A computational study of dielectric photonic-crystal-based accelerator cavities

by

C. A. Bauer

B.S., California Polytechnic State University, 2005

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2012

This thesis entitled:
A computational study of dielectric photonic-crystal-based accelerator cavities
written by C. A. Bauer
has been approved for the Department of Physics

John R. Cary

Albin J. Gasiewski

Tobin L. Munsat

Scott E. Parker

Gregory R. Werner

Date

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Bauer, C. A. (Ph.D., Physics)

A computational study of dielectric photonic-crystal-based accelerator cavities

Thesis directed by Prof. John R. Cary

Future particle accelerator cavities may use dielectric photonic crystals to reduce harmful wakefields and increase the accelerating electric field (or *gradient*). Reduced wakefields are predicted based on the bandgap property of some photonic crystals (i.e. frequency-selective reflection/transmission). Larger accelerating gradients are predicted based on certain dielectrics' strong resistance to electrical breakdown. Using computation, this thesis investigated a hybrid design of a 2D sapphire photonic crystal and traditional copper conducting cavity. The goals were to test the claim of reduced wakefields and, in general, judge the effectiveness of such structures as practical accelerating cavities. In the process, we discovered the following: (1) resonant cavities in truncated photonic crystals may confine radiation weakly compared to conducting cavities (depending on the level of truncation); however, confinement can be dramatically increased through optimizations that break lattice symmetry (but retain certain rotational symmetries); (2) photonic crystal cavities do not ideally reduce wakefields; using band structure calculations, we found that wakefields are increased by flat portions of the frequency dispersion (where the waves have vanishing group velocities).

A complete comparison was drawn between the proposed photonic crystal cavities and the copper cavities for the Compact Linear Collider (CLIC); CLIC is one of the candidates for a future high-energy electron-positron collider that will study in greater detail the physics learned at the Large Hadron Collider. We found that the photonic crystal cavity, when compared to the CLIC cavity: (1) can lower maximum surface magnetic fields on conductors (growing evidence suggests this limits accelerating gradients by inducing electrical breakdown); (2) shows increased transverse dipole wakefields but decreased longitudinal monopole wakefields; and (3) exhibits lower accelerating efficiencies (unless a large photonic crystal is used).

We have also developed a new electromagnetics algorithm for simulating dielectric objects using the finite difference technique. We explicitly tackled the issue of numerical accuracy and produced an algorithm that converges to exact electromagnetic eigenfrequencies with second-order error in the size of a grid cell. Although the algorithm is unstable in the time-domain, it remains useful as an eigensolver.

Dedication

To Dr. John A. Zoidberg, who teaches us humility in the face of academic recognition.

"...it's been years since medical school, so remind me. Disemboweling in your species: fatal or nonfatal?" [1]

– Dr. John A. Zoidberg

[1] From *Futurama: Why Must I Be a Crustacean in Love?, Feb. 6, 2000*, ©The Curiosity Company and 20th Century Fox Television

Acknowledgements

Thanks naturally go to my advisor, John Cary, whose unbridled optimism and fresh insights pulled our research through numerous challenges and setbacks, and whose leadership in the accelerator community gave our results widespread exposure. I am indebted to Greg Werner for the immense support in designing/running/debugging simulations, the late-evening discussions about the curiosities and frustrations of algorithms and electromagnetics, and above all the lessons learned by observing his calm, complete, and logical approach to not just scientific research, but all of life's interesting questions. Special thanks to Alexej Grudiev of CERN, who provided the CAD model and parameters for the CLIC cavity. Thanks to the Quantum Quartet for countless hours of classwork domination. Thanks to my family for all the trips and adventures that keep life interesting. Thanks to Tory. You make everything better. Your emotional support made this possible.

This work was supported by the U.S. Department of Energy grant DE-FG02-04ER41317. Most simulations were performed on a local computer cluster obtained with a funding supplement made possible by the American Recovery and Reinvestment Act. Simulations were also performed on Hopper (a 153,216-core Cray XE6 supercomputer) at the National Energy Research Scientific Computing Center; thanks are extended to Scott Parker for giving me processor-hours on this machine. The simulation framework VORPAL was used extensively in this thesis; thank you to the entire VORPAL team. I am extremely grateful for the TRILINOS team at Sandia National Laboratory, who build and maintain an amazing parallel linear algebra C++ library that enabled all the results in the final chapter of this thesis. Their mailing-list support is remarkably prompt, enthusiastic, and thorough. The creators of MIT Photonic Bands are also thanked profusely.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction: particle accelerators for high-energy physics

A particle accelerator accelerates (go figure) and guides charged particles using electromagnetic fields. In high-energy particle physics research, accelerators collide common particles such as electrons, positrons, protons, and ions to create heavier uncommon particles. The existence and properties of these uncommon particles shed light on the underlying laws of nature. The energies of the colliding particles are so large, they require accelerators tens of kilometers in length (and billions of dollars in cost). However, the enormity of these machines belies their level of precision; point-like particles traveling at the speed of light must collide head-on. The following traits are demanded of any collider: (1) the accelerating electric field must be maximized to minimize the length and cost of the machine; and (2) accelerated particles must be controlled with exquisite accuracy to maximize the probability of collisions. The dominant barriers to each goal are (1) electrical breakdown and (2) wakefields.

The subject of this thesis is a particle-accelerating structure (see Fig. 1.1) that may increase electric fields and collision probabilities above those offered by more traditional structures. The structure attempts (1) to increase the breakdown limit by using dielectric materials in place of conducting materials and (2) to reduce wakefields using the exotic electromagnetic properties of photonic crystals (PhCs). The remainder of this chapter fills in some details of the above arguments and concludes with an overview of the findings in this thesis. Enjoy.

1.1 Discovery and precision

Historically, hadron colliders like the LHC (proton-proton collisions) have been *discovery* instruments, uncovering new physics in ever higher-energy regimes. Electron-positron colliders on the other hand have been primarily *precision* instruments, built to scrutinize the discoveries of the hadron machines at slightly lower energies with "cleaner" collisions. This is because electrons and positrons are fundamental particles whereas hadrons are composite—made of quarks and gluons. However, it is also possible for discovery to arise from precision. For example, while the top quark was officially discovered by the Tevatron which collided protons and anti-protons at 1.8 TeV, its existence was also suggested by measurements in the Large Electron Positron (LEP) collider at energies of ≈ 100 GeV by examining the perturbative effects of *virtual* top quarks in common LEP collision events. In the end, the two techniques complement each other in pinning down the details of any new physics.

If the Higgs boson exists, the LHC will soon discover it; however, precision measurements of the Higgs' properties will be needed to build on the Standard Model. An electron-positron collider would enable Higgs production mechanisms that are not available to the LHC for such measurements. As a result, a linear e^+e^- collider would be able to measure directly the couplings between Standard Model particles and the Higgs, whereas the LHC can make accurate measurements of coupling *ratios*. Of course, there are questions beyond the existence of the Higgs, such as whether each Standard Model particle has a supersymmetric partner. A TeV e^+e^- collider would provide tests of this theory [14].

1.2 Linear vs circular

An e^+e^- collider at TeV energies demands a linear geometry for the following reasons. Accelerating charged particles emit radiation. In a circular collider like the LHC, the relativistic protons constantly undergo centripetal acceleration, inducing what is termed *synchrotron* radiation. The total power radiated from a highly-relativistic particle of mass m, energy E, and charge q moving

	CLIC	ILC
Energy (TeV)	3	0.5–1
E_{acc} (MV/m)	100	32
Frequency (GHz)	12	1.3
Length (km)	48	31
Wall-plug Power (MW)	560	300

Table 1.1: *Parameters of the two most likely candidates for a new linear e^+e^- collider.*

in a circular path of radius R goes like [37]

$$P \propto \frac{q^2 E^4}{R^2 m^4}. \tag{1.1}$$

For the same radius and particle energy, an electron loses approximately 10^{12} times more power to synchrotron radiation than a proton does. Thus, the circular collider geometry becomes prohibitive at lower energies for the lighter electrons and positrons.

Because a TeV e^+e^- collider requires a linear geometry, the maximization of the accelerating electric field E_{acc} is a high priority (to reduce the length, and thus cost, of the machine). Superconductors make great accelerating structures because of their low electromagnetic losses, but the quenching of superconductivity at high surface magnetic fields fundamentally limits the maximum E_{acc}. Very pure copper structures can withstand more intense electromagnetic fields, giving sustainable accelerating electric fields of ≈ 100 MV/m, but exhibit much higher electromagnetic losses compared to superconducting cavities. For any material, losses and high-E_{acc} limits also depend on the frequency of the electromagnetic fields (especially true for electromagnetic losses in conductors); therefore, along with material type, one of the first choices to be made in accelerator design is the frequency of the accelerating electromagnetic wave.

The two most likely candidates for a future linear collider are the Compact Linear Collider (CLIC) and the International Linear Collider (ILC). The CLIC design is based on copper accelerating cavities while the ILC uses superconductors. The basic parameters are given in Table 1.1.

1.3 Luminosity

Oscillating electromagnetic fields accelerate the particles in a collider; therefore, particles are accelerated in clusters (or *bunches*). A train of bunches is called a *beam*. Luminosity is basically the particle collision rate. Higher luminosity means better statistics and a greater chance of observing interesting rare collision events. It is defined by

$$\mathcal{L} = \frac{n_b N^2 f_{\text{rep}}}{A} H_b \tag{1.2}$$

where N is the number of particles in a bunch, n_b is the number of bunches in the beam, f_{rep} is the repetition rate of the machine (beams per second), A is the transverse cross-sectional area of a bunch, and H_b is called the beam enhancement factor. The rep rate of linear colliders is much less than that of circular colliders (because circular colliders reuse bunches). However, the bunch cross-sections in a linear collider can be much smaller since each bunch is used in only a single collision (denser bunches disrupt each other more during a collision). Therefore, in the next linear collider, minimization of bunch cross-section is a high priority (this is related to minimization of the bunch *emittance*—loosely, the bunch phase-space volume). Minimized cross-sections bring up the issue of particle control. Denser bunches demand higher precision from the accelerator and, therefore, lower wakefields.

1.4 Wakefields

Wakefields are the electromagnetic fields created by and left behind a relativistic bunch as it interacts with various accelerator components along the beam path. Because the particle bunches are small, electromagnetic pulses are emitted, which consist of many electromagnetic frequencies. These fields are generally unwanted, since they tend to disrupt (i.e. increase the emittance of) both the bunch that creates them and any trailing bunches in the beam. The level of the wakefields behind a bunch are determined largely by the charge of the bunch, misalignments of the bunch, and the geometry of the accelerating structure (higher bunch charges, larger bunch displacements, and more dramatic changes in accelerator-component cross-sections lead to higher wakefields).

The majority of wakefields in a linear accelerator are caused by the beam's interaction with the accelerating cavities (mostly because the accelerating cavities are the most numerous component). Therefore, minimizing wakefields is a major goal of accelerator cavity design.

Wakefields are inevitable; in cavity design, one tries to minimize their excitation (usually at the expense of accelerating efficiency) and/or damp them quickly. Because superconducting cavities have such low losses, the high efficiency of acceleration can be somewhat sacrificed for lower wakefield excitation. Normal conducting structures (that allow higher E_{acc}) have much greater losses, and therefore require heavy damping of wakefields. The goal in wakefield damping is to couple all wakefields out of the accelerating structure while keeping the accelerating fields strongly confined to the beam path.

1.5 Photonic crystal accelerator cavities

In principle, PhC accelerator cavities made of dielectric materials could increase E_{acc} *and* reduce wakefields compared to more traditional conducting cavity designs. Empirical evidence suggests that dielectrics such as sapphire or silicon could withstand larger surface electromagnetic fields than conductors; and, PhCs can selectively scatter electromagnetic fields based on their oscillation frequency. Dielectric PhC cavities can thus confine the accelerating fields (which oscillate at a single frequency) efficiently and at a high strength while letting harmful wakefields propagate away. Additionally, while not as efficient as superconductors, dielectrics may incur lower electromagnetic losses than normal conductors, potentially increasing beam acceleration efficiency. Such structures are the topic of this thesis.

1.6 Overview of this thesis

This thesis analyzes a special class of PhC-based accelerating cavity that incorporates traditional conducting cavity elements. Specifically, a sapphire PhC is used to confine the accelerating fields in the transverse direction (perpendicular to the beam axis), while normally-conducting plates confine in the beam direction. The essential structure is displayed in Fig. 1.1. This design shows

promise as a method of testing the proposed advantages of PhC-based cavities while still being somewhat reasonable to fabricate.

The first major result of this thesis is the optimization of the cavity shown in Fig. 1.1 to improve accelerating mode confinement using fewer sapphire elements. We found that by relaxing the lattice symmetry (but retaining the six-fold rotational symmetry) of a truncated version of Fig. 1.1, the radiative losses of the accelerating mode could be reduced by two orders of magnitude. In addition, this optimization reduced wakefields when compared to the lattice-based structure. Two potential disadvantages accompanied the optimizations: (1) conducting-surface magnetic fields were increased where the inner rods touch the plate and (2) radiative losses became very sensitive to rod positions. The optimizations were performed in 2D; simulation in 3D showed that the beam tube perturbation (absent in 2D) was enough to increase radiation losses significantly. Optimizations in 3D will be required to realize a practical design.

The next major result arises from a comparison between the PhC-based cavities (optimized and lattice) and the highly-optimized conducting cavities designed for CLIC. In this study, we sought to form a complete comparison of all relevant accelerator cavity figures of merit: maximum surface fields, accelerating efficiency, and wakefields. To summarize, we found that:

(1) Accelerating efficiencies are lower in PhC-based cavities than in the CLIC cavity (except in the lattice structure where the number of layers is large).

(2) Maximum conducting-surface magnetic fields (that likely limit E_{acc}) are reduced in the lattice structure but increased in the optimized structures as compared to CLIC.

(3) Transverse wakefields are higher in all PhC-based cavities.

(4) Longitudinal wakefields are lower in all PhC-based cavities.

The increased wakefields in the lattice PhC-based cavity are explained by the presence of low-group-velocity PhC modes, which provides some suggestions for improvement.

The final chapter includes an investigation into numerical algorithms for the more accurate

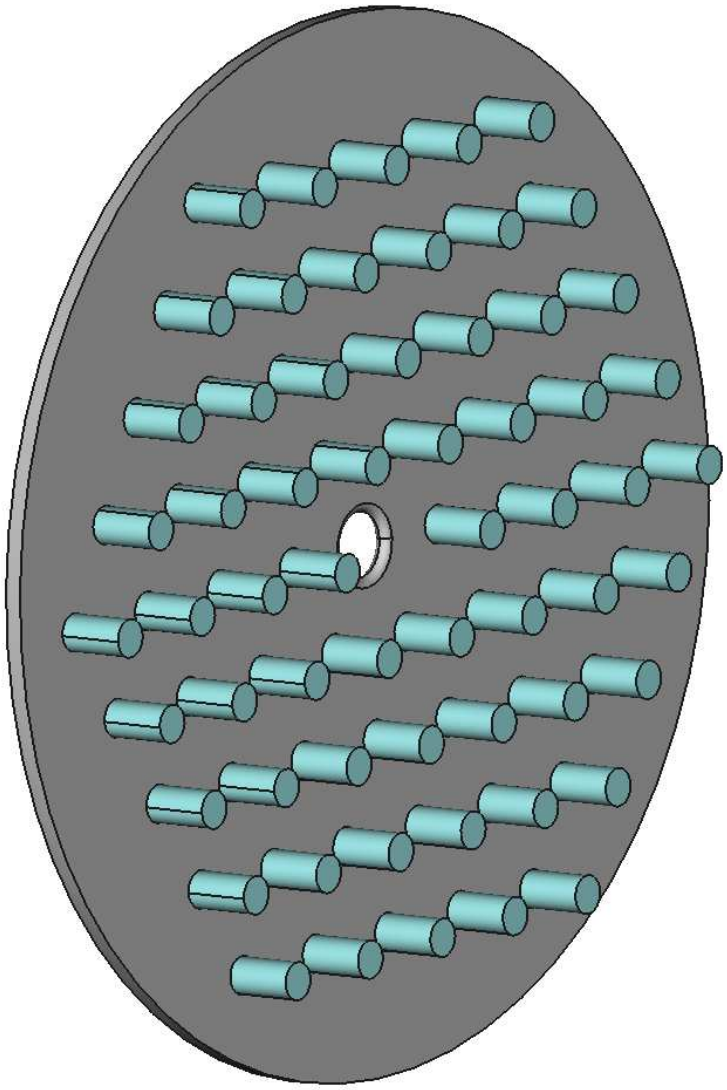

Figure 1.1: *The PhC-based accelerating structure of interest to this thesis. The circular plate with the beam hole is copper while the cylinders are sapphire. A multicell accelerating cavity is formed by stacking the above structure and capping with a final conducting plate. The particle beam moves along the z-axis, which is parallel to the rod axes and located at the center of the structure.*

simulation of dielectric materials in the finite-difference method. This work describes the first implementation of a finite-difference dielectric boundary algorithm to obtain reliable second-order error in resonant frequencies.

Chapter 2

Introduction to accelerator cavities

This chapter gives some background on the important characteristics of an accelerator cavity. In the ideal case, these are: unlimited $E_{\rm acc}$, zero electromagnetic losses (that is, all input electromagnetic energy goes into increasing the beam momentum), and the support of only one mode (the accelerating mode) that couples strongly to a particle beam. The practical cavity necessarily falls short in each of these categories; the goal in cavity design is to minimize each of these shortcomings. We will begin with a description of the simplest of all accelerator cavities, the perfectly-conducting cylindrical pillbox. Next, we generically discuss high-$E_{\rm acc}$ limitations, power losses, and wakefields. The introduction to wakefields will have the most relevance to the results in this thesis. Finally, the physics of multicell cavities is discussed, so that comparisons with the CLIC accelerator cavities can be made in later chapters. Throughout this chapter, many figures of merit will be introduced that help determine the usefulness of a given cavity design; these measures are used throughout the thesis.

2.1 Pillbox cavity

The cylindrical pillbox is the canonical example of an accelerating cavity, with which most cavities in modern accelerators share a strong resemblance. Thus, in discussing the figures of merit that determine the efficacy of a given accelerator cavity, it is useful to provide the corresponding values for the pillbox cavity as a reference. Just the basic characteristics of the pillbox cavity will be presented in this section; some examples of important figures of merit using the pillbox geometry

will be postponed until they are discussed in greater detail in later sections.

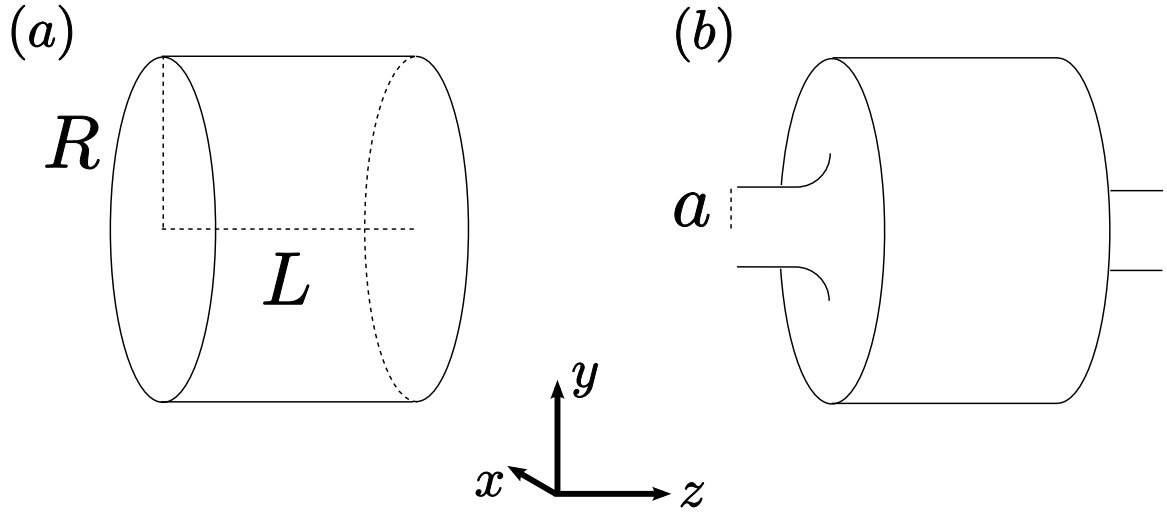

Figure 2.1: *Pillbox cavities. Cavity radius R, cavity length L, and beam tube radius a.*

2.1.1 Resonant modes

A pillbox cavity of radius R and length L is shown in Fig. 2.1. For perfectly-conducting walls, the resonant modes of the pillbox cavity are described in terms of their polarization and azimuthal, radial, and longitudinal wavenumbers (integers m, n, and p, respectively). The polarization is either transverse electric (TE) or transverse magnetic (TM); TE (TM) modes are identified by a lack of longitudinal electric (magnetic) field component (the z-direction in Fig. 2.1)—$E_z = 0$ ($B_z = 0$) for TE (TM) modes. The spatial dependence of pillbox resonant modes is sinusoidal in the azimuthal and longitudinal directions (wavenumbers m and p, respectively) and goes as a Bessel-function in the radial direction (radial wavenumber n). In accelerator cavities, the modes of interest are TM-polarized, since these are the modes that can accelerate (or be excited by) a bunch traveling parallel to the z-axis. In cylindrical coordinates (ρ, ϕ, z), a generic TM_{mnp} mode has the

following spatial dependence [64]

$$E_z = E_0 \cos\left(\frac{p\pi z}{L}\right) J_m\left(\frac{x_{mn}\rho}{R}\right) \cos\left(m\phi\right) \tag{2.1}$$

$$E_\rho = -E_0 \frac{p\pi R}{L x_{mn}} \sin\left(\frac{p\pi z}{L}\right) J'_m\left(\frac{x_{mn}\rho}{R}\right) \cos\left(m\phi\right) \tag{2.2}$$

$$E_\phi = -E_0 \frac{mp\pi R^2}{\rho L x_{mn}^2} \sin\left(\frac{p\pi z}{L}\right) J_m\left(\frac{x_{mn}\rho}{R}\right) \sin\left(m\phi\right) \tag{2.3}$$

$$B_z = 0 \tag{2.4}$$

$$B_\rho = iE_0 \frac{m\omega_{mnp} R^2}{c^2 \rho x_{mn}^2} \cos\left(\frac{p\pi z}{L}\right) J_m\left(\frac{x_{mn}\rho}{R}\right) \sin\left(m\phi\right) \tag{2.5}$$

$$B_\phi = iE_0 \frac{\omega_{mnp} R}{c^2 x_{mn}} \cos\left(\frac{p\pi z}{L}\right) J'_m\left(\frac{x_{mn}\rho}{R}\right) \cos\left(m\phi\right) \tag{2.6}$$

where x_{mn} is the nth root of the mth-order Bessel function ($J_m(x_{mn}) = 0$), J'_m is the derivative of the Bessel function with respect to its complete argument, and the resonant frequencies are

$$\omega_{mnp} = c\sqrt{\left(\frac{x_{mn}}{R}\right)^2 + \left(\frac{\pi p}{L}\right)^2} \tag{2.7}$$

Throughout this thesis, we treat resonant modes with a complex time-varying part. Thus, for a resonant mode n, the fields are

$$\mathbf{E}_n(\mathbf{x}, t) = \mathbf{E}_n(\mathbf{x}) e^{-i\omega_n t} \qquad \mathbf{H}_n(\mathbf{x}, t) = \mathbf{H}_n(\mathbf{x}) e^{-i\omega_n t} \tag{2.8}$$

The time-averaged energy of a resonant mode is then

$$U_n = \frac{1}{4} \int \left[\mathbf{E}_n(\mathbf{x}) \cdot \mathbf{D}^*(\mathbf{x}) + \mathbf{H}_n(\mathbf{x}) \cdot \mathbf{B}^*(\mathbf{x})\right] d^3x. \tag{2.9}$$

In the case of lossless materials, one obtains [37]

$$U_n = \frac{1}{2} \int \varepsilon(\mathbf{x}) |\mathbf{E}_n(\mathbf{x})|^2 \, d^3x \tag{2.10}$$

$$= \frac{1}{2} \int \mu(\mathbf{x}) |\mathbf{H}_n(\mathbf{x})|^2 \, d^3x. \tag{2.11}$$

The above is also a useful approximation when dealing with very low loss materials.

The TM_{010} mode is the accelerating mode, since, for the proper cavity length, it has the strongest coupling to a particle bunch traversing the cavity on axis. The length is chosen based on

the frequency of the TM_{010} mode; the bunch energy gain is maximized when it spends exactly one half-period of the TM_{010} mode oscillation in the cavity. In the case of a highly relativistic bunch traveling at speed c, the length of the pillbox cavity is thus set to $L = \pi c/\omega_{010}$ where ω_{010} is from Eq. 2.7. The spatial dependence of the TM_{010} accelerating mode is

$$E_{z,010}(\rho, \phi, z) = E_0 J_0(k_{01}\rho) \tag{2.12}$$

$$B_{\phi,010}(\rho, \phi, z) = -\frac{E_0}{c} J_1(k_{01}\rho) \tag{2.13}$$

where $k_{01} = 2.405/R$.

2.1.2 E_{acc} and Surface fields

The accelerating electric field (or accelerating gradient) E_{acc} is defined as the average on-axis longitudinal electric field a relativistic particle feels over the length of the cavity

$$E_{\text{acc}} = \frac{1}{L}\int_0^L E_z(\rho = 0, z, t = z/c)\, dz. \tag{2.14}$$

Two important figures of merit for accelerating cavities are the ratios of the maximum cavity surface fields to the above accelerating gradient; namely, $E_{\text{surf,max}}/E_{\text{acc}}$ and $cB_{\text{surf,max}}/E_{\text{acc}}$. This ratio determines the maximum particle energy gain based on the high-field limitations of the cavity materials. The peak surface electric field identifies regions that are most vulnerable to electrical breakdown and the peak surface magnetic field identifies regions of intense surface heating. In the case of the pillbox TM_{010} mode (where $E_{\text{surf,max}}$ is at $\rho = 0$),

$$\frac{E_{\text{surf,max}}}{E_{\text{acc}}} = L \Big/ \int_0^L \sin\left(\frac{\omega_{010} z}{c}\right) dz \tag{2.15}$$

$$= L \Big/ \int_0^L \sin\left(\frac{\pi z}{L}\right) dz \tag{2.16}$$

$$= \frac{\pi}{2} \approx 1.57. \tag{2.17}$$

Similarly, for the TM_{010} mode, $cB_{\text{surf,max}}/E_{\text{acc}} \approx 0.91$.

2.1.3 Beam tubes

To grant a particle bunch passage through a cavity, a hole or *beam tube* must be cut through each of the endplates (see Fig. 2.1b). The beam tube is most often a short cylindrical waveguide of radius a which transports the bunch to the next accelerator component. The fields at the beam tubes differ from those of the closed pillbox; the larger the beam tube radius, the less the mode resembles its closed pillbox counterpart. The beam tube radius affects acceleration in the following ways. As a increases, the accelerating gradient $E_{\rm acc}$ generally diminishes (as measured relative to other important cavity properties such as power losses to the cavity walls, the energy in the cavity, or the peak surface fields). This reduces accelerating efficiency and raises surface field to $E_{\rm acc}$ ratios. However, to the benefit of beam stability, an increasing a is accompanied by decreasing wakefields—the fields left behind by a particle bunch traversing the cavity (discussed at length in Sec. 2.4). This decreases bunch emittance, potentially leading to higher luminosities. The resulting trade-off makes the beam tube radius a critical parameter in cavity optimization.

2.2 High-$E_{\rm acc}$ limitations

The walls of an accelerator cavity are subject to tremendous electric and magnetic fields. At high enough strengths, the fields will melt, ionize, and generally destroy the surface. This field limit determines the minimum length of linac required to reach a target particle energy; and since cost is proportional to length, shorter is usually better. A theoretical upper bound can be provided based simply on the electric field strength binding an outer-shell electron to a nucleus, $\sim$ 10 GV/m (or $\sim$1 Volt per Angstrom) [11]. In practice, the limiting field is much lower, $\sim$ 100 MV/m. At these field strengths, materials begin to emit charged particles, the subsequent motion of which sucks the energy from the cavity fields, limiting their peak values. The record for peak surface fields is continually increasing, however, as new materials are tested and well-known materials are manufactured with higher purities and fabricated into structures with increased precision. Complicated (and mostly empirical) dependencies on electromagnetic

frequency or pulse length are also exploited to gain higher peak surface fields. In the end, cavity geometries can also be optimized to reduce the ratio of surface fields to accelerating fields.

In this section, we briefly overview the physical mechanisms (or the current understanding of them) that contribute to high-gradient limitations in accelerator cavities. We begin with the high-gradient limit in superconductors ($E_{\rm acc} \approx 50$ MV/m) because it is the lowest and because the reason for the limit (the quenching of the superconducting state) is the most straightforward (despite the nontrivial theoretical underpinnings). Next, more general mechanisms are discussed that can apply to a variety of wall materials; however, the focus is on normal conductors. The CLIC project anticipates that their copper cavities will withstand surface electric fields of approximately 250 MV/m at GHz frequencies (putting $E_{\rm acc} \approx 100$ MV/m for typical cavity designs) [26]. Finally, the poorly-known high-gradient limits of dielectrics (still in their infancy as a material for accelerator cavities) are discussed. Recent studies at GHz frequencies have shown a resistance to breakdown in dielectric-lined metallic cavities up to surface electric fields of tens of MV/m.

2.2.1 $B_{\rm c}$ and superconductors

Although superconducting cavities suffer the same barriers to high gradients as normal conducting cavities (detailed in the following sections), there is a theoretical limit specific to superconducting cavities that tends to occur at lower fields. This limit is the critical surface magnetic field, beyond which the superconducting state is quenched and the advantages are lost. Niobium is the material of choice in modern superconducting cavities with a critical temperature of 9.2 Kelvin and a critical magnetic field of 0.2 Tesla (this is the magnetostatic value—the critical field for oscillating magnetic fields can be slightly higher) [44, 64]. For the pillbox $\rm TM_{010}$ mode with maximum surface magnetic field at $\rho = 0.765R$, the limiting on-axis electric field is (using Eq. 2.13)

$$E_{0,\rm max} = cB_{\rm c}/J_1(1.841) \tag{2.18}$$

where $B_{\rm c}$ is the superconductor's critical magnetic field. In the case of Niobium, $E_{0,\rm max} \approx 100$ MV/m; thus, in principle, the accelerating gradient is limited to $E_{\rm acc} \approx 64$ MV/m. Of course,

optimizations to the cavity shape can decrease the peak surface magnetic field for a given on-axis electric field, making the ratio $cB_{\mathrm{surf,max}}/E_{\mathrm{acc}}$ an important figure of merit for superconducting cavities.

2.2.2 Volume currents

More generally, accelerating gradients in accelerator cavities are limited by anomalous electrical currents within the cavity volume. These stray volume currents often occur as *breakdown* events, where a violent release of current from an isolated region of the cavity wall rapidly depletes the cavity field energy. Other mechanisms are subtler, evidenced by a threshold electromagnetic input power past which the field levels in the cavity do not increase (an increase in input power past the threshold proportionately increases the volume current densities). In all cases, the undesirable volume currents result from the emission of electrons/ions from the cavity surface and the subsequent acceleration of the charges by the cavity fields.

Pushing to higher E_{acc} requires an understanding of the physical mechanisms that produce these free charges. Arguably, the most harmful candidate is *field emission*—the tunneling of electrons out of a material in the presence of strong surface electric fields. Field emission is usually attributed to the onset of several debilitating volume current effects—most notably, *breakdown*. Secondary emission is another charge-producing mechanism, where an accelerated charge that bombards a cavity surface produces more free charges on impact (of course, secondary emission often requires some other form of surface emission to get started, like field emission). When resonant, this effect is called *multipacting*, and leads to the more subtle field limit. Finally, heating of the walls due to large surface magnetic fields can cause structural damage, which in turn increases losses and can encourage breakdown at lower cavity fields. In the following subsections, we discuss these precursors to electric field limits in greater detail.

2.2.3 Field emission

For large electric fields perpendicular to a material surface, electrons can tunnel out of a solid. This is called field emission and is often a precursor to (and therefore a likely cause of) breakdown events in accelerator cavities. It can also limit E_{acc} through the "dark current" effect, where field-emitted electrons are caught and accelerated down the beam line, sucking energy from the accelerating mode [76].

The Fowler-Nordheim theory is an oft-quoted explanation of field emission in conducting cavities [23]. In short, it treats the electrons in a metal as a Fermi gas, and the metallic surface (in the presence of a static electric field) as a quantum potential barrier (the barrier lowers and thins as the surface electric field is increased). The tunneling probability of the electrons through the barrier is used to calculate an emitted current density; at room temperature, a greatly simplified form for this current density is

$$J \propto f(\phi) E^2 \exp\left(-\frac{6.53 \times 10^9 \phi^{1.5}}{E}\right) \tag{2.19}$$

where ϕ is the work-function of the metal, E is the surface electric field, and $f(\phi)$ is a function of ϕ omitted for simplicity [76]. The general dependence of the current density on the surface field in Eq. 2.19 is verified experimentally; however, the magnitudes of J and E often disagree. This discrepancy is usually accounted for with the simple replacement, $E \to \beta E$, where β is referred to as the *enhancement factor*.

For perfectly smooth surfaces, the maximum electric field anywhere on a cavity surface is on the order of the accelerating gradient. However, true surfaces may be riddled with protrusions, which tend to enhance the local electric field. The enhancement factor, β is a quantitative measure of this effect, and can be loosely defined as the ratio of the electric field on the surface of a protrusion to the background electric field (i.e. the electric field in the limit of perfectly smooth surface). Experimentally, the enhancement factor can be estimated using the Fowler-Nordheim relation with knowledge of the surface-emitted current, macroscopic cavity electric field, and work-function of the wall material. Breakdown events are most likely to occur at protrusions due to

the enhanced field emission; thus, an important stage of accelerator cavity construction involves "smoothing" these bumps out. The most effective way to eliminate troublesome protrusions has been to condition the cavity using high-power electromagnetic pulses. The highest-β protrusions are self-selectively ablated by the enhanced local electric fields, leaving (one hopes) a smoother surface behind. In practice, the point at which electromagnetic conditioning ceases to improve the maximum sustainable gradient in copper cavities tends to correspond to a β between 40 and 100, even though scanning electron microscope images show protrusions with $\beta \approx 10$ or lower [76]. For cavity fields of ≈ 200 MV/m, these experimental values of β put the actual limiting surface field on a protrusion nearer the theoretical limit. The explanation of these inflated values of β is currently under investigation in the accelerator community.

2.2.4 Multipacting

Multipacting is a resonant process leading to the buildup of free electrons above a cavity surface. The effect is produced by an interplay between the oscillating cavity fields and a surface phenomenon called secondary electron emission, whereby an electron incident on a material surface produces more free electrons (or secondary electrons) upon impact. When an electron is emitted from a cavity surface, the cavity fields can accelerate and steer the electron such that it collides with the surface (possibly at a different location). If the collision produces one or more secondary electrons, the process can continue, and a cloud of free electrons can quickly form. The buildup of electron current drains the field energy. However, for the process to be resonant and deleterious, the path of the cascading multipacting electrons should be somewhat localized. Otherwise, the cascading electrons can drift to a region where their trajectories do not produce an emitting surface impact, and multipacting ceases. This instability of the process is usually exploited to eliminate multipacting in a given cavity; when multipacting is observed and located, the geometry of that area can be finely adjusted to remove the multipacting resonance. Experimentally, multipacting is observed when an increase in electromagnetic power input to the cavity is not accompanied by an increase in cavity field levels (in superconducting cavities, thermal maps of the cavity during

operation can help pinpoint the multipacting region [64]).

The secondary emission coefficient (material dependent), $\delta(K)$, determines the average number of secondary free electrons produced by an incident electron with kinetic energy, K. The dependence of the coefficient on incident electron energy for a given material generally behaves like the curve shown in Fig. 2.2. The shape of the curve (namely that $\delta(K \to 0, \infty) \to 0$) is easily explained: at low incident energies, the collisions are too weak to produce secondaries; at high energies, the incident electron is more likely to penetrate deep into the material before losing a significant amount of its energy in a collision, thus confining scattered electrons to the solid. For most materials, the curve is above $\delta = 1$ somewhere between incident electron energies of ~ 100 eV and ~ 1000 eV [64]. Thus, raising the field levels in an accelerator cavity is a potential defense against multipacting (of course, this might just relocate the multipacting region to a surface closer to an electric field node). Unfortunately, $\delta_{\max}$ can be high for dielectrics like sapphire, which are of interest to this work.

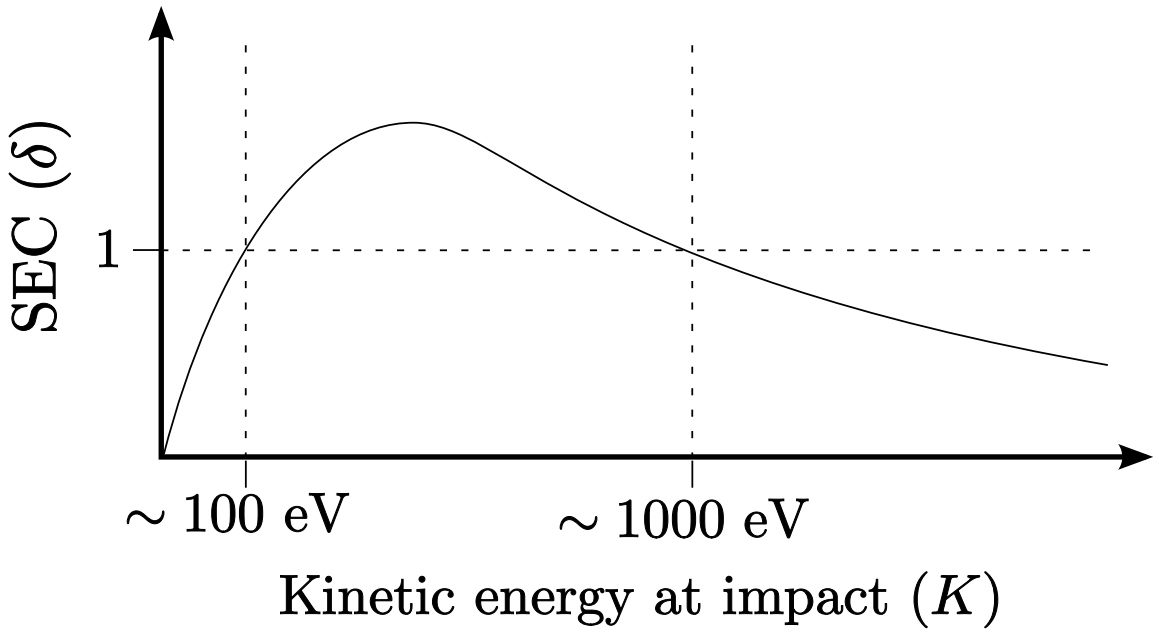

Figure 2.2: *Generic dependence of the secondary emission coefficient (SEC) on impacting particle kinetic energy. Figure based on Fig. 10.4 of [64].*

2.2.5 Heating

Heating of conducting cavity walls due to finite resistivity is another barrier to high-gradients. In this case, $E_{\rm acc}$ is limited by the surface magnetic fields, which rapidly heat the cavity walls by inducing currents. The rapid (and repetitive in pulsed operation) heating can cause microscopic structural damage due to thermal stresses. The microscopic damage (in the form of dislocations or cracking) increases losses by effectively increasing the surface resistance [68, 69]. Regions of intense heating are also more prone to breakdown as observed in a recent study of a copper photonic bandgap structure [55]; in that study, it was postulated that intense ohmic heating resulted in surface protrusions, which then enhanced the local electric field and caused arcing. The results of Ref. [19] also indicate that peak surface magnetic fields (and their duration) correlate more strongly with breakdown levels than do peak surface electric fields in copper pillbox cavities.

2.2.6 Strong fields in dielectrics

Dielectrics comprise an alternate diverse class of material currently being considered in the pursuit of high gradients. They have seen simple use as linings of traditional conducting cavities (to decrease the field at the conductor) or in more complicated configurations where they are the only material confining the accelerating mode (e.g. photonic bandgap cavities). Unfortunately, since the free-electron gas model is inapplicable to insulators, a theory for field emission from dielectrics is much more complicated and remains an active area of research; to date, speculations on the high-gradient limits for dielectrics come primarily from experiment.

Some experimental evidence suggests that at higher frequencies (relative to GHz), dielectrics may be able to support accelerating gradients at the few GV/m level. This is certainly suggested by the prevalence of dielectric optics over metal mirrors in high-power laser research. Explicit tests of breakdown that impinge high-power laser pulses on insulators have a long history. Nearly 40 years ago, a maximum breakdown field of ~ 1 GV/m was established for 10 ps laser pulses at a wavelength of 10.6 μm incident on alkali-halides such as NaCl, KCl, KF, etc. [7]. More recent laser

breakdown studies at wavelengths of 800 nm give similar breakdown thresholds at pulse lengths of $\sim$ 100 ps in SiO_2 [21, 75].

At lower frequencies, only a handful of breakdown experiments have been performed using more accelerator-related designs; therefore, knowledge of breakdown limits in dielectric accelerating structures at GHz frequencies (of most interest to this thesis) is limited. In the THz frequency range, breakdown events were recently observed in hollow-core SiO_2 cylindrical waveguides where the fields were excited by a high-energy electron bunch [74]. The study reported wakefields of $\sim$ 1-10 GV/m before witnessing breakdown events, but did not fully diagnose the pulse shape (100 ps pulse lengths were surmised). A related design (termed a dielectric loaded accelerator or DLA) was recently used to study dielectric breakdown in the GHz frequency range [66]. In this case, a klystron delivered $\sim$ 100 ns pulses of electromagnetic power into a waveguide consisting of an alumina-lined copper cylinder. Field levels reached $\sim$ 10 MV/m without breaking down; however, they were limited by multipacting—nearly half of the input power was absorbed by multipacting electrons at the highest power level. A later study showed that multipacting could be significantly reduced by coating the inner surface of the dielectric with a thin layer of TiN [38].

2.2.7 Summary

Currently, the highest sustained E_{acc} have been demonstrated in copper accelerator cavities at GHz frequencies. Therefore, this class of cavity underlies the CLIC project. Experiments with dielectrics have found an even greater resistance to electrical breakdown, but only at THz/optical frequencies and with short pulses; breakdown limits at GHz frequencies are still uncharacterized due to multipacting issues. The design explored in this thesis offers another testbed (unique from those mentioned in Sec. 2.2.6) for dielectric breakdown experiments at GHz frequencies.

2.3 Electromagnetic power losses in accelerator cavities

Power loss is another important consideration in cavity design, for reasons of both feasibility and economics. Cost obviously rises with power consumption and heat dissipation is power that

cannot be reclaimed. Also, the heating of accelerator components can seriously degrade their performance (e.g. by increasing material resistivities, and thus compounding the losses problem).

RF power (usually generated by a klystron) enters an accelerator cavity and begins filling the TM_{010} mode with energy. Because the walls of the cavity are lossy, they begin to heat and drain energy from the TM_{010} mode. When a particle bunch passes through the cavity, it gains energy from the TM_{010} mode, but also loses energy by radiating wakefields. In this section, we define the accelerator cavity figures of merit (namely, the quality factor and the shunt impedance) relating to the power lost to the cavity materials. Discussions on the calculation of power losses for different materials will also be presented along with examples using the pillbox geometry.

2.3.1 Quality factor

The quality factor, Q, is a measure of the lifetime of an electromagnetic mode in a lossy resonant cavity. It is defined for the nth resonant mode of a cavity by

$$Q_n = \frac{\omega_n U_n}{P_n} \tag{2.20}$$

where ω_n is the resonant frequency of the nth mode, U_n is the stored energy, and P_n is the power loss. Since the power loss $P_n = -dU_n/dt$, then $U_n(t) = U_{n,0}\exp(-\omega_n t/Q_n)$, so that the quantity, $Q_n/2\pi$, is the number of oscillation periods before the energy of that mode in the cavity falls to $1/e$ of its original value. A higher Q reduces the input power required to sustain a given electromagnetic energy in the cavity (for the accelerating mode). The Q-factor is also an important indicator of troublesome non-TM_{010} resonant modes (called *higher-order modes* or HOMs) that contribute to the wakefields; HOMs with high Qs are likely to cause problems as long-range wakefields (wakefields that affect later bunches in a bunch train) because of their longevity.

2.3.2 Shunt impedance

The shunt impedance, R_{shunt}, for the accelerating mode is a measure of the power required to produce a given accelerating voltage. It is defined by

$$R_{\text{shunt}} = \frac{V_{\text{acc}}^2}{P} \tag{2.21}$$

where P is the power lost to the cavity materials and V_{acc}, the accelerating voltage, is

$$V_{\text{acc}} \equiv \int_0^L E_z(\rho = 0, z, t = z/c)\, dz = E_{\text{acc}} L. \tag{2.22}$$

Obviously, a larger shunt impedance is desirable. The much higher losses in copper cavities (as opposed to superconducting) demand shape optimizations to maximize the shunt impedance. For example, increasing the beam tube radius generally reduces the shunt impedance, which explains why it is usually smaller in normal conducting cavities than in superconducting cavities (also, a larger beam tube can help suppress wakefields—superconducting cavities can afford the associated decrease in shunt impedance).

2.3.3 Electromagnetic power loss in conducting walls

For conducting cavities, the Q-factor of the accelerating mode is determined by the ohmic losses in the walls, which depend on the surface magnetic fields. The differential loss per area at a conducting surface (for "good" conductors such as copper) can be estimated by

$$\frac{dP}{da} = \frac{1}{2\sigma\delta}|\mathbf{H}_s|^2 = \frac{R_s}{2}|\mathbf{H}_s|^2 \tag{2.23}$$

where σ is the conductivity, δ is the skin depth (the depth over which the fields inside the conductor are appreciable), $\mathbf{H}_s$ is the magnetic field at the surface, and $R_s = 1/\delta\sigma$ is called the surface resistance [37, 64]. For good conductors, where the fields do not differ much from the perfectly conducting case, $\mathbf{H}_s$ is usually taken to be the surface field for the perfect conductor. This technique is used extensively in computer codes, since perfectly conducting boundary conditions are much easier to simulate. The total ohmic losses for a cavity are then given by the integral of Eq. 2.23

over the entire cavity surface. By Eq. 2.20, the Q-factor is then

$$Q_n = \frac{\omega_n U_n}{P_n} = \frac{\mu_0 \omega_n \int_V |\mathbf{H}|^2 \, dv}{R_s \int_A |\mathbf{H}_s|^2 \, da} \tag{2.24}$$

where A is the cavity surface and V is the cavity volume. Since the ratio of the integrals scales as $1/\omega_n$, Eq. 2.24 is often written as

$$Q_n = \frac{G_n}{R_s} \tag{2.25}$$

where G_n is the "geometry" factor for mode n (the form of which can be inferred from Eqs. 2.24 and 2.25. As its name indicates, the quantity G_n is independent of the cavity size (that is, a proportional scaling of the entire cavity). The geometry factor for the TM_{010} mode in the pillbox is 257 Ω. Thus, for a surface resistance of 0.028 Ω (copper at 12 GHz), $Q_{010} = 9300$. A superconducting pillbox cavity with a surface resistance of 20 nΩ would have a Q_{010} of 1.3×10^{10} [64]. Similarly, integrating Eq. 2.23 using the analytic TM_{010} surface fields, the shunt impedances for the copper and superconducting pillboxes are $R_{010} = 1.8 \times 10^6$ Ω (copper at 12 GHz) and $R_{010} = 2.5 \times 10^{12}$ Ω ($R_s = 20$ nΩ), respectively.

2.3.4 Electromagnetic power loss in dielectrics

Dielectrics also incur losses, quantified by the *loss tangent*, or imaginary part of the permittivity. If we define an isotropic complex permittivity as $\varepsilon(\mathbf{x}) = \varepsilon_r(\mathbf{x}) + i\varepsilon_i(\mathbf{x})$, then the power lost to the heating of dielectrics for mode n is calculated from

$$P_{\mathrm{diel}} = \frac{1}{2}\omega_n \int_V \varepsilon_i(\mathbf{x}) |\mathbf{E}_n(\mathbf{x})|^2 \, d^3x \tag{2.26}$$

where V is the cavity volume and $\varepsilon_i(\mathbf{x}) = 0$ outside the dielectric [37]. Using Eq. 2.20, Eq. 2.26, and the low-loss approximation to the stored energy, Eq. 2.10, the Q-factor due purely to dielectric heating is

$$Q_{\mathrm{diel}} \approx \frac{\int_V \varepsilon_r(\mathbf{x}) |\mathbf{E}_n(\mathbf{x})|^2 \, d^3x}{\int_V \varepsilon_i(\mathbf{x}) |\mathbf{E}_n(\mathbf{x})|^2 \, d^3x} \tag{2.27}$$

$$\approx \frac{1}{\tan\delta} \frac{U_n}{U_{n,\mathrm{diel}}} \tag{2.28}$$

Temperature (K)	300	77	4.2
$\tan(\delta)_{\|\|}$	4.8×10^{-6}	2.7×10^{-8}	5.0×10^{-9}
$\tan(\delta)_{\perp}$	9.1×10^{-6}	5.9×10^{-8}	5.0×10^{-9}

Table 2.1: *Loss tangents of single-crystal sapphire for different temperatures near 10 GHz resonant frequency. Subscripts* $(||)$ *and* $(\perp)$ *refer to responses in the direction of the c-axis of the crystal and transverse to the c-axis, respectively [48].*

where $U_{n,\text{diel}}$ is the energy in the dielectric volume and $\tan\delta \equiv \varepsilon_i/\varepsilon_r$ is called the *loss tangent*, the inverse of which is seen to be a lower limit on the Q_{diel} of the cavity (if the dielectric completely fills the cavity or if the electric fields for mode n are confined to the dielectric volume, then $Q_n \to 1/\tan(\delta)$). An example of a very low-loss dielectric is sapphire with loss tangents near 10^{-5} for GHz frequencies at room temperature—decreasing with temperature (see Table 2.1). In principle, if a TM_{010} mode can be confined with sapphire as efficiently as it can be confined with copper, Q-factors and shunt impedances 100 times larger are feasible.

2.3.5 Section summary

Wall losses in superconducting cavities are miniscule; thus, the acceleration efficiencies (shunt impedances) are enormous. Unfortunately, the critical magnetic field enforces a strict limit on the maximum accelerating gradient, and thus the length of a superconducting accelerator. Low-loss dielectrics like sapphire may be the next best bet, since their losses can be orders of magnitude lower than those of conductors and their E_{acc}-limit could be larger (especially under the right conditions—high frequency, short electromagnetic pulses). Confining radiation with conductors, however, is much easier than confining radition solely with dielectrics. Photonic crystals may alleviate this problem and are discussed in great detail in the next chapter.

2.4 High-order modes and wakefields

The fields due to a constant-velocity, highly relativistic particle bunch in vacuum are strongly concentrated in the plane perpendicular to the direction of motion (see Fig. 2.3a) [37]. For a bunch moving parallel to the axis of a cylindrical waveguide, the fields are modified slightly to satisfy the

boundary conditions (Fig. 2.3b). When the waveguide is perfectly-conducting, these fields travel self-consistently with the bunch; i.e., there is no scattering of the fields and the momentum of the bunch remains constant [84]. Thus, a uniform perfectly-conducting waveguide empty of electromagnetic fields remains empty after the passage of a bunch. However, when a bunch encounters a change in cross-section of the waveguide, the fields are scattered, and the otherwise empty waveguide fills with radiation behind the bunch (finite wall conductivity can also induce radiation from the bunch). These scattered fields are called the *wakefield* of the bunch and are of great concern in accelerator design since they tend to disrupt both the exciting bunch itself and any trailing bunches. The wakefield thus constrains important accelerator parameters such as inter-bunch spacing and bunch emittance (ultimately limiting luminosity).

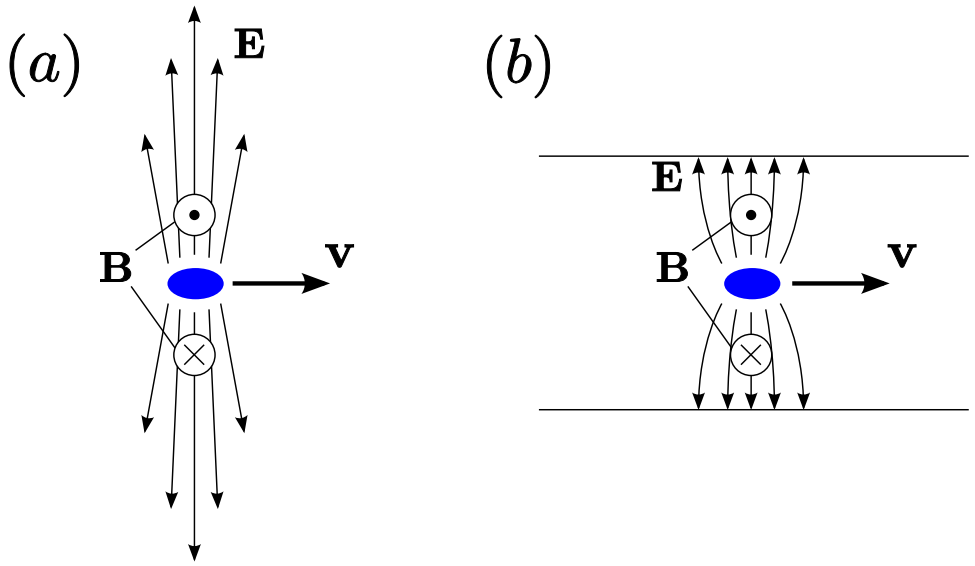

Figure 2.3: *The electromagnetic fields due to a relativistic bunch of positive charges in (a) vacuum and (b) a conducting cylindrical waveguide.*

An accelerator cavity is a dramatic change in waveguide cross-section for a bunch and thus induces significant wakefields. Specifically, at the cavity's beam entry and exit ports, a bunch generates pulses of radiation that continue to resonate in the cavity after its passage. These pulses contain a broad selection of frequencies. The higher-frequency fields quickly exit the cavity through side-couplers or down the beam tubes but can still disturb the tail of the exciting bunch. Lower

frequencies corresponding to trapped cavity eigenmodes can resonate for many oscillations after the bunch has exited, posing a potential problem for trailing bunches. These modes are often called *high-order modes* and are specifically and sometimes individually targeted for suppression in accelerator cavities.

