and R. Baets. Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers. Optical and Quantum Electron., 33(4–5):327–341, 2001.
- [23] T. A. Birks, J. C. Knight, and P. St.J. Russell. Endlessly single-mode photonic crystal fiber. Opt. Lett., 22(13):961–963, 1997.
- [24] N. Bloembergen. Nonlinear Optics. W. A. Benjamin, New York, 1965.
- [25] Marc Bonnet. Boundary Integral Equation Methods for Solids and Fluids. Wiley, 1999.
- [26] J.P. Boyd. Chebyshev and Fourier Spectral Methods. Springer, second edition, 1989.
- [27] John P. Boyd. The optimization of convergence for Chebyshev polynomial methods in an unbounded domain. J. Comp. Phys., 45:43–79, 1982.

- [28] Robert W. Boyd. Nonlinear Optics. Academic Press, London, UK, 1992.
- [29] W. Briggs, V. Henson, and S. McCormick. A multigrid tutorial. Philadelphia, PA: SIAM, 2nd edition, 2000.
- [30] Laurent Brilland, Fréderic Smektala, Gilles Renversez, Thierry Chartier, Johan Troles, Thanh Nguyen, Nicholas Traynor, and Achille Monteville. Fabrication of complex structures of holey fibers in chalcogenide glass. *Opt. Express*, 14(3):1280–1285, 2006.
- [31] Stanley Cabay and Dong-Koo Choi. Algebraic computations of scaled padé fractions. SIAM J. Computing., 15(1):243–270, 1986.
- [32] Y. Cao, S. Li, L Petzold, and R. Serban. Adjoint sensitivity analysis for differential-algebraic equations: the adjoint dae system and its numerical solution. SIAM J. Sci. Comput., 24(3):1076–1089, 2003.
- [33] F. Chagla, C. Cabani, and P.M. Smith. Perfectly matched layer for fdtd computations in piezoelectric crystals. *IEEE Ultrasonics Symp.*, pages 517–520, 2004.
- [34] T.F. Chan, J. Cong, T. Kong, and J. Shinnerl. Multilevel optimization for large-scale circuit placement. Proc. IEEE/ACM Int. Conf. on Computer-Aided Design, pages 171–176, November 2000.
- [35] Shih-Hui Chang and Allen Taflove. Finite-difference time-domain model of lasing action in a four-level two-electron atomic system. Opt. Express, 12(16):3827– 3833, 2004.
- [36] Zhiming Chen and Xuezhe Liu. An adaptive perfectly matched layer technique for time-harmonic scattering problems. SIAM J. Num. Anal., 43:645–671, 2005.
- [37] H. Cheng. Advanced Analytic Methods in Applied Mathematics, Science and Engineering. Luban Press, Boston, 2006.
- [38] R.L. Chern, C.C. Chang, C.C. Chang, and R.R. Hwang. Large full band gaps for photonic crystals in two dimensions computed by an inverse method with multigrid acceleration. *Phys. Rev. E*, 68(026704), 2003.
- [39] W. C. Chew and J. M. Jin. Perfectly matched layers in the discretized space: An analysis and optimization. *Electromagnetics*, 16(4):325–340, 1996.
- [40] W. C. Chew and W. H. Weedon. A 3d perfectly matched medium from modified Maxwell's equations with stretched coordinates. *Microwave and Optical Tech. Lett.*, 7(13):599–604, 1994.
- [41] Weng Cho Chew, Jian-Ming Jin, Eric Michielssen, and Jiming Song, editors. *Fast and Efficient Algorithms in Computational Electromagnetics*. Artech House, 2000.

- [42] G. Chietera, A.H. Bouk, F. Poletti, F. Poli, S. Selleri, and A. Cucinotta. Numerical design for efficiently coupling conventional and photonic-crystal waveguides. *Microwave and Optical Tech. Lett.*, 42(3):196–199, 2004.
- [43] A. G. Chirkov. On the adiabatic theorem in quantum mechanics. Tech. Phys. Lett., 27(2):14–21, 2001.
- [44] A. Christ and H. L. Hartnagel. Three-dimensional finite-difference method for the analysis of microwave-device embedding. *IEEE Trans. Microwave Theory Tech.*, 35(8):688–696, 1987.
- [45] K. Chun and J. Ra. Fast block matching algorithm by successive refinement of matching criterion. Proc. Int. Soc. for Opt. Engineeirng, 1818:196–199, 1992.
- [46] B.-K. Chung and H.-T. Chuah. Modeling of rf absorber for application in the design of anechoic chamber. *PIER*, 43:273–285, 2003.
- [47] P.J.B. Clarricoats and R.A. Waldron. Non-periodic slow-wave and backwardwave structures. J. Electron. Contr., 8:455–458, 1960.
- [48] C. Cohen-Tannoudji, B. Din, and F. Laloë. *Quantum Mechanics*. Hermann, Paris, 1977.
- [49] F. Collino and P. B. Monk. Optimizing the perfectly matched layer. Comput. Methods Appl. Mech. Engrg., 164:157–171, 1998.
- [50] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. Introduction to Algorithms. MIT Press, 3rd edition, 2009.
- [51] S.J. Cox and D.C. Dobson. Maximizing band gaps in two-dimensional photonic crystals. SIAM J. Appl. Math., 59(6):2108–2120, 1999.
- [52] S.J. Cox and D.C. Dobson. Band structure optimization of two-dimensional photonic crystals in h-polarization. J. Comp. Phys., 158(2):214–224, 2000.
- [53] S. A. Cummer. Perfectly matched layer behavior in negative refractive index materials. *IEEE Antennas and Wireless Propagation Lett.*, 3:172–175, 2004.
- [54] Alexei Deinega and Ilya Valuev. Subpixel smoothing for conductive and dispersive media in the finite-difference time-domain method. Opt. Lett., 32(23):3429– 3431, 2007.
- [55] S. Dey and R. Mittra. Efficient computation of resonant frequencies and quality factors of cavities via a combination of the finite-difference time-domain technique and and the padé approximation. *IEEE Microwave Guided Wave Lett.*, 8(12):415–417, 1998.
- [56] Supriyo Dey and R. Mittra. A conformal finite-difference time-domain technique for modelling cylindrical dielectric resonators. *IEEE Trans. Microwave Theory Tech.*, 47(9):1737–1739, 1999.

