2012-12-07

Modeling Radial Bernstein Modes in a Finite-Length Axisymmetric Non-Neutral Plasma

Mark Andrew Hutchison
Brigham Young University - Provo

Modeling Radial Bernstein Modes in a Finite-Length Axisymmetric Non-Neutral Plasma

Mark Hutchison

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Bryan G. Peterson, Chair
Grant W. Hart
Ross L. Spencer

Department of Physics and Astronomy

Brigham Young University

December 2012

ABSTRACT

Modeling Radial Bernstein Modes in a Finite-Length
Axisymmetric Non-Neutral Plasma

Mark Hutchison
Department of Physics and Astronomy, BYU
Master of Science

Axisymmetric radial Bernstein modes are known to exist in non-neutral plasmas and have been studied theoretically and computationally in 1D, but detection of these modes has still proven to be difficult due to self-shielding. To help advance the work on this front we created a 2D particle-in-cell (PIC) code that simulates a non-neutral plasma in a Malmberg-Penning trap. A detailed description of the PIC code itself has been included that highlights the benefits of using an r^2–z grid and how it can be tested. The focus of the PIC simulation was to discover how best to drive and detect these modes. While it is improbable that radial Bernstein modes will be detected in long plasmas, we show that it may be a possible due to the axial nodal structure in the potential and electric field generated by confining plasmas of any finite-length. Additionally, we find that for a short plasma the strongest detection signal along the trap wall occurs at the plasma's midpoint rather than near the ends. Results show that oscillating the confinement potentials is sufficient to excite the fundamental radial Bernstein mode, but not any of the higher order modes. The higher order modes can be seen in the simulation, however, by sinusoidally driving the radial electric field. Unfortunately, the individual modes are difficult to isolate which we suspect is due to mode mixing. Finally, we report frequencies and mode shapes for the fundamental mode and the (lower) first higher order mode.

Keywords: axisymmetric, Bernstein modes, finite-length, Malmberg-Penning, non-neutral, particle-in-cell, radial

ACKNOWLEDGMENTS

I want to graciously thank my advisor, Bryan Peterson, and the other two members in my graduate committee, Grant Hart and Ross Spencer, for all of their help in planning, executing, fixing mistakes, and analyzing the results of this project. I could not have done this on my own.

I would also like to thank Jean-Pierre Moreau for making his FORTRAN subroutines available, particularly his linear banded matrix solver with LU decomposition and his Bessel's function which I have used in my code. Additionally, I would like to thank Stephan Barlow and Alan Rominger for their helpful discussions.

Contents

List of Figures

Chapter 1

Introduction

Background

Non-neutral plasmas are often studied in plasma physics because they can be confined for relatively long periods of time when compared to confinement times of neutral plasmas. A basic difficulty in neutral plasmas is the recombination between electrons and ions themselves. But even at temperatures when this is small, electrostatic attraction between electrons and ions cause them to reorient themselves in order to preserve neutrality. Because electrons are much lighter than ions and, thus, have a higher thermal velocity, they tend to diffuse through the plasma faster than the ions [1]. This imbalance in charge creates an electric field that then pulls the ions along. This of course slows the diffusion of the electrons, but the overall diffusion of the plasma is increased. Furthermore, this complicated interaction between electrons and ions inside internal and external electromagnetic fields give rise to a large number of fluid and kinetic instabilities making long confinement times extremely difficult to obtain.

By extracting one species of charge, the complexity decreases significantly and confinement is much easier. Four of the more important features of non-neutral plasmas were outlined by Peurrung *et al.* [2]. First, non-neutral plasmas can exist at any temperature because recombination

does not occur within the plasma. Second, long confinement times allow the plasma to easily reach thermal equilibrium. Third, non-neutral plasmas respond well to electric fields, which makes it possible to move or modulate the shape of the plasma. Fourth, there are already in existence many diagnostic tools for analyzing non-neutral plasmas. Fortunately, despite the complexity loss, non-neutral plasmas still have many of the same characteristics as neutral plasmas including drifts, wave propagation, and Debye shielding. Therefore, non-neutral plasmas offer an ideal environment for experimentation with and observation of plasma phenomena.

Confinement of non-neutral plasmas is often achieved using the Malmberg-Penning trap consisting of a strong axial magnetic field to provide radial confinement and a deep electric potential well for axial confinement. Because perturbations in the plasma can be achieved through a simple modulation of the end potential, the Malmberg-Penning trap is ideal for studying wave propagation and/or normal modes. Detection of these modes is done with an electrically isolated ring (hereafter referred to as the diagnostic ring) where oscillations in the induced charge are recorded and analyzed for mode frequencies in the plasma. This method is particularly good at measuring axial modes because the longitudinal compression and rarefaction of plasma under the diagnostic ring leaves a clear signature.

Of particular interest to us are the axisymmetric radial Bernstein modes (waves that propagate perpendicular to the external magnetic field). These radial modes can also be detected using the method described above, especially in short plasmas where there exists a strong coupling between the radial and axial motion; squeezing the plasma longitudinally forces the plasma to expand radially (the opposite also being true). However, as the plasma increases in length, the coupling between radial and axial motion weakens, making the radial modes more difficult to detect. An important aspect of detecting radial modes is deciding where to place the diagnostic ring. In the past we assumed that this was near the end of the plasma, but there is no *a priori* reason for this to be the case. In fact, the sensitive region(s) may well change depending on the length of the plasma.

The radial Bernstein modes are not to be confused with *azimuthal* Bernstein modes which also propagate perpendicular to the external magnetic field and are often simply referred to as "Bernstein" modes. Azimuthal Bernstein waves have been detected and studied extensively [3–8], particularly the $m_\theta = 1$ azimuthal mode, commonly called the diocotron mode, which is used in Fourier Transform Ion Cyclotron Resonance Mass Spectrometry (FTICR-MS). There have also been a number of studies done on $m_\theta = 2, 3,$ and 4. But it seems rather surprising that comparatively little has been done on the $m_\theta = 0$ breathing mode (one exception is the work done by Barlow *et al.* [9]). It is also interesting to note that almost all of these studies assume an axial wave number of $k_z = 0$.

Motivation

The study of normal modes in plasma physics is important. Knowing of and understanding all of the possible phenomena in a plasma is significant in and of itself, but normal modes have an added significance in that they carry actual information about the interior of the plasma. Due to self-shielding effects it is difficult to measure quantities of the plasma without using destructive methods. While diagnostics using destructive methods are useful and important, nondestructive methods of gathering information are usually preferred. Plasma modes offer many nondestructive as well as destructive diagnostics for plasmas [10]. Radial modes are no exception. One known application for the $m_\theta = 0$ radial mode, or breathing mode, in a short plasma is to very gently acquire accurate mass spectrum measurements using ion cloud calorimetry [2]. Limited research could be hindering other applications from coming to light.

The plasma group at Brigham Young University has already performed much of the work in simulating and developing a kinetic theory for these modes in a 1D infinite length plasma[1] and efforts have been made to create a Malmberg-Penning trap suitable for making experimental

[1] At the time of writing, this work has not yet been published.

measurements of the modes as well. This thesis is primarily concerned with performing a 2D particle-in-cell (PIC) simulation of an axisymmetric finite-length non-neutral plasma to work in conjunction with the developing experiment by investigating detection signatures of the radial Bernstein modes. Specifically we are interested in the mode frequencies, shapes, and amplitudes as well as a description of how they can be driven and the location in which they can best be measured.

Chapter 2

Methodology

2.1 Initialization

We start by generating a thermal rigid-rotor equilibrium density on an r–z grid using a successive over-relaxation code EQUILSOR. This density is then linearly interpolated onto an r^2–z grid and particles are distributed to individual cells according to the average density of the cell's four bounding grid points. From there the particles are placed at random inside their respective cells and the particles are assigned r, z and θ thermal velocities from a Maxwell-Boltzmann distribution calculated using the Box-Müller method (see Appendix A). In addition to the thermal velocity, the plasma must have an overriding rotational velocity about the center of the trap in order to maintain radial confinement. In the limit of zero density, the equilibrium rotation velocity is just the $E \times B$ drift velocity, but as the density increases the rotation velocity must be modified [8]:

$$\omega_R = \frac{\omega_c}{2}\left(1 \pm \sqrt{1-\beta}\right), \tag{2.1}$$

where ω_R is the plasma's rotational velocity, ω_c is the cyclotron frequency, and β is the ratio of the plasma density to the Brillouin density limit $n_B = (\varepsilon_0 B^2/2m)$, the maximum density that can

be trapped in a given magnetic field (see Figure 2.1).[1] We use the smaller of the two equilibrium velocities in our simulation simply because the higher frequencies are more difficult to work with experimentally.

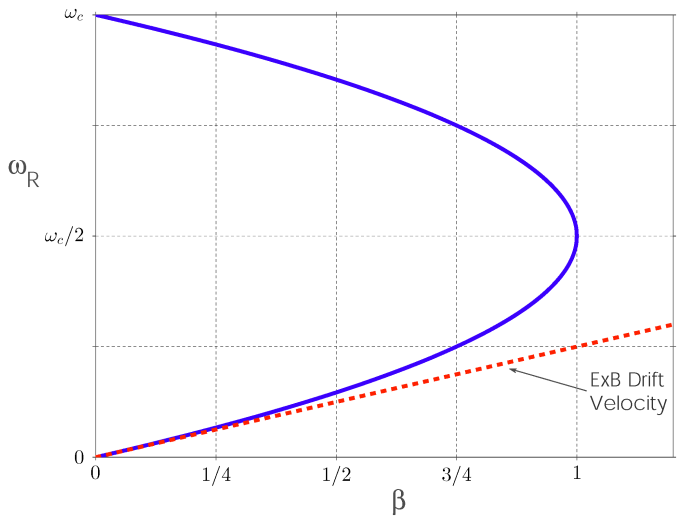

Figure 2.1. The equilibrium rotation velocity is plotted as a function of β. Above the Brillouin limit ($\beta = 1$) the magnetic field is no longer able to provide radial confinement. In the limit of $\beta \to 0$ the equilibrium rotation velocity is just the $E \times B$ drift velocity.

2.2 Particle-in-cell algorithm

Leapfrog method

To model the radial modes, we have written a 2D particle-in-cell (PIC) code that employs the leapfrog method to move the charged particles through space and time. The leapfrog method is a standard integration technique used in solving the equations of motion by centering the finite-differenced derivatives of the position and velocity at two different times separated by half a time

[1] β can also be calculated in terms of frequencies rather than densities: $\beta = 2\omega_p^2/\omega_c^2$, where $\omega_p = \sqrt{nq^2/\varepsilon_0 m}$ (the plasma frequency) and $\omega_c = qB/m$ (the cyclotron frequency).

step $\Delta t/2$. For example, suppose we have the equation of motion $\ddot{z} = a(t)$. For the leapfrog method to work, it is imperative that the function $a(t)$ not be a function of the velocity; this will be apparent below. First separating the second order differential equation into two first order differential equations, we then center the derivative in velocity at time t while the derivative in position is centered at $\Delta t/2$:

$$\begin{cases} \dot{z} = v & \longrightarrow & \dfrac{z_{t+\Delta t} - z_t}{\Delta t} = v_{t+\Delta t/2} \\ \dot{v} = a(t) & \longrightarrow & \dfrac{v_{t+\Delta t/2} - v_{t-\Delta t/2}}{\Delta t} = a_t. \end{cases} \tag{2.2}$$

We then solve for the advanced variables $z_{t+\Delta t}$ and $v_{t+\Delta t/2}$,

$$\begin{aligned} z_{t+\Delta t} &= z_t + \Delta t\, v_{t+\Delta t/2} \\ v_{t+\Delta t/2} &= v_{t-\Delta t/2} + \Delta t\, a_t, \end{aligned} \tag{2.3}$$

and use these formulas to advance the position and velocity forward in time at every time step using the previous values. Now we can see that if $a(t)$ were dependent on the velocity, we would be adding values centered at two different times which could lead to grave instabilities in the simulation.

To begin, initial conditions for both the velocity and position must be provided at $t = 0$ from which the fields are calculated. Typically the fields are then used to extrapolate the velocity back in time by half a time step to $v_{-\Delta t/2}$ to initialize the leapfrog like advancement of information in time; however, when perturbing an equilibrium this extrapolation is unnecessary since by definition an equilibrium is unchanging in time (at most the velocities would only differ in thermal noise). A cyclic process ensues in which, first, the velocity is advanced forward by one full step to $v_{\Delta t/2}$ using the fields and the acceleration equation. Second, the position is advanced forward to $z_{\Delta t}$ using the velocities and, third, the new fields are calculated based upon the new positions of the charges (see Figure 2.2). The process is repeated as many times as necessary.

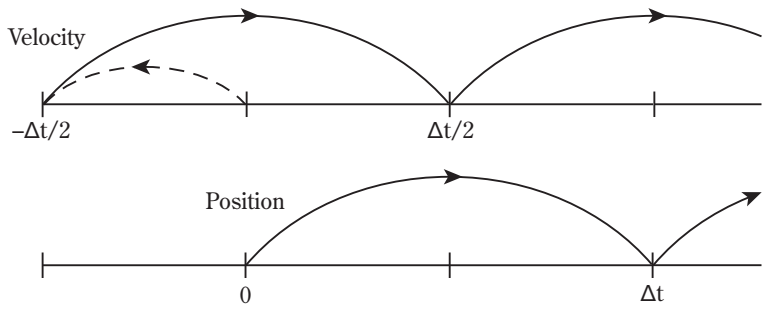

Figure 2.2. This illustrates the leapfrog like advancement of velocity and position that is used in the leapfrog method. Extrapolating the velocity back half a time step initializes the process and allows the calculations to be centered. [11]

Interpolation/extrapolation scheme

Because PIC codes use a discrete computational grid on which the density, potential, and electric field are kept, the leapfrog method described above requires an interpolation/extrapolation scheme by which one can transfer information back and forth between particles and grid. This is the mechanism that allows the field to push particles and conversely for the particles to update the field. We use a bilinear interpolation/extrapolation technique that effectively transforms the point-particles into rectangular "tents" of charge with definite shape and dimension. The tent is centered about the particle and extends across a rectangular area equivalent to four grid cells. It has a peak charge of q in the center and tapers to zero at the edges. The shape is determined by the equation

$$S = (dr - |r|)(dz - |z|), \tag{2.4}$$

where $-dr \leq r \leq dr$ and $-dz \leq z \leq dz$ (see Figure 2.3). At any given time the particle is in contact with four grid points[2] to whom it contributes a charge equal to the height of the tent at their respective locations. Mathematically this is equivalent to subdividing the cell into four rectangular areas whose vertices are defined by the charge and its four nearest grid points. Then each grid

[2]Unless it is located exactly at a grid point in which case there is only one grid point affected.

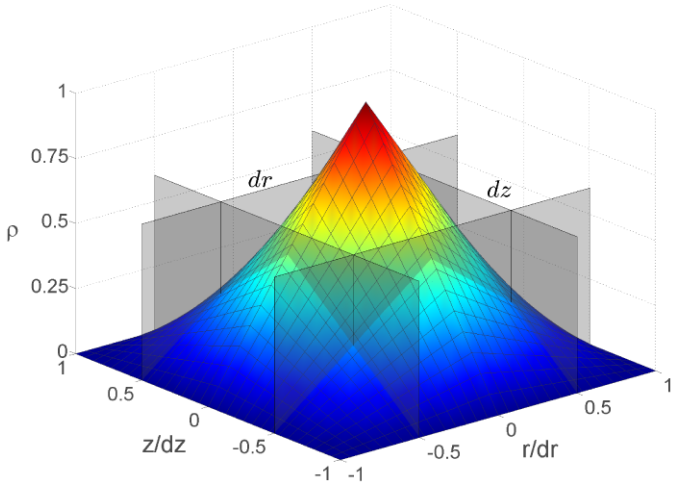

Figure 2.3. Shape of the effective charge (normalized to 1) when using the bilinear method. The computational grid is also included for perspective. Grid points receive a fraction of the total charge equal to the height of the tent at their respective locations.

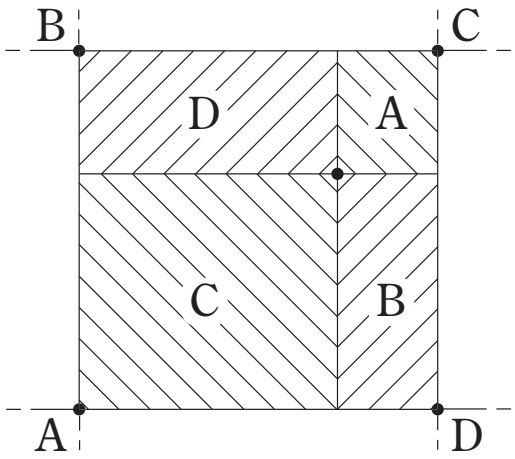

Figure 2.4. Extrapolation of the charge density to the grid points is done by splitting the cell into four rectangles and calculating the fractional area of the cell contained in each. Each grid point receives a fractional charge equal to the fractional area diagonally opposite to the grid point [11].

point is assigned a fraction of the charge equal to the fractional area in the opposite corner (see Figure 2.4).

As one might expect, this method is inadequate for modeling small-scale interactions between charges or finding trajectories of individual ions, but the large-scale physics is still very well preserved. While more accurate methods exist for extrapolating charge density to grid points, they involve higher order functions and cost a significant amount of time per cycle [11]. This bi-linear method of assigning charge is fast and efficient and is sufficient for the needs of this project.