Wakefields are categorized by their region of influence relative to the exciting bunch (*short-range* vs. *long-range*) and by their effect on witness charges (*longitudinal* vs. *transverse*). A bunch of finite length can interact with its own wakefield (the tail of the bunch feels the wakefield of the head of the bunch); this region is referred to as the short-range wakefield. Minimizing the short-range wakefield improves single-bunch stability and can thus lead to higher luminosities. The long-range wakefield describes the region of trailing bunches, which is usually several cavity lengths behind the exciting bunch. Lowering the long-range wakefield can increase beam current and efficiency by reducing the inter-bunch spacing. For each region, the wakefield is further classified by the direction of the force it imparts on a trailing charge. Very simply, the longitudinal wakefield imparts momentum along the beam axis (z-direction) whereas the transverse wakefield acts perpendicularly. The integrated effect of the wakefield on a trailing particle is called the *wake potential* and is related to the total change in momentum of the charge due to the wakefield. In this section, we detail the above concepts and review some common methods of calculating the wakefields for a given structure. Since long-range wakefield suppression is the main subject of this thesis, many of the ideas presented here will be used throughout.

2.4.1 Longitudinal and transverse wakefields

The longitudinal wake potential at transverse position (r, ϕ) due to an ultrarelativistic exciting point charge ($\mathbf{v} = c\hat{\mathbf{z}}$) at transverse position (r', ϕ'), is defined as

$$W_z(r, \phi, s; r', \phi') = \frac{1}{q'} \int E_z \left(r, \phi, z', t = (z' + s)/c; r', \phi'\right) \, dz' \tag{2.29}$$

where z' is the location of the exciting bunch, $s > 0$ is the distance behind the exciting bunch at which the wakefield is sampled, and q' is the charge of the exciting bunch ($W_z = 0$ for $s < 0$ by

causality). The change in momentum of a trailing particle with charge q at relative position s is then simply

$$\Delta p_z(r, \phi, s; r', \phi') = \frac{qq'}{c} W_z(r, \phi, s; r', \phi'). \tag{2.30}$$

High-frequency (compared to TM_{010}) longitudinal wakefields can cause energy spread in particle bunches, which as mentioned in the introduction, lowers luminosity.

Wakefields generated by off-axis exciting charges can give a transverse kick to trailing off-axis particle bunches, thus increasing the emittance of the bunch and likely leading to stronger transverse kicks in subsequent accelerator components. The integrated transverse momentum change of an off-axis bunch is directly proportional to the *transverse* wake potential (a la Eq. 2.30) which is defined as

$$\mathbf{W}_\perp(r, \phi, s; r', \phi') = \frac{1}{q} \int \left(\mathbf{E}_\perp + c\hat{\mathbf{z}} \times \mathbf{B}\right) \left(r, \phi, z', t = (z' + s)/c; r', \phi'\right) \, dz' \tag{2.31}$$

where $\perp$ indicates restriction to the transverse plane. Of course, Eq. 2.31 assumes that the transverse position of the trailing bunch (the bunch that feels the wakefield) remains constant as the bunch traverses the cavity; this is usually a good approximation near the axis for high-quality bunches. Transverse wakefields are almost never calculated directly. Instead, the Panofsky-Wenzel Theorem is used, which quite generally states:

$$\nabla_\perp W_z = \partial_s \mathbf{W}_\perp \tag{2.32}$$

where the derivatives act on the unprimed variables.

Because of the cylindrical symmetry of most accelerator cavities, wakefields are often decomposed into azimuthal harmonics (e.g. monopole, dipole, and quadrupole modes with azimuthal wavenumbers $m = 0$, 1, and 2, respectively). Monopole modes (which have a nonzero E_z along the cavity axis) dominate the wakefield spectrum since a high-quality beam is by definition tightly confined to the cavity axis and is highly cylindrically symmetric. Modes with $m > 0$ (which have an E_z node on axis) can be weakly excited by beam asymmetries or off-axis trajectories. In the theoretical case of a cylindrically symmetric cavity with infinitely long beam tubes, the radial de-

pendence of the mth-order wake potential is separable and simply r^m (for r less than the radius of the beam tubes) [59], allowing for the following general form of the wakefields:

$$W_{z,m}(r, r', \phi, s) = \partial_s X_m(s) r^m r'^m \cos m\phi \tag{2.33}$$

$$\mathbf{W}_{\perp,m}(r, r', \phi, s) = X_m(s) r^{m-1} r'^m \left(\cos m\phi\, \hat{r} - \sin m\phi\, \hat{\phi}\right) \tag{2.34}$$

where ϕ' has been set to zero in the above (since only the difference in azimuth matters for cylindrical symmetry) and $m > 0$ for Eq. 2.34. The transverse wake potential is zero when $m = 0$; only modes with $m > 0$ mediate transverse kicks in this case [59]. The longitudinal and transverse wake potentials are dominated by monopole and dipole wakefields, respectively. The expressions for these contributions are particularly simple:

$$W_{z,0} = \partial_s X_0(s) \tag{2.35}$$

$$\mathbf{W}_{\perp,1} = \mathbf{r}' X_1(s). \tag{2.36}$$

Notice that both expressions are uniform throughout the beam tube region and that the dipole transverse wake potential is in the direction of the drive beam transverse offset.

Because of the scaling in Eqs. 2.33 and 2.34, wake potentials are often expressed in r-normalized units (for example, the strength of the dipole transverse wake potential is usually expressed in units of V/pC/mm). The fact that $X(s)$ is shared between longitudinal and transverse wakefields is a direct result of the Panofsky-Wenzel Theorem (Eq. 2.32). The forms of Eqs. 2.33 and 2.34 can be good approximations of the wake potentials in cavities that slightly break cylindrical symmetry (especially for wake potentials near the axis) and simplify wake potential simulations since the wakefields need only be sampled at a single radius.

2.4.2 Short- and long-range wakefields

Short-range refers to the wakefields within the exciting bunch (i.e. for $s \lesssim \sigma_z$ where σ_z is the exciting bunch length). This section of the wake potential is characterized by a broad frequency spectrum, since, in this short time, high frequencies have not yet been damped or escaped

down beam tubes or side-couplers. The time-domain picture is more straightforward—short-range wakefields are essentially the "reflection" of the bunch field pulse off the beam tube port. This implies that the short-range wakefields depend largely on the beam tube port geometry. In fact, for a beam tube aperture radius a, the amplitude of the short-range wake potential is approximately proportional to $1/a^2$ for monopole modes and $1/a^3$ for dipole modes [84, 64]. Enlarging the beam tube radius is therefore one of the first steps taken toward reducing short-range wakefields. However, there is a trade-off, since larger beam tubes tend to decrease the shunt impedance of the accelerating mode, thereby reducing acceleration efficiency.

Long-range refers to the wakefields that act on trailing bunches (i.e. where $s > c/\omega_{010}$). In contrast to short-range, long-range wakefields are a narrow-band effect; that is, they are essentially the collective oscillation of all cavity-bound modes (i.e. modes with frequencies below the cutoff frequencies of the beam tubes and any side-couplers) that have been excited by the bunch. Frequencies above cutoff disperse quickly (within a few fundamental periods) and can be considered negligible after several fundamental oscillations. The amplitude at which a cavity mode is excited by a bunch depends on 3 properties of the bunch/cavity system: 1) the frequency content of the bunch, 2) the mode field pattern (TM-like modes are excited much more strongly than TE-like modes, and monopole TM modes more strongly than higher-order multipole TM modes), and 3) the transverse position of the bunch. The details of these ingredients are discussed in the following section. Once excited, each mode in the long-range wakefield resonates sinusoidally and decays according to its Q-factor. Therefore, to reduce long-range wakefields given a bunch shape and maximum transverse bunch offset, the goal in cavity design is first to minimize the excitation levels of HOMs and second to damp quickly those HOMs that cannot be avoided.

For a closed conducting or superconducting cavity, the losses due to wall heating are insufficient for the damping of HOMs. Therefore, couplers must be included in the cavity design. HOM couplers are basically waveguides (terminated by electromagnetic absorbers) built in to the cavity walls or attached to the beam tubes just off the ends of cavity. They are designed to couple strongly to unwanted HOMs while minimally perturbing the fundamental accelerating mode. This effect is

usually achieved by selecting the cutoff frequency of the coupler just above the fundamental and placing the coupler in a position where the ratio of troublesome HOM to fundamental field strength is highest. Couplers can introduce problems due to the increased complexity of the cavity geometry. For example, joints that arise from the construction of HOM couplers can have high-curvature features that may enhance electromagnetic fields and promote electrical breakdown.

2.4.3 Fundamental theorem of beam loading

The fundamental theorem of beam loading gives the energy left behind in (or the amplitude of) a given cavity mode by an exciting bunch. A well-known trick simplifies the derivation by considering two identical charges passing through the same resonant cavity. In short, the wakefield of the first charge is cancelled by the wakefield of the second charge; conservation of energy gives the energy transferred between charges, and therefore the energy left in the wakefield by the first charge. The details of the argument follow.

Imagine a lossless cavity empty of electromagnetic energy and two identical point charges each with charge q and energy U_q (both charges are highly relativistic, so that their speeds are unchanged by cavity interaction). Suppose that the first charge enters the cavity at time $t = 0$ and induces an on-axis voltage of $V_{n,1}(t) = V_{n,q}e^{-i\omega_n t}$ in the nth cavity mode ($V_{n,q}$ can be complex). Since electromagnetic power loss is given by

$$\int \mathbf{E} \cdot \mathbf{J}\, d^3x, \tag{2.37}$$

the energy loss of the first charge is expected to be proportional to its charge and the self-induced voltage, i.e.

$$\Delta U_{q,1} = -f|q||V_{n,q}| \tag{2.38}$$

where f is the unknown fraction of the self-induced voltage felt by the first bunch.

We choose the second charge to traverse the cavity exactly one half-period (π/ω_n) after the first charge, such that it sees an *accelerating* voltage $V_{n,1}(\pi/\omega_n) = -V_{n,q}$ due to the first charge.

By superposition with its own induced voltage, the energy gain of the second charge is

$$\Delta U_{q,2} = (1-f)|q||V_{n,q}|. \tag{2.39}$$

The voltage induced by the second charge is $V_{n,2}(t) = V_{n,q}e^{-i\omega_n(t-\pi/\omega_n)} = -V_{n,1}(t)$. Therefore, the total voltage after the passage of the second charge is $V_{n,1}(t) + V_{n,2}(t) = 0$ and the nth cavity mode is annihilated. By conservation of energy, $\Delta U_{q,1} + \Delta U_{q,2} = 0$, and we find that $f = 1/2$. Thus, the fundamental theorem of beam loading says that an exciting particle bunch feels a decelerating voltage at 1/2 the magnitude of the voltage it induces in the cavity.

Regardless of amplitude, the energy of the cavity mode is related to the square of the voltage by a constant, α_n, which is determined purely by the mode pattern (i.e. $U_n = \alpha_n|V_n|^2$). If a charge leaves energy $q|V_{n,q}|/2$ in mode n, then the voltage magnitude can be expressed as

$$|V_{n,q}| = \frac{1}{2\alpha_n}q = 2k_n q \tag{2.40}$$

where we have defined the mode *loss factor* k_n given explicitly by

$$k_n = \frac{1}{4}\frac{|V_n|^2}{U_n}. \tag{2.41}$$

The mode energy left by the charge is then simply

$$U_{n,q} = k_n q^2. \tag{2.42}$$

Out of convenience, the loss factor is often expressed in terms of the shunt impedance and the quality factor of mode n using the relation

$$k_n = \frac{\omega_n R_{n,\text{shunt}}}{4Q_n}. \tag{2.43}$$

Given Eq. 2.40, the total wake potential can be expressed as a sum on modes. Including the phase and losses for each mode, the on-axis longitudinal wake potential due to an ultrarelativistic on-axis point charge (called the *impulse* wake potential) can be written

$$\hat{W}_z(s) = 2\sum_n k_n e^{-\alpha_n s/c}\cos\frac{\omega_n s}{c} \tag{2.44}$$

where $\alpha_n = \omega_n/2Q_n$. For a line charge distribution rather than a point charge, Eq. 2.44 can be used as a kernel, or Green's function, to find the total wake potential via

$$W_z(s) = \frac{\int_{-\infty}^{s} \hat{W}(s-s')\lambda(s')\,ds'}{\int_{-\infty}^{\infty} \lambda(s')\,ds'} \tag{2.45}$$

where λ is the line charge density.

As an example, consider an on-axis line Gaussian bunch, i.e.

$$\lambda(s) = \frac{q}{\sqrt{\pi}\sigma_z} e^{-s^2/2\sigma_z^2} \tag{2.46}$$

Combining Eqs. 2.44, 2.45, and 2.46 and integrating to a value of s that is several σ_z behind the bunch (so that the upper integration limit can be approximated as ∞), we get

$$W_z(s) = 2\sum_n k_n e^{-(\omega_n^2-\alpha_n^2)\sigma_z^2/2c^2} e^{-\alpha_n s/c} \cos\frac{\omega_n s}{c} \tag{2.47}$$

$$= \sum_n \hat{W}_{n,z}(s) e^{-(\omega_n^2-\alpha_n^2)\sigma_z^2/2c^2} \tag{2.48}$$

which shows how the contribution to the wake potential falls off rapidly with increasing frequency above the bunch "cutoff" frequency, $\sim c/\sigma_z$ (the damping factor α_n is usually only a fraction of ω_n). The contribution increases with damping because the spectral response of a mode widens as damping increases, allowing more frequency components of the bunch to help excite the mode.

Under certain geometric circumstances, Eq. 2.44 can be generalized to off-axis transverse excitation and witness positions [84]. In these cases, the expressions for the longitudinal and transverse impulse wake potentials become

$$\hat{W}_z(\mathbf{r}', \mathbf{r}, s) = 2\,\mathrm{Re}\sum_n k_n(\mathbf{r}', \mathbf{r}) e^{(i\omega_n - \alpha_n)s/c} \tag{2.49}$$

$$\hat{\mathbf{W}}_\perp(\mathbf{r}', \mathbf{r}, s) = 2\,\mathrm{Re}\sum_n \boldsymbol{\kappa}_n(\mathbf{r}', \mathbf{r}) e^{(i\omega_n - \alpha_n)s/c} \tag{2.50}$$

where

$$k_n(\mathbf{r}', \mathbf{r}) = \frac{V_n^*(\mathbf{r}')V_n(\mathbf{r})}{4U_n} \tag{2.51}$$

$$\boldsymbol{\kappa}_n(\mathbf{r}', \mathbf{r}) = \frac{cV_n^*(\mathbf{r}')\nabla_\perp V_n(\mathbf{r})}{4(i\omega_n - \alpha_n)U_n} \tag{2.52}$$

and the voltages in the above are complex, i.e.

$$V_n(\mathbf{r}) = \int E_{n,z}(\mathbf{r}, z) e^{-i\omega_n z/c} dz. \tag{2.53}$$

2.4.4 Wakefield calculations

For all but the simplest cavity geometries, the wakefields must be determined through computation. The two most common approaches are: 1) the Condon method (a frequency domain method) and 2) time-domain simulation. Any computational approach is of course limited by the resolution of the simulation, which in the case of wakefield calculations places a lower bound on the bunch length that can be approximated. The cutoff frequency of the bunch must be selected below those electromagnetic frequencies (i.e. wavelengths) that cannot be simulated or are poorly approximated on the computational grid. Fortunately, for long-range wakefields (which are of most interest in this thesis), higher frequencies are less of a concern.

In the Condon method, the wake potential is pieced together using the sum-on-modes approach as written in Eqs. 2.49 and 2.50. Modes are found via an electromagnetic eigensolver code, which finds eigenmodes and eigenfrequencies of the electromagnetic wave equation

$$\nabla \times (\varepsilon^{-1}, \mu^{-1}) \nabla \times (\mathbf{H}, \mathbf{E}) = \frac{\omega^2}{c^2} (\mu \mathbf{H}, \varepsilon \mathbf{E}). \tag{2.54}$$

An advantage of the Condon method is that, once a resonant mode is found, the *impulse* wake potential for that *single* mode can be easily calculated at arbitrary transverse drive and witness offsets. The impulse wake potential can then be used to determine the mode contribution to a wake for an arbitrary bunch profile. A disadvantage is that the number of modes required to approximate the full wakefield spectrum grows quickly with frequency; in 3D resonant cavities, the electromagnetic mode density is proportional to ω^2. This presents a significant challenge to numerical eigensolvers, which must store each mode in memory for the duration of the calculation. However, when most HOMs are effectively damped by the structure, the first few decades of modes can well-approximate the long-range wake.

Time-domain simulations are commonly used to determine wakefields (especially short-range wakefields), since they intrinsically treat the cumulative effect of all cavity modes. In a time-domain wakefield simulation (see Fig. 2.4), a drive bunch is injected into the cavity and followed by regularly-spaced extremely low-charge test particles (test particles retain the charge-to-mass ratio of the particles of interest—electrons). Low charge ensures that the test particles do not create wakefields of their own. The net momentum change of each test particle gives the wake *potential* as a function of trailing distance, s. In general, time-domain simulations are specific to a single bunch profile and transverse offset (however, for cylindrically symmetric systems, a single simulation can provide the wake potentials at arbitrary transverse offsets for identical bunch profiles based on the form of Eqs. 2.33 and 2.34).

How should the simulation resolution be chosen for a given driving bunch profile? The answer depends mostly on the bunch cutoff frequency. In the standard finite-difference method, frequencies are approximated with second-order error in the grid cell size. Thus, to second order, the simulation frequency of mode n, $\tilde{\omega}_n$, is

$$\tilde{\omega}_n \approx \omega_n \left(1 + A\frac{\Delta z^2}{\lambda_n^2}\right) \tag{2.55}$$

where ω_n is the true frequency, $\lambda_n = 2\pi c/\omega_n$, and A is of order unity (in fact, let $A = 1$ for our purposes). One simple choice of resolution arises from a desired accuracy for the highest frequency simulated. For a Gaussian bunch of half-width σ_z, Δz is set to

$$\Delta z \approx \sigma_z \sqrt{\epsilon} \tag{2.56}$$

where $\sigma_z = \lambda_{\text{cutoff}}$ and $\epsilon = (\omega_{\text{cutoff}} - \omega_{0,\text{cutoff}})/\omega_{0,\text{cutoff}}$ is the desired relative error at the cutoff frequency of the bunch (c.f. Eq. 2.48).

A more careful resolution is chosen by considering not only the error in frequency, but how that error affects the accuracy of any point along the wake potential. For a cavity of length L and a maximum wake potential position $s_{\max}$, the time of interest T that the cavity rings with

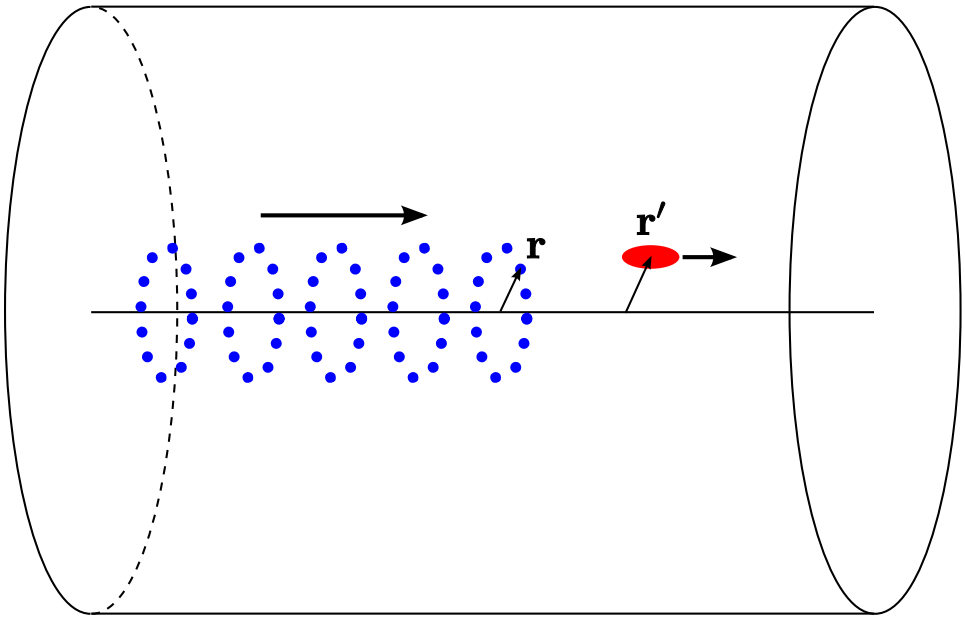

Figure 2.4: *A time-domain wakefield simulation. Excitation bunch (in red) is at transverse position* $\mathbf{r}'$*. Test particles with the electron charge to mass ratio but with greatly reduced absolute charge trail the excitation bunch. The test particles are arranged in rings about the cavity axis in order to extract the azimuthal multipole contributions to the wake potential.*

electromagnetic waves is $T = (L + s_{\max})/c$. Thus, phase error for frequency ω after time T is

$$\Delta\phi = (\omega - \omega_0)T \tag{2.57}$$

$$\approx \frac{2\pi(L + s_{\max})\Delta z^2}{\lambda_0^3}. \tag{2.58}$$

To simulate the wake potential accurately up to $s_{\max}$, the phase error should be small for all frequencies (i.e. $\Delta\phi \ll 2\pi$ for wavelengths $\lambda_0 \leq \sigma_z$). As a very rough upper bound, one can enforce $\Delta\phi \leq 2\pi$ for $\sigma_z \leq \lambda_0$; thus, from Eq. 2.58

$$\Delta z \leq \sqrt{\frac{\sigma_z^3}{L + s_{\max}}}. \tag{2.59}$$

The above result ensures that the *phase* of the wake potential is correct out to $s_{\max}$ (note that Ref. [9] does not include $s_{\max}$ in their recommendation, only L). This accuracy is essential for the analysis of accelerator cavities that rely on detuning to quell wakefields (discussed in the upcoming section on multicell cavities). However, Eq. 2.56 should suffice for determining the damping of higher-order modes.

Figure 2.5 shows a comparison between simulated and analytic wake potentials in a closed pillbox cavity [80]. The wake potentials were driven by a Gaussian bunch. The analytic wake potentials were calculated by the Condon method; i.e. for each pillbox mode, the impulse wake potential was convoluted with the Gaussian bunch profile via Eq. 2.45. For the short-range wake ($s < 8.5\sigma_z$), the sum was truncated at $\omega \sim 300c/\sigma_z$ and for the long-range wake, at $\omega \sim 10c/\sigma_z$. The upper plots show that for Eq. 2.59 treated as an equality with $s_{\max} = 0$, the phase error grows to near unity by the end of the calculation (here $L = 15$mm and $s_{\max} = 250$mm, the sum of which represents $\approx$7 TM_{010} oscillations). Doubling the resolution (lower plots) gets closer to the strict enforcement of Eq. 2.59 and shows much improvement in phase error.

2.4.5 HOM/Wakefield mitigation strategies: waveguide damping

In general, two approaches to eliminating HOMs/wakefields arise in accelerator cavity design: damping and detuning. The first is straightforward; couple all HOMs out of the structure and into

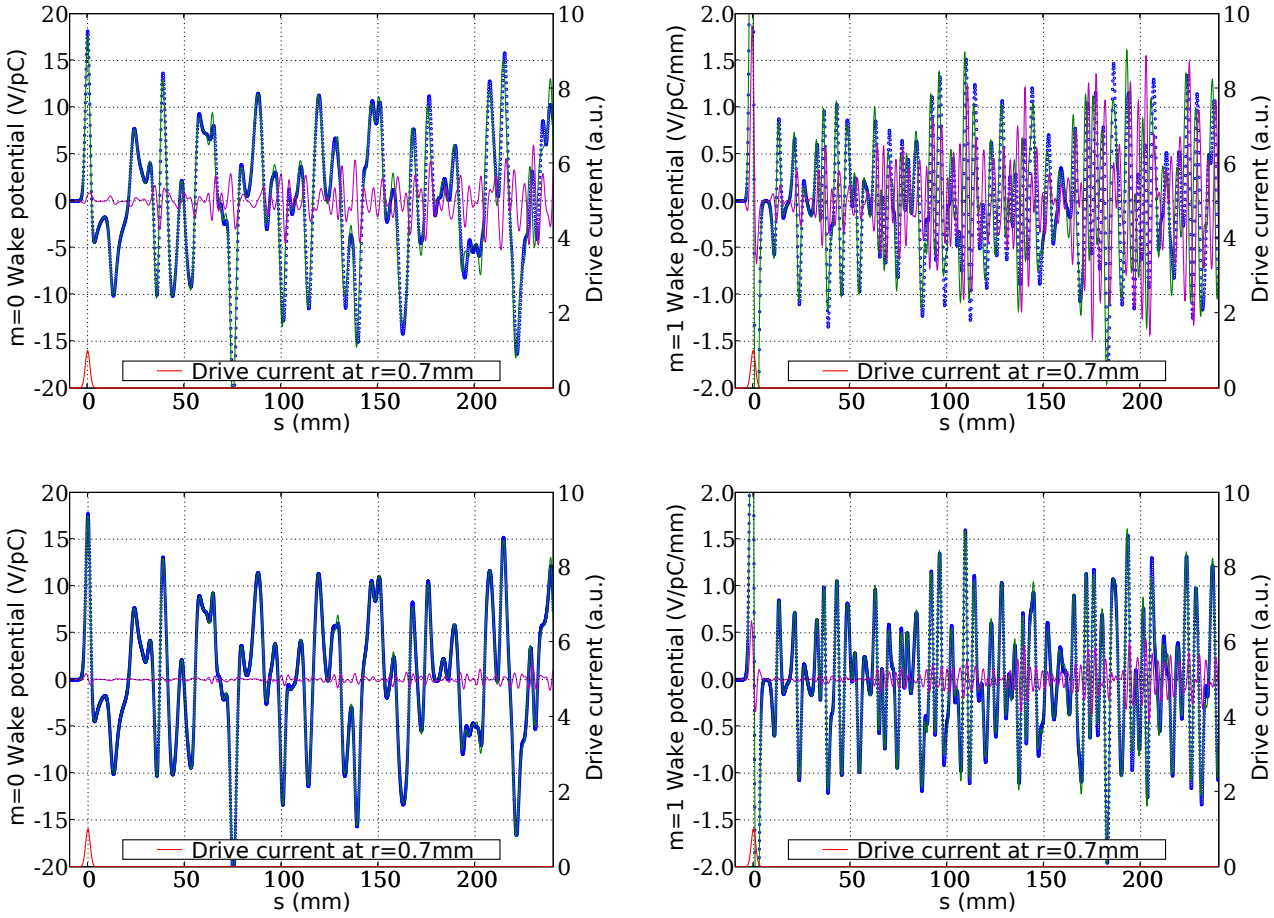

Figure 2.5: *Analytic (green) and simulated (blue circles) wake potentials and their difference (magenta) for a closed pillbox cavity of dimensions* $R = 11.5mm$ *and* $L = 15.0mm$ *due to a bunch of length* $\sigma_z = 2L/25 = 1.2mm$ *at a radial offset of* $0.7mm$*. Subfigures: the monopole and dipole wake potentials are shown in the left- and right-hand columns, respectively (dipole wake potentials are divided by the test particle radius); the upper and lower plots show simulations at* $\Delta z = 0.33mm$ *(the value suggested by Eq. 2.59) and* $\Delta z = 0.17mm$*, respectively.*

absorbers. In the recent CLIC cavity designs, wakefields are strongly damped by waveguides which are coupled directly to the radial cavity walls (see Fig. 2.6). Each waveguide is designed such that its cutoff frequency lies between the accelerating mode frequency and the first dipole HOM frequency. Wakefields in the CLIC cavities will be shown in Chapter 4 where comparisons are drawn with the photonic-crystal-based cavities designed in this thesis. Detuning is the canceling of the effect of one HOM with another (by destructive interference), and is usually an option restricted to multicell cavities. Detuning can lead to dramatic decreases in the wake potential close behind the exciting bunch; however, the recoherence of the wake at larger s can be a problem unless damping is also employed. Also, detuning is most effective when the wake potential is dominated by only a few HOMs, which is generally not the case. Very recent redesigns of the CLIC cavities are attempting to strike an optimal balance between damping and detuning [43].

Since comparisons with the CLIC cavity will eventually be drawn, we focus on the characteristics of waveguide damping (in fact, the concepts presented here will also aid in discussing wakefields within PhC-based cavities). Consider a rectangular waveguide with its axis in the x-direction (perpendicular to the beam axis) and side lengths L_y and L_z (we choose this unconventional coordinate system because the beam axis has already been assigned to the z-direction—c.f. Fig. 2.6). Just as in the pillbox resonant cavity, waveguide modes fall into the TE/TM classification; in the coupled cavity/waveguide system of Fig. 2.6, *cavity* TM modes couple strongly to *waveguide* TE modes (more specifically, to those TE waveguide modes with dominant E_z component).

The frequencies of the TE waveguide modes (i.e. $E_x = 0$) are given by

$$\omega_{mn}(k_x) = c\sqrt{k_x^2 + \frac{\pi^2 m^2}{L_y} + \frac{\pi^2 n^2}{L_z}} \tag{2.60}$$

where k_x is the wavenumber along the waveguide axis (all modes are proportional to $e^{ik_x x}$) and $m = 0, 1, 2, \ldots$ and $n = 1, 2, \ldots$ (or vice versa) [37]. The waveguide mode of most interest is the TE_{10} mode with E_z uniform in the z-direction; this is the waveguide mode that couples strongest to the cavity TM dipole mode, which is the most troublesome HOM in the transverse wakefield.

TE-mode *cutoff* frequencies are given by $\omega_{mn}(k_x \to 0)$. For each mn pair, the associated 2D

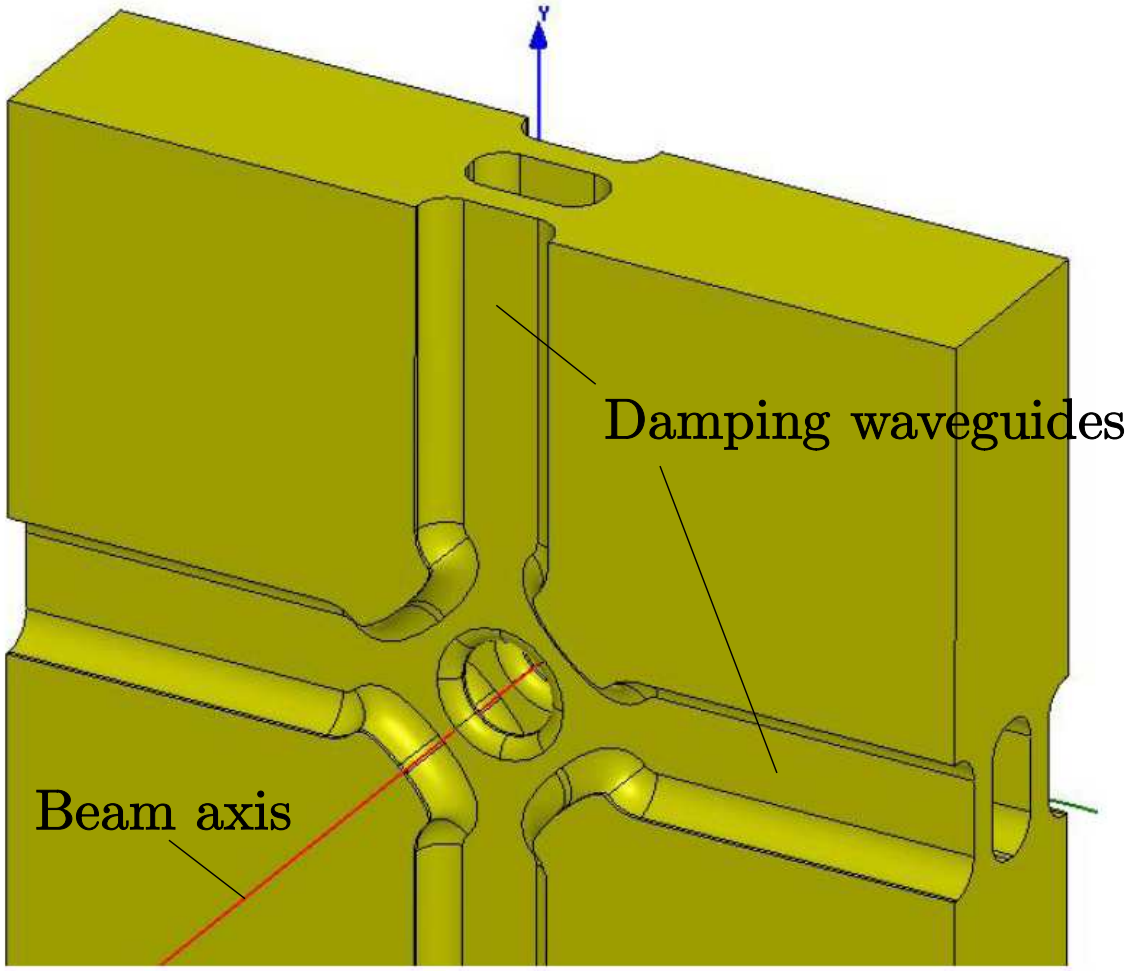

Figure 2.6: *Half (in the z-direction) of an accelerating cell from the CLIC design [27]. The four radial waveguides, which are terminated by electromagnetic absorbers, strongly damp HOMs; the cutoff frequency of each waveguide is just above the accelerating mode frequency.*

field pattern (in the y-z plane) cannot propagate in the waveguide below this frequency. The lowest cutoff frequency in the rectangular waveguide is the TE_{10} cutoff frequency, given by $\omega_{10,\text{cutoff}} = \pi c/L_y$. *No* propagating modes exist in the waveguide below this frequency. In the CLIC cavity, L_y is chosen such that the accelerating mode remains confined to the cavity (because its frequency is below $\omega_{10,\text{cutoff}}$) while maximizing coupling to the lowest dipole mode. For the cavity tested in Chapter 4, $L_y = 11$mm, giving $f_{10,\text{cutoff}} = 13.6$ GHz. Without damping waveguides, the TM dipole resonant frequency for an average CLIC cell is approximately 21 GHz, placing it well above $f_{10,\text{cutoff}}$; therefore, the added waveguides damp this mode effectively. The accelerating mode has a frequency of 12 GHz which is less than $f_{10,\text{cutoff}}$, meaning it must remain confined to the center of the cavity (because it cannot propagate within the waveguides).

Waveguide damping was investigated theoretically using a simple circuit model in Refs. [47, 52] (basically an LC resonator circuit coupled to a transmission line). The relevant parameters in the model are the undamped resonant frequency ω_0 (e.g. the undamped TM cavity dipole frequency) and the waveguide cutoff frequency ω_c (e.g. $\omega_{10,\text{cutoff}}$ for the rectangular waveguide). The response of the model to a delta-function voltage pulse is shown in Fig. 2.7 along with the response of an undamped resonator and a uniformly damped resonator (e.g. an RLC circuit); the responses in Fig. 2.7 can be likened to the impulse wake potential of a given mode in a true cavity-waveguide system. The shape of the resonance in Fig. 2.7a is intuitively expected; below cutoff, the waveguide provides no damping, thus the resonance matches the undamped resonance. Above cutoff, the waveguide provides damping which increases in effectiveness as the drive frequency is increased. Just above cutoff, the damping is ineffective; this is because the waveguide modes at these frequencies have vanishing group velocities, and thus slowly transport energy down the waveguide, away from the cavity. As the gap between ω_c and ω_0 is widened, damping improves. Figure 2.7b shows the amplitude of the time-domain signal (or wake potential); the lingering oscillations are due to the resonance peak at cutoff. This characteristic tail is consistently observed in the wakefield results of this thesis.

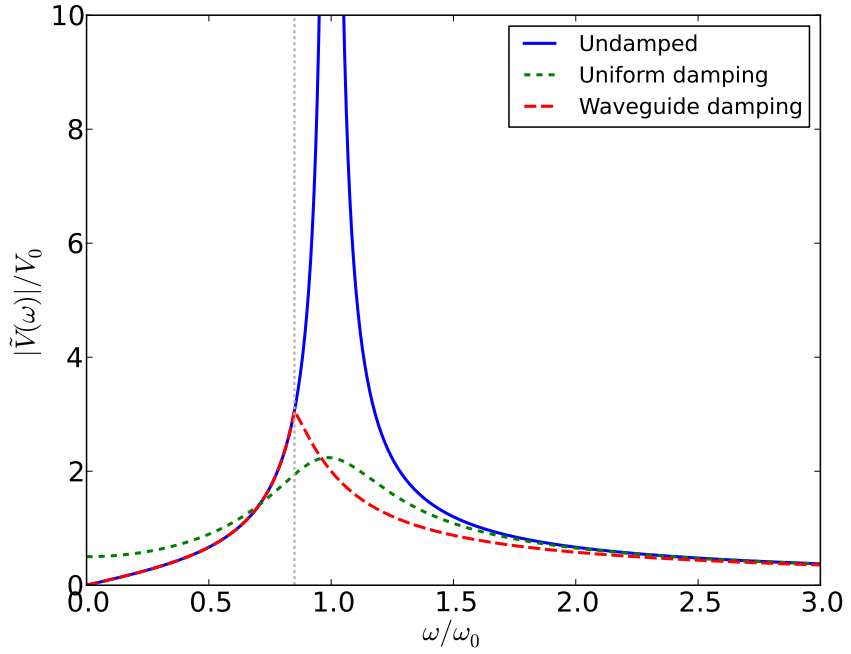

(a)

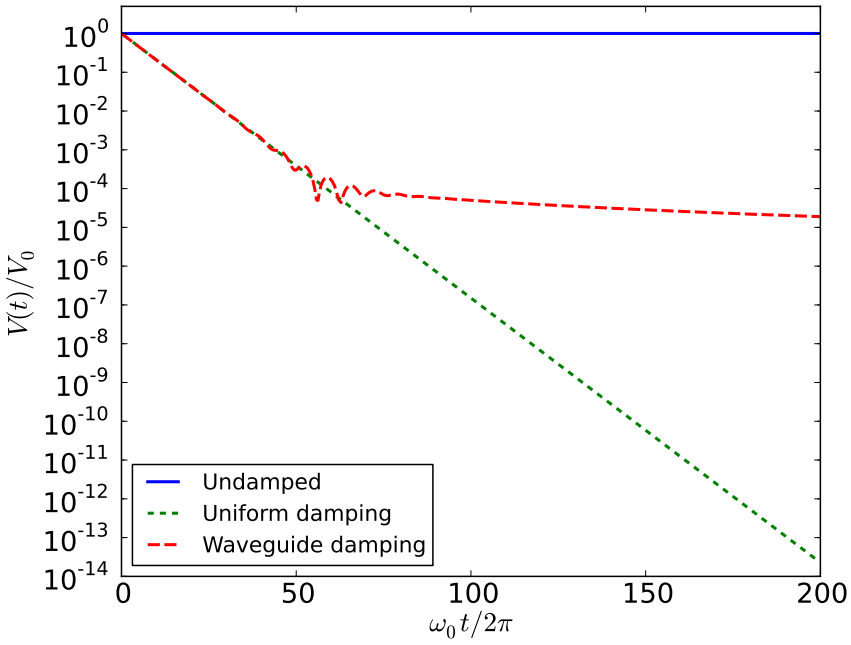

(b)

Figure 2.7: *The simplified (a) Fourier transform and (b) amplitude of the impulse wake potential for a mode with different damping mechanisms (from the circuit models described in the text). The vertical dashed line in (a) shows the location of the cutoff frequency, ω_c/ω_0, below which the waveguide does not provide any damping. The tail of the time-domain waveguide-damped curve in (b) is given by $t^{-3/2}$ decay. The Q of the uniformly damped oscillator was chosen to match the initial exponential decay of the waveguide-damped system in (b).*

2.4.6 Summary

Wakefields should always be minimized in accelerator cavity design since smaller wakefields increase luminosity. Wakefields are reduced by selectively damping HOMs and/or using interference of HOMs to lower their effect at target distances behind a bunch.

2.5 Multicell cavities

In linacs, accelerating cavities like the pillbox discussed above are chained together into multicell structures (or multicell cavities) (see Fig. 2.8). This increases the average acceleration per length of the structure as well as reduces the number of input/output coupling ports required. In multicell cavities, each pillbox-like cell is connected to the next by a very short beam tube, usually termed an *iris*. The size of the iris determines how strongly the individual cells interact or couple; the stronger the coupling, the less each acts like an isolated pillbox. As the number of cells increases, the terminology used in the study of electromagnetic waveguides becomes more relevant. In fact, another way to think of multicell cavities starts with the cylindrical waveguide. Since all traveling wave modes of a cylindrical waveguide have a phase velocity greater than c, a beam would continually slip between accelerating and decelerating phases of each wave and would therefore gain no net energy. To slow a cylindrical waveguide mode, one can introduce periodically-spaced metallic disks with holes bored through their centers (irises) [77].

This section reviews some of the theory of multicell accelerator cavities. First, we discuss the periodic single cell (i.e. infinite chain of identical cells) which introduces important concepts such as Bloch eigenmodes and particle synchronicity. Then, we review some properties of more general multicell cavities that are allowed to change their cell-to-cell geometries. Figures of merit for multicell cavities will be introduced.

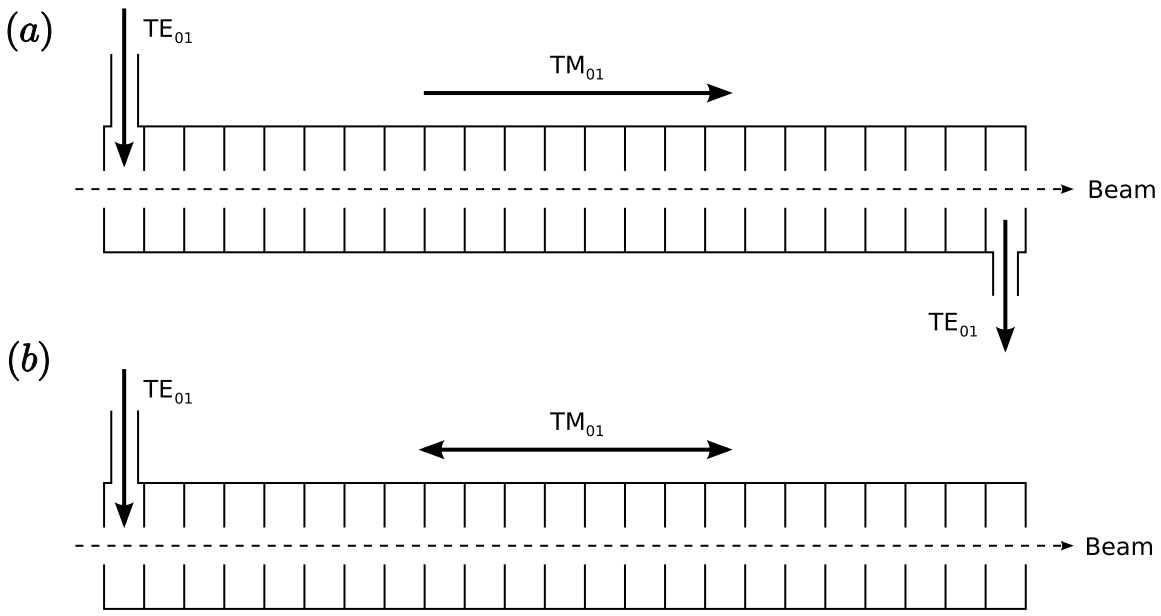

Figure 2.8: *Generic multicell accelerating cavities; (a) traveling-wave cavity and (b) standing-wave cavity. (Figure based on Ref. [83] Fig. 5.9).*

2.5.1 Modes in a periodic single cell

In the simplest case, a multicell cavity consists of an infinite chain of identical cells. This case is simple because discrete translational symmetry reduces the analysis of the entire structure to a single cell. For a periodic single cell of length L, the electromagnetic eigenmodes take the following form by the Bloch-Floquet theorem,

$$\mathbf{E}_n(k, \mathbf{x}, t) = \mathbf{E}_n(k, \mathbf{r}, z)e^{i(kz-\omega_n(k)t)} \tag{2.61}$$

where $-\pi/L < k < \pi/L$ is a continuous parameter related to (but *not* equal to) the wavenumber in the z direction, n is called the *band index*, and $\mathbf{E}_n(k, \mathbf{r}, z)$ is a complex function that is periodic in z, i.e.

$$\mathbf{E}_n(k, \mathbf{r}, z + L) = \mathbf{E}_n(k, \mathbf{r}, z). \tag{2.62}$$

(A more complete treatment of Bloch theory is presented in Section 3.1.2 in which photonic crystals are introduced.)

The parameter k is rarely used in accelerator literature; it is more common to speak of the *phase advance* (per cell), which is simply

$$\phi = kL \tag{2.63}$$

and runs from $-\pi$ to π (since accelerating cells are usually mirror symmetric in z, ϕ is further restricted to $[0, \pi]$ to eliminate redundancies). The name "phase advance per cell" is evident given that

$$\mathbf{E}(\phi, \mathbf{r}, z = L, t) = e^{i\phi}\mathbf{E}(\phi, \mathbf{r}, z = 0, t). \tag{2.64}$$

We will henceforth prefer ϕ in place of k.

The dispersion of these eigenmodes is described by $\omega(\phi)$. The phase velocity is $v_{\text{ph}} = \omega(\phi)L/\phi$ and the group velocity is

$$v_g(\phi) = L\frac{\partial\omega(\phi)}{\partial\phi}. \tag{2.65}$$

For typical cavity cells of interest in this work (e.g. CLIC cells), the group velocities are low: $v_g/c \sim 1$ % (that is, for the accelerating mode); this is an indication of weak intercell coupling. In

this case, the field pattern in any given cell is similar for all possible phase advances.

For a finite chain of identical cells (with closed end cells), the standing-wave solutions resemble the linear combinations of forward- and backward-traveling periodic modes. Roughly, they can be described by an amplitude (of a fixed field pattern; e.g. TM_{010}-like for the accelerating mode) in each cell $A_{p,q}$ where p is the mode number (applying to a particular phase advance), and q is the cell number. The modes then take the form

$$A_{p,q} \sim \cos\left(\frac{pq\pi}{N-1}\right) \tag{2.66}$$

where $0 \leq p, q \leq N$ [77].

2.5.2 Voltage gain in a periodic cell

Before detailing eigenmodes in periodic structures any further, a general analysis of the interaction between a relativistic particle and an eigenmode of a periodic cell can be performed based on the general Bloch solutions. The results (which follow) show that, on average, a relativistic particle interacts with only those modes that have a phase velocity equal to the speed of light. This seems intuitive, since, given enough time, any other phase velocity would cause a relativistic particle to experience both accelerating and decelerating phases of the wave.

In the context of accelerators, it can be useful to Fourier-expand the periodic part in Eq. 2.61, giving

$$\mathbf{E}_n(\phi, \mathbf{x}, t) = \sum_{l=-\infty}^{\infty} \mathbf{E}_{n,l}(\phi, \mathbf{r}) e^{i[2\pi l z/L + \phi z/L - \omega_n(\phi) t]}. \tag{2.67}$$

The field associated with each index l is called a *space harmonic*. This construction is useful because speed-of-light particles tend to interact with only one space harmonic, as we now demonstrate.

The voltage gain of a speed-of-light charged particle interacting with a Bloch eigenmode (Eq. 2.61) at a transverse position $\mathbf{r}$ is

$$V_n(\phi, \mathbf{r}) = \int_{-\infty}^{\infty} E_{n,z}\left(\phi, \mathbf{x}, t = z/c\right)\, dz. \tag{2.68}$$

Recall that the above is a complex quantity, where only the real part is physical; the phase position of the particle is crucial, i.e., the maximum voltage gain occurs when $\mathrm{Re}\, V_n(\phi, \mathbf{r}) = |V_n(\phi, \mathbf{r})|$. Using

the space harmonic decomposition of Eq. 2.67 in the above

$$V_n(\phi, \mathbf{r}) = \sum_{l=-\infty}^{\infty} E_{n,l,z}(\phi, \mathbf{r}) \int_{-\infty}^{\infty} e^{i[2\pi l + \phi - \omega_n(k)L/c]z/L}\, dz \tag{2.69}$$

$$= E_{n,l,z}(\phi, \mathbf{r}) L \delta\left[2\pi l + \phi - \omega_n(\phi)L/c\right] \tag{2.70}$$

and we see that the only nonzero contribution to the voltage is due to the *synchronous* mode; i.e. when $2\pi l + \phi - \omega_n(\phi)L/c = 0$ for some integer l.

2.5.3 Figures of merit

Figures of merit in multicell cavities are normalized to the cell length when appropriate. For instance, the shunt impedance R_{shunt}, scales with the number of cells (since the voltage gain V and the power loss P scale with length NL); thus, a shunt impedance *per unit length* is defined as

$$r_{\text{shunt}} = \frac{R_{\text{shunt}}}{L} \tag{2.71}$$

where R_{shunt} in the above is the shunt impedance for a single cell calculated in the usual way. Similarly, the loss factor (and wake potential) scale with length. The loss factor per unit length will be defined as

$$\hat{k} = k/L. \tag{2.72}$$

2.5.4 Multicell cavity design basics

To further detail the properties of modes in periodic single cells, we turn to the periodic pillbox example. Though simple, the periodic pillbox can still be used to describe the essential features of most modern-day cavities (in fact, the accelerating cavities of the SLAC linac, still in operation, are multicell pillbox structures). Figure 2.9 shows two cells of a periodic pillbox cavity. The geometry is parameterized by: the iris radius a, the iris thickness d, the cavity radius R, and the cell length L. In general, the free parameters in the design of the periodic pillbox are the above geometry descriptors and the phase advance of the accelerating mode ϕ.

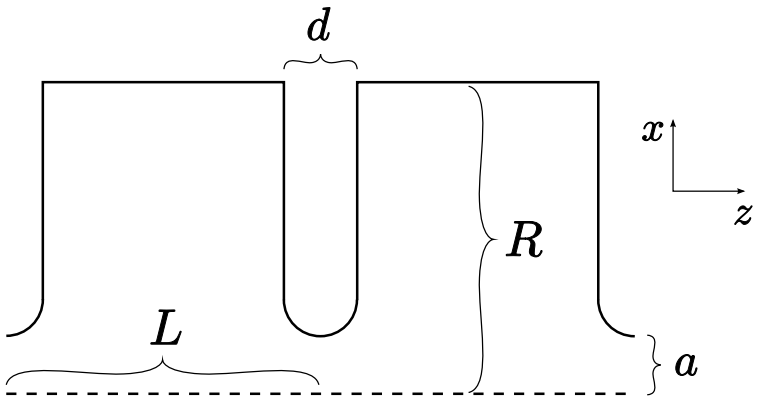

Figure 2.9: *Dimensions of a general periodic pillbox cavity. The structure is cylindrically symmetric about the dotted line.*

ϕ	L	Q	r_{shunt}	r_{shunt}/Q
$\pi/2$	6.2 mm	5300	$8.9 \times 10^7\ \Omega/m$	$1.7 \times 10^4\ \Omega/m$
$2\pi/3$	8.3 mm	6800	$1.1 \times 10^8\ \Omega/m$	$1.6 \times 10^4\ \Omega/m$
π	12.5 mm	8800	$9.8 \times 10^7\ \Omega/m$	$1.1 \times 10^4\ \Omega/m$

Table 2.2: *Figures of merit for the periodic pillbox cavity at three different phase advances. For each phase advance, the cavity length was adjusted to ensure synchronicity with a speed of light particle. The frequency in each case was 12 GHz.*

The first constraint (in no particular order) on the pillbox parameters is that the phase velocity of the accelerating mode must be synchronous with a particle moving at the speed of light; i.e. from the previous section, the following condition must be satisfied

$$2\pi l + \phi - \omega L/c = 0 \tag{2.73}$$

for some integer l (usually zero). Since the accelerating mode frequency is fixed by the overall accelerator design, the cell length L and phase advance ϕ are therefore linked by the above synchronicity condition. Figure 2.10 shows the z-component of the electric field of the accelerating mode for three different values of the phase advance; in each case, the cell length is adjusted to satisfy Eq. 2.73. Figures of merit for the periodic pillbox at these different phase advances are shown in Table 2.2. Differences in the shunt impedance are minor, yet maximized for a phase advance of $2\pi/3$; this is the phase advance most commonly chosen in current designs for CLIC.

The next constraint is the operating frequency. The frequency of the accelerating mode in the periodic pillbox is strongly affected by the cavity radius and only slightly affected by the iris dimensions, phase advance, and cell length. When changing any of these parameters, the others must be adjusted such that the operating frequency is maintained. The following perturbation theory provides an approximate picture of the effects of these parameters.

For small iris radii, the accelerating mode of the periodic pillbox can be treated as a perturbation of the closed pillbox cavity TM_{010} mode (actually, this approximation works well for practical iris radii as well). Based on Bethe's theory of small apertures in conducting walls (where small means the radius of the aperture is less than the wavelength of the mode), the perturbation

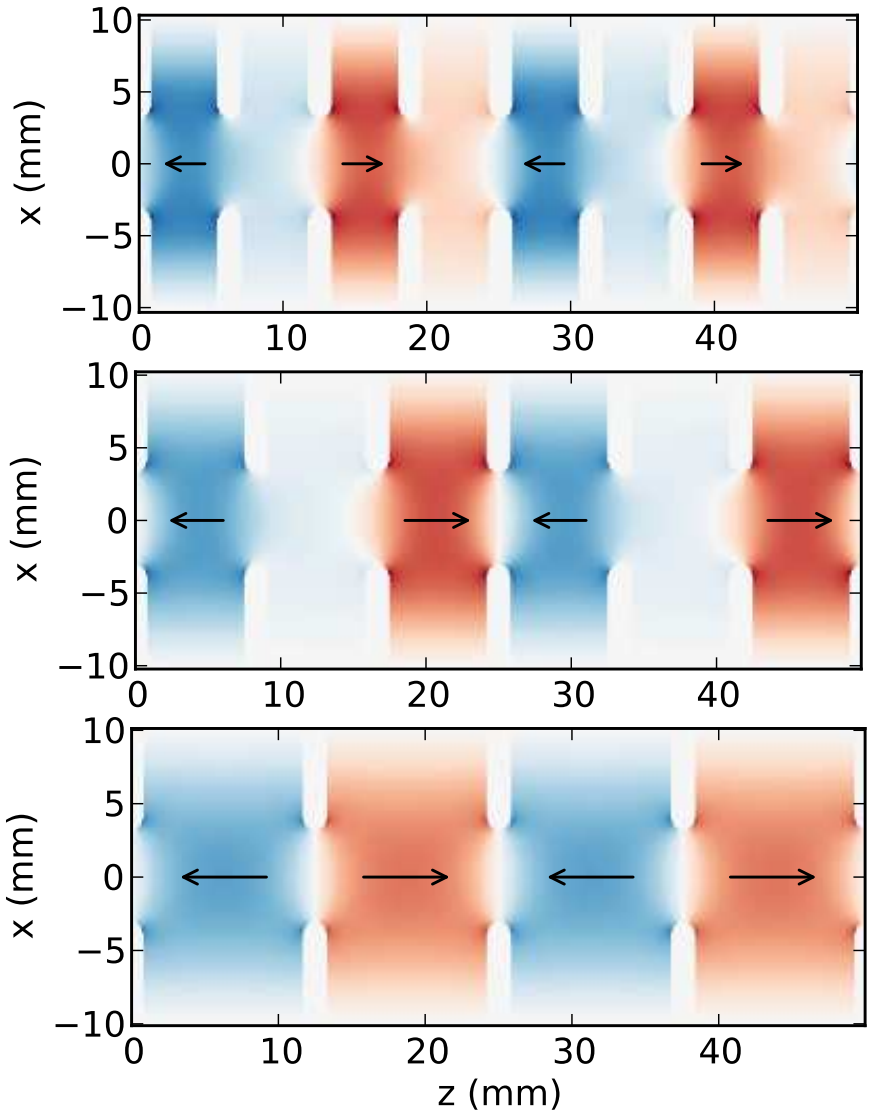

Figure 2.10: *Accelerating mode E_z for three different phase advances (from top to bottom: $\phi = \pi/2$, $2\pi/3$, and π). The frequency of each mode is 12 GHz. The lengths of the individual cells were adjusted to satisfy the synchronicity condition. The arrows indicate the on-axis field direction in each cell. The field patterns move to the right with time.*

of the pillbox TM_{010} resonant frequency due to the iris is [6, 12, 24]

$$\frac{\omega - \omega_p}{\omega_p} \approx \frac{1}{3}\frac{a^3 \varepsilon_0 E_0^2}{U}\left(1 - e^{-\alpha d}\cos\phi\right). \tag{2.74}$$

where ω is the new frequency, $\omega_p = 2.405c/R$ is the unperturbed pillbox frequency, U is the stored energy of the unperturbed mode, E_0 is the on-axis unperturbed electric field, and α is defined as

$$\alpha = \sqrt{\frac{2.405^2}{a^2} - \frac{\omega_p^2}{c^2}}. \tag{2.75}$$

The above decay factor is based on the cutoff frequency of the TM_{01} cylindrical waveguide mode for a waveguide of radius a (i.e. the iris). Some details leading to the above perturbation expressions are found in Appendix A.