- [57] A. Ditkowski, K. Dridi, and J. S. Hesthaven. Convergent cartesian grid methods for Maxwell's equations in complex geometries. J. Comp. Phys., 170:39–80, 2001.
- [58] X. T. Dong, X. S. Rao, Y. B. Gan, B. Guo, and W. Y. Yin. Perfectly matched layer-absorbing boundary condition for left-handed materials. *IEEE Microwave* and Wireless Comp. Lett., 14(6):301–303, 2004.
- [59] X.T. Dong, X.S. Rao, Y.B. Gan, B. Guo, and W.Y. Yin. Perfectly matched layer-absorbing boundary condition for left-handed materials. *Microwave and Wireless Components Letters*, *IEEE*, 14(6):301–303, June 2004.
- [60] K. Dossou, L.C. Botten, C.M. de Sterke, R.C. McPhedran, A.A. Asatryan, S. Chen, and J. Brnovic. Efficient couplers for photonic crystal waveguides. *Opt. Comm.*, 265(1):207–219, 2006.
- [61] A. Dranov, J. Kellendonk, and R. Seller. Discrete time adiabatic theorems for quantum mechanical systems. J. Math. Phys., 39(3):1340–1349, 1998.
- [62] C. Elachi. Waves in active and passive periodic structures: A review. Proc. IEE, 64(12):1666–1698, 1976.
- [63] David Elliott. The evaluation and estimation of the coefficients in the Chebyshev series expansion of a function. *Mathematics of Computation*, 18:274–284, 1964.
- [64] Atef Elsherbeni and Veysel Demir. The Finite Difference Time Domain Method for Electromagnetics: With MATLAB Simulations. SciTech, Rayleigh, NC, 2009.
- [65] S. Fan, J.N. Winn, A. Devenyi, J.C. Chen, R.D. Meade, and J. D. Joannopoulos. Guided and defect modes in periodic dielectric waveguides. J. Opt. Soc. Am. B, 12(7):1267–1272, 1995.
- [66] Shanhui Fan, P. R. Villeneuve, and J. D. Joannopoulos. Theoretical investigation of fabrication-related disorder on the properties of photonic crystals. *J. Appl. Phys.*, 78:1415–1418, 1995.
- [67] Shanhui Fan, Pierre R. Villeneuve, J. D. Joannopoulos, and H. A. Haus. Channel drop tunneling through localized states. *Phys. Rev. Lett.*, 80(5):960–963, 1998.
- [68] J. Fang and Z. Wu. Generalized perfectly matched layer for the absorption of propagating and evanescent waves in lossless and lossy media. *IEEE Trans. Microwave Theory Tech.*, 44(12):2216–2222, 1998.
- [69] A. Farjadpour, D. Roundy, A. Rodriguez, M. Ibanescu, P. Bermel, J.D. Joannopoulos, S.G. Johnson, and G.W. Burr. Improving accuracy by sub-pixel smoothing in the finite-difference time domain. *Opt. Lett.*, 31:2972–2974, 2006.

- [70] X.-P. Feng and Y. Arawaka. Off-plane angle dependence of photonic band gap in a two dimensional photonic crystal. *IEEE J. Quantum Electron.*, 32(3):535– 542, 1996.
- [71] M. Folk, R. E. McGrath, and N. Yeager. HDF: An update and future directions. In Proc. 1999 Geoscience and Remote Sensing Symposium (IGARSS), Hamburg, Germany, volume 1, pages 273–275. IEEE Press, 1999.
- [72] The MPI Forum. MPI: A message passing interface. In Supercomputing '93, pages 878–883, Portland, OR, November 1993.
- [73] Matteo Frigo and Volker Strumpen. The memory behavior of cache oblivious stencil computations. J. Supercomputing, 39(2):93–112, 2007.
- [74] S.G. Garcia, J. Juntunen, R.G. Martin, A.P. Zhao, B.G. Olmedo, and A. Raisanen. A unified look at berenger's pml for general anisotropic media. *Microwave* and Optical Tech. Lett., 28(6):302–304, 2001.
- [75] S.D. Gedney. An anisotropic perfectly matched layer-absorbing medium for the truncation of fdtd lattices. *IEEE Trans. Anten. Prop.*, 44(12):1630–1639, 1996.
- [76] P. Gill, W. Murray, and M. Saunders. Snopt: an sqp algorithm for large-scale constrained optimization. SIAM Review, 47(1):99–131, 2005.
- [77] M. Grant and S. Boyd. CVX: matlab software for disciplined convex programming. http://cvxr.com/cvx.
- [78] B.E. Griffith and C.S. Peskin. On the order of accuracy of the immersed boundary method: Higher order convergence rates for sufficiently smooth problems. *J. Comp. Phys.*, 208(1):75–105, 2005.
- [79] Wei-Hua Guo, Wei-Jun Li, and Yong-Zhen Huang. Computation of resonant frequencies and quality factors of cavities by fdtd technique and padé approximation. *IEEE Microwave and Wireless Comp. Lett.*, 11(5):223–225, 2001.
- [80] T. Haas, A. Hesse, and T. Doll. Omnidirectional two-dimensional photonic crystal band gap structures. *Phys. Rev. B*, 73(045130), 2006.
- [81] A. Hakansson, J. Sanchez-Dehesa, and L. Sanchis. Inverse design of photonic crystal devices. *IEEE J. Sel. Areas in Comm.*, 23(7):1365–1371, 2005.
- [82] A. Hakansson, P. Sanchis, J. Sanchez-Dehesa, and J. Marti. High-efficiency defect-based photonic-crystal tapers designed by a genetic algorithm. J. Lightwave Tech., 23(11):3881–3888, 2005.
- [83] Thomas D. Happ, Martin Kamp, and Alfred Forchel. Photonic crystal tapers for ultracompact mode conversion. Opt. Lett., 26(14):1102–1104, 2001.