Mover

The equations of motion used in our leapfrog algorithm are obtained using the Lagrangian formulation. In cylindrical coordinates, the Lagrangian for a particle with charge q in an axisymmetric electromagnetic field is given by

$$L = \frac{m}{2}\left(\dot{r}^2 + r^2\dot{\theta}^2 + \dot{z}^2\right) + q\mathbf{A}\cdot\mathbf{v} - q\phi(r,z), \tag{2.5}$$

where $\mathbf{A}$ is the magnetic vector potential, $\mathbf{v}$ is the velocity, and ϕ is the electric scalar potential. The magnetic potential can be obtained from the uniform external magnetic field directed along the z axis via the equation

$$B_z = [\nabla\times\mathbf{A}]_z = \frac{1}{r}\frac{\partial}{\partial r}(rA_\theta), \tag{2.6}$$

which yields

$$A_\theta = \frac{1}{2}Br. \tag{2.7}$$

Because $\mathbf{A}$ only has a $\hat{\theta}$ component, the dot product picks off the azimuthal velocity $v_\theta = r\dot{\theta}$ and the resulting Lagrangian is now

$$L = \frac{m}{2}\left(\dot{r}^2 + r^2\dot{\theta}^2 + \dot{z}^2\right) + \frac{qr^2B}{2}\dot{\theta} - q\phi(r,z). \tag{2.8}$$

The equations of motion are obtained from the Euler-Lagrange equations

$$\frac{\partial L}{\partial q_i} - \frac{d}{dt}\frac{\partial L}{\partial \dot{q}_i} = 0, \tag{2.9}$$

where q_i and $\dot{q}_i$ are the individual coordinates and coordinate velocities respectively.

Since θ does not explicitly show up in the Lagrangian it is an ignorable coordinate with a conserved conjugate momentum that can immediately be written as:

$$p_\theta = \frac{\partial L}{\partial\dot{\theta}} = mr^2\dot{\theta} + \frac{qr^2B}{2}. \tag{2.10}$$

Meanwhile, equation (2.9) dictates that the equation of motion in r is

$$\ddot{r} = r\dot{\theta}^2 + \omega_c r\dot{\theta} - \frac{q}{m}\frac{\partial\phi}{\partial r}, \tag{2.11}$$

where ω_c is the cyclotron frequency qB/m. Substitution from equation (2.10) allows us to replace $\dot{\theta}$ in favor of an expression involving the constant p_θ while the derivative of the electric potential can likewise be exchanged for the radial component of the electric field (i.e. $\partial\phi/\partial r = -E_r$). Finally, upon separating the second order differential equation into two first order differential equations we get

$$\begin{cases} \dot{r} &= v \\ \dot{v} &= \dfrac{p_\theta^2}{m^2 r^3} - \dfrac{p_\theta \omega_c}{mr} + \dfrac{\omega_c^2}{4} + \dfrac{qE_r}{m}. \end{cases} \tag{2.12}$$

The $1/r^3$ dependence creates a strong centrifugal barrier around the central axis. Physically, this is due to conservation of angular momentum which prohibits particles with nonzero angular momentum from ever reaching $r = 0$. Despite being physical, numerically a third order singularity poses a significant threat to the stability of the simulation. Figure 2.5 shows a cartoon of the radial position of a particle versus time. As the particle nears the origin, the particle makes a sharp, but smooth reversal in direction. Far from this turning point, the motion of the particle can be captured accurately by relatively large time steps. However, as the particle nears the turning point, the time step must decrease significantly in order to accurately step through the tight curves. Of course a smaller time step can be employed everywhere, but the code will run impossibly slow. Even using adaptive time steps does not gain much due to the number of particles constantly making this turnaround. The trouble near the origin can be alleviated significantly if the radial grid spacing is changed from r to r^2. Changing to an r^2 coordinate system effectively enlarges the tight curve near the origin while essentially preserving the trajectory further away. Figure 2.6 illustrates this expansion near the origin. The uniformity in the trajectory now allows us to use a large constant time step. To this end, returning to the Lagrangian in equation (2.8), we employ a change of variables by defining $x \equiv r^2$. Using this definition, and the resulting relation $\dot{r} = \dot{x}/(2\sqrt{x})$, the Lagrangian can be expressed in terms of the new coordinates x, θ, and z:

$$L = \frac{m}{2}\left(\frac{\dot{x}^2}{4x} + x\dot{\theta}^2 + \dot{z}^2\right) + \frac{qxB}{2}\dot{\theta} - q\phi. \tag{2.13}$$

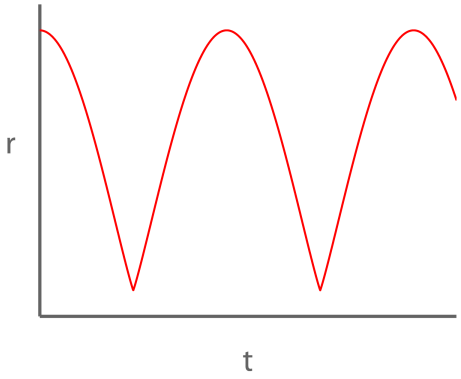

Figure 2.5. Cartoon showing the radial position vs time of a single particle in an r–z coordinate system. Near the origin, the particle makes a sudden reversal in direction. The large change in such a short period of time necessitates a small time step in order to capture the motion of the particle.

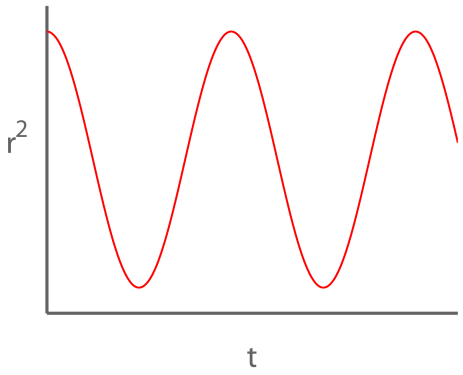

Figure 2.6. Cartoon showing the radial position vs time of a single particle in an r^2–z coordinate system. The coordinate change enlarges the turning points near the origin making the curve more uniform throughout. This then allows a uniform time step to be used everywhere in the trajectory.

Using a derivation similar to the one used above to obtain the equation of motion in r, it is straightforward to show that

$$\ddot{x} = \frac{\dot{x}^2}{2x} + \frac{2p_\theta^2}{m^2 x} - \frac{\omega_c^2 x}{2} - \frac{4qx}{m}\frac{\partial \phi}{\partial x}. \tag{2.14}$$

As before, we would like to replace the potential with the radial electric field. This can easily be done, but the new grid significantly changes the shape of the radial electric field near the origin. In r, the electric field was fortuitously linear making our bilinear interpolation to the particles more accurate, but now the electric field is proportional to $\sqrt{x}$. Thus, it would be in our best interest to preserve this linear relationship in our equations. To this end we define a new variable $G \equiv rE_r$ that preserves the linearity on the new grid and use this to replace the potential. Substituting $\partial\phi/\partial x = -G/2x$ into equation (2.14) for the potential, the set of first order differential equations

involving the new coordinate and its generalized velocity is

$$\begin{cases} \dot{x} &= v \\ \dot{v} &= \dfrac{v^2}{2x} + \dfrac{2p_\theta^2}{m^2 x} - \dfrac{\omega_c^2 x}{2} + \dfrac{2qG}{m}, \end{cases} \tag{2.15}$$

where v and $\dot{v}$ now represent $\dot{x}$ and $\ddot{x}$ respectively.

Although the centrifugal barrier from before is still present, it is much weaker (a first order pole as opposed to a third order pole). Encouraged by this result we proceed by finite differencing the new equations of motion. The acceleration equation becomes

$$\frac{v_{t+\Delta t/2} - v_{t-\Delta t/2}}{\Delta t} = \frac{1}{2x_t}\left(\frac{v_{t+\Delta t/2} + v_{t-\Delta t/2}}{2}\right)^2 + \frac{2p_\theta^2}{m^2 x_t} - \frac{\omega_c^2 x_t}{2} + \frac{2qG_t}{m}, \tag{2.16}$$

where we have been careful to average the generalized velocity term to preserve time centering. This equation is now solved for $v_{t+\Delta t/2}$ so as to provide a method for advancing the generalized velocity in time. Because the generalized velocity is quadratic, there exist two roots; the correct root is chosen by testing which root yields $v_{t-\Delta t/2}$ in the limit that $\Delta t \to 0$:

$$v_{t+\Delta t/2} = \frac{4x_t}{\Delta t} - v_{t-\Delta t/2} - 4\sqrt{\left(\frac{1}{\Delta t^2} + \frac{\omega_c^2}{4}\right)x_t^2 - \left(\frac{v_{t-\Delta t/2}}{\Delta t} + \frac{qG_t}{m}\right)x_t - \frac{p_\theta^2}{m^2}}. \tag{2.17}$$

For comparison, the equation for advancing the radial velocity in a standard r–z grid is

$$v_{t+\Delta t/2} = v_{t-\Delta t/2} + \Delta t\left(\frac{p_\theta^2}{m^2 r_t^3} - \frac{\omega_c p_\theta}{m r_t} + \frac{\omega_c^2}{4} + \frac{qE_r}{m}\right). \tag{2.18}$$

At first glance it appears that the singularity in position has simply been replaced by one in time. Since the error in the leapfrog method vanishes as $\Delta t \to 0$, this would still be troublesome.[3] However, when equation (2.17) is expanded in small Δt the time singularity is again replaced by the original singularity in x. Nonetheless, apart from a slight increase in computation time per particle due to the square root and increased complexity of the equation, the change of variable succeeds in reducing the order of the singularity by two—an impressive simplification!

[3] Of course there are always round-off errors when dealing with computer simulations and care must be taken not to choose a time step that is too small.

There is another advantage of using an r^2–z grid. Making a uniform grid in x creates equivolumetric cells that contain approximately equal numbers of particles. This helps to reduce numerical noise—especially near the central axis. Furthermore, the cell width is inversely proportional to r, making a finer grid near the edge of the plasma. Therefore, we have the advantage of resolving the fine detail in the edge without paying for the enormous computation cost associated with such fine grids.

Field solver

The third and final element of our leapfrog algorithm is the field solver. The field solver first calculates the potential from the charge density and boundary conditions and then uses the potential to calculate the electric field.[4] Rewritting Poisson's equation in terms of our new variable x we obtain

$$4\frac{\partial}{\partial x}\left(x\frac{\partial \phi}{\partial x}\right)+\frac{\partial^2 \phi}{\partial z^2}=-\frac{\rho}{\varepsilon_0}, \tag{2.19}$$

whereupon finite differencing yields

$$\frac{4}{\Delta x}\left(x_{i+1/2}\frac{\phi_{i+1,j}-\phi_{i,j}}{\Delta x}-x_{i-1/2}\frac{\phi_{i,j}-\phi_{i-1,j}}{\Delta x}\right)+\frac{\phi_{i,j+1}-2\phi_{i,j}+\phi_{i-1,j}}{\Delta z^2}=-\frac{\rho_{i,j}}{\varepsilon_0}, \tag{2.20}$$

where $\phi_{i,j}$ is used as shorthand for $\phi\left(x_i, z_j\right)$ (similarly for ρ). Because the indices i and j so closely resemble the indices of a matrix, it is convenient to use matrix operations to solve for the potential. However, the math is more transparent if the potentials and densities are relabeled using a single index. This can be done by simply assigning each grid point a number (i.e. counting the grid points) starting at, say, the bottom left corner of the grid and ending at the top right corner of the grid as shown in Table 2.1. Thus, if $1 \leq i \leq N_r$ and $1 \leq j \leq N_z$, then the total number of grid points through which the new index ranges is $1 \leq k \leq N = N_r N_z$. The problem can now be written using matrices

$$\mathbf{A}\boldsymbol{\phi}=\mathbf{b}, \tag{2.21}$$

[4]The magnetic field produced by the motion of the charges is negligible [1].

Table 2.1. Relabeling the potentials using a single index allows equation 2.5 to be put into matrix form. This table illustrates one method of replacing the double index with a single index by counting the grid points starting from the bottom left corner to the top right corner.

$i = N_r$	$\phi_{N_r,1} = \phi_{N_r}$	$\phi_{N_r,2} = \phi_{2N_r}$	$\phi_{N_r,3} = \phi_{3N_r}$	$\cdots$	$\phi_{N_r,N_z} = \phi_{N_rN_z}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
$i = 3$	$\phi_{3,1} = \phi_3$	$\phi_{3,2} = \phi_{N_r+3}$	$\phi_{3,3} = \phi_{2N_r+3}$	$\cdots$	$\phi_{3,N_z} = \phi_{N_r(N_z-1)+3}$
$i = 2$	$\phi_{2,1} = \phi_2$	$\phi_{2,2} = \phi_{N_r+2}$	$\phi_{2,3} = \phi_{2N_r+2}$	$\cdots$	$\phi_{2,N_z} = \phi_{N_r(N_z-1)+2}$
$i = 1$	$\phi_{1,1} = \phi_1$	$\phi_{1,2} = \phi_{N_r+1}$	$\phi_{1,3} = \phi_{2N_r+1}$	$\cdots$	$\phi_{1,N_z} = \phi_{N_r(N_z-1)+1}$
	$j = 1$	$j = 2$	$j = 3$	$\cdots$	$j = N_z$

where

$$\boldsymbol{\phi} = \begin{pmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \\ \vdots \\ \phi_N \end{pmatrix}, \qquad \mathbf{b} = \begin{pmatrix} -\frac{\rho_1}{\varepsilon_0} \\ -\frac{\rho_2}{\varepsilon_0} \\ -\frac{\rho_3}{\varepsilon_0} \\ \vdots \\ -\frac{\rho_N}{\varepsilon_0} \end{pmatrix}, \tag{2.22}$$

and $\mathbf{A}$ is an $N \times N$ banded matrix built directly out of the remaining quantities in equation (2.20). The band has a width of $2N_r + 1$, but has at most five nonzero values in each row because the potential at any given point depends only on its four nearest neighbors. A typical row in $\mathbf{A}$ will have the following form

$$\begin{pmatrix} & & \ddots & & & & & & \ddots & & & & \ddots & & & & \\ 0 & \cdots & 0 & \frac{1}{\Delta z^2} & 0 & \cdots & 0 & \frac{4x_{i-1/2}}{\Delta x^2} & -\frac{4}{\Delta x^2}\left(x_{i+1/2}+x_{i-1/2}\right) - \frac{2}{\Delta z^2} & \frac{4x_{i+1/2}}{\Delta x^2} & 0 & \cdots & 0 & \frac{1}{\Delta z^2} & 0 & \cdots & 0 \\ & & & \ddots & & & & & \ddots & & & & & \ddots & & & \end{pmatrix} \tag{2.23}$$

although it should be noted that the symmetry is altered somewhat for points located at the corners or along the edges of the grid since the number of neighboring points decreases. Also note that the index i has been left explicit on the radial positions to emphasize the fact that the values are cycled

every N_r rows.

While the matrix $\mathbf{A}$ is very large and relatively expensive computationally to assemble, it is important to realize that it only needs to be done once in the entire simulation and thereafter allows all of the potentials in a particular time step to be solved directly in one operation, $\boldsymbol{\phi} = \mathbf{A}^{-1}\mathbf{b}$. Furthermore, only the diagonals of the band need be stored because the rest of the matrix is filled with zeros. This reduces the number of values to be stored from N^2 down to $N(2N_r+1)$.[5]

Boundary Conditions

The boundary conditions for an axisymmetric grounded Malmberg-Penning trap are intuitive for an r–z grid—the potential is zero along conducting boundaries and has zero radial derivative along the $r = 0$ axis. By swapping coordinate grids no change occurs in boundary conditions along conducting surfaces, but there *is* a subtle change at the $r = 0$ boundary. While continuity across $r = 0$ must be conserved, slope need not be. Therefore, on our new grid, the potential has a finite radial slope along the axis. To illustrate, consider equation (2.19) with no z dependence

$$4\frac{\partial}{\partial x}\left(x\frac{\partial \phi}{\partial x}\right) = -\frac{\rho}{\varepsilon_0}. \tag{2.24}$$

If we do a power series expansion of the potential ($\phi = c_1 + c_2 x + c_3 x^2 + \ldots$) and plug this into the equation above we get

$$c_2 + 4c_3 x + 9c_4 x^2 + \ldots = -\frac{\rho}{4\varepsilon_0}. \tag{2.25}$$

[5]To illustrate the advantage of using the banded matrix over the full matrix, consider a grid with 100×1000 grid points. With each matrix element requiring 8 bytes of storage, the banded matrix only uses 153 MB while the full matrix surprisingly requires over 74 GB of space. Even though 74 GB of storage is becoming more feasible with modern technology, this information must be accessed every time step; sifting through such an enormous amount of memory would make the code run impossibly slow. There are other methods that are even more efficient than the banded matrix solver, but require an enormous amount of work for very little gain in speed.

Now solving for the derivative of the potential, evaluated at $x = 0$, we find

$$\left.\frac{\partial \phi}{\partial x}\right|_{x=0} = c_2 = -\left.\frac{\rho}{4\varepsilon_0}\right|_{x=0}, \tag{2.26}$$

where the second equality follows after substitution from equation (2.25). Thus we see that the derivative boundary condition is indeed nonzero and in fact depends on the density at the origin.

The fact that the derivative depends on a quantity at the origin brings up an important question as to the type of grid that will be used in the simulation. There are two main types of rectangular grids—cell-center and cell-edge. Cell-center grids with ghost points usually have an advantage when defining Neumann (derivative) boundary conditions because the boundaries can easily be defined using a finite differenced derivative involving quantities already on the grid. Conversely, cell-edge girds are advantageous with Dirichlet (fixed) boundary conditions because the boundary can be directly assigned without recourse to interpolation/extrapolation. While we technically have a Neumann boundary condition at the origin, the derivative depends on a non-cell-center grid quantity that fluctuates in time. However, if we persist in using a cell-center grid with ghost points, at the origin (i.e. $i = 1$) we set $x_{1/2} = 0$ and $x_{3/2} = \Delta x$ in equation (2.20) and the resulting relation simply regurgitates the boundary condition—we get no new information. This, coupled with the fact that all of the other boundaries are grounded (Dirichlet), has led us to chose a cell-edge grid in our PIC code.