A more useful form of Eq. 2.74 is achieved by noting that the energy of the unperturbed TM_{010} mode scales as $U \sim \lambda_p^2 L = c^2 L/\omega_p^2$. Thus, we can rewrite Eq. 2.74 as

$$\frac{\omega - \omega_p}{\omega_p} = \frac{1}{3}\frac{\omega_p^2 a^3}{c^2 L}\Gamma\left(1 - e^{-\alpha d}\cos\phi\right). \tag{2.76}$$

where $\Gamma = \frac{c^2 L \varepsilon_0 E_0^2}{\omega^2 U}$ is a dimensionless constant (related to the loss factor) that is determined purely by the TM_{010} field pattern. Furthermore, since the TM_{010} fields in the pillbox are independent of z, Γ can be written

$$\Gamma = \frac{4c^2 \varepsilon_0 E_0^2}{\omega^2 \int (\varepsilon |\mathbf{E}|^2 + \mu |\mathbf{H}|^2)\, dx\, dy}. \tag{2.77}$$

The above expression is now general for any 2D cavity (or 3D waveguide) with a TM_{010}-like mode (since any waveguide becomes a resonant cavity by adding conducting planes with normals along the waveguide axis); thus, the above perturbation theory can be applied to more complex transverse geometries, such as 2D photonic crystal cavities.

Given that ω is fixed, Eq. 2.74 represents (approximately) the second constraint on the periodic pillbox parameters. Since the cavity radius is mostly inconsequential (in terms of the usual figures of merit—shunt impedance, wakefields, etc.) and has the greatest effect on ω, it is usually considered the parameter set by the choice of operating frequency. Equation 2.74 also shows that the opening of the irises introduces a dispersion relation of the form

$$\omega(\phi) \approx \omega(\pi/2) - \frac{v_g(\pi/2)}{L}\cos\phi \tag{2.78}$$

with group velocity

$$v_g(\phi) = L\frac{\partial\omega}{\partial\phi} \approx v_g(\pi/2)\sin\phi. \tag{2.79}$$

Relating Eq. 2.79 to the ϕ-derivative of 2.76 shows that the group velocity is very sensitive to both the iris radius and the iris thickness.

Figure 2.11 shows the dispersion of the accelerating mode for two different periodic pillbox iris radii: $a = 3.15$ mm and $a = 6.3$ mm (iris thickness was 1 mm). The target TM_{010} frequency was 11.994 GHz and the target phase advance was $\phi = 2\pi/3$ (CLIC cavity values). First, L was set by the synchronicity condition, Eq. 2.73. Then, the cavity radius R was calculated using perturbation theory (Eq. 2.76) to achieve the target frequency; i.e. $\omega(\phi = 2\pi/3) = 2\pi 11.994$ GHz. For the smaller iris, perturbation theory produced the target frequency to $\sim$1%. For the larger radius, perturbation theory failed considerably. Even so, the solid lines of Fig. 2.11 show that the cosine dependence of the dispersion is still quite accurate. Therefore, once synchronicity and the target frequency are achieved for the target phase advance, the dispersion curve (and therefore the group velocity, bandwidth, etc.) can be calculated with confidence from only one additional simulation (at a different phase advance). This is how we calculate group velocities in the analysis of Chapter 4.

The remaining free parameters are the iris geometry and the phase advance. These are the parameters that can be varied to optimize certain figures of merit. For example, the iris radius can be increased for the purpose of reducing short-range wakefields, but only at the expense of shunt impedance (accelerating efficiency). The iris thickness can also be adjusted for the desired group velocity. These parameters will become more relevant in the following.

2.5.5 Cell-to-cell variations

As discussed in the following sections, real multicell cavities are rarely an array of identical cells; for example, the iris radius is often tapered to increase acceleration efficiency. Without translational symmetry, a periodic single-cell cannot accurately characterize the entire structure. In the latest designs of the CLIC cavities, the number of cells per structure is 26. While the simulation

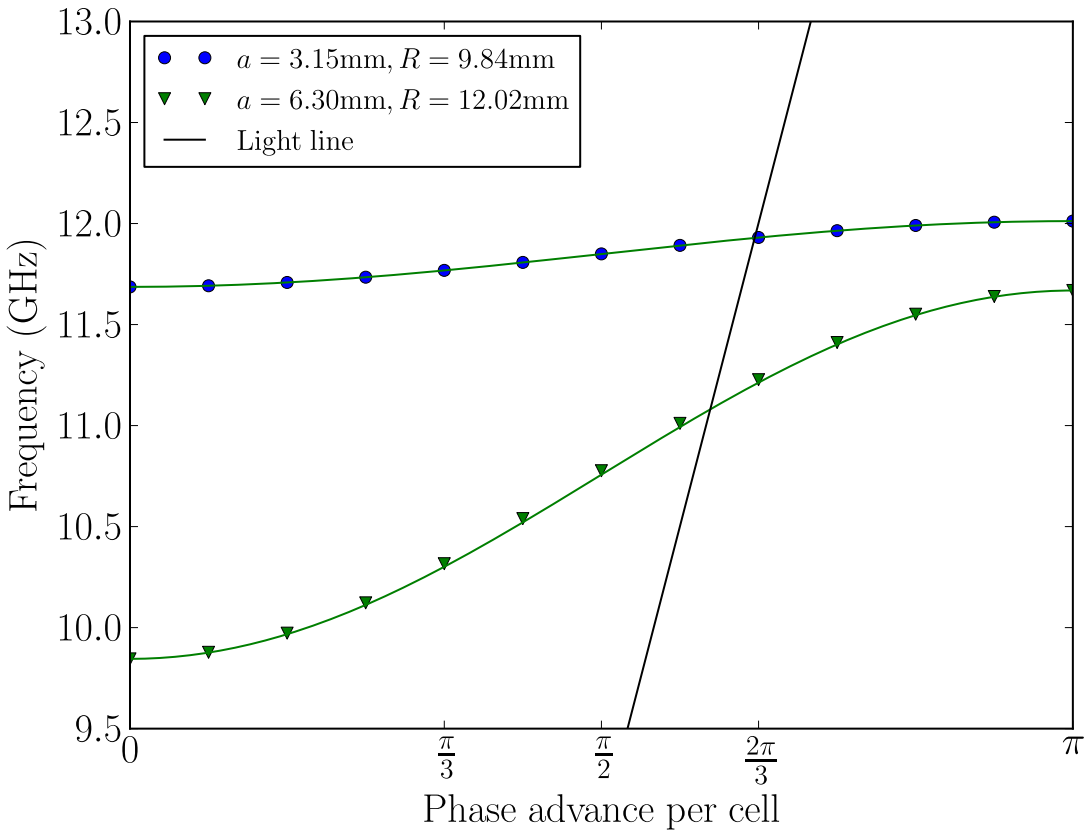

Figure 2.11: *These dispersion curves were generated from simulations of two pillbox cavities; for each, the cavity radius R was determined by perturbation theory given the iris radius a and target frequency $f(\phi = 2\pi/3) = 11.994GHz$. Symbols are frequencies calculated from simulations, the overlaid curves are best-fit cosine functions, and the light line represents the synchronicity condition (modes that lie on the light line accelerate relativistic particles, since their phase velocity equals c). The plot illustrates several concepts: (1) for $a = 6.3mm$, perturbation theory failed, i.e., $f(\phi = 2\pi/3) \neq 11.994GHz$; (2) even though perturbation theory failed for the large iris radius, the cosine dependence was well-preserved; (3) as the iris radius increases, coupling increases, and the frequency range widens (increasing v_g).*

of a single 26-cell cavity is computationally feasible, an optimization study requiring the simulation of many different 26-cell structures becomes computationally prohibitive. Wakefield/HOM simulations make this especially difficult due to mesh resolution constraints (c.f. Section 2.4.4). Fortunately, for the class of accelerating cell considered in this thesis (that is, cells separated by conducting irises), the coupling between cells is weak, and many quantities of interest can be approximated using uncoupled techniques.

Uncoupled methods do not treat a multicell cavity in its entirety, but rather approximate the cavity as a chain of single cells whose properties are determined by periodic single-cell calculations (at the synchronous phase advance for the mode in question). This is expected to be a good approximation because (1) for slowly varying geometries along the multicell cavity, the fields in neighboring irises are well-approximated by a periodic single cell at the average of the iris geometries and (2) coupling introduces frequency shifts that are on the order of the frequency variation in the mode's band, which is small relative to the center-band frequency. The next section describes an uncoupled calculation important to characterizing the accelerating mode in a multicell cavity.

2.5.6 Power flow in traveling-wave multicell cavities

In the following, all quantities are assumed to apply to the accelerating mode in the steady state. For cell number j, define an incoming power $P_{j,\text{in}}$ (power from cell $j-1$) and an outgoing power $P_{j,\text{out}}$ (power to cell $j+1$). The difference is the power lost in that cell due to heating of the cavity walls and the acceleration of a particle beam, i.e.

$$P_{j,\text{in}} - P_{j,\text{out}} = P_{j,\text{wall}} + P_{j,\text{beam}} \tag{2.80}$$

$$= \frac{\omega U_j}{Q_j} + IV_j \tag{2.81}$$

$$= \frac{\omega U_j}{Q_j} + 2I\sqrt{k_j U_j} \tag{2.82}$$

where U_j is the energy, V_j is the accelerating voltage, Q_j is the quality factor, and k_j is the loss factor in cell j (Eq. 2.41 was used to obtain the final expression). Since beams come in bunches, the beam current looks like: $I = \nu q$ where ν is the bunch frequency (i.e. number of bunches passing

a given point per second) and q is the charge of a single bunch.

Dividing Eq. 2.82 by the cell length L and taking the limit $N \gg 1$, we get the following differential equation for the power flow along the cavity as a function of cell number (or simply position z):

$$\frac{dP(z)}{dz} = -\frac{\omega u(z)}{Q(z)} - 2I\sqrt{k(z)u(z)} \tag{2.83}$$

where we have defined a spatially-dependent energy per unit length $u(z) \approx U_j/L$, loss factor per unit length $k(z) \approx k_j/L$, and quality factor $Q(z) \approx Q_j$. The linear energy density can be expressed in terms of the power flow via

$$u(z) = \frac{P(z)}{v_g(z)} \tag{2.84}$$

where $v_g(z)$ is the group velocity. Equation 2.83 becomes

$$\frac{dP(z)}{dz} = -\frac{\omega}{v_g(z)Q(z)}P(z) - 2I\sqrt{\frac{k(z)}{v_g(z)}}\sqrt{P(z)}. \tag{2.85}$$

The cavity parameters $k(z)$, $v_g(z)$, and $Q(z)$ are determined by periodic single-cell simulations at frequency ω and the target phase advance for cell geometries equal to those at position z in the multicell cavity.

For multicell cavities of identical cells (usually termed *constant impedance* cavities) it is well-known that (for small beam currents) the power decays exponentially along the structure (as is evident from the form of Eq. 2.85 for small enough I). In this case, the accelerating gradient $E(z) = 2\sqrt{kP(z)/v_g}$ also decays exponentially, which means that the constant impedance structure is a poor use of linac real estate. Ideally, the accelerating gradient would be constant along the cavity, at a level slightly below the breakdown threshold of each cell; indeed, this is the type of cavity used in high-energy accelerators today (called a *constant gradient* cavity).

A constant gradient is achieved by decreasing the group velocity along the length of the cavity, so that electromagnetic energy has more time to build up in later cells (where the power flow is weaker) before "leaking" into subsequent cells. Assuming a constant Q-factor and loss factor per unit length, (approximately true in most cases), the following is an extremely simple,

constant-gradient solution to Eq. 2.85:

$$P(z) = P_{\text{in}} - \frac{P_{\text{in}} - P_{\text{out}}}{NL} z \tag{2.86}$$

$$v_g(z) = v_{g,\text{in}} - \frac{v_{g,\text{in}} - v_{g,\text{out}}}{NL} z \tag{2.87}$$

where P_{in} is the power put into the first cell of the cavity, P_{out} is the power removed from the last cell, and $v_{g,\text{in}}$ and $v_{g,\text{out}}$ are the group velocities of the first and last cells, respectively. The constant electric field is given by

$$E = 2\sqrt{kP(z)/v_g(z)} \tag{2.88}$$

which gives the constraint

$$\frac{P_{\text{in}}}{v_{g,\text{in}}} = \frac{P_{\text{out}}}{v_{g,\text{out}}} = \frac{P(z)}{v_g(z)}. \tag{2.89}$$

Plugging the solution into Eq. 2.85 gives

$$IV = \eta_{\text{cg}} P_{\text{in}} \tag{2.90}$$

where $V = NLE$ (ultimately a function of P_{in} and $v_{g,\text{in}}$) and η_{cg} is the constant-gradient steady-state efficiency:

$$\eta_{\text{cg}} = \frac{\Delta v_g}{v_{g,\text{in}}} - \frac{\phi N}{Q} \frac{c}{v_{g,\text{in}}} \tag{2.91}$$

where $\Delta v_g = v_{g,\text{in}} - v_{g,\text{out}}$ and ϕ is the phase advance per cell. The first term describes a baseline efficiency due to the power exiting the cavity at its final cell; obviously, this is wasted power (it is sent to an absorbing load). However, it is necessary, since reducing the group velocity anywhere along the structure increases the *fill time* of the cavity (that is, the time between the input of RF energy and the injection of the beam)—the fill time is unaccounted for in the above steady-state analysis. The second term is the hit in efficiency given by the cavity material losses.

The above constant gradient solution gives only a rough idea of the behavior of a more realistic cavity. Two major complications arise. First, the constant gradient structure described above assumed a steady-state with beam loading; however, during the filling of the cavity, a beam is not present. Thus, the fields in the structure will be higher than predicted by the above, and the

cavity will be more likely to break down. Knowledge of both the loaded (with beam) *and* unloaded (without beam) gradient as a function of cell number for a given structure is required. Second, the loss factor per unit length and the Q-factor change with cell geometry. These complications necessitate the direct numerical solution of Eq. 2.85. Regardless, it is clear that a multicell cavity should be designed with a *tapering* group velocity.

2.5.7 Wakefields

Wakefield simulations require the calculation of frequencies up to the bunch cutoff frequency, which can be several times that of the fundamental. For these HOMs, weak coupling may no longer apply. Even so, an uncoupled calculation can be performed to give a rough idea of the wake potential for a chain of slowly varying cells. It was shown in [3] that for the lowest dipole modes of a multicell pillbox cavity, the uncoupled calculation can be accurate for many accelerating mode periods behind an exciting bunch. Examining the highest group velocity in the wakefield frequency range of interest can give an estimate of a trailing distance $s_{\max}$ below which the wake potential should be accurate in an uncoupled calculation. For a HOM excited in one cell, the time it takes for the next cell to feel its effects is on the order of L/v_g where v_g is the group velocity of the HOM; thus $s_{\max} \sim Lc/v_{g,\max}$ where $v_{g,\max}$ is the highest group velocity within the HOMs of interest (in fact, to the advantage of the method, Ref. [3] showed this to be a large underestimate).

In the uncoupled wakefield calculation, the net wake potential is determined by an average over the wake potentials in each cell, where each cell is treated as periodic. The net wake potential per unit length is

$$W_z(s, \mathbf{r}, \mathbf{r}') = \frac{1}{N} \sum_i^N W_z^{(i)}(s, \mathbf{r}, \mathbf{r}') \tag{2.92}$$

where $W_z^{(i)}(s, \mathbf{r}, \mathbf{r}')$ is the wake potential per unit length for a periodic single cell with the geometry of cell i. The above method is commonly used to characterize new multicell cavity designs since the periodic single cell calculations need be performed only once (also, results can often be interpolated between very similar geometries).

To calculate the wake potential for a periodic single cell $W_z^{(i)}(s, \mathbf{r}, \mathbf{r}')$ to be used in Eq. 2.92, either the Condon method or time-domain method can be used. In the Condon method, the loss factors of all *synchronous* modes need to be calculated up to the bunch cutoff frequency. Then, $W_z^{(i)}(s, \mathbf{r}, \mathbf{r}')$ for cell i is reconstructed from the impulse wake potentials. In the time-domain method, with sufficient computational resources, a multicell cavity of identical cells (with the geometry of cell i) can be simulated to approximate the wake potential in the periodic case. As the number of cells increases, the effect of the end cells diminishes [31]. Figure 2.12 shows the asymptotic effect on the wake potential of increasing the number of simulated cells in the case of the pillbox. In all simulations, particles are emitted from and absorbed by conducting plates in the beam tubes of the end cells.

2.6 Chapter segue

In the next two chapters, we introduce PhC accelerator cavities and compare our designs against the latest CLIC copper cavity. The following concepts from this chapter guide the analysis. With normal conductors, the surface magnetic field seems to be the primary indicator of E_{acc} breakdown limits (based on recent empirical evidence); since the PhC-based cavity designs use conductors, the ratio $cB_{\text{surf,metal}}/E_{\text{acc}}$ will be an important point of comparison. The shunt impedance is the major indicator of accelerating efficiency; by using dielectrics, electromagnetic losses to conducting surfaces may be reduced, which could raise the shunt impedance. We will also directly calculate the steady-state accelerating efficiency in constant-gradient multicell PhC cavities by solving the power-flow equation numerically (Eq. 2.85). Finally, wakefields will be calculated for a representative periodic single cell using time-domain simulation (c.f. Section 2.5.7, third paragraph). The theory of waveguide damping (Section 2.4.5) will aid the explanation of the unexpected higher wakefields in the PhC-based cavities. Onward!

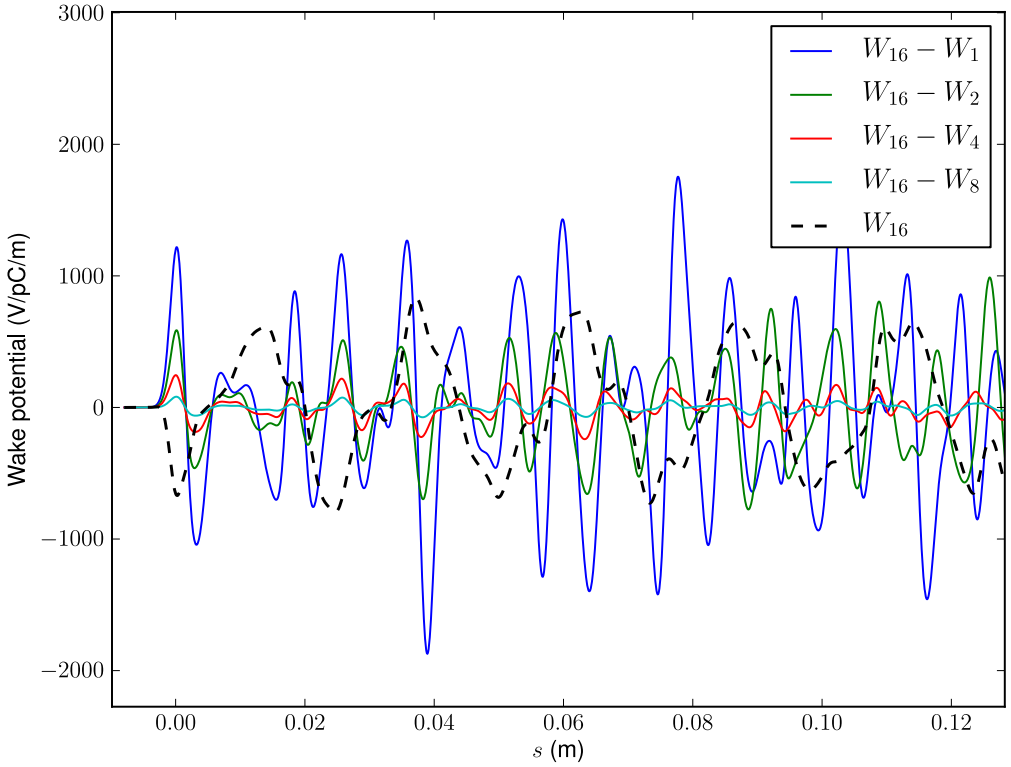

Figure 2.12: *A comparison between the monopole wake potentials in several N-cell pillbox cavities where each cell is identical. W_N indicates the wake potential per unit length for a structure consisting of N pillbox cells. Cell geometry was: $a = 3.15mm$, $d = 1.67mm$, $L = 8.33mm$, and $R = 9.98mm$. The Gaussian excitation bunch length was $\sigma_z = 1.0mm$ located at a transverse offset of 1mm.*

Chapter 3

Photonic crystal accelerating structures

In traditional pillbox-based (super)conducting accelerator cavity design, unwanted HOMs are removed by strategically-coupled waveguides. Photonic crystals (PhCs) offer an alternative approach. PhCs can act as frequency-selective reflectors; and by "selecting" the accelerating mode frequency, a PhC-based resonant cavity can confine the accelerating mode while allowing other modes to propagate out of the structure. In addition, this property of PhCs enables mode-confinement using only dielectric materials, which, as discussed in the previous chapter, may sustain higher surface fields and suffer lower electromagnetic losses.

This Chapter begins by introducing the generic properties of PhCs, shifting focus quickly to frequency-selection and accelerator cavity design. After presenting the well-known defect-based PhC resonant cavity, we discuss alternatives based on liberal structure optimizations. Specifically, using an optimization process that ignored lattice symmetry (but retained certain rotational symmetries), we have discovered that dramatic improvements (orders of magnitude) in the radiative Q-factor of the accelerating mode are possible. This work has resulted in three different PhC-based cavity designs that are the subject of numerical scrutiny in the following Chapter (one lattice-based PhC cavity and two PhC-based optimized cavities). We finish this Chapter with some theoretical predictions of the overall performance of these cavities.

3.1 Photonic crystals background

In general, photonic crystals (PhCs) are periodic arrays of electromagnetic wave scatterers. Under this slightly redundant definition, most PhCs are rather uninteresting; however, a subset of them has attracted interest in recent years due to the discovery of *photonic bandgaps* (or just bandgaps). A photonic bandgap is, in general, a range of frequencies over which no propagating electromagnetic waves can exist. Thus, for a PhC with a bandgap, no electromagnetic waves oscillating at frequencies within that bandgap can propagate through the structure (from this point on, a "PhC" will be assumed to have a bandgap unless explicitly stated otherwise). This property is most often exploited in the form of a frequency-selective reflector; for example, a semi-infinite PhC will necessarily reflect impinging electromagnetic waves whose frequencies lie within the bandgap (since they cannot propagate inside the crystal), but will transmit all others (lossless materials are assumed).

3.1.1 Dielectric mirror

The canonical example of a PhC is the dielectric mirror, which is simply a stack of plates of alternating dielectric constant. To reflect a wave with frequency ω (and vacuum wavelength λ), each plate is given the thickness $d = \lambda/4\sqrt{\varepsilon}$, where ε is the relative dielectric constant of the plate. This thickness causes destructive interference between the transmitted and reflected parts of the wave at each interface (adjacent plates must have different dielectric constants to produce such partial reflections). By this argument, one would expect that only a single frequency would be reflected off of the stack when, in fact, a bandgap forms and reflects a range of frequencies (of course, this is only true in the semi-infinite case—a finite-sized stack will indeed reflect the targetted frequency best, but not perfectly).

The dielectric mirror is an example of a 1D PhC because there is only one dimension of periodicity. A 1D PhC does not reflect frequency ω over all electromagnetic wave propagation directions (no variation in material properties along the direction of propagation means no wave

reflection). The dependence of the reflection on propagation direction is determined by examining the *band diagram* for the PhC. The most interesting case is that of a *complete* bandgap—one which exists for all electromagnetic wave propagation directions (and polarizations in the case of electromagnetism).

3.1.2 Band diagrams and bandgaps

[2] In a periodic medium, translation by a *lattice vector* $\mathbf{R}$ leaves the equations of motion unchanged (e.g. if a periodic dielectric medium has the property $\varepsilon(\mathbf{x}+\mathbf{R}) = \varepsilon(\mathbf{x})$, then Maxwell's wave equation is unchanged after translation by a lattice vector). Bloch's Theorem (applied to electromagnetism) says that the solutions to Maxwell's wave equation in a periodic medium can be written as

$$\mathbf{E}(\mathbf{x},\mathbf{k}) = \tilde{\mathbf{E}}(\mathbf{x},\mathbf{k})e^{i\mathbf{k}\cdot\mathbf{x}} \tag{3.1}$$

where $\tilde{\mathbf{E}}$ has the periodicity of the medium (i.e. $\tilde{\mathbf{E}}(\mathbf{x}+\mathbf{R},\mathbf{k}) = \tilde{\mathbf{E}}(\mathbf{x},\mathbf{k})$ for any lattice vector $\mathbf{R}$). There exists a unique *reciprocal* lattice vector, $\mathbf{K}_i$, such that $\mathbf{K}_i \cdot \mathbf{R}_j = 2\pi\delta_{ij}$. The wave vector, $\mathbf{k}$ in Eq. 3.1, can be rewritten in terms of reciprocal lattice vectors in the following way

$$\mathbf{k} = \sum_{i=1}^{3} n_i \mathbf{K}_i + \mathbf{k}' \tag{3.2}$$

where the n_i are integers and $\mathbf{k}'$ is restricted to a closed region in k-space about the origin, called the *first Brillouin zone*. Eq. 3.1 can now be written as

$$\mathbf{E}(\mathbf{x},\mathbf{k}) = \tilde{\mathbf{E}}(\mathbf{x},\mathbf{k})e^{i\sum n_i \mathbf{K}_i\cdot\mathbf{x}}e^{i\mathbf{k}'\cdot\mathbf{x}} \tag{3.3}$$

Like $\tilde{\mathbf{E}}$, the exponential involving the reciprocal lattice vectors has the periodicity of the medium (because $\mathbf{K}_i \cdot \mathbf{R}_j = 2\pi\delta_{ij}$), we can combine it with $\tilde{\mathbf{E}}$ to write

$$\mathbf{E}_n(\mathbf{x},\mathbf{k}') = \tilde{\mathbf{E}}_n(\mathbf{x},\mathbf{k}')e^{i\mathbf{k}'\cdot\mathbf{x}} \tag{3.4}$$

where, again, $\mathbf{k}'$ is restricted to the first Brillouin zone. The solutions are now indexed by the reduced $\mathbf{k}'$ *and* the integers n_i. In the above, the n_i have been grouped by a single index n,

commonly referred to as the *band* index. For each $\mathbf{k}'$, then, there is an infinite discrete set of solutions indexed by n. This "folding" of the solutions onto the first Brillouin zone is extremely useful for the analysis of periodic structures, especially when seeking a bandgap.

As an example, consider the 2D case of TM electromagnetic waves propagating through a triangular lattice of dielectric disks (see Fig. 3.1a; TM means $E_x = E_y = B_z = 0$). This particular structure is the basis for the work in this thesis. The center-to-center distance between nearest-neighbor disks is a and each disk has radius r. As shown in Fig. 3.1a, the lattice vectors are $\mathbf{R}_1 = (a, 0)$ and $\mathbf{R}_2 = (a/2, \sqrt{3}a/2)$. The first Brillouin zone for this lattice is the hexagon shown in Fig. 3.1b. However, using additional symmetries of the lattice (namely rotational and reflection symmetries), the essential region of k-space can be reduced further to the shaded wedge in Fig. 3.1b. All of k-space is "folded" into this wedge, yielding the band diagram in Fig. 3.1c (with the further simplification that only solutions along the bounding box of the zone are calculated; this is usually sufficient to determine the existence of a bandgap). Figure 3.1c shows that the TM spectrum for this 2D medium has a prominent *complete* (irrespective of propagation direction) bandgap between the first and second bands. In the parlance of the dielectric mirror, all waves with frequencies in this range are reflected perfectly off of a semi-infinite triangular lattice of sapphire disks.

3.1.3 PhC resonant cavities

Research into PhC resonant cavities has proceeded for decades and is gilded by the following list of highly-cited references [86, 65, 62, 1]. Largely driven by telecommunication applications and cavity quantum electrodynamics experiments, this research has concentrated mostly on the optical-frequency regime; however, the concepts and results apply equally well to the microwave and radio frequencies used in today's accelerators. Given a PhC with a complete bandgap, a resonant cavity is commonly formed by simply removing a single element from the PhC lattice. This can allow a mode to exist which oscillates at a frequency within the bandgap and is spatially confined to the crystal defect (since the bandgap prohibits propagation through the surrounding PhC). An example of such a cavity mode is displayed in Fig. 3.2 using the same triangular lattice of sapphire disks

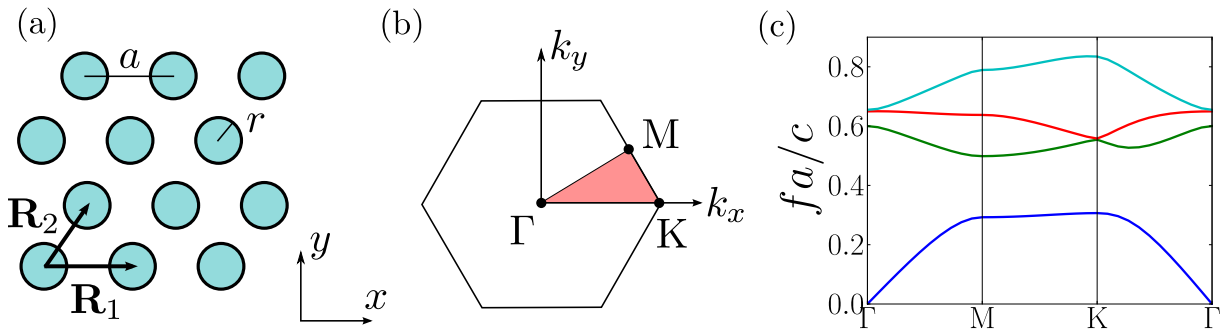

Figure 3.1: *TM characteristics of a 2D triangular lattice of sapphire discs. Propagation in the lattice (a) is forbidden for electromagnetic waves with frequencies near* $0.4c/a$ *because of the bandgap (c). Lattice vectors* $\mathbf{R}_1$ *and* $\mathbf{R}_2$*, inter-disc spacing* a*, and disc radius* r *are defined in (a). The first Brillouin zone of the reciprocal lattice is identified in (b) by the hexagon. The midpoints of the hexagon sides bisect the segments connecting the origin and the nearest reciprocal lattice points; thus,* $\Gamma = (0, 0)$*,* $\mathrm{M} = (\pi/a, \pi/\sqrt{3}a)$*, and* $\mathrm{K} = (4\pi/3a, 0)$*. The irreducible Brillouin zone is shaded and represents the entire Brillouin zone by symmetry. The dispersion curves for the first 4 bands along the path outlining the shaded region in (b) are shown in (c) for* $r = 0.17a$*. Calculations were performed using the MIT Photonic Bands simulation code [39].*

introduced in the previous section. The mode of Fig. 3.2 resonates at a frequency of $f = 0.41c/a$ which is at the center of the bandgap shown in Fig. 3.1c.

The simulated structure in Fig. 3.2 is actually a truncated lattice surrounded by absorbing boundaries. Nevertheless, it illustrates the PhC resonant cavity effect quite well. The difference between Fig. 3.2 and a defect cavity in an infinite lattice is that the mode in the infinite case oscillates indefinitely (assuming lossless dielectric) whereas the mode in Fig. 3.2 has a finite lifetime—it "leaks" electromagnetic energy through the surrounding layers of lattice. This is because, unlike a perfectly conducting cavity which confines all electromagnetic energy, the fields of trapped modes in PhC cavities penetrate into the surrounding lattice with an exponential falloff. A comparison of the field amplitudes inside the cavity to those just outside the last layer of lattice is a good measure of how quickly the mode decays in time. To quantify the losses due to radiation leakage from PhC cavities, we introduce the *radiative* quality factor Q_{rad} (from Eq. 2.20, $Q_{\text{rad}} = \omega U/P_{\text{rad}}$ where P_{rad} is the "leaking" electromagnetic power). In the case of Fig. 3.2, $Q_{\text{rad}} = 24,000$, which says that after $24,000/2\pi$ oscillations, the energy in the cavity is about 1/3 its original level. Because of the exponential falloff of the fields in the PhC, Q_{rad} increases exponentially with the number of layers of lattice surrounding the cavity.

3.2 Photonic crystal accelerators

The selective reflection property of PhCs is of particular interest to accelerator cavity designers because it suppresses HOMs. A PhC accelerator cavity should trap a TM_{010}-like mode (by bandgap frequency exclusion) but allow all other modes to propagate freely out of the structure. Photonic crystals (hence PhC resonant cavities) are easier theorized than fabricated. For this reason, only a handful of fundamentally different PhC accelerating structures have been (and are being) considered in the community. This Section seeks to review the most common design strategies employed over the last decade.

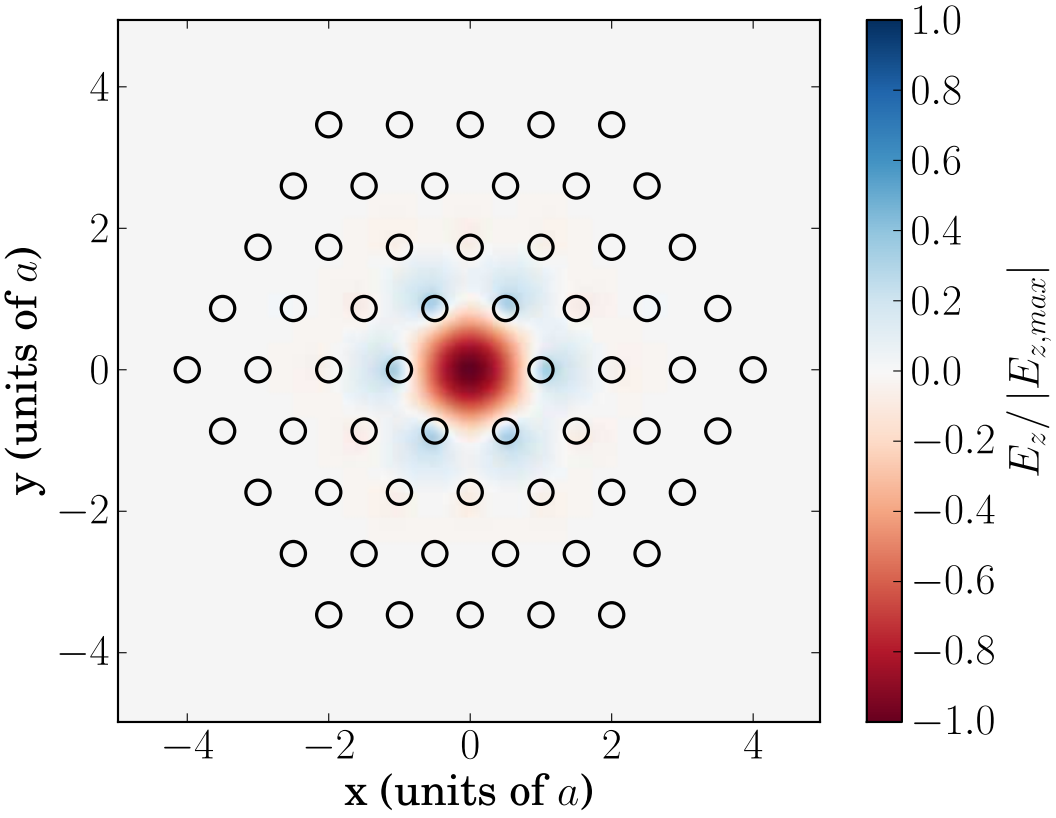

Figure 3.2: *Resonant TM defect cavity mode in a triangular lattice of lossless sapphire disks ($r = 0.17a$). The lattice is truncated at 4 layers (60 disks), giving a radiative Q-factor of $Q_{\mathrm{rad}} = 24000$. Q_{rad} increases exponentially with the number of layers.*

3.2.1 HOM mitigation

To protect against HOMs in PhC accelerator cavity design, one should choose a structure that has only a single bandgap at low frequencies (to trap a TM_{010} mode) so that higher frequencies propagate out. Figure 3.3 shows the location of 2D TM bandgaps as a function of the disc radius to lattice spacing ratio in the triangular lattice of sapphire discs. Highlighted by the vertical line is the ratio chosen for the cavity considered in this thesis. It is strategically located in a region where the bandgap which traps the accelerating mode is wide (ensuring good confinement for a truncated lattice) and there are few significant bandgaps at higher frequencies.

As detailed in Sec. 2.4.5, one should also be mindful of low-v_g modes, since energy transport in these modes is slow. Unfortunately, in a PhC with a bandgap, there will necessarily exist flat regions in the band diagram at the gap edges (for example, see Fig. 3.1c). Modes in this region of the dispersion will have low group velocities and thus could give rise to troublesome wakefields. In fact, a main result of this thesis shows that these are indeed troublesome for the cavity based on the triangular lattice of Fig. 3.2. The details are left to the comparison between PhC-based accelerator cavities and the CLIC cavity discussed in the next Chapter.

3.2.2 Hybrid PhC accelerator cavities

In (arguably) the simplest approach to 3D PhC cavity design (and the associated fabrication difficulties), a hybrid technique is used that combines a 2D PhC cavity (c.f. Fig. 3.2) with a standard conducting cavity [46, 73]. In this scheme, a 2D PhC cavity confines the accelerating mode in the transverse (to the beam axis) plane while conducting walls confine the mode longitudinally. The conducting walls are necessary for trapping a mode that resembles the pillbox TM_{010}, since they can terminate perpendicular electric fields (recall that TM_{mn0} modes have no variation in the z-direction). This scheme is sufficiently mature that a few designs have been built and tested in recent years; the work in this thesis is based on this concept.

Pioneering work was recently done on the design, construction, and testing of all-metal hy-

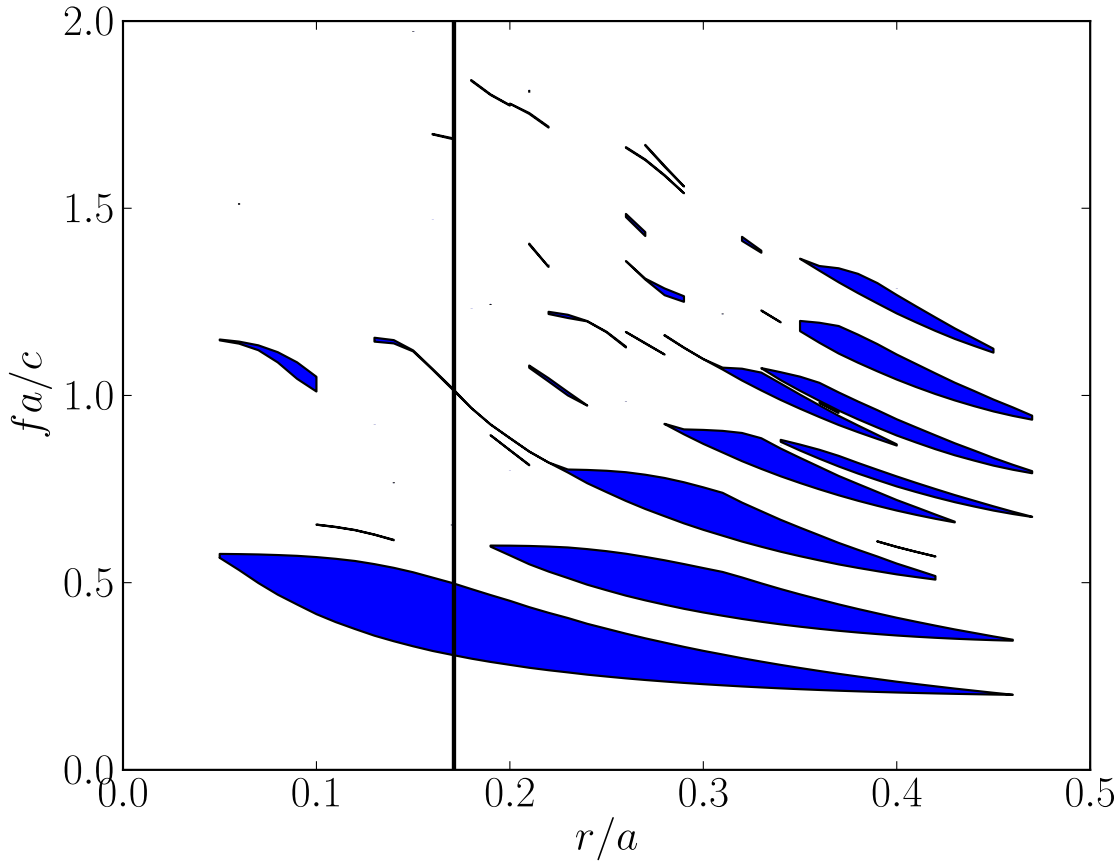

Figure 3.3: *Complete 2D TM bandgaps as a function of the disc radius to lattice spacing ratio in the triangular lattice of sapphire discs. The vertical line shows the ratio for the cavity considered in this thesis (this ratio defines the structure in Fig. 3.2). Calculations were performed using the MIT Photonic Bands simulation code [39]. Some of the very thin gaps are due to numerical error.*

brid PhC cavities (endplates *and* rods are metal) [72, 54, 55]. A transverse cross-section of the structure and the electric field patterns for the accelerating and lowest dipole modes are shown in Fig. 3.4. In the first experiment, acceleration of an electron beam at a gradient of $E_{\rm acc} = 35$ MV/m was demonstrated in a 6-cell structure where each cell was a hybrid PhC cavity connected to its neighboring cavities by a circular beam port [72]. In the next experiment, the structure was tested for HOM wakefields under excitation by a beam from the MIT linac [54, 55]. Wakefields were measured by frequency-analyzing the power escaping the structure after passage of the exciting beam. Theoretical predictions and simulations were validated for beam losses into the fundamental accelerating mode. Only monopole HOMs were observed in the power spectrum. Finally, high-gradient limitations were explored in high-RF-power tests at SLAC [54, 55]. These tests showed that gradients in these types of structures are limited by intense Ohmic heating of the inner ring of rods due to high surface magnetic fields. This is in contrast to the more common case of high surface electric field limits in other copper cavity geometries.

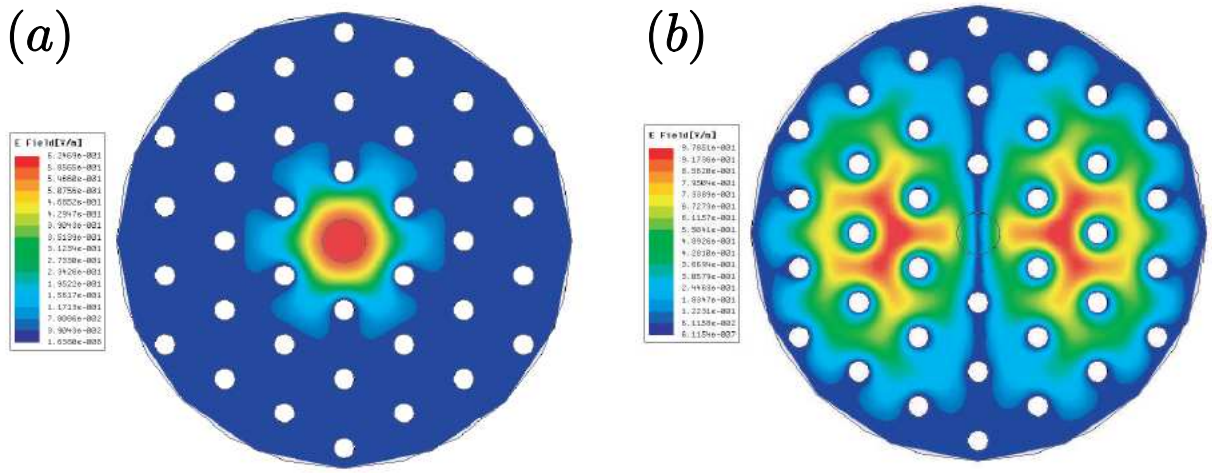

Figure 3.4: *Magnitude of the electric field for the (a) TM_{010} (fundamental accelerating) mode and the (b) TM_{110} (lowest dipole) mode in the metal-rod hybrid PhC accelerating cavity built and tested by MIT (plots taken from [54]).*

Another collaboration has been investigating hybrid dielectric/metal PhC cavities (dielectric rods sandwiched between copper plates) both theoretically and experimentally (see Fig. 3.5)

[56]. Their experiments have validated the predicted fundamental mode existence and Q-factor. Theoretical work found that some cavities formed from photonic quasi-crystals (rotationally symmetric arrangements of rods without translational symmetry) can give higher Q_{rad} values than lattice structures with similar extent [17]. Also, replacing outer dielectric rods with metal rods could substantially increase Q_{rad} without increasing losses in the rods (since the high fields are still concentrated on the inner ring of dielectric rods) [25]. However, the effect of this on HOMs is unknown.

3.2.3 All-dielectric PhC waveguide accelerators

A distinct effort is also being made in the accelerator community to explore particle acceleration at optical frequencies in purely dielectric PhCs. These schemes would have several advantages: (1) accelerating gradients could be much higher than those in metallic and superconducting structures because of the increased breakdown thresholds in dielectrics at optical frequencies, (2) high-power optical frequency lasers are a reliable and widely-available power source, and (3) lithographic techniques for micron-sized silicon structures are advancing rapidly. One scheme uses hollow-core photonic crystal fibers to guide a speed-of-light mode [51], whereas the other more complicated undertaking involves the lithographic construction of a PhC with a complete 3D bandgap, called the woodpile structure [13]. The above techniques are theoretically attractive; however, experiments have yet to verify the predictions.

3.3 Optimized high-Q_{rad} HDPhC cavities

Power losses determine important figures of merit in accelerator cavities (namely Q and r_{shunt}). The losses due to radiation leakage in finite PhC cavities (referred to via Q_{rad} from here on) introduce a method of power dissipation not found in traditional metallic accelerator cavities and must be minimized. One natural approach to this problem is to add more layers of lattice to the PhC cavity, guaranteeing an exponential increase in confinement with structure extent. Unfortunately, the number of layers can become impractically large, especially for weak scatterers

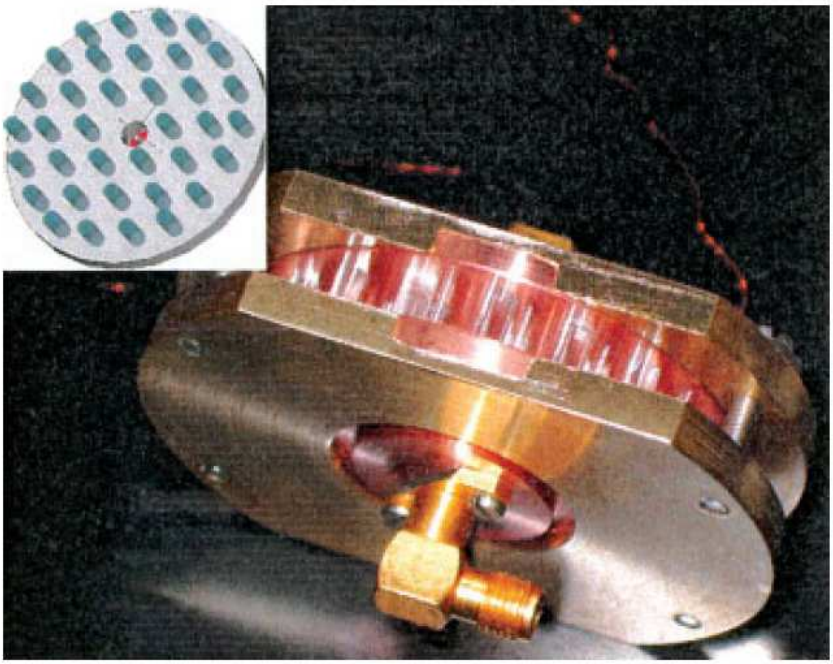

Figure 3.5: *Hybrid PhC cavity using dielectric rods constructed and tested at the University of Naples Federico II [56].*

(even sapphire, with a fairly large dielectric constant of $11.4\varepsilon_0$ takes 4 layers—60 rods—to give a Q_{rad} of 24,000, as seen in the previous Sections) or difficult fabrication techniques (such as nanolithography in the optical regime). An alternative method was suggested in Ref. [25], where outer-layer dielectric rods are replaced by metal rods. It was shown that this can greatly increase Q_{rad} while avoiding the issues associated with high surface fields on metal rods. This Section discusses yet another solution (introduced by us) to the Q_{rad} problem: the optimization of finite PhC cavities into arbitrary structures with improved mode confinement. We have shown that this technique works well for 2D dielectric cavities (and thus HDPhC cavities), exhibiting orders-of-magnitude increases in Q_{rad} without adding rods to the structure [4]. This Section comprises the details of our findings.

3.3.1 Model for optimization

Fig. 3.6 shows the TM_{010} mode trapped by a truncated triangular lattice of 18 rods with $\varepsilon = 10\varepsilon_0$ (alumina) in vacuum. By varying the ratio r/a, where r is the rod radius and a the lattice constant, we have found that maximum confinement occurs when $r/a = 0.18$, where $Q_{\text{rad}} = 130$. The resonant frequency of the mode lies near the center of the bandgap of the infinite triangular lattice for the same r/a (see Fig. 3.1). Because the optimizations in this study produce irregular structures, we express length scales in units of c/ω_0 where ω_0 is the resonant frequency of the TM_{010} mode. In these units, the rod radius in Fig. 3.6 is $r = 0.47c/\omega_0$ and the radial position of the innermost ring of rods is, $R = a = 2.6c/\omega_0$. Thus, for a target frequency of 12GHz, the dimensions are: $r = 1.9$mm and $R = 10.3$mm.

3.3.2 Optimization techniques

To maximize Q_{rad}, an optimization routine explored alternative configurations by gradually moving rods away from their original lattice positions of Fig. 3.6. Each new configuration in the optimization was tested by FDTD simulation of the decaying mode. From each simulation, the resonant frequency was extracted and used with the decaying cavity field data to calculate Q_{rad}.

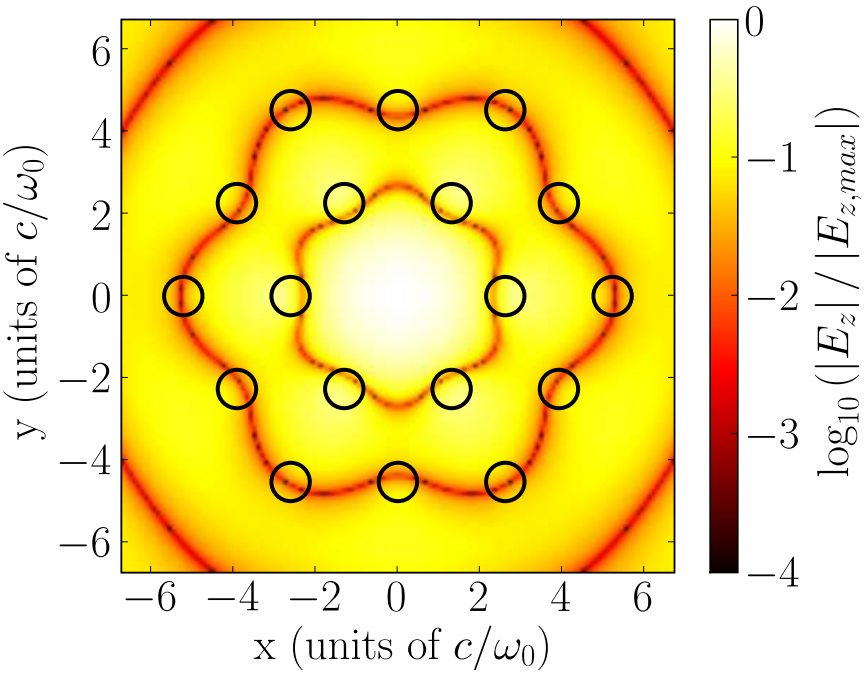

Figure 3.6: *Normalized electric field of the accelerating mode in a truncated triangular-lattice PhC cavity (initial stage of optimization). Cylinders have a dielectric constant* $\varepsilon_{\mathrm{r}} = 10$ *and a radius* $r = 0.47c/\omega_0$*, where* ω_0 *is the resonant angular frequency. The quality factor is* $Q_{\mathrm{rad}} = 130$ *and the lattice spacing is* $a = 5.59r$*. The radial position of the innermost ring is* $R = a = 2.61c/\omega_0$*.*

Simulations of each configuration in the optimization were performed with the FDTD electromagnetics code, VORPAL [60]. Dielectric cylinders were resolved to 10 grid cells per diameter, and fields were absorbed at the edges of the simulations using perfectly matched layers (PMLs) [5]. In grid cells only partially filled with dielectric, the dielectric constant was averaged according to

$$\varepsilon = f\varepsilon^{\text{alumina}} + (1-f)\varepsilon^{\text{vacuum}}, \tag{3.5}$$

where f is the fraction of the grid cell filled with dielectric. This averaging is important for detecting continuous cylinder displacements (i.e. displacements smaller than a grid cell). The above averaging is accurate only for electric fields parallel to the dielectric-vacuum interface [40].