- [84] Jiang-Qi He and Qing-Huo Liu. A nonuniform cylindrical fdtd algorithm with improved pml and quasi-pml absorbing boundary conditions. *IEEE Trans. Geoscience Remote Sensing*, 37(2):1066–1072, 1999.
- [85] John L. Hennessy and David A. Patterson. Computer Architecture: A Quantitative Approach. Elsevier, San Francisco, CA, 3rd edition, 2003.
- [86] R. Hettich and K. Kortanek. Semi-infinite programming: theory, methods, and applications. SIAM Review, 35(3):380–429, 1993.
- [87] Yeuk L. Hoo, Wei Jin, Chunzheng Shi, Hoi L. Ho, Dong N. Wang, and Shuang C. Ruan. Design and modeling of a photonic crystal fiber gas sensor. *Appl. Opt.*, 42(18):3509–3515, 2003.
- [88] W.P. Huang, C.L. Xu, W. Lui, and K. Yokoyama. The perfectly matched layer boundary condition for modal analysis of optical waveguides: leaky mode calculations. *IEEE Photon. Tech. Lett.*, 8(5):652–654, 1996.
- [89] Yingyan Huang and Seng-Tiong Ho. Computational model of solid-state, molecular, or atomic media for FDTD simulation based on a multi-level multi-electron system governed by pauli exclusion and fermi-dirac thermalization with application to semiconductor photonics. *Opt. Express*, 14(8):3569–3587, 2006.
- [90] M. Ibanescu, S.G. Johnson, D. Roundy, C. Luo, Y. Fink, and J.D. Joannopoulos. Anomalous dispersion relations by symmetry breaking in axially uniform waveguides. *Phys. Rev. Lett.*, 92:063903, 2004.
- [91] Mihai Ibanescu, Steven G. Johnson, Marin Soljačić, J. D. Joannopoulos, Yoel Fink, Ori Weisberg, Torkel D. Engeness, Steven A. Jacobs, and M. Skorobogatiy. Analysis of mode structure in hollow dielectric waveguide fibers. *Phys. Rev. E*, 67:046608, 2003.
- [92] T. Inui, Y. Tanabe, and Y. Onodera. Group theory and its applications in physics. Springer, 1990.
- [93] T. Inui, Y. Tanabe, and Y. Onodera. Group Theory and Its Applications in Physics. Springer-Verlag Terlos, 1996.
- [94] J. D. Jackson. Classical Electrodynamics. Wiley, New York, third edition, 1998.
- [95] J. B. Jensen, L. H. Pedersen, P. E. Hoiby, L. B. Nielsen, T. P. Hansen, J. R. Folkenberg, J. Riishede, D. Noordegraaf, K. Nielsen, A. Carlsen, and A. Bjarklev. Photonic crystal fiber based evanescent-wave sensor for detection of biomolecules in aqueous solutions. *Opt. Lett.*, 29:1974–1976, 2004.
- [96] C. Jiang, M. Ibanescu, J.D. Joannopoulos, and M. Soljacic. Zero–group-velocity modes in longitudinally uniform waveguides. *Appl. Phys. Lett.*, 93:241111, 2008.

- [97] J. Jiang, J. Cai, G.P. Nordin, and L. Li. Parallel microgenetic algorithm design for photonic crystal and waveguide structures. *Opt. Lett.*, 28(23):2381–2383, 2003.
- [98] Jianming Jin. The Finite Element Method in Electromagnetics. Wiley-IEEE Press, 2nd edition, 2002.
- [99] J. D. Joannopoulos, S. G. Johnson, R. D. Meade, and J. N. Winn. *Photonic Crystals: Molding the Flow of Light*. Princeton Univ. Press, 2008.
- [100] Sajeev John. Strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett.*, 58:2486–2489, 1987.
- [101] S.G. Johnson. The NLopt nonlinear-optimization package. http://abinitio.mit.edu/nlopt.
- [102] S.G. Johnson, P. Bienstman, M.A. Skorobogatiy, M. Ibanescu, E. Lidorikis, and J.D. Joannopoulos. Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E*, 66(066608), 2002.
- [103] Steven G. Johnson and al. Vectorial eigensystem perturbation theory for optical waveguides. To be published in 2002.
- [104] Steven G. Johnson, Peter Bienstman, M. Skorobogatiy, Mihai Ibanescu, Elefterios Lidorikis, and J. D. Joannopoulos. Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E*, 66:066608, 2002.
- [105] Steven G. Johnson, M. Ibanescu, M. A. Skorobogatiy, O. Weisberg, J. D. Joannopoulos, and Y. Fink. Perturbation theory for Maxwell's equations with shifting material boundaries. *Phys. Rev. E*, 65:066611, 2002.
- [106] Steven G. Johnson, Mihai Ibanescu, M. Skorobogatiy, Ori Weisberg, Torkel D. Engeness, Marin Soljačić, Steven A. Jacobs, J. D. Joannopoulos, and Yoel Fink. Low-loss asymptotically single-mode propagation in large-core OmniGuide fibers. *Opt. Express*, 9(13):748–779, 2001.
- [107] Steven G. Johnson and J. D. Joannopoulos. Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis. Opt. Express, 8(3):173– 190, 2001.
- [108] Steven G. Johnson, M. L. Povinelli, M. Soljačić, A. Karalis, S. Jacobs, and J. D. Joannopoulos. Roughness losses and volume-current methods in photoniccrystal waveguides. *Appl. Phys. B*, 81:283–293, 2005.
- [109] J. S. Juntunen, N. V. Kantartzis, and T. D. Tsiboukis. Zero reflection coefficient in discretized PML. *IEEE Microwave and Wireless Comp. Lett.*, 11(4):155–157, 2001.