Using a cell-edge grid, however, requires extrapolation from interior points in order to get the derivative at the boundary. Linear extrapolation is fast and easy to implement, but with a little more effort, and a very small loss in speed, we can get better accuracy using quadratic extrapolation. We start by letting $\phi = ax^2 + bx + c$, where a, b, and c are constants yet to be determined.[6] We then take the derivative of ϕ, evaluate it at $dx/2$, and set it equal to its equivalent finite difference derivative. Similarly, we can obtain a second equation for the point $3dx/2$. The resulting system of equations

[6] Actually, we don't need to find c because we are only interested in the derivative of ϕ where c is no longer present.

is

$$\begin{cases} \frac{\phi_3-\phi_2}{dx} = b + a\,dx \\ \frac{\phi_4-\phi_3}{dx} = b + 3a\,dx, \end{cases} \tag{2.27}$$

from which a and b can now be found. Upon substituting these constants back into our derivative and evaluating our equation at the boundary (i.e. $r = 0$), we get our desired quadratic extrapolation formula:

$$\phi_1' = -\frac{3}{2dx}\phi_2 + \frac{2}{dx}\phi_3 - \frac{1}{2dx}\phi_4. \tag{2.28}$$

This, together with the charge density and the corresponding z components from equation (2.20), are what make up our final axial boundary condition.

2.3 Analysis

The PIC code itself is not written to do any of the analysis of the data. Instead we save the root-mean-square (RMS) radial and axial positions of the plasma as well as the grid's fluid velocity and charge density at every time step to analyze after the simulation has finished. However, as was mentioned in the introduction (see page 2), the only quantity that can be measured in the experiment is the induced charge on a diagnostic ring placed near the end of the plasma.[7] Therefore, as a final diagnostic, we also output the charge induced on a small section of the conducting wall near one of the confinement rings. The RMS values, fluid velocity, and charge density help in identifying the modes, but it is this charge density that ultimately identifies how well the experiment will be able to pick up on these radial Bernstein modes.

We initially look for resonance peaks in the fast Fourier transforms (FFT) of the RMS values and induced charge. Then, once a potential mode frequency has been located, the fluid velocities

[7] Fortunately, the surface charge density (σ) is easy to compute from the difference of inner and outer radial electric fields at the conducting wall ($\mathbf{E}_1$ and $\mathbf{E}_2$ respectively), in particular $\sigma = \varepsilon_0 (\mathbf{E}_2 - \mathbf{E}_1) \cdot \mathbf{n}$. This is further simplified by the fact that the trap walls are conductors and, hence, the field inside is identically zero.

are used to construct the individual mode shapes. This is done by looking at the radial velocity flux. In practice this is done by performing a least squares fit (LSF) of the amplitude, frequency and phase of the function

$$f(t) = A\cos(\omega t + \phi). \tag{2.29}$$

Even though we already have an idea of the frequency, the exact frequency may lie in between points in our FFT. Therefore, the LSF allows us to determine more accurately the frequency and phase. Replacing the frequency and phase in equation (2.29) and taking into account the 90° phase shift between velocity and displacement, we then proceed to perform a LSF of the amplitude at every grid point. The amplitude of the radial velocity flux, being sensitive to the nodal structure of the mode, creates a distinct surface we term the mode shape. This allows us to identify the modes we are actually detecting.

Chapter 3

Testing

An important aspect of computational work is ensuring that the code is error free and that the resulting simulation actually relates to the physical world. However, because PIC codes are actually just an ensemble of smaller functions working together, it is difficult to come up with a single test that verifies the entire code is working. Often errors will not exhibit themselves outright until just the right set of parameters are chosen, making them difficult to pinpoint in a large code. Fortunately, once one section of code has been tested, it can be used in evaluating another whether as a tool itself or as an element of a larger function. For this reason we chose to test our field solver and mover first. Then, once these had been verified to be working correctly, we went about evaluating our interpolation/extrapolation technique and whether or not it produced unphysical self-forces.

3.1 Field Solver

The basic solver was first created and tested in Matlab in order to take advantage of the pre-defined matrix operations, particularly those relating to banded matrices. The solver was tested against exact solutions to Poisson's equation ($\nabla^2\Phi = \rho$) for arbitrary (not necessarily physical)

functions of ρ and found to give normalized root mean square errors (NRMSE) on the order of 10^{-5} for a sufficiently fine grid. Of course the real difficulty is in writing the banded matrix solver itself. For this we used a Fortran subroutine written by J. P. Moreau [12, 13] that uses LU decomposition without pivoting to solve linear systems with nonsingular linear banded system matrices. Comparing the results of the two field solvers we found them to be identical, modulo roundoff errors, for both single and multiple particle calculations.

3.2 Mover

The mover itself is very basic and does not involve a lot of testing because it is simply a finite differenced equation derived from an analytic solution. However, depending on the time step chosen, negative quantities can occasionally occur beneath the radical due to roundoff, thus resulting in imaginary velocities and positions. The size of the time step can also cause particles to occasionally cross the $r = 0$ boundary resulting in a negative position. In both of these cases, it is sufficiently accurate to simply force the respective negative quantities to be positive by taking their absolute value. Additionally, care must be taken to remove any particles from calculation that pass through any of the conducting boundaries. A simple diagnostic has been coded to ensure that this is true.

3.3 Interpolation/Extrapolation

Interpolation and extrapolation are used extensively throughout the code.[1] The bilinear method has been used for decades and known to work, but it was somewhat unclear whether the interpolation/extrapolation could be done without loss of precision using the fractional pseudo-areas in

[1] Although interpolation and extrapolation are technically different techniques, the bilinear method uses the same weighting factors in each case and can be tested simultaneously.

the r^2–z coordinates or if it was absolutely necessary to use the fractional areas computed in r–z. A simple example illustrates the difference (for this and the discussion that follows we will assume that the particle is always located at the midpoint between axial grid points since the axial dimension is the same in both cases). Suppose we space our grid points equally in r^2, but place a charge at the center of a particular cell in r (see Figure 3.1). If we were to distribute the charge

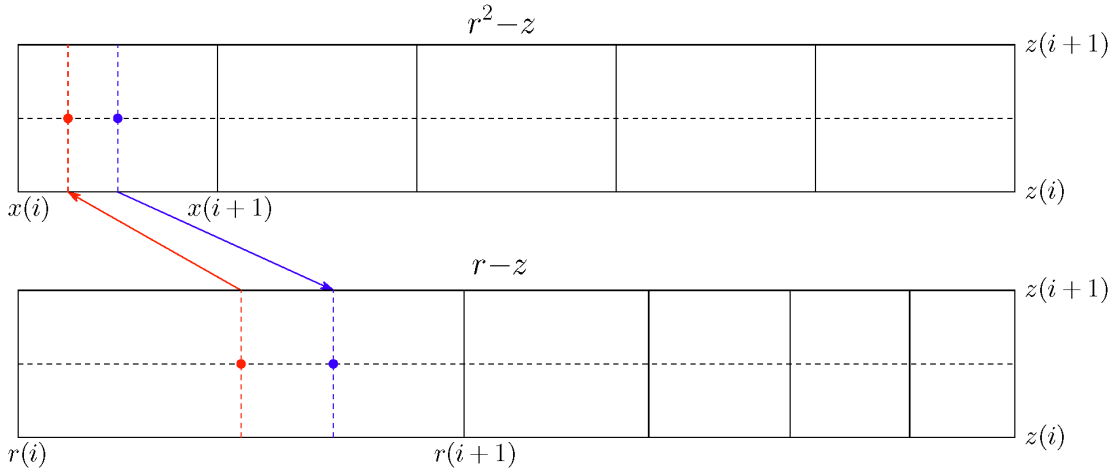

Figure 3.1. The two sets of coordinate grids (top: r^2–z; bottom: r–z) show the difference between placing a particle at the center of a cell in r (shown in red) as opposed to placing it at the center in x or r^2 (shown in blue). Depending on which coordinate system is used, the surrounding grid points will be assigned different charge densities according to their relative proximity to the charge.

of this particle to the surrounding grid points in r, each would receive exactly $1/4$ of the charge. Alternately, we can calculate the position of this particle on the r^2 grid and find that the particle does not sit at the center of the cell in this coordinate system. Therefore, when we distribute the charge bilinearly in this frame we distribute more charge to the grid points on the left than to the grid points on the right. More explicitly the left grid points receive

$$\frac{x(i+1)-\frac{1}{4}\left[x(i)+x(i+1)+2\sqrt{x(i)x(i+1)}\right]}{2[x(i+1)-x(i)]}, \tag{3.1}$$

while the grid points on the right receive

$$\frac{\frac{1}{4}\left[x(i)+x(i+1)+2\sqrt{x(i)x(i+1)}\right]-x(i)}{2[x(i+1)-x(i)]}. \tag{3.2}$$

More generally we can extend this to a particle at an arbitrary radial position within the cell. To make the equations more transparent and because we are only interested in fractional areas anyway, let α be the fractional length of the cell that represents the distance to the particle as measured from the left edge of the cell. Bilinear extrapolation in r would dictate that the left and right grid points would receive the fractional densities $(1-\alpha)/2$ and $\alpha/2$ respectively. In contrast, when we use the same method in r^2, equations (3.1) and (3.2) become

$$\left(\frac{1-\alpha}{2}\right)\left[\frac{(1-\alpha)\sqrt{x(i)}+(1+\alpha)\sqrt{x(i+1)}}{\sqrt{x(i)}+\sqrt{x(i+1)}}\right] \tag{3.3}$$

$$\left(\frac{\alpha}{2}\right)\left[\frac{(2-\alpha)\sqrt{x(i)}+\alpha\sqrt{x(i+1)}}{\sqrt{x(i)}+\sqrt{x(i+1)}}\right], \tag{3.4}$$

for the left and right grid points respectively. Thus, by switching from r to r^2 we pick up a multiplicative factor in the fractional areas. While these equations are good for *illustrating* the difference between coordinate systems, they are not conveniently written for *testing* because α is defined in the r coordinate system and all the variables are defined in r^2. If we define β to be the analogue of α in r^2 and then reverse the derivation by going from r^2 to r, the equations are much easier to code. For reference they are as follows:

$$\frac{\sqrt{x(i+1)}-\sqrt{(1-\beta)x(i)+\beta x(i+1)}}{2[x(i+1)-x(i)]} \tag{3.5}$$

$$\frac{\sqrt{(1-\beta)x(i)+\beta x(i+1)}-\sqrt{x(i)}}{2[x(i+1)-x(i)]}. \tag{3.6}$$

It is well known that bilinear interpolation/extrapolation works in r; therefore, one way we can test the method of interpolation/extrapolation in r^2 is by merely reproducing the results of the calculations in r. Obviously minor differences in numerical values can be expected, but small

deviations are permitted since the bilinear method itself—while very efficient—is far from being the most accurate method available. In fact, the method is only guaranteed to work for small grid spacings. Thus we can expect that the greatest discrepancy will occur near $r = 0$ where the cells in r will be the largest. First, we compared calculations of the electrostatic potential and field generated by a single particle located at various radii using the two weighting systems. Typically the two methods had a normalized root mean square deviation (NRMSD) of about 10^{-5}. The only exception occurred, as expected, when the particle was within the first radial grid point due to its increased size in r. However, even though the error was large near the point charge, the NRMSD for the entire grid was still on the order of 10^{-3} indicating that long range effects still showed a fair amount of agreement. Fortunately, as we increased the number of particles contained in the cells along the axis the discrepancy between the two methods diminished significantly.

3.4 Self-Forces

In addition to testing the interpolation/extrapolation technique on the new grid, it is important to verify that the new interpolation/extrapolation scheme does not produce any unphysical self-forces on the particles. Conceptually, these self-forces are most easily found by analytically solving for the total force exerted on a single particle in the region of interest and then comparing it to the force produced in the simulation. However, this is one case where having fewer dimensions complicates the calculations rather than simplifying them. To illustrate, consider a point charge in 3D free space. While the charge creates a field that permeates all of space, it cannot exert a force on itself; the self-force is zero. Even when a grounded conducting boundary is added and the self-force is not identically zero, this force can be calculated (see page 56 in Appendix B).

On the other hand, a point charge in 2D free space physically represents an extended object in 3D to account for the assumed symmetry in the vacant dimension. Therefore, in an r–z grid every

point charge physically represents a ring of charge and a ring of charge does produce a repulsive self-force that causes the ring to expand outward. While the electric field generated by a ring of charge can be calculated using Elliptic integrals, the equation contains a pole at the location of the ring. Furthermore, the electric field produced by the differential length in question would have to be somehow removed. Unfortunately, even by removing this "extra" field, the force retains its logarithmic pole. This can be demonstrated by expanding the potential in a Taylor series about the wire. The potential at a point (ρ, z) produced by a uniformly charged ring with radius R and centered about the z axis is given by (see Appendix C)

$$\Phi(\rho,z) = \frac{\lambda R}{\pi\varepsilon_0\sqrt{(R+\rho)^2+z^2}}\,K\left[\frac{4\rho R}{(R+\rho)^2+z^2}\right]. \tag{3.7}$$

If we expand this equation about the coordinates $\rho = R$ and $z = 0$, we obtain an approximation to the field in the vicinity of the ring. The expansion is difficult to do correctly because of the elliptic integral. It is easiest to perform the two expansions separately and combine the two results "by hand." Expanding about $\rho = R$ and equating $z = 0$ yields

$$\Phi_{(\rho\to R,z=0)} = \frac{2\lambda}{4\pi\varepsilon_0}\left\{\ln\left(\frac{8R}{\rho-R}\right)+\frac{\rho-R}{2R}\left[1-\ln\left(\frac{8R}{\rho-R}\right)\right] + \frac{(\rho-R)^2}{16R^2}\left[-7+5\ln\left(\frac{8R}{\rho-R}\right)\right]+O[(\rho-R)^5]\right\}, \tag{3.8}$$

while expanding about $z = 0$ and equating $\rho = R$ gives

$$\Phi_{(\rho=R,z\to0)} = \frac{2\lambda}{4\pi\varepsilon_0}\left\{\ln\left(\frac{8R}{z}\right)+\frac{z^2}{16R^2}\left[-7+5\ln\left(\frac{8R}{z}\right)\right]+O(z^7)\right\}. \tag{3.9}$$

Evidently only odd orders in $\rho - R$ (0, 1, 3, 5, ...) and only every fourth order in the z (0, 3, 7, 11, ...) survive. For the orders that are present, there is a corresponding factor of z for every factor of $\rho - R$. Naively we may guess that the general expression for the expansion in both coordinates would therefore be

$$\Phi_{(\rho\to R,z\to0)} = \frac{2\lambda}{4\pi\varepsilon_0}\left\{\ln\left(\frac{8R}{r}\right)+\frac{\rho-R}{2R}\left[1-\ln\left(\frac{8R}{r}\right)\right]+\cdots\right\}, \tag{3.10}$$

where $r = \sqrt{(\rho - R)^2 + z^2}$ now represents a radial distance from the ring. Validation for this choice comes from the excellent fit it provides near the ring and the fact that it noticeably outperforms any other expansion we have derived.

The real value in this expansion is not in the fit however, but rather its form. Notice that the zeroth order term in the expansion is simply the potential from a uniformly charged *straight* (infinite) wire. This reflects the fact that a test charge very close to the loop sees a straight wire with no curvature. The higher order terms carry the entirety of the curvature in the rest of the ring. Thus by removing the zeroth order potential we are ridding ourselves of the portion of the potential (and thereby the field) that does not contribute to the self-force. The hope was that this small fraction of the field would be enough to keep the self-force finite, but it is plain to see from the expansion that all higher order terms carry the same logarithmic pole. Therefore, a direct comparison with analytical results is unavailable.

Alternatively, we can an use an amalgamation of part simulation and part theory in an effort to validate, at least to some degree, the results from the PIC simulation. Because the simulation uses the electric field interpolated from four grid points to calculate the force on a particular charge, we can mimic the simulation analytically by using four uniformly charged rings and a point charge placed within the rectangular "torus" generated by these four rings (see Figure 3.2). This then sidesteps the issue with the pole by removing it to another location altogether and allows the force on the charge to remain finite. This, of course, assumes that the point charge remains at a respectable distance from the rings. While this is potentially problematic as the grid size decreases, it at least provides some insight into how well the simulation is performing.