A gaussian-modulated (in time) sinusoidal current excited each configuration at a single point near the cavity center. The excitation was truncated at four standard deviations on both sides of the gaussian, and, afterward, the excited modes were left to resonate undisturbed for a few hundred oscillations to rid the simulation of transients. The width of the corresponding spectral gaussian was $\Delta\omega/\omega = 6.4\%$ and was centered at the best guess for the TM_{010} frequency to be calculated. The timestep, Δt, was set to $0.99\Delta t_{\text{CFL,2D}}$, where $\Delta t_{\text{CFL,2D}}$ is the 2D Courant-Friedrichs-Lewy time step.

The complex resonant frequencies (oscillation and decay constant) were extracted using the modified filter diagonalization method of Ref. [82]. The method requires sufficient isolation of mode frequencies through appropriate spectral filtering. For our gaussian-modulated excitations truncated at four standard deviations, this method was consistently able to extract ω_0 to a relative error of 10^{-4}.

The quality factor, Q_{rad}, was calculated from the time decay of the freely-oscillating electric field sampled at the center of the cavity. In all simulations, ω_0 was the dominant frequency in the time series and therefore the electric field at the cavity center had the form (from Eq. 2.20)

$$E_z(t) = E_{z,0}e^{-i\omega_0 t}e^{-\omega_0 t/2Q_{\text{rad}}}. \tag{3.6}$$

The damping rate was obtained by a linear fit to $\ln(|E_z(t)|^2)$, and Q_{rad} calculated using the extracted ω_0.

Incorporating the methods described above, the optimization routine used a Nelder-Mead simplex minimization algorithm that manipulated the rod positions to increase Q_{rad} [67]. The cost function to be minimized was simply

$$g(\mathbf{p}) = g(x_1, \ldots, x_N, y_1, \ldots, y_N) = -Q_{\mathrm{rad}}, \tag{3.7}$$

where $\mathbf{p}$ is a vector describing one configuration of cylinders. Initially, the optimization built a simplex of configurations,

$$S = \{\mathbf{p}_1, \ldots, \mathbf{p}_{2N+2}\}, \tag{3.8}$$

each a small deviation from Fig. 3.6. At each iteration, a new configuration was constructed, excited, left to resonate, and analyzed. If the new configuration was an improvement, it replaced the worst configuration in S and the extracted ω_0 was used as the excitation frequency for the next trial configuration.

This feedback mechanism kept the excitation near resonance with the cavity and resisted the excitation of nearby (in frequency) modes. A restriction returned $g = 0$ if other mode frequencies were present, but was never encountered in any of the optimizations. Optimization terminated when the configurations stored in the last iteration returned

$$|g(\mathbf{p}_{\mathrm{i}}) - g(\mathbf{p}_{\mathrm{j}})| < 10 \quad \forall(\mathbf{p}_{\mathrm{i}}, \mathbf{p}_{\mathrm{j}}) \in S. \tag{3.9}$$

The Nelder-Mead routine will not guarantee a global minimum, although it has the ability to skip over shallow local minima by storing several configurations at each iteration. The simplicity of the routine slows convergence, but can be tolerated in this case because of rapid FDTD simulation and Q_{rad} calculation. The routine may also be highly dependent on the initial configuration; however, the results for the triangular lattice are sufficiently dramatic and so other possibilities are neglected here.

3.3.3 Results

Figure 3.7 shows the result of a sixfold rotationally symmetric optimization ($N = 3$ in Eq. 3.7) for non-overlapping cylinders. From Fig. 3.6 to Fig. 3.7, Q_{rad} went from 130 to 11000 and

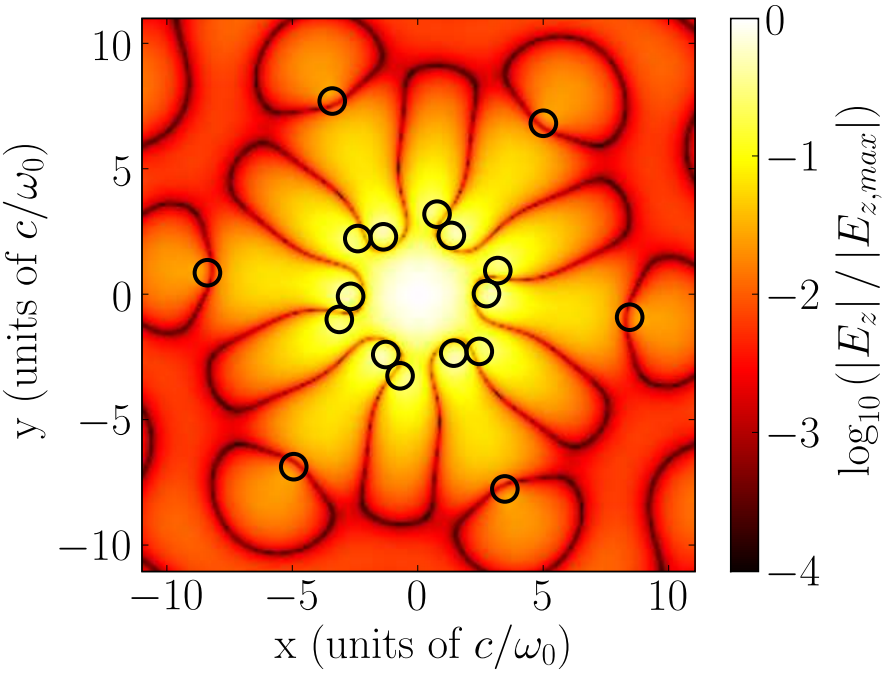

Figure 3.7: *Normalized electric field of the confined mode in the 18-rod sixfold rotationally symmetric, optimized accelerating cavity.* $Q_{\mathrm{rad}} = 11,000$, $r = 0.51c/\omega_0$, *and* $R = 2.71c/\omega_0$. *Rods do not overlap.*

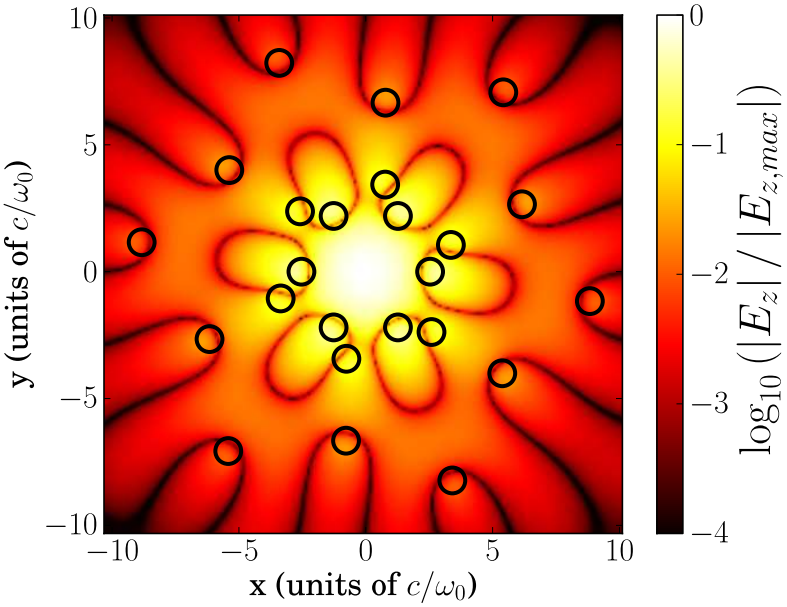

Figure 3.8: *Normalized electric field of the confined mode in the 24-rod sixfold rotationally symmetric, optimized accelerating cavity.* $Q_{\mathrm{rad}} = 11000$, $r = 0.51c/\omega_0$, *and* $R = 2.71c/\omega_0$. *Rods do not overlap.*

required 350 Nelder-Mead iterations (500 cost-function calls, or configuration simulations). For the triangular lattice to attain the same Q_{rad}, 4 layers (60 cylinders) are required. Enforcing symmetry on the optimization not only decreased the computation time significantly, but also yielded a higher Q_{rad} as compared to an asymmetric optimization ($N = 18$) which converged to a local minimum with a smaller Q_{rad}. An 18-rod pattern very similar to Fig. 3.7 was also found using sapphire rods ($\varepsilon = 11.4\varepsilon_0$) and resulted in $Q_{\mathrm{rad}} = 25,000$.

An optimization of a 24-rod sapphire structure ($N = 4$ in Eq. 3.7) was also performed using the optimized 18-rod sapphire configuration as a guide for a favorable initial condition. The first configuration in this optimization was constructed by placing one rod in each of the azimuthal vacancies of the outer ring of the 18-rod optimized pattern. The addition of these rods destroyed the previously high Q_{rad}, but it was quickly regained through optimization—the ultimate Q_{rad} was 192,000 and the associated rod pattern is shown in Fig. 3.8.

Symmetry breaking of various forms appears to be a consequence of these optimizations. First, the final configurations lack any resemblence to the original truncated lattice. Second, as the optimization proceeds, modest Q-factors were found until parity symmetry was broken. Then much higher Q-factors were obtained. Simulation confirmed that the mirror image configurations also have the same Q_{rad}.

The following additional observations were made. Because of the increased concentration of dielectric material in high-field regions, power loss to dielectric heating doubled. For an optimization allowing rod overlap, the inner rods partially merged in pairs, but without significant change in Q_{rad} compared to the nonoverlapping case. Grid effects proved to be insignificant—a high-resolution simulation of the optimized structure showed no change in Q_{rad}.

3.3.4 Sensitivity analysis

To test the sensitivity of the optimized structure, we computed the frequency and Q-factor for configurations with small random displacements from the optimized structure. Unlike the optimized structure, the displaced configurations were not necessarily sixfold symmetric, but used

the full configuration vector (except for the rod radius, which we kept constant), $\mathbf{p}_{\rm opt}$, with $N = 18$ in Eq. (3.7), which is a 36-dimensional vector. For each test, we formed a displacement vector, $\Delta\mathbf{p}$, where each $\Delta p_{\rm i}$ was normally distributed with mean 0 and standard deviation σ.

One thousand configurations ($\mathbf{p}_{\rm opt} + \Delta\mathbf{p}$) were tested for three different σ ($0.02r$, $0.05r$, and $0.10r$). One thousand different directions in the function space of $g(\mathbf{p})$ for 36 dimensions is sufficient for exploring along each dimension individually plus all diagonals spanning two dimensions at a time. (However, a restriction that prohibited rods from overlapping limited exploration in some directions of the 36-dimensional space.) Figure 3.9 shows $Q_{\rm rad}$ and ω_0 as a function of the total combined displacement of the inner 12 rods,

$$|\Delta\mathbf{p}_{\rm inner12}| = \sqrt{\sum_{i=1}^{24} \Delta p_i^2} = \sqrt{\sum_{i=1}^{12} (\Delta x_i^2 + \Delta y_i^2)}. \tag{3.10}$$

Displacements of the outer 6 rods showed comparatively insignificant effects, decreasing $Q_{\rm rad}$ by less than 3% for individual cylinder displacements $\sim 0.10r$.

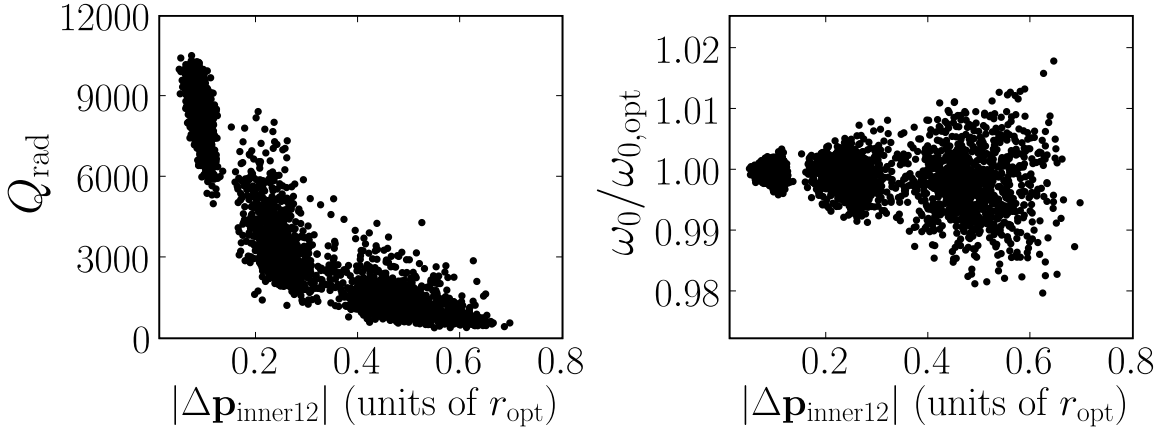

Figure 3.9: *Sensitivity of the structure as a function of the total combined displacement of the inner 12 rods,* $|\Delta\mathbf{p}_{\rm inner12}|$. *From left to right in each plot, the groups of points are from the ensembles with* $\sigma = 0.02r$, $0.05r$, *and* $0.10r$, *respectively.*

3.3.5 Conclusion

For high dielectric constant scatterers, optimization through FDTD simulation has yielded dramatic improvements in mode confinement for PhC-based cavities, increasing the Q_{rad} of a confined mode by 2 orders of magnitude. This increase in Q-factor benefits PhC accelerating structures by reducing the structure size and material requirements while retaining confinement properties that trap the desired mode but not unwanted HOMs.

The breaking of lattice symmetry in the optimized structures suggests an investigation of other types of discrete rotational symmetries besides the six-fold performed here. Additional studies of fully asymmetric optimizations are also warranted.

3.4 Expectations of a hybrid dielectric PhC cavity

The work in this thesis is based on the hybrid dielectric PhC (HDPhC) cavity (lattice-based and optimized) mentioned in the previous sections. Before turning to simulation results in the next Chapter, we first summarize some expectations of the performance of these novel structures.

3.4.1 Accelerating mode Q-factor

Losses (excluding those due to an accelerated beam) in HDPhC cavities have three main constituents: (1) radiative losses, (2) heating losses in metal endplates, and (3) heating losses in dielectric PhC elements.

The only conductors in a HDPhC cavity are the endplates (or irises in a multicell cavity). In the absense of beam tubes (closed cavity), the calculation of Q_{metal} for the accelerating mode becomes trivial since

$$\int |\mathbf{H}(\mathbf{x})|^2 dv = L \int |\mathbf{H}(\mathbf{x})|^2 da \tag{3.11}$$

and Eq. 2.24 gives the following for Q_{metal}:

$$Q_{\text{metal}} = \frac{\mu_0 \omega L}{2R_s} \tag{3.12}$$

$$= \frac{\mu_0 c \pi}{2R_s} \tag{3.13}$$

where the second expression is obtained by setting the cavity length to $L = c/2f$ (to maximize particle energy gain). For copper at 12 GHz, the surface resistance is $R_s = 0.028$ Ohms; therefore, $Q_{\text{metal}} \approx 2.1 \times 10^4$. The pillbox cavity suffers a lower Q_{metal} because of the high surface magnetic fields on the cylindrical wall; using the same R_s, the pillbox cavity gives $Q_{\text{metal}} = 9.2 \times 10^3$. The elimination of the copper walls reduces the heating losses by about a factor of 2.

It is hoped that most of the accelerating mode energy will be contained in the vacuum region of the PhC (based on Fig. 3.2). Nevertheless, an upper bound on the losses incurred by heating of the dielectric rods is given simply by the inverse of the dielectric loss tangent (i.e. as if all electric field energy were inside the dielectric). For the case of sapphire at room temperature, this puts a lower bound on the dielectric Q-factor of $Q_{\text{diel}} \geq 10^5$. Practically, the dielectric rods will heat under high RF power, which will likely increase the loss tangent (based on the trend in Table 2.1); therefore, when calculating figures of merit for the structures of interest in this thesis, a more conservative value of $\tan\delta = 10^{-4}$ is chosen for sapphire.

Finally, a significant fraction of the total losses in HDPhC cavities can come from radiation leakage through the PhC. Several counteractions are possible, including lattice extension (not always a practical option), replacement of weak scatterers with strong scatterers (e.g. replacing dielectric rods with metal rods), or high-Q_{rad} optimization as discussed in the previous Section.

The total Q-factor is calculated via

$$\frac{1}{Q_{\text{tot}}} = \frac{1}{Q_{\text{metal}}} + \frac{1}{Q_{\text{diel}}} + \frac{1}{Q_{\text{rad}}} \tag{3.14}$$

In the best-case scenario, Q_{diel} and Q_{rad} dominate over Q_{tot}; thus the total Q-factor is bounded by the losses to the metal—i.e. $Q_{\text{tot}} \approx Q_{\text{metal}}$. Therefore, the highest Q-factor one can expect from a HDPhC is about twice that of the pillbox.

3.4.2 Shunt impedance, loss factor

The shunt impedance of a HDPhC accelerating mode is, of course, related to the previous discussion on power losses. Recall that the Q-factor measures the ratio of the energy to the power loss, the shunt impedance measures the ratio of the accelerating gradient to the power loss, and the loss factor connects the two. If the power losses are roughly halved going from a pillbox to a HDPhC, the shunt impedance has the potential to double. This depends on the spatial distribution of the mode energy. In an HDPhC, the energy is not confined in the transverse plane as tightly as in the pillbox; thus the loss factor is likely to be lower. If the loss factor in an HDPhC can be kept above half of the pillbox loss factor, gains in shunt impedance can be expected.

3.4.3 Wakefields

Short-range wakefields are dominated by the interaction between a bunch and the beam tubes. Therefore, the short-range wakes in a HDPhC cavity should mirror those in a standard pillbox cavity for identical beam tube geometries. Loss factors ($\propto V^2/U$) of HOMs may decrease in general due to the more open nature of a HDPhC cavity compared to the CLIC cavities (because of diluted energy densities). However, mode densities may increase for the very same reason. The following Sections address more detailed attributes of wakefields in HDPhC cavities; namely, parallel-plate vs. rectangular waveguide damping and guided modes in dielectric rods.

3.4.4 Parallel-plate waveguide damping

Somewhat like the CLIC cavities, HDPhC cavities are waveguide-damped, where the waveguide in this case consists simply of two (infinite) conducting plates (we will refer to this geometry as a parallel-plate waveguide). Recall that in a rectangular cross-section waveguide with the x-axis along the waveguide axis and side lengths L_y and L_z, the frequencies of the TE modes (i.e. $E_x = 0$) are

$$\omega_{mn}(k_x) = c\sqrt{k_x^2 + \frac{\pi^2 m^2}{L_y} + \frac{\pi^2 n^2}{L_z}} \tag{3.15}$$

where, $m = 0, 1, 2, \ldots$ and $n = 1, 2, \ldots$ (or vice versa) [37]. In the CLIC cavity, $L_y = 11$mm, giving a TE_{10} cutoff frequency of 13.6 GHz. The undamped TM dipole resonant frequency for an average CLIC cell is approximately 21 GHz, placing it well above cutoff and encouraging strong damping based on the analysis of Section 2.4.5.

In a parallel-plate waveguide of gap L_z, the frequencies of modes with nonzero E_z are

$$\omega_m(k_x, k_y) = c\sqrt{\frac{\pi^2 m^2}{L_z} + k_x^2 + k_y^2} \tag{3.16}$$

where $m = 0, 1, 2, \ldots$. In contrast with the rectangular waveguide, modes with uniform E_z (i.e. $m = 0$) in the parallel plate waveguide have no cutoff frequency. Thus, the potential for stronger damping of the lowest dipole mode (when compared to rectangular waveguide damping) exists in HDPhC cavities. Modes *with* z-variation do have cutoff frequencies in the parallel-plate waveguide; the wakefields due to these modes will exhibit the persistent behavior discussed in Section 2.4.5 but will occur at higher frequencies (the rectangular waveguide also suffers from higher cutoff frequencies).

3.4.5 Dielectric-guided modes

HDPhC cavities have the potential to trap HOMs in the dielectric rods by total internal reflection. These modes would have high Q-factors (set by $\tan^{-1}\delta$) and thus a negative impact on wakefields. For this reason, we have calculated the analytic modes in sapphire rods (see Appendix B). The model considered here is of a single sapphire cylinder of radius r spanning a gap of length ℓ between two infinite perfectly-conducting plates (cylinder axis and conducting plate normals are along z, as in the HDPhC cavities). The modes of concern (because of their significant E_z component) are labeled by HE_{mnp} and TM_{0np} where m, n, and p are the azimuthal, radial, and longitudinal mode numbers, respectively (relative to the rod-centered coordinate system). HE modes have both E_z and H_z components and $m > 0$; TM modes have $H_z = 0$. The lowest-frequency modes of this group are HE_{111}, HE_{112}, TM_{011}, and HE_{212} (see Figs. 3.10 and 3.11). Frequencies for rod radius $r = 1.8$mm and two common values of ℓ are given in Table 3.1.

Given the analytic mode solutions, we were able to calculate loss factors for each. The loss

ℓ (mm)	HE_{111} (GHz)	HE_{112} (GHz)	TM_{011} (GHz)	HE_{212} (GHz)
6.66	17.4	22.2	22.0	31.6
7.33	16.9	21.3	20.4	31.0

Table 3.1: *Frequencies of guided modes in sapphire rods of length ℓ and radius $r = 1.8mm$ where each rod end abuts an infinite conducting plate.*

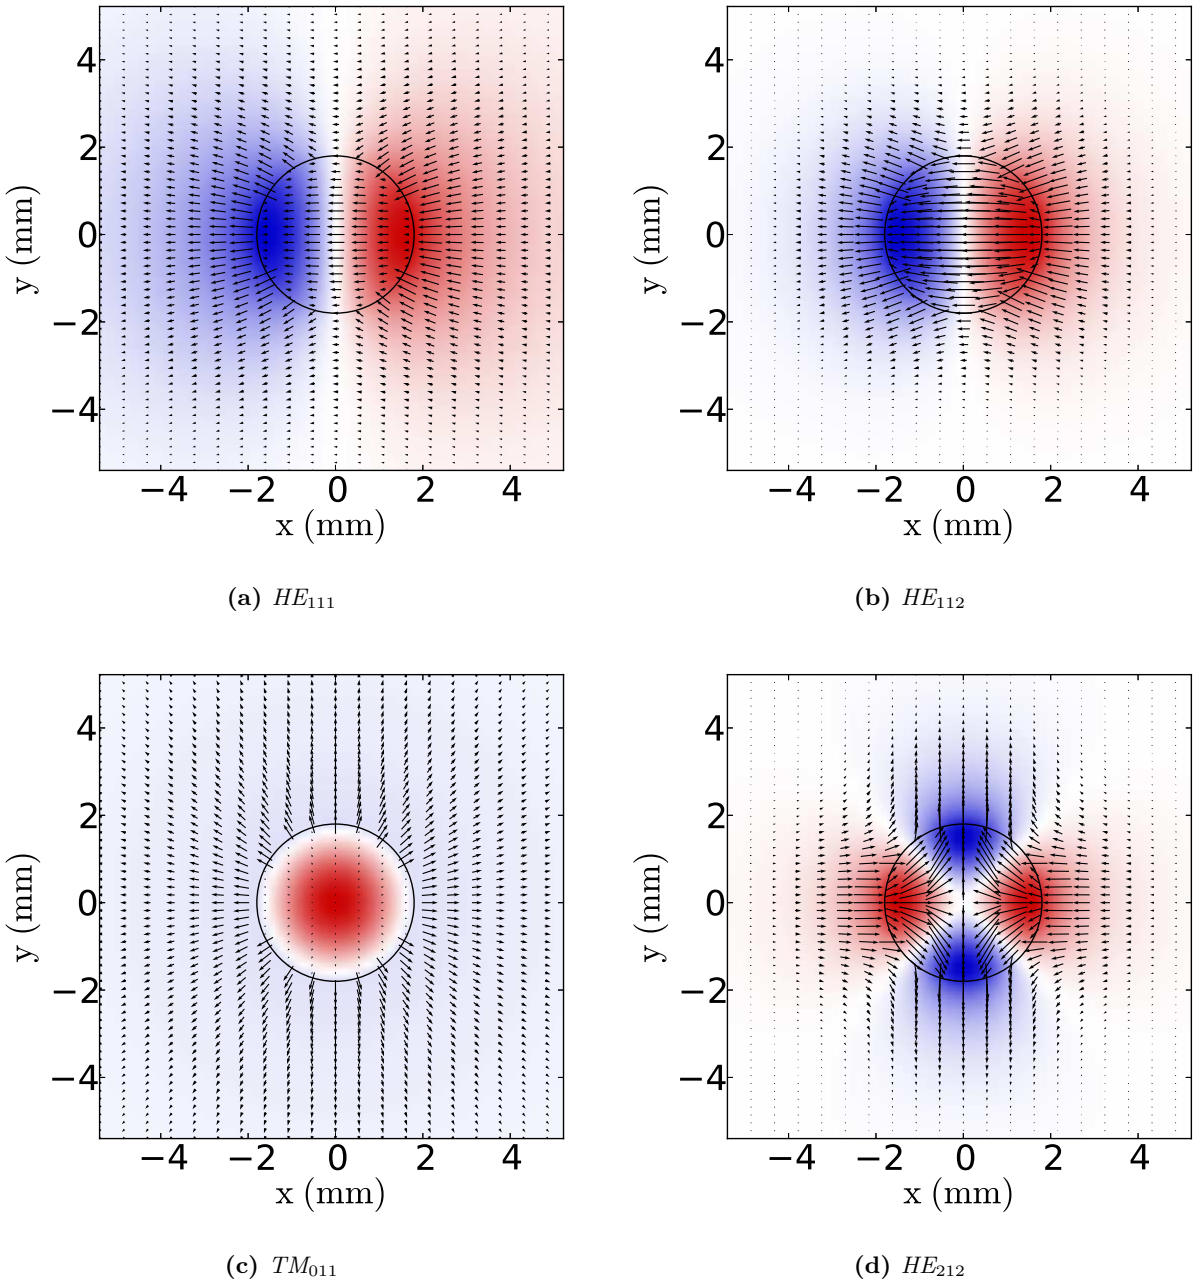

(a) *HE_{111}*

(b) *HE_{112}*

(c) *TM_{011}*

(d) *HE_{212}*

Figure 3.10: *Guided mode electric fields for $\ell = 6.66mm$ and $r = 1.8mm$. Color shows E_z while arrows show $\mathbf{E}_\perp$.*

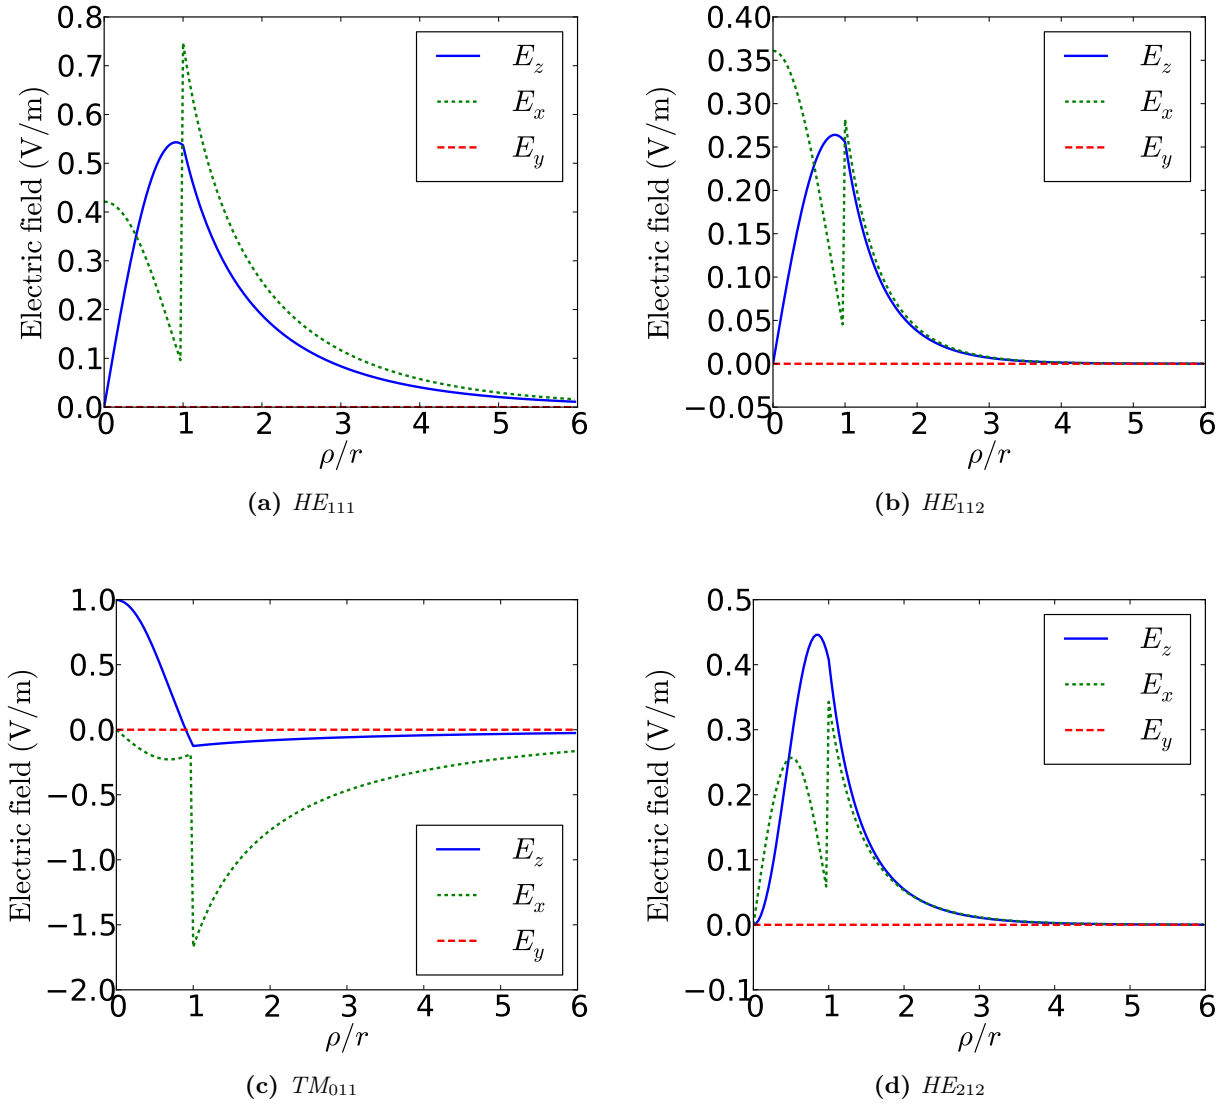

(a) HE_{111} **(b)** HE_{112}

(c) TM_{011} **(d)** HE_{212}

Figure 3.11: *Guided mode electric fields for* $\ell = 6.66mm$ *and* $r = 1.8mm$ *along the* $y = 0$ *line.*

factors shown in Fig. 3.12 were calculated assuming a speed-of-light charge moving parallel to the rod axis. The vertical line shown in the figure represents the distance to the beam axis from an inner rod in the triangular-lattice-based HDPhC cavity. At this distance, only the HE_{111} and TM_{011} modes are of concern; their loss factors are near unity (the accelerating mode loss factor is $\sim 10^2$). Their high Q-factors, however, could lead to problems in the long-range wakefield. Coupling between rod-guided modes in a HDPhC cavity will change the results presented here, especially in the optimized HDPhC cavities where inner rods almost touch. One consequence could be the lowering of guided-mode frequencies due to the increased fraction of dielectric (i.e. two rods could act like one bigger rod). In the next Chapter, we will see that PhC modes—unguided—present a larger problem than these weakly-excited guided modes. Therefore, we ignore the contribution of the guided modes in the calculated wakefields.

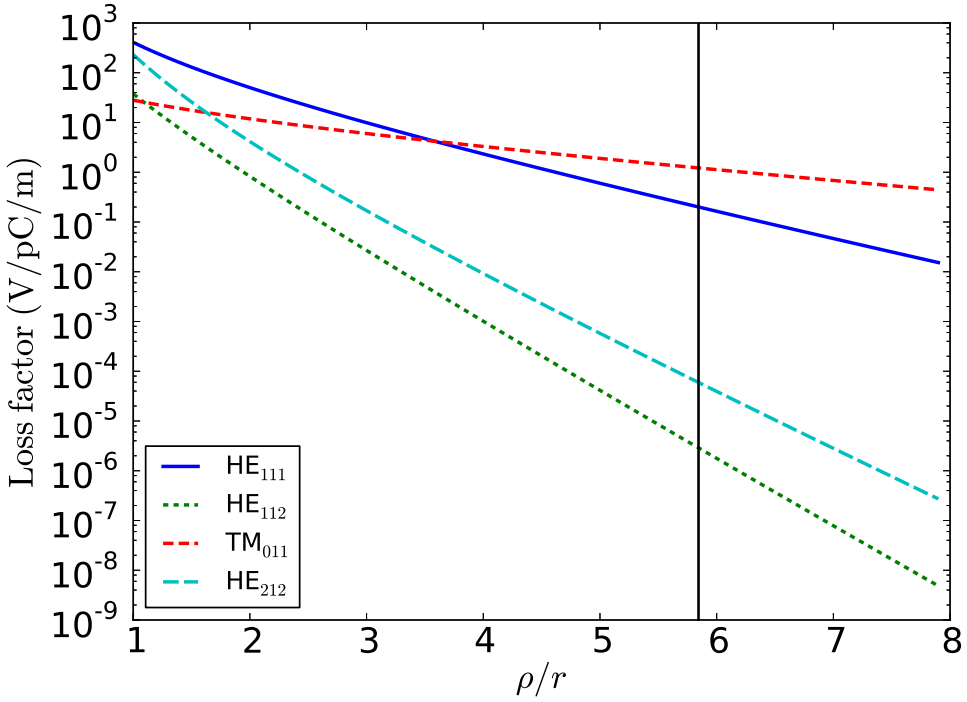

Figure 3.12: *Loss factors of the lowest-frequency guided modes for a charge traveling parallel to a sapphire rod sandwiched between two infinite conducting plates, where $\ell = 6.66mm$ and $r = 1.8mm$. The horizontal axis is the radial distance from the rod-center to charge. The vertical line indicates the distance to the beam axis relative to an inner rod in an HDPhC cavity.*

Chapter 4

Cavity performance analysis and comparison

This chapter presents and discusses simulation results for several different cavity types, including: the pillbox cavity (closed and with beam tubes), the triangular lattice and optimized HDPhC cavities (closed and with beam tubes), and the latest iteration of the CLIC cavity (beam tubes only). For each, we detail the figures of merit of the accelerating mode and the properties of the wakefields due to a $\sigma_z \approx 1$mm Gaussian electron bunch. The results show that there is a small loss in accelerating efficiency in HDPhC cavities as compared to the conducting cavities tested (given $\tan\delta = 10^{-4}$ for sapphire). However, this may be overcome by increasing Q_{rad} (by adding rods) and/or taking a more optimistic sapphire loss tangent. Wakefields are strongly suppressed in both the CLIC cavity and the HDPhC cavities. The CLIC cavity suppresses transverse wakefields more effectively than the HDPhC cavities, whereas HDPhC cavities damp longitudinal HOMs more strongly. The larger transverse wakes in the lattice-based HDPhC cavity are explained by examining the 2D TM band diagram.

The cavity geometries compared in this Chapter have all been mentioned previously and are shown together in 3D in Fig. 4.1. Proxies for each cavity type will be given as *CLIC*, *Tri-X-Sapphire* (meaning triangular lattice with X layers of sapphire rods—Fig.4.1 shows the Tri-4-Sapphire cavity), *Opt-18-Sapphire*, and *Opt-24-Sapphire*. A fifth cavity, the *empty* cavity, will also be analyzed. The empty cavity shares the same geometry as the HDPhC cavities, but, as its name suggests, does not contain any rods. The empty cavity will be referenced in Sections on wakefields, since its behavior in that regard is ideal for the classes of cavity considered in this thesis. The

undamped pillbox will also be referenced when discussing the properties of the accelerating mode since it tends to have the highest shunt impedance for conducting cavities of the same material.

4.1 Simulation difficulties

Before presenting the results of the simulations, we first discuss some primarily numerical issues that have somewhat dictated the simulation and analysis methods employed. First, we found it difficult (if not impossible) to perform a complete modal analysis of the wakefields because of the increased mode density in HDPhC cavities; the computational requirements involved in extracting the entire spectrum of interest are unreasonable. Second, the most effective electromagnetic absorbers were found to be numerically unstable in the CLIC cavity, forcing the simulations used in the comparison analysis to employ more physical conducting absorbers instead (at the expense of ideal absorption). These issues are described below in greater detail.

4.1.1 Mode-finding

Individual modes (such as the accelerating mode or troublesome HOMs) are usually calculated in one of two ways: frequency-domain eigensolvers or time-domain mode extraction. Frequency-domain eigensolvers are very efficient at finding low-frequency modes, but, to remain efficient, must calculate all modes up to the frequency of interest (once a set of the lowest-frequency modes are found, orthogonalization against them helps to find the next higher-frequency mode). This can pose a problem when the modes of interest are HOMs with frequencies much higher than the fundamental. In conducting cavities such as the pillbox, the spectrum is sparse for the modes of interest (the dipole mode has the second lowest frequency); thus, eigensolvers work well.

Simulated volumes in HDPhC cavities are much larger than those in conducting cavities; this increases the mode density at lower frequencies. Furthermore, in the case of HDPhCs, dielectrics tend to lower resonant frequencies, thus exacerbating the mode density problem. For example, in the pillbox, the accelerating mode *is* the lowest frequency mode whereas for the 2D Tri-4-Sapphire cavity, an eigensolver found hundreds of modes (resembling the lowest modes of an empty

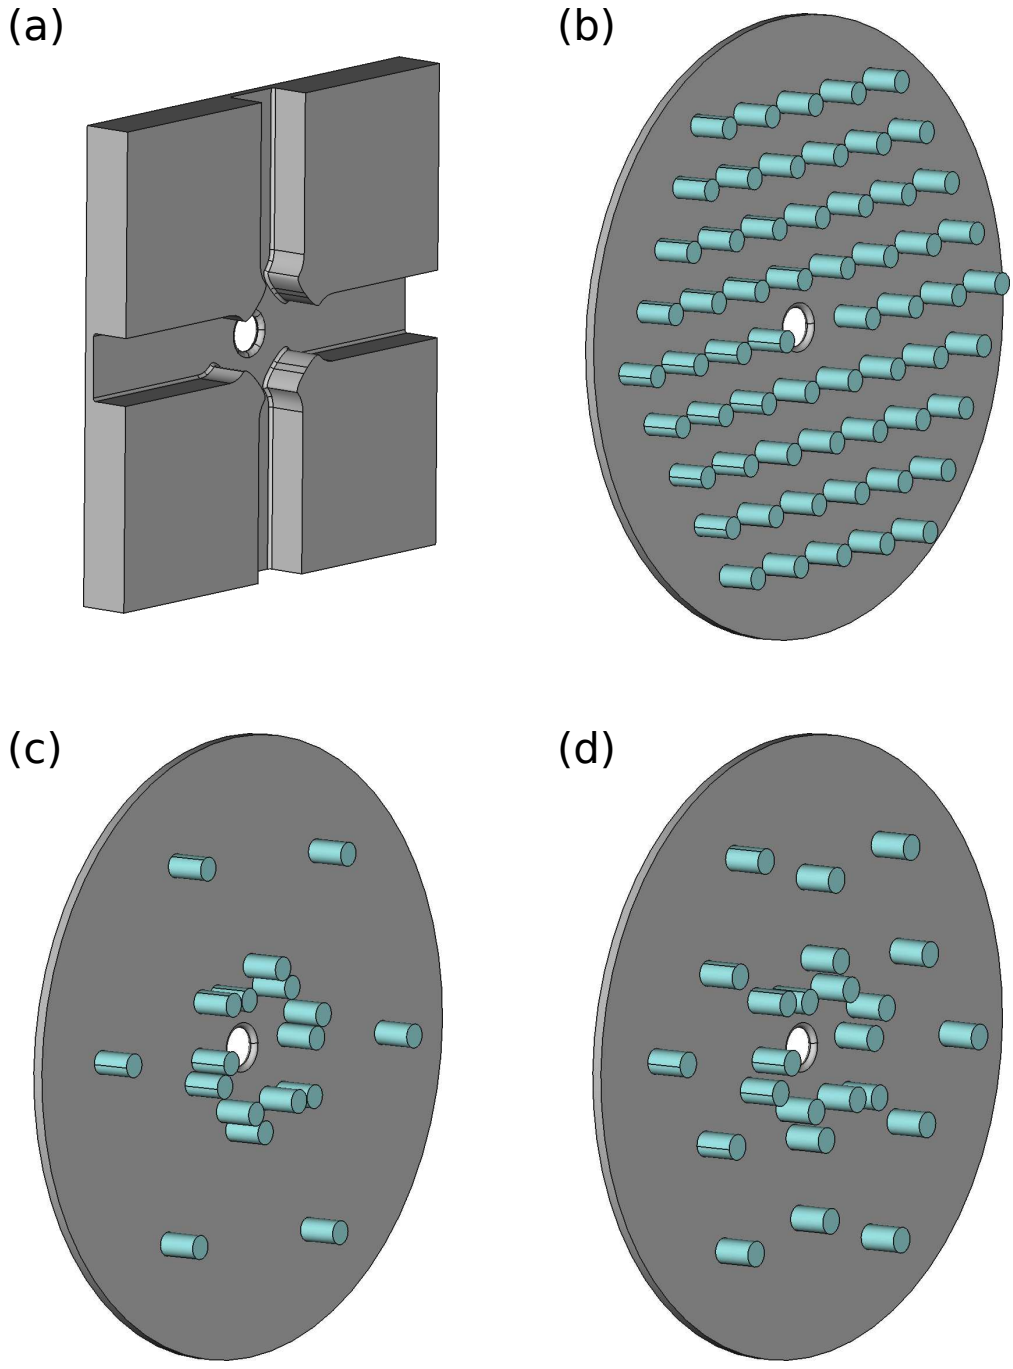

Figure 4.1: *Geometries of the cavities analyzed and compared in this Chapter (shown with beam tubes where $a = 3.15mm$ and $d = 1.67mm$). Shorthand for each of the cavities will be (a) CLIC, (b) Tri-4-Sapphire (triangular lattice, 4 layers of sapphire rods), (c) Opt-18-Sapphire (18-sapphire-rod optimized structure), and (d) Opt-24-Sapphire (24-sapphire-rod optimized structure). Also analyzed in this Chapter is the empty cavity; i.e. the geometry of (a), (b), and (c) without any rods. Multicell cavities are formed by layering the above structures and adding one iris plate to the final layer.*

simulation box of similar dimensions) below the trapped accelerating mode (these modes are heavily damped). There are eigensolver methods that, in theory, allow the calculation of interior eigenvalues without calculating the lowest frequency modes; however, in our experience, the convergence of these methods is lousy and/or inconsistent for the problems of interest.

The time-domain mode extraction technique has been the most useful for our simulations. In this method, electromagnetic fields are excited by a current source in a time-domain simulation. The spatial and frequency contents of the current source decide which modes in the simulation are excited to what amplitudes. The excited subset of eigenmodes can be extracted by analyzing the resulting oscillating fields in the simulation. This technique is well-suited to searching for interior-frequency eigenmodes through the use of current-source frequency filters (e.g. a Gaussian-modulated sinusoidal current excitation excludes frequencies that are several standard deviations from the central frequency). The accelerating mode for HDPhC cavities is found very effectively by this method. Of course this technique also has its drawbacks. For example, when searching for HOMs with low Q-factors, the mode density problem requires a narrow frequency filter; this in turn requires a long (in time) current excitation. In the time that it takes for the current source to finish, the mode of interest has decayed significantly and is difficult to extract.

4.1.2 Wakefield calculations

In practice, wakefield damping techniques require some kind of electromagnetically absorbing material; these materials should be simulated directly in order to characterize the wakefield for a given structure accurately. However, in most damping schemes (e.g. waveguide damping and PhC damping) the absorbers are far from the beam axis (and the extent of the accelerating mode), and thus should ideally absorb all radiation they encounter. Therefore, in simulation, a realistic absorber can be replaced by a more ideal (and less computationally costly) absorber and the question of experimental absorption can be left as a separate problem.

Unfortunately, our studies have shown that the most widely-used (and arguably best) absorber, the perfectly matched layer (PML), is subject to an insidious numerical instability. The

instability seems to manifest in problems that involve resonant cavities coupled to waveguides that are terminated by PMLs; i.e. the problems of interest to this thesis. The cause of this instability is still being debated in the literature. In some cases, the instability becomes a problem only at late simulation times, allowing the extraction of useful information up to the time that the fields from the instability dominate the physical fields. This has been the case in most HDPhC cavity simulations; therefore, PMLs were used in the closed cavity study presented next. Unfortunately, the PML instability in the CLIC cavity grows much faster, and so other techniques were devised to draw comparisons in the multicell cavity setting.

4.2 Closed cavity

We recently performed a study that compared the HDPhC designs with the conducting pillbox in the closed cavity case (i.e. without beam tubes) [81]. As a proof-of-principle, it showed that HDPhC cavities do provide significant damping for HOMs existing in the pillbox, resulting in greatly reduced wakefields. The results also suggested that the high-Q_{rad} optimized HDPhC cavities can suppress wakefields more than the triangular-lattice-based HDPhC structure.

4.2.1 Accelerating mode properties

Table 4.1 summarizes the properties of the accelerating mode in both closed HDPhC cavities and the closed pillbox. All data has been calculated from 2D simulations with a 10 GHz accelerating mode frequency. The Q_{metal} of each of the structures is as predicted from the analytic pillbox fields and Eq. 3.13 for the HDPhC cavities. For the HDPhC examples given, power losses generally increase compared to the pillbox (lowering the shunt impedance, hence accelerating efficiency); Q_{total} values remain similar because the stored energy in the HDPhC structures has also increased (based on $R_{\text{shunt}}/Q_{\text{total}}$ values, which we recall are proportional to V_{acc}^2/U).

A 2D study showed that adding layers to the triangular lattice past 3 does not increase $R_{\text{shunt}}/Q_{\text{total}}$ (the Q_{rad} rises exponentially as expected—see Table 4.2); thus, for very large Q_{rad}, the triangular lattice wins out (in terms of accelerating efficiency) over the current optimized

Table 4.1: *Figures of merit for the accelerating mode in relevant closed cavities (no beam tubes) at 12 GHz. Sapphire loss tangent was set to* 10^{-4}*. All values were calculated from 2D simulations.*

	Pillbox	Tri-4-Sapphire	Tri-X-Sapphire	Opt-18-Sapphire	Opt-24-Sapphire
Q_{metal}	9,000	21,000	21,000	21,000	21,000
Q_{rad}	∞	24,000	∞	25,000	193,000
Q_{diel}	∞	34,000	34,000	19,000	24,000
Q_{total}	9,000	8,000	13,000	7,000	11,000
R_{shunt} (MΩ)	1.8	1.1	1.6	0.6	1.1
$R_{\mathrm{shunt}}/Q_{\mathrm{total}}$ (Ω)	195	125	125	82	105
$cB_{\mathrm{surf,max}}/E_{\mathrm{acc}}$	0.91	1.32	1.32	1.60	1.53

structures. Data for this case is given by the Tri-X-Sapphire column in Table 4.1. The optimized structures have a larger energy density within the dielectric rods, lowering Q_{diel}; this becomes a problem for loss tangents near 10^{-4}. Also, the HDPhC structures tend to enhance the magnetic field in the vicinity of the inner rods (and thus on the metal endplates). This could be a disadvantage considering the growing evidence linking maximum surface magnetic fields to E_{acc} limits.

4.2.2 Wake potentials

A Gaussian bunch of $\sigma_z = 1.2$mm was used to excite the cavities at a transverse offset $r' = 1.2$mm, $\phi' = 0$. Rings of test charges followed the bunch (a la Fig. 2.4) in order to analyze the multipole contributions. It should be noted that for the closed cavities, the azimuthal multipole wake potential forms of Eqs. 2.33 and 2.34 do not apply [59, 80]. For example, in the closed cavity, even the accelerating mode gives a transverse kick to particles with some transverse offset. Practically, we care about beam tubes, in which case the transverse kick due to monopole modes becomes negligible. Therefore, in the following, we analyze only the monopole and dipole contributions to the longitudinal and transverse wake potentials, respectively.

Table 4.2: *The (exponential) dependence of* Q_{rad} *on the number of layers X for the 2D Tri-X-Sapphire cavity.*

Layers	3	4	5	6
Q_{rad}	2,300	24,000	250,000	2.6×10^6

The monopole wake potential is dominated by the accelerating mode; therefore, to analyze the harmful part (in the time domain), its contribution was subtracted. The subtractions were performed by fitting a decaying sinusoid (oscillating at the fundamental frequency) to the total monopole wake potential such that the difference was minimized. The resulting amplitudes and decay parameters of the sinusoids (for all simulated cavities) matched well with the figures in Table 4.1. Details and the effectiveness of the subtraction are shown in Fig. 4.2 for the Tri-4-Sapphire closed cavity. A value of $R_{\text{shunt}}/Q = 125$ for the Tri-4-Sapphire accelerating mode at 10 GHz gives a loss factor per unit length of $k = 130$ V/pC/m which matches well with the amplitude of the sinusoid shown in Fig. 4.2a (recall that the wake potential goes as $2k$).

Figure 4.3 compares the monopole longitudinal wake potentials in the pillbox and HDPhC closed cavities and Fig. 4.4 shows the impedance of each curve in Fig. 4.3. Figure 4.5 compares the transverse dipole wake potentials. Strong damping of HOMs in HDPhC cavities is evident when compared with the pillbox. Furthermore, the Opt-(18,24)-Sapphire structures show greater suppression of monopole HOMs than the Tri-4-Sapphire structure. Between the HDPhC closed cavities, there is not much difference in the damping of dipole transverse wakefields (this changes significantly when beam tubes are added). At large s, all HDPhC wakes show the characteristic flattening due to the persistent component (arising from waveguide cutoff frequencies).