- [110] Noriaki Kaneda, Bijan Houshmand, and Tatsuo Itoh. FDTD analysis of dielectric resonators with curved surfaces. *IEEE Trans. Microwave Theory Tech.*, 45(9):1645–1649, 1997.
- [111] C.Y. Kao, S. Osher, and E. Yablonovitch. Maximizing band gaps in twodimensional photonic crystals by using level set methods. *Appl. Phys. B*, 81:235– 244, 2005.
- [112] T. Kato. On the adiabatic theorem of quantum mechanics. J. Phys. Soc. Jap., 5:435–439, 1950.
- [113] T.K. Katsibas and C.S. Antonopoulos. An efficient pml absorbing medium in fdtd simulations of acoustic scattering in lossy media. *IEEE Ultrasonics Symp.*, pages 551–554, 2002.
- [114] E. Khoo, A. Liu, and J. Wu. Nonuniform photonic crystal taper for highefficiency mode coupling. *Opt. Express*, 13(20):7748–7759, 2005.
- [115] K.M. Kiang, K. Frampton, T.M. Monro, R. Moore, J. Tucknott, D.W. Hewak, D.J. Richardson, and H.N. Rutt. Extruded singlemode non-silica glass holey optical fibres. *Electron. Lett.*, 38(12):546–547, Jun 2002.
- [116] H.K. Kim, M.J.F. Digonnet, G.S. Kino, J. Shin, and S. Fan. Simulations of the effect of the core ring on surface and air-core modes in photonic bandgap fibers. *Opt. Express*, 12(15):3436–3442, 2004.
- [117] Ihn S. Kim and Wolfgang J. R. Hoefer. A local mesh refinement algorithm for the time domain-finite difference method using maxwell's curl equations. *IEEE Trans. Microwave Theory Tech.*, 38(6):812–815, 1990.
- [118] J. C. Knight, J. Broeng, T. A. Birks, and P. St.-J. Russell. Photonic band gap guidance in optical fibers. *Science*, 282:1476–1478, 1998.
- [119] T. Kolda, R. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review*, 45(3):385– 482, 2003.
- [120] Naoya Kono and Masanori Koshiba. General finite-element modeling of 2-D magnetophotonic crystal waveguides. *IEEE Photon. Tech. Lett.*, 17(7):1432– 1434, 2005.
- [121] S. Konorov, A. Zheltikov, and M. Scalora. Photonic-crystal fiber as a multifunctional optical sensor and sample collector. *Opt. Express*, 13:3454–3459, 2005.
- [122] M. Koshiba, Y. Tsuji, and S. Sasaki. High-performance absorbing boundary conditions for photonic crystal waveguide simulations. *IEEE Microwave and Wireless Comp. Lett.*, 11(4):152–154, 2001.

- [123] E. P. Kosmidou, T. I. Kosmani, and T. D. Tsiboukis. A comparative FDTD study of various PML configurations for the termination of nonlinear photonic bandgap waveguide structures. *IEEE Trans. Magn.*, 39(3):1191–1194, 2003.
- [124] C. Kottke, A. Farjadpour, and S.G. Johnson. Perturbation theory for anisotropic dielectric interfaces, and application to sub-pixel smoothing of discretized numerical methods. *Phys. Rev. E*, 77(036611), 2008.
- [125] W. Kuang, C. Kim, A. Stapleton, and J.D. O'Brien. Grating-assisted coupling of optical fibers and photonic crystal waveguides. *Opt. Lett.*, 27(18):1604–1606, 2002.
- [126] V. V. R. K. Kumar, A. George, J. Knight, and P. Russell. Tellurite photonic crystal fiber. Opt. Express, 11(20):2641–2645, 2003.
- [127] Karl S. Kunz and Raymond J. Luebbers. The Finite-Difference Time-Domain Method for Electromagnetics. CRC Press, Boca Raton, 1993.
- [128] T. Kurose, K. Tsuruta, C. Totsuji, and H. Totsuji. Fdtd simulations of acoustic waves in two-dimensional phononic crystals using parallel computer. *Memoirs* of Fac. Eng., Okayama U., 43:16–21, 2009.
- [129] L.D. Landau, L.P. Pitaevskii, and E.M. Lifshitz. *Electrodynamics of Continuous Media*. Butterworth-Heinemann, 2nd edition, 1984.
- [130] M. Lassas and E. Somersalo. On the existence and convergence of the solution of PML equations. *Computing*, 60:229–241, 1998.
- [131] B.H. Lee, J.B. Eom, J. Kim, D.S. Moon, U-.C. Park, and G-.H. Yang. Photonic crystal fiber coupler. *Opt. Lett.*, 27(10):812–814, 2002.
- [132] Jae-Yong Lee and Noh-Hoon Myung. Locally tensor conformal FDTD method for modeling arbitrary dielectric surfaces. *Microwave Opt. Tech. Lett.*, 23(4):245–249, 1999.
- [133] P.T. Leung, S.Y. Liu, and K. Young. Completeness and time-independent perturbation of the quasinormal modes of an absorptive and leaky cavity. *Phys. Rev. A*, 49:3982–3989, 1994.
- [134] R.J. Leveque and Z. Li. The immersed interface method for elliptic equations with discontinuous coefficients and singular sources. SIAM J. Num. Anal., 31(4):1019–1044, 1994.
- [135] Zhi-Yuan Li and Kai-Ming Ho. Light propagation in semi-infinite photonic crystals and related waveguide structures. *Phys. Rev. B*, 68(155101), 2003.
- [136] J. Limpert, T. Schreiber, S. Nolte, H. Zellmer, T. Tunnermann, R. Iliew, F. Lederer, J. Broeng, G. Vienne, A. Petersson, and C. Jakobsen. High-power air-clad large-mode-area photonic crystal fiber laser. *Opt. Express*, 11(7):818–823, 2003.