We start by generalizing equation (3.7) to a ring with nonzero height Z. This is done by simply replacing z by $Z - z$ everywhere in the equation. Only being interested in the radial electric field

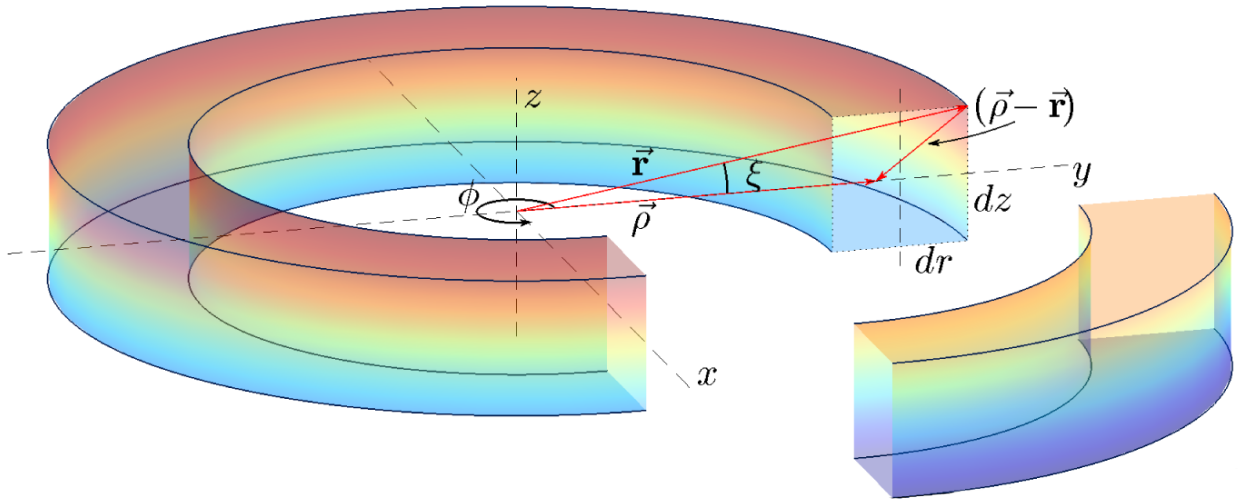

Figure 3.2. Setup for the force calculation using four uniformly charged rings located at the particle's cell grid points.

we take the derivative of Φ with respect to ρ:

$$E_r(\rho,z) = \frac{\lambda R}{2\pi\varepsilon_0\rho\sqrt{(R+\rho)^2+(Z-z)^2}}\left\{K\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right]\right.$$
$$\left.-\frac{R^2-\rho^2+(Z-z)^2}{(R-\rho)^2+(Z-z)^2}\,E\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right]\right\}. \quad (3.11)$$

Figure 3.3 shows the comparison between the self-forces obtained using the PIC code and the analytic solution. As expected, the analytic solution wildly oscillates from one side of the cell to the other. The points close to either edge are obviously wrong; at the right edge the force becomes negative and an expanding ring cannot pull itself back in towards the middle while the symmetry about the cell's midpoint suggests that the error near the left edge of the cell is just as egregious, but in the opposite direction. This confirms our assumption that the midpoint, if any, is the best point with which to compare our code. Surprisingly, the midpoint shows fair agreement with the PIC code even when the grid cells become quite small; we may ignore the image charge effects in the PIC calculation that occur near the conducting wall since the analytic solution was for a ring in free space. The discrepancy between the two solutions near the origin is probably a result of errors

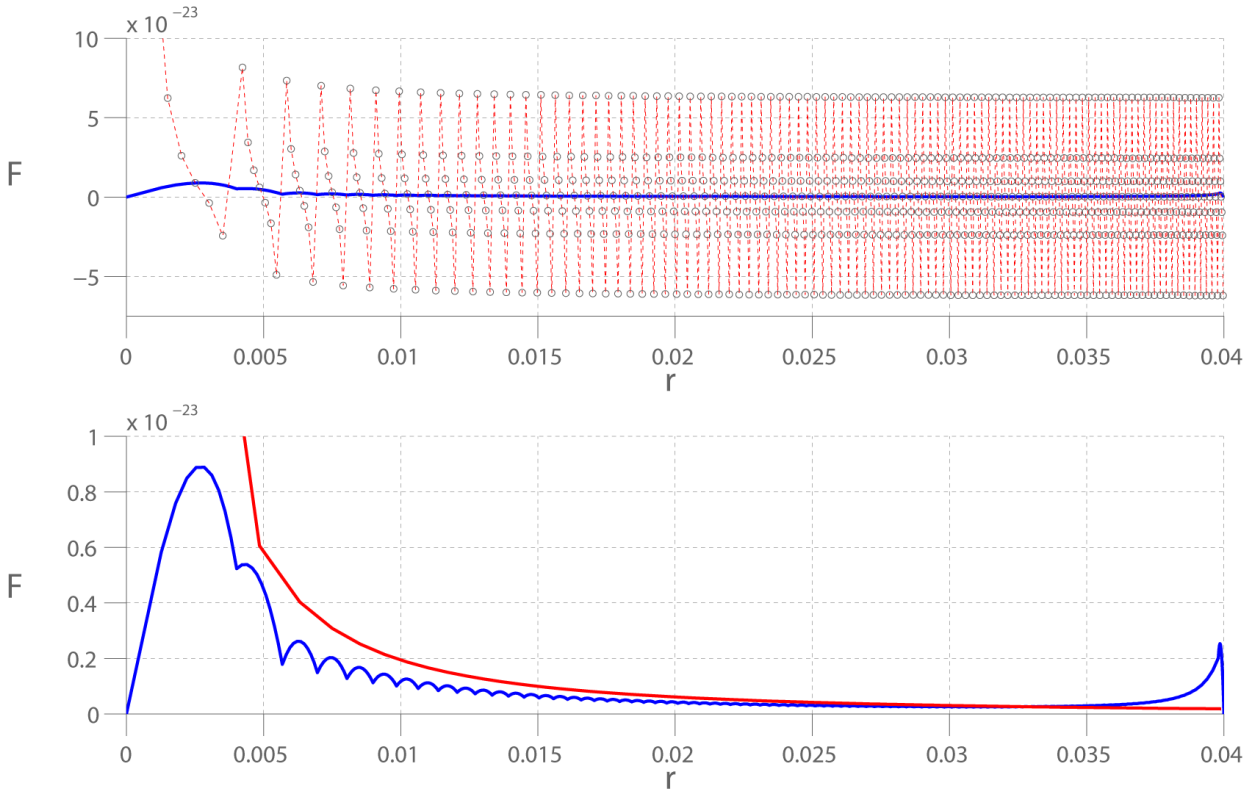

Figure 3.3. Self-force of a single particle plotted versus radius at z=0. The blue lines are the force calculated using the PIC code and the red lines are calculated using the analytic solution for four uniformly charged rings located at the four nearest grid points. The top graph shows the force calculated at seven locations within each cell and the bottom graph shows only the calculation at the middle of the cell.

in both calculations and it is difficult to say which one is more accurate. Needless to say, this is far from conclusive evidence that the simulation is correct, but it does lend credence to the accuracy of the self-force calculations obtained by the PIC code.

We also compared radial and axial self-forces calculated using r and r^2 weighting schemes as we did above. This was calculated at 100 evenly spaced positions per cell for all cells straddling the trap's axial midpoint and their nearest neighbors (see Figures 3.4 and 3.5). The focus on the midpoint was to ensure that the axial force correctly vanished everywhere along $z = 0$. As before, large deviations occur within the first radial grid point due to the increased cell size inherent in

the r^2 grid. Nonetheless, the excellent agreement outside of these cells suggest that there are no self-forces that arise from using r^2 weighting factors.

Finally, because our single particle calculations were somewhat inconclusive, we also tested the self-force using multiple particles. At best, a multiple particle calculation can only prove average effects, but since the simulation requires a large number of particles to work correctly anyway, this is sufficient for our needs. Smearing the charge in space causes the poles to disappear and the potential and field to become finite, but also complicates the integrals involved. One approach to the problem would be to use Green's functions. The advantage of using Green's functions over the other methods is the ease with which you can account for conducting boundaries, finite-length effects, and irregular shapes, thus making our test as accurate as possible.

Because there are sometimes convergence issues when integrating/summing Green's functions, we compared solutions from three different derivations to ensure that none of the features were caused by numerical errors (see Appendix B for derivations of these Green's functions). Although the different methods required a different number of terms to converge, they all converged to the same results. We placed one million equally spaced particles[2] inside a rectangular region thus forming a uniformly charged rectangular torus. Figure 3.6 shows the agreement between the electric field obtained by both the PIC code and the Green's function. The largest deviations occur at the boundaries of the torus, the corners being the worst; the NRMS error for this radial cross-section is 9.59×10^{-3}. This relatively large error is a result of the PIC simulation's inability to produce sharp boundaries (removing the boundaries reduces the error by about an order of magnitude). Fortunately, the plasma does not have an infinitely sharp boundary and this discrepancy from the analytic solution is not an issue.

[2]We also tested randomly placed particles inside the same region. The results were the same, but the randomly placed particles did produce a little more noise.

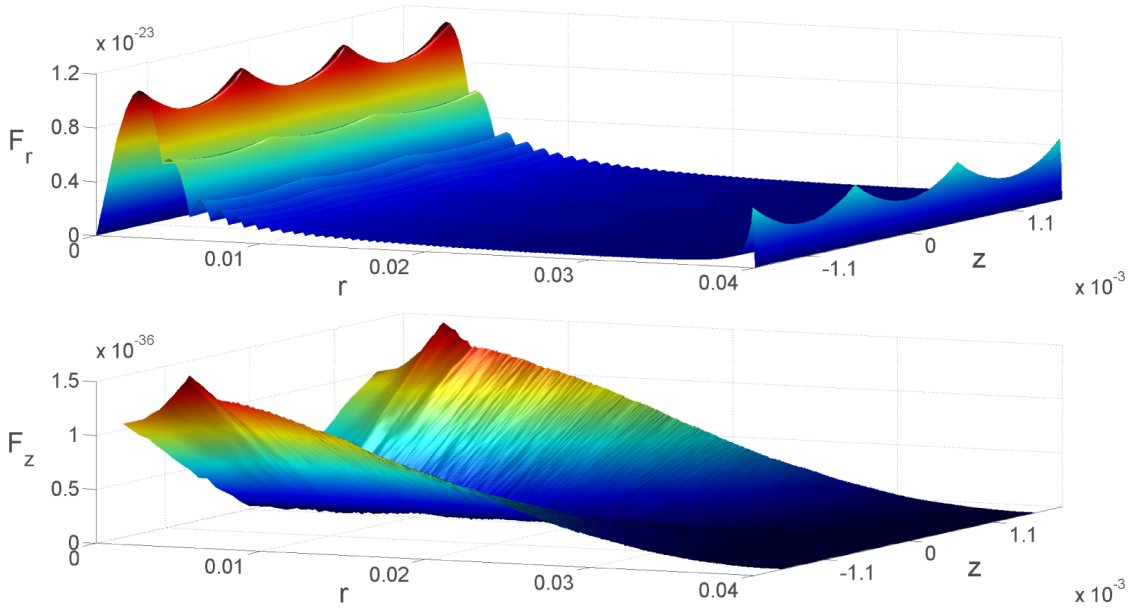

Figure 3.4. Radial self-force (top) and the absolute value of the axial self-force (bottom) of a single particle as a function of r and z near the midpoint of the trap using r^2–z weighting factors. The oscillations are an artifact of interpolation/extrapolation, but persist regardless of which weighting factors are used. Also visible is how the force becomes more and more smooth as the cells shrink in size at larger radii.

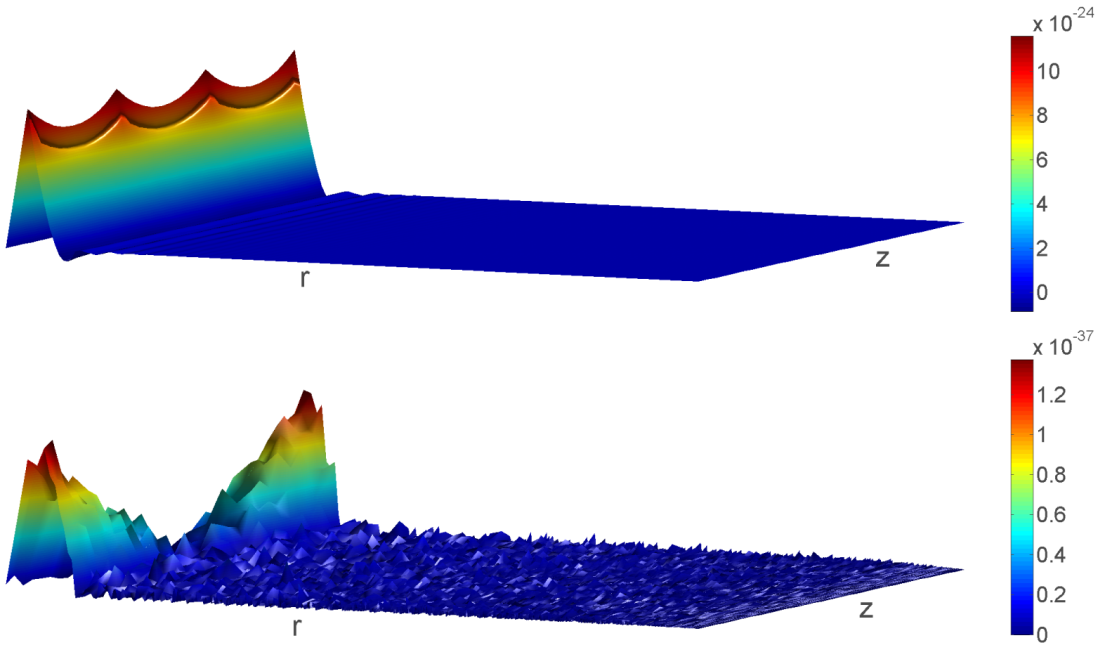

Figure 3.5. The difference between radial/axial self-forces (top/bottom) calculated using r–z and r^2–z weighting schemes. The two weighting systems yield almost identical results except near the axis.

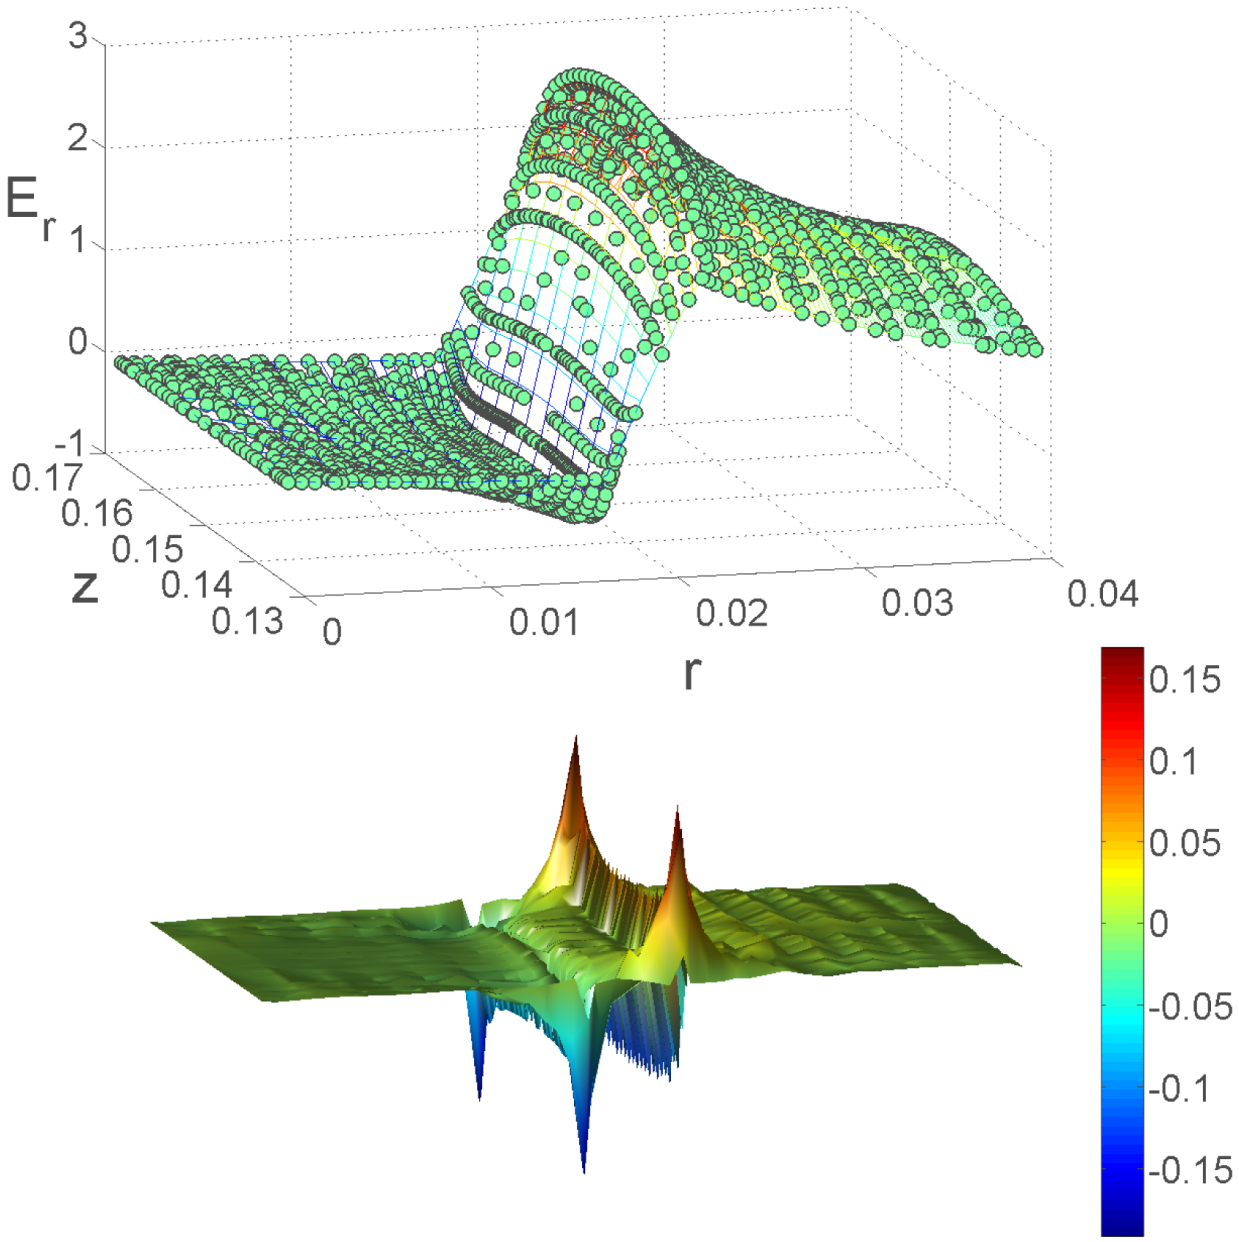

Figure 3.6. Top: a comparison between the electric fields obtained from the PIC simulation (mesh) and an analytic calculation using Green's functions (circles). Bottom: the difference between the two electric fields.