For the triangular-lattice-based structure, the dependence of the wake potential on the number of lattice layers is shown in Fig. 4.6. For each wake in Fig. 4.6, the associated frequency spectrum (or wake impedance) differs only by the widths of the peaks. This implies that the Q_{rad} of HOMs are increasing (along with the Q_{rad} of the accelerating mode). Consider the first HOM seen in the longitudinal monopole impedance of the Tri-4-Sapphire cavity (Fig. 4.2b). The time-domain mode extraction technique isolated the mode associated with this peak. The mode looks very much like the accelerating mode in the transverse direction, but with a nonzero $k_z = \pi/L$ (E_z changes sign once across the cavity). Examining the band structure of a triangular lattice of infinite-length rods in 3D showed that the accelerating mode bandgap persists to the above value of k_z (although the magnitude of the gap decreases). Therefore, increasing the layer count increases

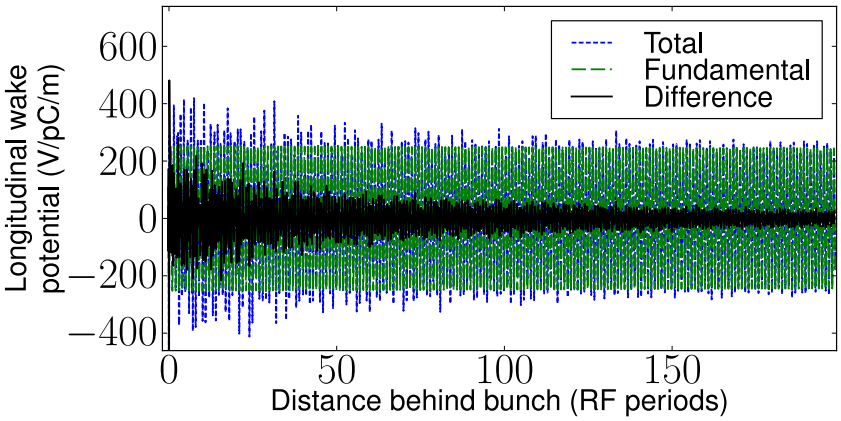

(a)

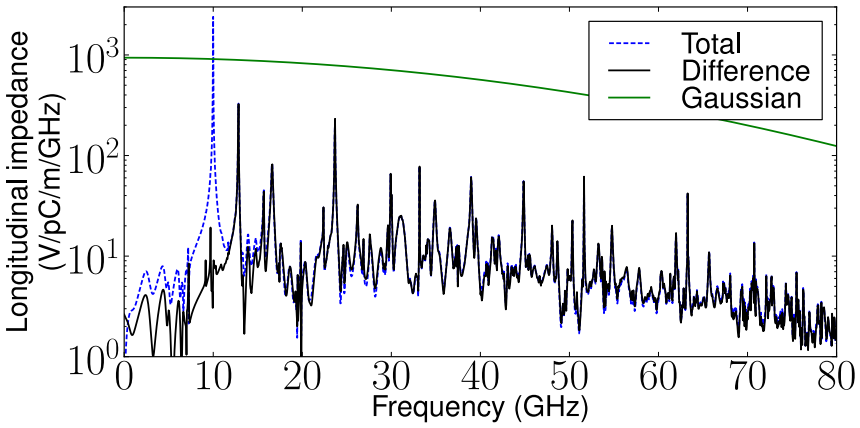

(b)

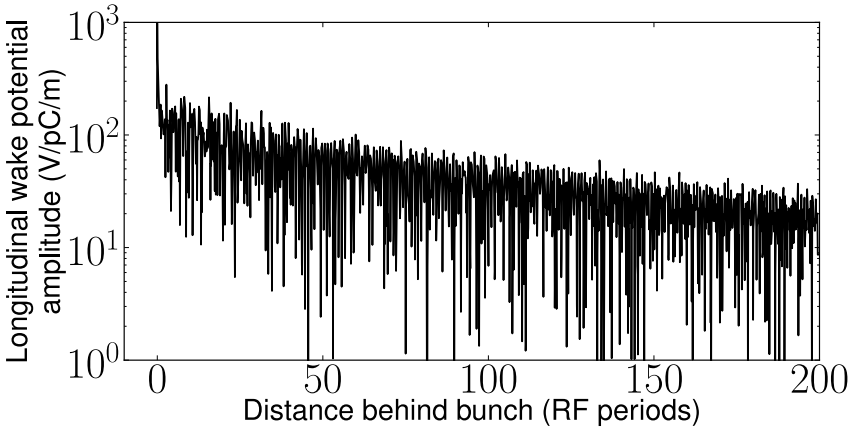

(c)

Figure 4.2: *The monopole wake potential in the Tri-4-Sapphire closed cavity is dominated by the accelerating mode. (a) The total monopole wake potential, the contribution of the fundamental (accelerating) mode, and the difference between the two. (b) The monopole longitudinal impedance (Fourier transform of (a))—total and with fundamental removed; the Gaussian frequency spectrum of the excitation bunch (arbitrary amplitude) is also shown. (c) Amplitude (absolute value of the zero-derivative points) of the difference wake potential.*

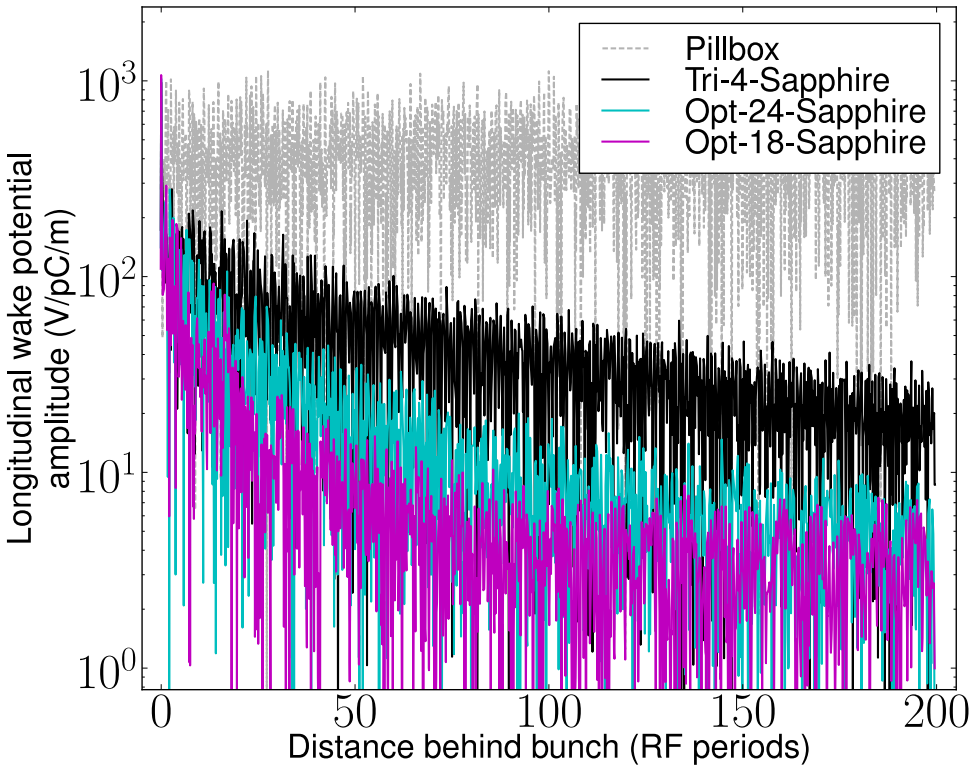

Figure 4.3: *Amplitude of the longitudinal monopole wake potential in the closed pillbox and the closed HDPhC cavities.*

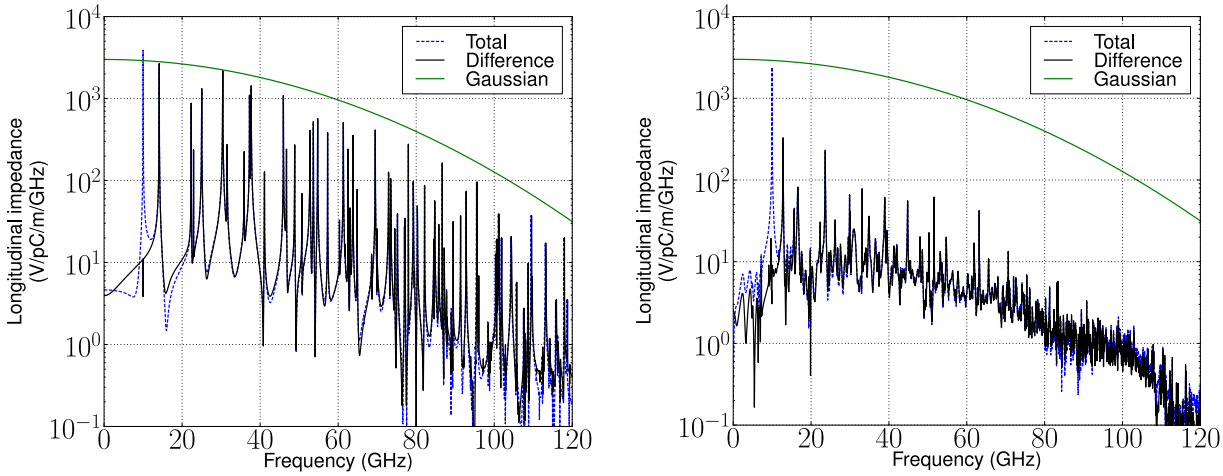

(a) *Pillbox*

(b) *Tri-4-Sapphire*

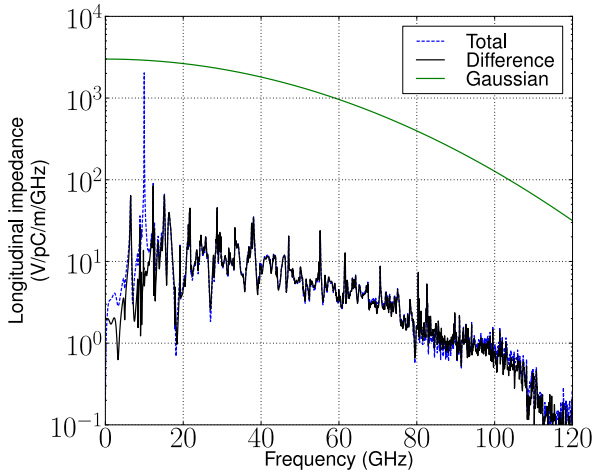

(c) *Opt-24-Sapphire*

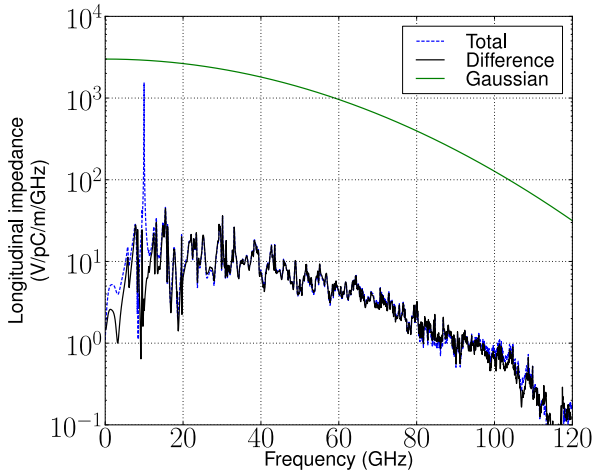

(d) *Opt-18-Sapphire*

Figure 4.4: *The monopole longitudinal impedances in the different closed cavities.*

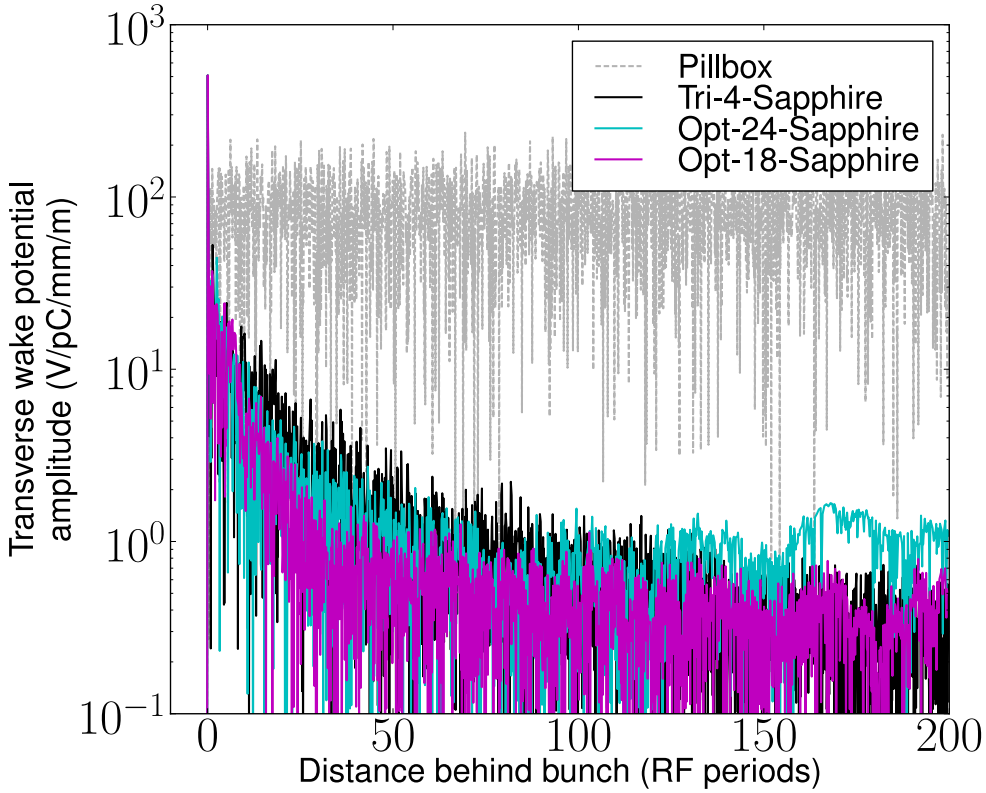

Figure 4.5: *Amplitude of the transverse dipole wake potential in the closed pillbox and the closed HDPhC cavities. A low-level late-time numerical instability can be seen affecting the Opt-24-Sapphire simulation.*

the confinement of this mode, which increases the wake potential. Other modes were too difficult to extract by the time-domain method for the reasons described previously.

4.3 With beam tubes/comparison with CLIC

The previous section showed that wakefields are greatly reduced in closed HDPhC cavities when compared with the cylindrical pillbox. However, a more useful comparison would test HDPhC wakefield damping against other damping mechanisms used in state-of-the-art metallic cavity design. To this end we have recently acquired the CAD geometry specifications for the baseline CLIC cavity design [28]. Comparison with the CLIC structure has forced us to examine the properties of HDPhC cavities in a more practical multicell setting where the effects of beam tubes and coupled cavities must be addressed. The results in this section are the main results of this thesis. In short, with reference to the CLIC cavity, we find that in HDPhC cavities: (1) accelerating efficiencies are lower (except in the lattice structure where the number of layers is large), (2) maximum surface magnetic fields are higher for the optimized structures but lower in the lattice structure, (3) dipole transverse wakefields are higher, and (4) monopole longitudinal wakes are lower.

4.3.1 CLIC cavity overview

The full 26-cell CLIC cavity including input/output power coupling cells and beam tubes is shown in Fig. 4.8. This structure is the result of a complex optimization process which tries to minimize surface fields, short- and long-range wakefields, and power input requirements while maximizing accelerating efficiency and beam current. It is a constant-gradient cavity where the iris geometry tapers linearly from $a = 3.15$mm and $d = 1.67$mm in the first cell to $a = 2.35$mm and $d = 1.0$mm in the final cell. As a result, the group velocity tapers nearly linearly from 1.65% to 0.83% c. Each cell is equipped with four identical (in cross-section) HOM damping waveguides, the cutoff frequency of which has been chosen to maximize the damping of the lowest dipole HOM and the loss factor (or shunt impedance) of the accelerating mode. Table 4.3 summarizes the relevant properties of the CLIC design and nominal electron beam it would accelerate [27].

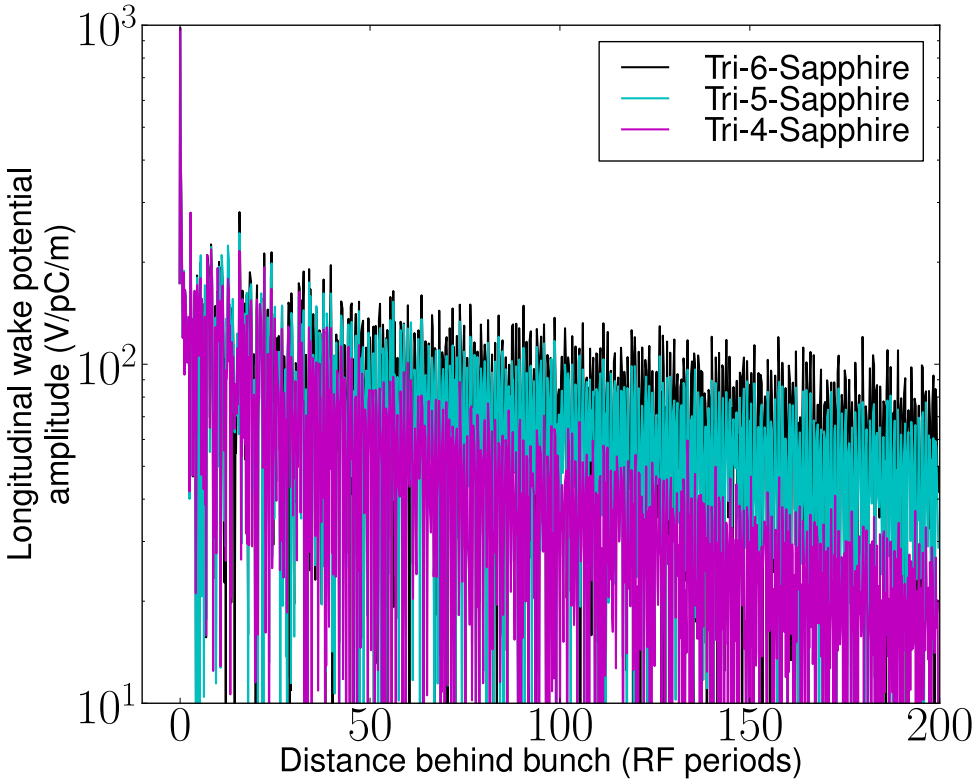

Figure 4.6: *Amplitude of the longitudinal monopole wake potential (accelerating mode has been subtracted out as in Fig. 4.2) in the triangular lattice structure (of sapphire rods) for different numbers of layers.*

Table 4.3: *Properties of the CLIC accelerating structure and beam [27].*

Frequency (GHz)	11.994
Phase advance	$2\pi/3$
Cell length (mm)	8.33
Cells	26
Average E_{acc} (MV/m)	100
Bunch population	3.72×10^9
Bunch spacing (RF periods)	6
Beam current (A)	1.19
Input power (MW)	61.3

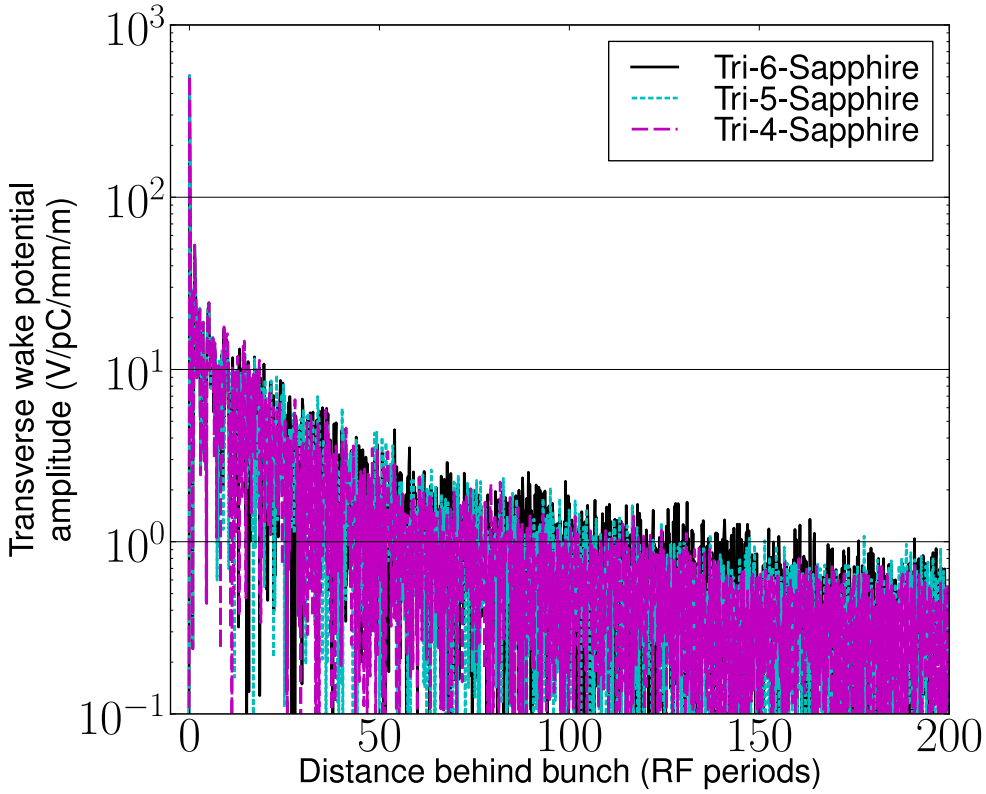

Figure 4.7: *Amplitude of the transverse dipole wake potential in the triangular lattice structure (of sapphire rods) for different numbers of layers.*

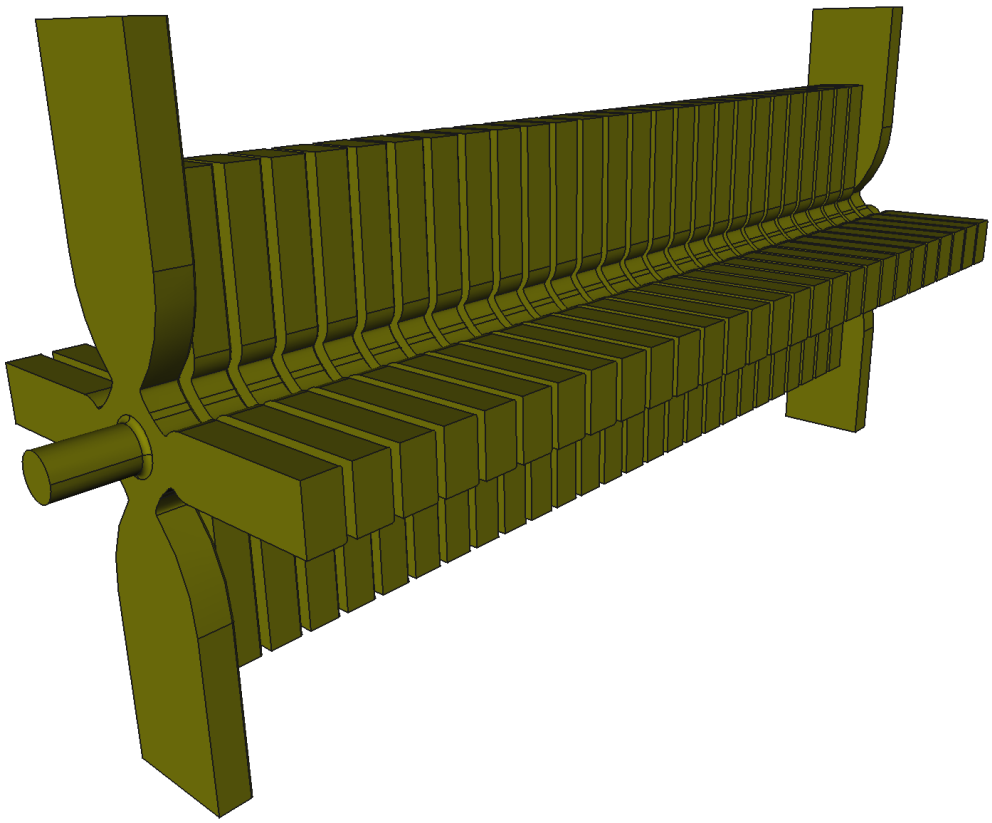

Figure 4.8: *The full CLIC accelerating cavity including power coupling end-cells and beam tubes.*

Table 4.4: *Figures of merit for the accelerating mode in relevant periodic single cell cavities.* $a = 3.15mm$, $d = 1.67mm$, $\phi = 2\pi/3$. *All simulations performed at* $\Delta z/d = 8$. *Used loss tangent of* 10^{-4}.

	Pillbox	CLIC	Tri-4-Sapphire	Opt-18-Sapphire	Opt-24-Sapphire
v_g/c (%)	1.83	1.65	1.16	0.78	0.95
Q_{metal}	6,700	5,900	11,400	11,400	11,400
Q_{rad}	∞	∞	26,600	3,800	16,600
Q_{diel}	∞	∞	67,000	39,000	45,500
Q_{total}	6,700	5,900	7,100	2,700	5,900
r_{shunt} (MΩ/m)	106	82	70	18	47
k (V/pC/m)	298	260	187	125	151
$E_{surf,metal,max}/E_{acc}$	1.93	1.96	1.93	1.93	1.93
$cB_{surf,metal,max}/E_{acc}$	1.0	1.54	1.49	1.73	1.76
$E_{surf,diel,max}/E_{acc}$	—	—	0.54	0.64	0.62
$cB_{surf,diel,max}/E_{acc}$	—	—	1.26	1.79	1.34
$E_{diel,max}/E_{acc}$	—	—	0.60	0.79	0.72

4.3.2 Accelerating mode comparison

Figures of merit for the accelerating mode were first compared for a periodic single cell with iris geometry $a = 3.15$mm and $d = 1.67$mm and are summarized in Table 4.4. This geometry applies to the first cell of the CLIC 26-cell structure. Figures 4.9 and 4.10 show the absolute values of E_z and $\mathbf{B}_\perp$ on the z-midplane of the relevant periodic single cell cavities. The z-dependences of the fields are very similar amongst the different cavity types, as expected based on the common iris geometry in each structure (refer to the pillbox fields shown in Fig. 2.10).

The maximum surface magnetic field occurs on the innermost radial walls of the CLIC cavity; the elliptic curvature of this feature was carefully chosen to minimize cB_{surf}/E_{acc}. In the HDPhC cavities, just as in the closed case, $B_{metal,surf,max}$ occurs where the innermost rods abut the conducting endplates. This surface field is higher in the HDPhC cavities for a given E_{acc}. Since the maximum occurs at the interface between dielectric and conductor, the method used to hold the rods in place will require careful consideration. For example, a brazing material could melt due to endplate heating, allowing the rods to shift; this could pose a problem for the optimized cavities which are very sensitive to inner rod displacements (one possible solution would be to secure the rods simply by endplate pressure, as in [56]). On the other hand, covering the region of conductor

suffering the maximum magnetic field with dielectric could suppress the breakdown mechanism. All of this remains to be seen in experiment. The maximum surface electric field occurs on the iris in all cases, explaining the uniformity of $E_{\mathrm{metal,surf,max}}/E_{\mathrm{acc}}$ across all cavity types.

The addition of beam tubes is seen to decrease the Q_{rad} of the optimized structures by a significant amount while the Q_{rad} of the 60-rod lattice is unaffected. For example (compare Tables 4.1 and 4.4), at a phase advance of $2\pi/3$, the 18-rod sapphire Q_{rad} drops from 24,000 to 3,800 (varying with ϕ from 7,300 at $\phi = 0$ to 3,200 at $\phi = \pi$). This is not such a surprise considering the sensitivity analysis summarized in Fig. 3.9. It remains to be seen whether further optimization in 3D can regain the original 2D Q_{rad}. Optimizations in the presence of beam tubes will be difficult because, for a fixed beam tube geometry, the target resonant frequency must be maintained at each iteration—this restriction is not easily incorporated into a cost function. If the target frequency is not maintained throughout, then a post-optimization transverse scaling will be required, resulting in a different beam tube radius.

Assuming that 3D HDPhC optimizations could be performed successfully, the majority of the power losses would then occur in the conducting iris plates. Replacing the Q_{rad} values in Table 4.4 with the original 2D optimized values and assuming that the changes in loss factors are negligible results in shunt impedances of $r_{\mathrm{shunt}} = 43$ MΩ/m and $r_{\mathrm{shunt}} = 70$ MΩ/m for the Opt-18-Sapphire and Opt-24-Sapphire cavities, respectively (see Table 4.5). Also included is the hypothetical shunt impedance for Tri-X-Sapphire cavity, where X is large enough such that the radiation losses of the accelerating mode are insignificant (i.e. $Q_{\mathrm{rad}} \to \infty$). Of the cavities with wakefield damping, this gives the Tri-X-Sapphire cavity the largest shunt impedance; however, as seen in the closed cavity results of the previous section (and in the HOM discussion at the end of this chapter), increasing the number of lattice layers could increase long-range wakefields.

Figures of merit for the different cavity types as multicell structures were also calculated using the hypothetical Q_{rad}. The predictions shown in Table 4.6 and Figs. 4.11 and 4.12 were made using the power-flow analysis of Section 2.5.6 and the CLIC beam parameters given in Table 4.3. A 26-cell constant-gradient structure was approximated for each cavity type based on a linear

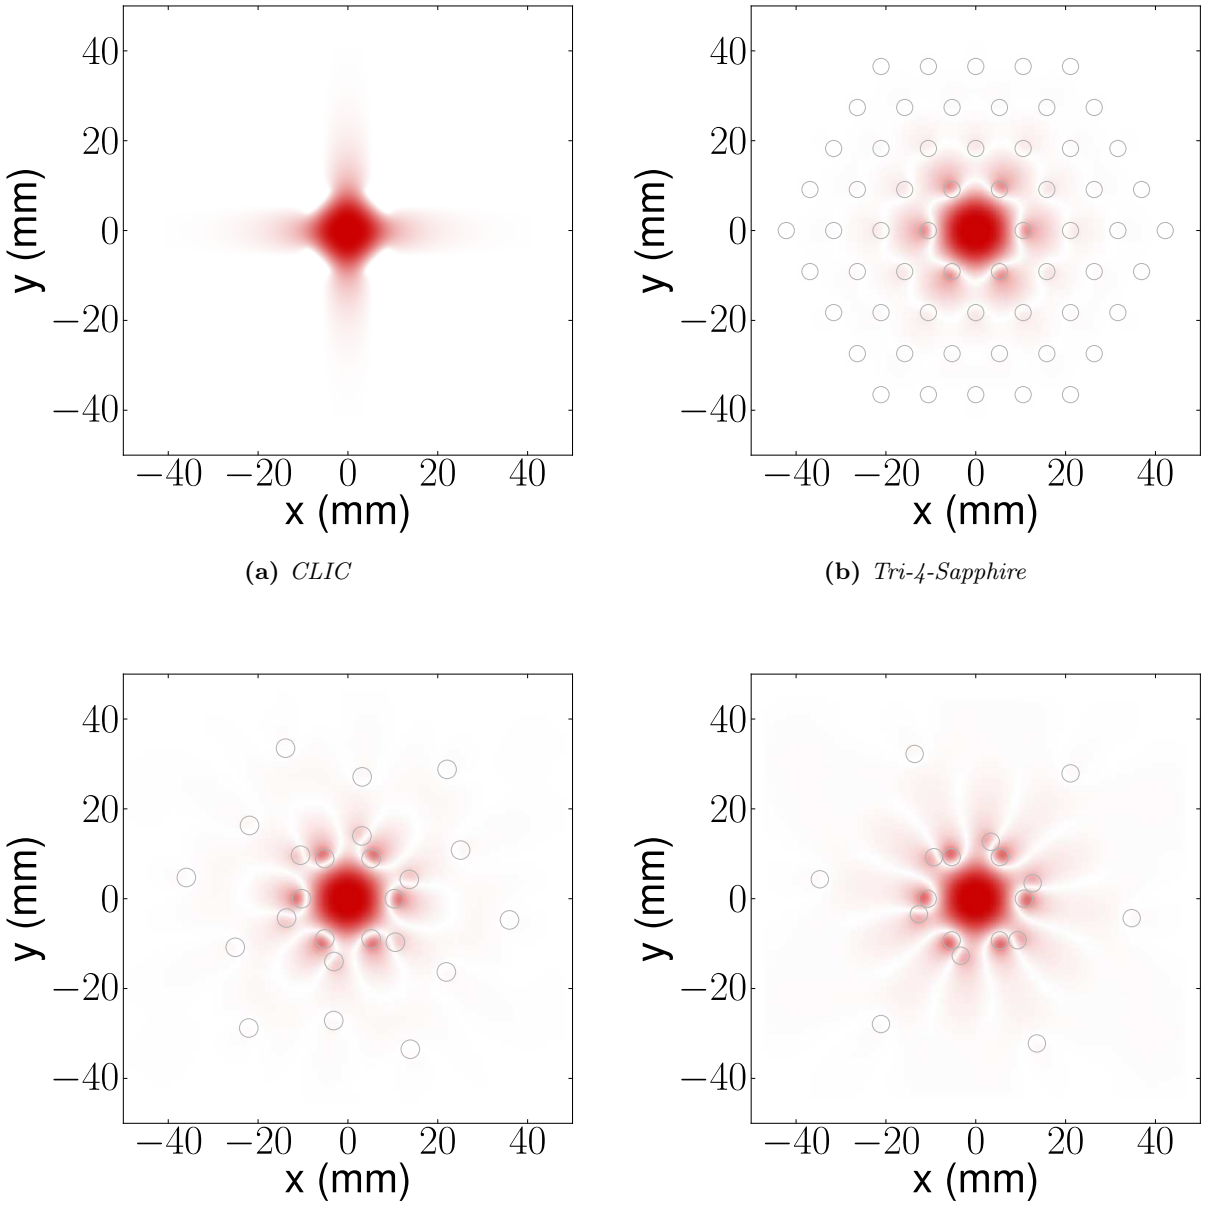

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.9: *Absolute value of E_z on the midplane (in z) of periodic single cell cavities.*

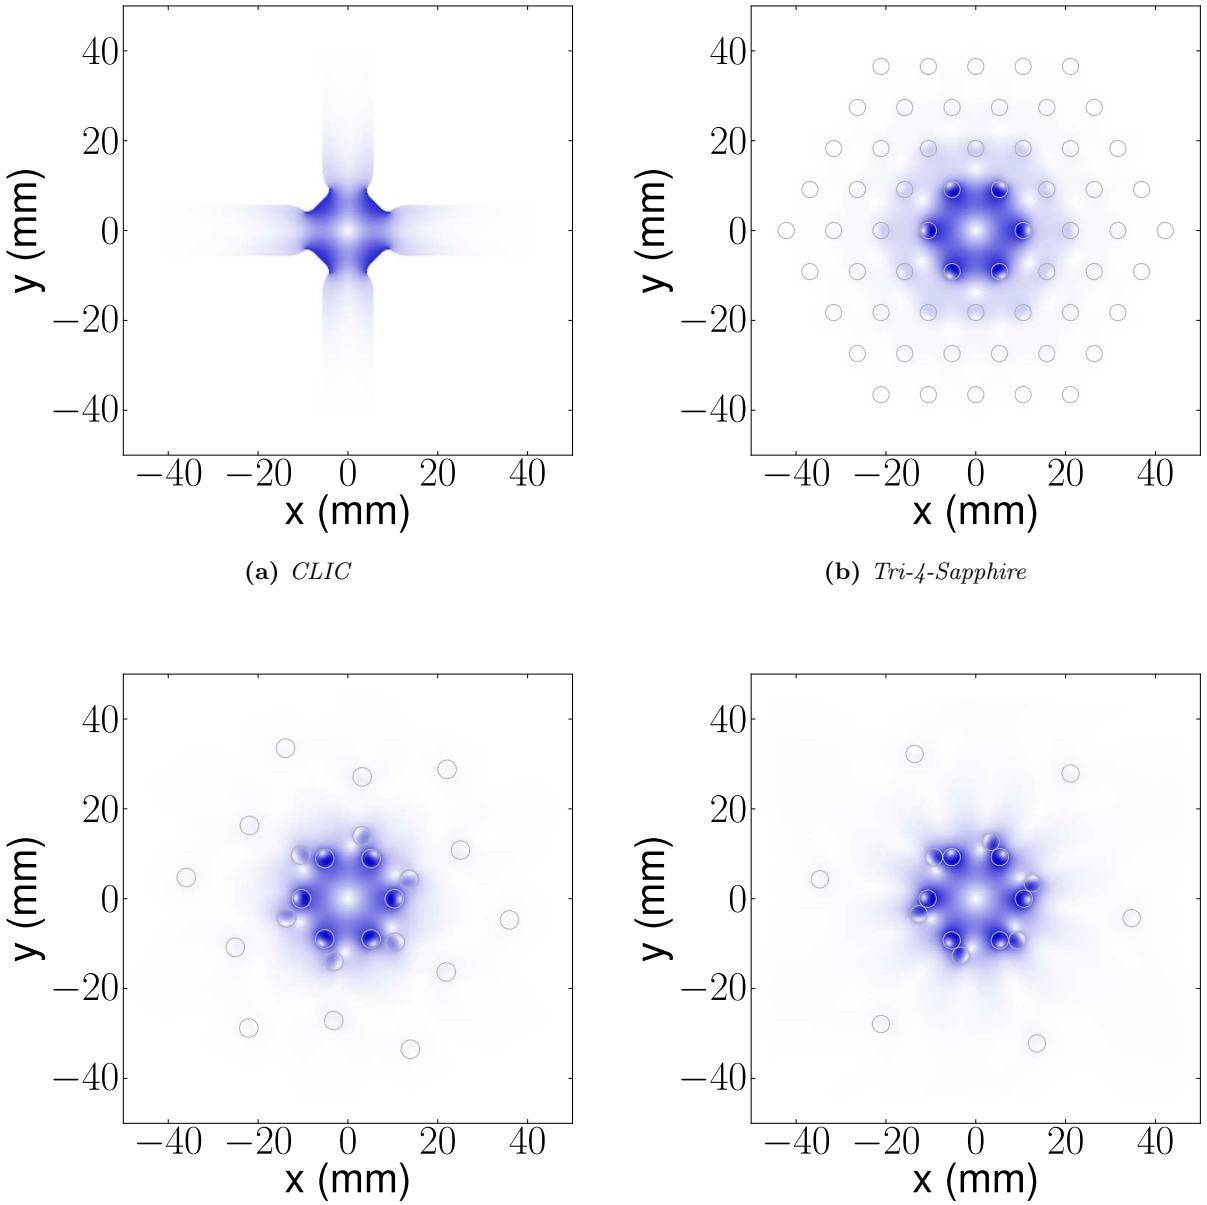

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.10: *Absolute value of* $\mathbf{B}_\perp$ *on the midplane (in z) of periodic single cell cavities.*

Table 4.5: *Hypothetical figures of merit for the accelerating mode given that X (for the Tri-X-Sapphire cavity) is large enough such that radiation losses are insignificant and that 3D optimizations reproduce 2D Q_{rad} values for the optimized cavities.*

	Pillbox	CLIC	Tri-X-Sapphire	Opt-18-Sapphire	Opt-24-Sapphire
$\mathrm{Q}_{\mathrm{metal}}$	6,700	5,900	11,400	11,400	11,400
$\mathrm{Q}_{\mathrm{rad}}$	∞	∞	∞	25,000	193,000
$\mathrm{Q}_{\mathrm{diel}}$	∞	∞	67,000	39,000	45,500
$\mathrm{Q}_{\mathrm{total}}$	6,700	5,900	9,700	6,500	8,700
r_{shunt} (MΩ/m)	106	82	97	43	70
k (V/pC/m)	298	260	187	125	151

interpolation of iris parameters and individual cell figures of merit between the first and last cell (the first and last cells for each cavity type were simulated). The iris parameters for the first and last cells matched those of the CLIC cavity. The power flow along each cavity was determined by a Runge-Kutta integration of Eq. 2.85 and is shown in Fig. 4.11; E_{acc} as a function of cell number was then derived from $E_{\mathrm{acc}} = 2\sqrt{kP/v_g}$ and is shown in Fig. 4.12 (the CLIC curves agree well with results in Ref. [27]). The steady-state accelerating efficiency η was determined from

$$\eta = \sum_{i=1}^{26} \frac{I E_{i,\mathrm{acc}} L}{P_{\mathrm{in}}} \tag{4.1}$$

where $E_{i,\mathrm{acc}}$ is the accelerating gradient in cell i.

The energy of the accelerating mode in the HDPhC cavities is more dispersed transversely than in the pillbox or CLIC cells, lowering the group velocity and loss factor (for the same iris geometry). The lower group velocities add significantly to the cavity fill time. While the steady-state efficiencies for the Tri-(4,X)-Sapphire and Opt-24-Sapphire cavities listed in Table 4.6 are near and above the CLIC efficiency, the increased fill time will make the HDPhC cavities less efficient in a practical accelerating scheme that must use a pulse of RF input power. Pulse durations for recent CLIC designs are approximately 240 ns; after accounting for fill time and pulse rise and fall times, the CLIC study quotes a total RF-to-beam efficiency of $\approx 28\%$ [28]. This suggests that a more optimal multicell HDPhC cavity design would have a larger average beam tube radius to raise group velocities; this is beneficial for short-range wakefields but would decrease the shunt impedance (and thus the steady-state efficiency) and increase the power input requirement.

Table 4.6: *Figures of merit for constant-gradient 26-cell versions of each cavity type. Tri-X-Sapphire and optimized structure results use the hypothetical Q_{rad} values from Table 4.5.*

	Pillbox	CLIC	Tri-X-Sapphire	Tri-4-Sapphire	Opt-18-Sapphire	Opt-24-Sapphire
η (%)	45	43	44	41	37	42
t_{fill} (ns)	52	58	83	83	123	101

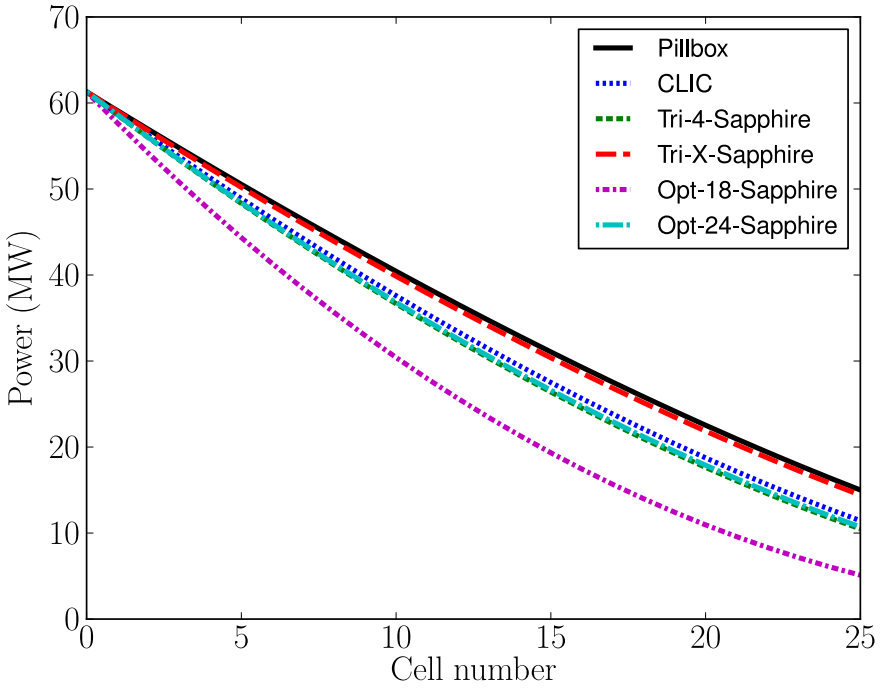

Figure 4.11: *Beam-loaded power flow along a 26-cell version of each cavity type, where iris geometries are identical to those in the CLIC cavity.*

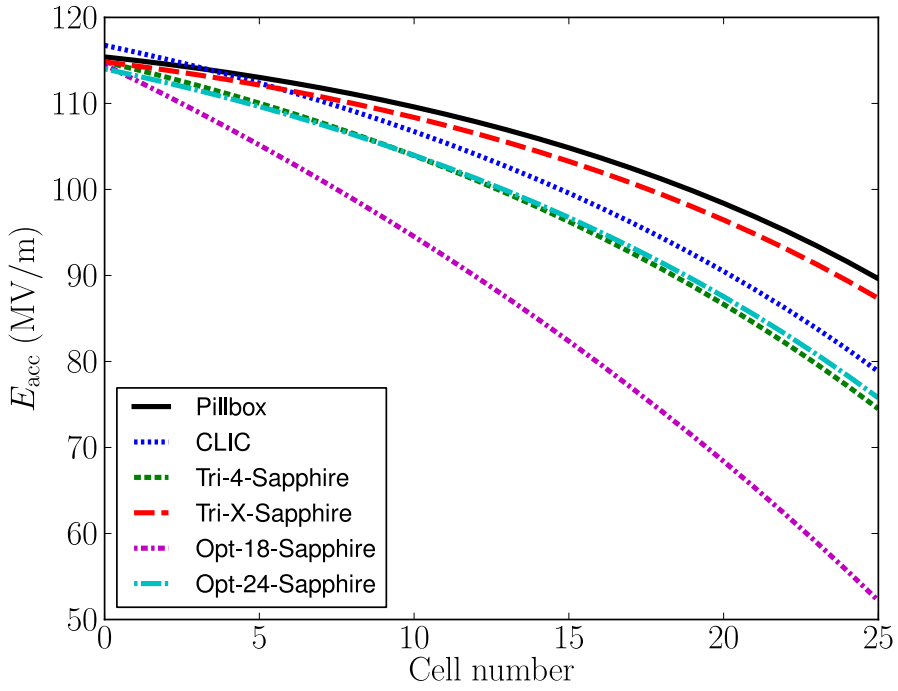

Figure 4.12: *Beam-loaded accelerating gradient (for the same input power) along a 26-cell version of each cavity type, where iris geometries are identical to those in the CLIC cavity.*

4.3.3 Wakefield simulations

Time-domain wakefield simulations were performed to approximate the wake potentials in a periodic single cell and were set up as follows. A Gaussian line excitation bunch of length $\sigma_z = 1\text{mm}$ at transverse position $r' = 1\text{mm}$, $\phi = 0$ was sent through an 8-cell cavity where each cell was identical. In all cases the beam tube geometry was $d = 1.67\text{mm}$ and $a = 3.15\text{mm}$ (this beam tube geometry corresponds to the first cell of the CLIC constant-gradient baseline structure). In all simulations, grid cells were cubic with $\Delta z = 0.3\text{mm}$ so that the excitation bunch cutoff frequency was simulated with $\approx 10\%$ accuracy according to Eq. 2.56.

Two different absorbing techniques were used in the wakefield simulations. The first was to extend the transverse simulation domain boundaries (as far as computational resources would allow) in order to emulate ideal "absorption" using the causality of the outwardly-propagating wakefields. The transverse boundaries (perfectly-conducting) were placed at a distance 125mm from the beam axis so that the wake potential in the range $0 < s < 250\text{mm}$ was unaffected by simulation-edge reflections; this distance is equivalent to approximately 10 accelerating mode oscillations (the CLIC cavity was designed to space bunches every 6 RF periods).

As mentioned previously in this Chapter, the CLIC cavity exhibited an overwhelming numerical instability when PMLs were placed in the damping waveguides. Thus, to examine longer-range wakefields (as compared to those acheivable using the extended domain technique), we employed numerically stable normal conductors. The ends of each damping waveguide (or the transverse edges of the HDPhC simulations) were filled with a block of conducting material with a special conductivity profile aimed at minimizing reflections off of the cavity-facing surface. The profile is given by

$$\sigma(w) = \sigma_{\max} \left(\frac{w}{W} \right)^2 \tag{4.2}$$

where w is the perpendicular distance from the inside surface within the conductor and W is the depth of the conducting block. The quadratic form resembles damping techniques used in practice [53], where a cone of absorbing material is placed at the end of a damping waveguide. Our

calculations used $W = 25\text{mm}$ and $\sigma_{\max}/\varepsilon_0 = 4.7\omega_{\text{acc}}$.

4.3.4 Extended-domain wakefield results

We first present the results from the extended-domain wakefield simulations. Total longitudinal wakes are shown in Figs. 4.13 and 4.14 (amplitudes); these plots represent the sum of all multipole contributions. Figure 4.13 details the subtraction of the accelerating mode contribution in each cavity type. For the mid-range longitudinal wake, the HDPhC cavities do not show a clear advantage.

Figure 4.15 shows the total radial wake potential directly behind the bunch (i.e. the total radial momentum kick normalized to the bunch charge and cavity length). While the HDPhC wakes are generally higher than those in the CLIC cavity, there are some interference effects between HOMs in the HDPhC cavities that decrease the transverse wake below that of the CLIC cavity very close to the excitation bunch. This raises an interesting question whether the greater mode-density in PhC cavities would benefit a detuning strategy for minimizing the wake potential at closely spaced multiples of an RF period behind an excitation bunch.

Finally, Figure 4.16 shows that (up to 10 RF periods), the dipole wakefields dominate the total radial kick behind the excitation bunch in all cavity types.

4.3.5 The effect of conducting absorbers

Conducting absorbers were placed at the ends of the damping waveguides in the CLIC cavity and at the transverse domain boundaries of the HDPhC simulations (see Fig. 4.17). The distance from cavity center to absorber edge was 40mm in the CLIC cavity. In each HDPhC simulation, a minimum distance of 4 rod radii was placed between the outermost rods and the absorber edges.

Conducting absorbers allow the wake potentials to be calculated out to large trailing distances (because of their numerical stability), but provide worse absorption when compared to the extended-domain technique. As a result, we have noticed that the wakefields in the CLIC structure change more dramatically than those in the HDPhC structures, indicating a greater sensitivity to the

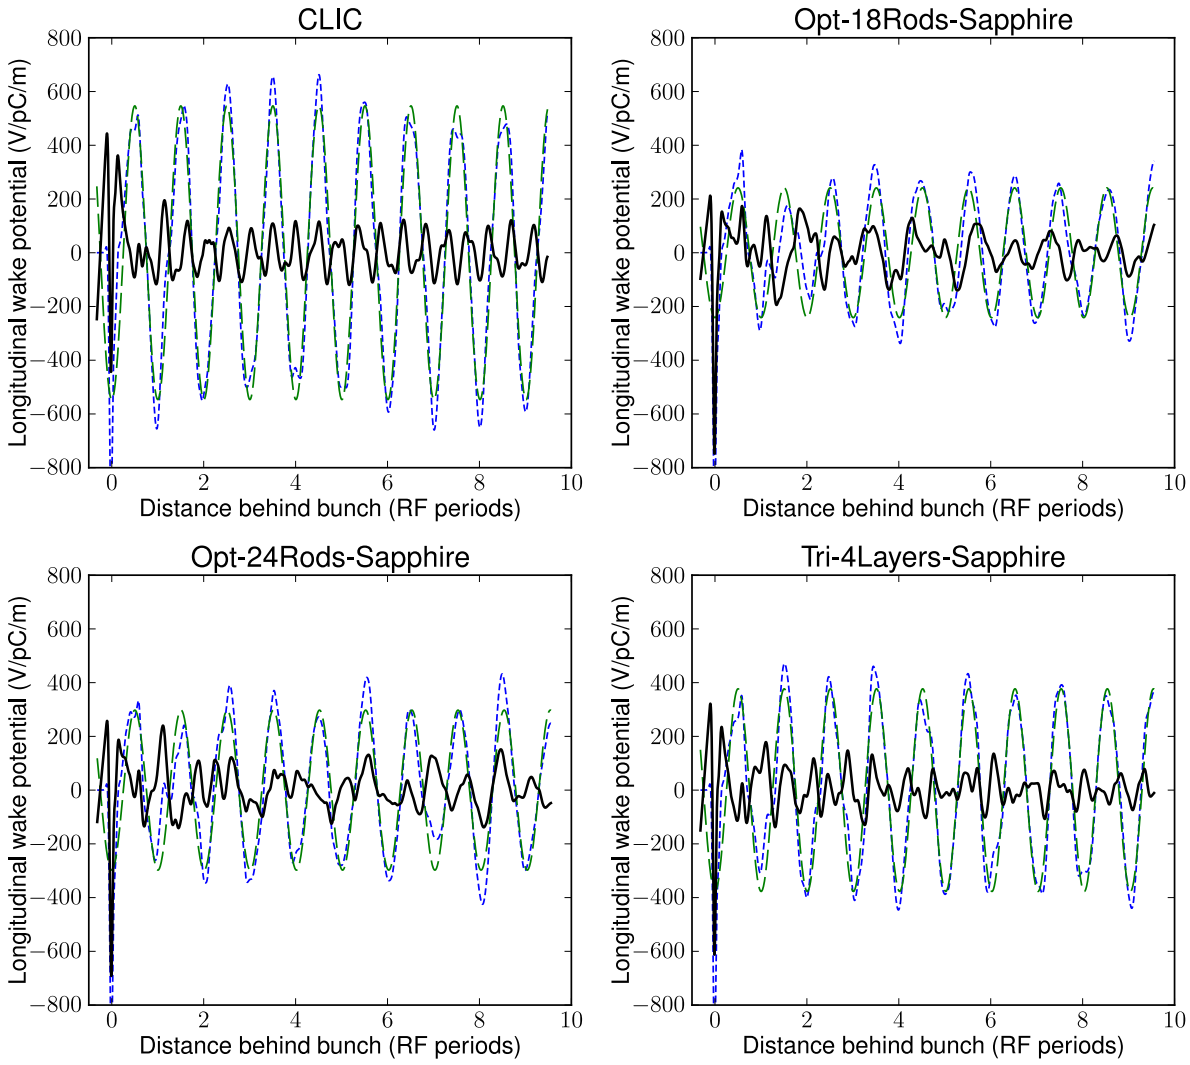

Figure 4.13: *In each plot, the blue short-dashed curve is the total longitudinal wake potential where $r = r' = 1mm$, $\phi = \phi' = 0$, and $\sigma_z = 1mm$; the green long-dashed curve is the best fit to the accelerating mode contribution; and the solid black line is the difference (i.e. the unwelcome portion of the longitudinal wake).*

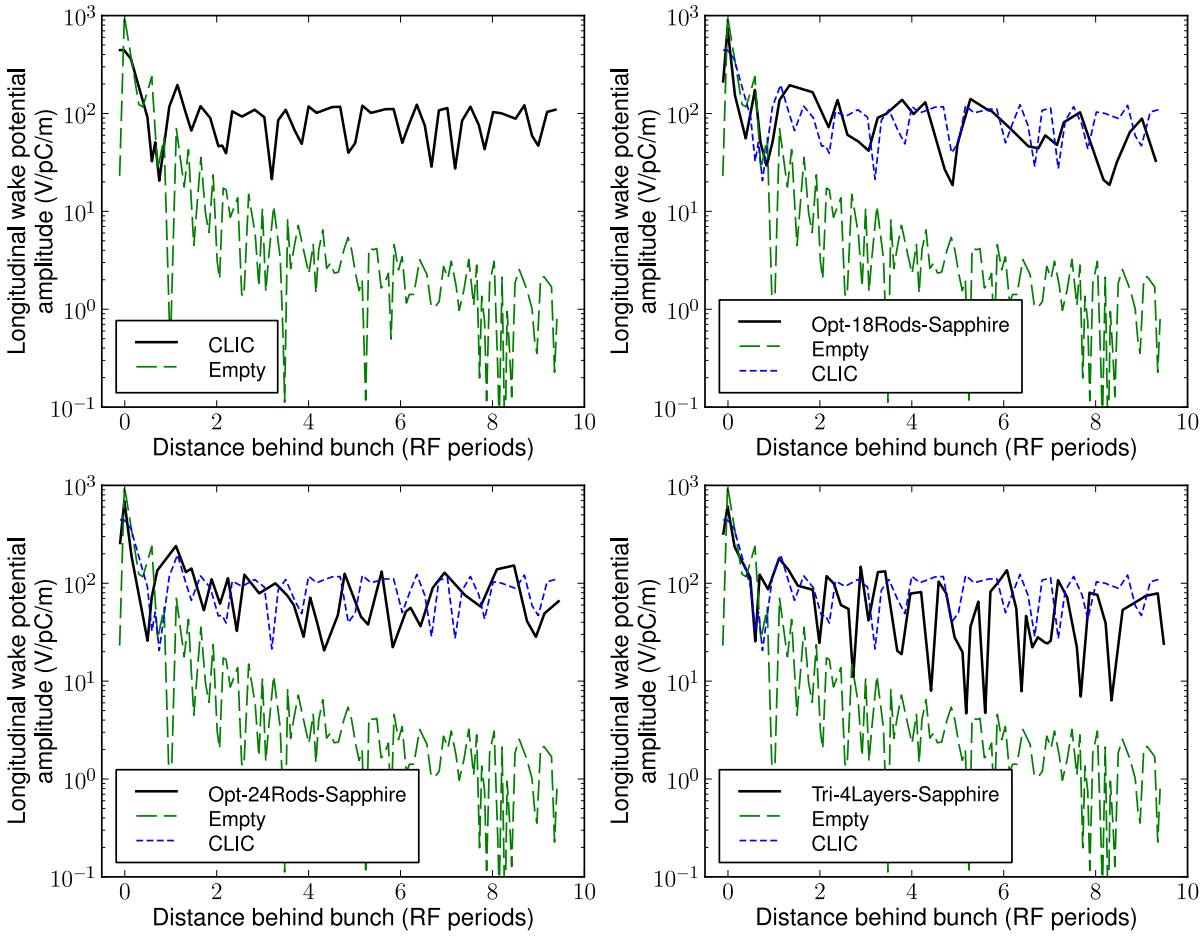

Figure 4.14: *A comparison of the amplitudes (zero-derivative points) of the subtracted longitudinal wake potentials from Fig. 4.13. Also included is the total longitudinal wake potential from the empty structure.*