- [137] Calvin Lin and Larry Snyder. Principles of Parallel Programming. Addison Wesley, 2008.
- [138] Q. H. Liu and J. Q. He. Quasi-pml for waves in cylindrical coordinates. Microwave and Optical Tech. Lett., 19(2):107–111, 1998.
- [139] Q.H. Liu. An fdtd algorithm with perfectly matched layers for conductive media. Microwave and Optical Tech. Lett., 14(2):134–137, 1998.
- [140] Q.H. Liu. Perfectly matched layers for elastic waves in cylindrical and spherical coordinates. J. Acoust. Soc. Am., 105(4), 1999.
- [141] Q.H. Liu. Pml and pstd algorithm for arbitrary lossy anisotropic media. IEEE Microwave Guided Wave Lett., 9(2):48–50, 1999.
- [142] Chiyan Luo, Mihai Ibanescu, Steven G. Johnson, and J. D. Joannopoulos. Cerenkov radiation in photonic crystals. *Science*, 299:368–371, 2003.
- [143] R.J. MacKinnon and G.F. Carey. Analysis of material interface discontinuities and superconvert fluxes in finite difference theory. J. Comp. Phys., 75(1):151– 167, 1988.
- [144] E. Magi, P. Steinvurzel, and B. Eggleton. Tapered photonic crystal fibers. Opt. Express, 12(5):776–784, 2004.
- [145] V. A. Mandelshtam and H. S. Taylor. Harmonic inversion of time signals and its applications. J. Chem. Phys., 107(17):6756–6769, 1997.
- [146] B. J. Mangan, L. Farr, A. Langford, P. J. Roberts, D. P. Williams, F. Couny, M. Lawman, M. Mason, S. Coupland, R. Flea, and H. Sabert. Low loss (1.7 db/km) hollow core photonic bandgap fiber. In *Proc. Opt. Fiber Commun.* (*OFC*)., Los Angeles, 2004. Optical Society of America. Paper PDP24.
- [147] A.A. Maradudin and A.R. McGurn. Out of plane propagation of electromagnetic waves in a two-dimensional periodic dielectric medium. J. Mod. Opt., 41(2):275–284, 1994.
- [148] D. Marcuse. Theory of Dielectric Optical Waveguides. Academic Press, San Diego, 2nd edition, 1991.
- [149] E. A. Marengo, C. M. Rappaport, and E. L. Miller. Optimum PML ABC conductivity profile in FDFD. *IEEE Trans. Magn.*, 35(3):1506–1509, 1999.
- [150] Alexander P. McCauley, Alejandro W. Rodriguez, John D. Joannopoulos, and Steven G. Johnson. Casimir forces in the time domain: II. Applications. arXiv.org e-Print archive, page arXiv:0906.5170, 2009.
- [151] J.F. McMillan, X. Yang, N.C. Panoiu, R.M. Osgood, and C.W. Wong. Enhanced stimulated raman scattering in slow-light photonic crystal waveguides. *Opt. Lett.*, 31(9):1235–1237, 2006.

- [152] K. O. Mead and L. M. Delves. On the convergence rate of generalized fourier expansions. IMA J. Appl. Math., 12(3):247–259, 1973.
- [153] R. D. Meade, A. M. Rappe, K. D. Brommer, J. D. Joannopoulos, and O. L. Alerhand. Accurate theoretical analysis of photonic band-gap materials. *Phys. Rev. B*, 48:8434–8437, 1993. Erratum: S. G. Johnson, *ibid.* 55, 15942 (1997).
- [154] R.D. Meade, K.D. Brommer, A.M. Rappe, and J.D. Joannopoulos. Existence of a photonic band gap in two dimensions. *Appl. Phys. Lett.*, 61(4):495–497, 1992.
- [155] A. Mekis, S. Fan, and J. D. Joannopoulos. Absorbing boundary conditions for FDTD simulations of photonic crystal waveguides. *IEEE Microwave Guided Wave Lett.*, 9:502–504, 1999.
- [156] A. Mekis and J.D. Joannopoulos. Tapered couplers for efficient interfacing between dielectric and photonic crystl waveguides. J. Lightwave Tech., 19(6):861– 865, 2001.
- [157] A. Messiah. Quantum Mechanics: Vol. II. Wiley, New York, 1976. Ch. 17.
- [158] Paolo Mezzanotte, Luca Roselli, and Roberto Sorrentino. A simple way to model curved metal boundaries in fdtd algorithm avoiding staircase approximation. *IEEE Microwave Guided Wave Lett.*, 5(8):267–269, 1995.
- [159] T. Miyashita. Acoustic defect-mode waveguides fabricated in sonic crystal: Numerical analyses by elastic finite-difference time-domain method. Japanese J. Appl. Phys., 45(5B):4440–4447, 2006.
- [160] Ahmad Mohammadi, Hamid Nadgaran, and Mario Agio. Contour-path effective permittivities for the two-dimensional finite-difference time-domain method. *Opt. Express*, 13(25):10367–10381, 2005.
- [161] T.M. Monro, Y.D. West, D.W. Hewak, N.G.R. Broderick, and D.J. Richardson. Chalcogenide holey fibres. *Electron. Lett.*, 36(24):1998–2000, Nov 2000.
- [162] Esteban Moreno, Daniel Erni, and Christian Hafner. Modeling of discontinuities in photonic crystal waveguides with the multiple multipole method. *Phys. Rev. E*, 66:036618, 2002.
- [163] D. Mori and T. Baba. Wideband and low dispersion slow light by chirped photonic crystal coupled waveguide. Opt. Express, 13(23):9398–9408, 2005.
- [164] S. Moskow, V. Druskin, T. Habashy, P. Lee, and S. Davidycheva. A finite difference scheme for elliptic equations with rough coefficients using a cartesian grid nonconforming to interfaces. *SIAM J. Num. Anal.*, 36(2):442–464, 1999.
- [165] A. Mutapcic, S. Boyd, A. Farjadpour, S.G. Johnson, and Y. Avniel. Robust design of slow-light tapers in periodic waveguides. *Engineering Optimization*, 41:365–384, 2009.

- [166] Jacek Nadobny, Dennis Sullivan, Waldemar Wlodarczyk, Peter Deuflhard, and Peter Wust. A 3-d tensor FDTD-formulation for treatment of sloped interfaces in electrically inhomogeneous media. *IEEE Trans. Anten. Prop.*, 51(8):1760– 1770, 2003.
- [167] A. S. Nagra and R. A. York. FDTD analysis of wave propagation in nonlinear absorbing and gain media. *IEEE Trans. Anten. Prop.*, 46(3):334–340, 1998.
- [168] J.A. Nelder and R. Mead. A simplex method for function minimization. The Computer Journal, 7(4):308–313, 1965.
- [169] G. Nenciu. On the adiabatic theorem of quantum mechanics. J. Phys. A, 13:L15–L18, 1980.
- [170] Junji Nishii, Toshiharu Yamashita, and Takashi Yamagishi. Chalcogenide glass fiber with a core—cladding structure. Appl. Opt., 28(23):5122–5127, 1989.
- [171] J. Nocedal and S. Wright. *Numerical optimization*. New York: Springer, 1999.
- [172] M. Notomi, K. Yamada, A. Shinya, J. Takahashi, and I. Yokohama. Extremely large group-velocity dispersion of line-defect waveguides in photonic crystal slabs. *Phys. Rev. Lett.*, 87:253902, 2001.
- [173] Michael Okoniewski, Ewa Okoniewska, and Maria A. Stuchly. Threedimensional subgridding algorithm for fdtd. *IEEE Trans. Anten. Prop.*, 45(3):422–429, 1997.
- [174] Alan V. Oppenheim and Ronald W. Schafer. Discrete-Time Signal Processing. Prentice Hall, 3rd edition, 2009.
- [175] Ardavan F. Oskooi, Chris Kottke, and Steven G. Johnson. Accurate finitedifference time-domain simulation of anisotropic media by subpixel smoothing. *Opt. Lett.*, 34(18):2778–2780, 2009.
- [176] Ardavan F. Oskooi, David Roundy, Mihai Ibanescu, Peter Bermel, J. D. Joannopoulos, and Steven G. Johnson. MEEP: A flexible free-software pack-age for electromagnetic simulations by the FDTD method. *Computer Physics Communications*, 181:687–702, January 2010.
- [177] Ardavan F. Oskooi, L. Zhang, Y. Avniel, and Steven G. Johnson. The failure of perfectly matched layers, and towards their redemption by adiabatic absorbers. *Opt. Express*, 16(15):11376–11392, 2008.
- [178] M. Palamaru and Philippe Lalanne. Photonic crystal waveguides: Out-of-plane losses and adiabatic modal conversion. *Appl. Phys. Lett.*, 78(11):1466–1468, 2001.
- [179] I.V. Perez, S.G. Garcia, R.G. Martin, and B.G. Olmedo. Extension of berenger's absorbing boundary conditions to match dielectric anisotropic media. *IEEE Microwave Guided Wave Lett.*, 7(9):302–304, 1997.