Chapter 4

Results

4.1 Chosing Parameters

The physical parameters used in the simulation are based on the actual experimental apparatus that will eventually be used to measure the radial Bernstein modes, but we have modified some of the parameters in order to simplify the simulation and increase our intuition of these modes. The biggest change was in the confinement region where we have doubled its length (now 20 cm) in order to produce a longer plasma column. This was done to limit the overlap and potential confusion and/or interaction between the Trivelpiece-Gould modes and the radial Bernstein modes in which we are interested. Doubling the length of the plasma lowers the frequency of the Trivelpiece-Gould modes without significantly altering the frequencies of the radial Bernstein modes. Once we have a good grasp on the frequencies of the radial modes, we can then revert back to the smaller confinement region and really assess the finite-length effects and the possibility of detecting these modes under these conditions. Another minor difference is that we have made the confinement rings exactly the same width (10 cm) whereas machining and trap design have given rise to slight asymmetries in their widths in the actual experiment. The diagnostic ring on the other hand, re-

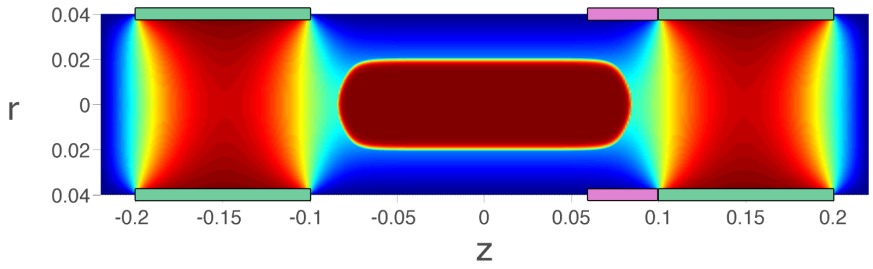

Figure 4.1. Overlaid colormap of the plasma density and confining potential inside the computation region. The green and purple rectangles represent the confinement (10 cm wide) and diagnostic (4 cm wide) rings respectively.

mains consistent with the experiment at a width of 4 cm and sits just inside one of the confinement rings.

The other parameters are a bit more arbitrary because there is no way of knowing the exact number density that will be produced in a given experimental shot. The only physical limitations are in the maximum magnetic field (.44 T) and confinement voltage (150 V) that we can produce in the experiment. From work done on the 1D infinite-length plasma, we know that the radial Bernstein modes achieve the greatest separation at about $6/7$ of the Brillouin limit. Therefore, we have chosen a peak central density of 1.4846×10^{13} particles/m^3 and a magnetic field of .2144 T which corresponds to $\beta = .85$.[1] With these parameters, we need the full 150 V to provide the proper axial confinement of the plasma. Finally, we use a relatively cold thermal equilibrium temperature of 0.2 eV in an effort to reduce damping of the modes. The combination of these parameters produces a plasma column that is approximately 17 cm in length and 2 cm in diameter (see Figure 4.1).

The plasma is driven by sinusoidally modulating the voltage with an amplitude of 5 V on both confinement rings for a total of 50 cyclotron periods. The drive is then turned off and we record the RMS radial and axial positions, the charge on the diagnostic ring, and the fluid velocity on the grid

[1] The particles are ^{7}Be ions for experimental reasons.

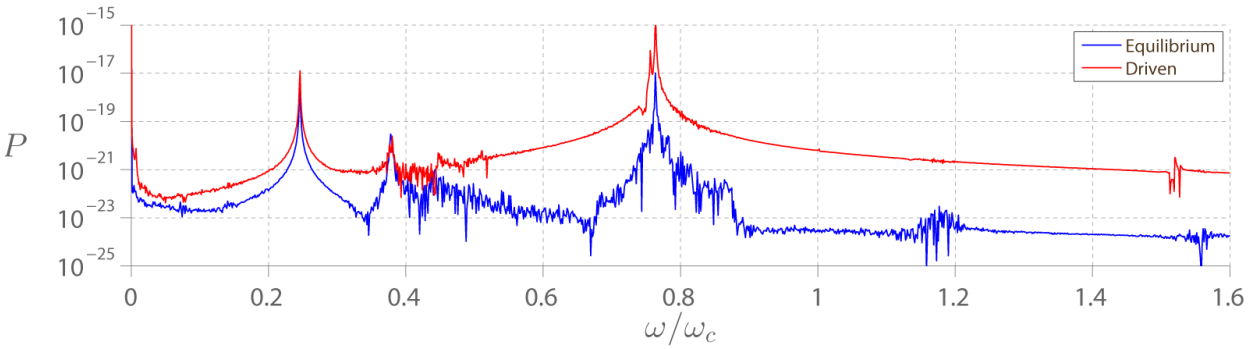

Figure 4.2. The equilibrium RMS radial FFT spectrum is shown in blue while the spectrum obtained from a plasma driven at 5 V is shown in red. The driven peak is two orders of magnitude larger than the equilibrium case and does not produce any nonlinear effects.

for 1000 cyclotron periods. The amplitude of the modulating potential was chosen by comparing FFT spectra produced by various drives with that of the equilibrium spectrum. If too large an amplitude was used, we found that non-linear effects dominate and shift the mode frequencies significantly. Conversely, an amplitude that was too small produced undetectable differences from the equilibrium spectrum. Therefore, we incrementally varied the amplitude from 1 to 15 V and found that the driving amplitude needed to be about 5 V in order to ensure detectable changes in the FFT spectrum while keeping the plasma in the linear regime. Figure 4.2 compares the equilibrium RMS radial FFT spectrum with that of a plasma driven with a 5 V potential.

4.2 Fundamental Mode

Driving the plasma at a frequency of $0.7608\,\omega_c$ we observed three main features in our FFT spectra (see Figure 4.3), all of which corresponded very closely to estimated frequencies we obtained using a 1D code (see Table 4.1).[2] First, we were able to identify the $m_z = 1, 2$, and 3 Trivelpiece-Gould modes. Interestingly, these were all slightly downshifted from the 1D case and

[2]The agreement between codes provides further evidence that the PIC code is working properly.

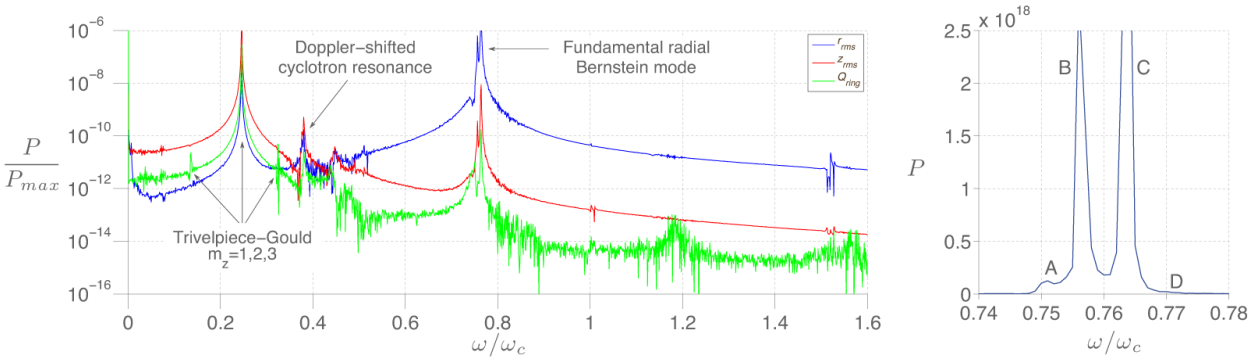

Figure 4.3. Left: the normalized FFT power spectrum for the RMS radial (blue) and axial (red) positions along with the charge on the diagnostic ring (green). The three prominent resonance peaks are identified as the Trivelpiece-Gould $m_z = 2$ mode, the Doppler-shifted cyclotron resonance, and the fundamental radial Bernstein mode. Right: a close up of the fundamental radial Bernstein mode obtained from the radial velocity flux illustrating the small cluster of modes in the vicinity of the fundamental (C). Although not particularly obvious here, it appears that modes with higher k_z occur on either side of the fundamental.

Table 4.1. A list of the frequencies we expected to find in our FFT spectra and the frequencies we actually observed. The first three frequencies correspond to three of the Trivelpiece-Gould modes and the other two frequencies are respectively the Doppler-shifted cyclotron resonance and vortex frequency. It so happens that for the β we are using $\omega_v \approx 2\widetilde{\omega}$.

	Expected Frequency (ω/ω_c)	Observed Frequency (ω/ω_c)
$m_z = 1$	0.144	0.136
$= 2$	0.265	0.246
$= 3$	0.356	0.325
$\widetilde{\omega}$	0.387	0.380
ω_v	0.758	0.764

the gap increased with higher mode number—probably a result of finite-length effects. The $m_z = 1$ and 3 modes were significantly weaker than the $m_z = 2$ mode because the drive was symmetric in z. Therefore, it was somewhat surprising that we could even see these modes at all. However, in hindsight, it makes sense that these modes showed up only on the diagnostic ring because it is not centered about the middle of the plasma.

The other two prominent features in the FFT spectra were the Doppler-shifted cyclotron resonance ($\widetilde{\omega}$) and the fundamental radial Bernstein mode which is resonant at $2\widetilde{\omega}$. While $\widetilde{\omega}$ was slightly downshifted from the 1D case, we observed a small increase in the frequency of the fundamental radial Bernstein mode. Again harmonics of both frequencies are easily identifiable. As can be seen in Figure 4.3, the fundamental radial Bernstein resonance is actually a compact set of modes with different axial wave numbers (k_z); essentially, the end potentials act as fixed nodes on a string and allow standing waves with odd half integer wavelengths to be erected in the axial direction. These standing waves must slightly change the resonant frequency giving rise to the tight cluster of distinct resonances in our FFT spectrum. The interesting thing is that they appear to straddle both sides of the fundamental frequency, although the peaks above the fundamental tend to be very noisy and difficult to isolate so it is not entirely clear whether these are real or just numerical artifacts of our analysis.

Mode shapes were obtained for each of the four frequencies listed in the closeup in Figure 4.3. The mode shapes are usually normalized by the number density, but this distorts the shape near the edge where the density decreases drastically. In effect this causes a sharp kick near the edge followed by a large amount of noise. Therefore, for clarity, we have chosen to display the unnormalized shapes.[3] Figure 4.4 shows the mode shape of the main peak with $k_z = 1/2$, Figures 4.5 and 4.6 are the two lower frequencies corresponding to $k_z = 3/2$ and $k_z = 5/2$ respectively, and

[3]This too can distort the shape; we have found this to be especially true of the axial shape of the modes. The normalized shapes tend to provide a much more uniform amplitude across the axial direction. However, the basic mode shape is most easily seen by the unnormalized shapes.

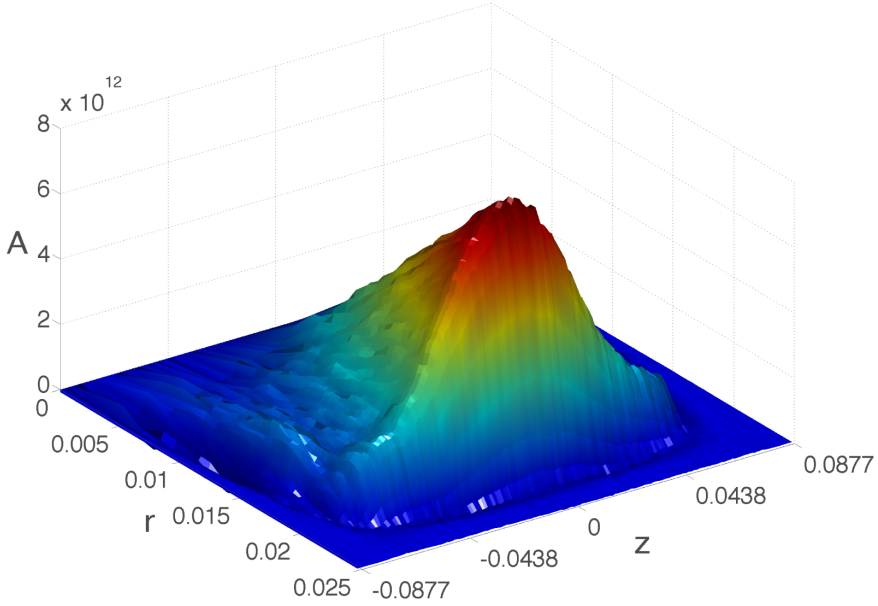

Figure 4.4. Mode shape for the main frequency peak with $k_z = 1/2$ (marked "C" in Figure 4.3).

Figure 4.7 is an upper frequency with $k_z = 3/2$.[4] From the 1D theory we expect the mode shapes to be a Bessel function of the first kind with a general form:

$$A J_1(k_r r), \tag{4.1}$$

where A is an arbitrary constant and k_r is the radial wave number. For the fundamental mode, k_r should be quite small and the mode shape should be approximately linear inside the plasma. To verify this we performed nonlinear least squares fits to radial cross-sections of the normalized mode shapes (see Figure 4.8). While only the largest frequency peaks provide clean data, the other modes have also been included for completeness.

The axial structure of these modes has important implications in where best to measure these modes with the diagnostic ring. Although we have assumed the same setup as the experiment in

[4]It should also be noted that fundamental radial Bernstein mode occurs very near the third harmonic of the $m_z = 2$ Trivelpiece-Gould mode. While the harmonics are relatively weak and we have not identified any adverse effects, this could potentially cause some interference.

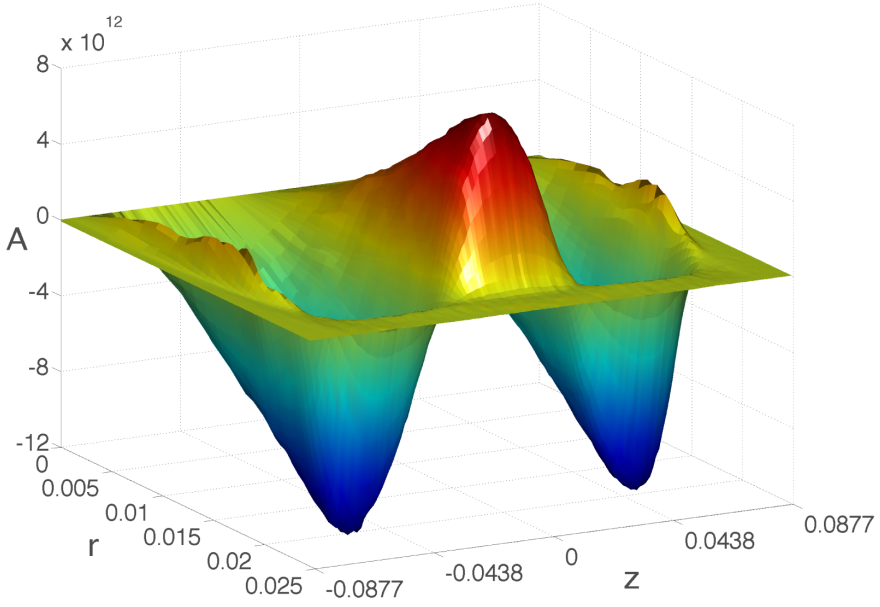

Figure 4.5. Mode shape for the first lower frequency with $k_z = 3/2$ (marked "B" in Figure 4.3).

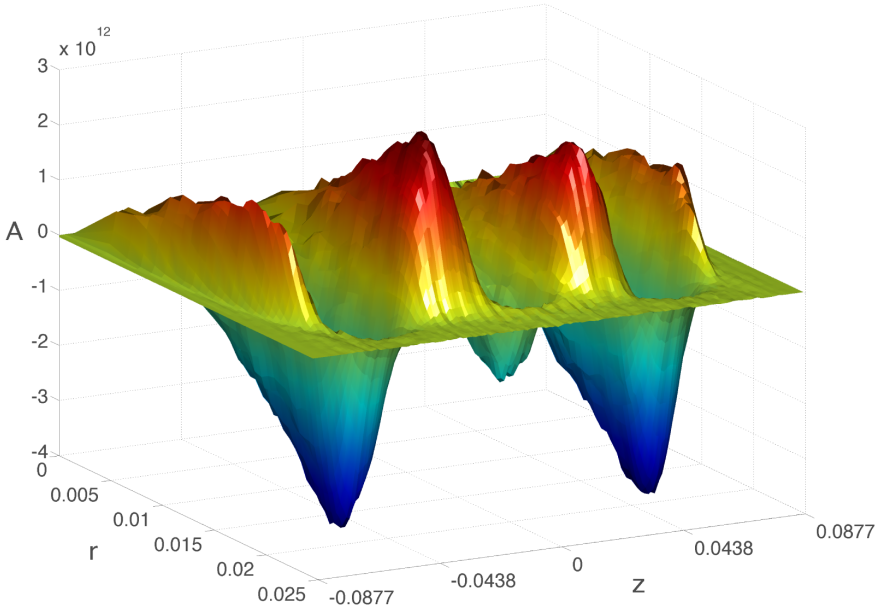

Figure 4.6. Mode shape for the second lower frequency with $k_z = 5/2$ (marked "A" in Figure 4.3).