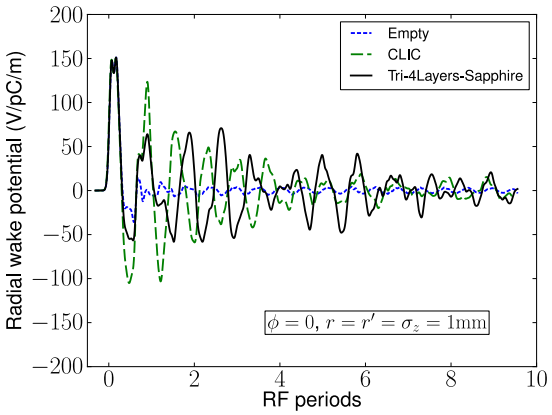

(a) *Tri-4-Sapphire*

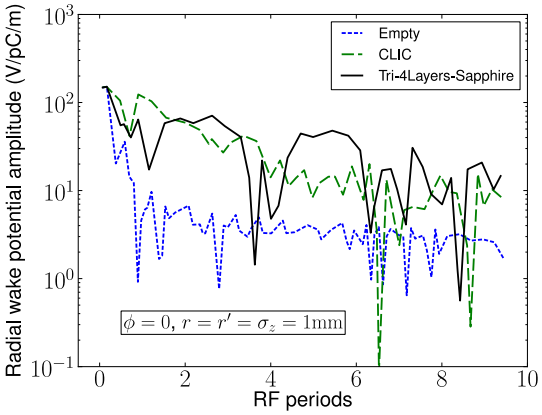

(b) *Tri-4-Sapphire*

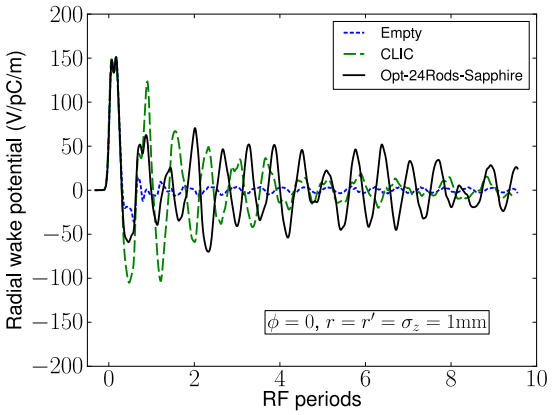

(c) *Opt-24-Sapphire*

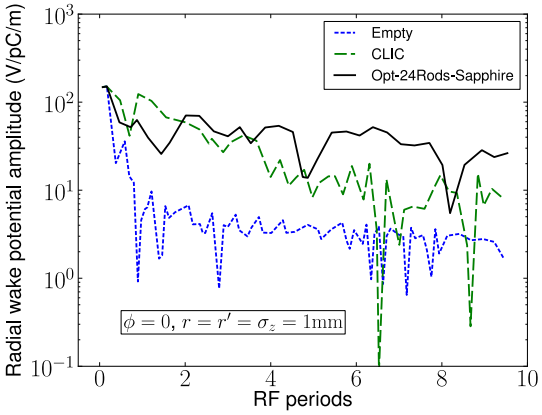

(d) *Opt-24-Sapphire*

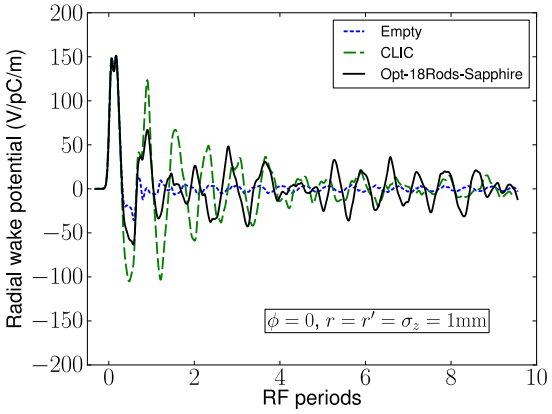

(e) *Opt-18-Sapphire*

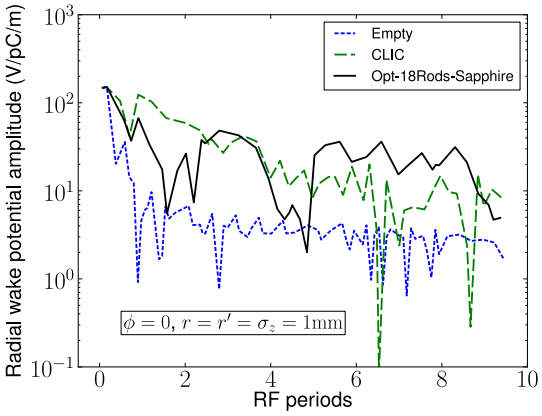

(f) *Opt-18-Sapphire*

Figure 4.15: *Total radial wake potential in extended-domain 8-cell cavities.*

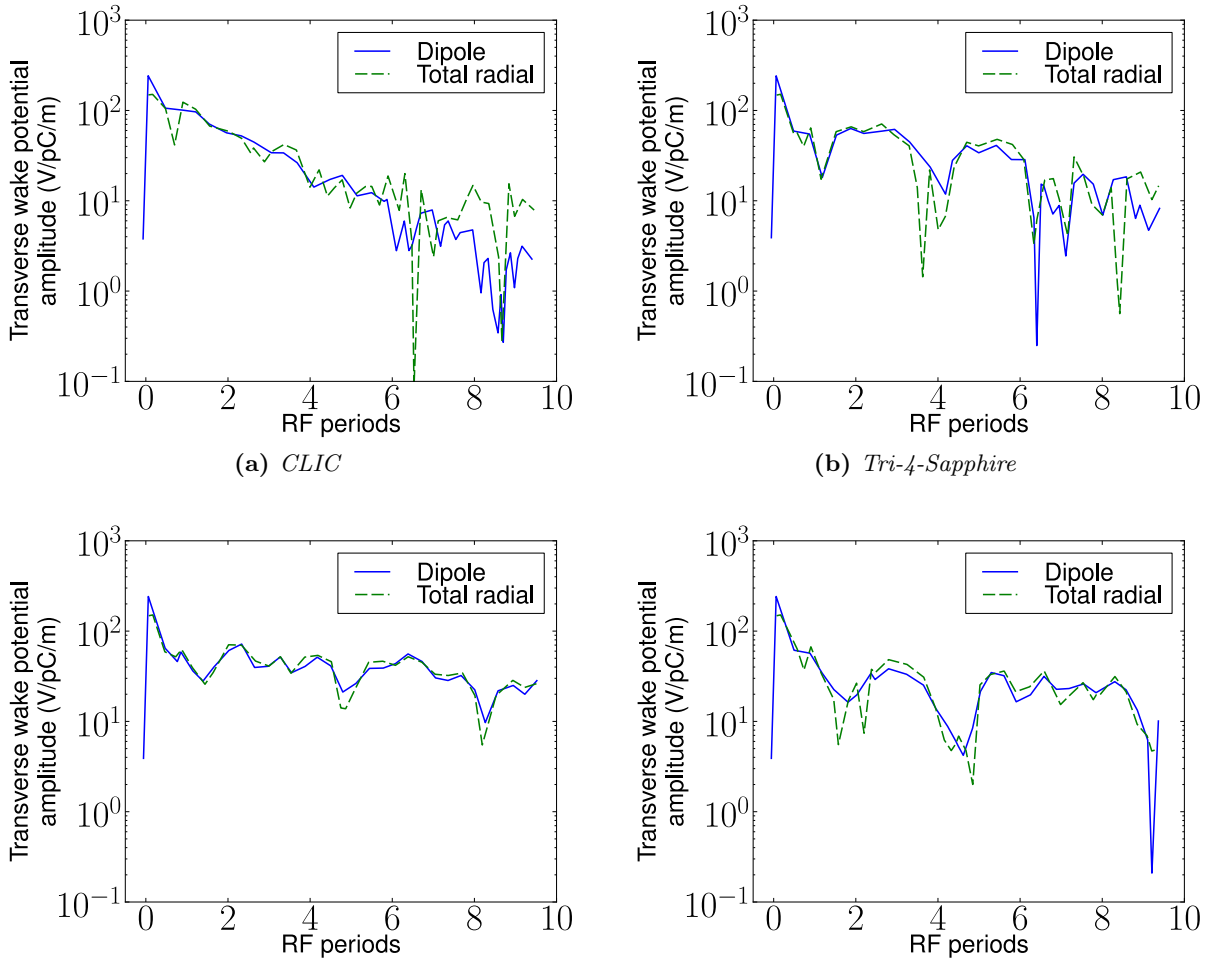

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.16: *Total radial vs. transverse dipole wake potentials in 8-cell extended-domain cavities showing the dominance of the dipole wakes in the first 10 RF periods behind the bunch.*

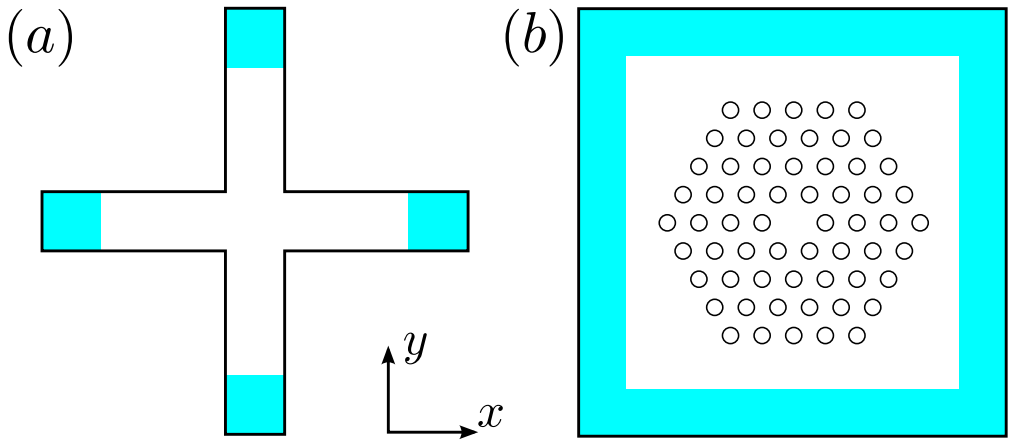

Figure 4.17: *Conducting absorber locations in (a) the CLIC cavity and (b) the HDPhC cavity simulations. Absorbers are shaded. There are 40mm between the CLIC cavity center and the inside absorber edges and at least 4 rod radii between the outermost rods and the absorber edges.*

properties of absorber.

Figure 4.18 shows the difference between the dipole transverse wake potentials in the extended domain cavities and the cavities equipped with conducting absorbers. The difference is quite small in all HDPhC cavities, whereas significant reflection is observed in the CLIC cavity. This reflection is reduced as the absorber depth W is increased, as expected, since the impedance mismatch along the absorber is decreased (assuming unchanging profile and $\sigma_{\max}$). The difference is likely due to the following.

The 1mm bunch strongly excites cavity TM modes that are uniform in the z-direction. In the CLIC structure, cavity TM modes are damped by coupling to rectangular waveguide TE modes. All rectangular waveguide modes have a cutoff frequency; near cutoff, these modes are reflected strongly by small changes in waveguide impedance (due to the flattening of the dispersion), and are therefore susceptible to large reflections off of the absorbers. Alternatively, in the HDPhC cavities, TM modes that are uniform in z have no cutoff frequency (the dispersion is plane-wave-like); thus, they penetrate farther into the absorbers and undergo weaker reflections (shorter bunches will excite modes *with* z-variation more strongly—these modes do have cutoff frequencies in the HDPhC cavities and would be more sensitive to absorber properties). Additionally, reflections off

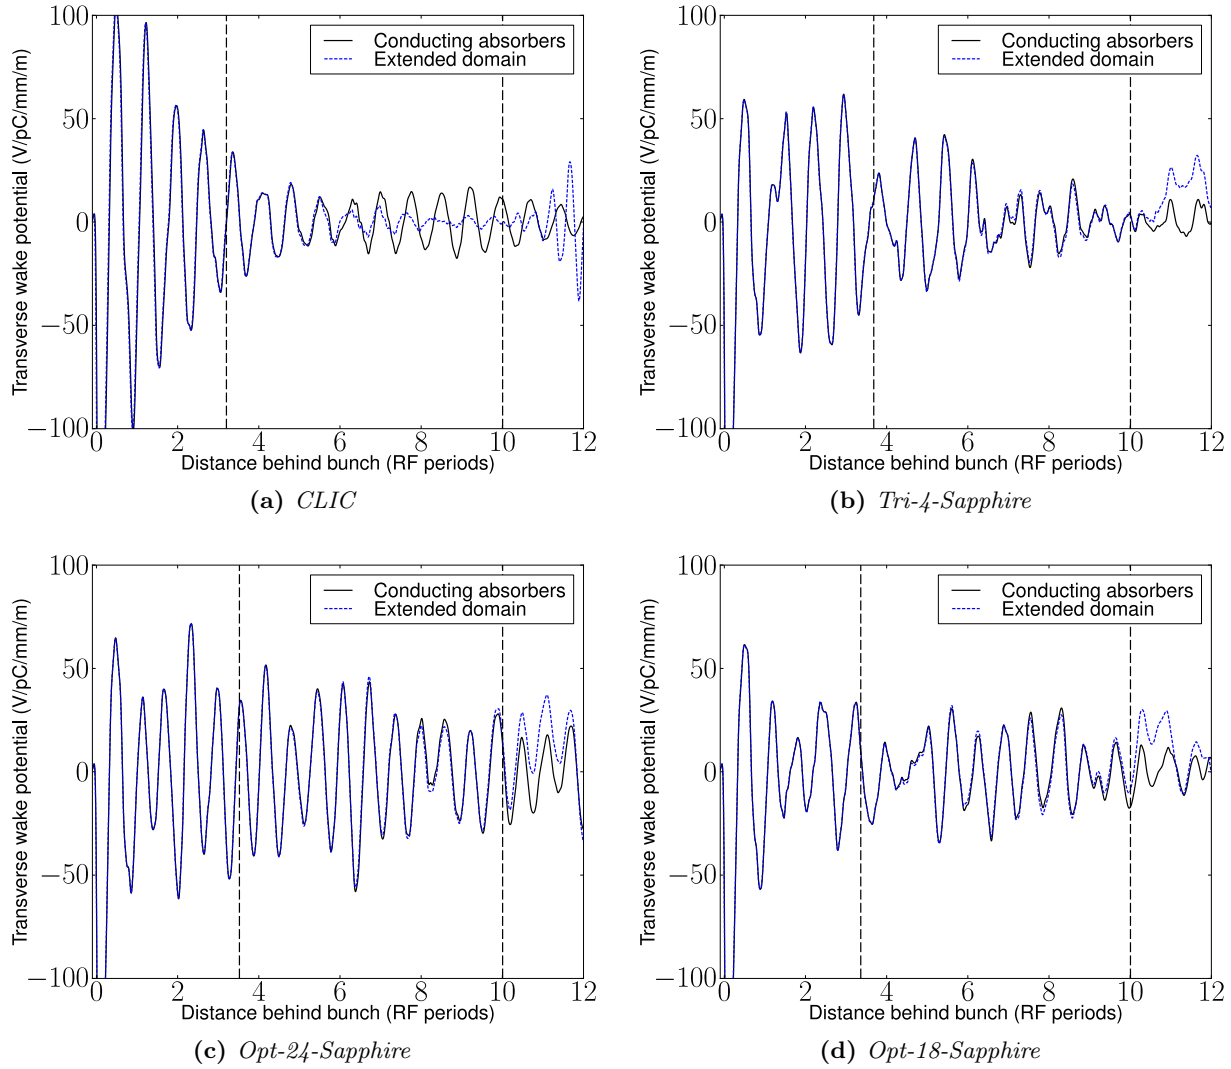

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.18: *Reflections off of the conducting absorbers cause differences in the wake potentials between the two vertical lines. Beyond the second line, reflections from the extended domain boundaries reach the beam axis. The absorbers work well in the HDPhC cavities, but produce significant reflection in the CLIC cavity. Reflections can be reduced in the CLIC cavity by increasing the depth of the conductors while retaining the original conductivity profile (thus reducing the impedance mismatch of the absorber).*

of the absorbers in the HDPhC cavities have the opportunity to encounter other absorbers without passing through the beam axis; in the CLIC structure, all absorber reflections are funneled back into the cavity.

Based on Fig. 4.18, reducing the impedance mismatch of the absorbers in the CLIC cavity will have a significant (beneficial) effect on the mid-range transverse wake, whereas such measures applied to the HDPhC cavities will have a relatively negligible effect.

4.3.6 Conducting absorbers wakefield results

Longitudinal monopole wake potentials are shown in Fig. 4.19 with the corresponding impedances shown in Fig. 4.20. These longer wakefield simulations reveal advantages for the HDPhC cavities. Because the accelerating mode in 8-cell cavities has an eight-fold near-degeneracy, we could not subtract its contribution from the monopole wake with enough accuracy. The "cleanliness" of the wake potentials in Fig. 4.19 can be qualitatively assumed from the smoothness of the envelopes. Alternatively, the impedances in Fig. 4.20 can be used to make a more quantitative assessment. The suppressed peaks above 12 GHz indicate stronger monopole HOM damping in HDPhC cavities, although the monopole mode density is significantly higher.

As the extended-domain results showed, the CLIC cavity (as it was designed to do) rapidly suppresses the transverse wake before the arrival of the first trailing bunch, and does so more effectively than the HDPhC cavities (where interference effects do not play a major role). This is because the lowest dipole mode is dominant in this region of the transverse wake, and the CLIC cavity was specifically designed to suppress it. As Fig. 4.21 shows, the CLIC cavity continues to damp effectively the dipole wakefields out to large trailing distances. Of the HDPhC structures, the Opt-18-Sapphire cavity damps dipole wakes the best, but (for the majority of the calculated wake region) still falls short of the CLIC cavity (recall that the CLIC damping can be improved further by reducing reflections off of the absorbers, whereas HDPhC damping cannot). The dipole wake impedances shown in Fig. 4.22 reveal the troublesome parts of the spectrum in the HDPhC cavities. The Tri-4-Sapphire cavity has the largest transverse wake which is due to a cluster of

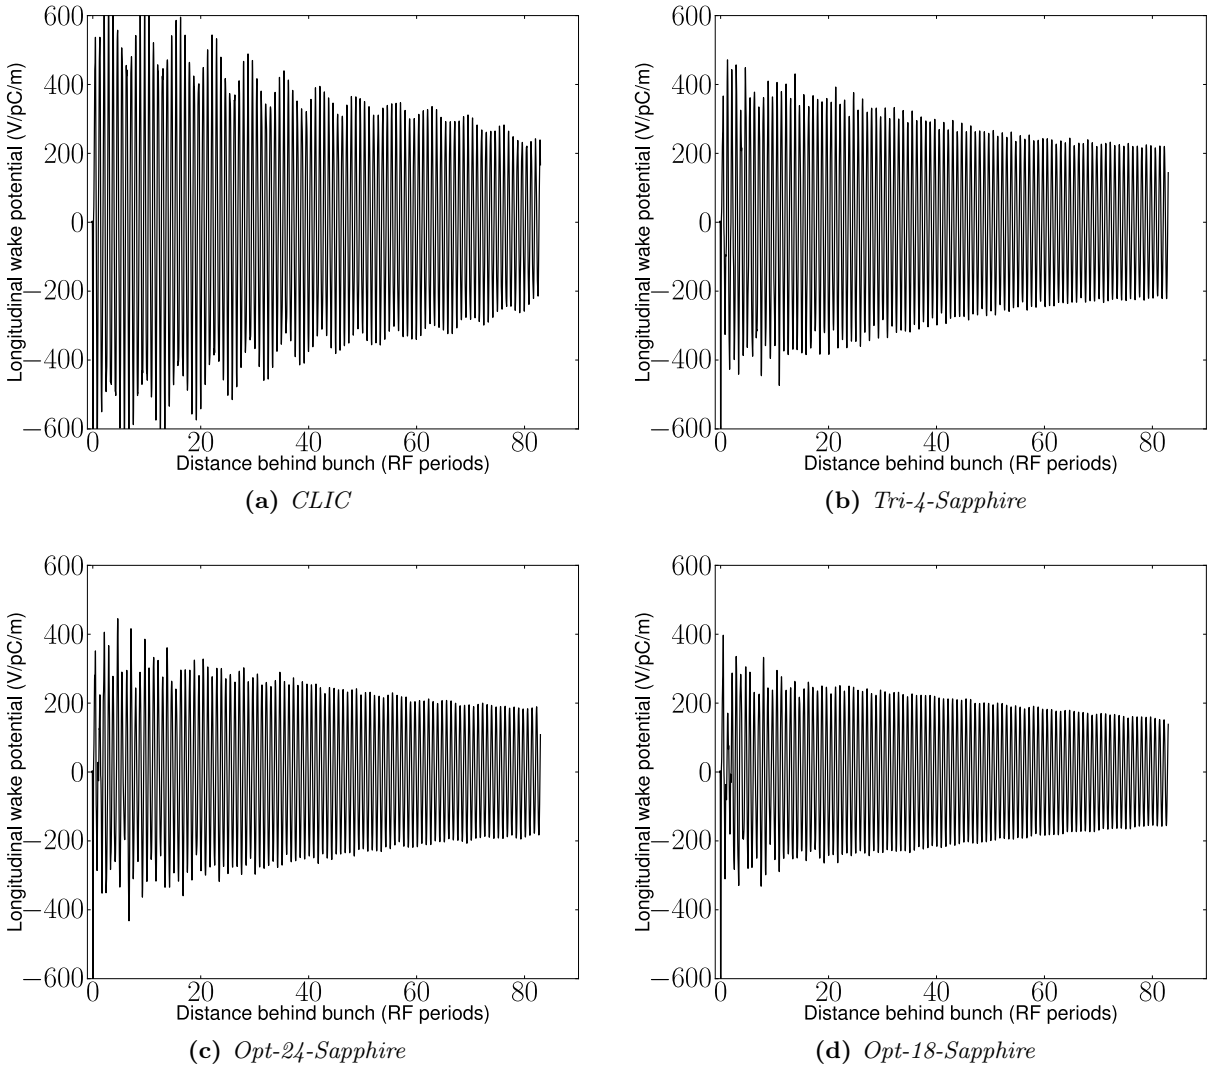

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.19: *Longitudinal monopole wake potentials in 8-cell cavities using conducting absorbers.*

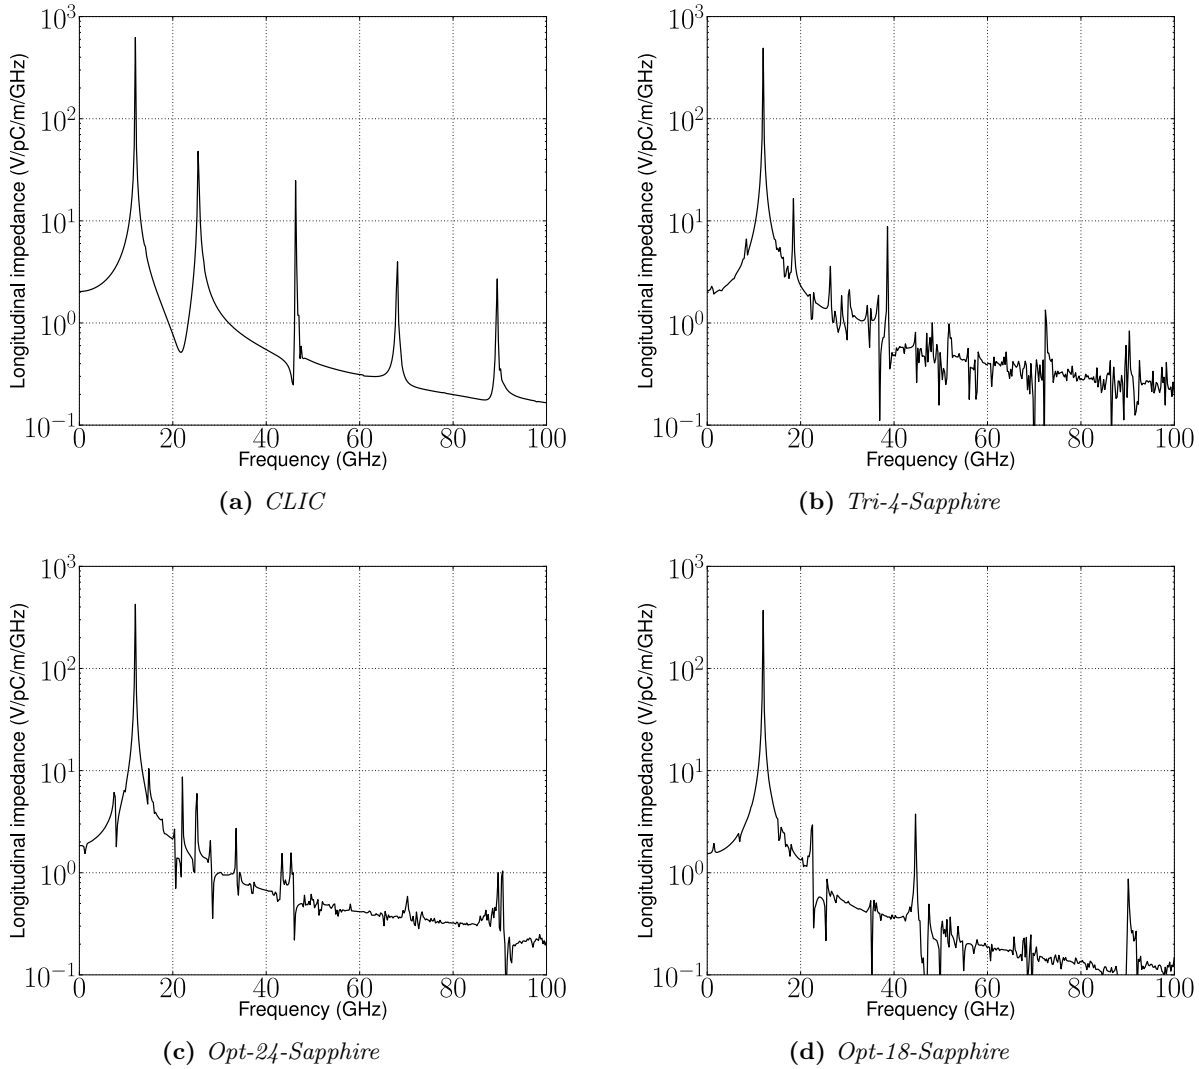

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.20: *Longitudinal monopole impedance in 8-cell cavities using conducting absorbers.*

high-Q dipole modes around 15 GHz.

4.4 Discussion: PhC HOMs

The transverse wake potential in the Tri-4-Sapphire cavity is surprisingly higher than in the CLIC cavity (we were also surprised to see relatively higher transverse wakes in the optimized structures, but are somewhat reassured by the longitudinal wakefield results). The transverse impedance shows that this is due to a cluster of high-Q modes near 15 GHz. In this section, we explain the presence of high-Q HOMs in the Tri-4-Sapphire cavity using the properties of the triangular-lattice band structure outside of any bandgaps—namely, that flat regions of the dispersion diagram imply low-group-velocity modes which sluggishly carry their energy through the lattice and thus can be considered "trapped." This issue was addressed briefly in Ref. [50] within a study of a square-lattice-based metal-rod accelerating structure.

We have seen in Section 2.4.5 on HOM mitigation strategies how waveguide damping is effective for coupling out frequencies only above cutoff (the further above cutoff the better). This is because waveguide modes with frequencies near cutoff have vanishing group velocities (because of the flattened dispersion at cutoff) and thus propagate slowly down the waveguide—effectively, they are trapped. In PhCs, the dispersion tends to flatten where the spatial variation of the modes matches the spatial variation of the dielectric. This flattening, while beneficial for the formation of bandgaps (and thus trapped accelerating modes), introduces PhC modes with vanishing group velocities.

Figure 4.23 shows the transverse dipole impedance of the Tri-4-Sapphire cavity in the frequency range 0-25 GHz (a zoom view of Fig. 4.22b) and matches it with the 2D TM band diagram for the triangular lattice of sapphire discs. The two most prominent peaks in the impedance clearly line up with the flat portions of the second band (the third peak can be matched with another flat band in the fully 3D band diagram). The annotated impedance peak (second largest) was investigated further using the time-domain mode extraction technique. Figure 4.24 shows the field pattern of the extracted mode (from a periodic single-cell simulation at phase advance $\phi \approx 3\pi/4$)

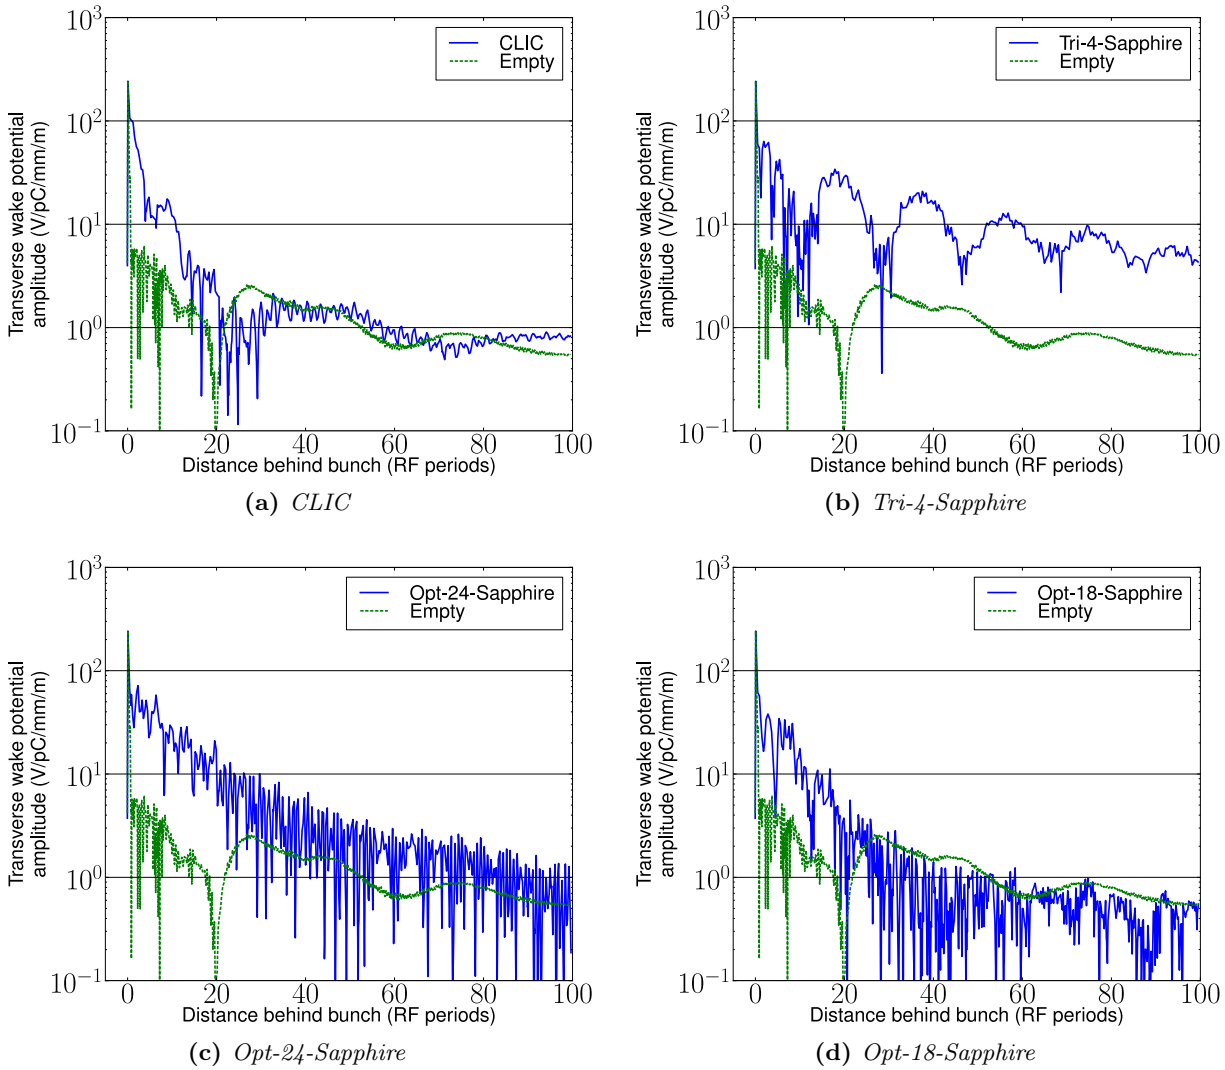

(a) *CLIC* **(b)** *Tri-4-Sapphire*

(c) *Opt-24-Sapphire* **(d)** *Opt-18-Sapphire*

Figure 4.21: *Transverse dipole wake potentials in 8-cell cavities using conducting absorbers.*

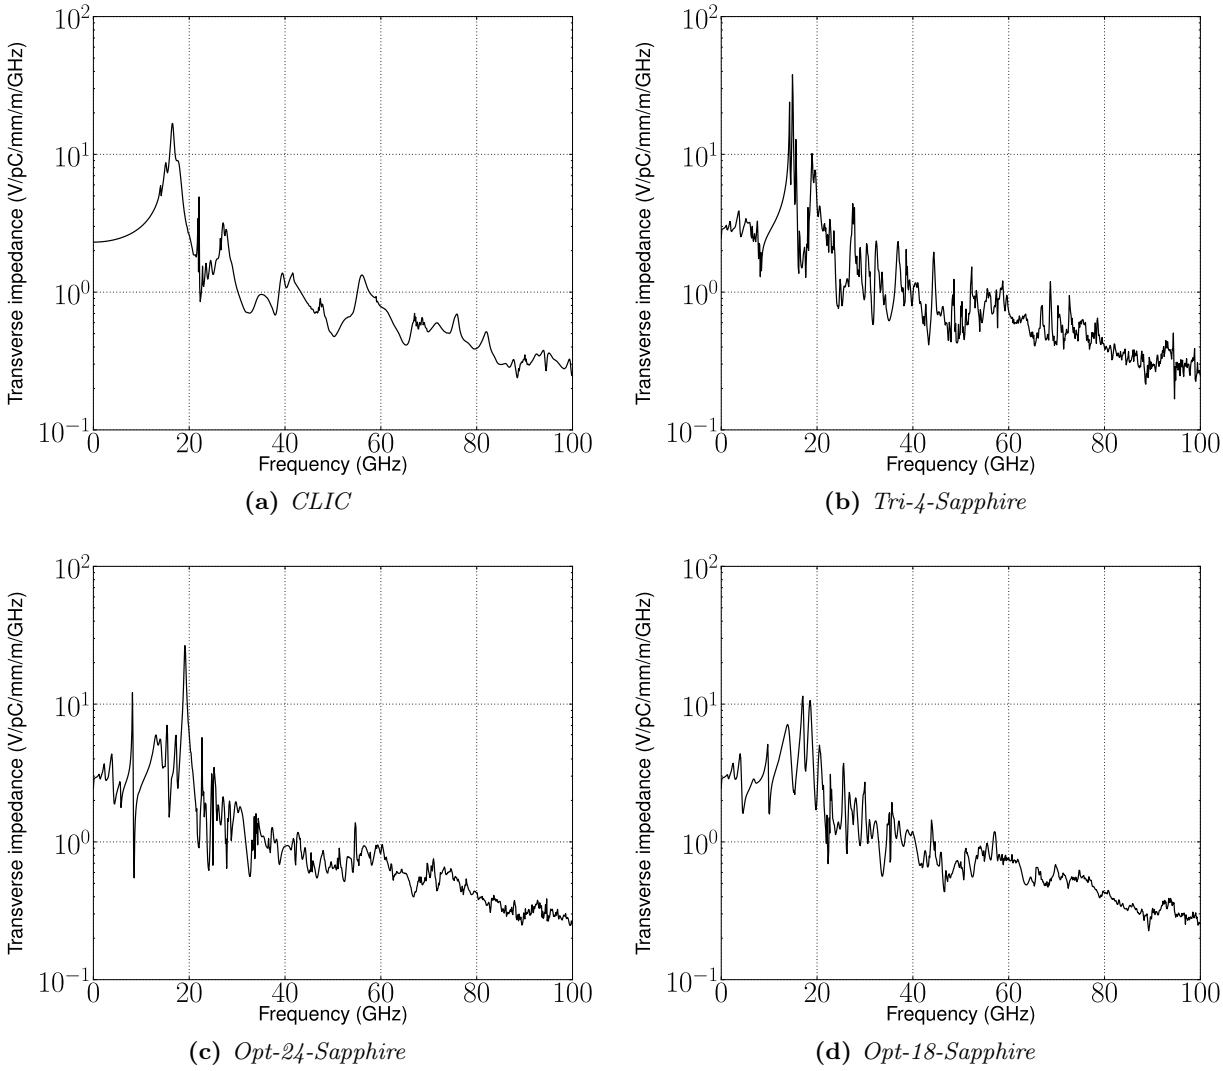

(a) *CLIC*

(b) *Tri-4-Sapphire*

(c) *Opt-24-Sapphire*

(d) *Opt-18-Sapphire*

Figure 4.22: *Transverse dipole impedance in 8-cell cavities using conducting absorbers.*

and compares it with the 2D PhC lattice mode at the M-point of the second band. The field patterns clearly match.

Both the damping waveguides and the triangular PhC suffer from low-group-velocity modes. Why then is the damping worse in the Tri-4-Sapphire cavity? The answer lies in the combined characteristics of the damping structure and the central cavity region. Consider the lowest dipole mode in the CLIC cavity. Section 2.4.5 showed that waveguide damping is more effective when the undamped resonant frequency is further above the cutoff frequency. The central cavity dimensions in the CLIC cavity determine a dipole resonant frequency; thus, the waveguide dimensions are selected such that cutoff is as far below that frequency as possible (without affecting the accelerating mode too much). The gap between the accelerating frequency and the dipole frequency is large enough such that this scheme results in effective dipole damping.

In contrast, the defect region of the Tri-4-Sapphire cavity is (by definition) highly commensurate with the geometry of the surrounding lattice. Also, the mode patterns shown in Fig. 4.24 have field nodes that pass through the dielectric, indicating that the removal of the central rod has a small effect on the mode (as compared to modes at the top of the first band, which have most of their field energy concentrated in dielectric). Thus, in the parlance of waveguide damping, the defect resonance is at the same frequency as the low-group-velocity lattice mode (or the "cutoff" mode), thus damping is ineffective. Put another way: the creation of the defect weakly affects the lattice mode; thus, the lattice mode retains a strong presence in the defect. This argument suggests pushing the inner layer of rods closer to the beam axis, which is likely to increase the dipole resonance frequency in the defect above the flat portion of the lattice band and thus increase coupling to higher-group-velocity lattice modes.

The situation may be exacerbated by the impedance mismatch at the outer layer of rods. Because of its low-group-velocity, the lattice mode shown in Fig. 4.24 may be highly susceptible to reflections off of the transition between lattice and vacuum, effectively increasing its Q_{rad}. This transition could be made smoother by slowly decreasing the radii of the rods in outer layers, but would add significantly to the transverse size of the structure.

Unfortunately, this analysis loses meaning in the optimized structures. A frequency-targeted mode search would characterize the less-than-ideal transverse wakefields in these cavities. It is worth recalling, however, that the accelerating mode Q_{rad} in the optimized structures is highly sensitive to rod displacements, indicating that a perturbative fix is an unlikely solution. Rather, transverse wakefield minimization should be an integral component of future PhC-based cavity optimizations.

4.5 Conclusion

Low wakefields require that the energy in HOMs be dissipated as quickly as possible. Thus, a damping mechanism (or structure) should have a mode spectrum without any low group velocities. In principle, the CLIC cavity suffers from low-v_g HOMs near the cutoff frequency of its damping waveguides; however, the sparse mode density of the conducting cavity allowed the placement of the cutoff frequency within an empty region of the spectrum, producing effective damping. We have found that lattice-based HDPhC cavities also suffer from low-v_g HOMs due to the flattening of bands near k-points of strong lattice symmetry. For the Tri-X-Sapphire design, the defect-based central cavity supported low-v_g dipole modes and thus showed large transverse wakes compared to the CLIC structure. Longitudinal wakes were lower in all HDPhC cavities when compared to CLIC.

Future work should focus on reducing wakefields in lattice-based HDPhC cavities. Because the lattice structure has the potential to reduce surface magnetic fields and increase accelerating efficiencies and PhC band theory provides an explanation for the existence of HOMs, it should take precedence over the optimized structures presented in this thesis in future designs. Possible routes to reducing wakes in Tri-X-Sapphire cavities include perturbing the central defect region to eliminate low-v_g resonances and/or reducing the impedance mismatch of all HOMs at the truncation of the lattice.

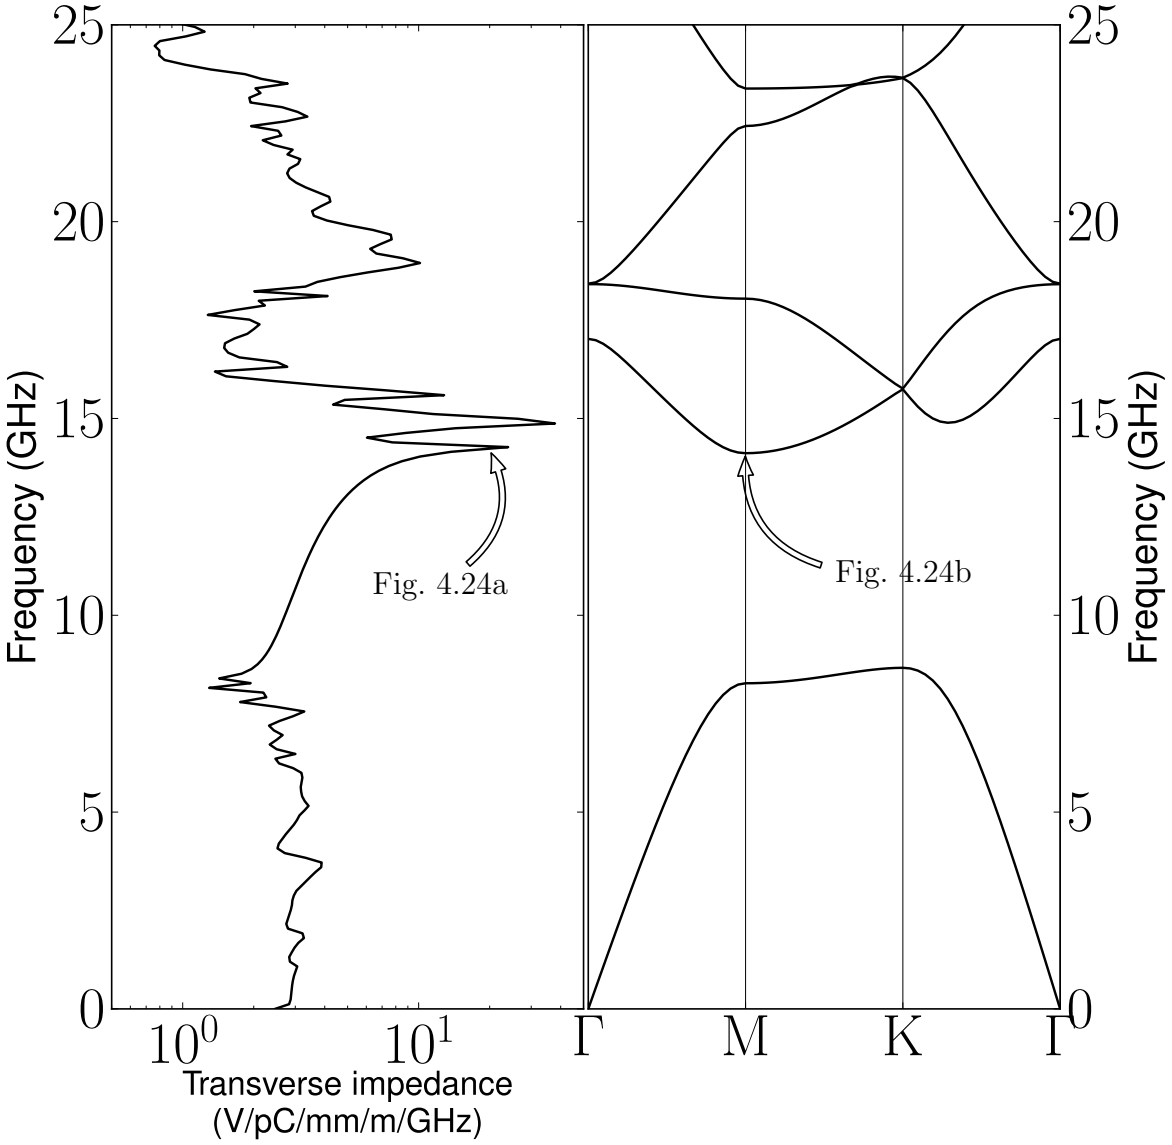

Figure 4.23: *The left plot shows a zoomed-in view of the troublesome part of the transverse dipole impedance (from Fig. 4.22b). The right plot is the 2D TM band diagram for the triangular lattice of sapphire discs. The annotated peak in the impedance is at the same frequency as the M-point of the second band (also see the mode patterns in Fig. 4.24). This correlation supports the idea that low-v_g modes pose a problem for wakefields.*

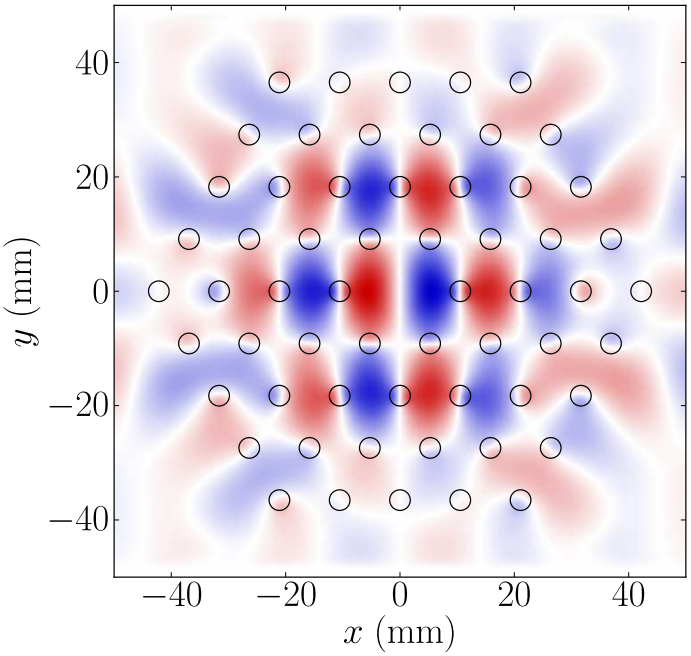

(a) *Tri-4-Sapphire mode*

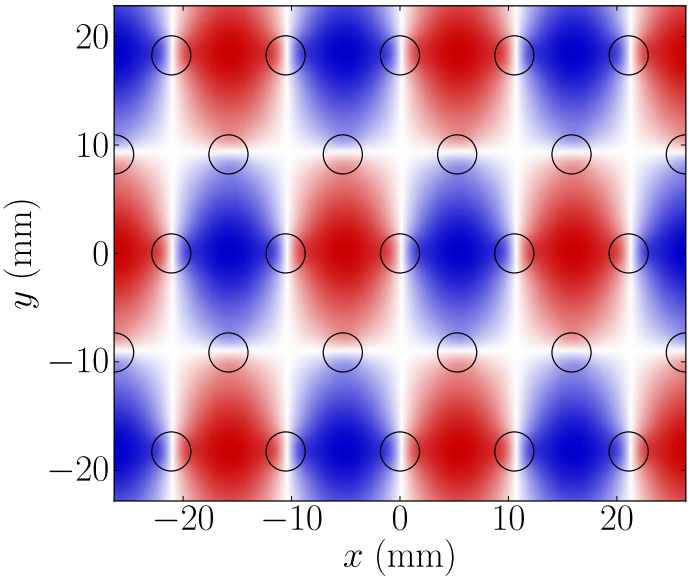

(b) *Infinite PhC mode*

Figure 4.24: *Analagous modes in an infinite 2D PhC and the Tri-4-Sapphire periodic single-cell cavity (the mode pattern in (b) is uniform in z). The PhC mode in (a) has a vanishing group velocity; thus, its counterpart in (b) is "trapped" and contributes significantly to the dipole wakefield. Computations used (a)* VORPAL *[60] and (b) MIT Photonic Bands [39].*

Chapter 5

A second-order 3D electromagnetics algorithm for curved interfaces between anisotropic dielectrics on a Yee mesh

A new frequency-domain electromagnetics algorithm is developed for simulating curved interfaces between anisotropic dielectrics embedded in a Yee mesh with second-order error in resonant frequencies. The algorithm is systematically derived using the finite integration formulation of Maxwell's equations on the Yee mesh. Second-order convergence of the error in resonant frequencies is achieved by guaranteeing first-order error on dielectric boundaries and second-order error in bulk (possibly anisotropic) regions. Convergence studies, conducted for an analytically solvable problem and for a photonic crystal of ellipsoids with anisotropic dielectric constant, both show second-order convergence of frequency error; the convergence is sufficiently smooth that Richardson extrapolation yields roughly third-order convergence. The convergence of electric fields near the dielectric interface for the analytic problem is also presented.

5.1 Introduction

The manipulation of light with dielectric structures (e.g. photonic crystal cavities and waveguides, dielectric accelerating cavities, fiber-optics, etc.) is a rapidly growing field of research requiring the accurate simulation of electromagnetic fields in the presence of dielectric boundaries. In finite-difference simulations, the Yee algorithm is often used because of its accuracy in homogeneous media (in this case, resonant frequencies converge with error that is second-order in the size of a grid cell Δx), its good dispersion characteristics, and the simplicity of its implementation [87, 30].

However, dielectric boundaries introduce discontinuities in the electric field and electric displacement (i.e. $\mathbf{D} = \varepsilon(\mathbf{x})\mathbf{E}$ where $\varepsilon(\mathbf{x})$ contains discontinuities) and significantly reduce the accuracy of the standard Yee algorithm (resonant frequencies tend to converge with only first-order error). To retain accuracy, the Yee scheme must be modified at dielectric boundaries.

An alternative to finite difference algorithms is the finite element approach, which can treat material boundaries using arbitrary meshes that conform to surfaces. Because of the many fundamental differences between finite element and finite difference algorithms, a comparison is difficult to draw. However, one compelling reason to continue improving finite difference algorithms is the advantage of short setup times; mesh generation in finite element codes can be a time-consuming and resource-intensive process, whereas the finite difference grid is simply implied at no computational cost.

Numerous papers have been written on second- and higher-order finite-difference dielectric boundary algorithms. The problem has proven to be an extremely difficult task of balancing accuracy, time-domain stability, generality, and ease of implementation. Many recent works have made significant advances in some of these areas, but most tend to focus on just one or two at a time. The following paragraphs offer a cursory overview.

Finite difference dielectric boundary algorithms tend to fall into two categories: those that use nonstandard approximations of the differential operators of Maxwell's equations at boundaries, and those that use standard finite differences but modify the constitutive material relations. The former often use one-sided finite differences or extrapolations, where all grid field values used to calculate a derivative are located within the same material. These types of algorithms could in principle achieve very high orders of accuracy while preserving C^n discontinuities at material interfaces, but they have been demonstrated only for very simple cases, such as 1D interface simulations, or planar geometries in 2D where the material interfaces are aligned with grid directions [90, 85, 88]. Second-order convergence has been achieved for 2D curved dielectric interfaces, but only for TM polarizations where all the fields are C^0 continuous across the interface [18, 8]. The applicability of these methods to general curved 3D geometries remains unknown.

The second category of algorithm approximates the constitutive material relations while using standard finite difference stencils to approximate derivatives. These methods often proceed by approximating the materials themselves. For example, "stair-stepping" (which generally gives first-order convergence) is the approximation of a curved object as a set of cubes (3D) or squares (2D) [35].

Effective dielectric methods, which use cell-averaged dielectric constants, can compensate for the poor accuracy of stair-stepping. In these methods, cells that are cut by the interface are replaced by cells filled with a new uniform dielectric. The new dielectric constant is calculated by averaging the true dielectric constants over the cut cell. Several algorithms using heuristic cell-averaging techniques show improved accuracy compared to stair-stepping, but do not achieve or show conclusively second-order convergence when simulating complex curved geometries [41, 42, 16, 49, 89, 36, 57, 70]. Some other schemes do show second-order or higher convergence, but only for very special-case geometries, such as grid-parallel interfaces in 1D or 2D, or for only TM fields, where the electric field is continuous across dielectric boundaries [20, 33]. A very recent paper has shown second-order convergence for 2D TE fields in the presence of a straight oblique dielectric interface [34]; unfortunately, the algorithm only applies to two specific interface angles with respect to the grid. A major advantage of these methods is their time-domain stability; because these algorithms are a kind of stair-stepping, they naturally preserve the symmetries that guarantee a stable algorithm [79].

Unfortunately, the ultimate order of convergence of effective dielectric algorithms is hard to determine theoretically. Two sources of error must be addressed: the error due to the approximation of the surface by the effective dielectric, and the error due to the discretization of the effective dielectric. A more rigorous treatment of the former is found in Ref. [45] where first-order perturbation theory is used to calculate the resonant frequency shift induced by a "smoothing" of the dielectric interface (i.e. making the transition from one dielectric constant to another continuous). With an appropriate smoothing, the frequency shift will be at most second-order in the volume of the smoothing region.

The algorithm of Ref. [63] combines the smoothed interface algorithm of Ref. [45] with the finite difference discretization scheme of Ref. [79]; however the errors introduced by the discretization are not discussed theoretically. A smoothly varying dielectric can be simulated with second-order error as long as the dielectric variation (hence field variation) across a single cell vanishes in the limit $\Delta x \to 0$ where Δx is the grid cell size [79]. For a given Δx, the algorithm of Ref. [63] creates a smoothed interface (a second-order approximation to a sharp interface) of approximate thickness Δx; thus, the dielectric variation across a grid cell at a boundary remains constant in the limit $\Delta x \to 0$. This would seem to incur a resolution-independent error for boundary cells, which would show up as first-order error in resonant frequencies [29]. Nevertheless, the results of Ref. [63] show second-order convergence for the cases presented. However, because their algorithm does not adhere to the strict criterion (put forth by this paper) of exactness of the surface electric fields in the infinite wavelength limit, we believe that ultimately the algorithm of Ref. [63] has first-order error (possibly obscured up to high resolutions or high dielectric constrasts).

The finite integration interpretation of the Yee scheme (which is the basis for our algorithm) also reduces the dielectric boundary problem to a modification of only the constitutive relations [78, 10]. This approach is exemplified by Ref. [15] which achieves second-order convergence for metallic boundaries embedded in the Yee mesh. For dielectric boundaries, the algorithm of Ref. [58] follows this approach, but (as we will discuss in Section 5.3.1) does not guarantee second-order convergence. Furthermore, it is unclear how to extract an order of convergence from the numerical results.

Our algorithm extends Ref. [58], focusing exclusively on accuracy for 3D curved interfaces between anisotropic dielectric materials (at the expense of time-domain stability); the algorithm shows unambiguous second-order convergence of resonant frequencies in the grid cell size. Underlying our derivation is the following principle: in the infinite wavelength limit (the limit of spatially constant fields), electromagnetic fields should be exactly simulated at planar dielectric boundaries. This basic principle, when consistently addressed in the finite integration formalism, it systematically leads to an algorithm with at most first-order error in the fields on a dielectric interface.

To achieve second-order convergence of resonant frequencies, we guarantee that in homogeneous regions, the algorithm has local second-order error [29].

We begin in the next section by reviewing the use of finite integrals to describe the Yee electromagnetic scheme and defining most of the notation convenient for our dielectric algorithm. Section 5.3.1 discusses a prescription for determining/ensuring accuracy for the discrete permittivity operator. Section 5.3.2 develops the algorithm that guarantees first-order error on dielectric boundaries and in anisotropic bulk dielectric regions. Section 5.3.3 describes a modification to the first-order algorithm to achieve second-order error in bulk anisotropic regions. Our numerical examples are presented in Sec. 5.4. We conclude the paper with an Appendix, which provides a severely condensed description of our algorithm in the form of a list of steps necessary for its implementation.