- [180] J. Le Person, F. Smektala, T. Chartier, L. Brilland, T. Jouan, J. Troles, and D. Bosc. Light guidance in new chalcogenide holey fibres from gegasbs glass. *Mater. Res. Bull.*, 41(7):1303 – 1309, 2006.
- [181] D. Pissoort, D.V. Ginste, and F. Olyslager. Including pml-based absorbing boundary conditions in the mlfma. *IEEE Antennas and Wireless Propagation Lett.*, 4:312–315, 2005.
- [182] Davy Pissoort, Bart Denecker, Peter Bienstman, Frank Olyslager, and Daniël De Zutter. Comparative study of three methods for the simulation of twodimensional photonic crystals. J. Opt. Soc. Am. A, 21:2186–2195, 2004.
- [183] Davy Pissoort and Frank Olyslager. Termination of periodic waveguides by PMLs in time-harmonic integral equation-like techniques. *IEEE Antennas and Wireless Propagation Lett.*, 2:281–284, 2003.
- [184] J.M. Pottage, D.M. Bird, T.D. Hedley, T.A. Birks, J.C. Knight, P.St.J. Russell, and P.J. Roberts. Robust photonic band gaps for hollow core guidance in pcf made from high index glass. *Opt. Express*, 11(22):2854–2861, 2003.
- [185] P. Pottier, I. Ntakis, and R.D. La. Photonic crystal continuous taper for lowloss direct coupling into 2d photonic crystal channel waveguides and further device functionality. *Opt. Comm.*, 223(4):339–347, 2003.
- [186] M. Povinelli, S. Johnson, and J. Joannopoulos. Slow-light, band-edge waveguides for tunable time delays. Opt. Express, 13(8):7145–7159, 2005.
- [187] D.W. Prather, J. Murakowski, S. Shi, S. Venkatraman, A. Sharkawy, C. Chen, and D. Pustai. High-efficiency coupling structure for a single-line-defect photonic-crystal waveguide. *Opt. Lett.*, 27(18):1601–1603, 2002.
- [188] S. Rao, D. Wilton, and A. Glisson. Electromagnetic scattering by surfaces of arbitrary shape. *IEEE Trans. Anten. Prop.*, 30(3):409–418, May 1982.
- [189] C. M. Rappaport. Perfectly matched absorbing boundary conditions based on anisotropic lossy mapping of space. *IEEE Microwave Guided Wave Lett.*, 5(3):90–92, 1995.
- [190] Y. S. Rickard and N. K. Nikolova. Enhancing the PML absorbing boundary conditions for the wave equation. *IEEE Trans. Anten. Prop.*, 53(3):1242–1246, 2005.
- [191] T. Ritari, J. Tuominen, H. Ludvigsen, J. Petersen, T. Sørensen, T. Hansen, and H. Simonsen. Gas sensing using air-guiding photonic bandgap fibers. *Opt. Express*, 12(17):4080–4087, 2004.
- [192] J.A. Roden and S.D. Gedney. Convolution pml (cpml): An efficient fdtd implementation of the cfs-pml for arbitrary media. *Microwave and Optical Tech. Lett.*, 27(5):334–339, 2000.

- [193] Alejandro Rodriguez, M. Ibanescu, J. D. Joannopoulos, and Steven G. Johnson. Disorder-immune confinement of light in photonic-crystal cavities. *Opt. Lett.*, 30:3192–3194, December 2005.
- [194] Alejandro Rodriguez, Marin Soljačić, J. D. Joannopoulos, and Steven G. Johnson. $\chi^{(2)}$ and $\chi^{(3)}$ harmonic generation at a critical power in inhomogeneous doubly resonant cavities. *Opt. Express*, 15(12):7303–7318, 2007.
- [195] Alejandro W. Rodriguez, Alexander P. McCauley, John D. Joannopoulos, and Steven G. Johnson. Casimir forces in the time domain: Theory. *Phys. Rev. A*, 80:012115, 2009.
- [196] P. Russell. Photonic crystal fibers. *Science*, 299(5605):358–362, 2003.
- [197] Z.S. Sacks, D. M. Kingsland, R. Lee, and J. F. Lee. A perfectly matched anisotropic absorber for use as an absorbing boundary condition. *IEEE Trans. Anten. Prop.*, 43(12):1460–1463, 1995.
- [198] Matthew N.O. Sadiku. Numerical Techniques in Electromagnetics. CRC, 2nd edition, 2000.
- [199] K. Saitoh, N.A. Mortensen, and M. Koshiba. Air-core photonic band-gap fibers: The impact of surface modes. Opt. Express, 12(3):394–400, 2004.
- [200] Peter H. Salus. A Quarter Century of UNIX. Addison-Wesley, Reading, MA, 1994.
- [201] P. Sanchis, P. Bienstman, B. Luyssaert, R. Baets, and J. Marti. Analysis of butt coupling in photonic crystals. *IEEE J. Quantum Electron.*, 40:541–550, 2004.
- [202] M.D. Settle, R.J.P. Engelen, M. Salib, A. Michaeli, L. Kuipers, and T.F. Krauss. Flatband slow light in photonic crystals featuring spatial pulse compression and terahertz bandwidth. *Opt. Express*, 15(1):219–226, 2007.
- [203] J.J. Shirron and T.E. Giddings. A finite element model for acoustic scattering from objects near a fluid-fluid interfaces. *Comp. Methods. in Appl. Mech. Engr.*, 196(1–3):279–288, 2006.
- [204] Gerard L. G. Sleijpen and Diederik R. Fokkema. BiCGSTAB(l) for linear equations involving unsymmetric matrices with complex spectrum. *Elec. Trans. on Num. Analysis*, 1:11–32, 1993.
- [205] Gerard L. G. Sleijpen, Hank A. van der Vorst, and Diederik R. Fokkema. BiCGstab(l) and other hybrid bi-CG methods. Num. Algorithms, 7:75–109, 1994.
- [206] Charlene M. Smith, Natesan Venkataraman, Michael T. Gallagher, Dirk Müller, James A. West, Nicholas F. Borrelli, Douglas C. Allan, and Karl W. Koch. Lowloss hollow-core silica/air photonic bandgap fibre. *Nature*, 424:657–659, 2003.