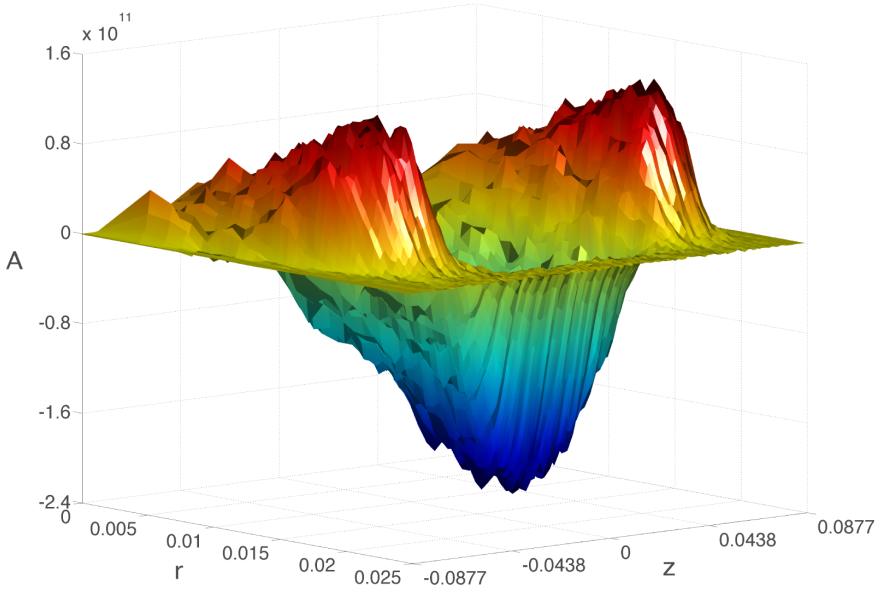

Figure 4.7. Mode shape for first upper frequency with $k_z = 3/2$ (marked "D" in Figure 4.3).

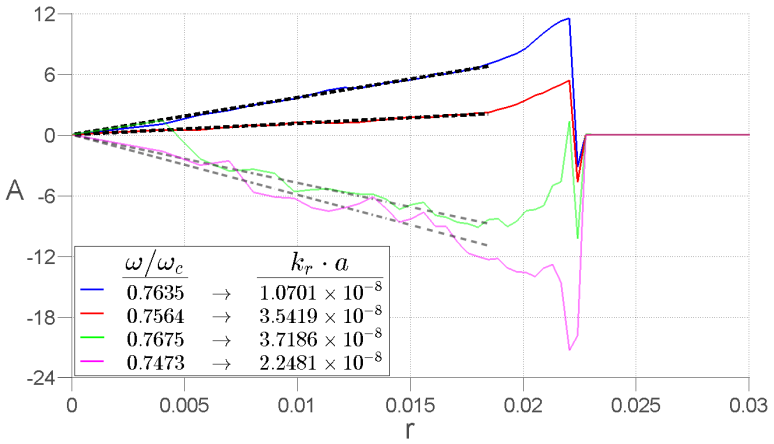

Figure 4.8. Nonlinear least squares fits to radial cross-sections of the normalized mode shapes for each of the four frequencies in the cluster of fundamental radial Bernstein modes shown in Figure 4.3. As expected, they exhibit a very small k_r and linear shape characteristic of the fundamental mode. The a referred to in the legend is the equilibrium radius of the plasma column (2 cm). Note the characteristic kick near the edge of the plasma that occurs when using normalized mode shapes. For this reason the fit does not extend completely out to the outer edge of the plasma.

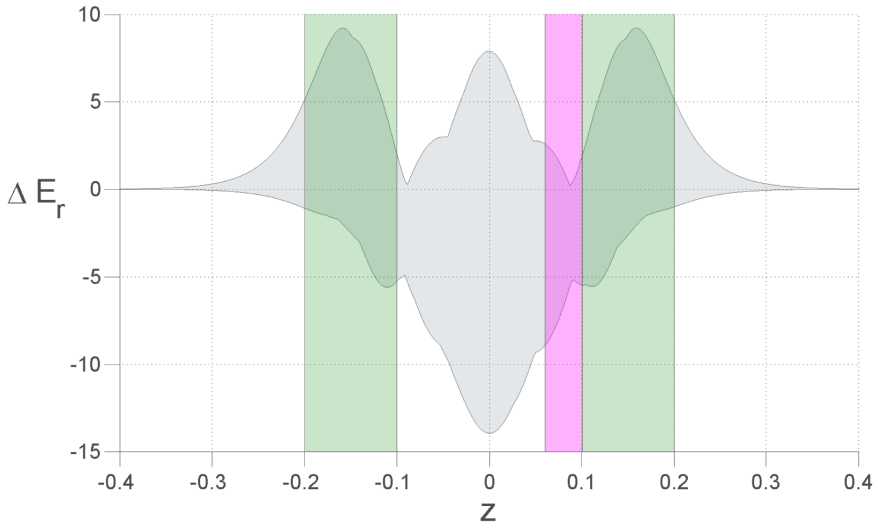

Figure 4.9. Shown in gray is the maximum/minimum change in radial electric field (from equilibrium) across the wall of the trap. The green sections indicate the location of the confinement rings while the purple section indicates the location of the diagnostic ring. The nodal structure in z is apparent and indicates that the ideal location for the diagnostic ring is not at the ends of the plasma, but near the middle.

our PIC code, the placement of the diagnostic ring in the experiment does not necessarily yield the optimal signal. In fact, the nodal structure we see near the ends of the plasma suggests that the signal on the wall may actually be the smallest in this region. To discover if this is true, we observed the radial electric field across the entire length of the wall at every time step and subtracted off the equilibrium field at the wall. We then looked for the place where the greatest change occurs. Figure 4.9 shows the maximum and minimum ΔE_r across the wall of the trap and indicates the current locations for the diagnostic and confinement rings. As suspected, the regions immediately adjacent to the confinement rings experience the smallest change in field. This means that the best location for the diagnostic ring is actually in the middle of the trap where the change in radial electric field is greatest.

4.3 Higher Order Modes

Unfortunately, we have tried driving the plasma at many other frequencies and have still been unsuccessful at exciting any of the higher order radial modes with just the end potentials. This means that the higher order modes will likely be completely hidden from the experiment. The flip side is that having fewer modes will make it easier to detect the fundamental mode. However, this being a simulation, it is important that we try to understand the other modes as well. Fortunately, we are able to do things in a simulation that would physically be impossible. Knowing that the radial Bernstein modes have equation (4.1) as their shape function allows us to *radially* drive the modes directly using the radial electric field without recourse to the axial motion. This is the radial equivalent of moving the axial potential walls in and out, but instead of just striking the ends of the plasma column, we are driving the plasma at every point along the radius. Because the amplitude of the drive is again somewhat arbitrary, we went through a process very similar to the one we used above for the axial drive amplitude. This time we found a good amplitude to be about 10 V/m.

The radial drive has the benefit of allowing us the possibility to seed a single mode, but in order for it to work we need the mode's precise frequency and k_r. Unfortunately, it appears that finite-length effects have a much stronger influence over the higher order radial modes; although we indeed see new frequency peaks appear on either side of the fundamental when using our radial drive (see Figure 4.10), they do not correspond very well to results from the 1D case. Instead these resonance peaks are shifted away from the fundamental and are very broad. To enhance the resolution we tried quadrupling the simulation's run time, but to very little effect. This suggests that the broadening of the frequency peaks is due to damping rather than just closely spaced resonances. The raw fluid velocity data for a single grid point indeed shows a rapid decay to a steady state (see Figure 4.11) Additionally we tried to vary the k_r in our drive in order to isolate individual peaks. Figure 4.12 shows two overlaid waterfall plots for two different driving frequencies and ten different radial wave numbers. Strangely the resonant frequencies shift back and forth as k_r

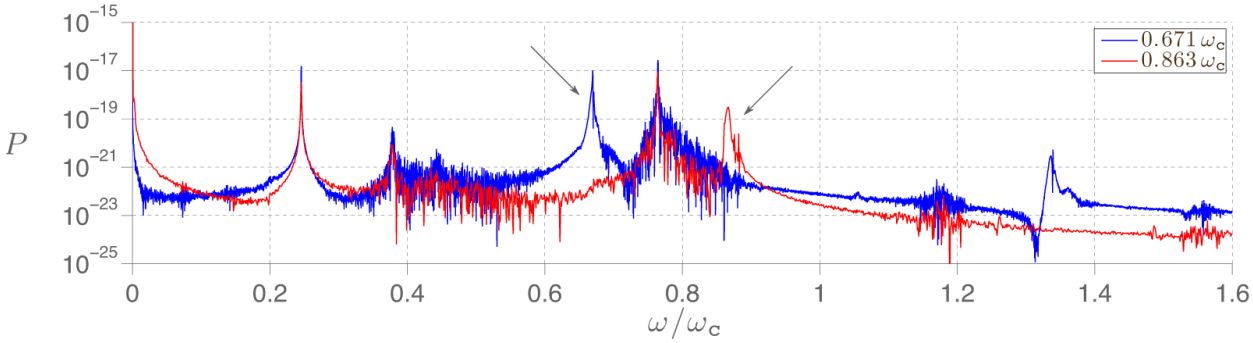

Figure 4.10. RMS radial position spectra showing two new resonance peaks that appear when using a radial drive as opposed to the axial drive. The blue curve used a driving frequency of $0.671\,\omega_c$ and the red curve used a driving frequency of $0.863\,\omega_c$. In both instances we used a $k_r \cdot a$ of 4.933.

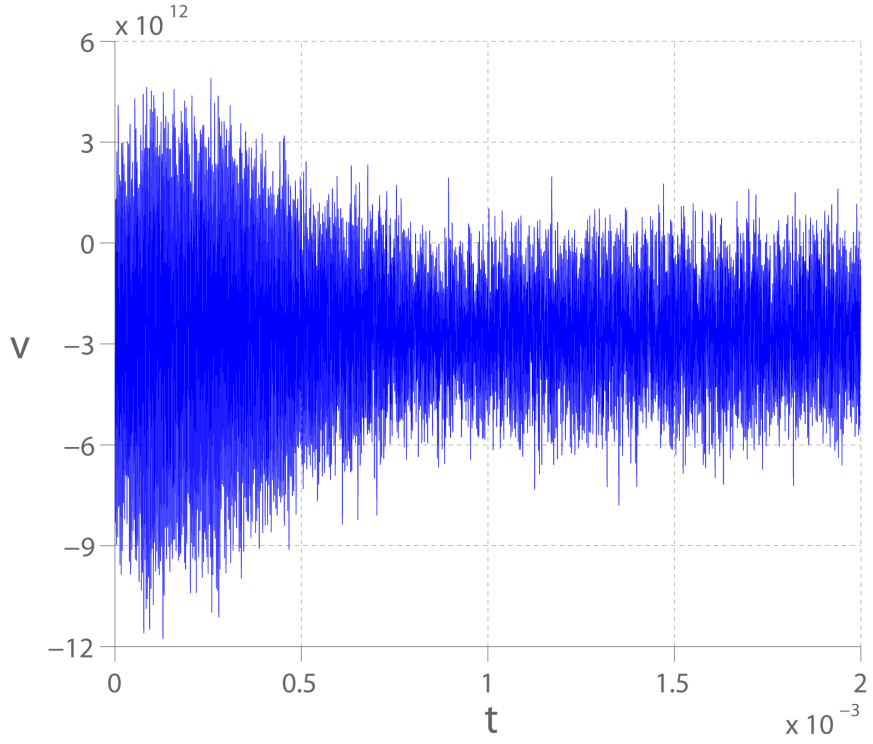

Figure 4.11. Raw fluid velocity data at a single grid point showing a rapid decay in signal to a steady state confirming that presence of damping in these modes. We believe damping is the cause for the broad frequency peaks we observe in our FFT signals for the higher order modes.

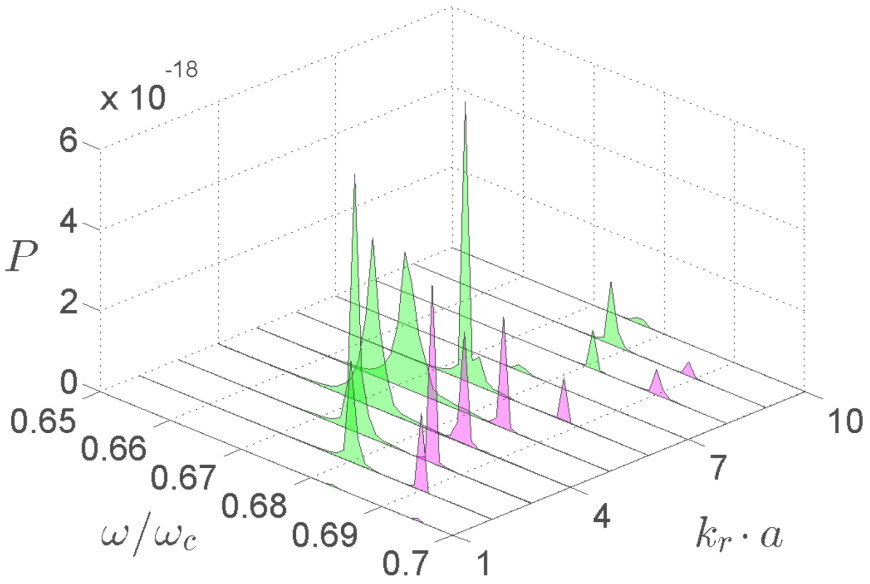

Figure 4.12. Waterfall plots of the RMS radial position FFT spectrum driven at $0.68\,\omega_c$ (green) and $0.7\,\omega_c$ (purple) for various radial wave numbers. In both cases, the peak frequencies undulate as k_r increases.

increases making it very difficult to correlate a wave number to a particular frequency.

Despite the difficulty of distinguishing the individual peaks in the FFT spectra, we have been able to isolate two mode shapes in the lower resonance that make sense. The data for these peaks comes from driving the plasma at $0.6715\,\omega_c$ and $k_r \cdot a = 4.462$. Counter intuitively, the first of these frequency peaks ($0.7029\,\omega_c$) lies considerably above the largest peak ($0.6694\,\omega_c$), but very close to where the 1D theory reports the frequency should be located ($0.7031\,\omega_c$). Being such a small peak, the shape is noisy, but resembles quite clearly the structure of the next radial mode with $k_z = 1/2$ (see Figure 4.13). While we were unable to locate a $k_z = 3/2$ mode, we were successful in finding the $k_z = 5/2$ mode at $0.6748\,\omega_c$ (see Figure 4.14); both the shape and spacing between the two peaks is comparable to what we observed in the fundamental mode. The radial structure is clearly different from the fundamental mode however. This time the mode shape in the interior

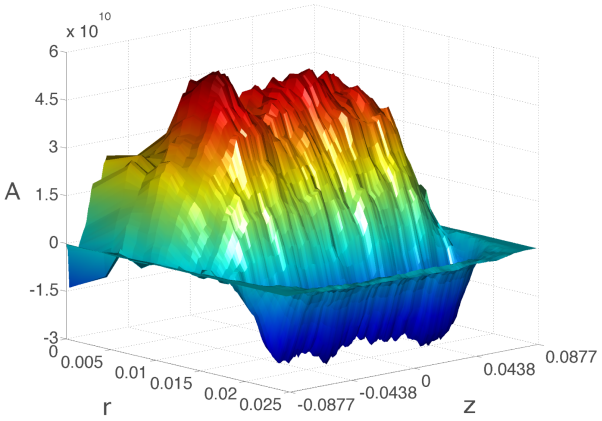

Figure 4.13. The first higher order radial mode with $k_z = 1/2$ at $\omega = 0.7029\,\omega_c$. The surface has been smoothed using a running average to show the general structure of the mode.

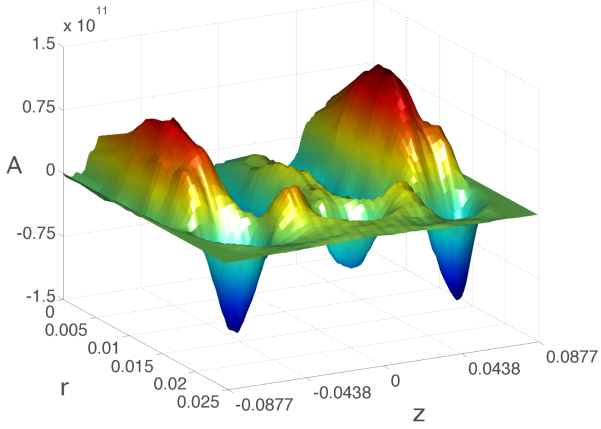

Figure 4.14. The first higher order radial mode with $k_z = 5/2$ at $\omega = 0.6748\,\omega_c$. This surface has also been smoothed, but the shape of this mode was much clearer than the $k_z = 1/2$. We were unable to locate the $k_z = 3/2$ mode.

of the plasma has a clear curvature. When we perform least squares fits for the two modes, as we did with the fundamental modes, we found $k_r \cdot a$ to be 4.5913 and 4.7669 for the two modes (see Figure 4.15). The 1D simulation gives $k_r \cdot a = 4.9331$; surprisingly, we still see good agreement between the 1D and 2D simulations. Although the resulting mode shapes for these two modes is somewhat questionable due to the amount of noise we see, the close agreement with the 1D simulations suggests that the radial structure is more than just numerical error.