5.2 Background

Our algorithm is based on the finite integration technique (FIT), which, for Cartesian meshes in homogeneous isotropic materials, reduces to the standard Yee scheme [87, 78, 10]. Maxwell's equations in integral form are

$$\frac{\partial}{\partial t}\int_A \mathbf{D}(\mathbf{x})\cdot d\mathbf{A} = \oint_C \mathbf{H}(\mathbf{x})\cdot d\mathbf{l} \tag{5.1}$$

$$\frac{\partial}{\partial t}\int_A \mathbf{B}(\mathbf{x})\cdot d\mathbf{A} = -\oint_C \mathbf{E}(\mathbf{x})\cdot d\mathbf{l} \tag{5.2}$$

where $\mathbf{D}(\mathbf{x}) = \varepsilon(\mathbf{x})\cdot\mathbf{E}(\mathbf{x})$ and $\mathbf{B}(\mathbf{x}) = \mu(\mathbf{x})\cdot\mathbf{H}(\mathbf{x})$ and $\varepsilon(\mathbf{x})$ and $\mu(\mathbf{x})$ are in general 3×3 symmetric tensors. The FIT associates finite areas with the Cartesian components of $\mathbf{D}$ and $\mathbf{B}$, and finite lengths with the Cartesian components of $\mathbf{E}$ and $\mathbf{H}$ (see Fig. 5.1). To define these finite regions, space is discretized into two overlapping Cartesian grids: the regular grid and the dual grid. In this paper, regular grid nodes will be defined as the points, $\mathcal{N}_{ijk} \equiv [i\Delta x, j\Delta y, k\Delta z]$ and dual grid nodes will be defined as the points, $\tilde{\mathcal{N}}_{ijk} \equiv [(i+0.5)\Delta x, (j+0.5)\Delta y, (k+0.5)\Delta z]$, where i, j, k are integers, $\Delta x, \Delta y, \Delta z$ are the dimensions of a grid cell, and the tilde indicates association with the dual grid.

Finite lengths will be labeled as $L_{\alpha|ijk}$ and $\tilde{L}_{\alpha|ijk}$ where α is one of x, y, or z. The regular line segment, $L_{\alpha|ijk}$, connects the regular grid node, $\mathcal{N}_{ijk}$ to the next regular grid node in the positive α-direction. For example, $L_{x|ijk}$ joins $\mathcal{N}_{ijk}$ and $\mathcal{N}_{i+1,jk}$. The dual line segment, $\tilde{L}_{\alpha|ijk}$, connects the dual grid node, $\tilde{\mathcal{N}}_{ijk}$ to the next dual grid node in the *negative* α-direction. For example, $\tilde{L}_{x|ijk}$ joins $\tilde{\mathcal{N}}_{ijk}$ and $\tilde{\mathcal{N}}_{i-1,jk}$. Some examples of these lengths are shown in Fig. 5.1.

Finite areas will be labeled by $A_{\alpha|ijk}$ and $\tilde{A}_{\alpha|ijk}$ where α now corresponds to the direction normal to the area. These are best described by examples, a few of which are shown in Fig. 5.1. Regular areas are defined by regular grid nodes. For example, $A_{x|ijk}$ is the rectangular area with corners at $\mathcal{N}_{ijk}$, $\mathcal{N}_{ij+1,k}$, $\mathcal{N}_{ijk+1}$, and $\mathcal{N}_{ij+1,k+1}$. Dual areas are similarly described by dual grid nodes, but in the negative sense, analogous to the dual length segments described above. For example, $\tilde{A}_{x|ijk}$ is the rectangular area with corners at $\tilde{\mathcal{N}}_{ijk}$, $\tilde{\mathcal{N}}_{ij-1,k}$, $\tilde{\mathcal{N}}_{ijk-1}$, and $\tilde{\mathcal{N}}_{ij-1,k-1}$. The finite areas are positioned such that the regular (dual) lengths $L_{\alpha|ijk}$ ($\tilde{L}_{\alpha|ijk}$) pierce the dual (regular) areas $\tilde{A}_{\alpha|ijk}$ ($A_{\alpha|ijk}$) at their midpoints.

The FIT then defines the following electromagnetic field quantities to be represented on the Yee mesh:

$$e_{x|ijk} \equiv \frac{1}{\Delta x} \int_{L_{x|ijk}} E_x(\mathbf{x})\, dl \tag{5.3}$$

$$d_{x|ijk} \equiv \frac{1}{\Delta y \Delta z} \int_{\tilde{A}_{x|ijk}} D_x(\mathbf{x})\, da \tag{5.4}$$

$$h_{x|ijk} \equiv \frac{1}{\Delta x} \int_{\tilde{L}_{x|ijk}} H_x(\mathbf{x})\, dl \tag{5.5}$$

$$b_{x|ijk} \equiv \frac{1}{\Delta y \Delta z} \int_{A_{x|ijk}} B_x(\mathbf{x})\, da \tag{5.6}$$

where the integrands are the unknown analytic electromagnetic fields (the corresponding definitions for the y and z components can be inferred from the above). This defines the association between grid field quantities and finite regions. The FIT has the convenient property that Maxwell's dynamical equations (Eqs. (5.1) and (5.2)) are represented exactly by the discretization [10]. As a result, we can isolate the dominant source of error in dielectric boundary problems to the discrete dielectric constitutive relation.

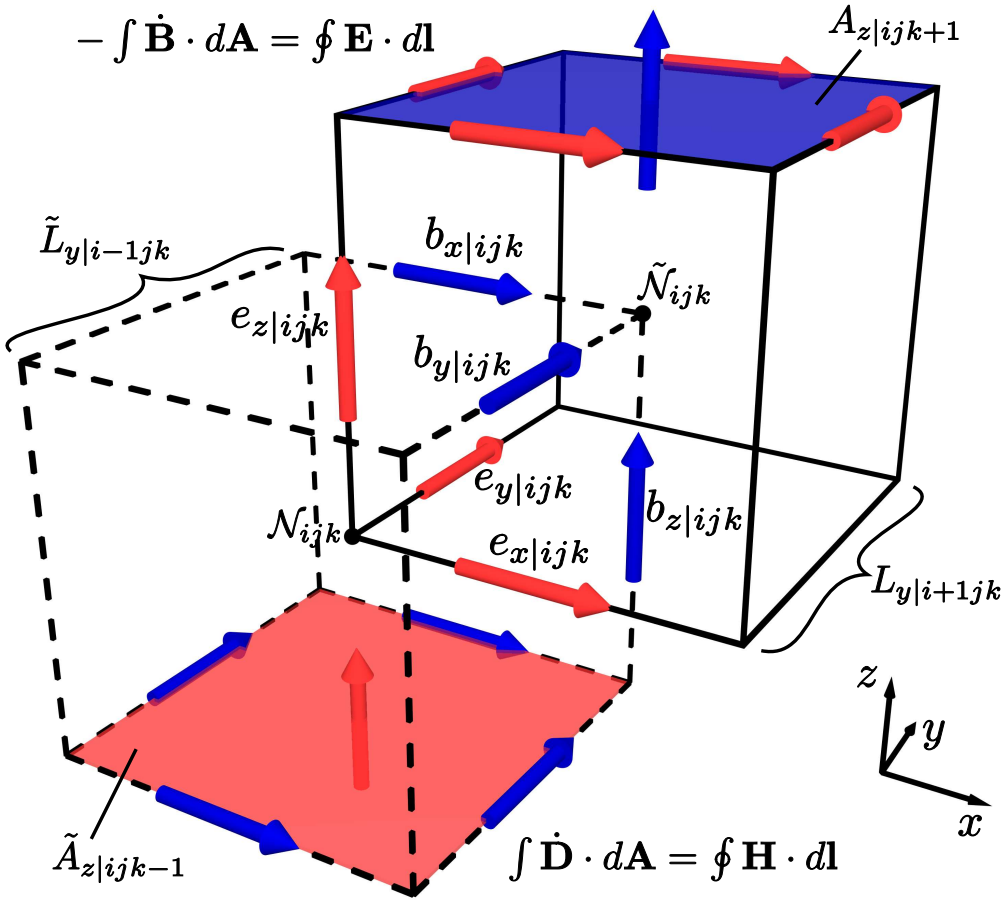

Figure 5.1: *Color online. The solid-line cube is the regular grid cell ijk and the dotted-line cube is the dual grid cell ijk. Field components are distributed by the usual Yee method. The figure shows how* **E** *(***D** *implied) field components (red) are associated with regular cell edges (dual cell faces) and how* **B** *(***H** *implied) components (blue) are associated with regular cell faces (dual cell edges). Examples of the finite integral Ampere and Faraday Laws are highlighted.*

Letting $\mathbf{e}$, $\mathbf{d}$, $\mathbf{h}$, and $\mathbf{b}$ be vectors of all $e_{\alpha|ijk}, d_{\alpha|ijk}, h_{\alpha|ijk}$, and $b_{\alpha|ijk}$ components on the Yee mesh, respectively, we can write the temporally-continuous, spatially-discretized Maxwell equations in the following matrix-vector form,

$$\frac{\partial}{\partial t}\mathbf{d} = \mathbb{C}\mathbf{h} \tag{5.7}$$

$$-\frac{\partial}{\partial t}\mathbf{b} = \mathbb{C}^{\mathrm{T}}\mathbf{e} \tag{5.8}$$

$$\mathbf{d} = \mathbb{M}_{\varepsilon}\mathbf{e} \tag{5.9}$$

$$\mathbf{b} = \mathbb{M}_{\mu}\mathbf{h} \tag{5.10}$$

where $\mathbb{C}$ is the usual Yee backward finite difference curl operator and $\mathbb{M}_{\varepsilon}$ and $\mathbb{M}_{\mu}$ are the discrete permittivity and permeability operators, respectively. The layout of the Yee mesh ensures that the curl operator of the electric field is the transpose of $\mathbb{C}$ (for a more detailed explanation see Ref. [79]).

In the next section, we show how to form an inverse permittivity matrix, $\mathbb{M}_{\varepsilon}^{-1}$, that has local first-order error on dielectric boundaries and second-order error everywhere else, such that resonant frequencies converge to the correct values with second-order error. Throughout the rest of the paper, the magnetic permeability will be set to the constant permeability of free space ($\mathbb{M}_{\mu} = \mu_0\mathbb{I}$). However, our algorithm applies equally well to simulating interfaces between anisotropic permeabilities by the symmetry of Maxwell's equations and the finite integration technique (one simply replaces $\mathbf{E}$, $\mathbf{D}$, and ε by $\mathbf{H}$, $\mathbf{B}$, and μ in all that follows).

5.3 Method

Our technique for simulating curved dielectric boundaries is described in two steps. First, we develop a dielectric algorithm that has local first-order error everywhere in the simulation (in the bulk and on arbitrary dielectric boundaries). Second, we describe how trivial variants of this algorithm can be combined such that second-order error is achieved in bulk regions (both isotropic and anisotropic), while first-order error remains on dielectric boundaries only. Since the dimensionality of the boundary is one less than that of the simulation volume, resonant frequencies

exhibit second-order convergence [29]. Reference [29] proved this for 1D, but common experience suggests this phenomenon holds in higher dimensions. For example, stair-stepping is known to incur zeroth-order error on dielectric boundaries, but shows first-order convergence of resonant frequencies.

5.3.1 Achieving the desired accuracy

The discrete dynamical equations in the FIT (Eqs. (5.7) and (5.8) with the grid field values as defined in Eqs. (5.3)–(5.6)) exactly represent Maxwell's equations; this is shown graphically in Fig. 5.1. The discretization error occurs in the constitutive relations, Eqs. (5.9) and (5.10), when converting from area-integrated quantities to line-integrated quantities. The accuracy of the discrete Maxwell equations as a whole is therefore determined by the accuracy of Eqs. (5.9) and (5.10). For $\mu(\mathbf{x}) = \mu_0$ everywhere, $\mathbb{M}_\mu = \mu_0 \mathbb{I}$ gives second-order error in homogeneous regions but only first-order error on dielectric boundaries because of the slope discontinuity in $\mathbf{B}(\mathbf{x})$. Again, as suggested by Ref. [29], this should only contribute second-order error to the calculation of resonant frequencies. It therefore suffices to focus exclusively on the error in $\mathbb{M}_\varepsilon$.

The following diagram illustrates the relationship between the discrete constitutive relation, Eq. (5.9), and the continuous one, $\mathbf{D}(\mathbf{x}) = \varepsilon(\mathbf{x}) \cdot \mathbf{E}(\mathbf{x})$,

$$\begin{array}{ccc} \mathbf{D} & \xrightarrow{\varepsilon^{-1}} & \mathbf{E} \\ \Big\downarrow \tilde{\mathcal{A}} & & \Big\downarrow \mathcal{L} \\ \mathbf{d} & \xrightarrow[\mathbb{M}_\varepsilon^{-1}]{} & \mathbf{e} \end{array} \tag{5.11}$$

where $\mathbf{D}(\mathbf{x})$ and $\mathbf{E}(\mathbf{x})$ are solutions to Eqs. (5.1) and (5.2) (i.e., boundary conditions are satisfied and $\nabla \cdot \mathbf{D}(\mathbf{x}) = 0$) and we have introduced the discretization operators, $\mathcal{L}$ and $\tilde{\mathcal{A}}$ which represent the integrations in Eqs. (5.3) and (5.4), respectively (the inverse permittivity is stressed because $\mathbf{d}$ is known from the discrete Ampere update, Eq. (5.7)).

The point of a dielectric algorithm is to build the operator, $\mathbb{M}_\varepsilon^{-1}$, such that Diagram 5.11 commutes. Formally, one can write

$$\mathbb{M}_\varepsilon^{-1} = \mathcal{L} \circ \hat{\varepsilon}^{-1} \circ \tilde{\mathcal{A}}^{-1}; \tag{5.12}$$

however, in practice, it is not clear what is meant by $\tilde{\mathcal{A}}^{-1}$, since a given discrete field $\mathbf{d}$ does not contain enough information to reconstruct $\mathbf{D}$ perfectly. Instead, one must form an approximate inverse, $\tilde{\mathcal{A}}_{\approx}^{-1}$, which maps a discrete field $\mathbf{d}$ to an approximation of $\mathbf{D}$. The discrete permittivity operator is then

$$\mathbb{M}_{\varepsilon}^{-1} = \mathcal{L} \circ \varepsilon^{-1} \circ \tilde{\mathcal{A}}_{\approx}^{-1}. \tag{5.13}$$

Once $\mathbb{M}_{\varepsilon}^{-1}$ is formed, the error can be determined by comparing the following two calculations of the discrete electric field (the two paths in Diagram 5.11) for arbitrary physical fields, $\mathbf{D}$:

$$\mathbf{e}' \equiv \mathcal{L} \circ \varepsilon^{-1} \circ \mathbf{D} \tag{5.14}$$

$$\mathbf{e} = \mathbb{M}_{\varepsilon}^{-1} \circ \tilde{\mathcal{A}} \circ \mathbf{D}. \tag{5.15}$$

The former is the reference calculation and the latter is the approximation. The error is then

$$\text{error} = \left\| \mathbf{e}' - \mathbf{e} \right\|. \tag{5.16}$$

To illustrate the above concepts in a simple case, we show how the Yee algorithm gives local second-order error for homogeneous isotropic dielectric ($\varepsilon(\mathbf{x}) = \varepsilon$ where ε is a scalar). To build $\mathbb{M}_{\varepsilon}^{-1}$, we start by defining $\tilde{\mathcal{A}}_{\approx}^{-1}$ to yield a piecewise constant field $\mathbf{D}^{\mathcal{P}}(\mathbf{x})$, where, for the case of the x-component, the value everywhere inside a grid cell box centered at $[(i+0.5)\Delta x, j\Delta y, k\Delta z]$ is equal to $d_{x|ijk}$ (similar definitions are made for the y and z components, with boxes centered on the grid field values). Equation (5.13) then gives

$$\begin{aligned}
\left(\mathbb{M}_{\varepsilon}^{-1} \circ \mathbf{d}\right)_{x|ijk} &= \left(\mathcal{L} \circ \varepsilon^{-1} \circ \tilde{\mathcal{A}}_{\approx}^{-1} \circ \mathbf{d}\right)_{x|ijk} \\
&= \frac{1}{\varepsilon \Delta x} \int_{i\Delta x}^{(i+1)\Delta x} D_x^{\mathcal{P}}(\mathbf{x})\, dx \\
&= \frac{1}{\varepsilon} d_{x|ijk}
\end{aligned} \tag{5.17}$$

which, when applied to all vector components, gives $\mathbb{M}_{\varepsilon}^{-1} = \varepsilon^{-1}\mathbb{I}$ where $\mathbb{I}$ is the identity matrix. This is the discrete permittivity operator used by the Yee algorithm.

The error introduced by Eq. (5.17) is now determined using Eq. (5.16). An arbitrary $\mathbf{D}$ field will be represented as a Taylor expansion about the point $\mathbf{x}_0 = [(i+0.5)\Delta x, j\Delta y, k\Delta z]$. We

concentrate on the error in the calculation of a single component $e_{x|ijk}$ from $\mathbf{D}$. To first order in $(\mathbf{x}-\mathbf{x}_0)$, Eq. (5.14) for a single component becomes

$$\begin{aligned} e'_{x|ijk} &\equiv \left(\mathcal{L}\circ\varepsilon^{-1}\circ\mathbf{D}\right)_{x|ijk} \\ &= \frac{1}{\varepsilon\Delta x}\int_{L_{x|ijk}}\left[D_x\big|_{\mathbf{x}_0}+\nabla D_x\big|_{\mathbf{x}_0}\cdot(\mathbf{x}-\mathbf{x}_0)+\cdots\right]dx \\ &= \frac{1}{\varepsilon}D_x\big|_{\mathbf{x}_0}+O(\Delta x^2) \end{aligned} \tag{5.18}$$

where the integration over the first-order term vanishes because $\mathbf{x}_0$ is at the midpoint of $L_{x|ijk}$. Similarly, for Eq. (5.15),

$$\begin{aligned} e_{x|ijk} &= \left(\mathbb{M}_\varepsilon^{-1}\circ\tilde{\mathcal{A}}\circ\mathbf{D}\right)_{x|ijk} \\ &= \frac{1}{\varepsilon\Delta y\Delta z}\int_{\tilde{A}_{x|ijk}}\left[D_x\big|_{\mathbf{x}_0}+\nabla D_x\big|_{\mathbf{x}_0}\cdot(\mathbf{x}-\mathbf{x}_0)+\cdots\right]dy\,dz \\ &= \frac{1}{\varepsilon}D_x\big|_{\mathbf{x}_0}+O(\Delta y^2,\Delta z^2) \end{aligned} \tag{5.19}$$

where this integration over the first-order term vanishes because $\mathbf{x}_0$ is also the midpoint of the area $\tilde{A}_{x|ijk}$. The error is clearly second-order since, for every component,

$$e'_{x|ijk}-e_{x|ijk}=O(\Delta x^2,\Delta y^2,\Delta z^2). \tag{5.20}$$

This second-order error is quite fortuitous, since the piecewise constant field we have defined by $\tilde{\mathcal{A}}_{\approx}^{-1}$ has first-order error at any given point (except for the midpoints of the integration regions where the error would be second-order). The cooperation of the operators $\mathcal{L}$ and $\tilde{\mathcal{A}}$ is responsible for this result; the collocation of the midpoints of the line and area integration regions ensures that any first-order terms vanish in the above error analysis.

In the Introduction, we claimed that the algorithm of Ref. [58] (which is similar to our algorithm) does not ensure second-order error in resonant frequencies; the above method shows that the algorithm of Ref. [58] incurs resolution-independent error at dielectric boundaries, and thus should give (at best) first-order error in resonant frequencies. A key step in the algorithm of Ref. [58] involves an equally-weighted interpolation of $d_{\alpha|ijk}$, even across dielectric boundaries

(where the goal is to find the integral of $D_\alpha(\mathbf{x})$ over some nonstandard finite area). While the interpolation gives second-order error for continuous fields, it gives zeroth-order error at dielectric boundaries because the fields are discontinuous. To see this, one can analyze the commutativity of the following diagram,

$$\begin{array}{ccc} \mathbf{D} & \longrightarrow & \mathbf{D} \\ \Big\downarrow{\scriptstyle \tilde{\mathcal{A}}} & & \Big\downarrow{\scriptstyle \tilde{\mathcal{A}}_{\mathcal{I}}} \\ \mathbf{d} & \xrightarrow[\mathcal{I}]{} & \mathbf{d}_{\mathcal{I}} \end{array} \tag{5.21}$$

where $\mathcal{I}$ is the interpolation operator and the subscript $\mathcal{I}$ indicates discretization on nonstandard finite integration areas ($\mathbf{d}_{\mathcal{I}}$ is the result of a translated version of Eq. (5.4)). To determine the error, the reference calculation is, $\mathbf{d}'_{\mathcal{I}} \equiv \tilde{\mathcal{A}}_{\mathcal{I}} \circ \mathbf{D}$, whereas the interpolation algorithm gives, $\mathbf{d}_{\mathcal{I}} = \mathcal{I} \circ \tilde{\mathcal{A}} \circ \mathbf{D}$. If one uses the equally-weighted interpolation of Ref. [58] as $\mathcal{I}$, then one finds that $\mathbf{d}'_{\mathcal{I}} - \mathbf{d}_{\mathcal{I}} = O(1)$ for field values at dielectric boundaries, because of the discontinuity in $\mathbf{D}$.

In this paper, our first goal is to find a permittivity operator that gives first-order error, even at boundaries between tensor dielectric constants. Although our first attempt will yield an operator with first-order error in homogeneous regions (as well as at boundaries), we will later modify it to ensure second-order error in homogeneous regions while retaining first-order error on boundaries. The work of Ref. [29] suggests that this will give second-order error in the calculation of resonant frequencies since the first-order error is restricted to a dimension one less than that of the simulation volume (the ratio of boundary grid cells to total simulation grid cells is $O(\Delta x)$).

5.3.2 A boundary algorithm for anisotropic dielectrics with first-order error

We will derive a first-order algorithm following the method described in Sec. 5.3.1 for homogeneous isotropic dielectric. First, we reformulate the problem to simplify the treatment of dielectric boundary conditions. Second, we find an approximate inverse operator that will map a given discrete field $\mathbf{d}$ to an approximation of a physical field that satisfies the dielectric boundary conditions to first-order error. Finally, after the operator $\mathbb{M}_{\varepsilon}^{-1}$ is formed, the first-order error of the algorithm is verified using Eq. (5.16).

Reformulation

To simplify the treatment of the fields near a boundary, we use a coordinate system that conforms to the dielectric interface. At each point on the surface, the local basis vectors are $\hat{\mathbf{n}}$ (the normal to the surface), $\hat{\mathbf{t}}$, and $\hat{\boldsymbol{\tau}}$ (two orthogonal unit vectors tangent to the surface). In this coordinate system, the field components, D_n, E_t, and E_τ are continuous across a dielectric boundary, since

$$\hat{\mathbf{n}} \cdot \mathbf{D}\big|_{S^+} = \hat{\mathbf{n}} \cdot \mathbf{D}\big|_{S^-} \tag{5.22}$$

$$\hat{\mathbf{n}} \times \mathbf{E}\big|_{S^+} = \hat{\mathbf{n}} \times \mathbf{E}\big|_{S^-} \tag{5.23}$$

where S is the interface and the $+/-$ signs indicate one-sided limits. The missing components, E_n, D_t, and D_τ are discontinuous across the boundary, but are completely determined from the continuous components through $\mathbf{D}(\mathbf{x}) = \varepsilon(\mathbf{x}) \cdot \mathbf{E}(\mathbf{x})$. It will be easier to work with D_n, E_t, and E_τ whenever possible. To this end, as in Ref. [45], we define the pseudo-field, $\mathbf{F} \equiv [D_n, E_t, E_\tau]^{\mathrm{T}}$, which we can write in a coordinate-independent way without using $\hat{\mathbf{t}}$ or $\hat{\boldsymbol{\tau}}$ (which are arbitrary for smooth interfaces):

$$\mathbf{F} \equiv \hat{\mathbf{n}}\hat{\mathbf{n}} \cdot \mathbf{D} + (1 - \hat{\mathbf{n}}\hat{\mathbf{n}}) \cdot \mathbf{E} = (\hat{\mathbf{n}}\hat{\mathbf{n}} \cdot \varepsilon + 1 - \hat{\mathbf{n}}\hat{\mathbf{n}}) \cdot \mathbf{E} \tag{5.24}$$

where $\hat{\mathbf{n}}\hat{\mathbf{n}}$ is the 3×3 projection matrix onto the normal.

The quantity $\mathbf{F}$ is especially useful for performing Taylor expansions at dielectric boundaries. For example, consider a planar dielectric boundary with normal $\hat{\mathbf{n}}$. We can expand $\mathbf{F}$ about a point $\mathbf{x}_0$ on the boundary to first order:

$$\mathbf{F}(\mathbf{x}) = \mathbf{F}\big|_{\mathbf{x}_0} + \Delta\mathbf{x}_S \cdot \nabla_S \mathbf{F}\big|_{\mathbf{x}_0} + \Delta x_n \left(\Theta(\Delta x_n) \frac{\partial \mathbf{F}}{\partial n}\Big|_{\mathbf{x}_0^+} + \Theta(-\Delta x_n) \frac{\partial \mathbf{F}}{\partial n}\Big|_{\mathbf{x}_0^-} \right) + \ldots \tag{5.25}$$

where $\Delta\mathbf{x}_S$ is the tangential displacement from $\mathbf{x}_0$, Δx_n is the displacement in the normal direction and Θ is the Heaviside step function. The above expansion naturally incorporates the proper boundary conditions, Eqs. (5.22) and (5.23).

We can include the $\mathbf{F}$ field in Diagram (5.11) in the following way,

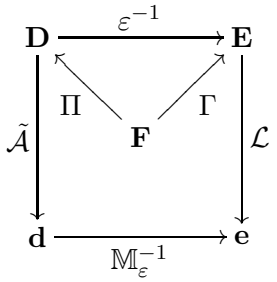

(5.26)

where Π and Γ are 3×3 matrices that give $\mathbf{D}(\mathbf{x}) = \Pi(\mathbf{x})\cdot\mathbf{F}(\mathbf{x})$ and $\mathbf{E}(\mathbf{x}) = \Gamma(\mathbf{x})\cdot\mathbf{F}(\mathbf{x})$, respectively, and can be written as

$$\Gamma = \left(\hat{\mathbf{n}}\hat{\mathbf{n}}\cdot\varepsilon + (1-\hat{\mathbf{n}}\hat{\mathbf{n}})\right)^{-1} \tag{5.27}$$

$$= 1 + \hat{\mathbf{n}}\hat{\mathbf{n}}\cdot\frac{1-\varepsilon}{\hat{\mathbf{n}}\cdot\varepsilon\cdot\hat{\mathbf{n}}} \tag{5.28}$$

$$\Pi = \varepsilon\cdot\left(\hat{\mathbf{n}}\hat{\mathbf{n}}\cdot\varepsilon + (1-\hat{\mathbf{n}}\hat{\mathbf{n}})\right)^{-1} \tag{5.29}$$

$$= \varepsilon\cdot\left(1 + \hat{\mathbf{n}}\hat{\mathbf{n}}\cdot\frac{1-\varepsilon}{\hat{\mathbf{n}}\cdot\varepsilon\cdot\hat{\mathbf{n}}}\right). \tag{5.30}$$

where we have kept the expressions coordinate-independent. Clearly, $\Pi\cdot\Gamma^{-1} = \varepsilon$, as required.

In analogy with Eq. (5.13), the discrete permittivity operator is constructed by,

$$\mathbb{M}_\varepsilon^{-1} = \mathcal{L}\circ\varepsilon^{-1}\circ\tilde{\mathcal{A}}_{\approx}^{-1} = \mathcal{L}\circ\Gamma\circ\left(\tilde{\mathcal{A}}\circ\Pi\right)_{\approx}^{-1}. \tag{5.31}$$

The problem is reformulated in this way to simplify the calculations in the following sections, because $\mathcal{L}\circ\Gamma$ and $\tilde{\mathcal{A}}\circ\Pi$ operate on $\mathbf{F}$ which is continuous across dielectric boundaries. Whereas in the Yee example, the essential step was finding $\tilde{\mathcal{A}}_{\approx}^{-1}$, now the essential step is finding $(\tilde{\mathcal{A}}\circ\Pi)_{\approx}^{-1}$, which we describe in the following section.

Constructing an approximate F from d

Following the Yee example, we define an approximate inverse operator $(\tilde{\mathcal{A}}\circ\Pi)_{\approx}^{-1}$, that maps $\mathbf{d}$ to a piecewise constant field, $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$. This field is constant on each "subdomain," $\mathcal{P}_{ijk}$; that is, $\mathbf{F}^{\mathcal{P}}(\mathbf{x}) = \mathbf{F}_{ijk}$ for $\mathbf{x}\in\mathcal{P}_{ijk}$. Because the subdomains we eventually employ are not simple cubes, the precise definition of the subdomains will be left until later in the derivation (for now, it can be said that the volume of each subdomain is that of a usual grid cell).

For dielectric objects with curved interfaces (no corners or edges), $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$ will be a first-order (in the size of the subdomains) approximation to the physical $\mathbf{F}(\mathbf{x})$ at any given point. This is obvious for subdomains that lie within regions of uniform dielectric constant, where the electric and displacement fields are continuous. More important, $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$ is a first-order approximation to $\mathbf{F}(\mathbf{x})$ within subdomains that are intersected by a dielectric boundary. For smooth boundaries, the portion of the surface that intersects subdomain $\mathcal{P}_{ijk}$ becomes *exactly* planar as the size of the subdomain vanishes. Since a constant $\mathbf{F}$ (e.g. $\mathbf{F}_{ijk}$ in $\mathcal{P}_{ijk}$) at a planar boundary is the zeroth order term in the Taylor expansion of $\mathbf{F}(\mathbf{x})$, it follows that $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$ approximates $\mathbf{F}(\mathbf{x})$ with first-order error in the size of the subdomains.

The operator $(\tilde{\mathcal{A}} \circ \Pi)^{-1}_{\approx}$ acts on a given $\mathbf{d}$ to produce $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$. As a first step in constructing this operator, we look at the simpler calculation, $\mathbf{d} = \tilde{\mathcal{A}} \circ \Pi \circ \mathbf{F}^{\mathcal{P}}(\mathbf{x})$ for a single component, $d_{x|ijk}$,

$$d_{x|ijk} = \frac{1}{\Delta y \Delta z} \int_{\tilde{A}_{x|ijk}} \hat{\mathbf{x}} \cdot \Pi(\mathbf{x}) \cdot \mathbf{F}^{\mathcal{P}}(\mathbf{x})\, dy\, dz. \tag{5.32}$$

If we assume that the integration region $\tilde{A}_{x|ijk}$ lies entirely within a single subdomain, say $\mathcal{P}_{ijk}$, then the above becomes

$$d_{x|ijk} = \hat{\mathbf{x}} \cdot \Pi_{x|ijk} \cdot \mathbf{F}_{ijk} \tag{5.33}$$

where we have made the definition,

$$\Pi_{x|ijk} \equiv \frac{1}{\Delta y \Delta z} \int_{\tilde{A}_{x|ijk}} \Pi(\mathbf{x})\, da. \tag{5.34}$$

We now have a linear relationship between the single component $d_{x|ijk}$ and $\mathbf{F}_{ijk}$. However, we still cannot form an inverse because the three vector components in $\mathbf{F}_{ijk}$ to be determined are matched by only one known $d_{x|ijk}$. To invert, we will need to relate $\mathbf{F}_{ijk}$ to two more knowns, $d_{y|i'j'k'}$ and $d_{z|i''j''k''}$, preferably adjacent to $d_{x|ijk}$.

At this point, it is helpful to define a "triplet" of grid values (see Fig. 5.2). A triplet will always contain all three Cartesian vector components and will be defined relative to a regular grid node. For example, $\mathbf{d}_{ijk}^{\pm\pm\pm}$ will refer to a triplet of $d_{\alpha|ijk}$ values adjacent to the node, $\mathcal{N}_{ijk}$, where the first, second, and third '$\pm$' signs indicate the relative positions (in the component directions)

of the x, y, and z components, respectively. There are eight possible triplets for any given ijk. As an explicit example, $\mathbf{d}_{ijk}^{+-+} = d_{x|ijk}\hat{\mathbf{x}} + d_{y|ij-1k}\hat{\mathbf{y}} + d_{z|ijk}\hat{\mathbf{z}}$.

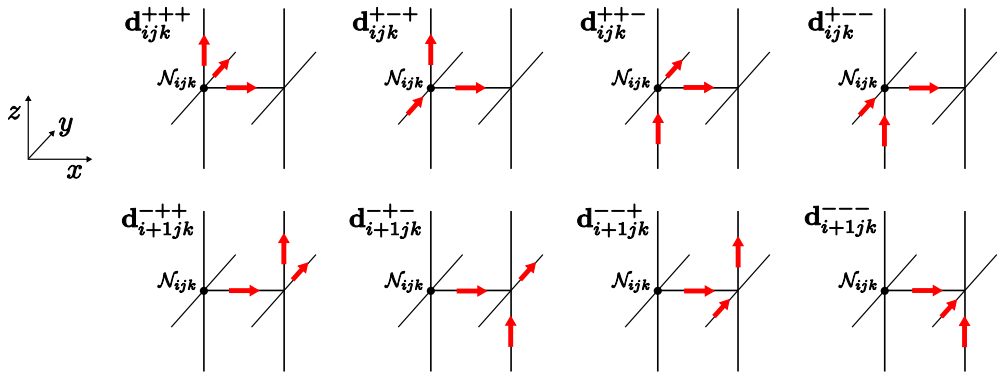

Figure 5.2: *Color online. Shown above are several examples of "triplets" of grid components. The one grid component common to each of the above triplets is* $d_{x|ijk}$.

We would now like to associate $\mathbf{F}_{ijk}$ with a triplet of $d_{\alpha|ijk}$ values. As a representative example, we will choose the triplet, $\mathbf{d}_{ijk}^{+++}$ (the choice of '$\pm\pm\pm$' is arbitrary). The one requirement we must make before linking $\mathbf{F}_{ijk}$ to the triplet $\mathbf{d}_{ijk}^{+++}$, is that all integration regions, $\tilde{A}_{(x,y,z)|ijk}$ must lie within the subdomain $\mathcal{P}_{ijk}$. Then, Eq. (5.33) can be applied to each $d_{\alpha|ijk}$ in $\mathbf{d}_{ijk}^{+++}$, and the resulting system of equations can be written in matrix form,

$$\mathbf{d}_{ijk}^{+++} = \Pi_{ijk}^{+++} \cdot \mathbf{F}_{ijk} \tag{5.35}$$

where the 3×3 matrix Π_{ijk}^{+++} is

$$\Pi_{ijk}^{+++} = \hat{\mathbf{x}}\hat{\mathbf{x}} \cdot \Pi_{x|ijk} + \hat{\mathbf{y}}\hat{\mathbf{y}} \cdot \Pi_{y|ijk} + \hat{\mathbf{z}}\hat{\mathbf{z}} \cdot \Pi_{z|ijk}. \tag{5.36}$$

Our goal is to find $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$ from $\mathbf{d}$, so we really want the inverse of Eq. (5.35); this can now be accomplished by inverting the 3×3 matrix, Π_{ijk}^{+++}. In our tests of the algorithm, Π_{ijk}^{+++} has never been singular. Invertibility can be proven for interfaces between isotropic dielectrics, in which case the determinant has the following analytical form,

$$\left|\Pi_{ijk}^{+++}\right| = n_x^2 \bar{\varepsilon}_{y|ijk} \bar{\varepsilon}_{z|ijk} + n_y^2 \bar{\varepsilon}_{z|ijk} \bar{\varepsilon}_{x|ijk} + n_z^2 \bar{\varepsilon}_{x|ijk} \bar{\varepsilon}_{y|ijk} \tag{5.37}$$

where $\bar{\varepsilon}_{x|ijk} \equiv \int_{\tilde{A}_{x|ijk}} \varepsilon(\mathbf{x}) da/\Delta y \Delta z$ and $n_x = \hat{\mathbf{n}} \cdot \hat{\mathbf{x}}$, etc. This determinant is nonzero and positive since $n_x^2 + n_y^2 + n_z^2 = 1$ and $\bar{\varepsilon}_{\alpha|ijk} > 0$.

We can now form the entire approximate inverse, $(\tilde{\mathcal{A}} \circ \Pi)_{\approx}^{-1}$, by inverting Eq. (5.35) for every triplet of the discrete field:

$$\mathbf{F}^{\mathcal{P}}(\mathbf{x}) = \sum_{ijk} \begin{cases} \Pi_{ijk}^{+++-1} \cdot \mathbf{d}_{ijk}^{+++} & \text{if } \mathbf{x} \in \mathcal{P}_{ijk} \\ 0 & \text{otherwise.} \end{cases} \tag{5.38}$$

The above is just one of several approximate inverses, since the choice of triplet is arbitrary (as long as the subdomains are defined appropriately). In fact, an inverse could be formed from a weighted linear combination of all eight '$\pm\pm\pm$' specifications for each ijk.

Constructing e from the approximate F

With the approximate inverse defined in the last section, we now complete the formation of $\mathbb{M}_{\varepsilon}^{-1}$ via Eq. (5.31). We start by considering the calculation $\mathbf{e} = \mathcal{L} \circ \Gamma \circ \mathbf{F}^{\mathcal{P}}(\mathbf{x})$ for the single component, $e_{x|ijk}$,

$$e_{x|ijk} = \frac{1}{\Delta x} \int_{L_{x|ijk}} \hat{\mathbf{x}} \cdot \Gamma(\mathbf{x}) \cdot \mathbf{F}^{\mathcal{P}}(\mathbf{x})\, dx. \tag{5.39}$$

Just as in the previous section, we will assume that the integration region $L_{x|ijk}$ lies within the subdomain $\mathcal{P}_{ijk}$, so that the above becomes

$$e_{x|ijk} = \hat{\mathbf{x}} \cdot \Gamma_{x|ijk} \cdot \mathbf{F}_{ijk} \tag{5.40}$$

where

$$\Gamma_{x|ijk} \equiv \frac{1}{\Delta x} \int_{L_{x|ijk}} \Gamma(\mathbf{x})\, dx. \tag{5.41}$$

We would now like to associate the same triplet of $e_{\alpha|ijk}$ values with $\mathbf{F}_{ijk}$ as we did for $d_{\alpha|ijk}$ values in the last section (i.e. if $\mathbf{d}_{ijk}^{+++}$ was linked with $\mathbf{F}_{ijk}$, then we should also link $\mathbf{e}_{ijk}^{+++}$ with $\mathbf{F}_{ijk}$). This requires that the integration regions $L_{(x,y,z)|ijk}$ all lie within subdomain $\mathcal{P}_{ijk}$. We can then write (similar to Eq. (5.35)),

$$\mathbf{e}_{ijk}^{+++} = \Gamma_{ijk}^{+++} \cdot \mathbf{F}_{ijk} \tag{5.42}$$

where Γ^{+++}_{ijk} is

$$\Gamma^{+++}_{ijk} = \hat{\mathbf{x}}\hat{\mathbf{x}} \cdot \Gamma_{x|ijk} + \hat{\mathbf{y}}\hat{\mathbf{y}} \cdot \Gamma_{y|ijk} + \hat{\mathbf{z}}\hat{\mathbf{z}} \cdot \Gamma_{z|ijk}. \tag{5.43}$$

This completes the requirements for the subdomains; in summary, subdomain $\mathcal{P}_{ijk}$ must contain,

(1) $\tilde{A}_{x|ijk}$, $\tilde{A}_{y|ijk}$, and $\tilde{A}_{z|ijk}$ so that Eq. (5.35) follows from Eq. (5.32), and

(2) $L_{x|ijk}$, $L_{y|ijk}$, and $L_{z|ijk}$ so that Eq. (5.42) follows from Eq. (5.39).

An illustration of these subdomains in 2D is shown in Fig. 5.3.

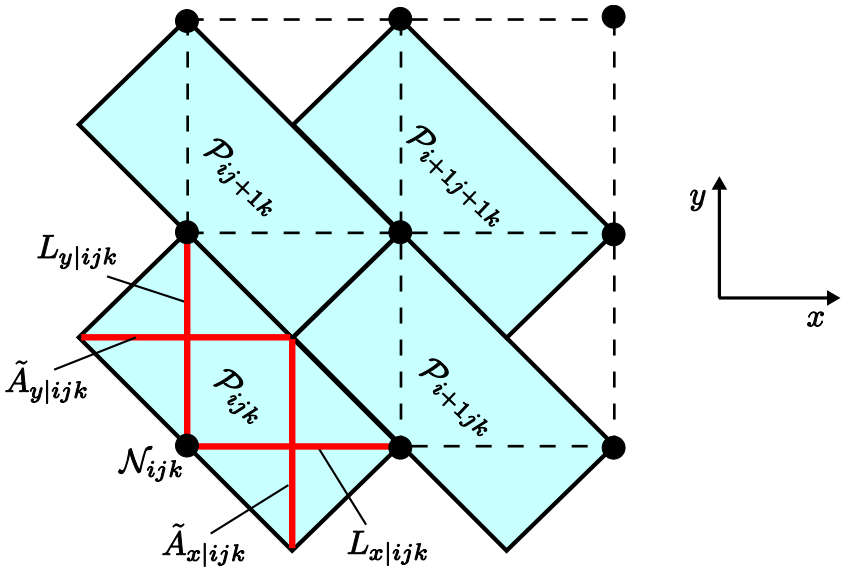

Figure 5.3: *Color online. When* $\mathbf{e}^{+++}_{ijk}$, $\mathbf{d}^{+++}_{ijk}$, *and* $\mathbf{F}_{ijk}$ *are linked, the subdomains* $\mathcal{P}_{ijk}$ *in 2D take the shape shown in the above illustration. Notice that the subdomain* $\mathcal{P}_{ijk}$ *contains all integration regions associated with* $\mathbf{e}^{+++}_{ijk}$ *and* $\mathbf{d}^{+++}_{ijk}$.

Using Eqs. (5.35) and (5.42), we can eliminate $\mathbf{F}_{ijk}$ altogether to find the operator Ξ^{+++}_{ijk} that calculates $\mathbf{e}^{+++}_{ijk}$ from $\mathbf{d}^{+++}_{ijk}$:

$$\mathbf{e}^{+++}_{ijk} = \Xi^{+++}_{ijk} \cdot \mathbf{d}^{+++}_{ijk} \equiv \Gamma^{+++}_{ijk} \cdot {\Pi^{+++}_{ijk}}^{-1} \cdot \mathbf{d}^{+++}_{ijk}. \tag{5.44}$$

The action of the operator Ξ^{+++}_{ijk} is illustrated in Fig. 5.4, which highlights the importance of the finite integral representation (in which Ξ^{+++}_{ijk} takes area-integrated quantities into line-integrated quantities). When subdomain $\mathcal{P}_{ijk}$ lies entirely within the same region of uniform tensor dielectric

constant ε, we have $\Xi_{ijk}^{+++} = \varepsilon^{-1}$, which continues to give first-order error when ε is nondiagonal because of the asymmetric distribution of the triplet components. When ε (and ε^{-1}) is diagonal, the usual Yee operation with second-order error is recovered (e.g. when the dielectric is isotropic or has its principal axes aligned with the grid directions).

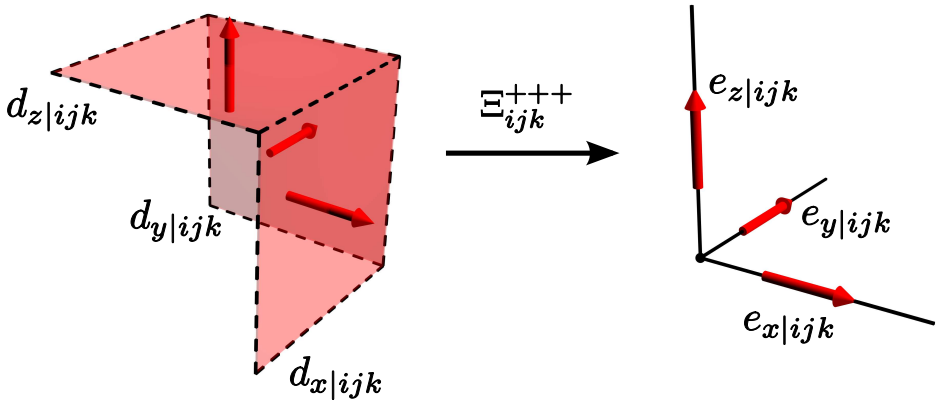

Figure 5.4: *Color online. An illustration of the operator Ξ_{ijk}^{+++} in 3D, which computes a triplet of line-integrated electric field values from the same triplet of area-integrated displacement field values, even in the presence of dielectric boundaries.*

Finally, the discrete permittivity operator can be written as the 3×3 block-diagonal matrix,

$$\mathbb{M}_{\varepsilon}^{-1} = \mathbb{M}_{\varepsilon}^{+++-1} \equiv \begin{pmatrix} \ddots & & \\ & \Xi_{ijk}^{+++} & \\ & & \ddots \end{pmatrix} \tag{5.45}$$

where each block is for a unique set of indices ijk. Again, the choice of '$\pm\pm\pm$' is arbitrary; $\mathbb{M}_{\varepsilon}^{\pm\pm\pm-1}$ are all discrete permittivity matrices with first-order error (in practice, for the matrix to appear in block-diagonal form as in Eq. (5.45), a very specific ordering of the components in the vectors $\mathbf{e}$ and $\mathbf{d}$ is required).

In Sec. 5.3.1 we discussed how the only approximation in discretizing Maxwell's equations using the FIT is in the constitutive relations (that is, in finding $\mathbf{e}$, $\mathbf{h}$ from $\mathbf{d}$, $\mathbf{b}$). Equation (5.45) describes that relationship; its accuracy determines the accuracy of the entire algorithm. The

formation of the Ξ_{ijk}^{+++} operator so that Eq. (5.45) has $O(\Delta x)$ error (instead of $O(1)$ error)—even on dielectric boundaries—is the most important advance of this work.

Error

We now demonstrate the first-order error of the algorithm, Eq. (5.45), by calculating a single component of $\mathbf{e}$ at a dielectric boundary for an arbitrary field, $\mathbf{F}(\mathbf{x})$ (as in Sec. 5.3.1). We first take the limit $\Delta x \to 0$, so that the local surface is a plane and the Taylor expansion (by Eq. (5.25)) to zeroth-order is simply $\mathbf{F}(\mathbf{x}) = \mathbf{F}|_{\mathbf{x}_0} + O(\mathbf{x} - \mathbf{x}_0)$, where the expansion point is some point on the boundary near the component, $e_{x|ijk}$. Then, we perform the reference calculation, which gives

$$e'_{x|ijk} \equiv (\mathcal{L} \circ \Gamma \circ \mathbf{F}(\mathbf{x}))_{x|ijk} \tag{5.46}$$

$$= \hat{\mathbf{x}} \cdot \Gamma_{x|ijk} \cdot \mathbf{F}\big|_{\mathbf{x}_0} + O(\Delta x, \Delta y, \Delta z) \tag{5.47}$$

The algorithm calculates

$$e_{x|ijk} = \left(\mathbb{M}_{\varepsilon}^{-1} \circ \tilde{\mathcal{A}} \circ \Pi \circ \mathbf{F}(\mathbf{x})\right)_{x|ijk} \tag{5.48}$$

$$= \left(\Gamma_{ijk}^{+++} \cdot \Pi_{ijk}^{+++\,-1} \cdot \Pi_{ijk}^{+++} \cdot \mathbf{F}\big|_{\mathbf{x}_0}\right)_{x|ijk} + O(\Delta x, \Delta y, \Delta z) \tag{5.49}$$

$$= \hat{\mathbf{x}} \cdot \Gamma_{x|ijk} \cdot \mathbf{F}\big|_{\mathbf{x}_0} + O(\Delta x, \Delta y, \Delta z) \tag{5.50}$$

so that

$$e'_{x|ijk} - e_{x|ijk} = O(\Delta x, \Delta y, \Delta z) \tag{5.51}$$

since the first-order parts do not necessarily cancel in this case. Therefore, our algorithm has ensured the perfect calculation of the zeroth-order parts of the fields in the limit $\Delta x \to 0$, as expected (i.e., in the limit of infinite wavelength and planar boundary, the algorithm is exact).

5.3.3 Restoring second-order error in bulk anisotropic dielectric regions

The algorithm described thus far has second-order error in bulk isotropic dielectric regions, and first-order error in bulk anisotropic dielectric regions and on dielectric boundaries. However, we have actually found eight different algorithms (each with a different '$\pm\pm\pm$') with first-order error.

We now show that a superposition of all eight algorithms can give second-order error in regions of homogeneous anisotropic dielectric, while retaining first-order error on dielectric boundaries.

A second-order method for simulating bulk anisotropic dielectrics is

$$\begin{aligned} e_{x|ijk} = \xi_{xx} d_{x|ijk} + \\ \frac{1}{4}\xi_{xy}\left(d_{y|ijk} + d_{y|i+1,jk} + d_{y|ij-1,k} + d_{y|i+1,j-1,k}\right) + \\ \frac{1}{4}\xi_{xz}\left(d_{z|ijk} + d_{z|i+1,jk} + d_{z|ijk-1} + d_{z|i+1,jk-1}\right) \end{aligned} \tag{5.52}$$

where $\xi \equiv \varepsilon^{-1}$. The above is a centered interpolation of the nearest-neighbor components (Fig. 5.2 shows all components involved in the above) [79].

As mentioned in Sec. , the operators $\Xi_{ijk}^{\pm\pm\pm}$ reduce to ξ in regions of uniform dielectric constant. Therefore, we can recover Eq. (5.52) for these regions using the following linear combination of $\mathbb{M}_{\varepsilon}^{\pm\pm\pm\,-1}$ matrices:

$$\begin{aligned} \mathbb{M}_{\varepsilon}^{-1} = \frac{1}{8}\Big(\mathbb{M}_{\varepsilon}^{+++\,-1} + \mathbb{M}_{\varepsilon}^{++-\,-1} + \mathbb{M}_{\varepsilon}^{+-+\,-1} + \mathbb{M}_{\varepsilon}^{+--\,-1} + \\ \mathbb{M}_{\varepsilon}^{-++\,-1} + \mathbb{M}_{\varepsilon}^{-+-\,-1} + \mathbb{M}_{\varepsilon}^{--+\,-1} + \mathbb{M}_{\varepsilon}^{---\,-1}\Big). \end{aligned} \tag{5.53}$$

The above now has second-order error in regions of uniform (possibly anisotropic) dielectric (where it is the same as Eq. (5.52)), and retains first-order error on dielectric boundaries. The error analysis of Sec. is still valid for the above, which simply describes a superposition of different $\mathbf{F}^{\mathcal{P}}(\mathbf{x})$ fields, each guaranteeing first-order error in the electric and displacement fields at any given point. Equation (5.53) is the permittivity operator we test in the next section. For a succinct description of our algorithm, and the steps required to implement it, see the Appendix.

5.4 Numerical verification

We examined the second-order convergence of resonant frequencies for the above algorithm in two different situations. First, we simulated a dielectric sphere inside of and concentric with a spherical metal cavity. The dielectric sphere had an isotropic dielectric constant, admitting analytic solutions for the comparison of cavity mode frequencies and surface fields. Next we simulated a

photonic crystal of dielectric ellipsoids with anisotropic dielectric constant. The principal axes of the ellipsoids and the dielectric tensor were oblique to the grid directions. Mode frequencies were calculated by eigensolving the wave equation for the magnetic field,

$$\mathbb{C}^{\mathrm{T}}\mathbb{M}_{\varepsilon}^{-1}\mathbb{C}\mathbb{M}_{\mu}^{-1}\mathbf{b} = \omega^2\mathbf{b}. \tag{5.54}$$

where ω is the angular frequency and $\mathbb{M}_{\mu}^{-1} = \mu_0^{-1}\mathbb{I}$. We used the TRILINOS software libraries for all sparse matrix computations in these problems [32].

5.4.1 Convergence of resonant frequencies: dielectric in spherical cavity

The dielectric sphere had a radius $a = 0.37$m and the spherical metal cavity had a radius of $b = 0.49$m. The isotropic dielectric constant of the dielectric sphere was $\varepsilon = 10\varepsilon_0$ and the rest of the cavity was at the vacuum permittivity ε_0. To reduce any beneficial effects of grid symmetries, the coincident centers of the spheres were offset from any grid nodes and the grid cells were given length ratios $\Delta y/\Delta x = 5/4$ and $\Delta z/\Delta x = 6/4$.

Analytic mode frequencies for this cavity can be found by expressing the fields in terms of vector spherical harmonics. The solutions come in two classes, transverse electric (TE_{nlm}) and transverse magnetic (TM_{nlm}) where n is the radial wavenumber and l and m are the polar and azimuthal wavenumbers, respectively. The mode frequencies are determined by the radial and polar wavenumbers only, giving each frequency a $(2l+1)$-fold degeneracy [37].

To simulate the curved metallic boundary, we implemented the Dey-Mittra algorithm which also exhibits second-order convergence for resonant frequencies [15, 61]. As a side-note, the Dey-Mittra algorithm introduces several complications to standard numerical eigensolving routines. The methods we eventually employed to circumvent these difficulties are beyond the scope of this paper. However, we will be detailing these methods in a forthcoming publication.