- [207] Marin Soljačić, Mihai Ibanescu, Steven G. Johnson, Yoel Fink, and J. D. Joannopoulos. Optimal bistable switching in non-linear photonic crystals. *Phys. Rev. E Rapid Commun.*, 66:055601(R), 2002.
- [208] Marin Soljačić, Mihai Ibanescu, Steven G. Johnson, J. D. Joannopoulos, and Yoel Fink. Optical bistability in axially modulated OmniGuide fibers. *Opt. Lett.*, 28:516–518, 2003.
- [209] Marin Soljačić, Steven G. Johnson, Shanhui Fan, Mihai Ibanescu, Erich Ippen, and J. D. Joannopoulos. Photonic-crystal slow-light enhancement of non-linear phase sensitivity. J. Opt. Soc. Am. B, 19:2052–2059, 2002.
- [210] Marin Soljačić, Elefterios Lidorikis, Mihai Ibanescu, Steven Johnson, J. Joannopoulos, and Yoel Fink. Optical bistability and cutoff solitons in photonic bandgap fibers. *Opt. Express*, 12:1518–1527, 2004.
- [211] J. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. Optimization Methods and Software, 11:625–653, 1999.
- [212] Dennis M. Sullivan. Electromagnetic Simulation Using the FDTD Method. Wiley–IEEE Press, New York, 2000.
- [213] J.-H. Sun and T.-T. Wu. Propagation of surface acoustic waves through sharply bent two-dimensional phononic crystal waveguides using a finite-difference timedomain method. *Phys. Rev. B*, 74:174305, 2006.
- [214] K. Svanberg. A class of globally convergent optimization methods based on conservative convex separable approximations. SIAM J. Optim., 12(2):555–573, 2002.
- [215] Allen Taflove and Susan C. Hagness. Computational Electrodynamics: The Finite-Difference Time-Domain Method. Artech, Norwood, MA, 3rd edition, 2005.
- [216] A. Talneau, Ph. Lalanne, M. Agio, and C.M. Soukoulis. Low-reflection photoniccrstal taper for efficient coupling between guide sections of arbitrary widths. *Opt. Lett.*, 27(17):1522–1524, 2002.
- [217] F. L. Teixeira and W. C. Chew. Systematic derivation of anisotropic pml absorbing media in cylindrical and spherical coordinates. *IEEE Microwave Guided Wave Lett.*, 7(11):371–373, 1997.
- [218] F. L. Teixeira and W. C. Chew. General close-form PML constitutive tensors to match arbitrary bianisotropic and dispersive linear media. *IEEE Microwave Guided Wave Lett.*, 8(6):223–225, 1998.
- [219] F.L. Teixeira and W.C. Chew. A general approach to extend berenger's absorbing boundary condition to anisotropic and dispersive media. *IEEE Trans. Anten. Prop.*, 46(9):1386–1387, 1998.

- [220] F.L. Teixeira and W.C. Chew. General closed-form pml constitutive tensors to match arbitrary bianisotropic and dispersive linear media. *IEEE Microwave Guided Wave Lett.*, 8(6):223–225, 1998.
- [221] F.L. Teixeira and W.C. Chew. Complex space approach to perfectly matched layers: a review and some new developments. Int. J. Numer. Model., 13(5):441– 455, 2000.
- [222] F.L. Teixeira, C.D. Moss, W.C. Chew, and J.A. Kong. Split-field and anisotropic-medium pml-fdtd implementations for inhomogeneous media. *IEEE Trans. Microwave Theory Tech.*, 50(1):30–35, 2002.
- [223] Ulrich Trottenberg, Cornelius W. Oosterlee, and Anton Schuller. *Multigrid.* Academic Press, 2000.
- [224] Yasuhide Tsuji and Masanori Koshiba. Finite element method using port truncation by perfectly matched layer boundary conditions for optical waveguide discontinuity problems. J. Lightwave Tech., 20(3):463–468, 2002.
- [225] K. Umashankar, A. Taflove, and S. Rao. Electromagnetic scattering by arbitrary shaped three-dimensional homogeneous lossy dielectric objects. *IEEE Trans. Anten. Prop.*, 34(6):758–766, June 1986.
- [226] Y.A. Vlasov, M. O'Boyle, H.F. Hamann, and S.J. McNab. Active control of slow light on a chip with photonic crystal waveguides. *Nature*, 438:65–69, 2005.
- [227] W. Wadsworth, R. Percival, G. Bouwmans, J. Knight, and P. Russell. High power air-clad photonic crystal fibre laser. Opt. Express, 11:48–53, 2003.
- [228] S. Wang, R. Lee, and F.L. Teixeira. Anisotropic-medium pml for vector fetd with modified basis functions. *IEEE Trans. Anten. Prop.*, 54(1):20–27, 2006.
- [229] Zheng Wang, Y. D. Chong, John D. Joannopoulos, and Marin Soljačić. Reflection-free one-way edge modes in a gyromagnetic photonic crystal. *Phys. Rev. Lett.*, 100:013905, 2008.
- [230] A. J. Ward and J. B. Pendry. Refraction and geometry in Maxwell's equations. J. Mod. Opt., 43(4):773–793, 1996.
- [231] A.D. Waren, L.S. Lasdon, and D.F. Suchman. Optimization in engineering design. Proc. IEE, 55(11):1885–1897, 1967.
- [232] W. D. Warters. WT4 millimeter waveguide system: introduction. Bell Syst. Tech. J., 56:1825–1827, 1977. The introduction to a special issue with many useful articles.
- [233] W. H. Weber, S. L. McCarthy, and G. W. Ford. Perturbation theory applied to gain or loss in an optical waveguide. *Appl. Opt.*, 13(4):715–716, 1974.