It is also interesting to note that the largest frequency peak in this lower resonance ($0.6692\,\omega_c$) sits very close to where the 1D simulation says the second higher order mode should be located ($0.6718\,\omega_c$). However, the mode shape from the most prominent peak (see Figure 4.16) makes it quite clear that it is not the second higher order mode. It appears to have the radial structure of the first higher order mode, but it contains a rather strange axial structure. This could quite possibly be explained by a superposition of the $k_z = 1/2$ and $k_z = 3/2$ modes, which have similar amplitudes.

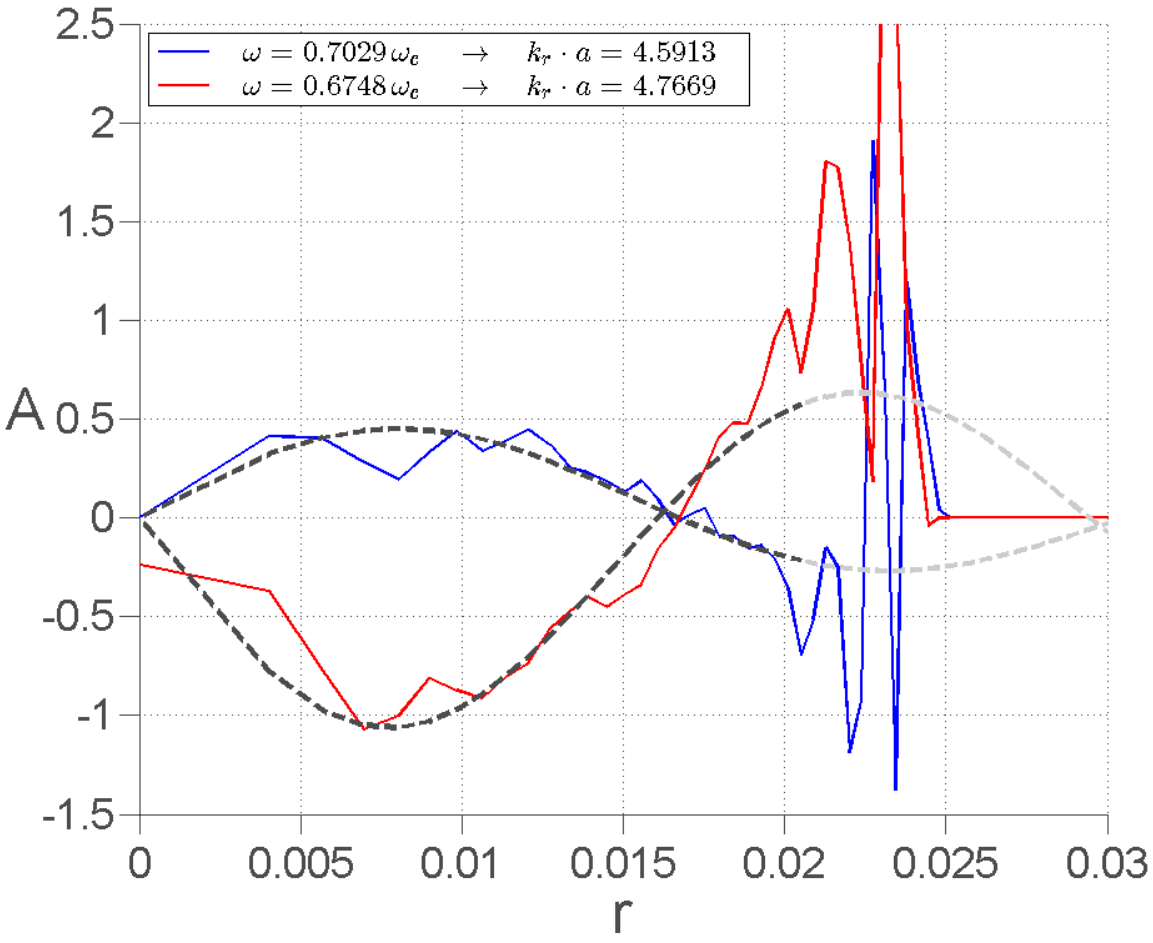

Figure 4.15. Nonlinear least Squares fits to radial cross-sections of the normalized mode shapes for the first higher order radial modes with $k_z = 1/2$ and $k_z = 5/2$. The resulting $k_r \cdot a$ values are markedly different than the fundamental mode and match closely to the 1D simulation. Again, the fits are not carried out all the way into the edge of the plasma due to edge effects and problems with normalization in this region.

Moving further down in frequency we find mode shapes do have more radial structure and that resemble what we would expect in the second higher order mode (see Figure 4.17), but again the structure is more complicated than what we observed for the fundamental and first higher order mode. We see equally confusing mode shapes in the upper resonance (see Figures 4.18 and 4.19). It is likely that these shapes represent a mixture of many different modes due to our current inability to seed any one particular mode. Careful experimentation of the drive parameters is needed in order to sort out exactly where these higher order radial Bernstein modes reside.[5]

[5]Care should be taken when experimenting with the larger values of k_r because the large grid spacing near the origin can distort the shape of the radial drive. More radial grid points should be added for $k_r > 7$.

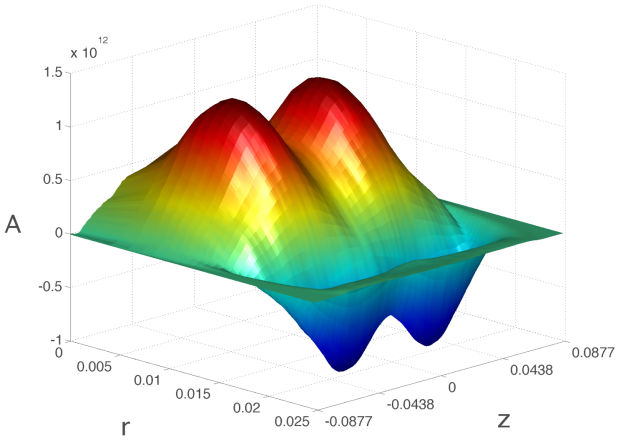

Figure 4.16. Mode shape for the largest frequency peak ($0.6692\,\omega_c$) in the resonance below the fundamental.

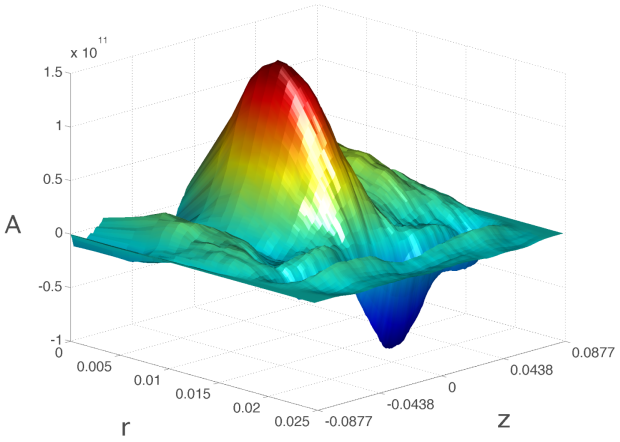

Figure 4.17. Mode shape for a frequency ($0.6636\,\omega_c$) lower than the largest resonance that shows more radial structure.

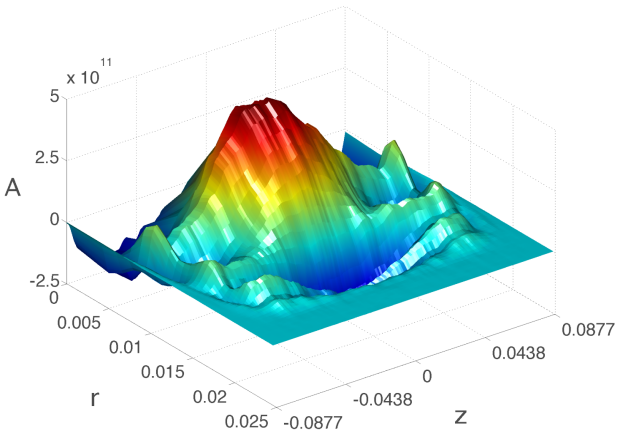

Figure 4.18. Mode shape for one of the two prominent peaks that shows up in the resonance above the fundamental at $\omega = 0.8612\,\omega_c$.

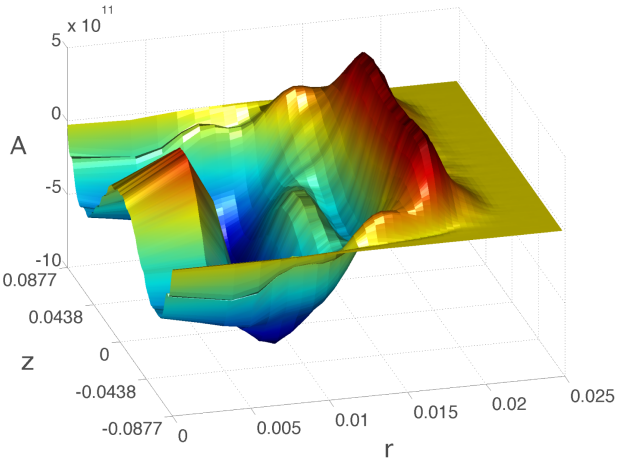

Figure 4.19. The mode shape for the second prominent peak that shows up in the resonance above the fundamental at $\omega = 0.8662\,\omega_c$.

Chapter 5

Conclusions

In this thesis, we have presented the benefits of using an r^2–z grid and describe how this type of grid can be implemented into a PIC code. We have also demonstrated effective (and non-effective) methods for testing the PIC code. In particular, we found that bilinear weighting in r^2 works effectively provided there are enough grid points to account for the large step size near the origin. We also found that single particle self-force calculations are very difficult to perform in a 2D r–z geometry due to the particle's inherent ringlike properties. Instead, we found that multiple particle calculations were necessary to confirm the absence of unphysical self-forces. Our PIC code showed good agreement with all of our theoretical tests and furthermore reproduces the frequencies calculated from a 1D code for the $m_z = 1, 2$, and 3 Trivelpiece-Gould modes, the Doppler-shifted cyclotron resonance, and the fundamental radial Bernstein mode.

The fundamental radial Bernstein mode was effectively driven by modulating the endcap potentials, but we found that the finite-length of the plasma allows axial standing waves to occur which then split the fundamental radial Bernstein mode into a tight cluster of resonant peaks. The FFT spectrum obtained from the induced charge on the diagnostic ring provides strong evidence that our experiment should be able to detect the fundamental mode. Furthermore, we found that placing the diagnostic ring near the center of the plasma would yield a stronger signal because the

nodal axial structure kills the signal near the ends of the plasma.

Finally, we found that driving the higher order radial modes with just the endcaps is ineffective and it is more than likely that these modes will be completely invisible to the experiment. However, we found that radially driving the plasma, although unphysical, produces new frequency peaks on either side of the fundamental. These frequencies indeed produce more radial structure like we would expect. Unfortunately, we were only able to successfully isolate two of the first higher order modes in the lower resonance. Although we did observe plenty of other mode shapes, the structure was too complicated to be attributed to any one particular mode. Most likely this is a result of exciting multiple modes simultaneously and we simply need to be more precise in choosing our drive parameters in order to effectively seed a single mode. We have also observed more damping with these non-fundamental frequencies and this could also be affecting our ability to see these other modes. At the very least, damping seems to be what is limiting our ability to resolve the frequencies in our FFT spectra when we run for longer time periods. Further work needs to be done to resolve the higher order radial modes. We also need to test whether reducing the confinement region by half will significantly alter the behavior of these modes and if the experiment will still be able to detect the fundamental mode under these new conditions.

Bibliography

[1] F. F. Chen, *Plasma Physics and Controlled Fusion*, 2 ed. (Springer, New York, 2006).

[2] A. J. Peurrung, R. T. Kouzes, and S. E. Barlow, Int. J. Mass Spect. Ion Proc. pp. 39–83 (1996).

[3] S. C. Neu and G. J. Morales, "Vlasov Equilibria and Stability of an Electron Gas," Phys. Rev. Lett. **78,** 4402–4405 (1997).

[4] R. W. Gould and M. A. LaPointe, Phys. Fluids B **4,** 2038–2043 (1992).

[5] R. W. Gould, "Theory of Cyclotron Resonance in a Cylindrical Non-neutral Plasma," Phys. Plasmas **2,** 1404 (1995).

[6] E. Sarid, F. Anderegg, and C. F. Driscoll, "Cyclotron Resonance Phenomena in a Non-neutral Multispecies Ion Plasma," Phys. Plasmas **2,** 2895 (1995).

[7] D. L. Book, "High-frequency Electrostatic Modes in Non-neutral Plasmas," Phys. Plasmas **2,** 1398 (1995).

[8] R. C. Davidson and N. A. Krall, "Vlasov Equilibria and Stability of an Electron Gas," Phys. Fluids **13,** 1543 (1970).

[9] S. E. Barlow, J. A. Luine, and G. H. Dunn, "Measurement of Ion/molecule Reactions Between 10 and 20 K," Int. J. Mass Spect. Ion Proc. p. 97 (1986).

[10] M. D. Tinkle, R. G. Greaves, and C. M. Surko, Phys. Plasmas **3,** 749–759 (1996).

[11] C. K. Birdsall and A. B. Langdon, *Plasma Physics Via Computer Simulation* (McGraw-Hill, New York, 1985).

[12] J.-P. Moreau, http://jean-pierre.moreau.pagesperso-orange.fr/Fortran/fbando_f90.txt, 1999 (Accessed June 2012).

[13] G. Engeln-MÏlges and F. Uhlig, *Numerical Algorithms with C* (Springer, New York, 1996).

[14] G. B. Arfken and H. J. Weber, *Mathematical Methods for Physicists*, 6 ed. (Elsevier, San Diego, 2005).

[15] M. D. Tinkle and S. E. Barlow, "Image Charge Forces Inside Conducting Boundaries," J. Appl. Phys. **90,** 1612–1624 (2001).

[16] S. E. Barlow, "Alternative Elctrostatic Green's Function for a Long Tube," J. Appl. Phys. **94,** 6221–6222 (2003).

[17] D. T. Lopes and C. C. Motta, "Electrostatic Force Between Rings and Discs of Charge Inside a Grounded Metallic Pipe Using the Green's Function Technique," J. Electrostatics **90,** 166–173 (2012).

[18] J. D. Jackson, *Classical Electrodynamics*, 3 ed. (Wiley, New York, 1999).

[19] E. G. Williams, "On Green's Functions for a Cylindrical Cavity," J. Acoust. Soc. Am. **102,** 3300–3307 (1997).

[20] O. Ciftja, A. Babineaux, and N. Hafeez, "The Electrostatic Potential of a Uniformly Charged Ring," Eur. J. Phys. **30,** 623–627 (2009).

Appendix A

Maxwellian Velocity Distribution

The Maxwell-Boltzmann distribution law is well known in molecular dynamics. Essentially it describes the partitioning of energy among the constituent elements of an isotropic fluid in thermal equilibrium. It assumes that the particles are identical but distinguishable and are freely moving except for brief elastic collisions. With these assumptions, the velocity distribution in any given direction is then equal to

$$P(v_i) = \frac{1}{\sqrt{2\pi}v_{th}} e^{-v_i^2/2v_{th}^2}, \tag{A.1}$$

where v_{th} is the thermal velocity $\sqrt{k_B T/m}$. Computers, however, are usually only equipped with a pseudo-random number generator only capable of producing a uniform distribution of random numbers between 0 and 1. Therefore we need a mapping of random numbers in a uniform distribution into a Maxwellian velocity distribution. A necessary and sufficient condition is that probability density must be conserved. This can be written in equation form as

$$\int_0^{\infty} P\,dv_i = \int_0^1 d\bar{R} = 1 \qquad \text{or} \qquad \int_0^{v_i} P\,d\bar{v}_i = \int_0^R d\bar{R} = R, \tag{A.2}$$

where R is a random number from the uniform distribution. Unfortunately, upon substitution from equation (A.1), the integral on the left hand side cannot be computed in terms of elementary functions. Even so, the integral can be done numerically (tabulated results are also commonly

used) and is known as the error function. This then must be inverted in order to obtain the velocity as a function of the random number. While the inverse error function can be computed, it is tedious to work with and lengthens computation time.

If a velocity distribution in more than a single dimension is needed, the Box-Müller method provides an economic alternative to the inverse error function. Box-Müller assumes a 2D distribution of velocities; therefore, the velocity v_i in equation (A.1) must be replaced by $v = \sqrt{v_x^2 + v_y^2}$. Transforming into polar coordinates, equation (A.2) becomes

$$\int_0^v \frac{1}{v_{th}} e^{-\bar{v}^2/2v_{th}^2} \bar{v} d\bar{v} = R. \tag{A.3}$$

Note the change in normalization constant from before. This integral is easily done by making the substitution $u = v^2/2v_{th}^2$. The result is

$$1 - e^{-u} = 1 - e^{-v^2/2v_{th}^2} = R, \tag{A.4}$$

whereupon solving for v yields

$$v = v_{th}\sqrt{-2ln(1-R)} = v_{th}\sqrt{-2ln(R_1)}, \tag{A.5}$$

where the final equality comes about because 1 minus a random number (between 0 and 1) is just another random number. The individual velocities v_x and v_y come from the polar equations

$$v_x = v\cos(2\pi R_2) \qquad \text{and} \qquad v_y = v\sin(2\pi R_2). \tag{A.6}$$

The only requirement for independent v_x and v_y is that R_1 and R_2 be uncorrelated random numbers. While the equations are derived assuming only two independent velocity distributions, more dimensions can be accommodated by simply repeating the method as many times as needed provided new uncorrelated random numbers are used each time. It should also be noted that there are alternative methods to Box-Müller that exist that are also very efficient (e.g. the ziggurat algorithm), but Box-Müller is fast and easy to implement.