Figure 5.5 shows the second-order convergence of the lowest five cavity mode frequencies as a function of the number of cells (in the z-direction) that resolve one vacuum wavelength (the vacuum wavelength is $\lambda_{\mathrm{vac}} = c/f$ where f is the reference eigenfrequency). Since $\Delta z > \Delta x, \Delta y$,

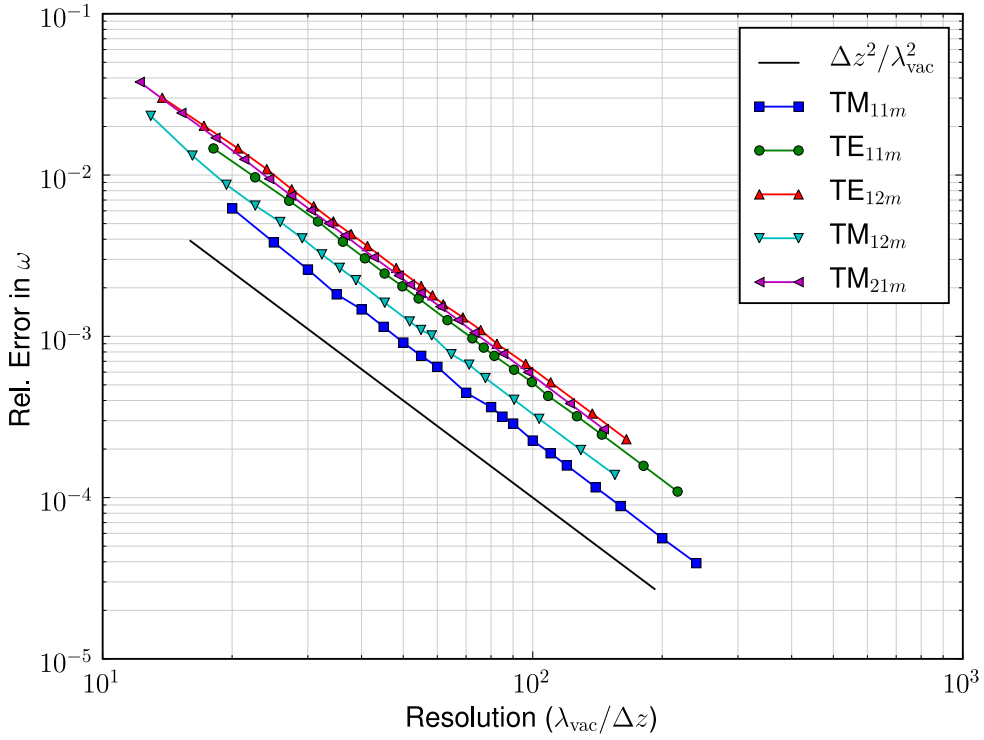

Figure 5.5: *Color online. Second-order convergence of cavity eigenmode frequencies for isotropic dielectric sphere inside and concentric with a perfectly conducting spherical metal cavity.* λ_{vac} *is the wavelength of each mode in vacuum* ($\lambda_{\mathrm{vac}} = c/f$ *where* f *is the eigenfrequency).*

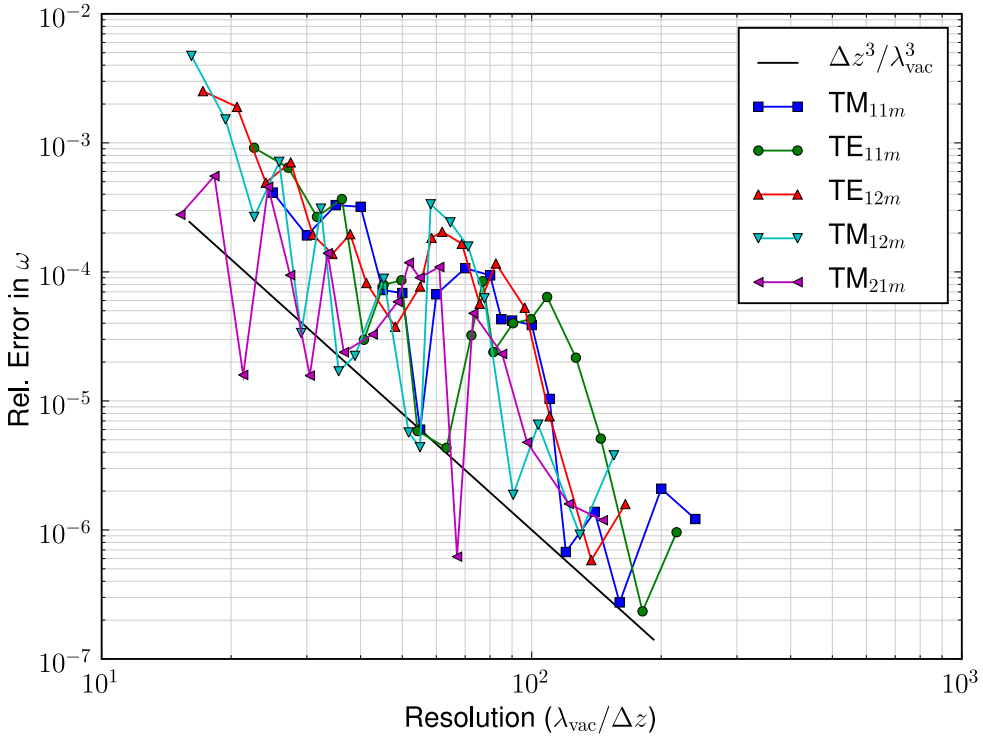

Figure 5.6: *Color online. Richardson extrapolation of cavity eigenfrequencies found with second-order algorithm leads to roughly third-order convergence.* λ_{vac} *is the wavelength of each mode in vacuum* ($\lambda_{\mathrm{vac}} = c/f$ *where* f *is the eigenfrequency). The resolution of each datapoint refers to the finer of the two resolutions used to Richardson extrapolate.*

the resolution in the z direction should be the dominant source of error. Because of the smooth convergence, we can also use Richardson extrapolation on neighboring frequency values to achieve roughly third-order convergence as shown in Fig. 5.6.

5.4.2 Convergence of surface fields: dielectric in spherical cavity

Electric fields near the surface of the dielectric sphere were calculated and compared against analytic solutions. For these simulations, parity symmetry was used to reduce the simulation domain to the positive x, y, and z octant (the spheres were centered at the origin). Grid cells were cubic ($\Delta x = \Delta y = \Delta z$). To compare simulations of different resolutions against the analytic solution, we compared field values interpolated to a number of points distributed throughout the region of interest. Since interpolation near the dielectric interface is equivocal and nontrivial (requiring a method based on the principles outlined in section 3.1 or one-sided extrapolation from the bulk), we defined the region of interest to be at least $\pm 3\Delta x$ from the interface. This ensured that, using a trilinear interpolation scheme, none of the integration regions associated with grid field values used in the interpolation were cut by the dielectric boundary.

The relative ℓ_p errors of the computed eigenmodes are,

$$\|\Delta \mathbf{E}\|_p / \|\mathbf{E}\|_p = \frac{\left(\sum_i \sum_{\alpha=1}^{3} \left| E_\alpha(\mathbf{x}_i) - \sum_m c_m E'_{m,\alpha}(\mathbf{x}_i) \right|^p \right)^{1/p}}{\left(\sum_i \sum_{\alpha=1}^{3} \left| E_\alpha(\mathbf{x}_i) \right|^p \right)^{1/p}} \tag{5.55}$$

where the $\mathbf{x}_i$ are some set of test points, α is the vector component, $E_\alpha(\mathbf{x})$ is a component of an analytic eigenmode evaluated at $\mathbf{x}$, and $E'_{m,\alpha}(\mathbf{x})$ is a component of the mth computed eigenmode interpolated to the point $\mathbf{x}$. For each analytic eigenmode, we chose c_m to minimize the ℓ_2 error of the bulk electric fields, where the bulk is defined by the volume

$$V_{\text{bulk}} = \{\mathbf{x} : 0 < |\mathbf{x}| < (a - 3\Delta x),\ (a + 3\Delta x) < |\mathbf{x}| < (b - 3\Delta x)\} \tag{5.56}$$

(as before, $a = 0.37$m is the dielectric radius and $b = 0.49$m is the metal radius). For $l_{2,\text{bulk}}$, then, the test points $\mathbf{x}_i$ were all contained in V_{bulk} (we used approximately 10^5 bulk test points,

equally spaced, for each eigenmode and resolution). The c_m were found by a linear least squares minimization.

The c_m that minimize the bulk errors were then used to calculate the error of the surface electric fields for the ℓ_1, ℓ_2, and ℓ_∞ norms. Figure 5.7 shows the convergence for electric fields evaluated on two shells a distance $\pm 3\Delta x$ from the boundary (note that these shells get physically closer to the boundary as the resolution increases). The same (approximately 5000) surface test points were used at each resolution. The nearly second-order convergence of the surface fields is unexpected since our dielectric update incurs first-order error on dielectric boundaries. It is possible that first-order behavior occurs at resolutions higher than those shown here.

5.4.3 Convergence of resonant frequencies for photonic crystal of anisotropic dielectric ellipsoids

The photonic crystal was a simple orthorhombic structure with lattice vectors $\mathbf{a}_1 = (1.2, 0, 0)$m, $\mathbf{a}_2 = (0, 1.5, 0)$m, and $\mathbf{a}_3 = (0, 0, 1.8)$m. The computational domain was the unit cell fitted with periodic boundary conditions. For these simulations, $N_x = N_y = N_z$, and so the grid cell length ratios were again $\Delta y/\Delta x = 5/4$ and $\Delta z/\Delta x = 6/4$. The center of the ellipsoid was again offset from any grid nodes. The semimajor axes of the ellipsoid were $a_{\text{ell}} = 0.45$m, $b_{\text{ell}} = 0.60$m, and $c_{\text{ell}} = 0.75$m. The ellipsoid axes a_{ell}, b_{ell}, and c_{ell} were first aligned with $\mathbf{a}_1$, $\mathbf{a}_2$, and $\mathbf{a}_3$, respectively, then the ellipsoid was rotated first about the x-axis by $\pi/8$ then about the y-axis by $\pi/9$ and finally about the z-axis by $\pi/10$. The tensor dielectric inside the ellipsoids was,

$$\varepsilon_{\text{in}} = \mathbb{R} \cdot \begin{pmatrix} 8\varepsilon_0 & & \\ & 10\varepsilon_0 & \\ & & 12\varepsilon_0 \end{pmatrix} \cdot \mathbb{R}^{\text{T}} \tag{5.57}$$

where $\mathbb{R} = \mathbb{R}_z(-\pi/6) \cdot \mathbb{R}_y(-\pi/5) \cdot \mathbb{R}_x(-\pi/4)$ and $\mathbb{R}_\alpha(\theta)$ is the transformation matrix for a rotation about the α-axis by an angle θ in the right-handed sense. The dielectric everywhere else was set to vacuum permittivity, ε_0.

Figure 5.8 shows the second-order convergence of our algorithm in this more general case.

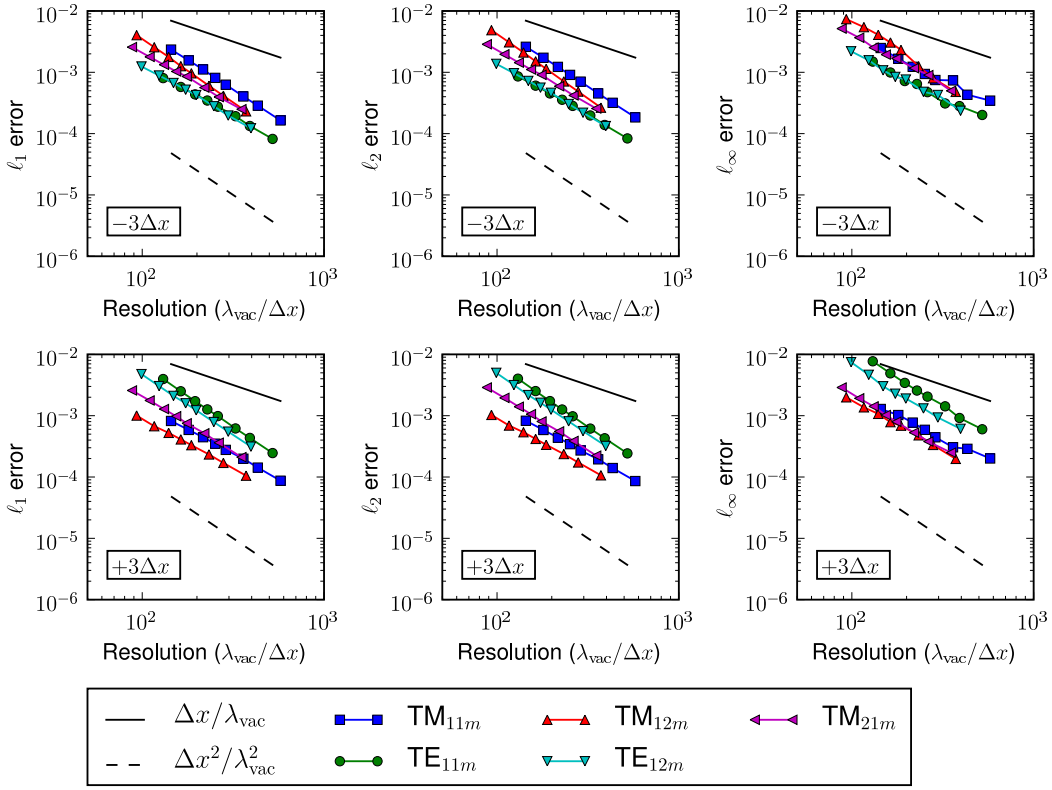

Figure 5.7: *Color online. Convergence of surface electric fields at a distance of* $\pm 3\Delta x$ *from the dielectric boundary for the concentric spheres problem.* $\lambda_{\rm vac}$ *is the wavelength of each mode in vacuum* ($\lambda_{\rm vac} = c/f$ *where* f *is the eigenfrequency).*

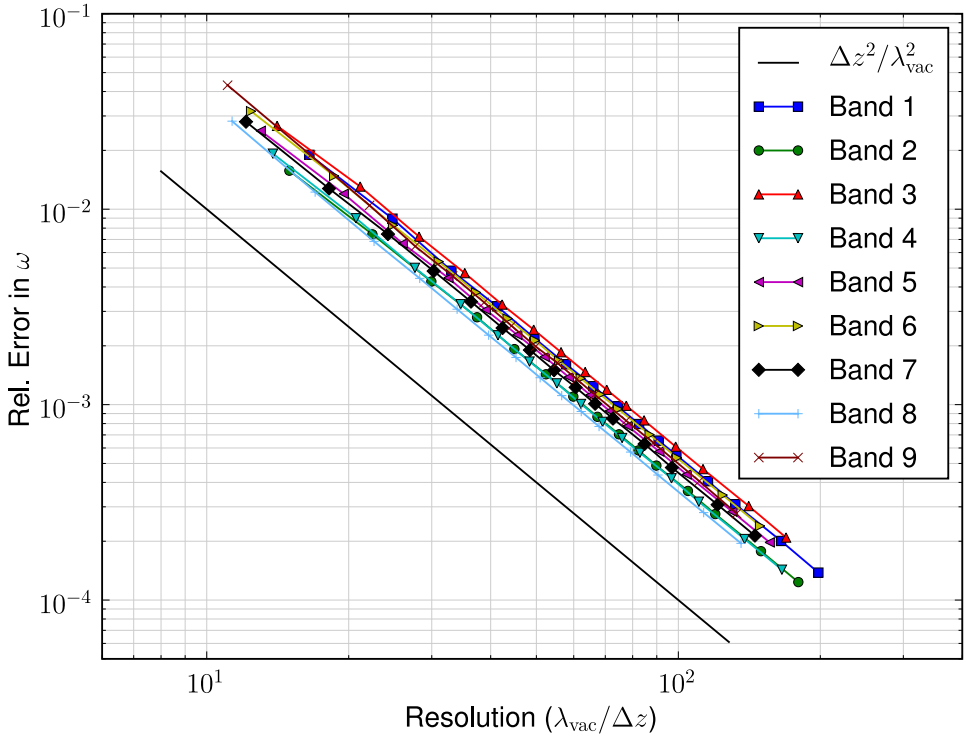

Figure 5.8: *Color online. Convergence of the lowest 9 band frequencies for the photonic crystal of ellipsoids with anisotropic dielectric constant.* λ_{vac} *is the wavelength of each mode in vacuum* ($\lambda_{\mathrm{vac}} = c/f$ *where* f *is the eigenfrequency).*

Band	Frequency (units of $c/\|\mathbf{a}_1\|$)
1	0.32300245
2	0.35626134
3	0.37826015
4	0.38601562
5	0.40978571
6	0.42944672
7	0.44031804
8	0.47125309
9	0.48171559

Table 5.1: *Normalized reference frequencies used for the photonic crystal convergence plot in Fig. 5.8. These are the Richardson extrapolated frequencies from simulations at resolutions of $N = 96$ and $N = 128$ and should be accurate up to the 6th decimal place. (c is the speed of light in vacuum).*

Because no analytic solution exists for this problem, the reference frequencies were calculated by Richardson extrapolating our own results at $N = 96$ and $N = 128$, which, based on the isotropic results shown above, would give relative errors of $\sim 10^{-6}$ (these frequencies are given in Table 5.1).

5.5 Conclusions and future work

We have described in this paper an algorithm for simulating discontinuous anisotropic dielectrics with second-order error in resonant frequencies. The algorithm is local, can be used everywhere in a simulation, and reduces to the second-order error algorithms ([87] and [79]) away from the boundaries. The algorithm relies on the extra order of accuracy gained by restricting a local update that is less accurate by one order to a $(d-1)$-dimensional surface in the d-dimensional simulation [29]. While we did not prove that this phenomenon is always true, the second-order convergence shown in our results convincingly supports it. The main advance of this work is the systematic approach to accuracy. Using the FIT, the important errors were condensed into a single step: the discrete permittivity relation. Then, to guarantee at most first-order error on dielectric boundaries, the discrete permittivity operator was forced to be exact in the limit of infinite wavelength; previous finite-difference algorithms have failed to guarantee this first-order error.

Dielectric corners/edges were not discussed in this paper, but are an important aspect of

photonics computation and present a significant challenge to finite-difference algorithms. Because our algorithm is designed to simulate finite field discontinuities, not singularities (electric fields diverge at dielectric corners), the results of simulating dielectric corners with our algorithm would likely suffer the same inaccuracies as detailed in Ref. [22]. However, the principles of Section 5.3.1 (specifically Diagram 5.11), should provide the necessary framework for attacking these inaccuracies. First, the asymptotic forms of the fields at the corner need to be determined. Second, a discretization scheme is chosen, e.g. the FIT. Finally, Eq. (5.13) is used to determine the dielectric corner update based on an approximate inverse of the discretization operator (e.g. $\tilde{\mathcal{A}}$ for the FIT); the approximate inverse is required to give the exact asymptotic form for the electric field at the corner in the limit $\Delta x \to 0$. We leave the details to further investigations.

For stability of a time-domain scheme, the matrix $\mathbb{C}^{\mathrm{T}}\mathbb{M}_{\varepsilon}^{-1}\mathbb{C}\mathbb{M}_{\mu}^{-1}$ must have all real eigenvalues and be positive semi-definite. Since $\mathbb{M}_{\mu} = \mu_0\mathbb{I}$, a symmetric $\mathbb{M}_{\varepsilon}^{-1}$ would ensure real eigenvalues. Unfortunately, the proposed algorithm does not guarantee symmetry of $\mathbb{M}_{\varepsilon}^{-1}$. Indeed, we have found complex eigenfrequencies upon solving Eq. (5.54); they occur mostly in the high-frequency spectrum (complex eigenfrequencies may also occur lower in the spectrum if degenerate or very nearly degenerate modes exist, but the imaginary part of these eigenfrequencies tends to be insignificant). As a consequence, this algorithm cannot be used in the time domain with stability guaranteed. Nevertheless, the algorithm is still very useful as an eigensolver for finding accurate, well-resolved modes. Work is currently underway to reformulate the method for time-domain stability without sacrificing our robust treatment of accuracy.

Acknowledgments

This work was supported by the U.S. Department of Energy grant DE-FG02-04ER41317.

Bibliography

[1] Y. Akahane, T. Asano, B.S. Song, and S. Noda. High-q photonic nanocavity in a two-dimensional photonic crystal. Nature, 425(6961):944–947, 2003.

[2] N. W. Ashcroft and D. N. Mermin. Solid State Physics. Brooks Cole, 1976.

[3] K. L. Bane and Robert L. Gluckstern. The Transverse wake field of a detuned x band accelerator structure. Part. Accel., 42:123–169, 1993.

[4] C.A. Bauer, G.R. Werner, and J.R. Cary. Truncated photonic crystal cavities with optimized mode confinement. Journal of Applied Physics, 104(5):053107–053107, 2008.

[5] J.P. Berenger. A perfectly matched layer for the absorption of electromagnetic waves. Journal of computational physics, 114(2):185–200, 1994.

[6] HA Bethe. Theory of diffraction by small holes. Physical Review, 66(7-8):163–182, 1944.

[7] N. Bloembergen. Laser-induced electric breakdown in solids. Quantum Electronics, IEEE Journal of, 10(3):375 – 386, March 1974.

[8] W. Cai and S. Deng. An upwinding embedded boundary method for Maxwells equations in media with material interfaces: 2D case. Journal of Computational Physics, 190(1):159–183, 2003.

[9] A. Chao. Handbook of accelerator physics and engineering. World Scientific Pub Co Inc, 1999.

[10] M. Clemens and T. Weiland. Discrete electromagnetism with the finite integration technique. Progress in Electromagnetics Research, PIER, 32:65–87, 2001.

[11] E.R. Colby. EM Structure Based and Vacuum Acceleration. In AIP Conference Proceedings, pages 39–46. IOP INSTITUTE OF PHYSICS PUBLISHING LTD, 2002.

[12] R.E. Collin. Field theory of guided waves. McGraw-Hill, 1960.

[13] B.M. Cowan. Three-dimensional dielectric photonic crystal structures for laser-driven acceleration. Physical Review Special Topics-Accelerators and Beams, 11(1):011301, 2008.

[14] S. Dawson. The physics case for a tev lepton collider. ALCPG Workshop, Eugene, Oregon, March 2011.

[15] S. Dey and R. Mittra. A locally conformal finite-difference time-domain (FDTD) algorithm for modeling three-dimensional perfectly conducting objects. Microwave and Guided Wave Letters, IEEE [see also IEEE Microwave and Wireless Components Letters], 7(9):273–275, 1997.

[16] S. Dey and R. Mittra. A conformal finite-difference time-domain technique for modeling cylindrical dielectric resonators. IEEE Transactions on Microwave Theory and Techniques, 47(9 Part 1):1737–1739, 1999.

[17] E. Di Gennaro, S. Savo, A. Andreone, V. Galdi, G. Castaldi, V. Pierro, and M.R. Masullo. Mode confinement in photonic quasicrystal point-defect cavities for particle accelerators. Applied Physics Letters, 93:164102, 2008.

[18] A. Ditkowski, K. Dridi, and J. S. Hesthaven. Convergent Cartesian grid methods for Maxwell's equations in complex geometries. Journal of Computational Physics, 170(1):39–80, 2001.

[19] V. Dolgashev, S. Tantawi, Y. Higashi, and B. Spataro. Geometric dependence of radio-frequency breakdown in normal conducting accelerating structures. Applied Physics Letters, 97:171501, 2010.

[20] KH Dridi, J.S. Hesthaven, and A. Ditkowski. Staircase-free finite-difference time-domain formulation for general materials in complex geometries. Antennas and Propagation, IEEE Transactions on, 49(5):749–756, 2002.

[21] D. Du, X. Liu, G. Korn, J. Squier, and G. Mourou. Laser-induced breakdown by impact ionization in SiO2 with pulse widths from 7 ns to 150 fs. Applied Physics Letters, 64:3071, 1994.

[22] A. Farjadpour, David Roundy, Alejandro Rodriguez, M. Ibanescu, Peter Bermel, J. D. Joannopoulos, Steven G. Johnson, and G. W. Burr. Improving accuracy by subpixel smoothing in the finite-difference time domain. Opt. Lett., 31(20):2972–2974, 2006.

[23] R.H. Fowler and L.W. Nordheim. Electron emission in intense electric fields. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, 119(781):173–181, 1928.

[24] J. Gao. Analytical formulas for the resonant frequency changes due to opening apertures on cavity walls. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 311(3):437–443, 1992.

[25] E.D. Gennaro, C. Zannini, S. Savo, A. Andreone, MR Masullo, G. Castaldi, I. Gallina, and V. Galdi. Hybrid photonic-bandgap accelerating cavities. New Journal of Physics, 11:113022, 2009.

[26] A. Grudiev. Update on structure optimization procedure, input and results. clic reference structure. CLIC-ACE, 2008.

[27] A. Grudiev. Possible improvement of the clic accelerating structure. CLIC-ACE, 2009.

[28] A. Grudiev. Clic baseline design. private communication, 2012.

[29] B. Gustafsson. The convergence rate for difference approximations to mixed initial boundary value problems. Mathematics of Computation, pages 396–406, 1975.

[30] S.C. Hagness and A. Taflove. Computational electrodynamics: The finite-difference time-domain method. Artech House Boston, 2005.

[31] H. Henke. Coupling impedances of a set of resonators. proceedings of EPAC88, page 711, 1988.

[32] Michael A. Heroux, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu, Tamara G. Kolda, Richard B. Lehoucq, Kevin R. Long, Roger P. Pawlowski, Eric T. Phipps, Andrew G. Salinger, Heidi K. Thornquist, Ray S. Tuminaro, James M. Willenbring, Alan Williams, and Kendall S. Stanley. An overview of the trilinos project. ACM Trans. Math. Softw., 31(3):397–423, 2005.

[33] T. Hirono, Y. Shibata, WW Lui, S. Seki, and Y. Yoshikuni. The second-order condition for the dielectric interface orthogonal to the Yee-lattice axis in the FDTD scheme. Microwave and Guided Wave Letters, IEEE, 10(9):359–361, 2002.

[34] T. Hirono, Y. Yoshikuni, and T. Yamanaka. Effective permittivities with exact second-order accuracy at inclined dielectric interface for the two-dimensional finite-difference time-domain method. Applied optics, 49(7):1080–1096, 2010.

[35] R. Holland. Pitfalls of staircase meshing. IEEE Transactions on Electromagnetic Compatibility, 35(4):434–439, 1993.

[36] K.P. Hwang and AC Cangellaris. Effective permittivities for second-order accurate FDTD equations at dielectric interfaces. IEEE Microwave and Wireless Components Letters, 11(4):158–160, 2001.

[37] J. D. Jackson. Classical Electrodynamics, 3rd Edition. John Wiley and Sons, Inc., 3 edition, 1999.

[38] C. Jing, W. Gai, J.G. Power, R. Konecny, W. Liu, S.H. Gold, A.K. Kinkead, S.G. Tantawi, V. Dolgashev, and A. Kanareykin. Progress toward externally powered x-band dielectric-loaded accelerating structures. Plasma Science, IEEE Transactions on, 38(6):1354–1360, 2010.

[39] 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.

[40] Steven G. Johnson and J. D. Joannopoulos. Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis. Opt. Express, 8:173, 2001.

[41] T.G. Jurgens, A. Taflove, K. Umashankar, and TG Moore. Finite-difference time-domain modeling of curved surfaces. Antennas and Propagation, IEEE Transactions on, 40(4):357–366, 1992.

[42] N. Kaneda, B. Houshmand, and T. Itoh. FDTD analysis of dielectric resonators with curved surfaces. IEEE Transactions on Microwave Theory and Techniques, 45(9):1645–1649, 1997.

[43] Vasim Khan. A damped and detuned accelerating structure for the main linacs of the compact linear collider. PhD thesis, University of Manchester, 2010.

[44] C. Kittel and P. McEuen. Introduction to Solid State Physics, volume 4. Wiley New York, 1986.

[45] Chris Kottke, Ardavan Farjadpour, and Steven G. Johnson. Perturbation theory for anisotropic dielectric interfaces, and application to subpixel smoothing of discretized numerical methods. Phys. Rev. E, 77(3):036611, 2008.

[46] N. Kroll, D. R. Smith, and S. Schultz. Photonic band gap structures: A new approach to accelerator cavities. In Advanced Accelerator Concepts, volume 279 AIP Conf. Proc., pages 197–211, 1993.

[47] N.M. Kroll and X. Lin. Persistent wakefields associated with waveguide damping of higher order modes. In Particle Accelerator Conference, 1993., Proceedings of the 1993, pages 3453–3455. IEEE, 1993.

[48] J. Krupka, K. Derzakowski, A. Abramowicz, M.E. Tobar, and R.G. Geyer. Use of whispering-gallery modes for complex permittivity determinations of ultra-low-loss dielectric materials. Microwave Theory and Techniques, IEEE Transactions on, 47(6):752–759, 2002.

[49] J.Y. Lee and N.H. Myung. Locally tensor conformal FDTD method for modeling arbitrary dielectric surfaces. Microwave and Optical Technology Letters, 23(4):245–249, 1999.

[50] D. Li, N. Kroll, DR Smith, and S. Schultz. Wake-field studies on photonic band gap accelerator cavities. In AIP Conference Proceedings, volume 398, page 528, 1997.

[51] X.E. Lin. Photonic band gap fiber accelerator. Physical Review Special Topics-Accelerators and Beams, 4(5):051301, 2001.

[52] X.E. Lin and N.M. Kroll. Minimum wakefield achievable by waveguide damped cavity. In Particle Accelerator Conference, 1995., Proceedings of the 1995, volume 3, pages 1809–1811. IEEE, 1995.

[53] M. Luong, I. Wilson, and W. Wuensch. Rf loads for the clic multibunch structure. In Particle Accelerator Conference, 1999. Proceedings of the 1999, volume 2, pages 821–823. IEEE, 1999.

[54] Roark A. Marsh. Experimental study of photonic band gap accelerator structures. PhD thesis, Massachusetts Institute of Technology, 2009.

[55] Roark A. Marsh, Michael A. Shapiro, Richard J. Temkin, Valery A. Dolgashev, Lisa L. Laurent, James R. Lewandowski, A. Dian Yeremian, and Sami G. Tantawi. x-band photonic band-gap accelerator structure breakdown experiment. Phys. Rev. ST Accel. Beams, 14(2):021301, Feb 2011.

[56] M. R. Masullo, A. Andreone, E. Di Gennaro, S. Albanese, F. Francomacaro, M. Panniello, V. G. Vaccaro, and G. Lamura. Study of hybrid photonic band gap resonators for particle accelerators. Microwave and Opt. Tech. Lett., 48(12):2486–2491, December 2006.

[57] A. Mohammadi, H. Nadgaran, and M. Agio. Contour-path effective permittivities for the two-dimensional finite-difference time-domain method. Optics Express, 13(25):10367–10381, 2005.

[58] J. Nadobny, D. Sullivan, W. Wlodarczyk, P. Deuflhard, P. Wust, C. C. V. Klinikum, and K. fur Strahlenheilkunde. A 3-D tensor FDTD-formulation for treatment of sloped interfaces in electrically inhomogeneous media. IEEE Transactions on Antennas and Propagation, 51(8):1760–1770, 2003.

[59] O. Napoly et al. A generalized method for calculating wake potentials. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 334(2-3):255–265, 1993.

[60] C. Nieter and J. R. Cary. VORPAL: a versatile plasma simulation code. J. Comput. Phys., 196(2):448, 2004.

[61] C. Nieter, J. R. Cary, G. R. Werner, D. N. Smithe, and P. H. Stoltz. Application of Dey–Mittra conformal boundary algorithm to 3D electromagnetic modeling. Journal of Computational Physics, 228:7902–7916, 2009.

[62] Susumu Noda, Alongkarn Chutinan, and Masahiro Imada. Trapping and emission of photons by a single defect in a photonic bandgap structure. Nature, 407(6804):608–610, October 2000.

[63] A. F. Oskooi, C. Kottke, and S. G. Johnson. Accurate finite-difference time-domain simulation of anisotropic media by subpixel smoothing. Optics Letters, 34(18):2778–2780, 2009.

[64] H. Padamsee, J. Knobloch, and T. Hays. RF Superconductivity for Accelerators. Wiley-VCH, 2 edition, 2008.

[65] O. Painter, R. K. Lee, A. Scherer, A. Yariv, J. D. O'Brien, P. D. Dapkus, and I. Kim. Two-dimensional photonic band-gap defect mode laser. Science, 284(5421):1819–1821, June 1999.

[66] JG Power, W. Gai, SH Gold, AK Kinkead, R. Konecny, C. Jing, W. Liu, and Z. Yusof. Observation of multipactor in an alumina-based dielectric-loaded accelerating structure. Physical review letters, 92(16):164801, 2004.

[67] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, et al. Numerical recipes, volume 547. Cambridge Univ Press, 1986.

[68] David P. Pritzkau. Rf pulsed heating. 2001. SLAC-R-577.

[69] David P. Pritzkau and Robert H. Siemann. Experimental study of rf pulsed heating on oxygen free electronic copper. Phys. Rev. ST Accel. Beams, 5(11):112002, Nov 2002.

[70] W. Sha, Z. Huang, X. Wu, and M. Chen. Application of the symplectic finite-difference time-domain scheme to electromagnetic simulation. Journal of Computational Physics, 225(1):33–50, 2007.

[71] J.C. Slater. Microwave electronics. New York, 1950.

[72] Evgenya I. Smirnova, Amit S. Kesar, Ivan Mastovsky, Michael A. Shapiro, and Richard J. Temkin. Demonstration of a 17-ghz, high-gradient accelerator with a photonic-band-gap structure. Phys. Rev. Lett., 95(7):074801, Aug 2005.

[73] D. R. Smith, Derun Li, D. C. Vier, N. Kroll, and S. Schultz. Recent progress on photonic band gap accelerator cavities. In Advanced Accelerator Concepts, volume 398 AIP Conf. Proc., pages 518–527, 1997.

[74] MC Thompson, H. Badakov, AM Cook, JB Rosenzweig, R. Tikhoplav, G. Travish, I. Blumenfeld, MJ Hogan, R. Ischebeck, N. Kirby, et al. Breakdown Limits on Gigavolt-per-Meter Electron-Beam-Driven Wakefields in Dielectric Structures. Physical review letters, 100(21):214801, 2008.

[75] An-Chun Tien, Sterling Backus, Henry Kapteyn, Margaret Murnane, and Gérard Mourou. Short-pulse laser damage in transparent materials as a function of pulse duration. Phys. Rev. Lett., 82(19):3883–3886, May 1999.

[76] J. W. Wang and G. A. Loew. Field emission and RF breakdown in high-gradient room-temperature linac structures. Talk given at Joint CERN-US-Japan Accelerator School: Course on Frontiers of Accelerator Technology: RF Engineering for Particle Accelerators, Tsukuba, Japan, 9-18 Sep 1996.

[77] T.P. Wangler. RF Linear accelerators. Wiley Online Library, 2008.

[78] T. Weiland. A discretization model for the solution of Maxwell's equations for six-component fields. Archiv fuer Elektronik und Uebertragungstechnik, 31:116–120, 1977.

[79] G. R. Werner and J. R. Cary. A stable FDTD algorithm for non-diagonal, anisotropic dielectrics. Journal of Computational Physics, 226(1):1085–1101, 2007.

[80] G.R. Werner. Analytical Wake Potentials in a Closed Pillbox Cavity. Arxiv preprint arXiv:0906.1007, 2009.

[81] Gregory R. Werner, Carl A. Bauer, and John R. Cary. Wakefields in photonic crystal cavities. Phys. Rev. ST Accel. Beams, 12:071301, Jul 2009.

[82] Gregory R. Werner and John R. Cary. Extracting degenerate modes and frequencies from time-domain simulations with filter-diagonalization. J. Comput. Phys., 227:5200–5214, 2008.

[83] K. Wille. The physics of particle accelerators: an introduction. Oxford Univ. Press, 2000.

[84] P.B. Wilson. Introduction to wakefields and wake potentials. In AIP Conference Proceedings, volume 184, page 525, 1989.

[85] Z. Xie, C. H. Chan, and B. Zhang. An explicit fourth-order staggered finite-difference time-domain method for Maxwell's equations. Journal of Computational and Applied Mathematics, 147(1):75–98, 2002.

[86] E. Yablonovitch. Inhibited spontaneous emission in solid-state physics and electronics. Physical review letters, 58(20):2059–2062, 1987.

[87] K. Yee. Numerical solution of inital boundary value problems involving maxwell's equations in isotropic media. IEEE Transactions on antennas and propagation, 14(3):302–307, 1966.

[88] A. Yefet and P. G. Petropoulos. A Staggered Fourth-Order Accurate Explicit Finite Difference Scheme for the Time-Domain Maxwell's Equations. Journal of Computational Physics, 168:286–315, 2001.

[89] W. Yu and R. Mittra. A conformal finite difference time domain technique for modeling curved dielectric surfaces. IEEE Microwave and Wireless Components Letters, 11(1):25–27, 2001.

[90] S. Zhao and GW Wei. High-order FDTD methods via derivative matching for Maxwell's equations with material interfaces. Journal of Computational Physics, 200(1):60–103, 2004.

Appendix A

Iris perturbation theory

Slater perturbation theory [71, 24, 77] relates the change in resonant frequency to the change in electromagnetic energy; specifically

$$\frac{\tilde{\omega}-\omega}{\omega} = \frac{1}{U}\left(\Delta U_B - \Delta U_E\right) \tag{A.1}$$

where $\tilde{\omega}$ is the perturbed frequency, ω is the unperturbed frequency, U is the stored energy of the unperturbed mode, and ΔU_B and ΔU_E are the changes in the time-averaged magnetic and electric energies. The quantities ω and U are known; the energy differences are calculated as follows using the theory of Bethe [6].

Bethe's theory says that the perturbing effects of apertures (with dimensions smaller than the wavelength in question) in cavity walls can be approximated by placing electric and magnetic dipoles at the centers of these apertures [6, 12, 24], where the strengths and directions of the dipoles are determined by the unperturbed cavity fields. For the special case of the TM_{010} mode in the periodic pillbox of Fig. 2.9, this technique amounts to placing an electric dipole at each end of the iris equal to [12]

$$\mathbf{P}^{(1,2)} = -\frac{2}{3}a^3\varepsilon_0\mathbf{E}^{(1,2)}(\rho=0) \tag{A.2}$$

where the superscripts 1 and 2 refer to the cavities on the left and right of the iris, respectively, and we let $\mathbf{E}^{(1,2)}(\rho=0) = E_0^{(1,2)}\hat{z}$ be the on-axis unperturbed electric field in each cavity.

The perturbative effect of the dipoles is determined by their interaction with the field differences. To understand this, consider an infinitesimally thin conducting wall dividing two identical

cavities. If an aperture is opened in the divider, the symmetric coupled mode satisfies the same boundary conditions enforced by the aperture-less divider; thus, the resonant frequency is unchanged. The antisymmetric mode, however, will have a shifted frequency since the fields must reverse sign (without discontinuity) through the aperture. The time-averaged electric energy change is then given by [24]

$$\Delta U_E = \frac{1}{4}\left(\mathbf{P}^{(1)} - \mathbf{P}^{(2)}\right) \cdot \left(\mathbf{E}^{(1)} - \mathbf{E}^{(2)}\right) \tag{A.3}$$

$$= -\frac{1}{3} a^3 \varepsilon_0 E_0^2 \left(1 - e^{-\alpha d} \cos\phi\right) \tag{A.4}$$

where the exponential factor comes from the cross terms since the fields decay exponentially through the iris (the iris can be considered as a waveguide with a cutoff frequency; the TM_{010} frequency falls well below this cutoff) and the cosine factor is due to the phase advance between cavities. The decay factor in the exponential is derived based on the cutoff frequency of the TM_{01} cylindrical waveguide mode, and is

$$\alpha = \sqrt{\frac{2.405^2}{a} - \frac{\tilde{\omega}^2}{c^2}}. \tag{A.5}$$

Since the on-axis magnetic field for the unperturbed TM_{010} mode is zero, we have $\Delta U_B = 0$.

The frequency perturbation is now

$$\frac{\tilde{\omega} - \omega}{\omega} = \frac{1}{3} \frac{a^3 \varepsilon_0 E_0^2}{U} \left(1 - e^{-\alpha d} \cos\phi\right). \tag{A.6}$$

The goal is to predict ω given a target frequency $\tilde{\omega}$ (in our case, this is usually $\tilde{\omega}/2\pi = 12$ GHz). For the unperturbed TM_{010} mode, the energy scales as $U \sim \lambda^2 L = c^2 L/\omega^2$. Thus, we rewrite the above equation as

$$\frac{\tilde{\omega} - \omega}{\omega} = \omega^2 \frac{1}{3} \frac{a^3}{c^2 L} \Gamma \left(1 - e^{-\alpha d} \cos\phi\right). \tag{A.7}$$

where $\Gamma = \frac{c^2 L \varepsilon_0 E_0^2}{\omega^2 U}$ is a dimensionless constant purely determined by the TM_{010} field pattern. Furthermore, since the TM_{010} fields are independent of z, Γ can be written

$$\Gamma = \frac{2 c^2 \varepsilon_0 E_0^2}{\omega^2 \int \varepsilon |\mathbf{E}|^2 \, dx \, dy}. \tag{A.8}$$

Equation A.7 is a cubic equation that can be solved for ω. The result will determine the transverse scaling of the unperturbed cavity needed to produce the target TM_{010} frequency, $\tilde{\omega}$, after the addition of the irises (Γ will be used to transversely scale our 3D photonic crystal structures based on 2D calculations).

Appendix B

Sapphire guided modes

With an anisotropic and spatially-dependent dielectric constant, Maxwell's equations read

$$\varepsilon_0 \varepsilon \cdot \partial_t \mathbf{E} = \nabla \times \mathbf{H} \tag{B.1}$$

$$-\mu_0 \partial_t \mathbf{H} = \nabla \times \mathbf{E} \tag{B.2}$$

where ε is the relative permittivity tensor. We restrict ε to have the form

$$\varepsilon = \begin{pmatrix} \varepsilon_\perp & & \\ & \varepsilon_\perp & \\ & & \varepsilon_{zz} \end{pmatrix} \tag{B.3}$$

where the scalars $\varepsilon_\perp$ and ε_{zz} vary in the x and y (transverse) directions only. In this case, the solutions take the form,

$$\mathbf{E}(\mathbf{x}, t) = \mathbf{E}(\mathbf{x}_\perp) e^{ikz - i\omega t} \tag{B.4}$$

$$\mathbf{H}(\mathbf{x}, t) = \mathbf{H}(\mathbf{x}_\perp) e^{ikz - i\omega t} \tag{B.5}$$

Assuming the above forms and splitting each field into its z and transverse component, Ampere's equation becomes

$$-i\omega\varepsilon_0 \varepsilon_{zz} E_z \hat{z} = \nabla_\perp \times \mathbf{H}_\perp \tag{B.6}$$

$$-i\omega\varepsilon_0 \varepsilon_\perp \mathbf{E}_\perp = ik\hat{z} \times \mathbf{H}_\perp - \hat{z} \times \nabla_\perp H_z \tag{B.7}$$

and Faraday's equation becomes

$$i\omega\mu_0 H_z\hat{z} = \nabla_\perp \times \mathbf{E}_\perp \tag{B.8}$$

$$i\omega\mu_0 \mathbf{H}_\perp = ik\hat{z} \times \mathbf{E}_\perp - \hat{z} \times \nabla_\perp E_z \tag{B.9}$$

where E, H_z, $\mathbf{E}, \mathbf{H}_\perp$ depend on only $x_\perp$.

Applying the transverse curl to B.7 and using $\nabla \cdot \mathbf{H} = 0$ gives

$$-i\omega\varepsilon_0 \nabla_\perp \times \mathbf{E}_\perp = ik\nabla_\perp \times \frac{1}{\varepsilon_\perp}\hat{z} \times \mathbf{H}_\perp \tag{B.10}$$

$$= \left(\nabla_\perp \frac{1}{\varepsilon_\perp}\right) \cdot (ik\mathbf{H}_\perp - \nabla_\perp H_z) - \frac{1}{\varepsilon_\perp}(\nabla_\perp^2 - k^2)H_z \tag{B.11}$$

Substituion into B.8 yields

$$\left(\nabla_\perp^2 + \frac{\omega^2}{c^2}\varepsilon_\perp - k^2\right) H_z = \varepsilon_\perp \left(\nabla_\perp \frac{1}{\varepsilon_\perp}\right) \cdot (ik\mathbf{H}_\perp - \nabla_\perp H_z) \tag{B.12}$$

The same is done for E_z, only that the divergence condition is slightly complicated by the spatially-varying dielectric; it reads

$$\nabla_\perp \cdot \mathbf{E}_\perp = -ik\frac{\varepsilon_{zz}}{\varepsilon_\perp}E_z - \frac{1}{\varepsilon_\perp}(\nabla_\perp \varepsilon_\perp) \cdot \mathbf{E}_\perp. \tag{B.13}$$

The analogue to B.12 for E_z using the above is then

$$\left(\nabla_\perp^2 + \frac{\omega^2}{c^2}\varepsilon_{zz} - k^2\frac{\varepsilon_{zz}}{\varepsilon_\perp}\right) E_z = -ik\frac{1}{\varepsilon_\perp}(\nabla_\perp \varepsilon_\perp) \cdot \mathbf{E}_\perp \tag{B.14}$$

Further substitutions can eliminate $\mathbf{E}_\perp$ and $\mathbf{H}_\perp$ in terms of both E_z and H_z; that is, the transverse components are determined entirely by the z-components. These relationships read

$$\mathbf{E}_\perp = \frac{i}{\varepsilon_\perp \omega^2/c^2 - k^2}\left(k\nabla_\perp E_z - \omega\mu_0 \hat{z} \times \nabla_\perp H_z\right) \tag{B.15}$$

$$\mathbf{H}_\perp = \frac{i}{\varepsilon_\perp \omega^2/c^2 - k^2}\left(k\nabla_\perp H_z + \omega\varepsilon_0\varepsilon_\perp \hat{z} \times \nabla_\perp E_z\right) \tag{B.16}$$

We now further restrict the problem to that of an infinite dielectric cylinder of radius a with its axis along z centered transversely at $x = y = 0$. Cylindrical coordinates are used, where ρ is

the distance from the origin and ϕ is the azimuth. The dielectric inside (outside) will be given by ε^{in} (ε^{out}). Let

$$\gamma_e^2 = \frac{\omega^2}{c^2}\varepsilon_{zz} - k^2\frac{\varepsilon_{zz}}{\varepsilon_\perp} \tag{B.17}$$

$$\gamma_h^2 = \frac{\omega^2}{c^2}\varepsilon_\perp - k^2 \tag{B.18}$$

and notice that both have the same root at $k = \varepsilon_\perp\omega/c$. The sign of $\gamma_{e,h}$ determines whether the solutions oscillate or decay in space. We are interested in modes trapped within the dielectric cylinder; thus, $\gamma_{e,h} > 0$ when $\rho < a$ and $\gamma_{e,h} < 0$ when $\rho > a$. Let $\beta_{e,h}^2 = -\gamma_{e,h}^2$ outside the cylinder.

Inside and outside the cylinder, the right-hand sides of B.12 and B.14 are zero. For guided modes, the solutions inside are oscillatory and the take the form

$$E_z = A_{e,m}J_m(\gamma_e\rho)e^{im\phi} \tag{B.19}$$

$$H_z = A_{h,m}J_m(\gamma_h\rho)e^{im\phi} \tag{B.20}$$

while the solutions outside are purely decaying (with increasing ρ):

$$E_z = B_{e,m}K_m(\beta_e\rho)e^{im\phi} \tag{B.21}$$

$$H_z = B_{h,m}K_m(\beta_h\rho)e^{im\phi}. \tag{B.22}$$

When $m = 0$, the solutions decouple into TM and TE modes; when $m > 0$, both E_z and H_z are nonzero.

Matching fields at the interface at $\rho = a$ gives an eigenvalue equation to be solved for the frequency as a function of the wave vector k. Four matching conditions are required to determine the constants $A_{e,m}$, $A_{h,m}$, $B_{e,m}$, and $B_{h,m}$. The simplest matching conditions enforce the continuity of E_z and H_z:

$$A_{e,m}J_m(\gamma_e a) = B_{e,m}K_m(\beta_e a) \tag{B.23}$$

$$A_{h,m}J_m(\gamma_h a) = B_{h,m}K_m(\beta_h a). \tag{B.24}$$

Next, we enforce the continuity of D_ρ. Using B.15, one obtains for both regions:

$$D_\rho^{\text{in}} = \frac{i\varepsilon_0\varepsilon_\perp^{\text{in}}}{\gamma_h^2}\left(A_{e,m}k\gamma_e J_m'(\gamma_e a) + A_{h,m}\frac{im\omega\mu_0}{a}J_m(\gamma_h a)\right) \tag{B.25}$$

$$D_\rho^{\text{out}} = -\frac{i\varepsilon_0\varepsilon_\perp^{\text{out}}}{\beta_h^2}\left(B_{e,m}k\beta_e K_m'(\beta_e a) + B_{h,m}\frac{im\omega\mu_0}{a}K_m(\beta_h a)\right). \tag{B.26}$$

Finally, the continuity of E_ϕ is enough to solve for the coefficients. The expressions for E_ϕ are:

$$E_\phi^{\text{in}} = \frac{i}{\gamma_h^2}\left(A_{e,m}\frac{imk}{a}J_m(\gamma_e a) - A_{h,m}\omega\mu_0\gamma_h J_m'(\gamma_h a)\right) \tag{B.27}$$

$$E_\phi^{\text{out}} = -\frac{i}{\beta_h^2}\left(B_{e,m}\frac{imk}{a}K_m(\beta_e a) - B_{h,m}\omega\mu_0\beta_h K_m'(\beta_h a)\right) \tag{B.28}$$

Combining B.23–B.28, one obtains the eigenvalue equation:

$$\left(\frac{\varepsilon_\perp^{\text{in}}x_e}{x_h^2}\frac{J_m'(x_e)}{J_m(x_e)} + \frac{\varepsilon_\perp^{\text{out}}y_e}{y_h^2}\frac{K_m'(y_e)}{K_m(y_e)}\right)\left(\frac{1}{x_h}\frac{J_m'(x_h)}{J_m(x_h)} + \frac{1}{y_h}\frac{K_m'(y_h)}{K_m(y_h)}\right) = m^2\left(\frac{\varepsilon_\perp^{\text{in}}}{x_h^2} + \frac{\varepsilon_\perp^{\text{out}}}{y_h^2}\right)\left(\frac{1}{x_h^2} + \frac{1}{y_h^2}\right) \tag{B.29}$$

where, for brevity, we have introduced

$$x_e^2 = \gamma_e^2 a^2 = \frac{\omega^2 a^2}{c^2}\varepsilon_{zz}^{\text{in}} - k^2 a^2\frac{\varepsilon_{zz}^{\text{in}}}{\varepsilon_\perp^{\text{in}}} \tag{B.30}$$

$$x_h^2 = \gamma_h^2 a^2 = \frac{\omega^2 a^2}{c^2}\varepsilon_\perp^{\text{in}} - k^2 a^2 \tag{B.31}$$

$$y_e^2 = \beta_e^2 a^2 = k^2 a^2\frac{\varepsilon_{zz}^{\text{out}}}{\varepsilon_\perp^{\text{out}}} - \frac{\omega^2 a^2}{c^2}\varepsilon_{zz}^{\text{out}} \tag{B.32}$$

$$y_h^2 = \beta_h^2 a^2 = k^2 a^2 - \frac{\omega^2 a^2}{c^2}\varepsilon_\perp^{\text{out}}. \tag{B.33}$$

The right-hand side of B.29 is always positive; therefore, the quantities in parentheses on the left-hand side are either both positive or both negative. When both are negative (positive), the modes are called HE_{mn} (EH_{mn}) and are characterized by a dominant E_z (H_z) component.

In the case of $m = 0$, E_z and H_z decouple and the eigenvalue equation is simplified greatly. For TM modes it is

$$\frac{\varepsilon_\perp^{\text{in}}x_e}{x_h^2}\frac{J_0'(x_e)}{J_0(x_e)} + \frac{\varepsilon_\perp^{\text{out}}y_e}{y_h^2}\frac{K_0'(y_e)}{K_0(y_e)} = 0. \tag{B.34}$$

Cutoff frequencies for the various modes are determined by solving the eigenvalue equation in the limit $\beta_{e,h} \to 0$, which indicates weak confinement (as $\beta_{e,h}$ goes through zero, the fields

outside the cylinder go from decaying to radiating). For example, from B.34, the lowest TM cutoff frequency is given by the smallest root of the equation

$$J_0(x_e) = 0, \tag{B.35}$$

i.e., $x_e = 2.405$. Using $y_e = 0$, the cutoff frequency is

$$\omega_{01,\text{cutoff}} = \frac{2.405c}{a\sqrt{\varepsilon_{zz}^{\text{in}} - \frac{\varepsilon_{\perp}^{\text{out}}\varepsilon_{zz}^{\text{in}}}{\varepsilon_{\perp}^{\text{in}}}}}. \tag{B.36}$$

For $m = 1$, similar analysis reveals that the HE_{11} has *no* cutoff frequency [37].

Appendix C

Prescription for algorithm

To ease the implementation of our algorithm, the following is an outline of the steps necessary to form the discrete permittivity matrix, $\mathbb{M}_{\varepsilon}^{-1}$.

For every grid node, $\mathcal{N}_{ijk}$, one needs to calculate the eight 3×3 matrices, $\Xi_{ijk}^{\pm\pm\pm}$, in order to form $\mathbb{M}_{\varepsilon}^{-1}$ according to Eqs. (5.45) and (5.53). As a representative example, consider building the matrix, Ξ_{ijk}^{+++}:

(1) Determine if any of the integration regions associated with the triplets of grid values $\mathbf{e}_{ijk}^{+++}$ and $\mathbf{d}_{ijk}^{+++}$ (see Sec. and Fig. 5.2) are cut by a dielectric boundary (in this case, these integration regions are: $L_{x|ijk}$, $L_{y|ijk}$, $L_{z|ijk}$, $\tilde{A}_{x|ijk}$, $\tilde{A}_{y|ijk}$, and $\tilde{A}_{z|ijk}$. See Sec. 5.2 and Fig. 5.1).

(2) *If* no integration regions are cut by a boundary, use the local value of ε^{-1} as Ξ_{ijk}^{+++} and begin the calculation of the next $\Xi_{ijk}^{\pm\pm\pm}$ matrix.

(3) *Else*, if any of the above integration regions are cut by a boundary, form Ξ_{ijk}^{+++} according to Eq. (5.44):

   (a) Find the local normal to the interface, $\hat{\mathbf{n}}$.

   (b) Calculate the matrices Π and Γ for both sides of the interface using $\hat{\mathbf{n}}$ and Eqs. (5.29) and (5.27) (use Cartesian coordinates: $\hat{\mathbf{x}} = [1, 0, 0]^{\mathrm{T}}$, etc.).

   (c) Calculate the averages, $\Pi_{x|ijk}$, $\Pi_{y|ijk}$, and $\Pi_{z|ijk}$ according to Eq. (5.34). (For sharp

interfaces between homogeneous dielectrics, this is simply, $f\Pi^{(1)} + (1-f)\Pi^{(2)}$ where f is the fraction of an integration region inside dielectric $\varepsilon^{(1)}$.)

(d) Construct and invert the 3×3 matrix, Π^{+++}_{ijk}, using Eq. (5.36).

(e) Calculate the averages, $\Gamma_{x|ijk}$, $\Gamma_{y|ijk}$, and $\Gamma_{z|ijk}$ according to Eq. (5.41).

(f) Construct Γ^{+++}_{ijk} according to Eq. (5.43).

(g) Multiply to get $\Xi^{+++}_{ijk} = \Gamma^{+++}_{ijk}\Pi^{+++^{-1}}_{ijk}$.