- [234] Andrew R. Weily, Levente Horvath, Karu P. Esselle, and Barry C. Sanders. Performance of PML absorbing boundary conditions in 3d photonic crystal waveguides. *Microwave and Optical Tech. Lett.*, 40:1–3, 2004.
- [235] G.R. Werner and J.R. Cary. A stable FDTD algorithm for non-diagonal anisotropic dielectrics. J. Comp. Phys., 226:1085–1101, 2007.
- [236] J. West, C. Smith., N. Borrelli, D. Allan, and K. Koch. Surface modes in air-core photonic band-gap fibers. Opt. Express, 12(8):1485–1496, 2004.
- [237] J.N. Winn, R.D. Meade, and J.D. Joannopoulos. Two-dimensional photonic band-gap materials. J. Mod. Opt., 41(2):257–273, 1994.
- [238] M. Wright. Direct search methods: once scroned, now respectable. Proceedings of the 1995 Dundee Biennial Conference in Numerical Analysis, pages 191–208, 1995.
- [239] Y. Xu, R.K. Lee, and A. Yariv. Propagation and second-harmonic generation of electromagnetic waves in a coupled-resonator optical waveguide. J. Opt. Soc. Am. B, 17(3):387–400, 2000.
- [240] Yong Xu, Reginald K. Lee, and Amnon Yariv. Adiabatic coupling between conventional dielectric waveguides and waveguides with discrete translational symmetry. Opt. Lett., 25(10):755–757, 2000.
- [241] E. Yablonovitch. Inhibited spontaneous emission in solid-state physics and electronics. *Phys. Rev. Lett.*, 58:2059–2062, 1987.
- [242] M. Yan and P. Shum. Improved air-silica photonic crystal with a triangular airhole arrangement for hollow-core photonic bandgap fiber design. Opt. Lett., 30(15):1920–1922, 2005.
- [243] M.F. Yanik and S. Fan. Stopping light all-optically. Phys. Rev. Lett., 92:083901, 2004.
- [244] A. Yariv. *Quantum Electronics*. Wiley, New York, third edition, 1989.
- [245] Kiyotoshi Yasumoto, editor. Electromagnetic Theory and Applications for Photonic Crstals. CRC, 2005.
- [246] Kane S. Yee. Numerical solution of initial boundary value problems involving maxwells equations in isotropic media. *IEEE Trans. Anten. Prop.*, 14(3):302– 307, 1966.
- [247] Ralph H. Young and Walter J. and Deal, Jr. Extensions of the adiabatic theorem in quantum mechanics. *Phys. Rev. A*, 1(2):419–429, 1970.
- [248] Wenhua Yu, Raj Mittra, Tao Su, Yongjun Liu, and Xiaoling Yang. Parallel Finite-Difference Time-Domain Method. Artech House, Norwood, MA, 2006.

- [249] I. A. Zagorodnov, R. Schuhmann, and T. Weiland. A uniformly stable conformal FDTD-method in cartesian grids. *Intl. J. Num. Modelling*, 16:127–141, 2003.
- [250] L. Zhang, J.H. Lee, A. Farjadpour, J. White, and S.G. Johnson. A novel boundary element method with surface conductive absorbers for 3-d analysis of nanophotonics. *Microwave Symposium Digest, 2008 IEEE MTT-S International*, pages 523–526, 2008.
- [251] Lei Zhang, Jung Hoon Lee, Ardavan Oskooi, Amit Hochman, Jacob K. White, and Steven G. Johnson. A novel boundary element method using surface conductive absorbers for full-wave analysis of 3-D nanophotonics. arXiv.org e-Print archive, April 2010.
- [252] A.N. Zhao. Generalized-material-independent pml absorbers used for the fdtd simulation of electromagnetic waves in 3-d arbitrary anisotropic dielectric and magnetic media. *IEEE Trans. Microwave Theory Tech.*, 46(10):1511–1513, 1998.
- [253] A.P. Zhao. The limitations of the perfectly matched layers based on e-h fields for arbitrary anisotropic media – reply to "comment on 'on the matching conditions of different pml schemes applied to multilayer isotropic dielectric media'". *Microwave and Optical Tech. Lett.*, 32(3):237–241, 2002.
- [254] A.P. Zhao, J. Juntunen, and A.V. Raisanen. Material independent pml absorbers for arbitrary anisotropic dielectric media. *Electron. Lett.*, 33(18):1535– 1536, 1997.
- [255] S. Zhao. High order matched interface and boundary methods for the helmholtz equation in media with arbitrarily curved interfaces. J. Comp. Phys., 229:3155– 3170, 2010.
- [256] C. Zheng. A perfectly matched layer approach to the nonlinear schrodinger wave equations. J. Comp. Phys., 227(1):537–556, 2007.
- [257] Richard W. Ziolkowski, John M. Arnold, and Daniel M. Gogny. Ultrafast pulse interactions with two-level atoms. *Phys. Rev. A*, 52(4):3082–3094, 1995.
- [258] Svetlana S. Zivanovic, Kane S. Yee, and Kenneth K. Mei. A subgridding method for the time-domain finite-difference method to solve maxwell's equations. *IEEE Trans. Microwave Theory Tech.*, 39(3):471–470, 1991.
- [259] Lin Zschiedrich, Roland Klose, Achim Schädle, and Frank Schmidt. A new finite element realization of the perfectly matched layer method for helmholtz scattering problems on polygonal domains in two dimensions. J. Comput. Appl. Math., 188:12–32, 2006.