Appendix B

Cylindrical Green's Functions

The derivations presented in this appendix are all known and published in the literature; however, to the author's knowledge, they are not published collectively. Because Green's functions are often derived with a particular problem in mind, not all Green's functions work equally well in every circumstance. Therefore, these derivations are given here as a quick reference in hopes that it will provide the reader with a diversity of functions with which to solve problems involving cylindrical systems.

Green's functions provide a powerful tool for solving problems such as the Helmholtz equation

$$\nabla^2\psi(\mathbf{x}) + k^2\psi(\mathbf{x}) = -4\pi\rho(\mathbf{x}). \tag{B.1}$$

The corresponding Green's function $G(\mathbf{x},\mathbf{x}')$ satisfies the homogenous point source equation

$$\nabla^2 G(\mathbf{x},\mathbf{x}') + k^2 G(\mathbf{x},\mathbf{x}') = -4\pi\delta(\mathbf{x}-\mathbf{x}'), \tag{B.2}$$

while being subject to the same boundary conditions imposed on $\psi(\mathbf{x})$ [14]. These formulas and their solutions simplify considerably when considering only electrostatic phenomena ($k = 0$). The equations to be solved are, thus, Poisson's equation and the corresponding point source equation

for the Green's function

$$\begin{aligned} \nabla^2 \psi(\mathbf{x}) &= -4\pi\rho(\mathbf{x}) \\ \nabla^2 G(\mathbf{x},\mathbf{x}') &= -4\pi\delta(\mathbf{x}-\mathbf{x}'). \end{aligned} \tag{B.3}$$

Because G satisfies a point source equation, physically G is proportional to the potential of a point charge. Therefore, with reference to Coulombs law, we can write down a general form for G as

$$G(\mathbf{x},\mathbf{x}') = \frac{1}{|\mathbf{x}-\mathbf{x}'|} + F(\mathbf{x},\mathbf{x}'), \tag{B.4}$$

where the first term on the right hand side is the contribution of the point charge and F is some function that arises due to having boundaries in the system. An important result of manually factoring out the point charge contribution is that F must satisfy the Laplace equation inside the volume:

$$\nabla^2 F(\mathbf{x},\mathbf{x}') = 0. \tag{B.5}$$

Physically, F is proportional to the potential produced by the induced charge residing on the boundary surfaces, a fact that is confirmed anytime the potential is obtained using the Method of Images [15]. We can then infer that F must also be negative since the induced image charge always carries a charge opposite to that of the volume charge. Now G will not always separate into two components like this; however, when it does, F can be used to calculate the image charge force on the volume charge.

To illustrate the usefulness of this method, consider a single point charge within a closed grounded surface. In this case, the only force acting on the charge is the image charge force. Unfortunately, if the full Green's function were used to calculate this force, G would need to be evaluated at $\mathbf{x}' = \mathbf{x}$ — a guaranteed pole! Of course the pole comes from expanding about $1/|\mathbf{x}-\mathbf{x}'|$, but this is the charge's own self-field and cannot contribute to the force anyway. Eliminating the self-field automatically eliminates the pole and allows the force to be calculated without problem [15, 16]. One important application of this method (which is particularly apropos to this paper) may be when

testing a three dimensional PIC code for unphysical self-forces by loading only a single particle and comparing the simulation's calculated force with that of the image charge force.

B.1 Expansion of Green functions

The Green's function is often obtained by expanding portions of both $G(\mathbf{x},\mathbf{x}')$ and $\delta(\mathbf{x},\mathbf{x}')$ using completeness relations of orthogonal functions. Herein lies a major difference between Green's functions; the Green's function inherits different properties depending on which functions are used and which dimensions are expanded. These differences can be beneficial especially when a particular Green's function leads to a solution that is unstable in one of its coordinates while another Green's function may lead to a solution that is stable [17].

The three dimensional delta function written in terms of its components becomes

$$\delta(\mathbf{x}-\mathbf{x}') = \delta(\rho-\rho')\frac{\delta(\phi-\phi')}{\rho}\delta(z-z'), \tag{B.6}$$

and the individual one dimensional delta functions can be expanded as follows [18]

$$\frac{\delta(\rho-\rho')}{\rho} = \sum_{n=1}^{\infty}\frac{2}{a^2 J_{m+1}^2(x_{mn})}J_m\left(\frac{x_{mn}}{a}\rho\right)J_m\left(\frac{x_{mn}}{a}\rho'\right) \tag{B.7a}$$

$$\delta(\phi-\phi') = \frac{1}{2\pi}\sum_{m=-\infty}^{\infty}e^{im(\phi-\phi')} \tag{B.7b}$$

$$\delta(z-z') = \frac{2}{L}\sum_{k=1}^{\infty}\sin\left(\frac{k\pi}{L}z\right)\sin\left(\frac{k\pi}{L}z'\right) \tag{B.7c}$$

$$= \frac{1}{2\pi}\int_{-\infty}^{\infty}dk\,e^{ik(z-z')} = \frac{1}{\pi}\int_{0}^{\infty}dk\cos[k(z-z')]. \tag{B.7d}$$

Likewise we can expand the Green's function using the same functions. However, the expansion of the left and right hand sides are only performed in two of the three dimensions, the advantage being that the resulting cancelation leaves us with an ordinary differential equation in the third dimension (as opposed to the original partial differential equation) that determines the ultimate

form of the Green's function. Finally, the appropriate boundary conditions are applied to determine the remaining constants.

Another popular technique for obtaining Green functions is by using an eigenfunction expansion. In this case, the Green's function is expanded in a series of normalized real orthogonal eigenfunctions that satisfy the homogeneous Helmholtz equation [14]. Thus, if φ_n is such an eigenfunction then it must satisfy

$$\nabla^2 \varphi_n(\mathbf{x}) + k_n^2 \varphi_n(\mathbf{x}) = 0, \tag{B.8}$$

and the Green's function can be written as

$$G(\mathbf{x}, \mathbf{x}') = \sum_{n=0}^{\infty} a_n(\mathbf{x}') \varphi_n(\mathbf{x}). \tag{B.9}$$

Upon substituting this into equation (B.2) and by expanding the delta function in its own eigenfunction expansion similar to what we did above, we can exploit the orthogonality of the eigenfunctions to solve for a_n:

$$a_n(\mathbf{x}') = \frac{\varphi_n(\mathbf{x})}{k_n^2 - k^2}. \tag{B.10}$$

Substitution of this back into equation (B.9) we obtain the Green's function eigenfunction expansion

$$G(\mathbf{x}, \mathbf{x}') = \sum_{n=0}^{\infty} \frac{\varphi_n(\mathbf{x}) \varphi_n(\mathbf{x}')}{k_n^2 - k^2}. \tag{B.11}$$

Note that, just as before, if the calculation is electrostatic then k can be set to zero.

While these are not the only methods for obtaining Green's functions, they probably have the most consistent derivation despite being long. Hence, due to the similarity between derivations, the details of obtaining each Green's function have been left out and only the final results are stated. Of course other methods of derivation exist and usually prove to be much shorter, but only by taking advantage of obscure relations (as done by Barlow [16]) or by exploiting symmetries as with image charge problems.

B.2 Infinite-Length Green functions

Using the completeness relations (B.7a) and (B.7b) to do the expansion of the left and right hand sides of (B.3), the final form of the Green's function becomes

$$G(\mathbf{x},\mathbf{x}') = \frac{2}{a}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}\frac{J_m\left(\frac{x_{mn}}{a}\rho\right)J_m\left(\frac{x_{mn}}{a}\rho'\right)}{x_{mn}J_{m+1}^2(x_{mn})}e^{im(\phi-\phi')}e^{-\frac{x_{mn}}{a}|z-z'|}, \tag{B.12}$$

whereas using the relations (B.7b) and (B.7d) instead yield the form

$$G(\mathbf{x},\mathbf{x}') = \frac{2}{\pi}\int_0^{\infty}dk\cos[k(z-z')]\sum_{m=0}^{\infty}(2-\delta_{0,m})\cos[m(\phi-\phi')] \times I_m(k\rho_<)\left[K_m(k\rho_>) - \frac{K_m(ka)}{I_m(ka)}I_m(k\rho_>)\right]. \tag{B.13}$$

Notice in this second form how the Green's function has been split into two pieces; one term is positive (volume charge contribution) and the other negative (image charge contribution).

B.3 Finite-Length Green functions

The derivation for the finite-length cases are very similar to the infinite-length cases and, therefore, the final forms of the equations look very similar. The trick is now to satisfy the boundary conditions of the end caps. Now if we use the radial and azimuthal expansions (B.7a) and (B.7b) the final form of G becomes

$$G(\mathbf{x},\mathbf{x}') = \frac{4}{a}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}J_m\left(\frac{x_{mn}}{a}\rho\right)J_m\left(\frac{x_{mn}}{a}\rho'\right)\frac{\sinh\left(\frac{x_{mn}}{a}z_<\right)\sinh\left[\frac{x_{mn}}{a}(L-z_>)\right]}{x_{mn}J_{m+1}^2(x_{mn})\sinh\left(\frac{x_{mn}L}{a}\right)}e^{im(\phi-\phi')}. \tag{B.14}$$

If expanding in the azimuthal and axial coordinates, we now opt in favor of using equation (B.7c) over (B.7d) because it lends itself better to the boundary conditions. The resultant Green's function

is now

$$G(\mathbf{x},\mathbf{x}') = \frac{4}{L}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}\sin\left(\frac{n\pi}{L}z\right)\sin\left(\frac{n\pi}{L}z'\right) \times I_m\left(\frac{n\pi}{L}\rho_<\right)\left[K_m\left(\frac{n\pi}{L}\rho_>\right) - \frac{K_m\left(\frac{n\pi a}{L}\right)}{I_m\left(\frac{n\pi a}{L}\right)}I_m\left(\frac{n\pi}{L}\rho_>\right)\right]e^{im(\phi-\phi')}. \quad \text{(B.15)}$$

We can generate a third Green's function by using an eigenfunction expansion. A suitable combination of eigenfunctions in cylindrical coordinates would be

$$\varphi_{nml}(\mathbf{x}) = \frac{\sqrt{2}}{aJ_{m+1}(x_{mn})}J_m\left(\frac{x_{mn}\rho}{a}\right)\frac{1}{2\pi}e^{im\phi}\sqrt{\frac{2}{L}}\sin\left(\frac{l\pi z}{L}\right). \quad \text{(B.16)}$$

Plugging this into equation (B.11) with $k = 0$ gives the following Green's function

$$G(\mathbf{x},\mathbf{x}') = \frac{8}{La^2}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}\sum_{l=1}^{\infty}\sin\left(\frac{l\pi}{L}z\right)\sin\left(\frac{l\pi}{L}z'\right)\frac{J_m\left(\frac{x_{mn}}{a}\rho\right)J_m\left(\frac{x_{mn}}{a}\rho'\right)}{\left[\left(\frac{x_{mn}}{a}\right)^2+\left(\frac{l\pi}{L}\right)^2\right]J_{m+1}^2(x_{mn})}e^{im(\phi-\phi')}. \quad \text{(B.17)}$$

An important thing to note here is that the eigenfunction expansions always carry an extra sum compared to the previous method and, therefore, could possibly result in longer calculation times when employing them in a problem.

B.4 Helmholtz Green functions

There may be circumstances when it is inappropriate to assume $k = 0$ and the full Helmholtz equation is needed. Although there is no direct application to the subject matter in this thesis, for completeness we simply reproduce the results of E. G. Williams [19] without motivation and refer the interested reader to his 1997 paper.

B.4.1 Dirichlet

$$G(\mathbf{x},\mathbf{x}') = \frac{2\pi}{L}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}\sin\left(\frac{n\pi}{L}z\right)\sin\left(\frac{n\pi}{L}z'\right) \times J_m(k_n\rho_<)\left[\frac{Y_m(k_na)}{J_m(k_na)}J_m(k_n\rho_>) - Y_m(k_n\rho_>)\right]e^{im(\phi-\phi')}, \quad \text{(B.18)}$$

where

$$k_n = \sqrt{k^2 - \left(\frac{n\pi}{L}\right)^2}. \quad \text{(B.19)}$$

$$G(\mathbf{x},\mathbf{x}') = \frac{4}{a^2}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}J_m\left(\frac{x_{mn}}{a}\rho\right)J_m\left(\frac{x_{mn}}{a}\rho'\right)\frac{\sin(k_{mn}z_<)\sin[k_{mn}(L-z_>)]}{k_{mn}J_{m+1}^2(x_{mn})\sin(k_{mn}L)}e^{im(\phi-\phi')}, \quad \text{(B.20)}$$

where

$$k_{mn} = \sqrt{k^2 - \left(\frac{x_{mn}}{a}\right)^2}. \quad \text{(B.21)}$$

B.4.2 Neumann

$$G(\mathbf{x},\mathbf{x}') = \frac{\pi}{L}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}(2-\delta_{0,n})\cos\left(\frac{n\pi}{L}z\right)\cos\left(\frac{n\pi}{L}z'\right) \times J_m(k_n\rho_<)\left[\frac{Y_m'(k_na)}{J_m'(k_na)}J_m(k_n\rho_>) - Y_m(k_n\rho_>)\right]e^{im(\phi-\phi')}, \quad \text{(B.22)}$$

where

$$k_n = \sqrt{k^2 - \left(\frac{n\pi}{L}\right)^2}. \quad \text{(B.23)}$$

$$G(\mathbf{x},\mathbf{x}') = -\frac{4}{a^2}\sum_{m=-\infty}^{\infty}\sum_{n=1}^{\infty}\frac{\overline{x_{mn}}^2 J_m\left(\frac{\overline{x_{mn}}}{a}\rho\right)J_m\left(\frac{\overline{x_{mn}}}{a}\rho'\right)}{(\overline{x_{mn}}^2 - m^2)J_m^2(\overline{x_{mn}})}\frac{\cos(\overline{k_{mn}}z_<)\cos[\overline{k_{mn}}(L-z_>)]}{\overline{k_{mn}}\sin(\overline{k_{mn}}L)}e^{im(\phi-\phi')}, \quad \text{(B.24)}$$

where

$$\overline{k_{mn}} = \sqrt{k^2 - \left(\frac{\overline{x_{mn}}}{a}\right)^2}. \quad \text{(B.25)}$$

Appendix C

Potential/Field Calculation of a Uniformly Charged Ring

The potential for a uniformly charged ring is probably most easily obtained by direct integration,

$$\Phi(\mathbf{x}) = \frac{1}{4\pi\varepsilon_0}\int \frac{dq}{|\mathbf{x}-\mathbf{x}'|}, \tag{C.1}$$

where $\mathbf{x}$ is the vector from the origin to the field point and $\mathbf{x}'$ is the vector from the origin to the source point. If we assume cylindrical coordinates centered somewhere along the axis of the ring, we can arbitrarily orient our axes such that the source point lies in the x–z plane. Then we can rewrite the denominator in the integral above using the Law of Cosines as

$$|\mathbf{x}-\mathbf{x}'|^2 = (\rho^2+z^2)+(R^2+Z^2)-2\sqrt{(R^2+Z^2)(\rho^2+z^2)}\cos(\xi), \tag{C.2}$$

where $\mathbf{x} = \rho\cos(\phi)\hat{x}+\rho\sin(\phi)\hat{y}+z\hat{z}$, $\mathbf{x}' = R\hat{x}+Z\hat{z}$, and ξ is the angle between them. The last term on the right hand side can be recognized, using the definiting of the dot product, as $2\mathbf{x}\cdot\mathbf{x}' = 2(\rho R\cos(\phi)+zZ)$. Combining this with the fact that dq for a ring in cylindrical coordinates is

$\lambda R\,d\phi$, equation (C.1) can be written like

$$\Phi(\rho,z) = \frac{\lambda}{4\pi\varepsilon_0}\int \frac{R\,d\phi}{\sqrt{\rho^2+R^2-2\rho R\cos(\phi)+(Z-z)^2}}. \tag{C.3}$$

This integral cannot be carried out in terms of elementary functions, but the integral can be written in terms of the complete elliptic integral of the first kind [14]:

$$K(m) = \int_0^{\pi/2} \frac{d\theta}{\sqrt{1-m\sin(\theta)}}. \tag{C.4}$$

This can be done using the substitution $\phi = \pi + 2\theta$ [20]. Noting that $\cos(\pi+2\theta) = -1+2\sin^2(\theta)$ the factor inside the square root in equation (C.3) becomes

$$(R+\rho)^2 + (Z-z)^2 - 4\rho R\sin^2(\theta), \tag{C.5}$$

whereupon factoring out $(R+\rho)^2+(Z-z)^2$ from inside the square root yields the desired form with $m = 4\rho R/[(R+\rho)^2+(Z-z)^2]$. The final result for the potential in terms of the complete elliptic integral of the first kind is

$$\Phi(\rho,z) = \frac{\lambda R}{\pi\varepsilon_0\sqrt{(R+\rho)^2+(Z-z)^2}}\,K\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right]. \tag{C.6}$$

The electric field can now be obtained by taking the negative gradient of the potential, which yields

$$E_r(\rho,z) = \frac{\lambda R}{2\pi\varepsilon_0\rho\sqrt{(R+\rho)^2+(Z-z)^2}}\left\{K\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right] - \frac{R^2-\rho^2+(Z-z)^2}{(R-\rho)^2+(Z-z)^2}\,E\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right]\right\}, \tag{C.7}$$

and

$$E_z(\rho,z) = -\frac{\lambda R(Z-z)\,E\left[\frac{4\rho R}{(R+\rho)^2+(Z-z)^2}\right]}{\pi\varepsilon_0\left((Z-z)^2+(R-\rho)^2\right)\sqrt{(R+\rho)^2+(Z-z)^2}}. \tag{C.8}$$

Index